

**THE COMPLETE
TECHNICAL PAPER PROCEEDINGS**
FROM:



Co-Owned and Published By:



These papers were presented at the SCTE·ISBE EXPO on October 17-20, 2017.



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

ISBN Number: 0-940272-56-3

©2017, NCTA – The Internet and Television Association and SCTE·ISBE.

All rights reserved.

Automating Success

Moderator: Jose Lugo, Altice USA

Ken Gold; Jim Gayton - EXFO Inc.
*Business Services in a DOCSIS Network:
Strategies for Supporting Service Level
Agreements (SLAs).....* 1

Darren Pralle - Spirent
*Automating Small Cell Turn-Up & SLA
Management.....* 9

Ethernet Provisioning and Activation

Moderator: Jose Lugo, Altice USA

Mark Gibson - Amdocs
*Ethernet Activation Goes Digital: The Next-
Generation of Automated Service-Activation
Systems and Tools for MSOs.....* 19

John Hawkins - Ciena
*Automating Service Creation and Provisioning
for Creative Ethernet Services: Delivering
Ethernet Service at Scale 43*

Ready or Not, 5G Is Coming: Understanding the Backhaul Requirements for 5G

Moderator: Guy McCormick, Cox

Jon Baldry - Infinera
*Cell Backhaul – Building the 5G-Ready Network
of the Future, Today.....* 56

Keith R. Hayes - Broadband Advisors Group, LLC
The Intersection of HFC and 5G..... 64

Service Assurance: Are You up to the Test?

Moderator: Steve Nocella, Charter
Communications

Robert J. Flask - Viavi Solutions
*Testing for SLA Compliance of Business
Services Over DOCSIS 3.1 77*

Sean Yarborough - Spirent
Zero Touch Service Assurance 91

The Impact of Virtualization on Business Services

Moderator: Dean Brewster, Charter
Communications

Ajay Manuga - Benu Networks
*Virtualizing Managed Business Services
for SoHo/SME Leveraging SDN/NFV
and VCPE 104*

Fady Masoud - Infinera
*Delivering High-Performance Business Services
Over a Dynamic Optical Infrastructure 112*

Finding Success in Software-Defined Business Services

Moderator: Dan Templin, Mediacom

Etienne Martel; Gregory Spear - Accedian
*How to Succeed with SD-WAN Using
Virtualized Service Assurance 124*

Ralph Santitoro - Fujitsu Network Communications
*SD-WAN and Beyond: Delivering Virtual
Network Services 151*

Hans Vanderstraeten; Erwin Six; Mihai Fagadar;
Willem Acke - Nokia
*Network Service Descriptors as a Factory: Agile
Networking for Differentiating MSO Business
Services 162*

The Journey From vCCAP to vHUB

Moderator: Jorge Salinger, Comcast

Asaf Matatyaou - Harmonic Inc.
*Real-World Deployment of a Virtual Cable Hub
..... 181*

Mark Szczesniak - Casa Systems, Inc.
*Getting Real Performance from a Virtualized
CCAP 197*

Ruobin Zheng; Wenle Yang - Huawei Technologies Co., Ltd.
Cable Access Network Virtualization: Headend Re-Architected as a Data Center 209

Upgrading Video and Cloud Through SDN

Moderator: David Grubb III, ARRIS

Srini Akkala - ARRIS
Media Processing on Cloud: Scalable, Manageable, and Cost Effective 226

John Jason Brzozowski; Mark Brittingham; Chris Luke; Zheng Yin - Comcast
Cloud Overlay (CLOVER): Extending the Cloud Virtually 238

Shlomo Ovadia; Jenson Thottian - Charter Communications
Cloud-DVR Real-Time Splunk-Based Monitoring and Alerting System 250

Virtualization—Stacks and Schedulers

Moderator: Dean Stoneback, SCTE - ISBE

Tong Liu - Cisco Systems Inc.
Interference Group Discovery for FDX DOCSIS 267

YuLing Chen; Alon Bernstein - Cisco Systems Inc.
Bridging the Gap Between ETSI-NFV and Cloud Native Architecture 282

Keith Alan Rothschild; Guy Meador III - Cox Communications; Brian Kahn – Sea Street Technologies
Fungible Virtualization Stacks: Refocusing on Optimization of Underlying Resources 307

Virtualizing Service Assurance

Moderator: Will Berger, Charter Communications

David S. Early; Jason K. Schnitzer - Applied Broadband, Inc.; Paul E. Schauer - Comcast
A Practical Approach to Virtualizing DOCSIS 3.1 Network Functions 325

Andrew Sundelin - Guavus
Leveraging Machine Intelligence and Operations Analytics to Assure Virtualized Networks and Services 344

John Holobinko; Todd Greene - Cisco Systems, Inc.; Ron Zimmerman – Cox Communications
Access Network Operations Savings Through Extending Automation and Orchestration Beyond Remote PHY 355

Digital Identity Meets Digital Devices

Moderator: Mark Hess, Comcast

Doug Fantuzzi; Hadar Sharon; Ira Kogan - Amdocs Media & Entertainment Solutions
Using Digital Identity to Drive Personalization, User Experience and Monetization 368

Arsham Hatambeiki - Universal Electronics Inc.
Smart Entertainment in the Smart Home: Reducing Friction in Content and Service, Discovery and Consumption, Across Devices at Home 388

Machine Learn From Both the Plant and the Home

Moderator: Martin Marcinczyk, Comcast

Gary Cunha - ARRIS
We Have Arrived. Our Light Bulbs Finally Have IP Addresses! Approaches for Proactively Managing Customer Experience and Reducing OPEX in a Cable Operations Environment... 412

Karthik Sundaresan; Jay Zhu - CableLabs
Access Network Data Analytics: (Machine Learning Applied to Cable Access Data) 442

The Impact of Good Field Service on Customer Satisfaction

Moderator: Morgan Kurk, CommScope

Erin Hayes - Midco
GIS a Success Story - Facilitating a Customer Journey from Design to Connection Using GIS 464

David Troll - Glympse
The New CX Standard: Location Data-Based Models for Driving Cost Savings and Improving Customer Satisfaction in Field Service Customer Journeys 470

Turning PNM Automation Into Customer Satisfaction

Moderator: Jeff Gutterman, Charter Communications

Marc Bellini - Nokia
The New Customer Care Experience: Moving from Scripted Dialogs to Automation, Omni-Channel and Predictive Analytics 491

Franklin Lartey - Cox Communications
Proactive Network and Technical Facilities Monitoring Using Standardized Scorecards .. 506

Using Enhanced Analytics to Enhance Customer Satisfaction

Moderator: Sergio Gambro, Altice USA

Sanjay Dorairaj; Bernard Burg; Nicholas Pinckernell - Comcast Corporation; Chris Bastian – SCTE•ISBE
Simplifying Field Operations Using Machine Learning: Applications of Machine Learning to Multiple System Operators (MSOs) 536

Anis Cheikhrouhou; Jim Davenport; Anish Kelkar - Nokia, Bell Labs Consulting
The Imperative of Customer-Centric Operations 557

D3.1 Technology Advancement I

Moderator: Alon Bernstein, Cisco

Ayham Al-Banna; Tom Cloonan - Network Solutions, ARRIS
When Does the DOCSIS 3.1 TaFD Feature Increase the Capacity of My Network? 569

Venk Mutalik; Brent Arnold; Benny Lewandowski-ARRIS; Phil Miguelez - Comcast; Mike Cooper - Cox
1024 to 4096 Reasons for Using D3.1 Over

RFoG: Unleashing Fiber Capacity by Jointly Optimizing D3.1 and RFoG Parameters 585

D3.1 Technology Advancement II

Moderator: Ron Hranac, Cisco

Tom Williams - Cable Television Laboratories, Inc.
The Universality of Modulation 612

Zhuo Zhao; Jiayou Meng; Haibin Tang - Cisco System
Best Practices for DOCSIS 3.1 Phase Noise Design in the Remote PHY Node 629

Designing Deep Fiber Networks

Moderator: Fernando Villarruel, Cisco

Yuxin Dai - Cox Communications
Unified Architectures for Remote PHY Backhaul and 5G Wireless Fronthaul 647

Todd Loeffelholz - Alpha Technologies Inc.
Fiber Deep Networks and the Lessons Learned from the Field 660

DOCSIS® 3.1 Operational Management

Moderator: Joe Godas, Altice USA

John Downey - Cisco Systems Inc.
DOCSIS 3.1 Downstream Early Lessons Learned 673

Doug Jones - CableLabs
DOCSIS 3.1 Configurations for HFC and RFoG 690

Energy Management at Edge Facilities

Moderator: Arthur Moore, Altice USA

John Dolan – Rogers Communications; Daniel Howard – Hitachi Consulting; Arnold Murphy - SCTi; Ken Nickel – Quest Controls, Inc; Dave Smargon – AIRSYS North America
Guidelines for Cable Facility Climate Technology Optimization: Cooling Optimization for Edge Facilities 708

Daniel Marut; Michael Baseline - Comcast; Daniel Howard; George Gosko; Supriya Dharkar; Riebeeck van Niekerk; Tanner McManus - Hitachi; Gregory Baron – USAF

*Energy Conservation Measure
Recommendations for Cable Edge Facilities:
Energy Audits and Analysis of Ten Cable
Headends 737*

FDX DOCSIS®—How It Works and How to Get There

Moderator: Sangeeta Ramakrishnan, Cisco

Ayham Al-Banna; Tom Cloonan; Jeff Howe - ARRIS
*Network Migration Strategies for the Era of
DAA, DOCSIS 3.1, and New Kid on the Block...
Full Duplex DOCSIS! 773*

Tong Liu - Cisco Systems Inc.
IG Discovery for FDX DOCSIS 793

Fixed Mobile Convergence

Moderator: Craig Cowden, Charter Communications

J.R. Flesch; Bryan Pavlich; David Virag; Charles Cheevers - ARRIS; Belal Hamzeh; Dorin Viorel – Cable Labs
*Can a Fixed Wireless Last 100m Connection
Really Compete with a Wired Connection and
Will 5G Really Enable This Opportunity? 807*

Glenn Laxdal - Ericsson
*Fixed Mobile Convergence in the Transition to
5G..... 851*

IP Architecture Solutions for a New Generation of Devices

Moderator: Rajesh Khandelwal Sr., Altice USA

Darren Gamble - Shaw Communications
*Shaw Communications IPv6 Deployment:
Developing Company Momentum..... 866*

Jeffrey Tyre; Wendell Sun - ARRIS
*Addressing IP Video Adaptive Stream Latency
and Video Player Synchronization..... 877*

Is PON the Final Frontier?

Moderator: Randy Kinsey, Charter Communications

Phil Miguelez - Comcast
*Moving Towards the Light: Migrating MSO
FTTP Networks to a Distributed Access
Architecture 894*

Kevin Noll – Tibit Communications; Steve Burroughs; Brionna Lopez - CableLabs
*An Architecture for Distributed EPON Access
..... 914*

Mobile Backhaul Optimization

Moderator: Ahmed Bencheikh, Charter Communications

Jennifer Andreoli-Fang - CableLabs; John T. Chapman – Cisco
*Mobile Backhaul Synchronization Architecture
..... 935*

John T. Chapman - Cisco; Jennifer Andreoli-Fang - CableLabs
Mobile Backhaul Over DOCSIS 961

RPHY Backhaul: New Ideas and New Solutions

Moderator: John Chapman, Cisco

Harj Ghuman - Cox Communications
*DWDM Access for Remote PHY Networks
Integrated Optical Communications Module
(OCML) 991*

Zhensheng Jia; L. Alberto Campos; Chris Stengrim; Jing Wang; Curtis Knittle - CableLabs
*Digital Coherent Transmission for Next-
Generation Cable Operators' Optical Access
Networks 1008*

The Impact of HFC-Wireless Convergence

Moderator: Craig Cowden, Charter Communication

John Chamberlain; Mark Alrutz - CommScope
Competitive Advantages of HFC Networks for Wireless Convergence..... 1048

Hugo Amaral Ramos; John Ulm; Zoran Maricevic;
Joseph Tavares; Claudio Albano - ARRIS
Making More with Less! A Case Study in Converging Wireline and Wireless Network Infrastructures Using Distributed Access Architectures 1059

Traffic Engineering Optimization I

Moderator: Manish Jindal, Wireless R&D, Charter

Tom Cloonan; Tushar Mathur; Ben Widrevitz; John Ulm – ARRIS; Ruth Cloonan – Blue Opus
The Big Network Changes Coming with 1+ Gbps Service Environments of the Future ... 1080

Karthik Sundaresan - CableLabs
Accurately Estimating D3.1 Channel Capacity 1109

Traffic Engineering Optimization II

Moderator: Jeff Finkelstein, Cox

Claudio Righetti; Emilia Gibellini; Florencia De Arca; Carlos Germán Carreño Romano; Mariela Fiorenzo; Gabriel Carro; Fernando Rodrigo Ochoa - Cablevisión S.A.
Network Capacity and Machine Learning ... 1140

John Ulm; Tom Cloonan - ARRIS
Traffic Engineering in a Fiber Deep Gigabit World 1165

Checking the Signal: What's Next in Wi-Fi?

Moderator: John Bahr, CableLabs

Carol Ansley; Charles Cheevers - ARRIS
Are We Done Yet? Opportunities in Wi-Fi with 60 GHz..... 1190

David Brownell; Salman Naqvi - Shaw Communications
Automation of the Best Practices Used to Evaluate 802.11 Access Network..... 1206

Enhancing the Interactive Experience

Moderator: Mark Hess, Comcast

Charles Cheevers - ARRIS; Michael McCluskey - Espial
Pay TV Is Not Dead! Myth Busting 101: It's (NOT) Inferior to OTT Cost and Value Experience 1249

Chris Lintz - Comcast VIPER
Smart Recordings 1340

Full-Duplex DOCSIS® 3.1

Moderator: John Holobinko, Cisco Systems, Inc.

Hang Jin; John Chapman - Cisco Systems
Echo Cancellation Techniques for Supporting Full Duplex DOCSIS 1353

Richard S. Prodan - Broadcom Limited
Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture 1375

Futureproofing Today's Converged Networks

Moderator: Damian Poltz, Shaw Communications

Martin J. Glapa; Hungkei Chow; Werner Coomans; R. J. Vale; Enrique Hernandez-Valencia - Nokia Bell Labs Consulting
Future Proofing Access Networks Through Wireless/Wireline Convergence..... 1419

Wayne Hickey; Joseph Shapiro - Ciena Corporation
*How Integrated Photonics Enhances Capacity
and Scalability for Fiber Deep Networks* 1434

How Services & Analytics Are Driving Future Network Design

Moderator: Tom Cloonan, ARRIS International

Tony Kourlas - Nokia
*Insight-Driven Network Performance
Management and Protection in the Cloud/IoT
Era 1446*

Jeroen Wellen; Prudence Kapauan; Amit
Mukhopadhyay - Bell Labs Consulting/Nokia
*Sustained Throughput Requirements for Future
Residential Broadband Service: Traffic Model
for Bandwidth Estimates 1460*

Improving IoT Connectivity

Moderator: Bill Warga, Liberty Global plc

Mark Bugajski; Paul Moroney - ARRIS
*MSO's Health Over Cable: The Ways We Can
Add Value 1486*

Chris Kocks - Pure Integration, LLC
*Powerful LPWAN Solutions for IoT: How Low
Power Wide Area Networks Will Accelerate
Smart City and Connected Business Initiatives
..... 1501*

Arun Ravisankar - Comcast Corporation
IoT for Peace of Mind 1519

In Home Wi-Fi Optimization

Moderator: Mark Poletti, CableLabs

Steven Harris - SCTE - ISBE
*Deploying and Optimizing the Next Generation
Wireless Home 1537*

Narayan Menon; Angelo Cuffaro; Jean-Louis
Gauvreau; Todd Mersch - XCellAir
*Optimizing and Protecting the Value of
Unlicensed Spectrum 1560*

Moving IoT Into the Network and Synchronizing These New Technologies

Moderator: Ralph Brown, CableLabs

Shengbo Ge - Cisco Systems (China)
*Can EMTC IOT Be Supported Over the HFC
Network: A Constructional Proposal of
MOVING IOT into HFC 1572*

J. Clarke Stevens - Shaw Communications
*Principles for Interoperability in the Internet of
Things 1581*

Taking a Closer Look at HDR

Moderator: Craig Cuttner, HBO

Sean T. McCarthy - Sean McCarthy, Ph.D.
Consulting
*High Dynamic Range for HD and Adaptive
Bitrate Streaming 1589*

David Touze - Technicolor; Leon van de Kerkhof -
Philips
*Single-Layer HDR Video Coding with SDR
Backward Compatibility 1606*

The Latest in SDN and Network Segmentation Tech

Moderator: Marty Glapa, Nokia Bell Labs

Jan Ariesen - Technetix Inc
Upstream Challenges with DOCSIS 3.1 1625

Mohcene Mezhoudi; Benjamin Y. Tang; Jean-
Phillippe Joseph; Enrique Hernandez-Valencia - Bell
Labs Consulting
*A Proposed End-To-End SDN Architecture for
MSO 1648*

Steven Krapp - MaxLinear, Inc.
*Virtual Fiber – 100 Gbps Over Coax: Coaxial
Cable: The Once and Future King 1667*

Blockchaining: What Is It and How Can We Leverage It?

Moderator: RT Lu, Ubee Interactive Inc.

Steve Goeringer - CableLabs
A Simple Overview of Blockchains: Why They Are Important to the Cable Industry 1681

Zane Hintzman - CableLabs
Comparing Blockchain Implementations 1691

I Had a Data Breach? Emerging Practices on Prevention and Detection

Moderator: Mary Haynes, Charter Communications

Robert Gyori - Charter Communications
IT Data Security in an MSO Environment ... 1718

David Yates - Guavus
Assuring Security in the IoT: Implementing a Behavioral Analysis Approach to Thwart IoT Attacks 1738

IoT Security: Is It Really a Risk?

Moderator: Derek DiGiacomo Sr., SCTE - ISBE

Petr Peterka - Verimatrix
Adapting Proven Technology to Counter IoT Threats 1747

Brian A. Scriber - CableLabs
Device Risks to Network Operators from IoT: Exploring the Critical Aspects of Onboarding, Authentication, Authorization and Accountability 1765

The Future of Security Architectures

Moderator: Chris Bastian, SCTE - ISBE

Steve Goeringer - CableLabs; Indrajit Ray – Colorado State University
Security of Open Distributed Architectures: Yet Another SDN and NFV Security Paper 1774

Ivan Ong - Comcast
Enhancing Public WIFI Security 1787

You're Stealing My Business!

Moderator: Mary Haynes, Charter Communications

Egbert Westervelt; Edward Florendo; Dave Belt - Irdeto
Service Theft in DOCSIS Networks: Identifying the Hidden Leaks in Your System 1797

Lucas Catranis; Brian Yuan; Dave Belt - Irdeto
Automated Detection for Theft of OTT Services and Content: Identifying Your Content Out in the Wild 1804

Business Services in a DOCSIS network

Strategies for Supporting Service Level Agreements (SLAs)

A Technical Paper prepared for SCTE•ISBE by

Ken Gold

Leader, Solutions Marketing
EXFO Inc.

250 Apollo Drive, Chelmsford, MA 01824
404-436-5756
ken.gold@EXFO.com

Jim Gayton

Marketing Manager, End-to-End Service Assurance and Analytics
EXFO Inc.

250 Apollo Drive, Chelmsford, MA 01824
978-3670-5622
jim.gayton@EXFO.com

Introduction

The value enterprise customers place on their communications services is typically much greater than that of the residential customer. Unlike residential customers, their livelihood depends on their ability to stay connected, conduct commerce, and communicate with employees and business partners. Staying connected in the global, 24-hour economy requires more than just access to the internet. For many, this connectivity needs to meet increasingly stringent performance requirements—and best effort simply won't cut it.

Cable operators have been providing internet connectivity for many years now via the data over cable service interface specification (DOCSIS) protocol. Initially intended for residential internet connectivity, enhancements to the DOCSIS protocol have enabled operators to serve not only the increasingly demanding residential market but also the more lucrative enterprise market, with higher capacity (1Gbit/s or more) uplinks and downlinks that are more typically required by businesses, especially those who rely on internet connectivity for their livelihood.

However, the battle for business services is about more than just bandwidth. As businesses evolve, their connectivity needs are also evolving and raw bandwidth is now no longer the problem. Overall quality of experience (QoE) becomes a more meaningful consideration for these customers as reliability and end-user experience become more critical to their customers' overall satisfaction. A service level agreement (SLA) between an operator and a business customer is the traditional tool for defining and policing contractual obligations with regards to service performance. Businesses, for which connectivity is an essential part of their operations model, demand SLAs to ensure that operators deliver the best possible QoE, and when this does not occur, provide for financial remedies to compensate for lost revenue and damage to reputation.

This paper examines the role that active monitoring and analytics play in managing network and service performance—not only when SLAs are involved but also for tracking performance of all other services against published service level objectives (SLOs). Taking a continuous, proactive approach to managing services in their networks by utilizing active performance monitoring combined with deep analytics of the derived key performance indicators (KPI), operators can confidently provide SLA-based services knowing they have complete visibility into the performance of their DOCSIS networks.

Background

1. Delivering data services over coax

Cable operators, or multiple systems operators (MSOs), have been delivering data services over their coax access infrastructure for 20 years based on DOCSIS standards. Initially, DOCSIS was primarily used for best-effort residential internet access and was limited in bandwidth and Quality of Service (QoS) features. However, through continuous development and improvements, DOCSIS is now capable of access speeds of up to 10 Gbit/s and includes QoS capabilities to allow for prioritized traffic handling as well as constant bitrate (CBR) and variable bitrate (VBR) services. These improvements have opened the enterprise business services markets through initiatives, such as business services over DOCSIS (BSoD). Cable operators now compete actively against the traditional communications service providers (CSP) for small, medium and large enterprise customers.

DOCSIS is essentially a Layer 1/ Layer 2 solution that supports tunneling of customers' Layer 2 (Ethernet) and Layer 3 (IP) traffic transparently across the coax access infrastructure. From the perspective of business services, this is an important consideration. Today's business services are typically based on Metro Ethernet Forum (MEF) definitions of services as defined in the MEF Carrier Ethernet (CE) standards. Not only do these standards define the service creation (such as E-LINE, E-LAN and E-TREE) but they also include features for performance assurance monitoring, something essential to businesses that rely on internet access for their livelihood. DOCSIS does not currently support MEF service creation or performance assurance; however, since it is essentially transparent at Layer 2 and 3, it does not preclude or block CE payloads and features.

As cable operators continue to shift their DOCSIS networks towards higher bitrates and more sophisticated services for both the enterprise and residential customer, it is critical to know whether their networks can support these services as advertised; especially if growth is greater than expected. This is where features like active performance monitoring become an invaluable tool. Without this, operators are essentially blind to performance issues in their networks until customers call to complain, at which point it is too late.

Business services over Ethernet

2. Types of business services

Ethernet services for businesses can be broadly categorized into three class: basic internet (best effort), business Ethernet (with a defined SLO) and enhanced business Ethernet (with a contractual SLA). Furthermore, within each class, there can be many variations depending on the performance commitments for such things as available bandwidth, service availability, latency, jitter and packet loss. The cost of the service is directly proportional to the level of performance required by the business and the type of guarantee associated with the service.

2.1. What is a SLA?

So, what exactly is a service level agreement? Essentially, it is a contract between the operator and their customer that provides performance commitments with regards to the service itself and, quite often, with regards to how the operator will behave when the performance commitments are not being met. Typical SLA requirements include such things as:

- Maximum latency
- Maximum delay variation (jitter) and/ or maximum inter-packet delay variation
- Maximum packet loss
- Service availability
- Help desk availability
- Help desk responsiveness
- Timeframe for equipment replacement for failure
- Penalties for not meeting commitments

The last point is primarily how a SLA differs from an SLO. For an SLO, operators publish objectives for the performance of the service and for their responsiveness during outage or degradation events; however, there is no financial penalty associated with a failure to meet the objective.

2.2. How are SLAs managed in a MEF Carrier Ethernet network?

The market for business Ethernet services is a robust and highly competitive market which, in the past, has tended to favor the traditional telecom CSPs over cable operators. The reasons for this are mostly historical and include breadth of coverage, incumbency based on traditional telecom voice services and the ability to offer SLAs on critical, high revenue and margin services.

CSP business services are typically based on Carrier Ethernet standards from the Metro Ethernet Forum (MEF). These standards, along with complimentary standards from the IEEE, IETF and ITU-T, provide facilities within Carrier Ethernet to support performance monitoring at various layers of the protocol stack, including ITU-T Y.1731 for Layer 2 and IETF RFC 6349 Two-Way Active Measurement Protocol (TWAMP) for Layer 3. Both protocols are considered ‘active’ performance monitoring methods since they require coordinated support from both ends of the measurement path to derive their relevant KPIs, such as delay (or latency), delay variation (or jitter) and packet loss.

To provide highly precise measurements, especially one-way latency and jitter between any two points across the entire network and correlate these measurements with other measurements in the network, requires each measurement point to provide a timestamp referenced to a common clock source. The accuracy and resolution of this timestamp is highly dependent on both the method used for distributing timing synchronization and the processing speed of the device.

Timing distribution in today’s Ethernet networks is typically done through either a standardized packet timing protocol (PTP), such as IEEE 1588v2, or a global navigation satellite system (GNSS), such as the U.S. global positioning system (GPS) or the European version, GALILEO. Of greater importance, however, is the speed at which the endpoint can process packets and timestamps in the active measurement protocol. Traditionally, MEF-based service endpoints that supported highly accurate performance monitoring would implement the timestamping functionality in hardware as close to the PHY level as possible to avoid any additional latency or jitter caused by software processing delays. By the same token, in virtualized networks, based on network function virtualization (NFV), the implementation of the active protocol—TWAMP for example—as well as timestamp management would be located in the service chain as close to the MAC layer processing as possible to avoid any additional latency and jitter.

Managing the SLA for a Carrier Ethernet service typically involves two domains. First, the CSP would actively track the relevant KPIs and establish threshold crossing alerts to notify their network or service operations center anytime a service was trending towards or had violated a SLA condition. Additionally, the customers themselves would either receive a monthly report of the performance of their services or have access to an online portal where they could track performance of their services in real-time or near real-time.

2.3. Options for implementing SLAs in DOCSIS networks

The question is: can today’s DOCSIS networks support SLA-based Ethernet business services? As described above, to support a SLA, three things are required:

1. Support for an active protocol for end-to-end performance measurements
2. The ability to implement the protocol without adding significant additional latency or jitter
3. A synchronized and precise timestamp

Let's look at each of these requirements.

The standardized, active protocols used for performance monitoring operate at either Layer 2, in the case of ITU-T Y.1731, or Layer 3 for IETF TWAMP. Since we have already shown that DOCSIS can transparently carry Layer 2 and Layer 3 protocols, we can safely assume these protocols can be supported.

Implementing the protocol will require additional functionality at the customer premise (cable modem) and possibly the headend in the case of hosted services. Supporting one of the standardized performance monitoring protocols in the headend is likely not an issue as most routers and switches support them today. Support at the cable modem will be more of a challenge as most, if not all, cable modems do not support these protocols today.

Like the active protocol, support for timestamp synchronization will be the same, with the headend likely supporting both a GNSS or PTP solution—but not the cable modem. Cable modems may support the network timing protocol (NTP), but this protocol does not provide sufficient timestamp precision for today's SLA based services.

2.3.1. Options for enhancing DOCSIS networks for SLA support

There are four options for enhancing DOCSIS networks to support SLA-based services using either ITU-T Y.1731 or IETF TWAMP active performance monitoring protocols.

- **Option 1: Add additional inline equipment**
This is easiest and quickest way to add support for both Carrier Ethernet-based service creation and performance monitoring for SLAs. Simply connecting a traditional CSP-style network interface device (NID) to the customer Ethernet port of the cable modem will allow the generation and termination of both the MEF service itself (policing, shaping, etc.) and the performance monitoring protocols required by the SLA. The downside of this solution is it involves additional space, power and cost at the customer premise as well as an additional management burden for operations. The NID would also introduce an additional point of failure.
- **Option 2: Software only solutions**
It may be possible to upgrade some of the existing cable modems to support the active performance monitoring protocols and packet timing protocol; however, such a solution would likely suffer from unacceptable latency and jitter issues since the cable modem would probably not support real-time timestamp functionality, or be optimized for the required processing speed.
- **Option 3: Upgraded DOCSIS equipment**
This method would involve developing a new DOCSIS cable modem with active performance monitoring protocols and timestamp synchronization support. Such a solution would eliminate most of the negative issues associated with Option 1. However, the solution would still likely be cost-prohibitive from a network retrofit and operations integration perspective as well as the MSOs traditional modem upgrade cycle.
- **Option 4: Network Function Virtualization (NFV)**
This option holds the best promise for cost-effective implementation. By leveraging the software defined networking (SDN) and NFV transformation already being considered by many cable operators, the development of a cable modem that supports virtualized network functions (VNFs) would allow for the instantiation of a protocol, such as Y.1731, directly into the service chain at the customer premise. Such a cable modem could also support the required timestamp

functionality, through an additional VNF, positioned in the service chain to ensure minimal impact on latency and jitter measurements.

3. What will be the impact of network virtualization?

Network virtualization represents a fundamental change in the way communications networks are built and operated. By leveraging lower costs, high-performance server platforms, along with ‘white box’ edge devices based on industry-standard X86 processors, entire networks can be built using SDN and NFV. What this means to operators is a lower-cost network, based on open platforms, implemented in software, which can be readily scaled to meet demand and leverage the latest techniques for rapidly developing new services. Since the network and its’ services are built in software, new services, enhanced network functionality and upgrades are essentially handled as software upgrades. With regards to managing SLA-based services, new or existing services can be upgraded simply by adding a standards-based performance monitoring VNF wherever it is required.

Many cable operators are already considering how they can leverage this industry transformation to benefit their bottom line. Some of the key benefits of virtualized networking include:

- The ability to leverage economies of scale when purchasing networking hardware – i.e., servers, storage, white box edge devices
- The ability to define new services as software functions only, enabling an agile or dev-ops approach that save both time and money and does not require new hardware
- Existing services can be upgraded simply by adding an additional VNF into the service chain
- Many of the functions traditionally found on the customer premise, such as firewalls, traffic policing/shaping or other QoS management functions, can be aggregated back into the core of the network where they can be managed more consistently and securely
- SDN and NFV networks leverage automation and orchestration heavily, greatly reducing the need for manual intervention in most operations processes

3.1. SLA support comes for free

One of the features of a virtualized network is that the service layer is abstracted fully from the physical layer. The benefit of this separation is that the network and service topologies can be continually optimized to address changes in the network, such as traffic congestion or hardware failure. However, this also means it is no longer possible to infer service quality from network QoS metrics since the service can be moved at any time due to optimization. For this reason, the only way to manage end-to-end service quality is to implement performance monitoring at the service endpoints. By doing this, performance monitoring becomes part of the service itself and stays with the service—regardless of how the network is optimized. Consequently, every service is equipped to support a SLA, if required.

3.2. What about the impact of DAA?

While network virtualization may still be in the planning stages, most cable operators are investing in the Distributed Access Architecture (DAA) standards. This initiative will drive a great many changes in the access/coax part of the network. As a result, it presents an ideal opportunity to introduce X86-based cable modem devices, which could support DAA as well as standard-based performance monitoring—all with the longer-term goal of supporting virtualized networking once it is adopted.

Introducing performance monitoring capabilities on all services, not just SLA-based services, will allow cable operators to have complete visibility into the performance of their DOCSIS networks, allowing them to confidently offer SLA-based services knowing the network can fully support them.

Making your network smarter

4. Data-driven operations and marketing

An operators network is their greatest asset. Operators that can best leverage their networks will have a distinct advantage in a highly competitive market. By instrumenting most, if not all, of the services to generate KPIs on a continuous basis, operators can implement a big-data analytics environment to develop an extremely granular view of their networks—both from the operational standpoint and from the perspective of how their customers are using their services. Knowing if a SLA has been violated is certainly important; however, being able to predict the behavior of a service against a SLA gives operators the opportunity to remedy the problem before the customer notices. This saves not only the cost of a SLA violation, it also helps to maintain a positive public image and keep high-revenue customers happy. Understanding which services are in greatest demand as well as where they are located, and correlating this data against non-network information, such as economic growth indicators, can provide insight into future marketing campaigns, new service growth opportunities and network capacity planning.

4.1. The role of analytics in predictive SLA management

Analytics have become integral tools for managing SLAs. By being able to track KPIs at the network, service and customer level, and correlate them against issues found throughout the network or external to the network, such as extreme weather or natural disasters, operators can develop smarter insights that can alert them when a service is at risk of violating a SLA. Having this detail in advance of the violation, especially if automation plays a key role in orchestrating troubleshooting, allows operators to take a proactive stance in support of SLA management, rather than the more traditional approach.

4.2. Leveraging analytics for service innovation

This same big data analytics environment can also be leveraged to provide a very granular view of how each and every service in the network is being utilized. By analyzing such things as types of service, growth of these services, customer loyalty per service, price, cost or even external factors, such as new construction activity, economic growth, population demographics and more, operators can develop extremely granular and targeted marketing campaigns with a very high degree of confidence. And based on this information, investment programs can be developed, again with a very high degree of confidence of success.

Conclusion

Business services are an important part of cable operators' strategic growth initiatives. Being able to offer SLAs on these services and compete against CSP operators is critical to this growth. DOCSIS networks have come a long way in their ability to offer the bandwidth required by medium and large business; however, they still lack the facilities necessary for effective SLA management.

Several options exist for enhancing existing DOCSIS networks with standard-based active performance monitoring to generate the required KPIs for SLA management. Unfortunately, most of these options come with some serious drawbacks in terms of measurement accuracy or cost. However, the cable industry is currently going through significant architectural changes—both now and in the near future. This opens up the opportunity to enhance the DOCSIS network to support industry competitive SLAs.

Abbreviations

CSP	communications service provider
DAA	distributed access architecture
DOCSIS	data over cable service interface specification
GNSS	global navigation satellite system
GPS	global positioning system
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
KPI	key performance indicator
LLC	logical link control (layer)
MAC	media access control (layer)
MEF	Metro Ethernet Forum
MSO	multiple systems operator
NFV	network function virtualization
NID	network interface device
NTP	network timing protocol
PHY	physical (layer)
PTP	packet timing protocol
QoE	quality of experience
QoS	quality of service
RFC	request for comments
SDN	software defined networking
SLA	service level agreement
SLO	service level objective
TWAMP	two-way active measurement protocol
VNF	virtualized network function

Automating Small Cell Turn-up & SLA Management

A Technical Paper prepared for SCTE•ISBE by

Darren Pralle

Sr. Manager, Product Marketing
Lifecycle Service Assurance
Spirent
Darren.Pralle@spirent.com

Introduction: The Exponential Growth of Small Cell Backhaul & Ethernet Services

The telecommunications marketplace is experiencing a significant growth in small cell backhaul and Ethernet backhaul service deployment. This year alone, revenues for mobile data and mobile networks have increased by 16% while total units are up by 35%. And these numbers are only predicted to grow. By 2021, the small cell market will skyrocket to \$2.2B with a CAGR of 8.4% and more than 2 million units deployed in the network to support the surge in mobile traffic (source: IHS). The number of small cells being deployed is going to vastly outnumber those of macro cells.




Figure 1 – Worldwide Service Provider Small Cell Units CY 2016-2019 (Source: IHS)

With the increased deployment and growth of small cells comes an increased demand for small cell mobile services. However, to truly realize the benefits of a small cell deployment, the cell must first be connected into the mobile network infrastructure. This connectivity drives the demand for backhaul services. The industry tends to refer to these backhaul services as Ethernet services, implying a single technology; however, small cells use a variety of technologies for backhaul connectivity such as microwave, hybrid fiber and coax, A/C cable networks, dedicated fiber, and more.

Small cell backhaul equipment revenue is predicted to grow to \$1.2B in 2021 at a 5-year CAGR of 51%. No matter the backhaul technology in place, the bottom line is that reliable, low latency connectivity is imperative for these small cells so they can deliver the high quality mobile services that customers expect. To meet these high customer expectations, services need to be guaranteed with service level agreements (SLAs). SLAs govern acceptable levels of packet delay, jitter, loss, availability, and other key metrics for small cell connectivity to the core mobile network. The only way to effectively monitor backhaul SLAs is through continuous active testing of the links.




Figure 2 – Outdoor Small Cell Backhaul Equipment Revenue CY 2016-19 (Source: IHS)

While it is possible to turn-up and validate that backhaul services are working and small cells are functioning as designed, service providers need to be able to monitor these connections in real-time and in an ongoing fashion to ensure that SLAs are being met. Adding to the already complex nature of these deployments is the fact that many providers operate heterogeneous backhaul solutions using a plethora of technologies and equipment vendors. Services are needed in specific locations and the backhaul technology available to provide connectivity varies based on these locations. For example, some providers have invested heavily in fiber backhaul networks, while other providers use a mix of microwave, legacy TDM, cable modem, and hybrid fiber/COAX connections.

This paper explores key challenges service providers face as they deploy large volumes of small cells with heterogeneous backhaul connectivity. The paper explains how automation of small cell acceptance workflows can address these challenges and deliver significant benefits in terms of speed of deployment, cost savings and customer satisfaction.

Small Cell Deployment Challenges

The challenges for service providers deploying small cells are threefold:

- 1) Service delivery
- 2) Trouble/fault management
- 3) SLA management

The critical issue from a service delivery perspective is the installation and turn-up of the small cell, including service provisioning and validation. With thousands upon thousands of deployments, traditional turn-up methods are not scalable. Additionally, there is a certain percentage of sites which will be non-functional at turn-up and providers need a way to quickly identify these defective units. Service validation is another concern because failure to deliver the service that was ordered leads to increased customer dissatisfaction and churn.

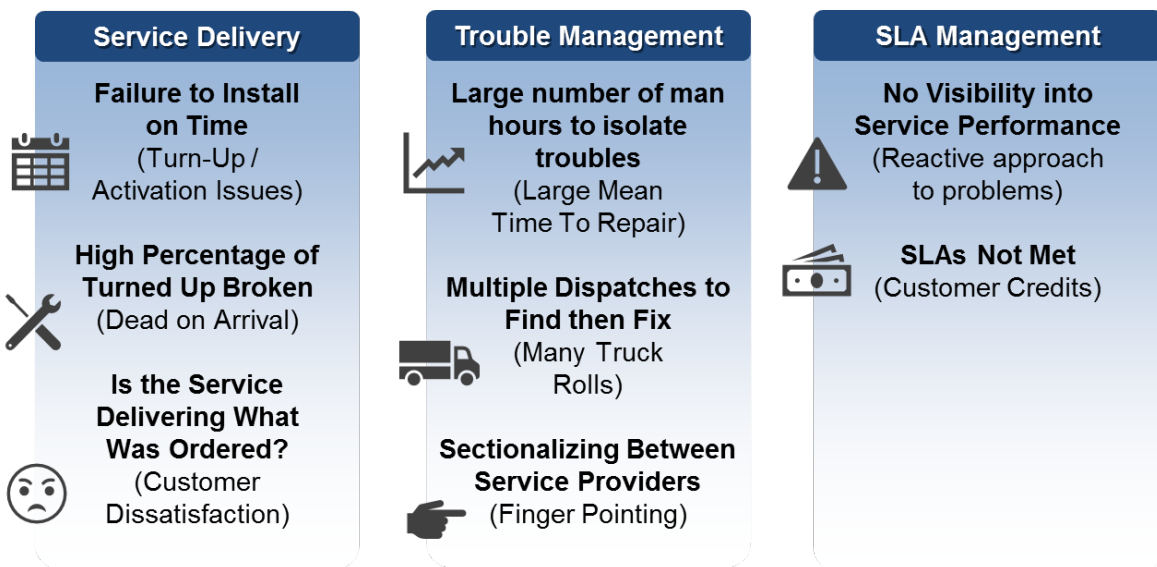


Figure 3 – Small Cell Deployment Challenges

Once the backhaul service is up and active, providers transition into SLA and trouble management mode. SLA management challenges include limited visibility into the service performance and the inability to proactively manage SLAs, both of which can result in significant issues. These challenges often lead to a reactive approach to fault resolution, which can have serious monetary consequences for the provider.

When dealing with fault resolution, an adhoc approach to troubleshooting results wasted time spent isolating the root cause of issues, leading to an increased time to fix the issue and remedy associated SLA violations. Additionally, multiple dispatches are often required to find and fix the issue which can be very costly. When the backhaul is provided by an off-net service provider, too much time is often spent trying to isolate issues between service providers. As most of these services are being delivered through a third-party vendor, alternate access vendor (AAV), or type II service provider, there tends to be a lot of finger pointing back and forth between providers to determine whose network is at fault. If providers are unable to rapidly determine who owns the problem, it further delays resolution, which exacerbates SLA violations, costly fees, and customer dissatisfaction.

Current Test & Assurance Approaches

Many service providers are deploying small cells and see the need for major network expansions consisting of thousands or tens of thousands of small/mini-macro cells. Providers must provision IP backhaul for each site and typically must support multiple network equipment vendors. Additionally, they have multiple access vendors or off-net providers with various topologies such as microwave, fiber, DOCSIS, and a public Internet with asymmetrical bandwidth. This creates significant challenges for standardizing turn-up, provisioning, and monitoring processes.

Furthermore, many service providers take a manual approach to the testing and assurance requirements of these processes. The acceptance workflow usually involves manual testing of backhaul connectivity and performance for every small cell immediately after turn-up, with no good way to ensure that all topologies and vendors use exactly the same process. If the KPIs are not acceptable and the validation process is incomplete or inconsistent, diagnosing the root cause and fixing the fault can be extremely difficult and time consuming.

A key challenge to this manual approach is that providers need to manage these deployments using their current resources. They want to avoid the cost of hiring additional staff or contractors, yet they need to significantly increase the scale of deployment. With traditional processes, providers can activate approximately 10 cells per day, but to deploy large numbers of sites, they need to be activating at least 100 cells per day with a process that is consistent and reliable. The reality is that a manual test approach can't be scaled fast enough to meet these objectives and it certainly can't be done with existing resources.

Automation of Turn-Up & SLA Monitoring

To meet the demands of rapid, high-volume small cell deployments, the acceptance testing and SLA monitoring workflows must be automated. Based on our history and experience with small cell backhaul deployments, about 80% of initial turn-ups will be successful while approximately 20% will have a material performance issue that requires troubleshooting. Service providers can implement “zero-touch” fully automated approaches for the 80% of initial turn-ups with no issues. The automation consists of the following elements: detection of new links, performance testing, creation and storage of turn-up records, and setup/initiation of SLA monitoring.

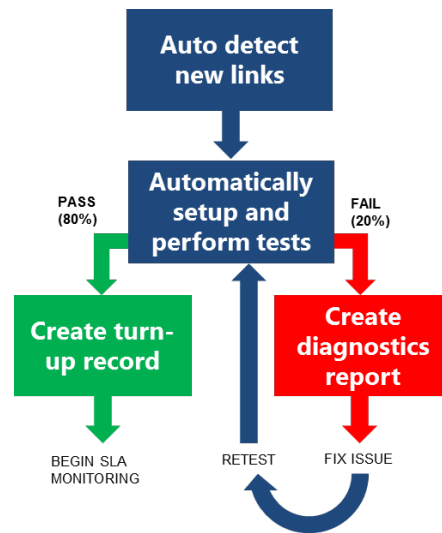


Figure 4 – Small Cell Deployment Automation Workflow

For the 20% of activations that require troubleshooting, automation can accelerate the troubleshooting process by providing detailed diagnostic reports that help isolate root causes. Diagnostic data associated with the turn-up failure, including isolation to a specific network segment or element, can also be pushed to the appropriate organization to drive faster fault resolution.

1. Automation Solution Components

There are several key components required to automate the process of small cell turn-up and SLA management:

- Centralized active test probes and/or Virtual Test Agents (VTAs)
- Service Assurance Test Controller
- Small cell network elements with support for industry test standards
- Inventory Interface
- Trouble Ticket Interface

The service assurance test controller is needed to configure and manage the active test probes/VTAs and communicate with network elements and small cells. Fortunately, most small cell network elements support Ethernet or IP service assurance industry standards (e.g., TWAMP, Ethernet OAM, MEF 48, MEF 46 standards) which allow standardized service assurance tests to be performed across backhaul connections between probes or agents and the small cell.

The Inventory Interface/Gateway detects newly activated small cells and configures and executes backhaul validation tests. An accurate inventory database is a must: automation is not possible without an accurate network inventory which provides a clear understanding of how the service is built, the service attributes, what network elements are included and how they are configured.

A Trouble Ticket Interface/Gateway is needed for the 20% of small cell backhaul deployments that typically fail on initial turn-up. This allows the test controller to communicate with back office systems, open a ticket, and send it to the correct organization.

2. Putting it All Together

Once the key components of the automation process have been assembled, they must be integrated to perform the following automated steps:

- 1) The Inventory Gateway detects a new live backhaul link and loads the network configuration to the Service Assurance Test Controller.
- 2) The Inventory Gateway builds service records to enable testing and sends test requests to various test heads, whether physical or virtual.
- 3) The test agents then send tests to the small cell and perform SAT tests, which can be a combination of MEF 48, RFC 5164, or RFC 2554, depending on the desired level of activation testing. Some providers have the need to test up to Layer 4 TCP throughput in order to validate the service from the core to the macro/small cell. If desired, such tests can be performed as well.
- 4) If the test is successful, a turn-up record is created and 24x7 active service monitoring begins.
- 5) If the test is unsuccessful, the Test Controller creates a diagnostic report and trouble ticket, both of which are routed to the appropriate organization. Based on the accuracy of the inventory data, the Test Controller can even perform a sequence of automated fault isolation tests to determine exactly where the fault is located and include this location in the trouble ticket for faster fault resolution.

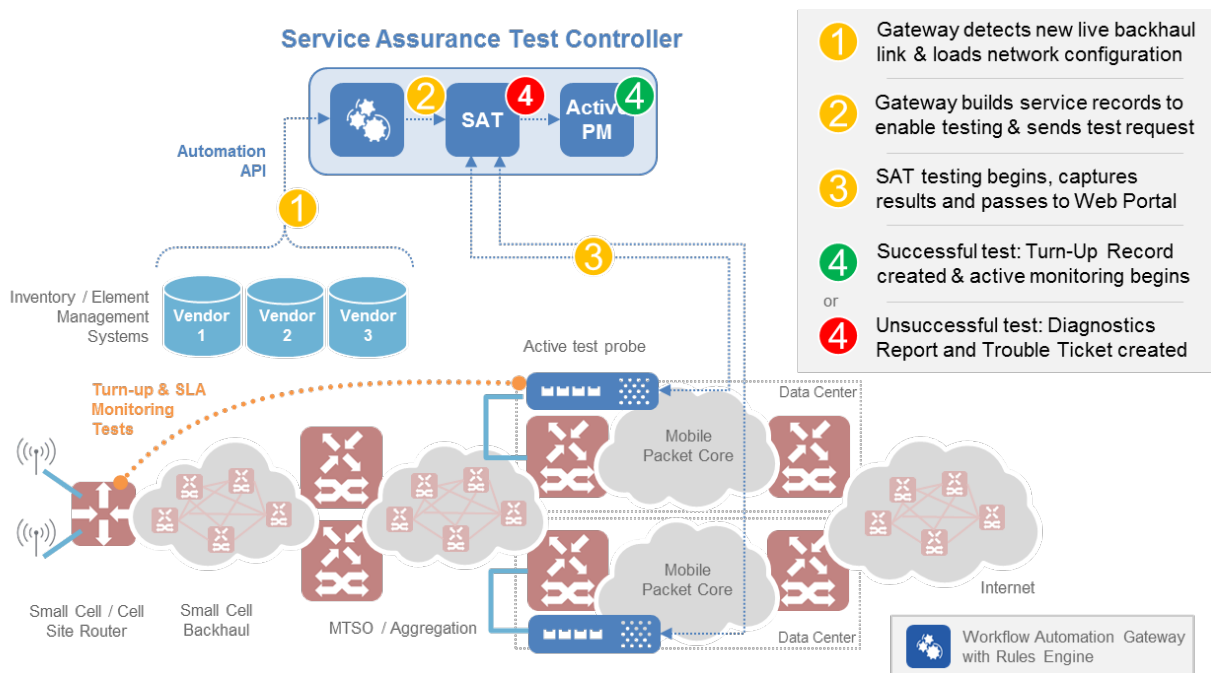


Figure 5 – Turn-up & SLA Monitoring Automation System Architecture

Conclusion: Automation Leads to ROI and Cost Savings

Service providers that have adopted the automation approach and other best practices outlined in this paper have achieved significant benefits including reduced deployment costs, faster deployments and improved service experience.



Figure 6 – Key Benefits of Automating Small Cell Acceptance

Automation of the small cell acceptance process, by leveraging a service assurance controller with tight integration to back office systems, enables providers to avoid hiring any additional staff or contractors, even as they turn-up an order of magnitude more small cell backhaul links. Costs are also reduced because this process is vendor agnostic, allowing for a more heterogeneous network of vendors and ultimately improving efficiency. Integration to back office systems reduces manual testing by 95% which also provides significant cost savings.

Leveraging automation, providers can accelerate time to revenue by activating more than 1,000 small cells per week. This process provides more than a 10x faster time to revenue when compared to traditional, manual processes. Manual processes that require 2 hours per small cell, now only take minutes, allowing new small cells to start carrying revenue-generating services faster.

Small cells provide coverage in critical, under-served areas. They also provide service to areas where congestion on a single tower or a single macro cell has reached a point where the service quality is unacceptable. Enabled by automation, rapid small cell deployments lead to a significantly improved service experience for customers in uncovered or congested areas. The faster a small cell can be turned-up, the higher quality of service the end user will experience. This higher quality of service ultimately leads to increased levels of customer satisfaction and provider loyalty.

Abbreviations

AAV	Alternate Access Vendor
DOCSIS	Data Over Cable Service Interface Specification
MTSO	Mobile Telephony Switching Office
PM	Performance Management
SAT	Service Activation Test
SLA	Service Level Agreement
TWAMP	Two-Way Active Monitoring Protocol
VTA	Virtual Test Agent

Bibliography & References

IETF RFC 5357: Two-Way Active Measurement Protocol (TWAMP), October 2008
(<https://tools.ietf.org/html/rfc5357>)

Technical Specification MEF 48: Carrier Ethernet Service Activation Testing (SAT), October 2014
(https://mef.net/Assets/Technical_Specifications/PDF/MEF_48.pdf)

Technical Specification MEF 46: Latching Loopback Protocol and Functionality, October 2016
(<https://wiki.mef.net/display/CESG/MEF+46++Latching+Loopback+Protocol>)

Ethernet Activation Goes Digital

The next-generation of automated service-activation systems and tools for MSOs

A Technical Paper prepared for SCTE•ISBE by

Mark Gibson

Director of Product Management, Amdocs Network Group
Amdocs
Mark.gibson@amdocs.com

Introduction

Metro-Ethernet activation involves much more than just turning up the connectivity –it’s also about turning up the support and active-monitoring functions of the service (such as fault management) to ensure MSOs can track and monitor the service over time once it’s live – and that’s not something that’s easily performed manually.

Activating a service on DOCSIS might not be simple but at least it offers a standard interface. On the other hand, Ethernet – partly due to its ubiquity –shows up across a range of equipment type, from optical switches to edge routers, each offering a variety of provisioning techniques.

It used to be a choice of TL1 or command line interfaces, but now these are augmented by YANG, XML/SOAP, MTOSI and a plethora of other methods, each aligned to a domain-specific standards organization. And as if this weren’t enough, modern Ethernet services are expected to come with active reporting on the health of the connection, giving in-life updates. In fact, frequent configuration of the monitoring actually results in more commands than the service itself.

If a network is to operate flexibly and support the rate of rapid service change that MSOs want to offer customers, automation is crucial. In a world where customers can press a button to change their service bandwidth, it’s unacceptable that this should result in a network engineer typing commands.

This paper shows how activation must be driven from a service design that reflects both the connection, monitoring, and operational KPIs for the service so that the whole lifecycle is activated.

It also shows how linking a service design to command generation enables quick command generation and removes human error.

Referencing a variety of customer use cases, the paper explains the required features of an Ethernet activation system for a digital customer experience, and also shares Amdocs’ new initiatives in MEF to harmonize network provisioning and reporting.

1. The problem for MSOs

Step into a modern office building and it can resemble something like a science fiction movie from the 90s. Workers sit in front of two, three, or even four screens, looking at colleagues wearing headsets from across the globe, while a live demo whirs with ghostly precision in the background and a chat stream critiques the content in a closed group. Other employees look at social-media postings of colleagues at an off-site meeting, while yet more are at the coffee station reading the latest WhatsApp group video post from their kid’s foreign holidays. Someone, somewhere, is even reading an email.

For customers looking for providers to support these work-environment services, there is great news – in many regions, there are tens of providers who will promise to hook up your premises and deliver you “unbeatable connectivity” at a market-leading price. But once the customer is signed up, keeping them happy and preventing them from looking elsewhere for a better deal critically depends on the MSO showing value to the customers they serve.

There are six main areas that an MSO needs to address to be competitive in the enterprise space:

- A flexible network foundation capable of delivering up to gigabit speeds.
- You must have a large geographical reach so that you can serve a customer wherever they have an office
- The connection that you offer must be reliable and when this is not the case, the customer must be made aware of this quickly
- The solution provided must be future proof and able to adapt to the changing needs and operational models of the customer
- The solution should be capable of outsourcing IT operations of the customer and providing a seamless work environment.
- The solution must be able to adapt to the customer's own environment including the network infrastructure that they operate.

While each of these relies to a greater or lesser extent on an operational process, they all result in a service or a device being activated.

In the following sections, we'll look at how the role of activation varies, and how, in many cases, it's more than just sending commands to a target device.

For customers, it's no longer good enough for an MSO to activate a service at the start of its life, and to only change it when there is planned maintenance. For an MSO to differentiate themselves in the marketplace, they don't only need to provision a service, but they also need to manage and maintain that service over its lifetime, so that they can reduce outages by either reacting as they occur, or predicting them *before* they occur. This then gives the MSO a toolkit that can be applied to more interesting and flexible offerings.

The days when an activation operation pushed commands on to the network and then handed over maintenance almost context-free to the network operations team are gone.

1.1. Strong Foundation

There's little doubt that Ethernet has won the race to be the Layer 2 interconnect of choice. And while as recently as 10 years ago, you might have heard any of a number of rival technologies being mentioned, **in today's networks, there is only Ethernet.**




Figure 1 – Ethernet is now a ubiquitous aggregation technology

Except... that's not quite the case in reality.

Nearly all Ethernet Services are described according to the service types laid out by the MEF (www.mef.net). The customer end of the service is described as the UNI (User to Network Interface) and it is at this port that customer receives their Ethernet traffic. However, there can be many types of Ethernet Service between different UNIs.

For example, in some cases, you see a L2VPN network which only needs edge configuration to enable the port-to-port connection, sending commands to the edge devices of the MPLS network to map the Ethernet traffic onto an MPLS connection.

Anyway, regardless of the technologies, there are many ways to enable the service, ranging from a command line interface (CLI) to a RESTful service-level interface that abstracts the service from the network details. There may even be some flavor of YANG being used.

In many cases, there are multiple generations of network equipment being used to deliver similar end customer services over dissimilar infrastructures – this means that the activation platform needs to adjust to each of the network systems that it finds. But however this is abstracted, it's a complexity that still needs to be addressed, and the orchestration layer above shouldn't have to be aware of how the network is implemented.

MEF's Third Network Vision sets out a candidate blueprint for how Ethernet can be used as the common foundation to support the diverse needs of an MSO and ensure normalized and predictable services which will leave MSOs free to concentrate on the value-added IP services that create more profit and customer interest.

At the core of the vision is the concept of Lifecycle Service Orchestration (LSO), which defines a methodology and practice for managing network connections at a service level. It looks at the whole service lifecycle from modelling the service to activating the network equipment to provisioning the whole support infrastructure needed to maintain the service.

However, as Figure 2 (Source: MEF Forum [0]) sets out, the vision of LSO extends to how you link the services that a typical MSO needs to provide into a coherent managed whole within an MSO. It also

looks at how you manage the co-operation between MSOs to offer global scale offerings outside of their geographic region.

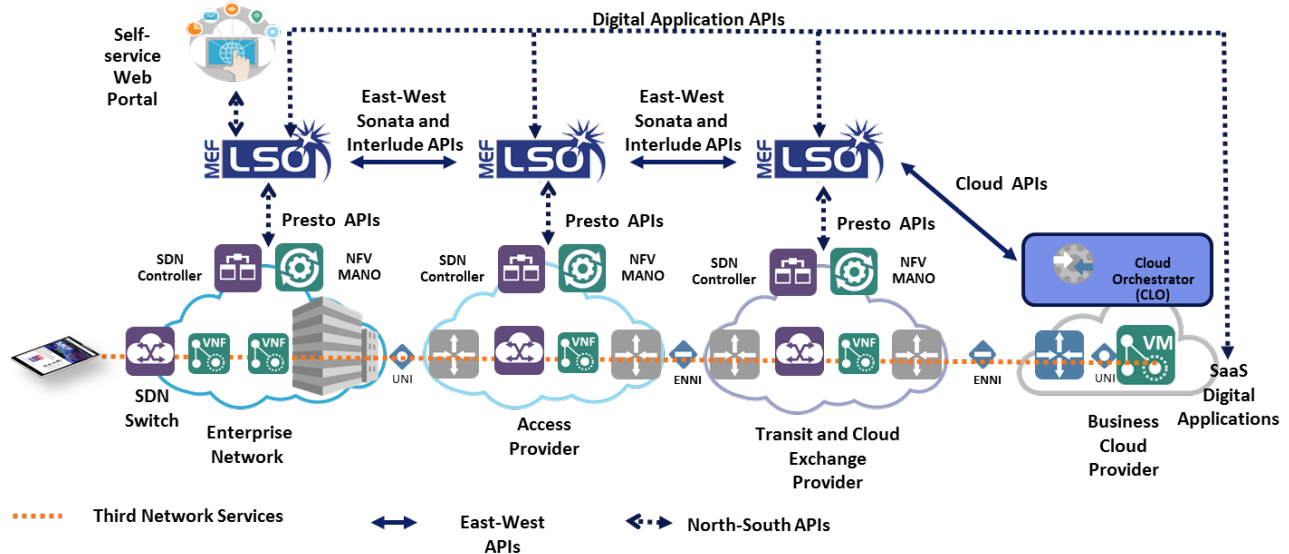


Figure 2 – MEF Third Network Architecture

In fact, in a recent survey [1] **over 90% of respondents said that dynamic, orchestrated Third-Network services will transform the Carrier Ethernet Market, having a significant impact in the time period of 2018+.**

1.2. Geo reach

Few things in life ever turn up out of the blue and just slot in to whatever plan or space you have handy. The same is true of enterprise customers asking an MSO to support a new service, especially when the enterprise is a sizeable one – the larger the customer, the more likely it is that their footprint will extend beyond the geographical boundary of their provider, and the more complicated the service turn-up process will be. A key measure for customers about whether the MSO will be able to support their request is how long it's going to take to get a new service up and running – an MSO that can overcome this first obstacle quickly can differentiate themselves from competitors.

In reality, it's likely that many enterprise customers, especially larger corporations, will require support that reaches beyond the limits of a single MSO (and a highly probable scenario in the US given the traditional geographical divide). But this is also the case in regions like Europe where many enterprises have established offices and data centers in cheaper locations – full end-to-end services can span multiple operating companies within the same provider, who will need to co-operate to provide services.

In the same survey, **over 91% of respondents cited orchestration over multiple provider networks as either a serious or major challenge** [0].

Recently, significant effort by industry bodies has been made to simplify the task of ordering between carriers. The Telemangement Forum (TMF) OpenAPI work stream [0] has developed a number of inter-carrier ordering APIs: Service Qualification; Quote and Service Ordering. This work is now being picked

up in MEF and used to generate industry standards for inter-carrier service ordering using normalized attribute lists and defining predictable outcomes. The first Release of the code (named LSO SONATA R1 after the inter-carrier ordering reference point) took place in July, with the open-source examples shared within the MEF community. While this first release was as much a proof of concept on the APIs themselves, work can now accelerate beyond the initial E-Access service use case to other inter-carrier services.

1.3. Reliable connections

Just because the network now makes it easier for the MSO to turn up and order, doesn't mean the enterprise customer will suddenly lower their expectations of performance and reliability. However, this equally doesn't prevent the MSO from differentiating itself at the point of sale either.

A MEF CE 2.0-certified service [0] will come with support for Operations and Maintenance (OAM) traffic, giving the MSO the ability to measure how the service is operating. This can be compared against the established KPIs for the service variant and allows for alarms and alerts to be generated when the service is non-conformant.

There are two points to consider here though:

The first is how, and when, those KPIs are established and linked to a service.

The traditional Fulfillment and Assurance model has treated them as separate operational stacks with no more than a “ships-in-the-night” passing acquaintance with each other. However, if full lifecycle orchestration is going to be achieved, that gap needs to disappear and the systems merge into a continuous whole. A better solution is that the service model that defines how to design the service instance also describes how to provision its measurement over time as well as the alarm levels it should have.

The second is what to do with the information that has been collected.

When a service has been ordered across operators and boundaries with a partner MSO, it raises the issue of operational information about congestion, packet loss and other OAM information – all of these are highly sensitive business intelligence. And while it might be laudable for an MSO to transparently share information with the purchasing partner about the operation of a leased connection, it is also an unlikely scenario. So setting the KPIs in the ordering process that will be reported against is critical to keeping the customer satisfied. The same is true of a simpler customer-MSO relationship for a more traditional Ethernet Service.

So regardless as to how the information is shared externally, the MSO must have a prepared action plan that is enacted when a critical KPI is hit. This is the final part of the activation process – to link the action plan to each OAM event that is of interest, (even if all that means is to generate an alert that goes into an operational dashboard).

Ensuring that the in-life response is ready when the service goes live ensures that the link between fulfilment and assurance is made.

1.4. Changing customer needs

A customer's needs regarding a connection often change over and the ability of the MSO to support this is becoming less of a “nice to have” feature and more of a service differentiator. This really boils down to one of two operations:

- Adding another service to an existing set
- Making a change to an existing service.

Adding a service will use the same sort of ordering process as the original connection was established but here a couple of alternatives are possible:

For example, when activating on the same UNI, it's important to establish that there's enough capacity to support the new service. The UNI service configuration helps with that, while still allowing the flexibility to support concepts such as overbooking. But this is a service design-level construct in the orchestration layer – endpoint activation is still enabling a VLAN and cross-connection to a connectivity layer such as MPLS.

The trick here is to be sure that the resource is available without creating a burden of reservation and management... Linking the service to a location at the point the order is taken and then linking that explicitly to a specific port often gets the right balance between knowing you can support the new service at order time and avoiding fallout during activation by picking the port too soon. (It also makes tidying up dormant orders simpler).

More interesting is how a customer makes a change to an existing service and the level of direct control an MSO wants to let a customer have over their network service. (And this will mainly depend upon how directly the options offered in the customer portal match against the resulting activation operation).



Figure 3 – Typical customer portal operations

For example, the portal might allow the customer to adjust the CIR by setting the exact rate. Here the command might take a fairly direct path to the activation layer to alter the port configuration. However, even to enable that, the change must be within policy – this means that the portal must only permit allowed values, and the change operation still needs to be represented in the service model and checked against the operational policy too. For example, if there is any overbooking at all on a physical port, the system needs to check that not all customers have turned up the CIR at once.

And what if the customer wants to make a change that's outside of the policy?

Well in this case, you're back to an ordering interaction since this becomes a billable interaction with the nature/level of the service changing. In this scenario, the portal should only offer the customer the options that can be supported at the location, based on the device capabilities. This needs a good integration between the activation layer and the service catalog. This will then be further refined by the BSS layer to constrain the options that can be offered to the customer.

So activation only occurs here once the change is agreed upon, (although for changes which *don't* require new physical resources, the order-to-activation interaction can be very rapid.)

Finally there's a question here regarding where the portal resides. A customer portal does not have to be in the customer domain – one common approach is for the MSO to host a portal that their customers log into to make service level changes. However it is also possible for the service provider to provide APIs on top of which the customer can build their own portals. (And as this allows direct access to MSO-owned devices, security of access is a key consideration).

1.5. Managed service

While giving a customer access to their service so they can make changes is definitely a step up from today's services, they don't enable the MSO to deliver extra value to their customers. If a customer has flexible connections, that customer still needs to run analytics and monitoring between their sites to understand how they are using their resources.

But here is how an MSO might be able to innovate, because rather than having just one connection service, they can offer customer choice of different connection levels:

- Standard Ethernet connection with only major outages being reported.
- Enhanced Ethernet connection where regular reports against KPIs are produced with a live alert stream in a customer portal
- Managed Ethernet connection where the MSO actively manages the service against KPIs and makes proactive suggestions about how better to manage the service. For example, if a customer has two connections and one is regularly hitting a capacity limit while the other remains under-utilized, the MSO could suggest adjusting the capacity limits on the two connections to even this out.

In all of the above models, the MSO could perform active service management, making changes to the network to sustain the service. But in the last case, **the MSO could start to offer proactive service management.**

For example, if a customer has a multi-site system but there is congestion to and from one site, they may choose to increase capacity to that location, while decreasing it elsewhere to provide a good customer

experience. This takes the service towards a closed-loop style of operation where the network layer refers to a higher-layer operational policy for guidance and there is an automated re-activation of the network to match current conditions. Whether this looks like traditional activation (e.g. command line changes to a configuration file), or an operational policy change (to an SDN Controller) followed by monitoring, will depend on the MSO's operational environment.

1.6. Customer environment

An extension to managing the service for the customer inside-out, is to offer to manage the customer's edge devices. This Customer Premise Equipment (CPE) can take many forms, and it is this very lack of uniformity that makes CPE onboarding such a challenge for an MSO. While ordering is moving towards becoming a fully automated process, managing a new set of CPE equipment for a new customer can pose a different problem each time.

For MSOs seeking to offer a full managed service, the time it takes to model and be able to activate CPE devices is one of the biggest delays in a service going-live.

2. Towards a flexible activation solution

Wind the clock back ten years, and IP and Ethernet service activation was divided into a set of clear categories.

- Command Line Interfaces directly to device configuration files
- XML over HTTP interfaces to an EMS
- MTOSI variant of XML over SOAP
- RMI
- CORBA
- TL1

Come back to the present day and... well, in some instances, you still find these provisioning interfaces.

Most MSOs will run equipment that is stable and supporting thousands of services well beyond the end of extended warranty periods, rather than go through the disruption of managing the migration of those services onto new equipment. But obviously, you still need to touch those old devices occasionally.

Network mediation does move on though, and **in this section, we will examine new trends: from new mechanisms to the new operational patterns they enable.**

2.1. Ethernet to the edge

One of the more striking developments of the last 10 years has been the convergence in architecture to the last mile of a customer connection. Ethernet has not only become the de facto standard way to build out a network aggregation architecture but it has also normalized what such a network looks like to the extent that the difference between an MSO and Connectivity Service Provider (CSP) aggregate network diminishes and common toolkits and approaches can be shared.

As a result, MSOs can now benefit from the work in the CSP community to build Ethernet services. The certification process by the MEF (www.mef.net) has led to Ethernet providing a predictable and reliable substrate for a raft of services. It provides the sort of well-behaved underlay that makes operations more straightforward and lets CSPs and MSOs alike concentrate their efforts on delivering the overlay services that generate more money.

All that being said there are still a number of pieces that need to be marshalled to realize seamless operation.

2.2. NETCONF and YANG

NETCONF has been around for a number of years and was available on routers as a configuration method from the turn of the decade. Early data models were variants of CLI commands but started to harden with the introduction of YANG models a few years after NETCONF was finalized.

NETCONF provides a great standard way of managing configuration, with built-in functions around rolling back those activation commands that used to keep activation developers awake at night. These sort of high-level EMS functions for configuration management are now available on individual devices.

However, while the push and pull of commands has been standardized, there has been a huge diversification of the flavors of YANG that are on offer, described recently by Cisco CTO Dave Ward as “50 shades of YANG” [0]. For the activation developer, this means that the problem shifts from one of building out a robust API with the activation target, to one of flexibly modelling the differences in YANG and choosing how much of this “vendor variation” to hide from the service design layer, (and how much needs to be exposed as a feature). In this respect, not much has moved on from the days of Cisco IOS and its multitude of variations.

Most forward looking architectures mention YANG extensively due in part to the fact that it stratifies into two distinct layers. The lower layer looks like a traditional CLI and uses recognizable commands structures for turning up network functions. However, YANG also offers a service layer too, which, while this is not used for devices access, is used as an entry point to what used to be called the NMS layer. This introduces another decision in how to build activation into an overall system.

While YANG is relatively prevalent in IP and Ethernet core switches it is by no means ubiquitous and the activation layer needs to be adaptable to not only manage YANG and non-YANG equipment but to (attempt to) insulate the service design layers from the method used to provision the network.

2.3. Software Defined Networks and Controllers (and Orchestration)

Few terms in modern networking have evolved so dramatically in the last ten years than “**Software-Defined Networking**” (SDN). What started out as a means of decoupling the control plane of networking equipment from the data-forwarding plane now encompasses all means of real-time, near-real-time and offline network traffic engineering. An SDN-controller used to mean that it interacted with devices solely using OpenFlow – now the prevalent use cases are based on a YANG configuration.

For the purposes of this section, we are focusing on SDN Controllers – the modern, generally open-source, replacements for the NMS. The classic example here is the Controller from the OpenDaylight project [0].

To an extent, the SDN-C has changed what activation means in a modern network. There have always been two flavors of activation into a network: device direct connection and NMS-based, but the SDN-C now adds a whole extra level of capability, and builds a new strata layer too. While an NMS built out a local model of the devices that were under its control and could present a service-like model, it generally could only derive attribute values based on algorithms encoded into the system.

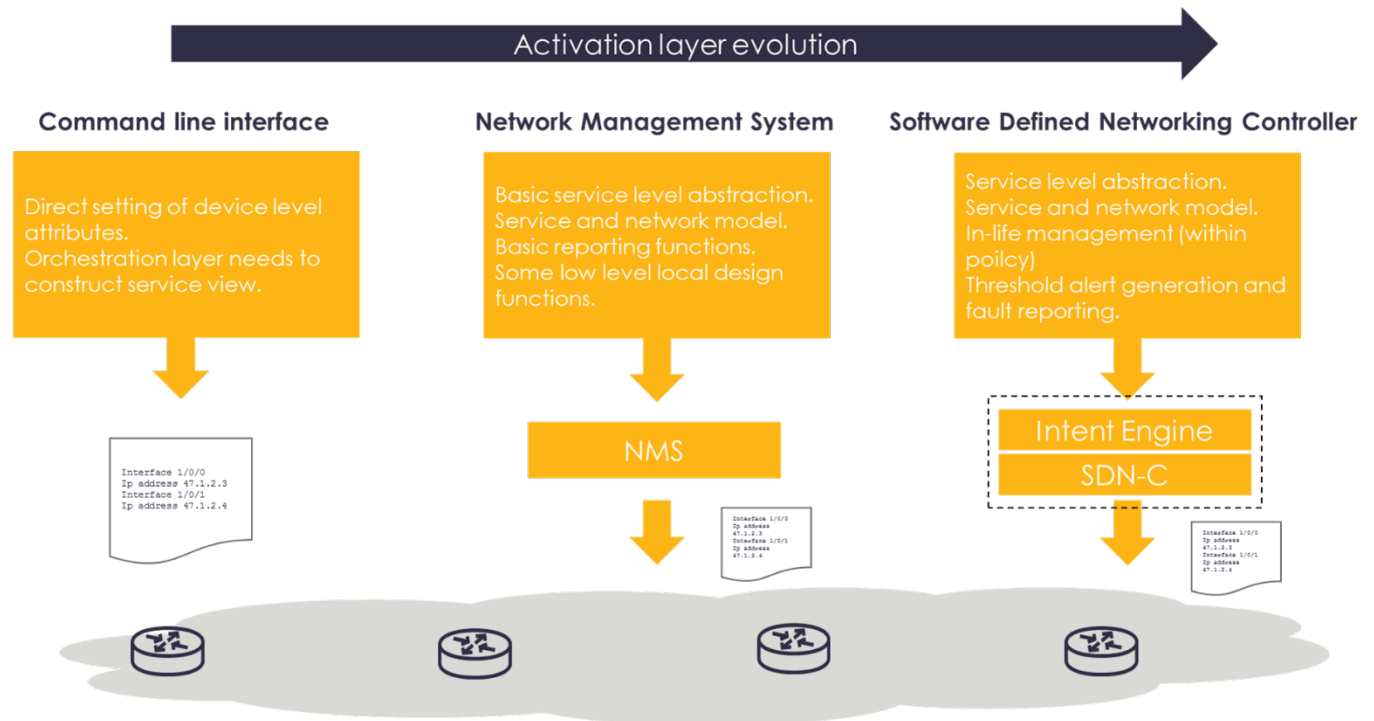


Figure 4 – Evolution of network activation

An SDN-C aims to abstract that further by using policies that indicate intent. In this model, the overlying OSS or orchestrator gives the SDN-C a fixed set of constraints for a new connection, but lets the SDN-C figure out the rest of the implementation detail. The aim here is to let the SDN-C, which has close proximity to the network and a local network model, figure out what's best for the connection over its lifetime and make adjustments as necessary.

However, in practice, the amount of freedom an SDN-C has will depend on:

- How the service is sold in the first place and what the UNI and NNI restrictions are.
- How many options there are in the network for moving the service between resources.

For an Access Ethernet Service on fiber-access network to a PE device on a fiber ring, there may be little practical flexibility.

An SDN-C tends to be able to exert more influence on the traffic streams that are supported by the Ethernet Services. As such, there is a growing market for deploying an SDN-C in the customer environment to make the best use of a leased set of Ethernet connections by adaptively tuning how IP traffic runs across this constrained infrastructure.

The ability to support some of these in-life operations means that most SDN-Cs will tend to also support some local orchestration capability. This is a decision point when deciding how to deploy an overall solution and where you want to build out the orchestration of a service.

So from one perspective an SDN-C is still technically activated, but in this case with a looser more intent-based description of a network connection, described as a YANG service. From another perspective, the SDN-C has now become the activation layer taking in requests to turn up a service and managing the whole process itself. Whatever the positioning, **the SDN-C is an important emerging architectural component.**

But it's not a universal panacea.

There are those who would position the SDN-C as a universal activation/orchestration layer with a single inbound API definition, the idea being here that with some extension modules to support legacy devices, all of your service activation needs can be delivered by a single solution.

In most activation systems, a key measure is how easy it is to add a new device or service type as this is a key dependency in time to market for a new offering. In many regards, the emerging SDN-C market is no more efficient than the existing mature activation market. In some cases, those mature solutions are significantly better having been invested with toolkits and configuration capabilities over the years. **In a hybrid network solution, an SDN-C might not always be the best answer. But there is a balance here between the number of platforms you need to support and the time to get to a working solution.**

2.3.1. Intent in focus

A recent development in the usage of SDN Controllers has been to entirely remove the Intent component of the SDN-C and make it a separate layer. Curiously this makes the rump of an OpenDaylight-based SDN-C resemble little more than a multi-vendor Open Source NMS. It also confuses the boundary of Intent somewhat.

Intent is an abstraction of design and implementation of a service where the specifics of an implementation is somewhat hidden from the way the service needs to operate. This is not really a new concept in Operational Support Systems, indeed often the only fixed point in a service design is the physical port that the customer connects to. All other aspects of the way data is transferred between these points is generally abstracted into a high level specification of the way the ports are connected (point to point, hub and spoke, ...) and a QoS policy that is pre-configured in the network. Many Orchestration systems on the market today use a similar sort of approach too.

In the longer term it may be that Intent is fully re-homed into the Orchestration layer or it may evolve into a separate and separable module alongside path computation. Much of this will depend on the use cases it is asked to solve and whether the separation adds significant benefit.

2.3.2. Other Controllers

Not all network controllers come with an SDN-heritage, (or at least they are perhaps more honest in that regard). Most multi-vendor EMS solutions have nearly all the characteristics of a classic SDN-C but without the OpenFlow support – they are also generally not open source. However, some offer the same sort of service abstraction offered by an SDN-C and many have more developed toolkits for supporting

new devices than are readily available in the OpenDaylight domain. Most seasoned activation vendors will have been challenged on the time it takes to add a new device and those that are successful will have overcome this problem.

2.4. Network inventory and configurations

It is a common mistake to assume that because you have a complete set of device configurations, you also have a network inventory. This isn't the case and can be proven by this simple thought experiment.

Suppose I (as the MSO) have a new card that I want to add to a device (which supports hot plug-in of the card). Once I have added the card, is it automatically configured with the way I want it to operate? No, you have to send the configuration down to run network services on this card.

In an ideal system, your network configuration and a discovered network inventory will sit side-by-side in the same Controller and will be able to be queried interchangeably. This is important as both are important data sources in decision-making prior to activation.

Let's say I'm trying to decide whether I can support a new service at a location. I have two master reference points for making a decision. The first is my top-down network service inventory, which tells me how many designed services I have terminating in that location, and also how many new planned services are waiting to go live – this is my demand. The second source is the network inventory from the activation layer, which tells me the current network situation of used ports – this is my supply. Though note that to get a true view of the supply, the network inventory needs to have looked at the Interfaces Management Information Base (MIB) which tracks all of the ports on a device, not just those that are active.

By subtracting the demand from the supply, you can make an accurate decision on whether you will choose to support a new service at that location. A qualifier here might be that at ordering time you might decide to allow overbooking of resources meaning that you will support a demand in excess of supply. And if you run this query immediately before provisioning, you can have a large impact on activation fallout.

There are other similar processes, but **the point is that having a local and accurate view of both network state in a data model and the device configuration gives the overlying orchestration layer invaluable insight that will make for more accurate overall processes.**

2.5. SD-WAN

Software Defined Wide Area Network (SD-WAN) is a classic example of a disruptive technology in the activation space. There are numerous whitepapers and web resources on what the technology does, but its main value for the MSO is:

- Rapid service enablement by piggy-backing on an existing Internet connection
- Capacity overflow management by handing off excess traffic to a backup connection
- Better customer management by placing priority traffic onto better links and handing off lower priority traffic onto wireless/Internet connections.

SD-WAN helps MSOs set up new connections fast and has quickly gained traction, but on the minus side, it complicates matters in the activation layer. Each SD-WAN offering on the market comes complete with

its own controller and API layer. As SD-WAN has exploded in popularity, vendors in the space haven't had time to wait around for API standards or to build on open platforms.

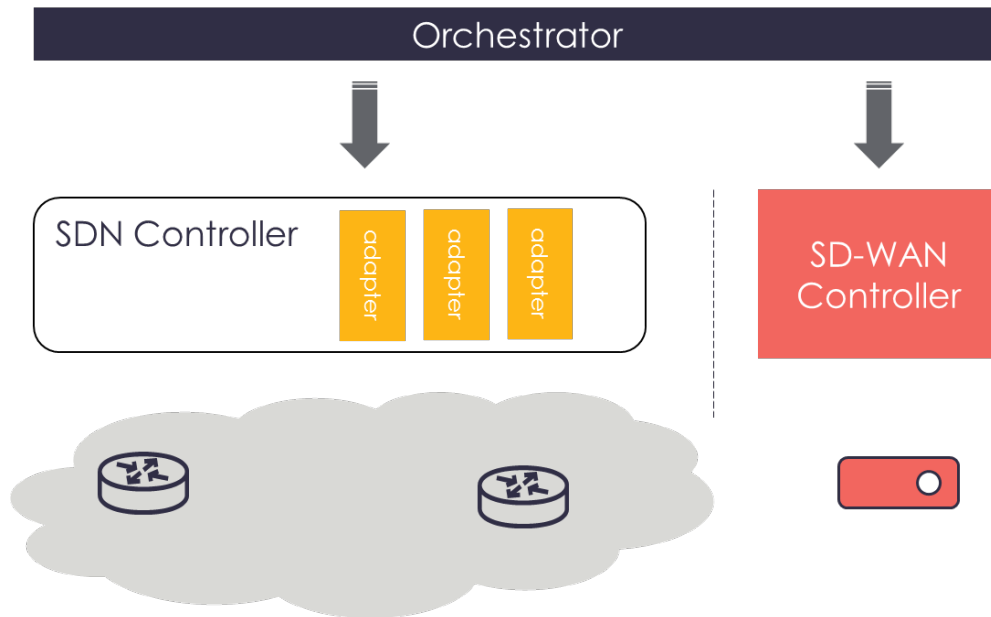


Figure 5 – SD-WAN new vertical

So how do you accommodate such a disruptive technology into an activation layer?

One approach is to position the SD-WAN controller under an overall activation system and build out a uniform northbound API that adds in the new SD-WAN solution.

Another approach is to go against accepted wisdom and add them as a new whole vertical under an orchestration umbrella. The orchestration layer now makes two calls to set up a service – one for the existing network in to the existing activation layer and another into the SD-WAN controller to activate that part of the service. (Note that for SD-WAN, there are likely to be many more interactions to make service policy changes than there are going to be changes to the underlying network connection).

On the surface, for a service like SD-WAN, the decision seems a little simpler since part of the point of SD-WAN is that it acts independently of the provisioned underlying network. In this case, it makes sense to run it in parallel and for the orchestration layer to consume the provided SD-WAN APIs and merge onto a consolidated platform as the need arises. (For example, when your activation provider releases a product upgrade or there is a standard interoperable API that you can converge upon).

In fact SD-WAN is a classic example of how new offerings are introduced and need to be incorporated into a solution space. When a new offering is identified, there is often a race to get a working implementation to market and the success can hinge on being able to provide a full working support infrastructure. In the world of software, this means providing enough of an API to activate/manage a service without worrying where in an existing support architecture this is accessed from. Most SD-WAN solutions ship with their own Controller which is purely responsible for configuring the SD-WAN technology, adding a new vertical that may or more likely may not migrate onto a common controller platform over time. Consider being an IT director making a case to migrate from a dedicated SD-WAN

controller to a common infrastructure without there being some sort of monetizable benefit. The most likely push will be the retirement from support of a foundational software element, rather than any proactive choice.

Architectural purity and alignment is often secondary to market success. And in truth this underlines the need for a flexible orchestration architecture that can flex the point where activation starts to accommodate more service-oriented APIs as a way into the network.

2.6. Virtual Network Functions (VNFs)

There is masses of information available about the way Virtual Network Functions (VNFs) are provisioned and the way the virtual infrastructure that they depend on needs to be managed. **This section will concentrate on the impact of VNFs on enterprise services and on the Ethernet connections needed to support them.**

Prior to the introduction of VNFs, the role of the Ethernet connection was relatively simple. It provided a single IP hop from an IP-enabled CPE device to an IP enabled Provider Edge (PE) device. In most cases this caused the activation system to need to model two device activation operations – one for each device type – to cross-connect the L2 interface at either end of the connection. There were only two target devices.

The separation of IP-forwarding components into VNFs instantly creates a larger number of activation targets, and each VNF is a new target. But there is also new infrastructure that needs to be managed, such as the vSwitch at the entry point of the Data Center (DC) that supports the VNFs. And there are also longer data paths with more IP hops too. The number of IP hops depends on the type of service being delivered and the location of data centers, but one thing is clear, there is not much more work to be done to both activate the VNF elements and to manage their configurations.

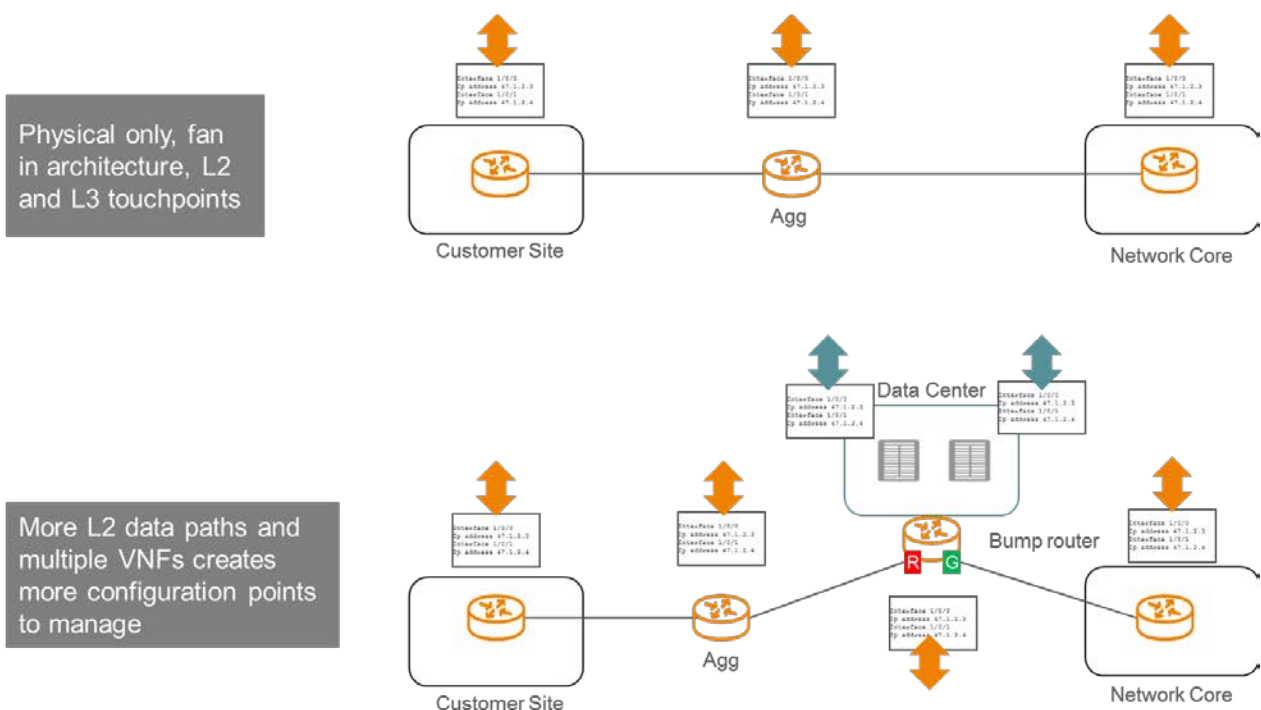


Figure 6 – The impact on network configuration of a simple VNF service

YANG has become the most common language for this configuration, but as noted above you still need a system that can flexibly model each different dialect that you encounter by VNF vendor.

This is also a common operational model where the controller of controllers model is going to be encountered. If there are a group of VNFs deployed in a single DC, then this service chain will be connected together in the DC by an SDN-Controller. The activation system will therefore not only have to describe the individual device configuration, but also the order in which to connect the VNFs so that the SDN-C can manage this.

2.6.1. Live forwarding plane design

As [Timon Sloane](#) noted in at MPLS SDN World 2017, deconstructing the IP components of routers into VNFs forces the forwarding plane to be re-constructed when the VNFs are provisioned. Activation now takes on another responsibility: **it's now not enough to accurately describe the configurations at activation time – you now have to get the service chain correct too.**

While in practice this chain will have been predesigned in the orchestration layer from components that have been loaded into the service catalog following lab trials, it still needs to be accurately passed into the activation layer, and the correct response to a failed activation operation or conditional success needs to be supported.

2.7. Assurance

What constitutes activation in current networks is already somewhat blurred and in some cases there might even be two or more levels of activation occurring. **What's important is that the fulfilment and**

assurance verticals are bridged to give a coherent network service-centric view that enables in-life measurement, adjustment and management.

Managing a network used to be the exclusive domain of the Network Operations team who acted as the gatekeeper to the smooth operation of the network and in some cases, regarded activation systems as just another potential source of disruption to a network. The requirement for activation systems to rigorously confirm “golden configs” for each network service being activated is still a benchmark for activation.

And yes, the Network Operation Centre (NOC) is still critical to a smoothly running service, but **an emerging new trend is that of the Service Operation Centre (SOC), which takes a service-level view and looks at performance and planning, rather than alarms and outage management.**

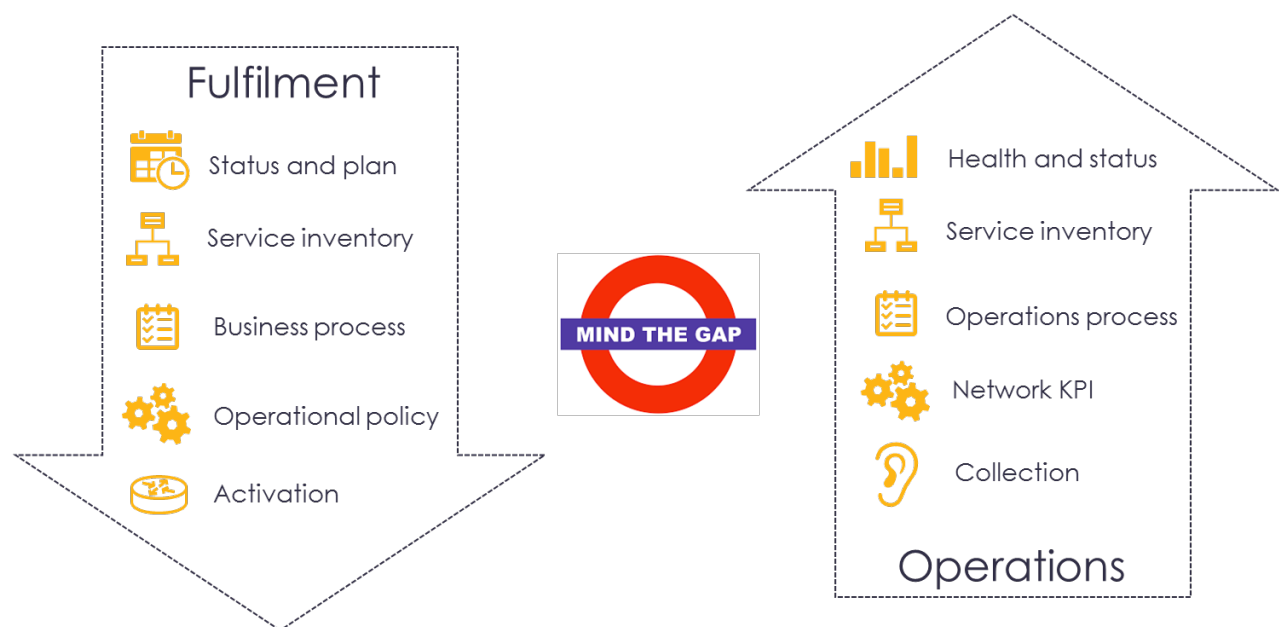


Figure 7 – Fulfilment and Operations gap

For the SOC to prosper, it’s critical that there is an integrated approach to defining the KPIs for each service. This means that activation needs to extend beyond simply activating a service, it must also activate the measurement functions to be used during the service lifetime and link the output alarms, alerts and reports to the operations functions. Ethernet has done a great job of defining OAM tooling for EVCs, and an important part of service activation is to properly link an Ethernet service to its correct OAM function.

While this sort of measurement gives a good indication of overall health, a more specific probe often needs to be enabled to investigate a problem further. This results in more activation work, to both configure the probe and then to toggle it between active and passive states. Note that it’s common for a probe to be deployed, but which is only turned on when accurate measurement is needed.

Though not strictly activation, setting up the SOC also extends to the configuration of the assurance and support infrastructure too, linking the KPIs for the service to the PM (Performance Management) and FM (Fault Management) infrastructure. While the collectors in a PM system will be set up so that they’re liberal in the way they collect data – especially as most assurance systems deploy some sort of big-data

warehousing structure to store incoming information – they are generally conservative in the way they generate alerts from that data. The reason is simple: large networks can generate millions of collected data points – however it’s only those that relate to an actionable event that are of interest to the SOC and often what is actionable is that a number of smaller events have occurred on a specific connection.

Here’s an example: for a given Ethernet connection, if there are five events of 80% capacity threshold crossing within a 24-hour period, this generates an alert for the MSO to investigate the connection since the customer has an SLA KPI that triggers an alert in their portal.

2.8. TOSCA Templates

Having looked at all the tools and levels available to an MSO to manage services, it makes sense to look at a prevalent means of linking it all together. What started as a template method for deploying cloud services, Topology and Orchestration Specification for Cloud Applications (TOSCA) is becoming widely adopted as a means of describing network and hybrid services.

For an IT application it is used to describe all of the artefacts needed to deploy an application, from the software image to the initial day zero configuration of the application. In a VNF context, this usage is extended to cover:

- A topology template that describes the options for deploying the VNFs in the service. This may be as simple as a single rigid example or may allow for context-based decisions. This template also provides links to all of the deployment artefacts needed to enable the service. This might include the initial configuration for standing up a VNF.
- The plan or process for deploying a service composed of multiple VNFs. For example, defining the dependencies and complex sequencing requirements of the VNFs over and above the topology template.

One of TOSCA’s biggest strengths is that it is file-based and thus able to support any configuration file or operational template/method that is in common usage in an organization, though this also rather detracts from driving common standards. However, within a single organization it can harmonize operations effectively and provide a common means of service definition both into and out of an orchestration layer.

While it can support initial configuration, services that are enabled by TOSCA still need to be activated and have their configuration managed over their lifetime using mechanisms like NETCONF and YANG. But TOSCA can crucially group all of the service artefacts together meaning that it establishes a practice of binding the Assurance of a service to its initial deployment in a single definition. Here its flexibility can be a boon to bind disparate components together.

A note of caution however. TOSCA is still in the throes of settling down into a pattern of usage. While early deployments used XML, YAML is seen as the current preferred method of describing human readable templates. Even more confusingly, YANG service descriptions and TOSCA templates are often advanced for the same use cases within standards bodies. However, as most MSO deployments will show some variation based on existing deployed operational components, having a choice of tools at each layer allows for a choice to be made of the best fit.

3. How to deploy the tools you have

This whitepaper has set out a number of tools and related domains for activation and indicated a number of points of flexibility for Ethernet service activation.

This section will provide some guidance on how best to use these elements to build a successful solution (with the caveat that each environment will vary, and what's right in one place might not work as well elsewhere).

Nevertheless, there are a few important decisions to work through:

3.1. Where to orchestrate

Orchestration spans all MSO processes, from order capture and network build, to activation workflow. In general, it's good practice to cut the number of different orchestration systems back to a minimum in any single system so that the number of orchestration platforms supported by an IT organization is minimized. This allows greater re-use and establishment of best practices when developing. It also helps that the user community are familiar with the engine across long and complex processes.

In practice, this generally means centralizing design decisions into the orchestration layer and minimizing activation layer operations for service design, (though leaving in place the orchestration needed to interact with the network).

Moving the design into the LSO layer also makes supporting inter-MSO services simpler. The placement decisions needed to design a service are also needed when determining whether a new service can be supported. Think back to the port supply/demand example. That's the sort of re-usable process block that you want to code once and re-use often.

Which leads us to...

3.2. Unified activation layers or multiple verticals

The choice here is related closely to the orchestration decision. Experience in this area has shown us that while unified activation layers sound great on paper, but they're very hard to build out in practice with long timelines. (And long timelines inevitably mean changing requirements, causing API changes which make the time to deliver a solution even longer).

That's not to say that each technology should have its own vertical. A typical activation layer in a current network would have 3 broad groupings:

- Fixed network activation for L2 and L3 devices. – this might be a “legacy” activation solution
- An SDN-C managing the VNFs and data center activation.
- An SD-WAN controller managing SD-WAN service configuration

Here the orchestration layer will be responsible for mastering the end-to-end service implementation according to the design in the service and network inventory, with each activation component dealing with the specifics of each technology domain. This lowest level of orchestration then effectively becomes the way into the activation layer.

3.3. Common models

It's far easier than it appears to end up with multiple databases in fulfilment, activation and assurance systems.

As a result it's important to draw sensible boundaries between the systems and get information via APIs when needed rather than duplicate locally. It's also relatively simple to spend your entire time merely synchronizing the data stored in overlapping locations rather than building out new solutions. As a rule of thumb you need:

- **A service and network inventory that models the network services that you have created and are managing and any service artefacts that are needed.** UNI port configuration is a good example of this, since it dictates the way that services can be multiplexed together at a single physical port.
- **A network inventory within the controller activation layer** – this will be segmented by the activation domain, and should be queryable by the upper orchestration layers during design-time operations. It should represent the “as-is” network and may include service models for use in discovery and background sync into the service and network inventory.
- **An assurance data warehouse that collects and collates the OAM data that the network generates.** Good systems will pre-process the inbound information and link it into the correct service and KPI, and this database will merge with the service and network inventory so that correlated operational events can be linked to specific service instances. (However, care should be taken to avoid too much duplication).

3.4. Catalog your processes

Definitely another important aspect, **the more you can make your process elements look like small functions in a catalog, the easier it will be to add new technologies or service variations.** And this is where having a common orchestration tool starts to pay even bigger dividends.

Consider adding support for a new VNF to an existing Ethernet service – for a catalog-based environment, once the VNF has been on-boarded into the catalog, you need to develop the new processes needed to support the activation of that component and define the rules for how it can be built into a service chain. Then you can build out a new service chain with the VNF included. Then it's ready to go, without having to re-build whole templates of service design.

4. The digital future

A recent whitepaper by Analysis Mason [0] looked at the key characteristics for a digitally-empowered enterprise in the future and found the following five big trends in the solutions that will enable those enterprises:

- Big Data Analytics, providing service-level information
- Robust inter-company interfaces for interconnection of ecosystem partners and services
- AI-driven software for design of customer networks and services
- Methodologies for structured specification of network and IT services
- Virtualized network slicing for highly customized, dedicated solutions.

There are two interesting sets of trends here:

The first is that many aspects of the solutions are based on making the process of delivering solutions in a simpler way and at a lower cost, which indicates that MSOs will still be looking to drive cost out of supporting an enterprise.

The other is that both the network and the infrastructure it's built upon will be more tailored to the needs of the enterprise, giving specific and managed behavior. What was also interesting in the report is that opinions strongly differ between whether a full managed service is desirable for the enterprise or whether a communication service provider (or MSO) wants to offer the customer full access via a portal.

This is a difference of opinion that has been repeated many times in the last few years. One MSO spent a number of months helping their customers get the best out of variable bandwidth Ethernet connections and avoiding bill shock. Another reckoned that variable bandwidth was just a differentiation exercise that was rarely used in practice.

Where there is more common ground is the need for an MSO to move up the value chain from simple connectivity provider towards all-in-one service and IT manager for a customer. The main demand from this tends to come from the smaller customer, for whom removing the need for an IT department to manage a bespoke service mix takes an overhead away from the business. This leaves an emerging business able to concentrate on building their offering and brand, rather than worrying about whether to use Microsoft Azure or Amazon Web Services to host their IT infrastructure.

While at the moment larger enterprises are still happy to manage over the top, a strong digital focused offer emerging from the management of smaller clients could persuade many to switch across over time.

Take a moment to look back at the vision laid out in Figure 2. This is the sort of offering that all MSOs are striving for and the good news is that many of them are pursuing this in bodies such as the MEF. This shared experience and ability to draw on some of the leading minds in the industry, coupled to the move towards building working prototypes that can be used and refined in practice means that there is a strong roadmap already being laid out. There is also a strong correlation with the work that is going on in the Open Network Automation Platform (www.onap.org), which is a linux-foundation hosted collaboration program born out of AT&T's ECOMP solution. MEF and ONAP have recently agreed to strong joint working practices and architectural alignment. As the code base evolves, ONAP is expected to provide solutions to the problems laid out in this whitepaper.

And new initiatives continue to emerge whose focus is on the operational practicalities. The first phase of a programme to explore the architecture needed to support Zero-Touch Orchestration of Network Service Management, which has developed from a workshop that Deutsche Telekom facilitated, has recently concluded and by the time of publication will have found a new home. This work looks at how to manage the lifecycle of a service automatically and starts to plot the balance between intent and orchestration.

Conclusion

What it means to activate a network has changed greatly over the last ten years. What started out as a device-direct mechanism that sent command line instructions to make changes to the config file of a device has evolved into a service level command set. What used to be a device-by-device operation has now become a service by service method. What used to stop at the point that a connection function had been activated now needs to extend into enabling the support operations for the service.

Likewise, the footprint of Ethernet has evolved. Ethernet used to just be a LAN frame type that could be transported over an MPLS L2 VPN. It has become the uniform means of building a connectivity service in the MSO and CSP space up to the edge of the network as a result of the hard work of the MEF. With the new focus on Orchestration in MEF, with certification of service orchestration planned and a direction that looks towards application management, Ethernet is positioned as the best technology to support the drive towards digital offerings.

Ethernet is, however, only a forwarding technology. For an MSO to evolve into a true Digital Service Provider, a full support infrastructure must be wrapped around the strong base that Ethernet provides. Services must consider from the bottom up how they are to be managed once they have been activated. Initial deployment cannot simply halt once the commands for the connection have been sent, activation must extend to include the measurement and assurance structure and then tie this back into the operational objectives of the service as a whole.

Many industry bodies are looking at parts of this solution space but only MEF is taking the holistic view of services needed to deliver this end-state. In the last 6 months ONAP has emerged as the pre-eminent place for discussing how the services that will be delivered over the top of this common foundation should be delivered. As MEF and ONAP strengthen their direction together, the joint outcome will provide a strong blueprint for the digital future.

Abbreviations

ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers
API	Application Programming Interface
CIR	Committed Information Rate
CLI	Command Line Interface
CPE	Customer Premises Equipment
CSP	Communication Service Provider
DC	Data Center
DOCSIS	Data Over Cable Service Interface Specification (www.cablelabs.com)
ECOMP	Enhanced Control Orchestration Management & Policy
EMS	Element Management System
EVC	Ethernet Virtual Connection
FM	Fault Management
IT	Information Technology
KPI	Key Performance Indicator
L2VPN	Layer 2 Virtual Private Network
LSO	Lifecycle Service Orchestration
MEF	MEF (www.mef.net)
MIB	Management Information Base
MPLS	Multi-Protocol Label System
MSO	Multiple System Operators
MTOSI	Multi-Technology Operation System Interface
NETCONF	Network Configuration Protocol
NMS	Network Management System
NOC	Network Operation Center
OAM	Operations Administration and Maintenance
ONAP	Open Network Automation Platform
PM	Performance Management
RESTful	Representational State Transfer (web services)
RMI	Remote Method Invocation
SDN	Software Defined Network
SDN-C	Software Defined Network Controller
SD-WAN	Software Defined Wide Area Network
SLA	Service Level Agreement
SOAP	Simple Object Access Protocol
SOC	Service Operation Center
TL1	Transaction Language 1
TMF	Tele Management Forum
TOSCA	Topology and Orchestration Specification for Cloud Applications
UNI	User to Network Interface
VNF	Virtual Network Function
VoIP	Voice over Internet Protocol
XML	eXtensible Markup Language
YANG	[not an acronym]

Bibliography & References

Vertical Systems Group “Emerging Third Network Services Enabled By LSO, SDN, NFV & CE 2.0”, State of the Industry Research Report, Jan 2017.

<https://www.tmforum.org/open-apis/>

[https://www.mef.net/Assets/Documents/Carrier Ethernet 2.0 Certification Blueprint -
_VERSION_1_1.pdf](https://www.mef.net/Assets/Documents/Carrier_Ethernet_2.0_Certification_Blueprint_-_VERSION_1_1.pdf)

<https://www.youtube.com/watch?v=cnllFPRyBFs>

https://wiki.opendaylight.org/view/OpenDaylight_Controller:Main

Analysis Mason “B2B User Experience Digitalisation” January 2017.

MEF Forum www.mef.net

Automating Service Creation and Provisioning for Creative Ethernet Services

Delivering Ethernet Service at Scale

A Technical Paper prepared for SCTE•ISBE by

John Hawkins

Product and Solutions Marketing

Ciena

1185 Sanctuary Parkway

Alpharetta, GA

678-867-3331

jhawkins@ciena.com

Introduction

Ethernet service uptake continues on the rise, especially among the small and medium business segments. The increasing interest in high-reliability and low-cost options drive operational challenges with deployment and activation headaches that can only be addressed by comprehensive automation of the network infrastructure in an intelligent and operationally sound model.

The MEF has proposed a Lifecycle Services Orchestration (LSO) architecture that allows for such comprehensive automation when coupled with technologies such as Software Defined Networks (SDN), Network Function Virtualization (NFV) and service orchestration. The approach allows for flexibility in service definition to meet varying qualities of service yet is cost-effective in terms of CapEx and (more pointedly) OpEx.

However, the approach breaks down at scale unless specific considerations are accounted for from the outset. The rise of popularity of the software defined WAN (SD-WAN) is a case in point, where an application drives the adoption of a promising technology whose business case depends on automated deployment and provisioning. Here we explore the technology components that make this possible in the cable operator context.

The Challenge of Scale

1. Unrelenting demand for data

The chairman of India's Reliance Industries conglomerate and the world's richest man recently opined that "data is the new oil." His perceptive remark points out that the world runs on data. Proof of this can be found in many forms. Consider the fact that the largest companies in the world today (by market capitalization) are Apple, Alphabet, Amazon, Microsoft and Facebook. Ten years ago, only one of these (Microsoft) made that list alongside the likes of Exxon Mobil, Petrochina and Royal Dutch Shell. It is no coincidence that these newcomers are also the top cloud data players in the market today and serves as proof that data, and quick access to it, drives not only business and commerce, but modern-day life as we know it.

The prevalence of cloud-based commerce makes for a challenge operators haven't experienced since the early days of telephony (or perhaps mass electrification): how to scale the network to meet the demand for instant, secure and reliable access to all that data. In the case of telephony or electricity, the rollout of those services was taken up country by country, region by region, whereas in this instance everyone wants access to their data, and want it now irrespective of locale or other context. In both business service and mobile backhaul situations, Ethernet and IP-based services have become the instruments of choice to allow for private network builds as well as access to the public Internet for cloud based application support.

Of course, making the business case for the massive infrastructure shift required is not simple. Observe that Operational Expense (OpEx) exceeds Capital Expense (CapEx) in most operator environments (see Figure 1), and owes this to the care and feeding of the growing infrastructure. New technologies have been pioneered to address this fact given that operational efficiencies deliver the best chance for increasing margins and justifying the new investment in infrastructure.

CFOs have mixed feelings about the financial implications. While CapEx allows for better up-front planning of costs and amortization of the investment over an extended period of time, the relatively short technology refresh cycle (as compared to former telephony or electrical technologies) tends to work against predictable budgeting as technology moves faster than operators can digest. In an OpEx environment, at least a portion of the technology refresh cycle can be outsourced to the vendor, who collects a usage-based fee for the latest and greatest technology and sees to it that it maintains currency with the state-of-the-art. At least that's the theory.




Figure 1 - Cost Breakdown for Telecom Operators

But shifting costs from CapEx to OpEx (or from operator to vendor) doesn't solve the fundamental underlying challenges of complexity and scale.

Into this complex mix, technology has been brought to bear to solve the problems it may have itself created. In fact, this challenge of "scaling out" the network is at the root of most of today's technology advances. Fundamentally, Software Defined Networks (SDN), Network Function Virtualization (NFV), Lifecycle Services Orchestration (LSO) all purport to address the high cost of operationalizing an expanding network. Automation is the bottom line and reducing OpEx the main goal.

2. Modern deployment and service commissioning

The processes and procedures employed to initially start up a new service at a new customer site are some of the most labor intensive and costly activities operators face. Thousands of end points must be brought

“on-net” and services provisioned properly and quickly before revenue can flow to the operator’s books. Inefficiency and error can delay revenues, increase customer complaints and lead to competitive disadvantage. The conventional approach involves scheduling several tasks:

- Customer premise equipment (CPE) must be delivered, installed in an appropriate environment, powered up, connected to the network, configured with the appropriate software, commissioned for the appropriate services and authorized to use the appropriate network resources.
- Live communication between the field technician and the network operation center (NOC) personnel is required to ensure both ends are communicating.
- The NOC or technician then initiates a series of service activation tests such as line rate testing, latency, packet loss, etc. to produce a benchmark for the service.
- The NOC then enables the end-to-end service per the order once the back-office tasks of performance measurement and billing have been set up.

One or more customer visits or “truck rolls” are required for all this to come together, and still the process remains fraught with the potential for error and further delay leading to grumbling customers.

To date several techniques have been perfected to address these service startup challenges.

2.1. Zero Touch techniques

Zero touch provisioning (ZTP), or low touch provisioning, has been used for over a decade to auto-provision CPE gear allowing for rapid service turn-up at newly installed locations. While fiber availability or installation is a pre-requisite for any service activation, turning up the service once the media is in place can also be a significant bottleneck if not addressed.

The process involves automating the provisioning tasks listed above. This might look something like Figure 2 and involve the following steps

1. An order arrives from the customer and is entered into the order flow.
2. The NOC confirms the order (payment terms, timeframes, etc) and generates a work order to the warehouse for CPE.
3. The NOC generates the appropriate service profile(s) and a software load with the appropriate configuration script(s). These are pre-loaded to a network server and awaits the device installation.
4. Meanwhile the appropriate device is shipped to the customer or alternatively delivered via truck roll with technician.
5. Once powered and connected to the network, the CPE device requests an IP address via a call to the DHCP server on the private VLAN. Once the device authenticates, the software and configuration scripts are downloaded to the device. Alternatively, a technician can scan a barcode on the device indicating its unique ID, allowing the NOC to initiate the transfer from the server once the correct device ID is confirmed.
6. Service templates are applied once the device is discovered by the NOC.
7. Line-rate service activation testing is completed and a “birth certificate” report can be sent to the customer indicating the service is up and running. Further “report cards” can be issued per the service level agreement (SLA).



Figure 2 - Zero Touch Technique

Automating these service deployment steps results in faster time to market, increased accuracy of provisioning, higher customer satisfaction, minimized training costs for field personnel, highly scalable growth in service turn-up, and ultimately lower OpEx.

2.2. Remote maintenance and upgrades

Maintaining the ongoing health and performance of the service also drives significant OpEx throughout the lifecycle of the service. Given it is far more expensive to capture a new customer than keep a satisfied one, it is critical that these expenses too are minimized by maintaining a high level of customer satisfaction. Maintenance upgrades and feature enhancements can be difficult to schedule and may result in service interruption. While these cannot be completely eliminated via automation, consistent and accurate software patch/upgrade distribution can deliver similar benefits to those from automated turn-up, namely lower incidence of human error, faster deployment and higher customer satisfaction. Similarly, automated system configuration backup and restoration can minimize downtime when errors or network failures do occur and the network state needs to be quickly restored.

To the extent these processes can be triggered, executed, and verified remotely (and admittedly, not all can) the OpEx challenge is minimized and downtime is minimized.

2.3. Orchestration

As packet services have become more complex involving multiple operators, often crossing service territories and frequently involving dynamic behavior such as route selection, virtualized function deployment and variable SLAs, orchestrating all of the moving pieces has been a critical component to facilitate mass rollouts (see Figure 3). Of course with such flexibility of services comes complexity, and complexity usually breeds expense.




Figure 3 - Multi-Domain Service Orchestration Concept

Network management system (NMS), Operations Support System (OSS) and Business Support System (BSS) coordination is one of the benefits of orchestration and much of the reason the technology has captured the industry's attention and generated a significant amount of hype. Automating NMS functions can enable significant cost savings, as well as adding considerable capabilities to operators' OSS/BSS systems, but should be regarded as a journey, not a destination. So far, large scale deployments have been restrained and lessons learned are being carefully watched by the industry as a whole. MSOs are among the most aggressive proponents of the orchestrated service approach.

The orchestration approach relies on a single point from which devices can be provisioned, monitored, and troubleshot. The benefit of visualizing the entire network from a single point of administration cannot be overstated. By collecting the network state in this manner, the network is better documented, making troubleshooting tasks shorter and more straightforward. This ultimately impacts metrics like Mean Time to Insight (the time it takes to correctly diagnose and triage new issues) and Mean Time to Repair.

Not only is the visibility to network bottlenecks and trouble-spots improved by this approach, configuration files and scripts can be greatly simplified, minimizing the possibility of human error which is the single largest source of network downtime.

2.4. Comprehensive model for automation

Evolving from existing and legacy networks to automated networks is among the primary hurdles many organizations are currently addressing. Few have the luxury of deploying technology in a greenfield environment. However, most are at least evaluating automation via orchestration in limited environments. A select few have implemented orchestration in production environments. Those early adopters are beginning to realize the benefits, as well the challenges involved when doing so at scale.

So far we have only implied the use of today's hottest network technologies: Software Defined Networks (SDN) and Network Function Virtualization (NFV). These are natural corollaries to the orchestrated network as they provide the ability to control the behavior and functionality of the network from a logically centralized controller allowing for considerable flexibility and cost-containment. Individual devices (whether core, edge or CPE) need not be as expensive and high-functioning as before. End-to-end service attributes can be pre-staged, verified against network policies, and pushed to the network as needed.

Ultimately the nirvana state for this virtualized, software-centric infrastructure is the so-called "self-driving network." The vision is one in which high-level software entities can address the network directly via well-defined Application Programming Interfaces (APIs) that allow access to key network attributes and functions. These applications are thereby empowered to request network resources (capacity, end-points, service types, performance tiers, etc.) and receive responses from the network as to whether the request can be fulfilled and at what cost. These interactions can be between consumers and service providers (in the case of a retail transaction) or between operators (in the case of a wholesale transaction). In combination, service providers can thereby "stitch together" end-to-end services using a combination of on-net and off-net infrastructure. While our focus here has been on the effect of automation on OpEx and CapEx, this has the effect of adding considerable value to the service provider's product offerings and hence revenue streams.

3. The MEF LSO framework

In this environment of rapid technology evolution, standards organizations such as the MEF (formerly Metro Ethernet Forum) and the TMF (formerly Tele-management Forum) have stepped in to offer a framework for network automation (see Figure 4). As we stated earlier, automation is not a new concept, but corralling the various industry efforts revolving around SDN and NFV including proprietary and open-source projects has been a challenge so far. The MEF Lifecycle Service Orchestration (LSO) reference model (MEF 55) is particularly useful and forms the basis for our discussion here.

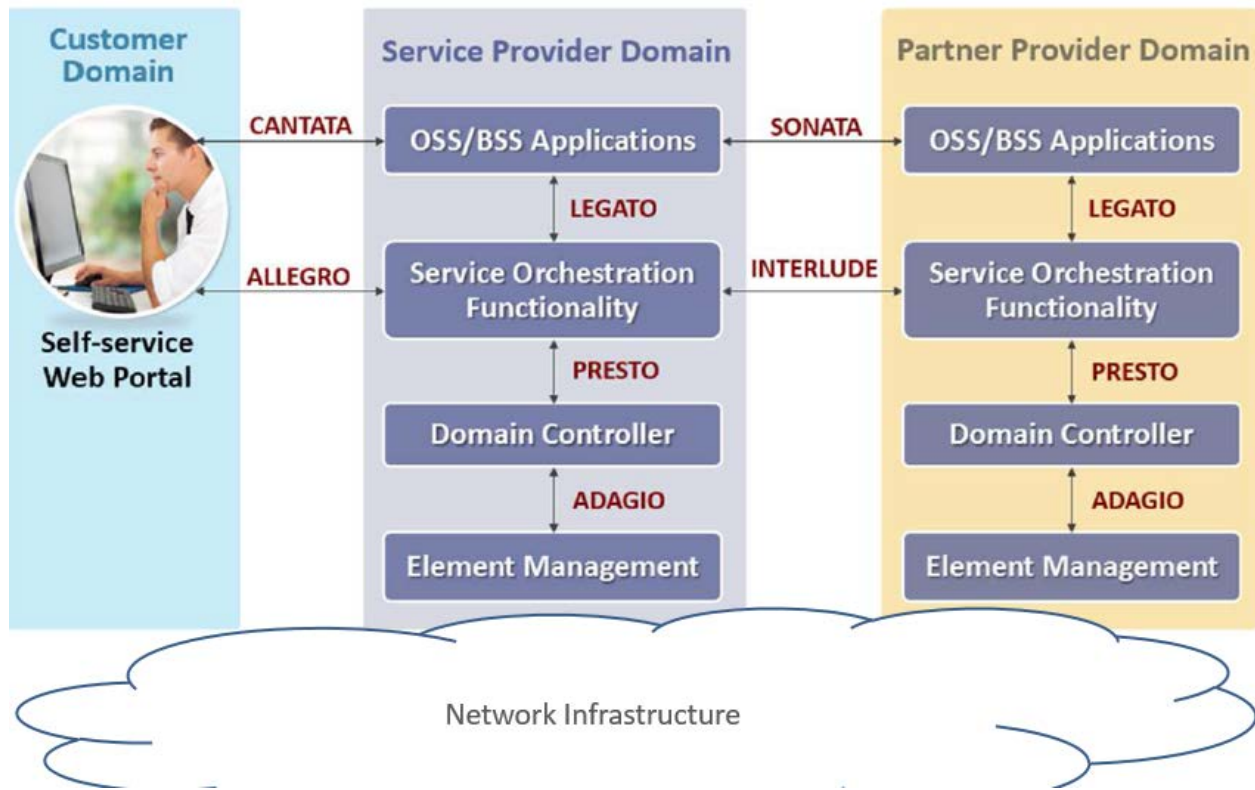


Figure 4 - MEF LSO Reference Architecture

3.1. Standards-based interfaces and operational processes

In order to facilitate interactions between customers, service providers and partner providers, the model defines each of these as individual domains, each with specific interfaces to allow a request/response between them. These interfaces (cleverly named for musical terminology) along with the APIs that give them their functionality, are in early stages of development within the Technical and Operations Committee of the MEF, but already useful early implementations have been floated in the industry via proof of concept and “hackathon” activities involving multiple operators and equipment vendors. Specifically, multiple demonstrations of the Presto interface between service orchestration and multiple domain-specific controllers have been demonstrated, creating MEF-compliant Carrier Ethernet services across more than one operator (also using the Interlude interface).

The ultimate goal is to allow for certified MEF services to be created using these interfaces along with value-added virtual functions. Layering additional managed services via NFV and orchestrating these such that they can be deployed, updated, monitored, and managed by the service provider opens entire new avenues of revenue than can compensate for the commoditization of the bandwidth-only service. Several use cases have been proposed by the MEF and are reviewed in section 4.

3.2. Standards-based connectivity services

The basis for any managed service is the connectivity service itself and the MEF service types continue to enjoy good market penetration worldwide. E-Line, E-LAN, E-Tree, E-Access, and E-Transit service types

are well understood by operators and a growing number of enterprises who increasingly demand them. These provide options of native Ethernet connections, either private or virtual that can address capacity, QoS, and resilience, requirements for access to data centers, branch offices, and the Internet itself. Given the challenges of access and metro network connectivity, along with the comparative expense of IP/MPLS services, these Ethernet based services have become the workhorses for those building networks to access cloud-based data.

Their value lies in the fact that these are standards-based, well-understood and increasingly widely-available services. In the past, a TDM-based service such as ATM or Frame Relay would have been the mainstay of enterprise connectivity. They were well understood, consistent in behavior and performance, and very cost-effective. Today, they are challenged to deliver the bandwidth required by modern applications and have largely been replaced with IP-based alternatives such as IP-VPNs based on MPLS technology. But further growth in bandwidth requirements and increasing complexity of those protocols have led to a resurgent interest in Carrier Ethernet services.

4. Ethernet business service use cases (a sampling)

4.1. Connectivity +

Operators are interested in introducing new value-added services that can command higher revenues than simple commoditized connectivity. Doing so at scale can lead back to complexity, and therefore requires careful planning and considerable automation to avoid jumping from the pan back into the fire. There is a need to rapidly introduce new services to market, test their viability (both economic and technical) and then scale out the network.

NFV based services allows for a software-centric model of doing just that. This “Connectivity +” or Virtual Network Function as a Service (VNFaaS) business model is in its infancy and is made possible by the holistic view of the lifecycle of the service now well captured in MEF standard MEF 55. Concepts such as “branch-in-a-box” are being tested in various markets with connectivity services bundled along with firewalls, encryption, virus scanners, WAN optimizers and the like. With this distributed NFV model, service providers can provide ongoing, subscription-based updates to the software, and avoid separate and costly appliances otherwise required for each service. The generic processor complex required for NFV also allows a future opportunity for added value services to be upsold for added revenue opportunities.

4.2. Wholesale Network as a Service (NaaS)

Figure 5 shows an example of an MEF E-Access service using two operator networks to access a datacenter in an “out-of-footprint” scenario. While the service provider cannot offer direct access to the desired datacenter, an E-Access service from the wholesale operator can complete the end-to-end connection. These wholesale arrangements are not uncommon. The reality, however, is that such arrangements are labor-intensive given both business and technical perspectives. Automating the interactions based on pre-arranged business-driven agreement can drastically reduce the time to provision and turn-up the service using the MEF Sonata interface and associated API. Reducing provisioning from weeks to hours, perhaps minutes makes such a service viable at scale and highly valuable to end-users large and small.

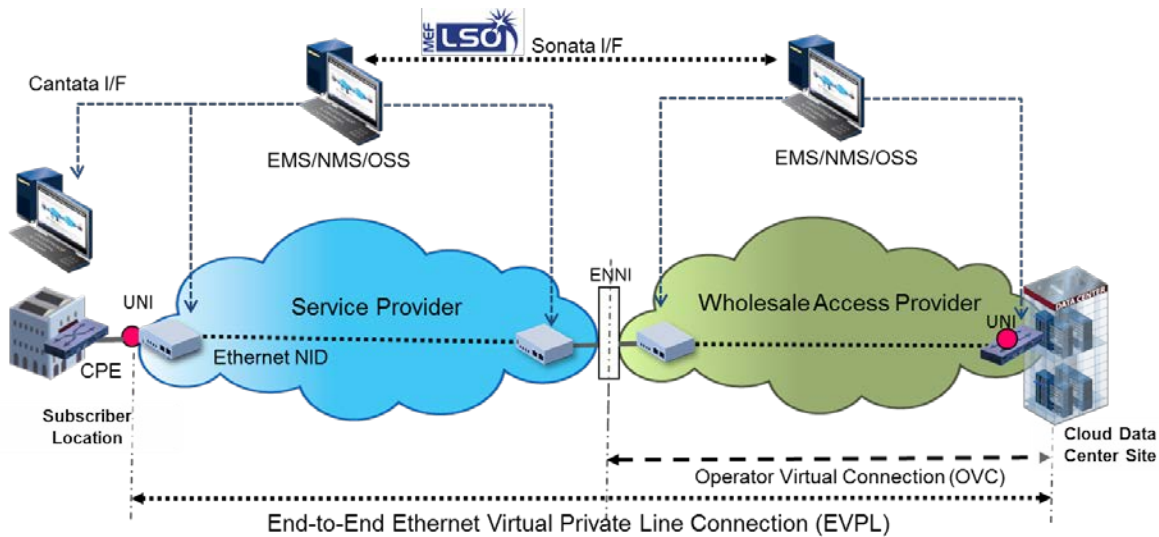


Figure 5 - Multi-operator Datacenter Access Service

4.3. SD-WAN

The software defined WAN (SD-WAN) is thought to be an even more compelling use case for enterprises wanting to connect branch locations to their vital sources of data. SD-WAN is intended to optimize the use of multiple links to a branch office based on policy and best-practice. For example, web access might be provided using a low cost Internet access links, whereas more vital IP telephony or video conferencing traffic might be steered towards a fiber-based, protected and private E-Line service. Wireless connections may also be leveraged in the event of poor performance or outages of landline services. Policy may also dictate that different connections be used at differing times of day to take advantage of cost savings or congestion avoidance.

A more ambitious implementation of SD-WAN envisions a centralized controller using SDN techniques selecting actual routes to minimize latency or jitter performance of a given link for a given performance target. This would involve active and dynamic management of the operator's available resources and would require sophisticated automation. From the point of view of the operator, path selection is already in use to route around network hot-spots (or downed links/nodes) so extending the value to the customer is an additional opportunity for incremental revenues.

4.4. Cloud exchange

Leveraging multiple services from multiple cloud providers offers another use case whereby an operator becomes a broker of cloud services. Similar to SD-WAN in application, multiple cloud services (such as Amazon, or Microsoft's cloud offers) can be provided via a single connectivity service via such a brokerage arrangement. While the network architecture for such connectivity (see Figure 6) is facilitated by MEF service constructs (end-to-end E-Line connectivity stitched together from multiple operator virtual connections), the orchestration of the services in an automated fashion relies on interactions between operators via the MEF Sonata API.

It is important to define methods for assembling services across multiple administrative domains in this model. Not only is interconnecting networks for transiting traffic critical, but also the interconnection

business resources to support the network services. In this way, the assemblage of resources appears to the end user as a single, end-to-end service, dedicated to his use and billed by a single entity, the service provider.

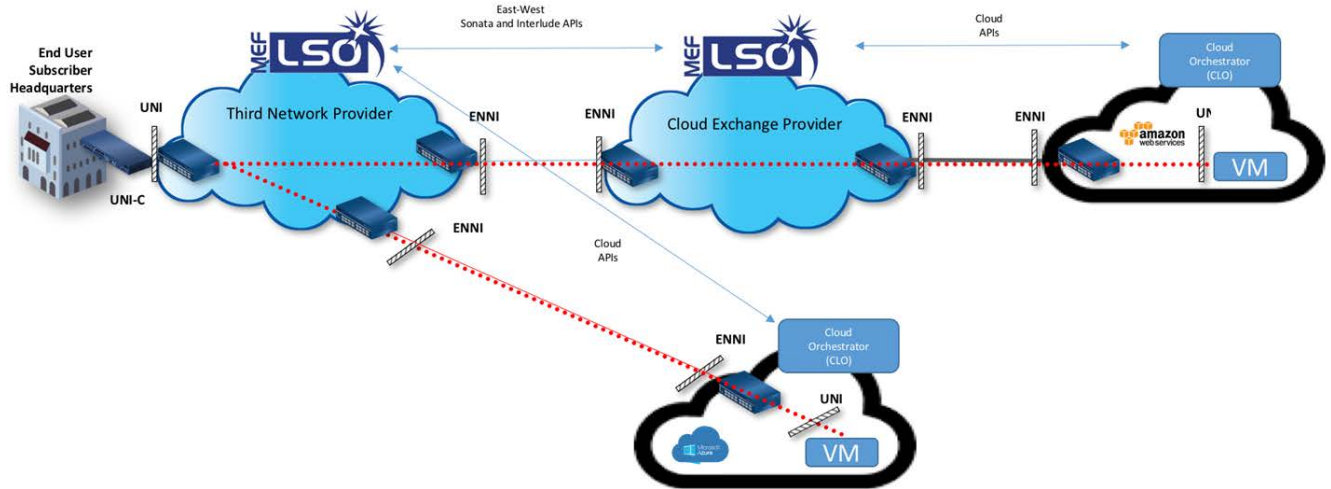


Figure 6 - Cloud Exchange Access

Conclusion

A seismic shift is underway in today's networking market. Connectivity services, while vital to the efficient operation of modern business, is becoming a commodity. Layering services in addition to connectivity is seen as a viable and valuable approach to deliver what end-users really want, however the opportunity for simply layering complexity onto an already complex environment constrains its rollout at scale. Techniques such as zero-touch-provisioning, service scripting, SDN/NVF, and lifecycle service orchestration address this complexity so these value-added services can be rolled out en-masse without breaking the service provider's business case.

Abbreviations

BSS	Business Support System
CaPex	Capital Expense
CE	Customer Edge
CEN	Carrier Ethernet Network
CPE	Customer Premise Equipment
EMS	Element Management System
EPL	Ethernet Private Line
EVPL	Ethernet Virtual Private Line
LAN	Local Area Network
OpEx	Operational Expense
OSS	Operations Support System
MAC	Media Access Control
MEF	MEF Forum (formerly, Metro Ethernet Forum)
NFV	Network Function Virtualization
NFV-O	Network Function Virtualization Orchestrator
NID	Network Interface Device
NMS	Network Management System
SDN	Software Defined Network
SLA	Service Level Agreement
UNI	User-Network Interface
VLAN	Virtual Local Area Network
VM	Virtual Machine
WAN	Wide Area Network

Bibliography & References

Shane, D. (2017, July 31). Reliance's Big Data Play. Barron's. Retrieved July 31, 2017, from <http://www.barrons.com/articles/reliance-makes-big-data-play-with-cheap-mobile-1501303073M>.

Walker, "A Growth Opportunity for Vendors: Telco Opex," OVUM, Oct. 31, 2012.

"Central Office Re-architected as a Datacenter (CORD)." (n.d.). Accessed June 01, 2017, from <https://wiki.opencord.org/pages/viewpage.action?pageId=1278047>

"Network Functions Virtualization — An Introductory White Paper," *SDN and OpenFlow World Congress, Oct. 2012*.

Berde, P., M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, and G. Parulkar "ONOS: Towards an Open, Distributed SDN OS." *ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN 2014)*. August 2014.

MEF Forum, *Service Operations Specification: Lifecycle Service Orchestration (LSO) Reference Architecture and Framework*, MEF 55, March 2016

2007.Mehta, Kumar. "The Evolution of SD-WAN to SD-Branch." *SDx Central*. July 14, 2017, Accessed July 15, 2017, from <https://www.sdxcentral.com/articles/contributed/evolution-sd-wan-sd-branch/2017/07/>

Tarplin, Johnathan. "Can The Tech Giants be Stopped?" *Wall Street Journal*, July 15, 2017, Review sec.

Whelan, Greg. "How the CORD Project Will Change the Economics of Broadband." *Linux.com*. August 10, 2016. Accessed June 1, 2017, from <https://www.linux.com/blog/how-cord-project-will-change-economics-broadband>

Aitor Gutierrez, et al March 31, 2015. "Value-based network management for telecoms: Building big-data capabilities to optimize capital and operating expenses." Accessed June 01, 2017, from <https://www.strategyand.pwc.com/reports/value-network-management-telecom>

Cell Backhaul – Building The 5G-Ready Network of The Future, Today

A Technical Paper prepared for SCTE•ISBE by

Jon Baldry

Director Metro Marketing

Infinera

125 Finsbury Pavement, London, EC2A 1NQ, UK

+44 7766 146 440

jon.baldry@infinera.com

Introduction

5G promises huge advances in wireless technology, with its greater bandwidth and lower latency that enable a wide array of new services and applications. Radical changes will be required throughout the network, from handset design to the architecture deep into the core of the network. 5G will bring significant challenges to the transport network as well, and also adds a degree of uncertainty as the 5G standards are still under development.

Nonetheless, mobile operators and MSO-based wholesalers who provide them with mobile transport services want to be able to evolve seamlessly from today's 4G-based fronthaul and backhaul environment to a future 5G environment while addressing evolving transport network requirements, including:

- Even higher demands on performance – Low latency, synchronization and higher capacity demands are a given with 5G.
- Ethernet evolution – 5G fronthaul will migrate to Ethernet, creating a hybrid fronthaul/backhaul environment sometimes called midhaul or crosshaul (X-haul). But Ethernet needs to adapt to support this new environment.
- Seamless coexistence of 4G and 5G - Whereas the transitions from 2G-3G and then 3G-4G were totally separate networks, 4G doesn't go away with 5G. 4G infrastructure remains a key element in the new network, which must coexist with the new 5G infrastructure.
- Virtualization of key network resources - The move to a software-defined network (SDN)-controlled and cloud-structured environment will help facilitate support for mobile edge computing (MEC), fog networking and virtualization of key network resources.

To address these challenges, network operators need networks that are flexible and open, and offer high performance. This paper will describe the challenges associated with the migration to 5G and show how MSOs must evolve their transport services to adapt and grasp the exciting opportunity that 5G presents to the industry.

Content

1. Evolution from 4G to 5G

5G promises huge advances in bandwidth and network performance that will enable an array of new mobile services and applications. The terms “4G” and “5G” are quite broad and cover a range of releases within the plan of the 3rd Generation Partnership Project (3GPP), which is the global wireless standards group that unites the standards development activities of the primary seven standards organizations within the wireless world. As a generalization, the term 4G covers 3GPP release 8 which covered Long Term Evolution (LTE) through LTE-Advanced (LTE-A) in releases 10 and 11 to release 14 which will be standardized in 2017 and prepares the ground for 5G networks. The term 5G covers release 15 onwards which is designed to support the new 5G requirements. 5G standardization activities started in the 3GPP in 2017 and it is anticipated that Phase 1 of the 5G spec will be standardized in R15 in the 2nd half of 2018 and Phase 2 will be concluded in R16 in late 2019-2020.

One key aspect of 5G networks is the of co-existence with 4G infrastructure. This differs from previous network transitions where 3G replaced 2G networks and 4G replaced 3G within a cell site once the

network was upgraded. We therefore need to understand current 4G trends to understand how transport networks will be required to support 5G networks in the future.

1.1. Current advances in 4G

4G networks brought significant advances over 3G networks and have continued to evolve since the introduction of the Release 8 standard in 2008 and the first commercial LTE services 12 months later by Telia Sonera in Stockholm and Oslo. The standard has evolved from R8 through to R14 with many advances including the extension from macro cells so a variety of small cell options and the inclusion of additional advanced functionality such as enhanced intercell interference coordination (eICIC) and coordinated multipoint (CoMP). These capabilities have enabled LTE/LTE-A networks to extend to heterogeneous networks (HetNets) containing overlapping cells of various sizes simultaneously working together to support each end user devices in a coordinated manner.

The ability to support better simultaneous interaction between multiple cell sites will be critical in future 5G networks as 5G cells will be smaller than current 4G cells, so more cells will be needed to provide coverage and this will create a mix of coverage from new 5G cells and “legacy” 4G cells. To put the relative sizes in perspective, if 4G cells are measured in kilometers then corresponding 5G cells will be measured in hundreds of meters.




Figure 1 - 4G and 5G Cell Coexistence. Source: Infinera

1.2. Migrating to Cloud-RAN Architectures

A significant trend within wireless networks in recent years has been the adoption of centralized-RAN and cloud-RAN architectures, both commonly abbreviated to C-RAN. These architectures take advantage

of the migration of the interface between the remote radio head (RRH), which connects to the antennae, and the baseband unit (BBU), which creates the radio frequency (RF) signal that each antennae/RRH transmits, from coaxial cables to fiber optics. This was initially done to reduce power consumption and cost but also then created the opportunity to move the BBU out of the cell site to a centralized “BBU hotel” where several BBUs could be co-located. The connectivity between the BBU and RRH then requires a new mobile fronthaul network which uses a protocol called common public radio interface (CPRI) to carry a digitized version of the analogue RF signal between the BBU and RRH. Some deployments use a very similar protocol called open base station architecture initiative (OBSAI).

In the C-RAN architecture the BBU still requires a backhaul link to provide connectivity to the evolved packet core (EPC) network, known as the S1 interface, and BBU to BBU connectivity over what is known as the X2 interface.

The migration of the BBU to a BBU hotel creates the initial step in this architectural shift with a centralized-RAN where the BBUs are centralized in a single location. This reduces power and space requirements in the cell site and makes BBU to BBU communications via the X2 interface easier, thereby assisting in advanced functionality such as CoMP or eICIC.

The next step in the architectural shift is to consolidate these collocated BBUs into a single larger BBU that can work across a collection of RRHs and cell sites to create a true cloud-RAN. This may be a larger BBU with the processing power to consolidate several BBUs or more likely a virtualized BBU (vBBU) in a network functions virtualization (NFV) environment, which effectively turns the BBU hotel into a mini data center.

Adoption of C-RAN varies around the globe, which Asia taking the early lead. There are two primary reasons for network operators:

1. Network economics and environmental reasons. The first C-RAN deployments that required mobile fronthaul networks were largely driven by economic reasons. This is well documented in the China Mobile Research Institute’s white paper “C-RAN - The Road Towards Green RAN”. In these cases, there was a clear economic business case centered around reducing power and space requirements in the cell site as BBUs were moved to the BBU hotel. This had the additional benefit of overall power consumption reduction which led to a reduction in the carbon dioxide footprint of the network.
2. Preparation for LTE-A and 5G. While the business case outlined above works in some regions of the world, due to differing economic and commercial factors it hasn’t really been viable in some regions such as North American and Europe. These regions are therefore behind Asia in terms of C-RAN and fronthaul deployments but many operators in all these regions are now looking at C-RAN and fronthaul as a mechanism to support some of the advanced functionality introduced in LTE-A and 5G that need better real-time coordination between cell sites.

1.3. New advances in 5G

Future 5G networks have the promise of considerable improvements in network performance with a drop of latency from 10 milliseconds to 1 millisecond and an increase in throughput to support services in the order of 1 gigabit per second. To support these performance advances significant changes will be required in the overall architecture of the network. The first change to consider is the “cloudification” of the mobile network. 5G will require sophisticated coordination between cell sites and will therefore require a C-RAN like architecture where the BBU hotel is essentially a mini data center. One of the principals of

the 5G network is to allow the network to offload functions from simpler wireless devices that can then preserve battery life or enable more complex operations on simpler, lower-cost hardware.

However, 5G networks will take the cloud model one step further with mobile edge computing (MEC) where a shared compute resource is placed at the edge of the mobile network, most likely in the same location as the vBBU. There is a very similar but subtly different architecture called fog networking where the cloud of compute and storage resources are distributed between a range of data centers from the vBBU location at the edge of the network, centralized core data centers and other mini data centers in other intermediate locations. MEC and fog will enable the network to provide a better user experience and also can optimize the network resources to match the performance requirements of the application with low latency applications using resources closer to the user and more latency tolerant applications using lower cost centralized resources elsewhere in the network.

The 5G architecture also utilizes network sliceability to create dedicated pools of network resources to create separate domains over the same physical infrastructure. Sliceability will enable applications to sit above a sliced control plane and sliced forwarding plane as if they were discrete networks with differing performance metrics. Operators can therefore carve out a slice of the network dedicated to a function, such as supporting autonomous vehicles, knowing it has the necessary resources, or for a network to be logically shared between multiple operators.

To support the trends outlined above 5G will be built using software-defined networking (SDN) control and NFV will play a major role in the optimization of the network. 5G networks will leverage the structural separation of HW and SW, as well as the programmability offered by SDN and NFV.

2. Protocol changes required for 5G mobile transport

In order to support the performance improvements and architectural changes required for 5G, the mobile transport network will undergo significant changes. Firstly, fronthaul and backhaul will merge into a single midhaul or crosshaul (X-haul) network that will be able to support fronthaul-like and backhaul-like transport over a common network. This will require adaptations to the current Ethernet standards to support this new environment, such as the ability to understand and respond to the latency sensitivity of traffic. Standard Ethernet switches traffic and makes decisions over which packets to forward or store based on the priority information so work is currently underway to add the ability to consider latency sensitivity to this decision making to ensure that applications needing lower latency can be prioritized correctly. This is known as time-sensitive Ethernet.

In parallel to this standardization work, the 3GPP has proposed a model covering the possible blending of fronthaul and backhaul capabilities with eight options for functional split between the processing in the distributed unit (DU) that will replace the RRH and the central unit (CU) that will replace the BBU in a fronthaul scenario supporting the vBBU and other NFV functions. At one extreme, there is a fronthaul-like split, which is essentially the transmission of a digitized version of the analog RF signal. This results in minimal processing in the DU at the cell site, but will have latency limits similar to CPRI based fronthaul today and requires significantly more bandwidth than other options. At the other extreme, there is Ethernet-like transmission and there are also many possible options between the two extremes. 3GPP is discussing how to take this model forward.

In addition to the changes in the protocol that mobile transport networks will be required to support, the changes to 5G performance will also have a knock-on effect to synchronization and latency requirements.

While the exact details are yet to be determined, it is clear that these performance parameters of the transport network will become even more critical than they are today.

3. Requirements for 5G-Ready mobile transport

As you can see, we are currently in pre-5G period of uncertainty. We understand the big picture of how 5G networks will be architected but we don't know the exact details yet and we won't for some time. This creates a challenge for wireless operators and wholesalers such as cable MSOs who sell connectivity services to wireless operators. These operators need to be able to continue to deploy backhaul services today and many are evaluating adding fronthaul services to support advanced LTE-A capabilities and to prepare the ground for 5G. But everyone wants to avoid the situation where fronthaul and backhaul deployments are quickly obsoleted as 5G is standardized.

Network operators therefore need flexible equipment for mobile fronthaul and backhaul that has the ability to support CPRI or Ethernet today and can be updated in the future via a software upgrade to support 4G cell sites that will be upgraded to support a mixed 4G and 5G environment.

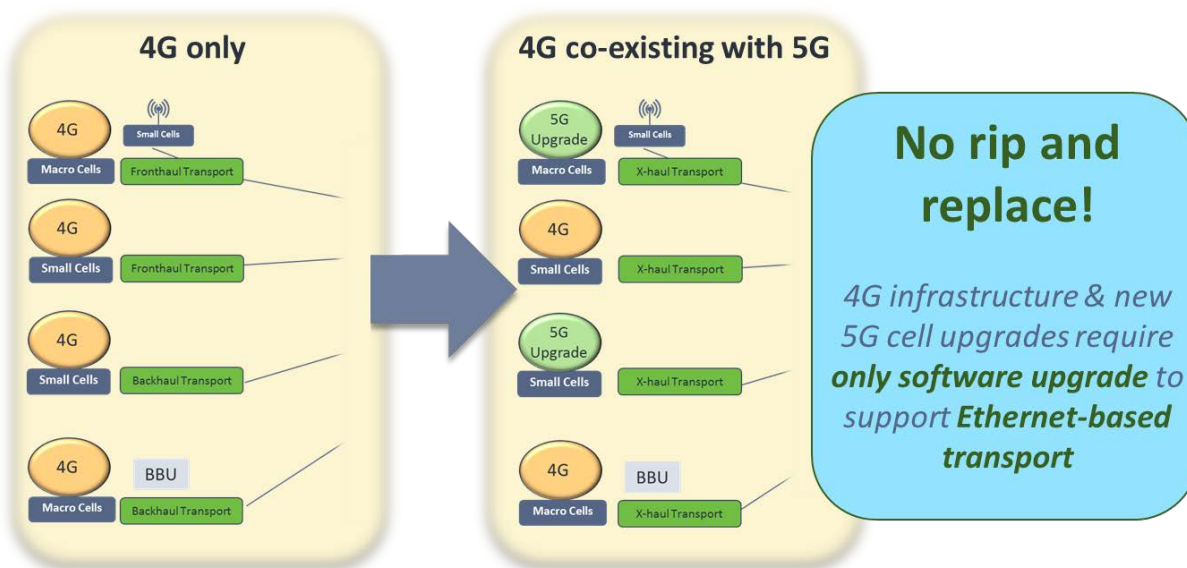


Figure 2 - Upgrade scenario that avoid "rip and replace". Source: Infinera

These systems are FPGA-based and allow for very flexible support of network protocols such as CPRI, OBSAI and Ethernet. By using FPGAs it is possible in the future to add new protocol framing to support the new variant of Ethernet that will be required for 5G via an in-service field upgrade. This means the operator can reuse the fronthaul hardware in a 5G network and avoid the need to rip and replace hardware.

5G will require new mobile transport hardware in all new 5G sites and any site performing layer 2 aggregation and switching but many current 4G cell sites can support fronthaul and backhaul services today and provide a smooth migration to 5G without the need to replace cell site mobile transport hardware.

Conclusion

The migration to 5G will require improvements to the overall network infrastructure in terms of performance, features and bandwidth. These improvements will drive new fiber builds, and fiber upgrades to an ever-growing number of cell sites, creating significant opportunity for cable MSOs and other wholesale operators to capture significant share of cell backhaul and fronthaul services for 4G and 5G mobile networks. For example, cable MSOs could create significant competitive advantage addressing current and potential cell fronthaul and backhaul services as they look to rearchitect their networks to support remote PHY, which is very well suited, from both performance and bandwidth perspectives, for these requirements.

A key consideration for the MSO community is how they can take advantage of the migration to 5G with managed services instead of simply providing dark fiber to wireless operators. While wireless operators will possibly take the initial view that their own network built using dark fiber from numerous sources, such as cable MSOs, is the best way forward, the MSO community should challenge this with a value proposition built around better network economics and performance. Wireless operators should focus their resources on differentiating their networks and services against their competition rather than all building “me-too” transport networks in parallel. The architects of the 5G standards already anticipate that network sharing will be key to 5G and are building support for this into the standards and architectures with capabilities such as network slicing.

MSOs should take advantage of the physical resources of fiber, HFC and real-estate and their field force to take away the pain of scaling networks from 4G to 5G with the massive proliferation of cell sites in geographies where all wireless operators will require transport. In these areas transport is a cost the wireless operators will all need to bear but not will gain competitive advantage by building their own.

Better economics can potentially be achieved by managed services from MSOs based on sharing a common MSO-based network between multiple wireless operators or taking advantage of other network transition projects such as packet-optical based remote-PHY transport to support wireless over a common infrastructure. In both these cases network slicing can potentially provide the wireless operator with the SDN-based control they desire without the need for them to build their own dark fiber based network. A further consideration is can the MSO also combine their network assets with their field force assets to enable the wireless operator to avoid the need to drastically increase their own field force to deal with the explosion of cell sites that 5G will require.

MSOs who can build a business case for business services instead of dark fiber should carefully consider future 5G requirements for both fronthaul and backhaul services to ensure future 5G migration can be accommodated as much as possible within current hardware to provide investment protection and to minimize network reengineering. Careful consideration should also be given to network performance in areas such as low latency and synchronization performance as means of differentiating the MSOs managed service performance against their own competitors.

Abbreviations

BBU	baseband unit
CoMP	coordinated multipoint
CPRI	common public radio interface
C-RAN	centralized or cloud radio access network
CU	central unit
DU	distributed unit
eICIC	enhanced intercell interference coordination
EPC	evolved packet core
FPGA	field-programable gate array
HetNet	heterogeneous network
LTE	long term evolution
LTE-A	long term evolution – advanced
MEC	mobile edge computing
NFV	network functions virtualization
RAN	radio access network
RRH	remote radio head
SDN	software defined networking
vBBU	virtual baseband unit

Bibliography & References

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

The Intersection of HFC and 5G

A Technical Paper prepared for SCTE•ISBE by

Keith R. Hayes

Principal

Broadband Advisors Group, LLC

175 Hog Farm Circle Canton, GA 30115

770-378-3595

Keith.hayes@broadbandadvisorsgroup.com

Introduction

5G. 5G. 5G. On billboards, in your inbox, at your favorite tech blog:

- 5G - **What** is it?
- 5G- **Where** is it?
- 5G - **WHEN** is it?
- 5G - **How** will it be different for users and the supporting infrastructure?
- 5G - **Why** do we need it?

One of the foundational differences of 5G from the previous mobile communications Generations is its density – perhaps 10x or more than 4G – driving the need for Location, Power, and Connectivity for the Radios that the Hybrid Fiber-Coax (HFC) network is potentially well-equipped to support. This paper will delve into the 5G topology to answer those questions and review performance expectations and how they might map into an HFC network with Data Over Cable Service Interface Specifications (DOCSIS) 3.0 and DOCSIS 3.1

Content

1. Wireless Mobility Evolution




Figure 1 - Cellular Standards Timeline Source: Broadband Advisors Group, LLC

Cellular communications as we know it today began in the early 1980s with the Advanced Mobile Phone System (AMPS), often colloquially referred to as Analog Mobile Phone Standard since the communication was indeed analog and could be listened to (one end of the call) by an eavesdropper with a scanner. The AMPS standard broke the spectrum bottleneck and lowered power requirements for the mobile stations by replicating the spectrum bandwidth through “cells”, radio transmitters/receivers serving a specified geographic area. Adjacent “cells” would use different frequencies to minimize interference, and the key to operational success was the Mobile Switching Telecommunications Office which through telemetry with the handset supported switching from one cell to the next rapidly enough that calls would stay connected and conversations would be able to be conducted normally. AMPS was launched in the 850 megahertz (MHz) spectrum, and was subsequently expanded into reclaimed ultra-high frequency (UHF) broadcast spectrum to add bandwidth.

As mobility technology developed in the late 1980’s, there was an effort in Europe to develop mobility standards that would be digital, work cross-platform, and more importantly enable international roaming between countries, and the Global System for Mobile (GSM) standard set came to life. GSM is a joint

Time Division Multiple Access (TDMA) / Frequency Division Multiple Access (FDMA) modulation schema, and also brought along the Simple Message System (SMS) with its 160-character limitation.

The first packet-switched cellular data service, General Packet Radio Service (GPRS) introduced data service to complement voice and SMS around 2000.

As data use ballooned, GSM continued its development with the Enhanced Data Rates for Global Evolution (EDGE) that was marketed by some operators as 3G for third-generation cellular capabilities. Concurrently a competing standard – Code Division Multiple Access (CDMA) was developed and in the United States some providers chose GSM (AT&T, T-Mobile) while the others chose CDMA (Verizon, Sprint, US Cellular) setting up a handset challenge as for many years handsets only had one chipset or the other in them and could not operate on the other network. Both standard sets supported higher data rates through denser orders of modulation.

CDMA was further enhanced with Evolution – Data Optimized (EV-DO) which introduced a separate data channel that would work in tandem with the voice channel. High-Speed Packet Access (HSPA) was the corollary in the GSM ecosystem.

Recognizing that there was both an insatiable demand for mobile data consumption and that handsets would soon be available with chipsets that could support both GSM and CDMA, the 3 Generation Partnership Project (3GPP), a consortium of international standards bodies charted a path of network development towards the 4th Generation environment, which they named Long Term Evolution (LTE), trademarked by the European Telecommunications Standards Institute – (ETSI) that would provide both platforms with download speeds of up to 300 Megabits per second (Mbps) per cell, along with support for additional frequencies and improved centralized communications management.

Release 8 LTE provided a pre-4G capacity and bandwidth network performance enhancement including carrier-width support from as little as 1.4 MHz to as much as 20 MHz and a much more spectrally-efficient air interface modulation – Orthogonal Frequency Division Multiple Access (OFDMA). The rapid improvement in data performance provided impetus for operators to market these enhancements as 4G though officially 4G was tied to release 10 that included carrier aggregation and Multiple In Multiple Out (MIMO) antenna support. Release 10 is marketed as “True 4G” or LTE Plus in Sprints case.

2. 5G Timeline and Foundational Pillars

As the mobile environment moves towards 5G, it is important to understand that there is are marketing and user-experience components that have driven consumer acceptance and expectations of the performance of the cellular networks. Two main factors have been in play:

1: Coverage – how much of the US the network covered – we all remember the Verizon television commercials with the “Can you hear me now?” guy – giving customers confidence their phone calls would connect and wouldn’t drop.

2: Connectivity – as in to the internet – with increasing performance; 3G, LTE, 4G, LTE-Plus, along with battling maps in TV commercials and on billboards.





Figure 2 – Verizon Wireless and ATT Coverage Map Source: Engadget.com

What was not messaged, marketed, or even remotely hinted at was the speed/bandwidth of the data connection, contrary to how wireline (or satellite) internet services are sold with claims of “Speeds up to...X Mbps” with the FCC Broadband America report and sites like Speedtest.net validating those claims. There are a couple of key differentiators that drive this dichotomy. First, 4G/LTE handsets and networks under optimal conditions can deliver speeds for a test of up to 100 Mbps, but in reality, given that one phone and user can only do one thing at one time, there is not a practical need for more than a few Mbps for typically video streaming. Second, except for Wi-Fi tethers, only one device, the mobile phone, is consuming bandwidth concurrently, unlike the home environment where dozens of devices may be connected and requiring bandwidth.

The 5G topology is designed to move away from the “Coverage/Connectivity” environment to one where much higher connectivity speeds, much lower latency, and a huge increase in the density of connected devices is expected.

A joint effort is under way led by the 3GPP and the International Telecommunications Union (ITU) to develop and promulgate standards that will support the planned capabilities of the International Mobile Telecommunications (IMT) system 2020, the 5G network we have all been hearing about.



M.2083-03

Figure 3 - Key Capabilities IMT 2020 (5G) vs IMT 2015 Source: ITU-R M.2083.0

The chart above depicts the expected increase in capabilities of the 5G network (IMT 2020) vs the 4.5 G network (IMT 2015). You likely have heard of peak data rates of 1 Gigabit per second (Gbps), along with Internet-of-Things (IoT) connection density, and latency between radio and network in the millisecond (ms) range, there are other key vectors including improved Network Energy Efficiency, Improved Spectrum Efficiency, and increase in Mbps-per-square-meter. The big question is – when will all these new capabilities come to fruition?

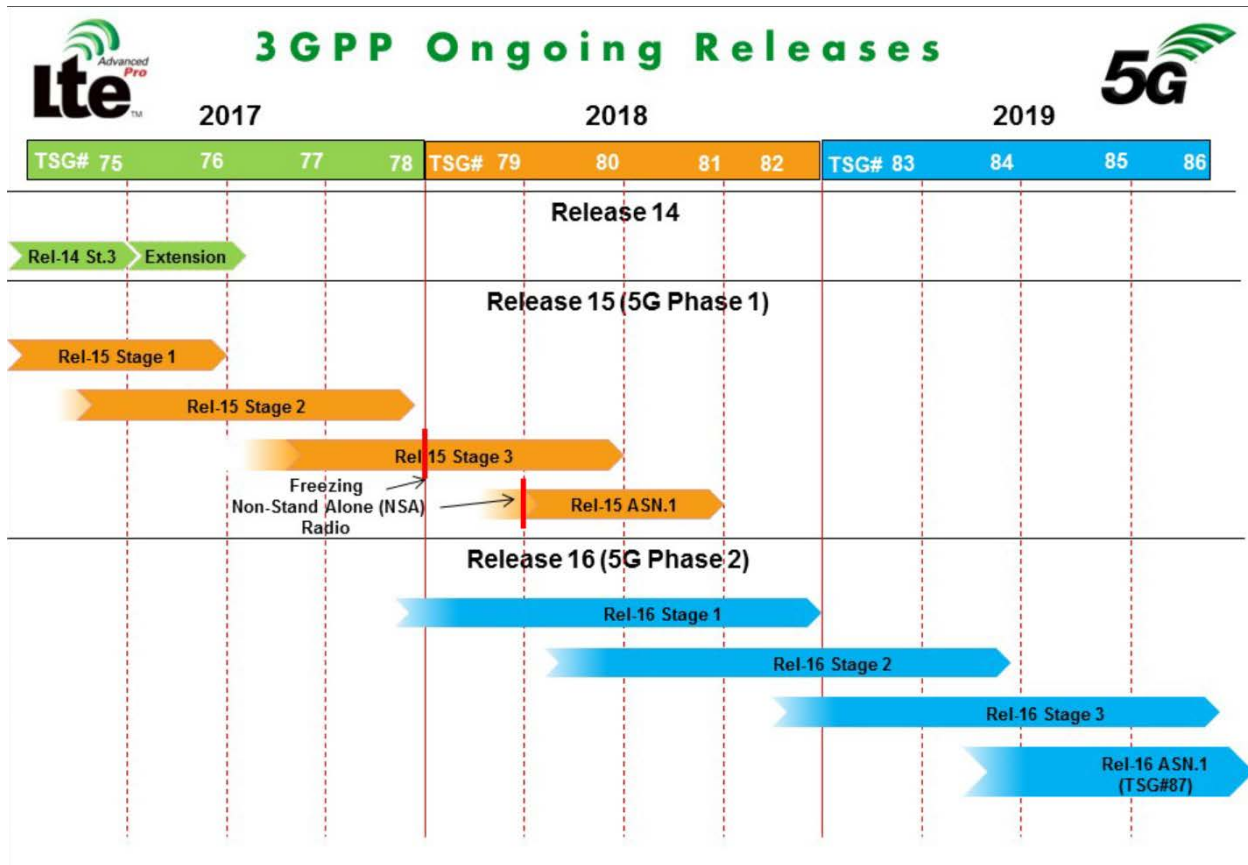


Figure 4 - 3GPP Release Calendar Source: 3GPP.org

To answer that question, we go to the 3GPP Release Calendar, which earlier this year anticipated locking down the 5G Standards in the middle of 2018. However, at the urging of all the large US mobility carriers and others, the 3GPP agreed to accelerate by about 6 month the standards for Non-Stand Alone (NSA) Radios.

This should allow silicon to be developed and integrated for initial deployment of NSA upgrades in 2019, with the reaming standards coming to market in 2020 but likely taking a couple of years to develop scale. One important point to note is that 4/4.5G does not go away in a 5G world, but will both interact with support, and be supported by 5G technologies.

Now that we have reviewed some of the platform expectations of the 5G network, how do we get to Gbps downloads with billions of connected devices?

3. 5G Spectrum and Radio Access Network (RAN) Densification

When more bandwidth is needed in a communications network there are three main levers to pull individually or in combination to develop the increase:

1: Obtain more spectrum

2: Increase the efficiency through denser modulations allowing the transport of more bits per hertz


3: Enable the re-use of spectrum through network densification

The 5G roadmap delivers on all three, all of which the cable industry has employed going from 300 MHz to 1 gigahertz (GHz), moving from Quaternary Phase Shift Keying (QPSK) to 256 Quadrature Amplitude Modulation (QAM) and beyond, and performing node recombines, node splits, and service group splits for spectrum re-use.

On the spectrum front, there are two very significant additions from the legacy cellular world:

1: Potentially using unlicensed or shared spectrum such as the 3.5 GHz Citizens Broadband Radio Service (CBRS) in the US.

2: Employing hundreds of MHz of spectrum at very high frequencies such as 28 GHz and 39 GHz. It should be noted that these higher frequencies are commonly called “millimeter wave”, however only the 39 GHz spectrum has sub-centimeter wavelength of ~7.6 millimeters (mm).



Nationwide, population-weighted average spectrum assets as of 2/7/17. These numbers are national averages and do not represent the spectrum assets in any specific market.

Figure 5 - Spectrum by Carrier US Source: Sprint

The figure above reflects approximate spectrum of the 4 major US carriers reflecting population-weighting (how much spectrum per average person) versus geographic coverage. Each of the spectrum areas are broken up in a number of blocks/channels, ranging from 10 to 30 MHz in width that are continually reused in alternating cells. Though mileage may vary, for the purposes of providing bandwidth scoping and to make the math easy we will assume throughput on the order of 5 bits per hertz, so a 10 MHz channel would support ~50 Mbps, a 20 MHz 100 Mbps, etc. By aggregating channels together, the carriers are able to deliver higher total throughput, but the challenge for handset manufacturers is adding more antennas and tuners while managing battery life.

The chart below provides the licensee status of the cellular spectrum blocks in the top 3 US Markets.

CHART OF LICENSEES

CMA	CMA Name	Cellular A Block	Cellular B Block	PCS A Block	PCS B Block	PCS C Block	PCS D Block	PCS E Block	PCS F Block	PCS G Block
1	New York, NY	AT&T	Verizon Wireless	AT&T, T-Mobile	Sprint Nextel	Verizon Wireless	T-Mobile	AT&T	Verizon Wireless	Sprint Nextel
2	Los Angeles, CA	AT&T	Verizon Wireless	Sprint Nextel	AT&T	T-Mobile, Verizon Wireless, Metro PCS	AT&T	Verizon Wireless	T-Mobile	Sprint Nextel
3	Chicago, IL	AT&T	Verizon Wireless	AT&T	US Cellular, Verizon Wireless	T-Mobile	Sprint Nextel	Sprint Nextel	AT&T	Sprint Nextel

Figure 6 - Cellular Block Licenses – top 3 US markets Source: FCC

The new cm/mm frequency blocks are where significant bandwidth can be added – the 28 GHz block goes from 27.5 to 28.35 GHz, and the 39 GHz block goes from 38.6 to 40 GHz with some license holders having as much as 600 MHz in certain cities. Though there is a lot of bandwidth in these bands, they are not without challenges. Higher frequency signals do not propagate as far, cannot penetrate buildings well, and are reflected more easily than the lower-band cellular channels. (For you cable folks that remember rain fade on Community Antenna Relay System (CARS) Band microwave links, these cm/mm wave frequencies “enjoy” tree fade and a simple brick/stud/sheetrock wall can introduce ~30dB of attenuation).

Innovative players like Vivint have developed platforms in the 28 GHz spectrum that use a hybrid approach – a high bandwidth point-to-point link is established between a tower or other transmitter location to a rooftop of a “hub” – a home with clear line-of-sight to the tower and an owner willing to allow antennas to be place on the roof (typically in exchange for free internet) which then can serve as many as 24 other homes via 5 GHz Wi-Fi. This example demonstrates two key concepts in the coming 5G world:

- 1: 5G will serve fixed as well as mobile applications
- 2: Creative uses of multiple frequencies will be commonplace and enabled by the 5G environment. In fact, 5G can and will be used in low, mid, and high band legacy cellular frequencies either in tandem or standalone from the cm/mm bands.

A third concept is important to understand – the “brain” that will allow all these different bands, applications, and systems to interoperate in a sub-millisecond latency environment by definition cannot be back at the core as in the legacy cellular networks connected to the Mobile Telephone Switching Office (MTSO) – it must be near the edge, be able to cache data, and have significant localized intelligence. 5G proposes to solve this challenge with a combination of Software Defined Networking (SDN) and Network Functions Virtualization (NFV). These technologies will enable a multi-level network with autonomous local routing directed from a SDN controller. An example of this would be in a vehicle-to-vehicle communications system instead of the packets having to go from car to tower to MTSO to tower to car,

the Remote Radio Head (RRH) would have the SDN-enabled capability to connect the two vehicles locally.

4. The 5G PAC (Power, Attachment, Connectivity) Conundrum

In the 4G world of 2017 there are 300,000 – 350,000 classic tower/building-top cellular base stations – known in the industry as Macros, and tens of thousands of micro, pico and femtocells. Each cell site could cover as much as several square miles to only a diameter of 1000 feet or so. In the 4G environment absent new spectrum (and at this point not including modulation density improvements) to double the bandwidth would require doubling the towers, but the aspirations of 5G are for a ten-times higher Mbps per square meter – **10X**! The answer is cell-densification, also known as “small cells”. How small? In some proposals only a few hundred feet wide.

To illustrate this the example below shows a 4G “Macro” site with a 1-mile radius, and a subset of 500-foot radius small cells.




Figure 7 - 500 foot small cells, 1 mile diameter Macro cell

Source: Broadband Advisors Group, LLC

If a 1 square mile macro site (please note a 1-mile diameter hex is not a square mile – illustrative example) was to be fully populated with 500-foot diameter small cells – there would be more than 100 of them, with all the ensuing challenges of power, how/where to attach to light poles, power poles, strand, and how to get network connectivity to them. That’s not the end game for 5G – for one, the Macro site doesn’t go away – it remains serving 4G customer and perhaps enhancing 4G performance and coverage through 5G-enabling technologies such as the dedicated fronthaul/backhaul channel to the Macro site employed by the Sprint Magic Box.

If a cellular operator faces bandwidth challenges in residential areas, there are not many options if additional spectrum is not available – current tower locations are often as close to homes as the municipalities and residents will allow, and adding additional towers, even if a smaller monopole, requires significant time for zoning, permitting, and a fiber extension would be required.

What if an approach modeled after outdoor Wi-Fi networks was taken that could improve bandwidth?

Keeping with our model of a 1 square mile cell, what else is roughly a square mile that the cable industry has hundreds of thousands of? HFC Nodes. Ranging from as small as a couple tenths of a square mile in dense urban areas to several square miles in rural zones, if one looks at road-miles-per-square mile from state and county Departments of Transportation in suburban areas there are 5-7 road miles per square mile – sound familiar? HFC nodes serving 500 homes passed are around 5 strand-miles.



Figure 8 – 500 foot small cells integrated with HFC network

Source: Broadband Advisors Group, LLC

The example above depicts part of node with ~500' diameter cellular base stations overlaid. Much like outdoor Wi-Fi deployments, these radios could solve the PAC conundrum by obtaining **Power** from the HFC network, being **Attached** via strand mount brackets, and by obtaining **Connectivity** through an embedded DOCSIS modem. Note that the overlays are not ubiquitous – one of the challenges with radio densification is managing interference, so the HFC-connected radios would be designed to cover areas with heavy bandwidth demand (a bus stop for example) or where coverage from the macro base needs to be improved.

For planning purposes, we will assume the radios would be similar to outdoor wireless equipment in terms of space, power consumption, and bandwidth – roughly the size of an amplifier or fiber splice

closure, drawing low tens of watts, and provisioned much like a residential or commercial cable modem with 50-100 Mbps downstream and 3-5 Mbps upstream.

In the power domain, if each radio consumed 20 watts, if 20 were deployed in a node, at 75 volts average input the current draw would increase by 5.3 amps, in most nodes requiring the installation of an additional power supply.

In the Connectivity domain, to make the math easy we will assume the modems are provisioned for 100 Mbps downstream (we will leave upstream out of this scenario as DS/US loads in the cellular domain are very similar those of wireline networks – much lower than downstream) – if the Macro site was congested, a dozen users waiting at the bus stop would have more than enough bandwidth to stream video from the local radio.

A second benefit is Internet-of-Things devices would be able to use dramatically less transmit power as the local radios would be hundreds of feet away versus thousands.

HFC operators would rightly be concerned about overloading their DOCSIS platforms as they were with outdoor Wi-Fi deployments, and certainly extensive traffic analysis and modeling would need to be undertaken prior to deployment. In this traffic modeling, as CEO of Charter Tom Rutledge has noted, more than 70% of the bits flowing down to a smartphone actually do not flow through the cellular network – they flow through the customers home Wi-Fi network, attached to the cable modem and then through the DOCSIS network. The outdoor HFC-fed radios would only be supporting small volumes of short-duration activity such as a car driving down the street with a child in the backseat streaming video or a jogger streaming Pandora as she made her neighborhood rounds, and likely would consume traffic much like a residential modem – adding another 20 customers in a 500 homes-passed node would not overload a typical DOCSIS Service Group.

If this scenario was deployed, now the operators have another cellular backhaul revenue stream and the improved cellular coverage would provide better experiences for the cable customers TV Anywhere-type streaming via the mobility network.

As cellular densification continues and the need for bandwidth exceeds the evolving capabilities of DOCSIS, (some 5G nodes are postulated to need as much as 80 Gbps), fiber could be extended to the radio for connectivity with power supplied by the coax network – yet another demonstration of the flexibility and capability of the Hybrid Fiber-Coax Architecture.

Some equipment vendors have already implemented LTE-backhaul optical circuits in their optical nodes, and some early DOCSIS 3.1 Remote-Phy Devices (RPD) also have backhaul circuits. The current versions provide 1 to 10 Gbps symmetrical but 80 Gbps or more is on the near-term roadmap. Where these nodes/RPD's are deployed, a more powerful radio consuming more bandwidth could be installed nearby (known in the industry as a Mini-Macro) with little or no outside plant modification.

Conclusion

Now that we have reviewed the evolution of the cellular network, the forthcoming 5G augments to it, and how the HFC network might intersect with it, what have we learned?

- 5G - **What** is it? – an exciting new set of standards for wireless communications, both mobile and fixed, that will dramatically increase connectivity speeds, support much denser end-device connections, reduce latency, and employ new licensed and unlicensed frequencies
- 5G - **Where** is it? – Nowhere, but just wait...
- 5G - **WHEN** is it? – 2020 - with some field trials before and mass scale by 2022
- 5G - **How** will it be different for users and the supporting infrastructure? – Higher connectivity speeds, lower power requirements, better support for machine-to-machine traffic, much denser cellular radio network
- 5G - **Why** do we need it? – to solve the insatiable demand for higher bandwidth for smartphones, overcome cellular spectrum challenges, enable exciting capabilities such as autonomous vehicles, and support the explosion of Internet-of-Things connected equipment

As 5G comes to fruition, cable operators have opportunities to continue to expand the fiber-optic backhaul for macro sites and sites that will be employing large swaths of cm/mm wave spectrum, but also can potentially employ the HFC network for Power, Attachment, and Connectivity for cellular network densification in the legacy cellular radio spectrum.

Abbreviations

3GPP	Third Generations Partnership Project
AMPS	Advanced Mobile Phone System
CBRS	Citizens Broadband Radio Service
CDMA	Code Division Multiple Access
DOCSIS	Data Over Cable System Interface Specifications
EDGE	Enhanced Data Rates for Global Evolution
ETSI	European Telecommunications Standards Institute
EV-DO	Evolution – Data Optimized
FDMA	Frequency Division Multiple Access
Gbps	Gigabit per second
GHz	Gigahertz
GPRS	General Packet Radio Service
GSM	Global System for Mobile
HFC	Hybrid Fiber-Coax
HSPA	High-Speed Packet Access
IMT	International Mobile Telecommunications
IoT	Internet of Things
ITU	International Telecommunications Union
LTE	Long Term Evolution
Mbps	Megabits per second
MIMO	Multiple In Multiple Out
MHz	Megahertz
mm	Millimeter
MTSO	Mobile Telephone Switching Office
NFV	Network Functions Virtualization

NSA	Non-Standalone
OFDMA	Orthogonal Frequency Division Multiple Access
PAC	Power Attachment Connectivity
QAM	Quadrature Amplitude Modulation
QPSK	Quaternary Phase Shift Keying
RAN	Radio Access Network
RPD	Remote Phy Device
RRH	Remote Radio Head
SDN	Software Defined Networking
SMS	Simple Message System
TDMA	Time Division Multiple Access
UHF	Ultra-High Frequency

Bibliography & References

Verizon/ AT&T Coverage map <https://www.engadget.com/2009/11/03/atandt-sues-verizon-over-theres-a-map-for-that-ads/>; Engadget.com

Enhancement of key capabilities from IMT-Advanced to IMT-2020; ITU-R M.2083-0 International Telecommunications Union – Radio

Spectrum by US Carrier; <http://newsroom.sprint.com/in-land-wireless-spectrum-is-king.htm>; Sprint

Cellular Block Licenses – Top Three US Markets; FCC – Federal Communications Commission

Testing for SLA Compliance of Business Services over DOCSIS 3.1

A Technical Paper prepared for SCTE•ISBE by

Robert J. Flask
Sr. Product Line Manager
Viavi Solutions
5808 Churchman Byp.
Indianapolis, IN 46203
317-614-8125
robert.flask@viavisolutions.com

Introduction

Communications and data service growth is driving the demand for Ethernet services. Many Ethernet services, such as cellular backhaul have stringent Service Level Agreement (SLA) requirements. SLA requirements are verified at time of turn-up and can be monitored in the background. With the extensive reach and coverage of the Hybrid Fiber Coax (HFC) networks, cable operators are well positioned to utilize DOCSIS as a delivery method for Ethernet services. Legacy DOCSIS systems including DOCSIS 3.0 have limitations that prevent it from being used for timing critical Ethernet services. DOCSIS 3.1 has added additional mechanisms that if designed and implemented can provide capability to provide Ethernet services to stringent services such as Advanced LTE backhaul. Best practices for Ethernet service SLA compliance can be implemented that include turn up and automation to simplify the process. There are various testing standards used for validating business service SLA's including RFC-6349, RFC-2544, Y.1564, Y.1731. Many of these standards weren't very relevant with DOCSIS 3.0 but now are likely to become part of the repertoire of DOCSIS 3.1 Ethernet service deployments.

Content

1. DOCSIS and Ethernet SLA Testing

DOCSIS is by nature a shared environment. When providing Ethernet services over DOCSIS, only a portion of the DOCSIS network capacity is being utilized for the customer. The DOCSIS network can be viewed as the User Network Interface (UNI). Multiple customers will use the DOCSIS network simultaneously, including residential data services and voice service. Providing a dedicated Ethernet service to a business entails the creation of an Ethernet Virtual Circuit (EVC) that utilizes a portion of the overall UNI. The most typical way of supporting Business Services Over DOCSIS (BSOD) today is via Layer 2 VPN support. CableLabs specifies guidelines for operators to offer MEF compliant Layer 2 Transparent LAN services (TLS) to commercial enterprises.

Note that the sum of all of the EVC's for guaranteed SLA service must be less than the overall UNI speed. This creates challenges for sharing a portion of a DOCSIS network with SLA business services and high capacity best-effort services.




Figure 1 - Multiple EVC's over the DOCSIS network

The DOCSIS network is only a portion of what the Ethernet service resides in. Rarely does a customer only need Ethernet service within the DOCSIS portion of the network. As shown in Figure 2 the DOCSIS portion is typically just the last mile.

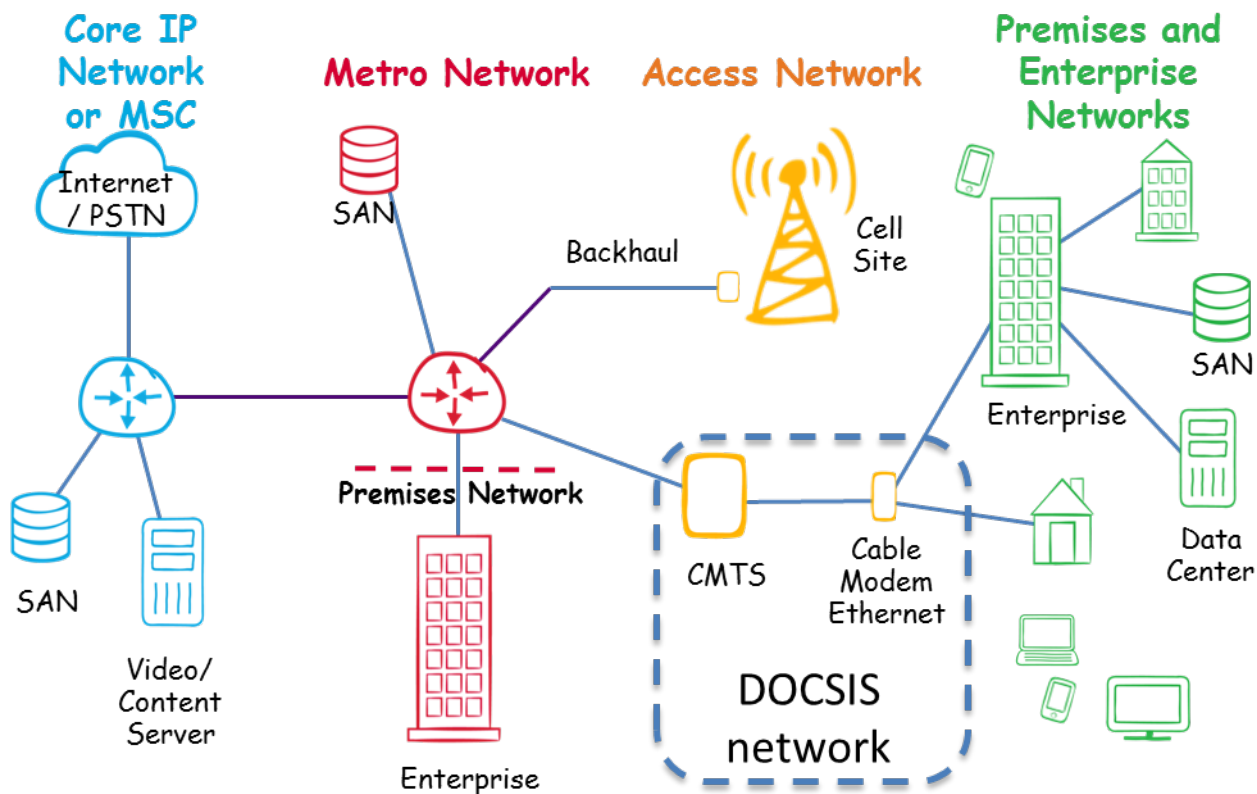


Figure 2 - DOCSIS Network is Last Mile Portion of the Ethernet Service

1.1. Legacy DOCSIS

Existing DOCSIS services up through DOCSIS 3.0 had capacity, architecture, and design limitations that limited the opportunities where DOCSIS could be effectively used for SLA based Ethernet business services. Asymmetric capacity limitations limited the speed of service that could be offered and performance metrics regarding latency (frame delay), and jitter (frame delay variation). The timing accuracy for DOCSIS 3.0 systems prevented the use of DOCSIS for cellular backhaul, LTE, and Advanced LTE.

1.1.1. Best Effort Services – SMB

DOCSIS fits well into the Small to Medium Business that mostly requires solid and consistent Internet services. Throughput tests with capability to validate best effort pipe speeds is critical to make sure impairments aren't impacting the capacity. Typically, operators will provision with about 10 % Excess Information Rate above the stated service level to ensure Speedtest throughputs meet or exceed the advertised upload/download speeds.

But with best effort services there typically aren't any penalty clauses regarding speeds. More typically, the SLA's for best effort service surround availability and Mean Time To Repair (MTTR).

To ensure real life customer experience for a best effort service, RFC-6349 testing is recommended. This will identify the customers actual layer 4 TCP throughput and also identify any configuration or TCP efficiency issues. RFC 6349 testing can show the effects of traffic shaping that may be occurring.

1.2. Metro Ethernet service

1.2.1. Key Components of Ethernet service activation SLA testing

Ethernet business services can vary from being a basic best effort internet service, like a traditional home cable modem, to a high-performance Ethernet service with guaranteed performance with SLA requirements for availability, MTTR, Latency/Frame Delay (FD), Jitter/Frame Delay Variation (FDV), Throughput (Committed Information Rate (CIR)/ Excess Information Rate (EIR), and frame loss. Ethernet business services typically require some level of SLA. A typical SLA for a local Metro Ethernet service is show in Table 1 for Ethernet services and standard mobile backhaul service.

Table 1 - Example of typical local Metro Ethernet service requirements

KPI	Ethernet Services	Mobile Backhaul services
Frame Delay	typical 5 ms - best effort up to 30 ms	< 8 ms typical 5 ms
Frame Delay Variation	< 2 ms	
Frame Loss	6.25 x 10 ⁻⁶	
Throughput	99.99%	
Availability	99.995%	
Mean-time to repair	2 hours (max 4 hours)	

Comparatively, with legacy DOCSIS 3.0 the SLA levels are much relaxed. Particularly the Frame Delay and the Frame delay variation. Table 2 shows a typical example of an SLA over an existing DOCSIS 3.0 network. Note that the Frame Delay allowance is even higher than typical best effort Ethernet services.

Table 2 - Example of typical local Ethernet SLA over DOCSIS 3.0

KPI	DOCSIS Ethernet Services
Frame Delay	<60ms
Frame Delay Variation	< 12 ms
Frame Loss	< 0.1%
Throughput	95%
Availability	99.9%
Mean-time to repair	4 hours

1.2.2. Ethernet KPI verification

Validation of SLA requirements require testing and monitoring. Common practice is for initial testing via a field meter at install and service turn up. Often the field testing is combined with automated or centralized test through remote test heads or software agents to make a faster and more consistent test practice.

1.2.3. *What are the key tests*

To verify and test to the SLA, various tests are used for Metro Ethernet services depending on the service type and SLA. Table 3 shows a list of common tests that are used for validating the varying SLA requirements.

Table 3 - Service Tests for validating Ethernet SLA's

Service Activation Test	Description	Comment
Connectivity, Throughput and Auto-Negotiation	Verify basic connectivity Verify best effort throughput Validate auto negotiation settings to identify half/full duplex limitations	Pre-Check prior to complete SLA testing.
RFC-2544 – Single Stream Pipe test	Industry-standard service activation test for single-service Ethernet and IP (i.e. “pipe” test) Measures key performance indicators and bandwidth profile such as: throughput, latency, packet Jitter, frame loss, and committed burst size (CBS)	Commonly done at service activation with loop-back of remote device. This is an out of service test Testing in both directions, both independently and concurrently (independent transmissions from both ends) for bi-directional testing is critical to determine which direction is not meeting the SLA. Single ended loopback modes cannot isolate which direction.
Y.1564	The industry standard service activation test for multi-service Ethernet and IP (“Triple Play”) Measures KPIs and bandwidth profile such as: • CIR, EIR (Throughput) • Frame Delay (FD), Latency • Frame Delay Variation (FDV), Jitter • Frame Loss Rate (FLR) • Committed Burst Size (CBS), Policing	Commonly done at service activation with loop-back of remote device. Testing in both directions both independently and concurrently (independent transmissions from both ends) for bi-directional testing is critical to determine which direction is not meeting the SLA. Single ended loopback modes cannot isolate which direction. This is an out of service test
Layer 2 Control Plane	Confirm transparent forwarding of Ethernet traffic through the providers network	Done at field turn-up. Typically instrument turn up with loopback from remote device
RFC-6349	Automated and repeatable TCP-throughput test per IETF RFC 6349 standards, including key performance metrics of TCP efficiency and Buffer delay	Can be virtualized with centralized reporting. This can be done in-service but utilizes full bandwidth
Y.1731	Performance monitoring and PM protocol. Can be used for service testing including loopback, frame delay, frame delay variation, frame loss	Useful in remote automation

1.2.4. Operation Administration and Monitoring (OAM)

802.11ag Connectivity Fault Management which is part of the Operational Administration and Maintenance (OAM) protocol allows maintenance and management of each EVC regardless of the transport layer. This allows operators to partition the network across the different boundaries and third-party operators.

Y.1731 Performance Monitoring OAM protocols allow for devices to have some continuous remote monitoring activities. Y.1731 provides in-service performance monitoring to measure KPI's like frame delay, frame delay variation, frame loss. Y.1731 can also provide fault management capabilities including remote loopback. Y.1731 can be an important part of automated turn-up testing as well.

1.3. How DOCSIS 3.1 enables improved Ethernet services

Gigabit services offerings are now available with DOCSIS 3.1 systems. A system that utilizes two 192 MHz wide OFDM downstream channels as well as two 96MHz wide OFDM-A upstream channels is capable of supplying over 2.5 Gbps downstream and 1Gbps upstream of shared capacity.

The CableLabs specification for DOCSIS 3.1 devices and services took into consideration the desire to deploy business services over DOCSIS. The new modulation schemes provide improved data efficiency and signaling, and the addition of the DOCSIS Timing protocol provides steady and consistent point to point timing and clock accuracy at the CM egress point.

1.3.1. DTP – DOCSIS Timing Protocol

The DOCSIS timing protocol provides the ability for the CMTS to reference a timing protocol source and provide the clock on the output of the CM to the end device. DTP is based on the Precision Timing Protocol (PTP) of IEEE-1588-2008 with a few modifications




Figure 3 - DTP Clock System

When implemented, DTP allow the CMTS and CM to appear as a single DOCSIS system to the outside world. This allows the DOCSIS system to appear as a Boundary clock. This allows the Ethernet system to compensate and minimize the effects of any system delays or frame delay variances. The system delays that are calculated are added to the timestamp, which effectively negates the delays. Delays and variances are issues that made legacy DOCSIS inoperable for most business class services.




Figure 4 - DOCSIS System Works Like Boundary Clock

DTP provides the capabilities and facilities for service providers to design and implement a system. Operators will need to do the DTP tuning and system optimization to meet their desired needs.

One consideration is the modem to modem/system to system skew. This variance in modem behavior can be one of the largest sources of timing error in the error budget.

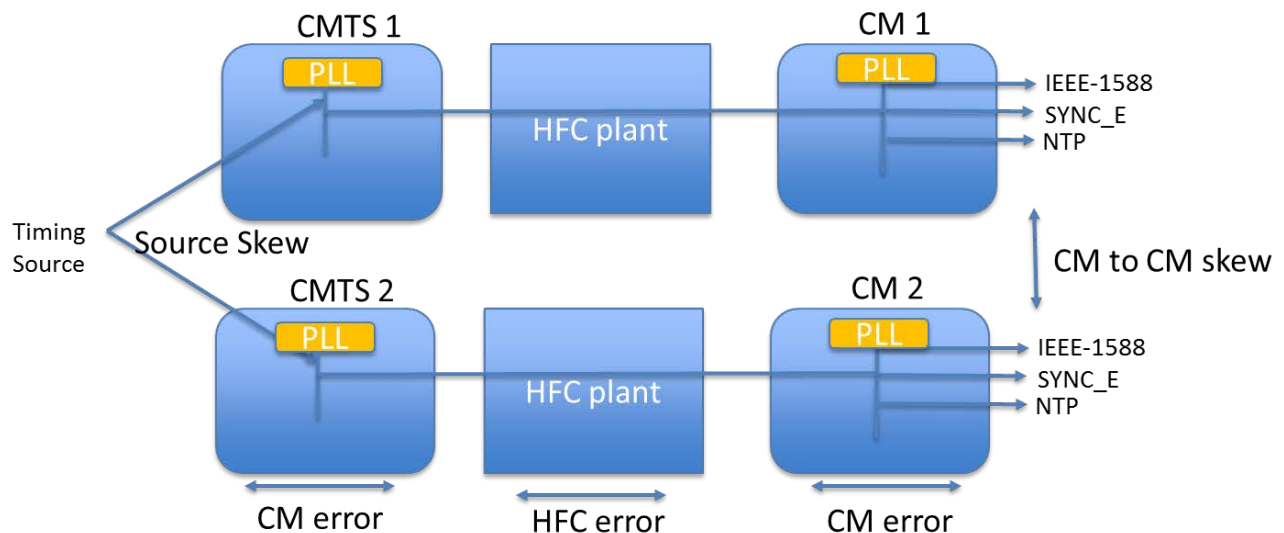


Figure 5 - DTP System Skew And Error Budgets

CableLabs has provided error budgets for the DTP systems to meet different levels of systems to support deployments such as LTE advanced. Table 4 shows system error budget for the different portions of where errors come in the DOCSIS network.

Table 4 - DTP System Timing Error Budget (table 10-9 CM-SP-MULPlv3.1-I11-170510)

Parameter	Level I System (GPS location requirements)	Level II System (Relaxed positioning)	Level III System (LTE Advanced Macro Cell + addnl.	Level IV System (LTE Advanced macro cells and small cells	Level V System – Current DOCSIS implementation
T-cmts-error	+/- 20 ns	+/- 40 ns	+/- 150 ns	+/- 200 ns	+/- 500 ns
T-cm-error	+/- 20 ns	+/- 40 ns	+/- 200 ns	+/- 300 ns	+/- 500 ns
T-docsis-error	+/- 40 ns	+/- 80 ns	+/- 350 ns	+/- 500 ns	+/- 1000 ns
T-source-skew	5 ns	10 ns	100 ns	200 ns	500 ns
T-hfc-error	+/- 7.5 ns	+/- 15 ns	+/- 50 ns	+/- 150 ns	+/- 250 ns
T-cm-cm-skew	100 ns	200 ns	900 ns	1500 ns	3000 ns

1.3.1.1. HFC error portion

Looking at Table 4 on error budget it can be seen that the HFC error is one of the smaller components. This is due to the consistent nature of HFC propagation delay. DTP/PTP account for the delays and variations from CMTS to CM due to fiber propagation delay, coax delay and remote node data conversion delays.

2. Best Practices for SLA testing

Best practices for SLA compliant Ethernet Service activation include physical layer through application layer testing. Full KPI testing including RFC-2544 or Y.1564 is recommended to validate proper EVC functionality. It is recommended to do testing of the Layer 2 Control Plane to validate that the control plane traffic flows transparently from end to end and ensure that the control plane traffic flows. As a test to validate the actual end user customer experience and optimize any customer equipment, RFC-6349 tests are recommended. The RFC-6349 tests throughput at layer 4 TCP traffic which is what actual client users will be experiencing.

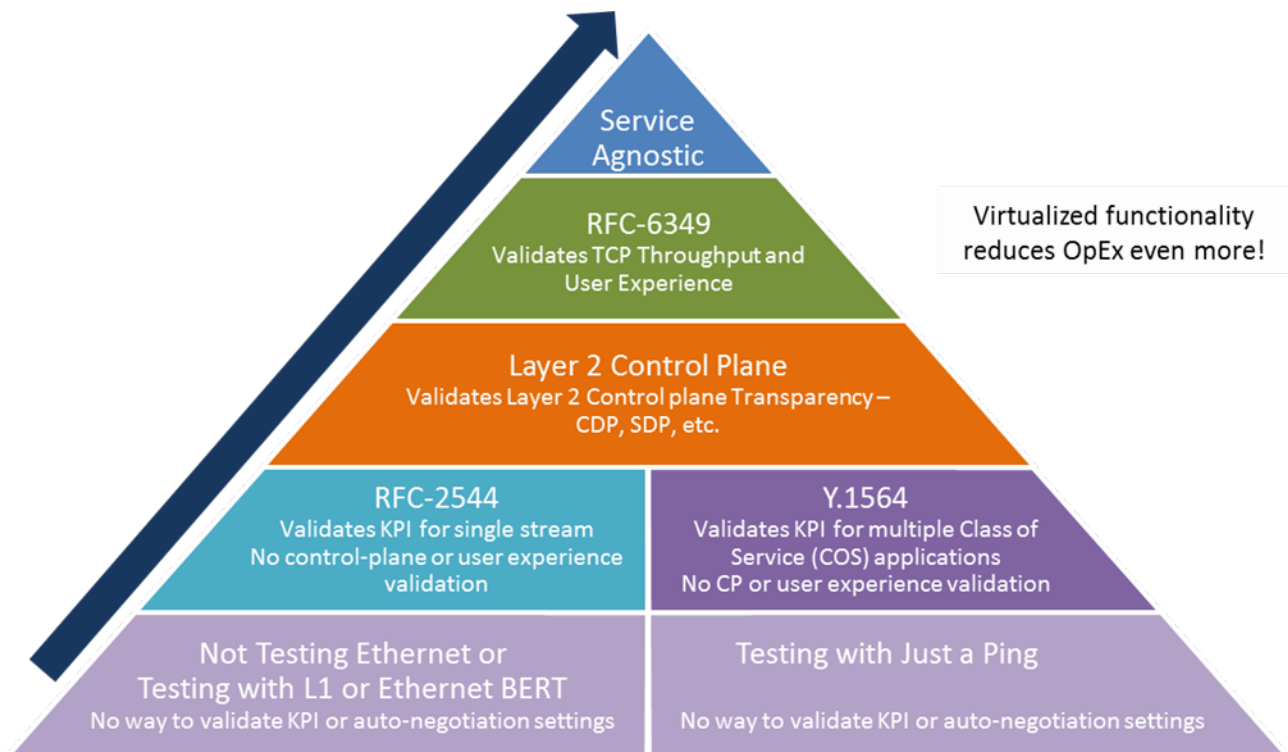


Figure 6 - Best practices for complete service activation

2.1. PTP/DTP timing

For technicians and engineers who must install Carrier Ethernet, Ethernet Backhaul, and PTP (IEEE 1588v2) circuits (like the DOCSIS DTP), testing the OWD, and 1588v2 performance referenced to CDMA and/or GPS receivers can save hours of troubleshooting by detecting asymmetric traffic ensuring proper handoffs for time-sensitive service-critical (mobile) applications. This solution can attain accuracies 10 times greater than most common SLAs, permitting service providers and operators to differentiate their offering and allowing network planners better understand delay tolerances affecting their applications.

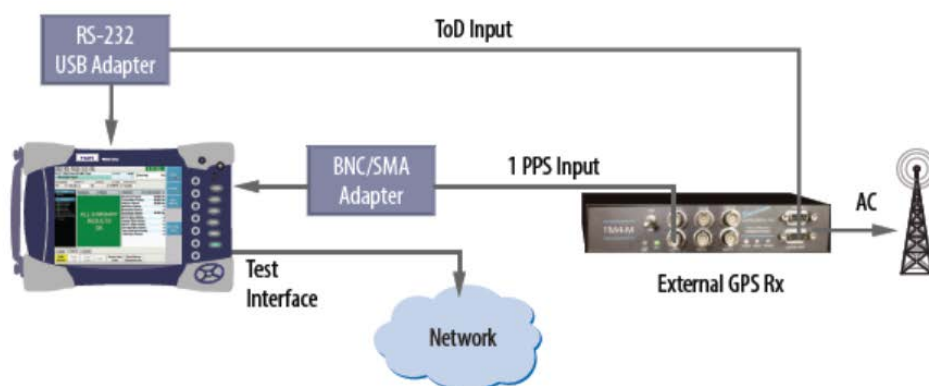


Figure 7 - OWD and 1588V2 testing connection using GPS source

2.2. OWD measurements for Mobile Backhaul SLA's and Carrier Ethernet

Measuring highly accurate one-way metrics, including OWD and packet jitter, in an Ethernet/IP backhaul scenario improves application troubleshooting and ensures thorough testing and verification of SLAs. Devices at the very edge of the network still can experience asymmetric delays. For example, in a mobile-voice application, increased delay may cause edge devices to buffer the information, thus smoothing out the speech. Unfortunately, unequal or asymmetrical delay can cause one side of the conversation to sound perfectly clear, while the other side appears to constantly talk over the speaker. Avoiding this requires verifying the OWD metric during installation and recording the measurement for future monitoring and troubleshooting as necessary. OWD measurements are important in validating the DOCSIS DTP system design supports the proper type of service and for example support an LTE advanced backhaul situation.




Figure 8 - PTP/DTP One Way Delay Testing with External GPS reference

2.3. Using automation with remote probe to speed up service activation

Service turn up normally involves sending a technician to the site to do the initial connection and activation. It is recommended to have the technician do an initial abbreviated connectivity and brief service test to ensure all EVC configurations and switch configurations are working properly. Lengthy testing can then be kicked off using virtual probes or physical hardware probes that can automate the RFC testing and be running multiple tests in parallel.




Figure 9 - Example of SLA Testing with Automation for Efficient Turn Up and Activation

Example of best practice for a typical workflow on cell site activation:

- Tech goes to a cell site
- Tech tests RF and DOCSIS network for proper operation and Throughput
- Tech install DOCSIS 3.1 Modem/Gateway
- Virtual Probe or HW probe is placed in wait for loopback mode
- Tech run's an abbreviated RFC-2544 test with a couple of frame sizes and only throughput testing to check for connectivity – 15minute timeframe
- Tech installs NID (if separate from D3.1 modem) & Cell Site Gateway
- Tech notifies the SYSTEM that there is connectivity and initiates test from Virtual Probe or HW probe.
- Full RFC-2544 test is run to NID/Customer equipment
 - NOTE 1: *This could be run from both directions: Portable to Virtual probe or Virtual Probe to Portable*
 - NOTE 2: *A quick test should be run on each configured VLAN*

Value of the workflow:

- Keeps Techs moving and working – Sites turned up faster
- Allows for full testing of EVC configuration (QOS, CIR, CBS, EIR, EBS, PIR, Frame Delay, Frame Delay Variation, and Frame Loss)

Conclusion

Legacy DOCSIS prior to DOCSIS 3.1 had limitations that prevented it from being a serious contender as a delivery mechanism for high performance SLA compliant Ethernet services. DOCSIS 3.1 provided additional signaling and timing mechanisms into the standard that equipment vendors and operators can leverage to create DOCSIS 3.1 systems capable of supporting key Ethernet services such as Advanced LTE. In particular, the DTP adoption of Ethernet PTP timing eliminates the problems associated with the delays and delay variations in legacy DOCSIS by making the entire DOCSIS system look like a border clock. When enabling SLA services, there are a set of best practices to validate the KPI's associated with the Ethernet circuit. Traditional tools, including RFC-2544 and Y.1564 are the cornerstone of EVC validation. Higher level Layer 2 control plane and RFC-6349 user experience validation should also be included. Service Activation for the EVC's can be optimized with automation using a testing solution that leverages virtual test probes and/or hardware test probes so technicians can rapidly progress from activation to activation.

Abbreviations

AP	access point
bps	bits per second
BERT	Bit Error Rate Test
BSOD	Business Service Over DOCSIS
CBS	Committed Burst Size
CIR	Committed Information Rate
CM	Cable Modem
CMTS	Cable Modem Terminating System
CPE	Customer Premise Equipment
DTP	DOCSIS Timing Protocol
EVC	Ethernet Virtual Circuit
EIR	Excess Information Rate
FD	Frame Delay
FDV	Frame Delay Variation
FEC	forward error correction
FLR	Frame Loss Rate
HFC	hybrid fiber-coax
Hz	hertz
ISBE	International Society of Broadband Experts
KPI	Key Performance Indicator
L2VPN	Layer 2 Virtual Private Network
LTE	Long Term Evolution (wireless service)
MEF	Metro Ethernet Forum
MTTR	Mean Time To Repair
NID	Network Interface Device
OAM	Operational Administration and Maintenance
OWD	One Way Delay

RFC	Request For Comment
SCTE	Society of Cable Telecommunications Engineers
SLA	Service Level Agreement
TLS	Transparent LAN service
UNI	User Network Interface
VLAN	Virtual Local Area Network

Bibliography & References

Data-Over-Cable Service Interface Specifications, DOCSIS® 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I11-170510, May 10, 2017

Data-Over-Cable Service Interface Specifications Business Services over DOCSIS®, L2VPN Development Guidelines Technical Report, CM-TR-L2VPN-DG-V01-121206

Metro Ethernet Forum: *Bandwidth Profiles for Ethernet Services*, Ralph Santitoro

One-Way Delay and PTP (IEEE 1588v2) Test Applications, Viavi Solutions
<http://www.viavisolutions.com/en-us/literature/one-way-delay-and-ntp-ieee-1588v2-test-applications-product-and-solution-briefs-en.pdf>

The Essentials of Ethernet Service Activation – Five important tests explained, Viavi Solutions
<http://www.viavisolutions.com/en-us/literature/essentials-ethernet-service-activation-overview-product-and-solution-briefs-en.pdf>

The Essentials of Ethernet Service Activation – Single-Service Enhanced RFC-2544, Viavi Solutions
<http://www.viavisolutions.com/en-us/literature/enhanced-rfc-2544-single-service-test-product-and-solution-briefs-en.pdf>

The Essentials of Ethernet Service Activation – Multi-Service Y.1564 SAMComplete, Viavi Solutions
<http://www.viavisolutions.com/en-us/literature/y1564-samcomplete-multi-service-test-product-and-solution-briefs-en.pdf>

The Essentials of Ethernet Service Activation – RFC 6349 TrueSpeed™ Layer 4 TCP Throughput, Viavi Solutions
<http://www.viavisolutions.com/en-us/literature/rfc-6349-test-truespeed-product-and-solution-briefs-en.pdf>

Zero Touch Service Assurance

A Technical Paper prepared for SCTE•ISBE by

Sean Yarborough
Sr. Director, Product Marketing
Lifecycle Service Assurance
Spirent
sean.yarborough@spirent.com

Introduction: Metro Ethernet Service Trends

Today's telecommunications marketplace is experiencing a significant growth in the number of deployments and the types of services being deployed in the metro Ethernet or business Ethernet space. For example, small cells are a driving factor behind the need for increased Ethernet backhaul services. One of the reasons for this surge is the promise that small cells will deliver higher quality voice, video, and data services than ever before, with lower deployment costs than macro cells. Likewise, the transition on cloud hosted applications, VoIP networks, and the overall increase in bandwidth utilization are driving demand from enterprises for higher speed services. In order to truly reap the benefits of the small cell promise, cloud hosted applications, or any new technology, providers must ensure the quality of service demanded by today's end users and carriers. This paper explores key metro Ethernet service trends, the challenges these trends create for service providers and the benefits of transitioning to zero-touch service assurance to help address these challenges.

Metro Ethernet Market Trends

1. Ethernet port growth

IHS Technologies is forecasting worldwide 100G port revenue to grow at a 137% compound annual growth rate (CAGR) from 2014 to 2019. This growth is a must in order to deliver the volume of data and services required in the network of the future.




Figure 1 – 100G Port Revenue Expected to Reach \$60B by 2019 (Source: IHS)

Another reason for the increased 100G surge is that the price per port has been declining over the past four years which is helping to driving demand. What used to cost tens of thousands of dollars, is now affordable to many providers, enabling them to offer cost-effective services above 10Gbps. And costs are only predicted to keep reducing, further increasing demand. As both the number of services being

deployed and the speed at which they are being deployed increase, providers must move to a holistic, automated, and intelligent approach to service assurance.

2. Stricter SLAs

According to a Global Service Provider Study by Infonetics entitled, Macrocell Backhaul Strategies and Vendor Leadership, “Latency is a very critical SLA metric, rated very important by 100% of respondents, followed by uptime/reliability, downstream bandwidth, jitter, and upstream bandwidth.”

Data services are becoming more and more mission critical, while the speeds at which they operate are only getting faster and faster. And with applications being outsourced and hosted in the cloud, the need for continuous connectivity is no longer a luxury; for many it is essential. Today’s consumers have clear expectations about the availability and the quality of their services and because of this, providers must track and be proactive about monitoring service level agreement (SLA) metrics.

3. Deployments of SDN/NFV Networks

Another issue that is driving increased usage and bandwidth consumption as well as an overall change in the market, is the migration to SDN and NFV. IHS interviewed current incumbent providers, competitive providers, independent/wireless providers, and cable operators who control 53% of the global telecom CAPEX and the universal position is that the industry is moving towards virtualization.




Figure 2 – 100% of Service Providers Surveyed by IHS Plan to Deploy SDN and NFV

No longer will applications run on dedicated, custom-built hardware. Instead they will be software-based micro-services running on a standard COTS “white box” platforms with compute resources. Virtualization will drastically change the way networks are deployed, managed, and maintained; therefore, a new level of service assurance will be needed to keep up with customer expectations in a virtual environment.

Service Provider Challenges

In this new paradigm, Service provider challenges to service activation and SLA management will fall into three categories. While these categories are not new, the way they must be addressed in this new environment is different. These challenges include:

- 1) Service Delivery
- 2) Trouble Management
- 3) SLA Management

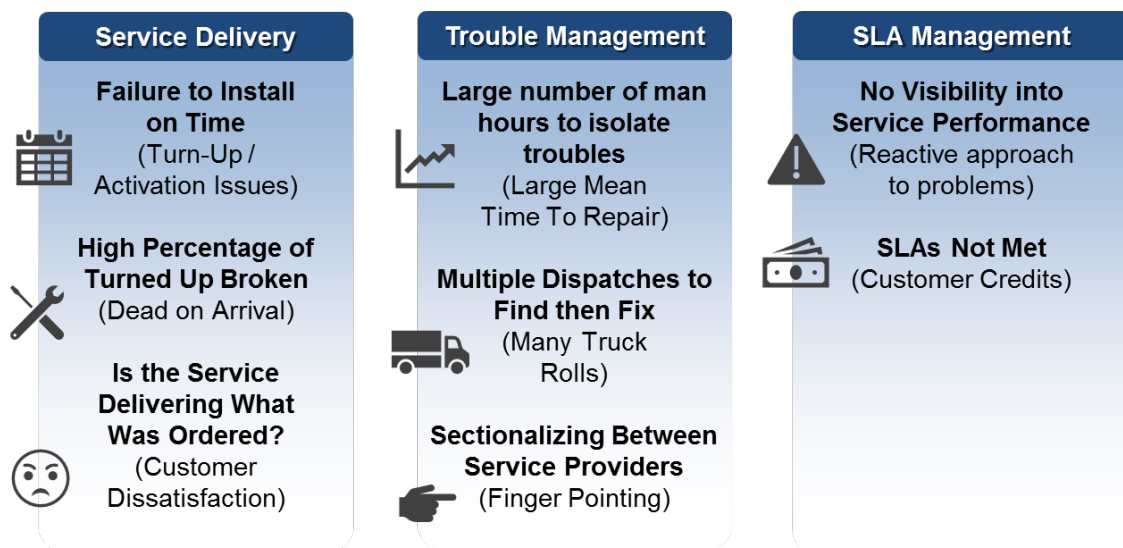


Figure 3 – Key Service Provider Challenges to Service Activation and SLA Management

1. Service Delivery Challenges

Service delivery challenges include failure to install on time (turn-up/activation issues), a high percentage of turned-up broken (dead on arrival) devices, and delivering the service that was ordered (preventing customer dissatisfaction).

Customers expect services to be turned-up very quickly (within hours or days, not weeks or months). This presents a challenge for traditional carriers who process orders by receiving a service order, putting that order into a queue, and scheduling a service technician dispatches to turn-up that service sometime in the next week. Today's customers want their services activated quickly and correctly the first time.

Another challenge with service delivery is trying to address the high percentage of services that are turned-up incorrectly or dead on arrival. As services become more complex, field technicians are not always equipped to install or troubleshoot services during their first customer visit. In other cases, the service may be “working”, but not to the performance/SLA guaranteed to the customer, thereby initiating a trouble call. Often equipment is replaced unnecessarily which is costly and not always an effective solution.

Providers need to prove to customers that the services they ordered have been delivered. Gone are the days where a technician would install the connection and then run a PING test to verify installation. Today, providers need to be able to present a clear indication of how the service is being delivered, what level of quality is present, and assurance that the SLA is being met.

2. Trouble Management Challenges

Trouble management challenges include a large number of man hours to isolate troubles (large mean time to repair), multiple dispatches to find and then fix (many truck rolls), and sectionalizing between service providers (finger pointing).

Service providers spend a great deal of time trying to isolate troubles and this results in many wasted man hours. Even the simplest of problems can take hours of wasted manual intervention. The lack of automated processes and automated workflows only add to extended time to repair. And in some situations, SLAs are being violated just because of the lack of actual data from the network.

Some providers still subscribe the theory that “when in doubt, dispatch out.” This mentality creates a lot of truck rolls and often these truck rolls involve unnecessary equipment replacement, which isn’t a cost-effective way to manage these services.

When you have services that are delivered through an alternate vendor, type II services for example, the idea of being able to clearly sectionalize between providers can be quite a challenge. When services go to a third-party, the visibility into that service becomes severely limited. This is even more so when you are delivering metro Ethernet or business Ethernet services because the cable infrastructure, cable modems, etc. don’t necessarily possess the capability to provide the same level of testability, service assurance capability, troubleshooting capabilities that traditional service provider premise equipment (NIDs, network terminating equipment) have historically provided.

3. SLA Management Challenges

The challenges with SLA Management include poor visibility into service performance and not meeting SLAs.

The lack of visibility into service performance leads to a reactive approach to problems. This reactive mode is time consuming and unproductive and usually end up leading to an SLA violation or customers waiting longer than necessary for service restoration.

SLAs have liquidated damage clauses, penalties, etc., if violated. Having to pay fees for SLA violations drastically cuts into the profitability of the service.

Current Approaches to Service Activation & SLA Management

So how are service providers dealing with service activation and SLA management challenges today? For service activation, many providers are still dispatching multiple personnel who use handheld devices that can only test one circuit at a time and can be costly and don't scale well.

SLAs are often monitored with NIDs, which can certainly be a viable methodology, however each network interface device (NID) manufacturer has a different element management system (EMS) which work in slightly different ways, have different KPIs, and different capabilities. As a result, there are multiple systems that have to be simultaneously managed to truly have full coverage of the network. It's also very difficult to sectionalize when SLAs are not met using this approach because while NIDs can test end-to-end or point-to-point, they are limited in their ability to segment the network and isolate faults. This makes troubleshooting problems difficult and ultimately leads to finger-pointing between the access vendor or type II provider and the provider.

Many providers use passive probes to manage SLAs. The problem with these tools is that they are very limited in terms of what type of data they can collect. Often, they are limited to the ingress UNI and do not provide an end-to-end view of the performance. Additionally, these probes only provide analysis when user traffic is available, so there is very limited visibility into issues such as loss, true availability, latency, jitter, and delay – issues that are extremely important to end-users. This solution also does not adequately support virtual and hybrid networks, which many providers are moving towards. The passive nature of these tools don't migrate well into a virtualized environment. The concept of “sniffing” a link in a virtual environment has some technical challenges.

Best Practices for Zero-Touch Automation

Spirent has compiled some recommendations for best practices which aim to achieve the goal of “zero-touch” service assurance for Ethernet and IP services. These best practices revolve around several key points:

- 1) Centralized, automated & intelligent Lifecycle Service Assurance
- 2) Use of global, industry-wide test standards (e.g., leverage CPE embedded features)
- 3) End-to-end visibility, troubleshooting & segmentation (e.g., “Dispatch to fix, not to find”)
- 4) Integration of all dependent systems within a single Service Assurance Test Controller
- 5) Scalability to handle drastic service growth w/o large increases in OPEX

1. Centralized, Automated & Intelligent Lifecycle Service Assurance

The first best practice is the concept of centralized, automated and intelligent Lifecycle Service Assurance. This model purports that service assurance should encompass the entire network lifecycle from design, to onboarding, deployment, operations, maintenance, and back to design. Providers see the

need to provide a closed loop from design into operations and back to design to effectively manage network services as they migrate to the virtual space.



Figure 4 – Complete Lifecycle Service Assurance

There are 4 key functions that must be in place in order to implement a complete Lifecycle Service Assurance platform throughout the service lifecycle in the production network, which includes: active service activation, active performance monitoring, and on-demand troubleshooting (with passive monitoring and active testing).

Active service assurance enables consistent and repeatable activation tests, centralized storage of the service birth certificate, automated network element control, and multiple test methodologies that can be embedded into a service activation workflow.

Active performance monitoring provides the ability to really monitor the performance and availability of the network in real-time, 24x7, with comparison against SLAs on a service-by-service basis, and alarming and thresholding accordingly and the ability to push this data to external systems whether that is done by traditional methods by pushing reports on a regular basis or streaming that telemetry into a data link that can be accessed from northbound systems for analytics, policy mapping, etc.

Active performance monitoring allows for scalable, 24x7 real-time analysis of the network, monitoring and reporting, SLA and availability monitoring enables SLA management, native web GUI and NB interface to existing OSS.

Once active monitoring is in place, passive monitoring and active on-demand troubleshooting are used to isolate faults. Because of this integration, if an issue arises, the system can automatically execute a troubleshooting workflow, identify where the problem resides, and isolate that segment of the network before a tech is dispatched. If this can be done in an auto fashion, the time a tech is engaged is reduced from hours to minutes. This could be the difference between violating an SLA, having to pay liquidated damages, or not.

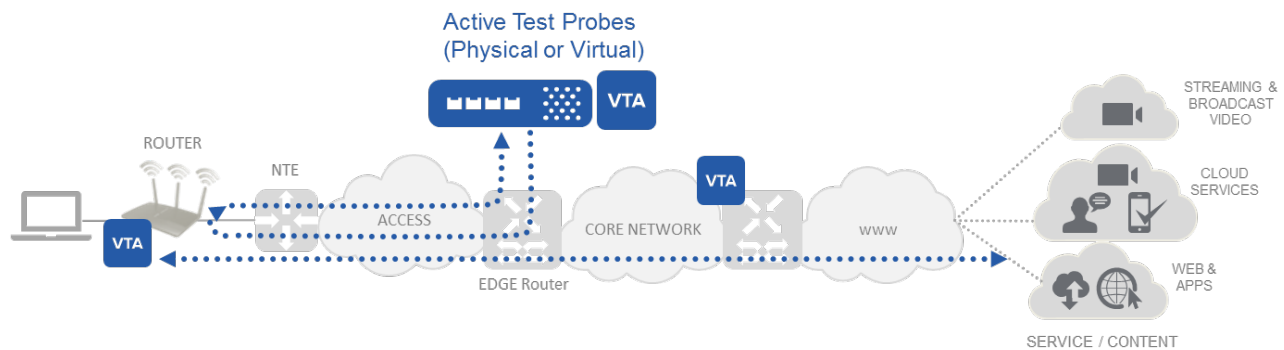


Figure 5 – Active test probes (physical or virtual test agents(VTAs)) for end-to-end service testing and testing of network segments.

2. Use of Global, Industry-wide Standards

The use of global, industry-wide test standards such as embedded CPE features and embedded network element features (e.g., MEF 46 latching loopback, Ethernet OAM, TWAMP, TR69, TR143, etc.) are crucial to creating a zero-touch service assurance platform.

In our global market, service providers, network owners and equipment vendors must work together to deliver telecommunication services to end users. Guiding them are the industry standards and requirements that help ensure services are reliable, cost effective and deployed properly in a timely manner.

The Metro Ethernet Forum (MEF) has developed a set of standards so that providers and vendors can work together in a harmonious way, using any vendors' equipment, and know that the same KPIs apply. Furthermore, this approach enables service providers to deploy service Assurance at the edge/core of their networks while leveraging already deployed CPE devices at the customer premises, providing a cost-effective solution for true end-to-end Service Assurance.

While some standards have been in use for years, like RCF 2544 for benchmarking and 802.1ag for testing loopback, delay and multicast loopback, there has been an evolution to more sophisticated testing to keep up the growing sophistication of the services and the supplication of the user's expectations.

Standards like Y. 1564 has been developed to replicate user traffic more realistically though a use of EMIX and burst testing and RC 6349 can perform TCP throughput testing. While these standards may have been developed for typical service activation, they also apply to service assurance for continuous network monitoring and guarantees that your service is delivering great quality, meeting customer expectations, and satisfying SLAs.

3. Automated End-to-end Troubleshooting & Fault Segmentation

By implementing a centralized test management system, providers can look through a single pane of glass to access all of their provisioning, trouble ticketing, analytics, reporting, and inventory systems. The test

manager is also automated and can automatically launch a set of standardized tests without human intervention.

True end-to-end visibility is also imperative to have a true view of the network. With access network and aggregation networks operated by local providers and the core network by a national provider, all using multi-vendor interfaces such as Accedian, Ciena, Cisco, ALU, Juniper, Arista, Rad, etc., there needs to be a seamless way to communicate with the variety of embedded test function in NIDs and other network elements.

By having true end-to-end visibility into troubleshooting and segmentation, faults can be isolated before a trouble ticket is generated. This allows any dispatch to have the goal “to fix, not to find” and creates faster time to resolution and ultimately higher customer satisfaction.

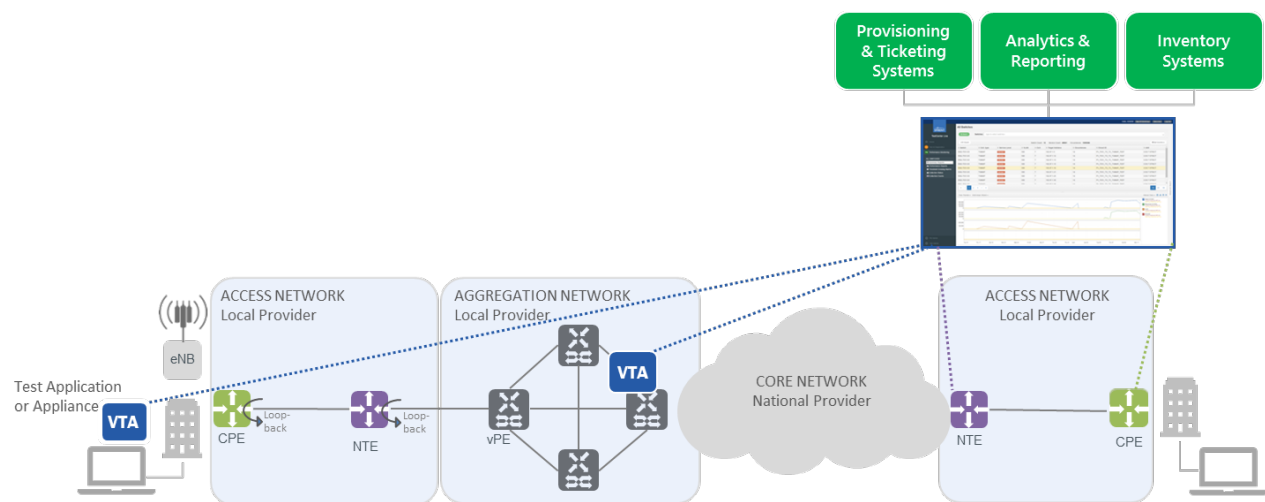


Figure 6 – Automation of End-to-End Troubleshooting & Fault Resolution

4. Integration of Systems with Service Assurance Test Controller

Having a single interface that controls all tasks simplifies the process of deploying, onboarding, operating, and maintaining services. Instead of applying a manual “swivel-char” approach the management of multiple systems from multiple locations, dependent systems can be integrated within a single Service Assurance Test Controller enabling automated ticketing and reporting and provide a complete “zero-touch” solution to provisioning and segmentation.

This also allows for a much more holistic approach to fault isolation and troubleshooting. This solution can automatically segment what area of the network the fault resides and automatically include this information in the trouble ticket which reduces time to repair or prevents a dispatch altogether saving time and money.

For example. inventory records can be retrieved by the service assurance system automatically and integrated with the alarm, trouble ticketing OS, customer trouble ticket portal, all reducing the amount of time to identify faults, isolate faults in a repeatable manner.

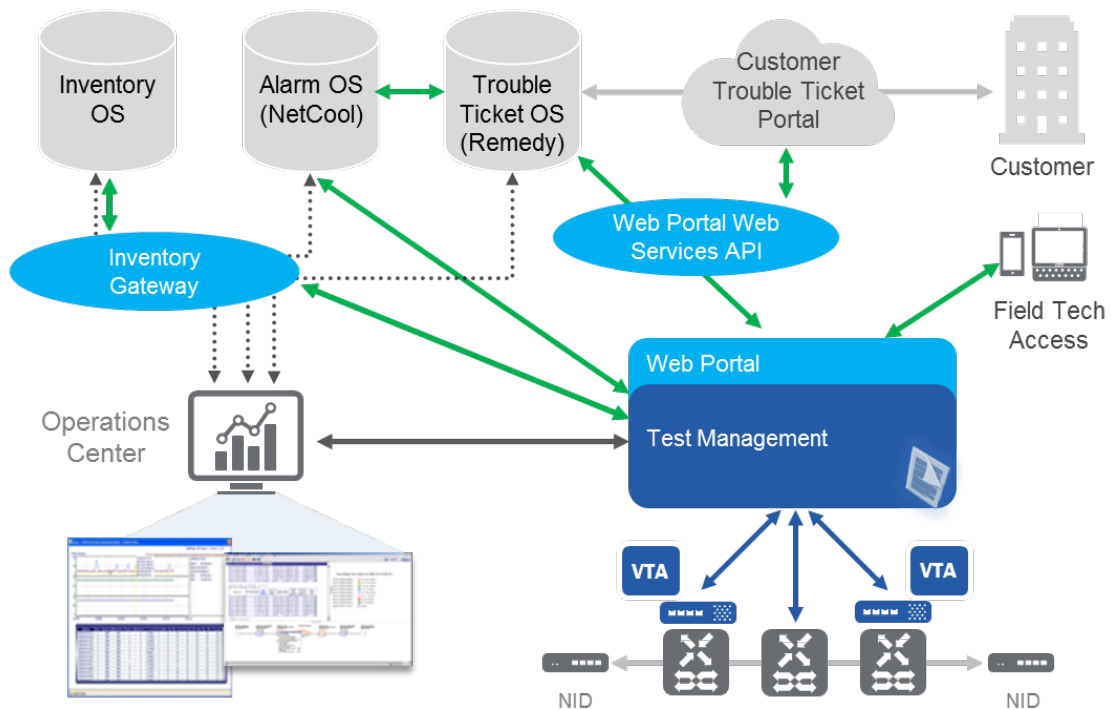


Figure 7 – Integration within a Single Service Assurance Test Controller

Integration of the service assurance test controller to all systems supporting an Ethernet service turn-up workflow enables workflow automation that can reduce test time more than 80% while also eliminating manual errors.

Table 1 – Test Controller Integration Enables Automation That Can Reduce Test Time by 87%

Task	Manual	Automated
Inventory Query	1 min	10 sec
Validate Configuration	20 min	20 sec
Enable Loopback on NTE	1 min	10 sec
Add Test Agent / Appliance to Service	5 min	30 sec
Execute Service Tests	4 min	4 min
Return Service to Original Config & Validate	10 min	40 sec
Total Time	41 min	5 min 50 sec

5. Scalability to Handle Drastic Service Growth

As Metro Ethernet services continue to grow and virtualization efforts accelerate this growth, the scale of service deployments and changes in the network will be exponential. With surge in the number of services and the frequency of changes, it will be impossible to manage without a large increase in OPEX unless a centralized and automated system is in place that can scale to the magnitudes needed.

While the promise of such a large-scale network deployment is exciting, providers are will have more users than ever before and various types of SLAs, so they will need a controller with the ability to simultaneously have multiple interfaces to active new customers and perform continuous active monitoring for existing customers.

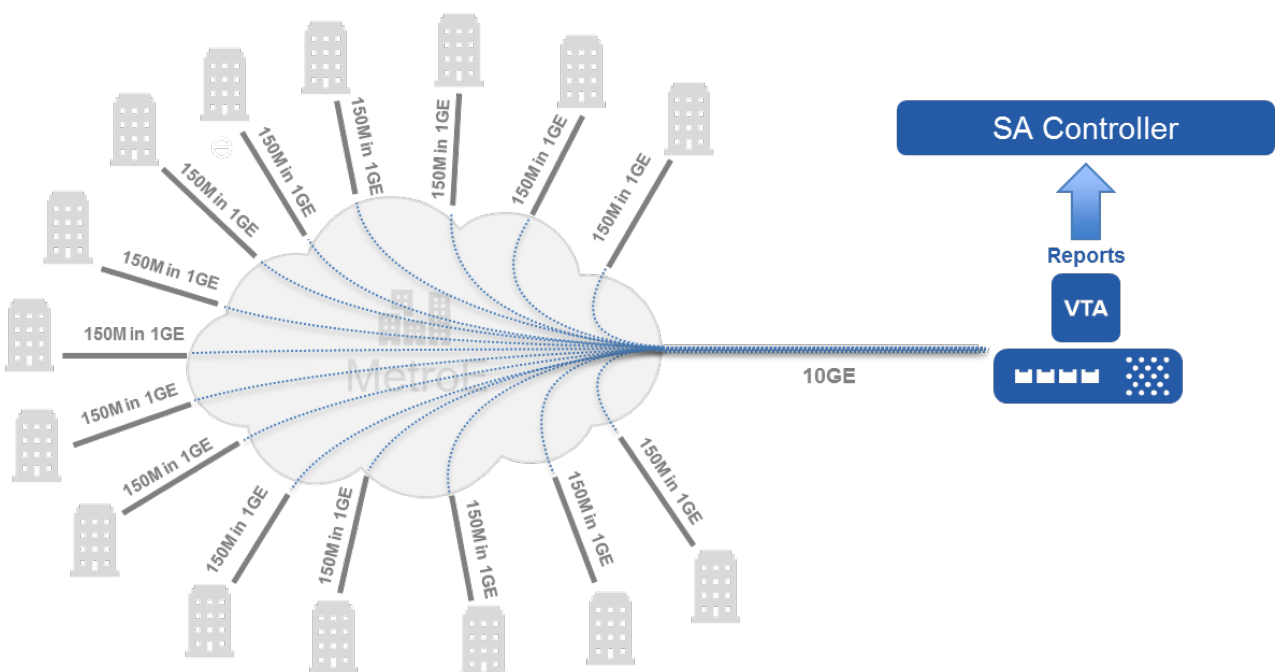


Figure 8 – Simultaneous Metro Ethernet Service Activation

Conclusion: Benefits & ROI

Spirent has worked closely with multiple tier-1 providers that have implemented automated service assurance systems and experience significant financial benefits. In one case, a tier 1 provider of Ethernet services implemented an automated service assurance system to monitor backhaul services provided to a mobile network operator. The mobile network operator experienced poor performance on the backhaul links and submitted a multi-million dollar SLA violation claim to Ethernet service provider. The Ethernet service provider could use SLA management and detailed diagnostic data from their service assurance system to prove the root cause of the SLA violation was in the mobile operator's network, saving millions.

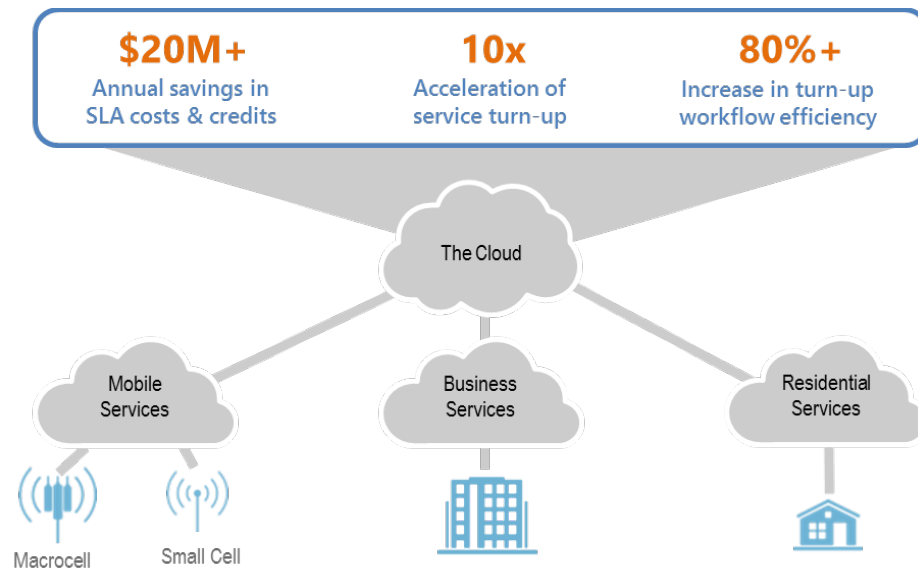


Figure 9 – Automated Service Assurance Delivers Significant Benefits.

In another case, a mobile network operator needed to roll out thousands of small cells and activate backhaul links for each of these new sites. The operator implemented a service assurance system to automate service activation testing of backhaul. As a result, the operator was able to accelerate the launch of small cells by an order of magnitude, from 100 to more than 1000 cells per week. In addition to improving the speed of deployment, the operator used automation to improve the efficiency of their activation workflows by 80%, enabling them to increase the rate of deployment without needing to add any additional resources to the activation teams.

As providers adopt NFV and SDN technologies, network configuration become much more fluid and dynamic. Workflows such as turn-up verification, SLA monitoring and issue resolution must be automated, as manual techniques simply won't run fast enough to keep up with network changes. In addition, since physical networks will persist for years, service assurance will need systems that unify these workflows across hybrid physical-virtual networks. As a result, automated service assurance is in the process of transforming from a highly beneficial, non-mandatory capability to an essential requirement.

Abbreviations

AAV	Alternate Access Vendor
SDN	Software-Defined Network
NFV	Network Functions Virtualization
PM	Performance Management
SAT	Service Activation Test
SLA	Service Level Agreement
TWAMP	Two-Way Active Monitoring Protocol
VTA	Virtual Test Agent

Bibliography & References

IETF RFC 5357: Two-Way Active Measurement Protocol (TWAMP), October 2008
(<https://tools.ietf.org/html/rfc5357>)

Technical Specification MEF 48: Carrier Ethernet Service Activation Testing (SAT), October 2014
(https://mef.net/Assets/Technical_Specifications/PDF/MEF_48.pdf)

Technical Specification MEF 46: Latching Loopback Protocol and Functionality, October 2016
(<https://wiki.mef.net/display/CESG/MEF+46+-+Latching+Loopback+Protocol>)

Virtualizing Managed Business Services for SoHo/SME Leveraging SDN/NFV and vCPE

A Technical Paper prepared for SCTE•ISBE by

Ajay Manuga
VP Engineering
Benu Networks
amanuja@benunets.com

Introduction

The cloud managed services market has been flourishing in recent years due to the advancements in cloud computing, virtualization and mobility services. This market will grow from \$35.54B in 2016 to \$76.73B by 2021, at an estimated CAGR of 16.6% from 2016 to 2021. Managed services allow businesses to outsource most aspects of their IT infrastructure to a service provider. This can reduce the recurring in-house IT costs by 30-40% and bring about 50-60% increase in efficiency.

While many of the managed service offerings have been focused on medium to large enterprises, there is a major opportunity for Multiple Systems Operator's (MSO) to offer a solution for the SoHo/SME market. The MSO can offer compelling services to this segment, like a custom guest Wi-Fi splash page, social network sign on, promotions management, customer analytics, multi-site VPN, and more.

This document covers a virtual CPE (vCPE) architecture that provides enterprise-class, cloud based value added services for the SoHo/SME market. Using technologies common in data center networking such as virtualization, SDN and overlay tunneling this real-world MSO deployment was able to leverage their existing business infrastructure and premise equipment on top of their DOCSIS network. This Virtual Service Edge (VSE) solution greatly amplifies the types of business service offerings for the operator. The platform is designed to be owned, controlled and maintained by the service providers, while reducing their time to roll out and manage these new enhanced service offerings. This architecture seamlessly integrates into a service provider's OSS, BSS, customer care processes and other back end systems and the offering can be combined with other existing services like residential homespot and outdoor Wi-Fi.

Managed Business Services

Figure 1 below shows current service provider architecture for SoHo/SME market where network intelligence and control functions are hosted by customer premise gateway. This architecture provides little to no visibility to the service provider of actual devices/users on the premise or their use and demand of network resources such as bandwidth; so, any value added services are delivered at a whole premise basis, i.e., parental control is usually implemented using DNS and everyone is subjected to the same level of restriction. Lack of intelligence about customer's usage of the network bandwidth prevents MSO's customized marketing efforts to achieve higher level of acceptance of value added services.

Furthermore, with this architecture, the service providers need to think long and hard about introducing new products and services, as they need to consider capabilities of CPE vendors and types of CPEs in the deployed base before they can determine if a new product or feature can be supported. This dependency has both cost in money and time implications to new product introductions, with typical new time taken being anywhere from 9 to 18 months.

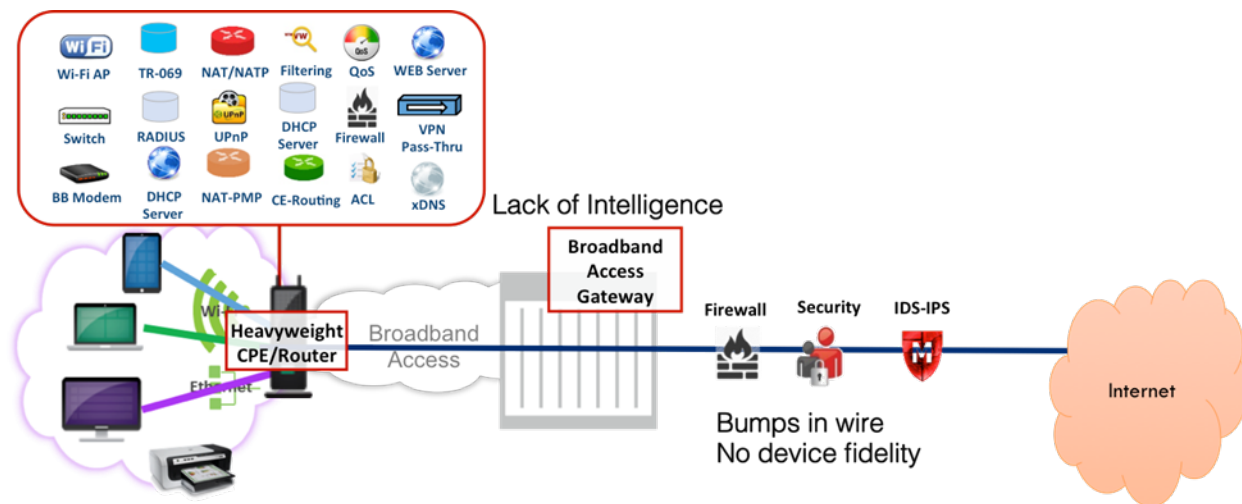


Figure 1 - Legacy Architecture

Virtual CPE solutions moves complex network functions from the physical CPE into a Virtual Network Function (VNF) running at provider edge. Benu VSE's vCPE architectural framework provides visibility into the devices and networks at the SoHo/SME premise, and allows greater agility and flexibility to roll out highly programmable, fully virtualized service. The VSE unlocks the IP router logic previously embedded in vendor-specific CPE, and consequently out of reach of service providers' control, and makes it open, available, and part of the service provider's native service offering. Conventional customer gateways are repurposed as lightweight, agile NTUs (network termination units).

The lightweight premise-based equipment requires minimal configuration, or other change; the services are fully instantiated in the service provider network. This simple but fundamental architectural transformation has a profound impact on the pace of innovation and the economics of new services.

This vCPE solution enables service providers to accelerate service velocity by reducing dependencies on CPE and helps providers improve the bottom line and accelerate time-to-market by avoiding expensive and time-consuming CPE evaluation, certification, and integration and support lifecycles. The solution also streamlines customer support by moving business logic from the customer premise into the service provider edge/cloud. This fundamental shift enables the service provider with full visibility all the way to the subscriber device.

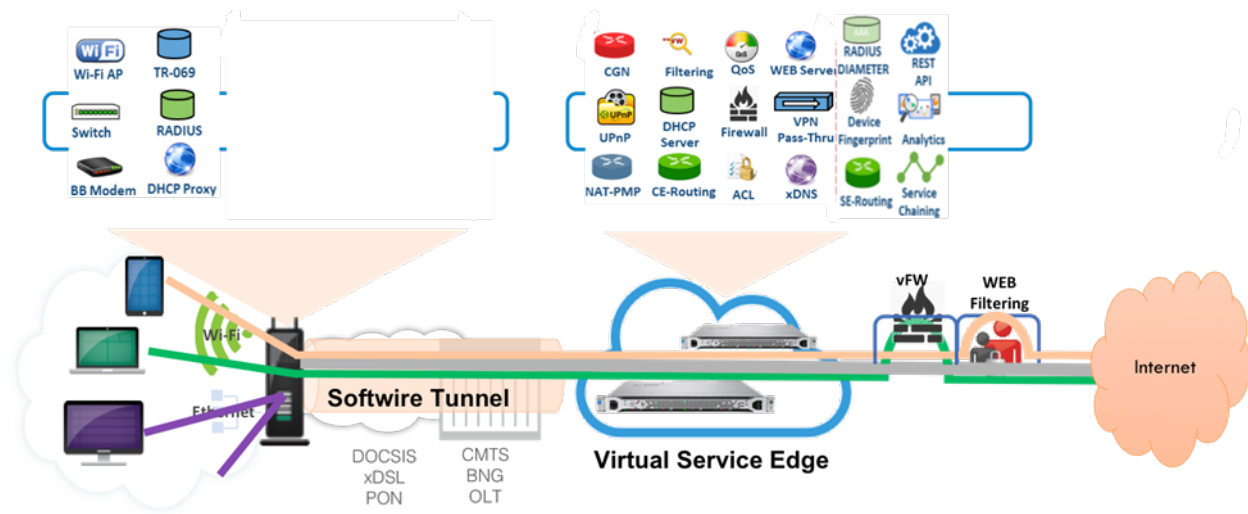


Figure 2 - vCPE Architecture

1. Network Architecture

Like any other network transformation, this architecture requires changes in the existing network elements, but not a lot. The network architecture consists of Smart bridge CPE deployed at customer premise and VSE solution deployed in the cloud or service provider edge network. The two elements transfer layer 2 frames via Soft GRE tunnel.

1.1. CPE

The CPE is enabled into a “smart bridge” mode. In this mode, CPE is acting like a switch on the LAN/wireless side, and simply puts any WAN or broadcast traffic into an overlay tunnel (ie; softGRE) toward the vMEG (Virtualized Multi-service Edge Gateway). Conversely it takes the incoming traffic from the overlay tunnel and forwards the traffic on the LAN/wireless side. The operations and lifecycle management on the CPE does not change at all. Simply a firmware upgrade to allow smart bridge functionality is required. By doing so, all the other CPE network functions are disabled. All broadcast traffic including DHCP packets are forwarded to the vMEG nodes running at the customer edge. The smart bridge CPE tags the layer 2 frames with different VLAN IDs, each ID corresponding to a specific segment or slice of the customer premise network. For example, all traffic to and from business or customer private devices is tagged with one VLAN ID whereas the traffic from visitors is tagged with another VLAN ID.

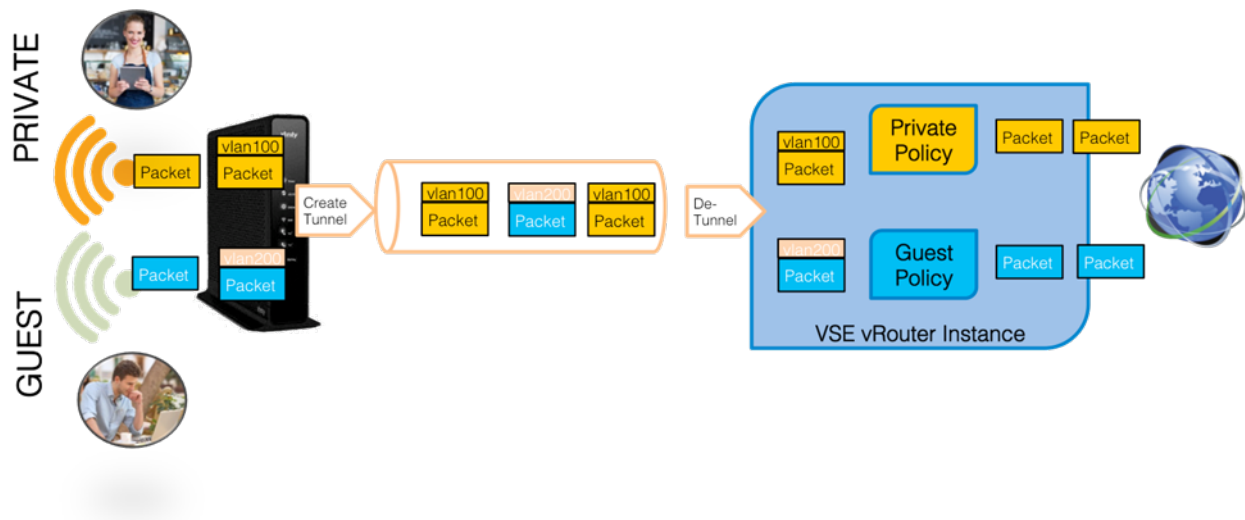


Figure 3 - SmartBridge and Soft GRE tunnel

1.2. User Portal

User portal provides customer ability to view and change existing services and subscribe to new ones. It is deployed in a typical web services model and front ended with an external firewall. When customers access the default home gateway using a browser, they are redirected to this portal, login to their accounts and manage their subscription, and perform all administrative functions for their business service. Bringing intelligence and analytics to the customers enables them to differentiate the value added services that works well from the ones that don't and tethers them to the superior service offerings and leads to greater overall satisfaction.

1.3. Service Provider Admin portal

Service provider has ability to view each business customer's network and capture packets. The service provider gets holistic view of the entire network like number of businesses connected, average bandwidth per business, total number of devices on the network and many more metrics. Using this portal, the service provider can implement policies or behaviors desired based on device type such as IOT and set up default business account settings for customers.

The portals are a natural place to view analytics by business customer at the customer premise level and by the operator at the customer premise level and also at operator network level.

1.4. Virtual Service Edge

Virtual Service Edge is a collection of Virtual Machines running together to provide a high scale multi-tenant vCPE solution in the NFV environment. It consists of mainly two VNFs/components: SSC and vMEG.

1.5. Subscriber Session Controller (SSC)

SSC is a database centric application running on Linux based VMs and acts as aggregate subscriber policy database and control functions for the service provider. SSC maintains subscriber accounts and their policies configured via REST API or portals. It also stores network and device usage analytics reported by vMEGs. It interacts with user and service provider portals and north bound service orchestrators. And it provides customer and device specific policy information to the policy enforcement engines (vMEG).

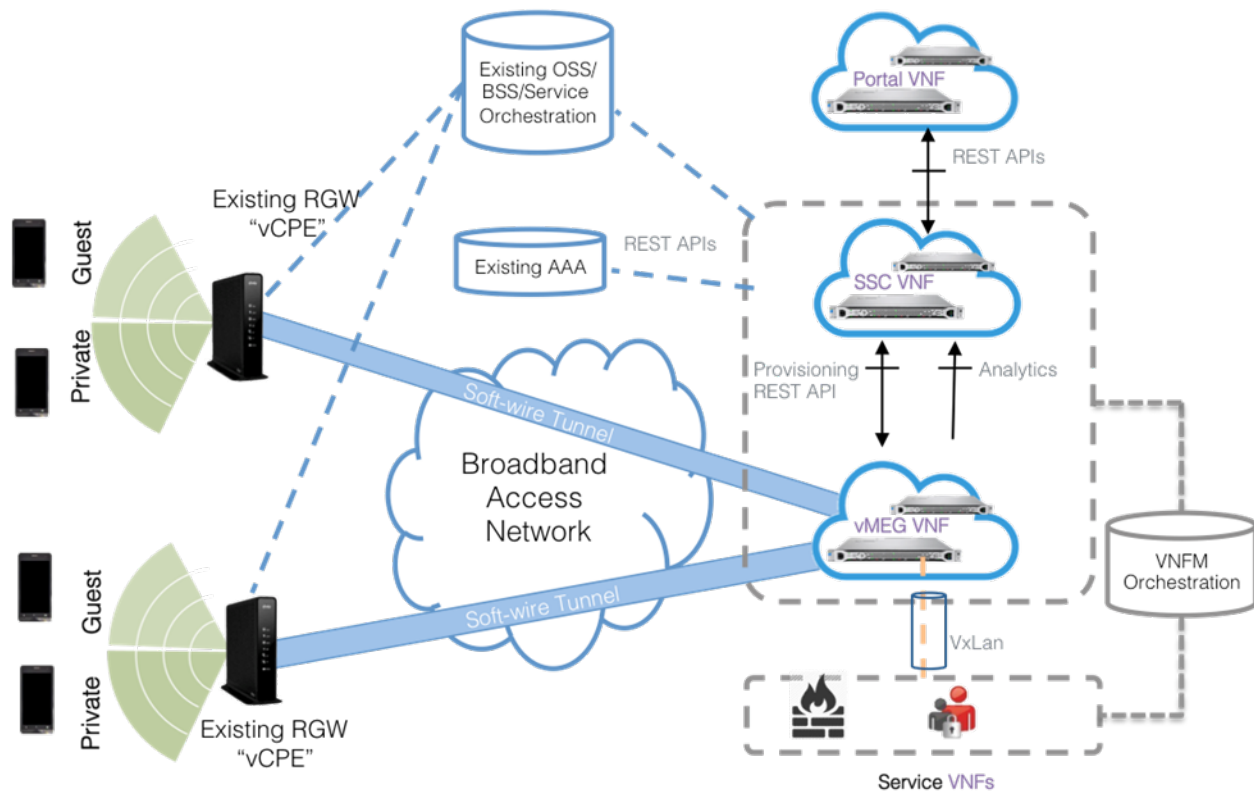


Figure 4 - VSE Solution

1.5.1. Virtual Multi-Service Edge Gateway (vMEG)

vMEG is the high performance and throughput packet processing and policy enforcement virtual machine simultaneously performing tens of millions of packets per second from and to multiple businesses. This high throughput data plane virtual machine is based on Intel DPDK framework, and especially designed to handle real time traffic such like gaming and voice/video at low latency and jitter.

Each vMEG instance can handle multiple tens of thousands virtual router functions in a single VM. Many of these vMEG virtual machines are clustered together as one VNF to create a very high capacity Managed Business Service solution in an NFV framework that can support millions of devices across multiple businesses.

1.6. Service VNFs

One of the key tenets of vCPE architecture is its ability to add and remove service/network functions to a device, VLAN or business on need basis. With this architecture new services are added and stitched using Service Function Chain(SFC) framework over VxLAN and NSH tunnel.

2. NFV Architecture

The following figure shows VNFs for the Managed Business Services in an NFV framework diagram. Each of the elements scale up and down independently to provide flexibility and agility in the service provider network.

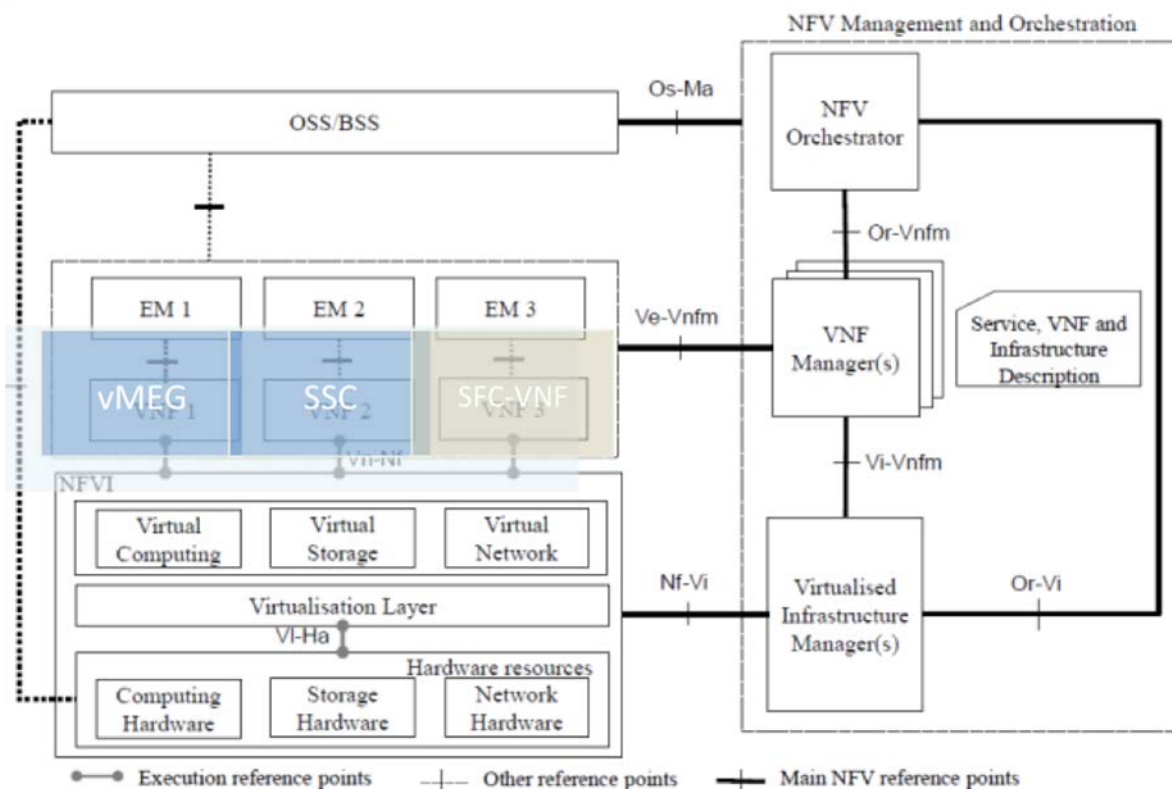


Figure 5 - NFV framework

Conclusion

Highly scalable vCPE solutions have moved from the labs to the real world deployment. Along the way, several complex issues have been worked out to make this a reality, with numerous improvements in

performance and scale in a virtualized dataplane environment with close collaboration with Intel and VMWare, and integration of the MSO's back-end OSS/BSS functions. All these innovations would amount to nothing if the customer experience has not been improved. Using this solution, small business owners can easily opt in for new value added services from the service provider, like guest network services, social media integration, parental control, enterprise grade security, promotions and vouchers, etc, and are armed with intelligence of their network like top site usage, repeat guests and their behaviors.

Service provider retains control over several key aspects of the network functions and their customer network intelligence, without requiring big changes in their existing underlay network.

Abbreviations

AP	Access Point
bps	bits per second
DPDK	Data Plane Development Kit
FEC	Forward Error Correction
HFC	Hybrid Fiber-Coax
HD	High Definition
Hz	Hertz
IOT	Internet Of Things
ISBE	International Society of Broadband Experts
MBN	Managed Business Networking
MSO	Multiple Systems Operator
NFV	Network Function Virtualization
NSH	Network Service Header
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networking
SME	Small Medium Enterprise
SSC	Subscriber Session Controller
vCPE	Virtual Customer Premise Equipment
VM	Virtual Machine
vMEG	Virtualized Multi-Service Edge Gateway
VNF	Virtual Network Function
VSE	Virtualized Service Edge

Delivering High-performance Business Services over a Dynamic Optical Infrastructure

A Technical Paper prepared for SCTE•ISBE by

Fady Masoud

Principal, Product and Technology Marketing
Infinera

555 Legget Drive, Suite 222, Tower B, Ottawa, ON, Canada K2K 2X3
fmasoud@infinera.com

Introduction

Driven by the need for greater productivity and lower costs, enterprises around the globe are moving their applications to the cloud. Today, nine out of 10 enterprises are using at least one cloud application to increase productivity and reduce cost, a fact that is not surprising when 84 percent of chief information officers report that they have cut application costs by moving to the cloud¹. Compounded by the need for scalable bandwidth driven by the proliferation of large volumes of digital content and applications, enterprises are turning to a new hybrid cloud network model, with applications delivered from an abstracted cloud services layer that bridges private and public cloud infrastructures. This article describes how cable operators can use the latest software and hardware innovation in optical networking to deliver advanced, secure and high-performance business services to the enterprise, such as software defined capacity (SDC), optical private network virtualization, data encryption and many others, to underpin the deployment of cloud enterprise applications. Common challenges and solutions for cable operators as they build and deliver key enterprise applications are also presented, such as business continuity/disaster recovery (BC/DR), data mirroring, off-hours/off site data backup and many others. This paper also describes best practices and emerging operational models to reduce first-in costs as well as recurring operational expenditures.

Content

1. An Industry Undergoing Massive Transformation

Cloud-based applications are impacting multiple aspects of the enterprise business - from the products they manufacture to the services they offer and even to the way their employees interact with each other or with customers and partners.

- **Connected products:** More products than ever are designed and manufactured to be online. 30 billion devices are expected to connect to the Internet in 2020². From connected cars to home appliances to smart sensors on city streets, products are becoming more sophisticated.
- **Connected services:** Services enterprises are also elevating their portfolios by leveraging wireless and wireline connectivity to provide advanced offerings such as remote monitoring and emergency response, intelligent home and business surveillance and many others.
- **Connected employees:** The workforce is more connected than ever before in day-to-day operations, relying on cloud applications for document management and sharing, social media-based employee interaction and video conferencing for meetings and training.

Enterprises need connectivity with the highest levels of performance (high capacity, low latency, high reliability, agility, etc.) across short (local area network, or LAN), medium (wide area network, or WAN)

and long distances, putting the transport network at the heart of the enterprise evolution to the cloud and the Internet of Things (IoT).

2. A Transformation That Brings New Challenges

The enterprise's journey through this transformation creates many challenges for cable operators offering enterprise services and applications. Better network performance, on-demand bandwidth consumption models, and the ability to combat cyber-attacks are a few examples of emerging requirements for cable operators. The following paragraphs briefly describe imperatives relating to these concerns and their impact on the cable operator's business and operational health.

- **Scale the network:** It can take 45 to 60 days for a simple bandwidth upgrade, not to mention the additional cost of equipment. This places significant pressure on cable operators' network planning teams, especially for unexpected network events.
- **Decrease the cost of operations:** Enterprise migration to the cloud and the evolution of operational and business models are fueling unprecedented demand for bandwidth. Traditionally, more bandwidth requires more money and more complexity, leading to significant hikes in capital and operational expenditures (CapEx/OpEx).
- **Protect enterprise data:** As more content is being pushed to the cloud, cyber-attacks and data breaches are becoming frequent occurrences. The annual damage to the U.S. economy caused by cyber-attacks is estimated to be up to \$100 billion³. Cable operators must protect enterprise customer data carried over the network from intruders and hacking tools.
- **Enhance network performance:** Latency, capacity and transactions per second are all key concerns cable operators must address when servicing enterprise customers. Networks must meet or exceed performance requirements dictated by enterprise applications and must be agile to support dynamic demand for bandwidth and any change in network topology, including connectivity to new offices or data centers.
- **Minimize downtime:** Cable operators' network outages can be disastrous to enterprises, resulting in significant loss of revenue, massive disruption to business operations and a major impact on customer loyalty. The average downtime costs vary across industries, from approximately \$90,000 per hour in the media sector to about \$6.48 million per hour for large online brokerages, according to Information Management magazine⁴.

Overcoming the above challenges requires intelligent high-capacity optical networks that offer the flexibility, scale and programmability to meet bandwidth demands and ensure the highest levels of availability and security while lowering operating costs.

3. Technology Enablers for Cable Operators

New advances in software and hardware are allowing cable operators to broaden their addressable enterprise markets by offering new services with the highest level of performance, flexibility and security, as well as advancing to dynamic and on-demand operational and service consumption models. The optical technologies that enable these new services provide the foundation for intelligent high-capacity optical networks and are described below:

- **High 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 and thus increase capacity and maximizing fiber utilization. However, the current growth in internet traffic and enterprise migration to the cloud are demanding a whole new level of scalability. A new innovation, called super-channels, 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 1. Super-channels overcome three fundamental challenges:
 - How to scale bandwidth without scaling operational procedures
 - How to optimize DWDM capacity and reach
 - How to support the next generation of high speed services such as 100 Gigabit Ethernet (GbE), 400 GbE, etc.

The use of super-channels is also transparent and seamless from the end user's (client services) perspective, as enterprises can hand over to service providers a mix of low (e.g. 10 GbE), medium (40 GbE and 100 GbE) and high (100 GbE and 400 GbE) bandwidth services without any change in their network infrastructure.

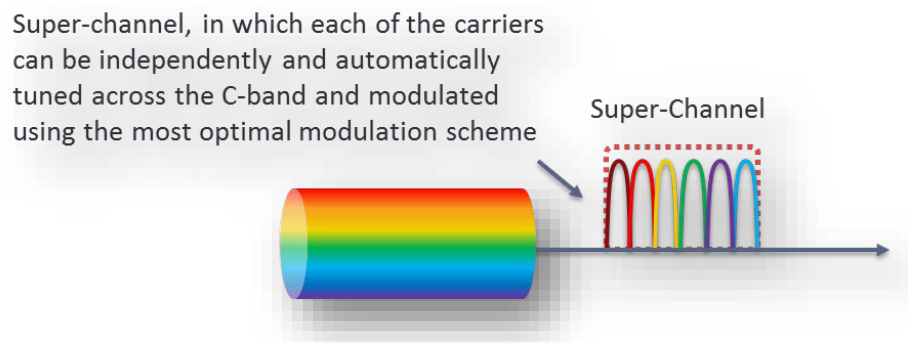


Figure 1 - High Capacity Through Super-channels

- Data protection through in-flight encryption:** An enterprise's success is heavily dependent on its ability to protect its own and its customers' data. Data breaches can trigger irreparable damage to the company's reputation and its ability to conduct business in the future, even driving enterprises to bankruptcy. Cable operators can protect enterprise data carried over the network using in-flight encryption as depicted in Figure 2, without the need for external boxes or complex setups. Different and flexible encryption schemes are available, such as Layer 1 encryption and Layer 2 Media Access Control Security (MACSec) encryption.




Figure 2 - Data Protection Through In-flight Encryption

- Purpose-built cloud direct connection platforms:** The enterprise shift to cloud is well underway and accelerating, driving the need for increased bandwidth and performance to connect to data centers and service providers. As more applications and content move to the cloud, including not only web-based customer-facing applications, but also mission-critical business applications, such as customer relationship management (CRM), enterprise resource planning and human resources, enterprises are relying more heavily on the networks that connect their sites to cloud service providers at carrier-neutral facilities across the world (Figure 3). A new breed of optical platforms provides the needed capacity for direct connection to cloud services with high security, low latency and simple scalability. These compact 1 or 2 rack unit (RU) platforms also provide open interfaces such as representational state transfer (REST) application programming interfaces (APIs) that allow seamless integration of these cloud services into the enterprise's existing IT environment and processes and the development of custom applications for traffic setup and monitoring, network optimization, etc.




Figure 3 - Cloud Direct Connection

- Dynamic bandwidth allocation:** New technology innovation in software and hardware allows bandwidth allocation to become dynamic and on-demand. Dynamic bandwidth allocation provides enterprise customers with on-demand provisioning of digital Optical Transport Network (OTN) and Metro Ethernet Forum (MEF)-compliant Ethernet services. Dynamic bandwidth can be deployed for a wide variety of use cases, including customer self-provisioned connectivity services and advanced policy-based services. Through the use of open APIs or the cable operator's graphical user interface, enterprise customers can provision services tailored around their needs, such as an ultra-low latency service that cannot exceed a certain latency threshold for delay-sensitive applications, or time-sensitive or time-of-day-based services for more efficient utilization of network assets.
- Optical networks virtualization:** New advancements in network abstraction capabilities can create virtualized networks at the packet, digital and optical layers and across metro and core domains. Parts of the cable operator's optical infrastructure can be logically partitioned based on each enterprise customer's needs, as depicted in Figure 4. Enterprise customers can have dedicated logical partitions of the network with complete visibility and control of their logical network and isolation from other enterprise customers, allowing them to customize their connectivity and services around their own applications. Similarly, such virtualization capability allows cable operators to maximize their return on assets and broaden their addressable markets without the additional capital often required to build private physical networks.




Figure 4 - Optical Network Virtualization

- Increased network uptime through intelligent software and optical hardware design:** Network downtime can translate into millions of dollars lost for enterprise customers, as well as serious damage to the cable operator's reputation and brand image. A network failure can have a disastrous impact on cloud applications such as data mirroring and backup. Cable operators can use 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. Network changes may include anything from multiple simultaneous failures to an increase in latency across any one of the network's critical spans. A control plane increases network availability and protects against failures, such as fiber cuts or hardware failures that would otherwise impact connectivity between an enterprise and its data centers. It can make enterprise networks autonomous and self-healing, even in the event of multiple fiber cuts or hardware failures. When coupled with the network virtualization capabilities described above, the control plane can partition an enterprise network by allowing specific links, wavelengths, subwavelengths or even nodes to be dedicated to a specified use, with preset thresholds for latency, bandwidth and resiliency. Some virtual enterprise networks are configured for high capacity, high resiliency and low latency for mission-critical applications (e.g. data mirroring). Other virtual networks are created with less stringent requirements for user access or other applications. Intelligent high-capacity optical networks make the infrastructure highly flexible to accommodate varying customer requirements.

4. Key Enterprise Applications

Most enterprises rely on a few key networking applications that are vital to their existence, all with varying requirements. Optical transport networks are built to accommodate these varying requirements

while maximizing performance and cost-efficiency. Some of these enterprise networking applications are listed below:

- **Business continuity/disaster recovery:** BC/DR describes a set of applications built to minimize the impact of downtime on an enterprise's operations in the aftermath of an emergency (natural disasters, terrorist attacks, major disruptions to the company's network, etc.). These applications consist of backup plans to transfer data and control access and offload other activities to one or more enterprise sites, including alternate data centers. Recovery time varies based on numerous factors, ranging from the few seconds needed to automatically transfer control and reroute traffic from primary to secondary sites once a failure or disruption is detected, to several minutes (Figure 5). Intelligent high-capacity optical networks play a vital role in BC/DR applications, providing the required bandwidth and alternative routes that meet stringent latency and capacity requirements in a very short period of time to minimize the impact on business operations. The successful deployment and operations of BC/DC applications are directly related to the performance of the optical network they rely on.




Figure 5 - Business Continuity/Disaster Recovery

- **Offsite/off-hours data backup:** This is a recurring operational procedure aimed to replicate or back up enterprise-critical data every day after hours to a remote data center (Figure 6). Typically, a large amount of data (terabytes) is automatically duplicated at another remote site/data center overnight during a backup window that spans anywhere from 30 minutes to several hours. There are several methods of data backup designed to save and keep an accurate history of data changes, such as full, incremental, differential, hybrid techniques and many others. The high capacity, low latency and task automation enabled by intelligent high-capacity optical networks make a significant positive impact on the time of execution, reducing the backup window from several hours to minutes, and enabling the overall success of this application.




Figure 6 - Offsite/Off-hours Data Backup

- Data mirroring:** Heavily used by enterprises that process millions of transactions daily (e.g. retail, finance, airlines, etc.), this application is used to instantaneously duplicate the same set of processes and transactions from a primary or master site (mainframe/data center) to a secondary volume or mirroring site (Figure 7). Keeping latency below a certain threshold is often a key factor in properly deploying this application. For example, a data mirroring application in a metro network is typically designed with round trip times (RTT) that range between 2 milliseconds (ms) and 20 ms. When deploying this application in a wide area network, the RTT can be as long as 200 ms⁵, hence a high-performance optical transport network with low latency and low jitter (a variation of latency) is a key building block of deploying any data mirroring application.



Figure 7 - Data Mirroring

5. Best Practices and New Operational Models

The landscape of the telecommunication industry is changing with new technologies, new end-customer requirements and new players. The following paragraphs highlight some of the best practices and emerging operational models that can help cable operators enhance customer experience, reduce operating costs, and increase their competitive edge.

- **Software defined capacity – on-demand service turn-up:** One of the main reasons behind customer dissatisfaction is the long time it takes to turn up a service. Slow turnarounds can easily escalate as a major source of conflict between the customer and network operator, and quite often it becomes the motivator for a customer to quit and go to a competitor that offers a faster response. Deploying a high-capacity network with super-channels, as described earlier, enables cable operators to turn up services faster, and often without truck rolls, hence reducing turn-up time and accelerating time to revenue. Pre-deploying high capacity does not require massive CapEx, as next-generation super-channel-based optical transport networks allow cable operators to add, activate and pay for additional bandwidth on existing network hardware in real time by making few clicks in a software application. The SDC model speeds up service activation from weeks to minutes and avoids the need for significant capital investment (e.g. pre-deployed chassis or idle line cards) to meet future growth.
- **Next-generation packet-optical platforms for lower CapEx/OpEx:** Next-generation transport platforms offer a compelling alternative to router-based solutions as they consume significantly less power and provide capacity and cost advantages for Carrier Ethernet services. Features aimed to simplify network implementation, design and operation such as plug-and-play setup and installation, zero-touch provisioning (ZTP), network auto-discovery, intuitive graphical user interfaces and easy, simultaneous network upgrades also contribute to lower OpEx.
- **Differentiated service offerings:** Enterprise customers have various requirements based on the industry they belong to and the applications they run. Cable operators can exploit the latest innovations in optical networking to offer differentiated services based on performance (bitrate), latency, security (e.g. encrypted links for financial institutions or government agencies) and survivability (e.g. protection against numerous simultaneous network failures due to a natural disaster). A broad service offering further widens the addressable market, increases competitive edge and augments revenue streams.
- **Enhanced network security:** Cyber-attacks and data breaches are frequent occurrences today. According to the recent Verizon Data Breach Investigation Report⁶, there were 3,141 confirmed data breaches in 2016. Cable operators can exploit the latest innovations in software and hardware to protect their networks and the enterprise customers they serve from intruders and hacking tools with features like in-flight Layer 1 and Layer 2 encryption, stringent access procedures and centralized authentication and authorization, to name just a few examples.

- **Task automation to streamline operations and eliminate human error:**
Enterprise migration to the cloud is dictating a whole new level of network agility, making traffic profiles and trends difficult to predict. Moreover, human error is often behind major network outages, hence the need to automate recurring tasks for better efficiency and reliability. Cable operators can use the latest developments in programmability, software tools and open interfaces such as REST APIs, NETCONF/YANG and others to simplify network management and automate recurring tasks. These intelligent software capabilities play a vital role in streamlining operations and evolving cable operators' and their enterprise customers' networks to the cloud by implementing task automation, proactive network monitoring, dynamic bandwidth allocation and much more.

Conclusion

Enterprises are undergoing a major shift in how they conduct business. Cable operators can benefit from enterprise evolution to the cloud by broadening their service portfolios and increasing their addressable markets while undertaking a continuous reduction in costs. The latest innovations in optical technologies provide the scalability, flexibility and programmability for cable operators to build intelligent high-capacity optical networks. These in turn help them to better serve enterprise customers to protect critical information, and provide the scalability required to meet the surging demand for bandwidth driven by cloud while reducing recurring costs.

Abbreviations

API	application programming interface
BC/DR	business continuity/disaster recovery
CapEx	capital expenditures
CRM	customer relationship management
DWDM	dense wavelength-division multiplexing
GbE	Gigabit Ethernet
IT	information technology
LAN	local area network
MACSec	media access control security
MEF	Metro Ethernet Forum
ms	millisecond
NETCONF	network configuration
OpEx	operational expenditures
OTN	optical transport network
REST	representational state transfer
RTT	round trip times
RU	rack unit
SDC	software-defined capacity

WAN	wide area network
YANG	yet another next gen
ZTP	zero-touch provisioning

Bibliography & References

- [1] http://www.irms360.com/blog_post/state_cloud_2015_supply_chain_adopters_reaping_roi_rewards
- [2] <http://www.mckinsey.com/industries/high-tech/our-insights/the-internet-of-things-sizing-up-the-opportunity>
- [3] <http://www.datacenterdynamics.com/security-risk/infographic-the-cost-of-cyber-attacks-in-the-us/96087.article>
- [4] http://www.information-management.com/infodirect/2009_133/downtime_cost-10015855-1.html
- [5] <https://technet.microsoft.com/en-us/library/cc917681.aspx#EGAA>
- [6] <http://www.verizonenterprise.com/verizon-insights-lab/dbir/2016/>

How to Succeed With SD-WAN Using Virtualized Service Assurance

An Operational Practice prepared for SCTE•ISBE by

Etienne Martel

Solution Manager

Accedian

2351 Blvd Alfred-Nobel, Suite N-410

Saint-Laurent (Montreal), Quebec, H4S 2A9, Canada

514-331-6181

emartel@accedian.com

Gregory Spear

Solution Manager

Accedian

2351 Blvd Alfred-Nobel, Suite N-410

Saint-Laurent (Montreal), Quebec, H4S 2A9, Canada

514-331-6181

gspear@accedian.com

Introduction

1. Executive Summary

To succeed in the enterprise market, cable multiple system operators (MSOs) must serve nationwide or global corporations that expect their providers to reach all their sites, so they do not have to assemble their own network. This means crossing multiple service areas, including outside the MSOs' footprint.

Software-defined wide area networks (SD-WAN) can help MSOs reach on and off-net sites, over any access media, uniformly—allowing them to leverage the scale and reach of their extensive DOCSIS and Carrier Ethernet services, augmented where required with third party access.

The pitfall: SD-WAN appliances do not offer standards-based test, turn up and monitoring functions required to offer service level agreement (SLA)-grade services. SD-WAN solutions use proprietary monitoring and reporting methods, which do not interoperate with existing network equipment. Because SD-WAN may only be required in certain customer locations, any implementation has to interact seamlessly with traditional service delivery methods.

This is not optional. All MSOs in North America offer SLA-backed services over fiber, and the majority over DOCSIS too, according to a 2017 Heavy Reading study¹; best-effort only services will not satisfy the enterprise market requirements for uniform services and stringent SLAs.

Virtualized test probes and test reflectors cost-efficiently replicate network interface device (NID) functionality, bringing standards based turn-up testing, monitoring and operations & maintenance (OAM) functions to SD-WAN endpoints. Virtualized instrumentation uplifts SD-WAN with carrier-grade functionality, making it interoperate with existing network infrastructure, operations procedures, and support systems.

In today's SD-WAN market, the two main deployment architectures are centralized or distributed. They will both benefit from service assurance solutions that can be deployed as software and optionally enhanced with NFV-powered hardware modules.

When selecting a service assurance solution, it is important to choose an industry-proven method for extending standards-based test, measurement, and OAM to virtualized environments for SD-WAN and x86 infrastructure to ensure satisfactory coverage and unified visibility. The interoperability provided by a standards-based solution enables centralizing performance monitoring data into existing reporting and fault-management systems to achieve the operational success of the technology. It also opens the possibility to optimizing the SD-WAN performance by leveraging the highly granular and micro-second accurate end-to-end monitoring data.

¹ Breznick, A. *Heavy Reading*, January 2017. How cable can conquer the enterprise market. Retrieved from https://accedian.com/wp-content/uploads/2017/01/HR_Accedian_Cable_Enterprise_WP_1-24-17.pdf

Content

2. SD-WAN adoption — Enablers and Drivers

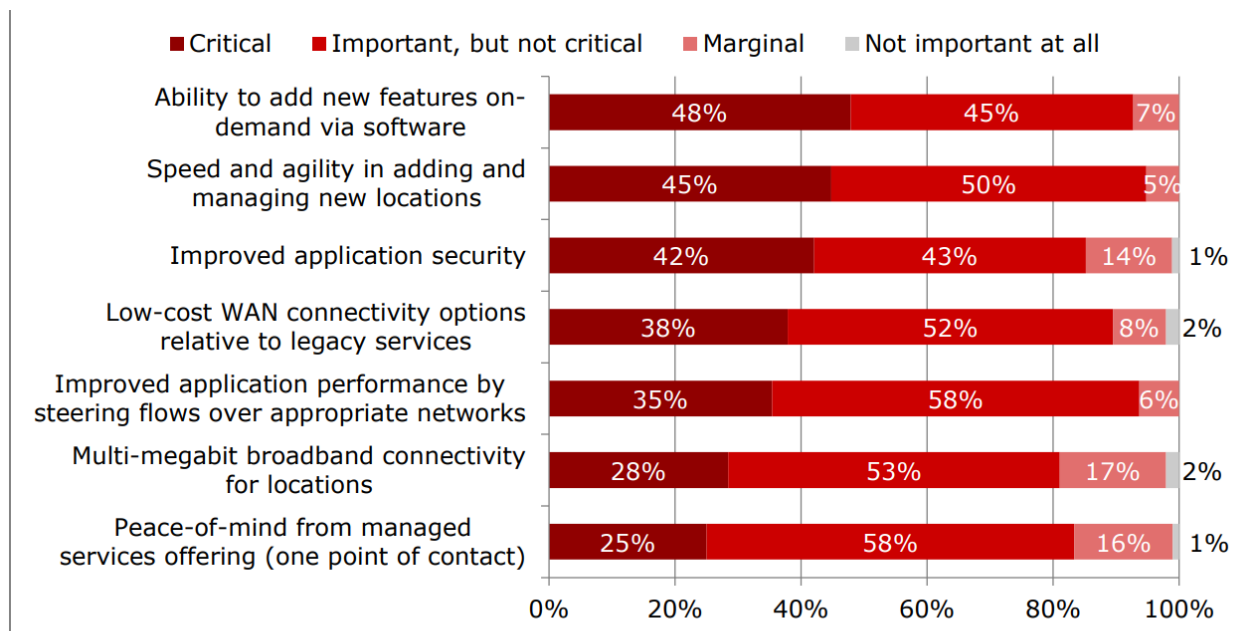
The ever-growing use of cloud-based applications by enterprises is making SD-WAN more relevant every day. Software-defined networking (SDN), initially reserved for data center applications—along with a number of other technology enablers—have set the table for SD-WAN to disrupt traditional WAN architectural models prevalent within most enterprises.

Technology enablers for enterprise SD-WAN include:

- Widespread availability of scalable and elastic cloud computing for the enterprise.
- Widespread availability of faster, more diverse, and more reliable wired and wireless broadband internet access technologies.
- SDN concepts: control plane and data plane separation.
- Application-aware routing with deep-packet inspection (DPI) instead of traditional IP routing.
- Availability of affordable x86 network appliances (commercial off-the-shelf hardware/COTS)

From hybrid models to full scale implementations, the advantages offered by SD-WAN cannot be ignored. SD-WAN offers increased network agility, better overall performance, lower cost per megabit, and a new WAN architecture that complies with the software as-a-service (SaaS) and cloud computing business model.

All enterprises are now considering SD-WAN solutions for their next WAN refresh; none plan to fully retain the traditional single WAN, hub and spoke model with centralized internet access at the hub. Enterprises are compelled to consider SD-WAN, if only for the ability to have branch locations locally break-out to the internet in a secure and controlled way instead of backhauling all internet traffic across the WAN to the hub.



N=96

Source: Heavy Reading December 2016 Operator Views on Emerging SD-WANs Survey, Sponsored by ADVA

Figure 1 - Most Important Expected Benefits for Operators' Customer²

From the operators point-of-view, SD-WAN is appealing for similar reasons, primarily to offer a low-cost bandwidth enhancement to offered services—but also because it unlocks the ability to turn up new features on-demand, enhancing agility and speed for service delivery and the initial service turn-up.

As shown in Figure 1, operators are deploying SD-WAN in their network to gain agility and flexibility first and foremost. The software automation at the heart of the SD-WAN solution will allow for the creation of fully dynamic networks and give end-users a control into the nature and level of services they require on an ongoing basis. From the MSO point-of-view, having the ability to both deploy and maintain features and services via software deployment is crucial. Running software on COTS servers dramatically lowers both risks and costs when compared with the traditional dedicated hardware appliance solutions that required extensive trials to approve and the trained personnel, space, power and cooling to run.

As it stands, SD-WAN deployments are still in the early adoption phase and do not offer the same service-level expectations as traditional business service WAN offerings. Now, as widespread adoption continues, operators find that SD-WAN managed services must deliver the same quality levels as traditional WAN offerings. Operators therefore need tools that offer the visibility and reporting capabilities to manage network performance and SLAs.

² Sterling, P. Heavy Reading. February 2017. Operator Success in the New Age of the Software-Defined WAN, Retrieved from <https://resources.ext.nokia.com/asset/201132>

3. Business Services Over SD-WAN Lifecycle Overview

Specific operational practices pertain to each phase of the business services over SD-WAN service lifecycle as illustrated below in Figure 2:

1. Provisioning and Turn Up: Deployment and service activation testing (SAT)
2. Performance Management: Performance monitoring and SLA reporting; collecting and presenting key performance metrics
3. Fault Management: techniques to identify, isolate, and troubleshoot service issues

These three phases are consistent with the Metro Ethernet Forum (MEF) definition of the Carrier Ethernet service lifecycle³, which serves as an established model for commercial connectivity.




Figure 2 - Metro Ethernet Forum Service Lifecycle⁴

This paper is structured to address the operational practices associated with each stage of the service lifecycle, as it applies to business services over SD-WAN.

³ MEF Forum. April 2012. Introducing the Specifications of the MEF. MEF 38: Service OAM Fault Management YANG Modules Technical Specification. Retrieved from: <http://slideplayer.com/slide/5687304/>

⁴ MEF Forum. March 2016. Service Operations Specification MEF 55: Lifecycle Service Orchestration (LSO): Reference Architecture and Framework. Retrieved from: http://dev.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf

4. Deployment Options Over Common SD-WAN Architectures and Hybrid Services

Any performance monitoring solution requires four key elements: 1) test session control; 2) a test packet generator; 3) a test packet reflector or receiver; and 4) precision timestamping. In a traditional WAN network, these elements were often customer premises equipment (CPE)-based network interface devices (NIDs) and centralized test suites, all of which required time to install and configure—something that virtualized network services have eliminated.

The operational practices can be implemented using a network-embedded architecture that employs small footprint, programmable service assurance hardware modules (vCPE modules) augmented by virtualized service assurance functions hosted on a centralized, virtualized performance assurance controller (vPAC). A lightweight, stand-alone orchestratable software agent is another viable way to instrument the network; this architecture can offer a complete software-only solution. However, it cannot rival the precision and the feature-set offered by vCPE modules. Section 4 introduces these architectures, as well as operational considerations that facilitate integration of these approaches with existing operational support systems (OSS), network management systems (NMS), and virtual network function (VNF) orchestrators.

4.1. Architecture Overview: Using Virtualized Performance Assurance Controller VNFs and Lightweight Stand-Alone Orchestratable Software Agent VNFs.

For deployments where service assurance using standard-based protocols is needed, but the added-benefits offered by the NFV-enabled modules are not required, stand-alone orchestratable software agents can be used to offer the reflection capabilities needed to complement a centralized performance monitoring approach using vPAC VNFs as probe generators.

The main benefit of this software-only architecture is the deployment speed and agility offered by being able to remotely and centrally deploy, configure, and run everything needed to instrument an existing network, on-demand and with minimal expense. Standards-based monitoring methods integrate the network itself into a ubiquitous instrumentation layer. With this visibility centralized in data centers shared with SDN control and big data analytics, providers have an integrated foundation to deliver a new level of customer experience.

The vPAC assumes all session setup, control, and sequencing functions, as well as results analysis and reporting to file servers. As a virtual network function (VNF), vPAC instances can be deployed and orchestrated seamlessly with the network service descriptors, allowing fully-automated setup and assurance of virtual service chains.

The lightweight software agent VNF offers reflection capabilities required to instrument the network with any orchestrator and can easily run un-privileged on any Linux based operating system.

The lightweight software agent VNF also enables bi-directional measurements, unrivaled metrics set, measurement granularity, and third party interoperability—features that are unavailable when using built-in standard open-source tools (such as ICMP ping) or even proprietary measurement methods offered by SD-WAN vendors.

4.2. Architecture Overview: Using Controller VNFs and NFV-Powered Hardware Modules.

In the context of this document, the term *vCPE* will refer to the strategy of virtualizing as many customer-located networking functions as possible, while retaining the minimum hardware necessary for service delivery, consistent with performance, reliability, and quality of experience (QoE) expectations. An example of a vCPE strategy—where onsite hardware appliances performing firewall, PBX, and routing functions have been virtualized—is illustrated in Figure 3, below. Virtualization is accomplished by transferring local networking functionality to software-based VNFs, which can be hosted on low-cost COTS servers or cloud infrastructure.

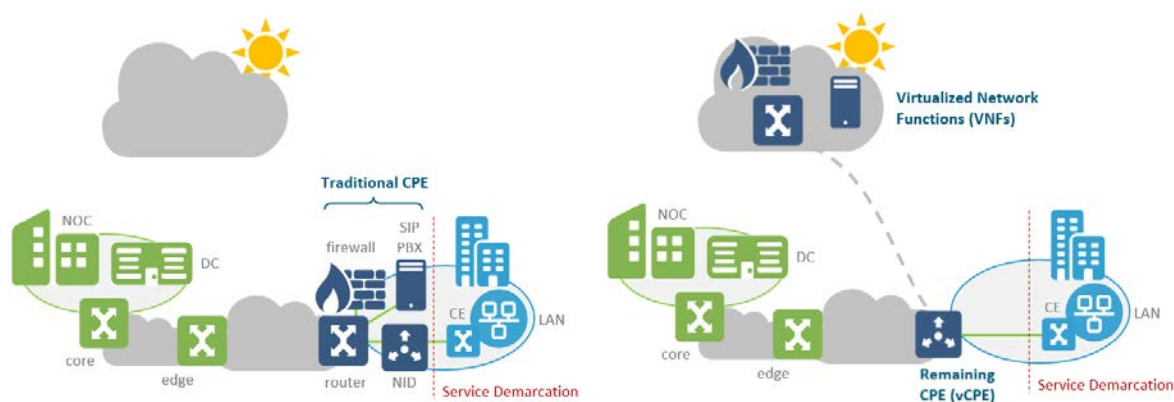


Figure 3 - vCPE: Traditional vs. Virtualized Customer Premises Equipment Example

In the context of SD-WAN, this approach can be used to introduce customer premises-located performance monitoring, turn-up test, service OAM (SOAM) and troubleshooting functionality, which—in the case of fiber business services—is normally provided using a NID. Reducing hardware appliances required at the branch site is a key benefit of SD-WAN; installing a standard NID along with the SD-WAN appliance is not normally a feasible CPE option.

NFV-powered hardware modules can offer the same level of performance monitoring precision, as well as loopback and full line-rate turn-up test capabilities at a fraction of the cost of a NID, making this approach an economically viable fit when deploying SLA-grade business services over SD-WAN. The solution delivers complete quality of service (QoS) and QoE insight without compromise. In addition to supplementing the SD-WAN appliance (or COTS server) with service assurance features, this approach has a number of other benefits:

1. Truck-rolls are reduced over the service lifecycle when compared to handheld test sets, as a single vCPE module can remotely perform turn-up testing, continuous monitoring, and on-demand troubleshooting.
2. Compatibility with existing hand-held Ethernet test sets and third-party centralized monitoring probes allows straightforward integration into existing operational practices and infrastructure.
3. By employing NFV, new functionality can be added to the vCPE module remotely, without impacting the service. This allows MSOs to introduce new, performance-assured commercial services without requiring new equipment on-site.

An example of how NID functionality can be virtualized using NFV is shown in Figure 4 below. On the left, you can see a traditional CPE NID, it contains both a control plane and a data plane. On the right, we have disaggregated the functionality into a layer of software to deliver the control plane in the form of a service assurance VNF and a layer of hardware in the form of vCPE Modules which contain just enough hardware to deliver the required data plane features at the site while leveraging the VNFs for any compute intensive job.

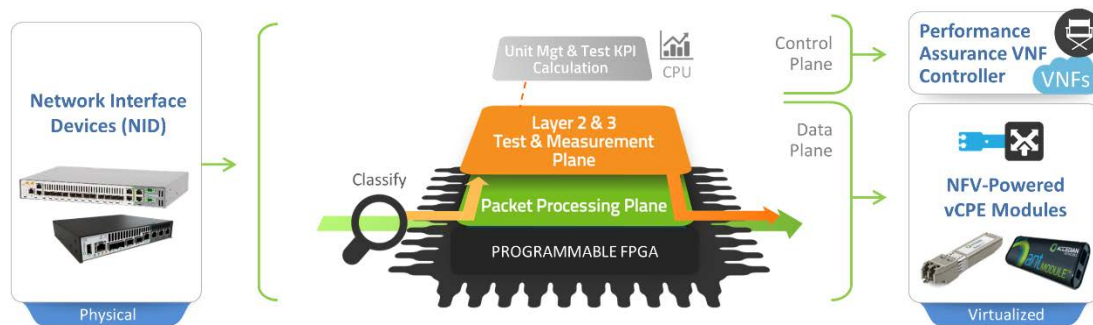


Figure 4 - Virtualization of NID Architecture Using NFV

The connection between the service assurance VNF controller and each module needs to be reliable, secure, and lossless (e.g. transmission connection protocol (TCP) based) to ensure the vCPE module can assume the same level of functionality as a traditional NID. As shown in Figure 5 below, this management ‘tunnel’ is critical to support service assurance VNFs, as raw data is returned to the controller for test results calculation, performance monitoring, and fault reporting, in addition to performance monitoring session control, module management, synchronization information, etc. In an NFV-based vCPE architecture, the ‘lossless’ control sessions allow each remote module to virtually become a remote ‘port’ of the controller, which is analogous to a virtualized NID that can support many remote endpoints.

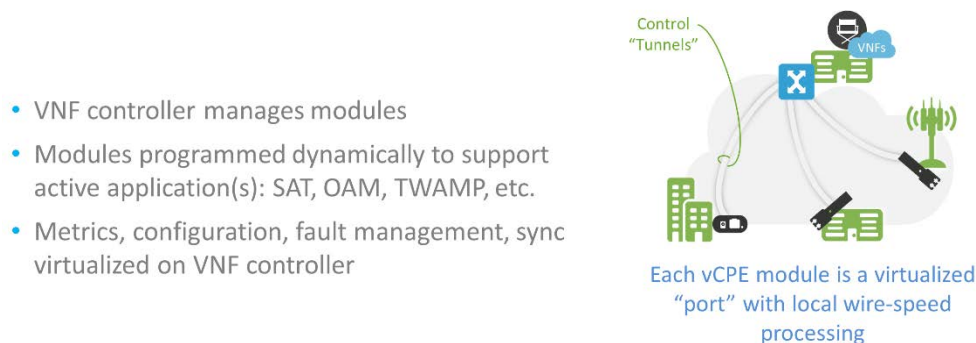


Figure 5 - NFV-based vCPE Control Tunnel & Function Locations

SD-WAN architectures virtualize some or all customer premises functions with a simple COTS server at the customer site. As part of their standard feature-set, SD-WAN solutions implement path monitoring and measurement. However, these measurements are typically lacking for managed business services over SD-WAN deployments because those service assurance functions implemented purely in software:

1. lack sufficient time stamping precision and packet transmission scheduling control to meet the requirements of:
 - a. Full line-rate test traffic generation and loopback for SAT and troubleshooting.
 - b. Precise traffic generation sequencing required by common turn-up test standards (where inter-packet delay needs to be controlled for burst testing, for example).
 - c. Microsecond-level latency measurement precision required to monitor and report on commercial services SLAs.
2. are subject to the resource-sharing of the x86 system. This causes additional uncertainty in the results by bundling the performance of the x86 system with the performance of the network itself.

In addition, SD-WAN solutions use proprietary monitoring and reporting methods which do not interoperate with existing network equipment (or other SD-WAN vendors). Because SD-WAN may only be required in certain locations, any service assurance implementation has to interact seamlessly with the traditional service delivery methods.

Relying on built-in SD-WAN monitoring also has the effect of creating a potential blind spot. This is especially true when considering service activation testing (SAT) such as RFC2544⁵ or ITU-T Y.1564⁶ which have no support from the SD-WAN vendors. Moreover, the SD-WAN built-in performance monitoring functions can also only provide a top-down view of performance—the over-the-top (OTT) path. This view presents no insight into why a specific path is operating badly, just that it is not performing. Complementing this top-down view with a bottom-up perspective provided by hop-by-hop or layer 2-3 path monitoring can add the missing pieces to more efficiently run an assured SD-WAN services, enabling detailed troubleshooting and measurable quality improvements.

Running a unified service assurance solution across both the incumbent part of the network and the SD-WAN part of the network also has the benefit of offering a unified level of precision and reporting intervals. As such, pin-pointing events and segmenting the network will ease troubleshooting and accelerate mean time to resolution (MTTR) when issues arise.

Aside from enabling these capabilities, vCPE modules also offer a number of other advantages over traditional test set and centralized probe solutions:

1. Modules can monitor and test between themselves: for site-to-site monitoring, end-to-end turn-up testing, and troubleshooting between customer service endpoints. Most probe-based solutions are limited to loopbacks or monitoring tests from a central location to a service endpoint. This ‘hub-and-spoke’ topology does not test the actual service path between customer locations, or between a customer and a remotely hosted data center, for example.

⁵ Bradner S., McQuaid J., March 1999, RFC2544. Benchmarking Methodology for Network Interconnect Devices. Retrieved from <https://www.ietf.org/rfc/rfc2544.txt>

⁶ ITU-T. February 2016. Y.1564: Ethernet Service Activation Test Methodology. Retrieved from: <https://www.itu.int/rec/T-REC-Y.1564/en>

2. Test sets require trained technician dispatch to each service endpoint requiring service activation testing or troubleshooting, which is much less responsive and much more costly than a remotely initiated test using vCPE modules that are initially installed during service provisioning.

4.3. Service Assurance Functions

NFV-based vCPE solutions must be capable of all service assurance functions required to support the business services lifecycle, as described in Section 3 and shown in the Figure 6 displayed below. These include, but are not limited to:

1. Standards-based SAT supporting commonly employed IEEE RFC-2544 and ITU-T Y.1564 turn-up testing approaches.
2. Ethernet connectivity fault management (CFM), as defined by IEEE 802.1ag⁷ to ensure service availability meets SLA definitions, and to measure continuity and latency using CCM and DMM/DMR messages, respectively.
3. Standards-based performance monitoring for Layer 2 (Ethernet) and Layer 3 (IP) services, typically implemented using ITU-T Y.1731⁸/ IEEE 802.3ah⁹ Ethernet SOAM and RFC-5357 Two-Way Active Measurement Protocol¹⁰ (TWAMP), respectively.
4. Bandwidth utilization monitoring, per port and per service flow (as defined by the MSO: VLAN, class of service/CoS, source or destination MAC or IP address, etc.) for usage-based billing, trending, and troubleshooting.

⁷ IEEE. December 2007. 802.1ag - Connectivity Fault Management. Retrieved from: <http://www.ieee802.org/1/pages/802.1ag.html>

⁸ ITU-T. August 2015. G.8013/Y.1731: OAM functions and mechanisms for Ethernet-based networks. Retrieved from: <https://www.itu.int/rec/T-REC-Y.1731>

⁹ IEEE. September 2015. IEEE Standard for Ethernet. Retrieved from: <http://standards.ieee.org/about/get/802/802.3.html>

¹⁰ Yum, K. IETF. October 2008. RFC-5357: A Two-Way Active Measurement Protocol (TWAMP) Retrieved from: <https://tools.ietf.org/html/rfc5357>

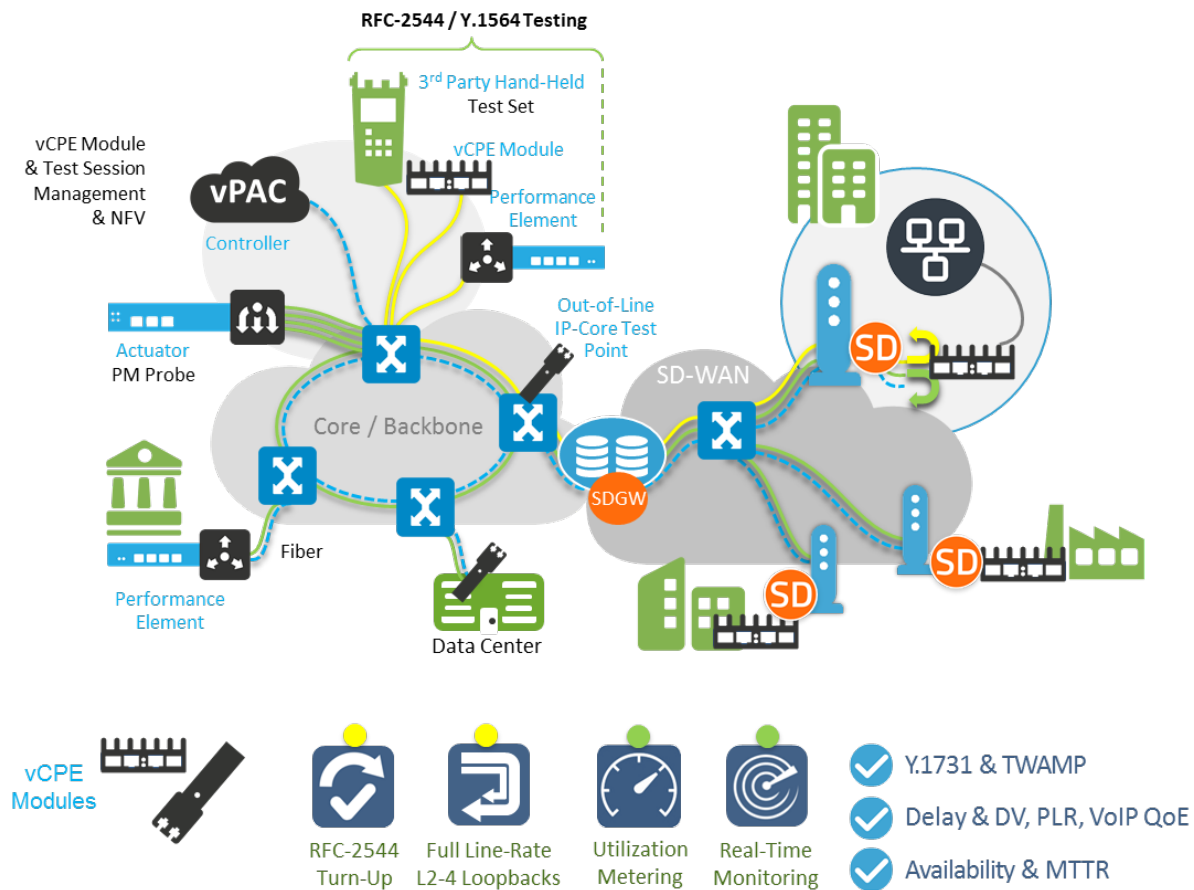


Figure 6 - Typical SD-WAN Service Assurance Functionality & Implementation

The full gamut of the Service Assurance solution can be seen in Figure 6. The provisioning and turn-up stage of the service lifecycle is delivered with the ability to generate and reflect Turn-up tests (yellow line) using traditional NIDs, Performance Elements, or hand-help test sets up to full line rate at any Ethernet supported packet size. The Performance Management stage of the service lifecycle is delivered through performance and SLA monitoring (green and blue lines) with micro-second accurate continuity and latency measurements and one second granularity bandwidth monitoring

4.4. Operations Integration Considerations

Implementation of an NFV-based solution should interwork with existing OSS to permit integration with existing management practices and procedures, and to make deployment of vCPE modules—as well as the monitoring and maintenance of the services they support—as operationally efficient as possible. Main areas to consider include:

- Deployment and management of the solution itself, to facilitate and automate element management of remote vCPE modules and SD-WAN service provisioning.
- Integration with SLA reporting platforms and fault management systems to harmonize monitoring, and reporting and within existing tools.

As introduction and general guidelines, the following operational aspects should be considered during solution selection and deployment.

4.4.1. Low-Touch vCPE Module Provisioning

Ideally, vCPE modules should be ready to install in ‘factory default’ configuration, without requiring pre-staging by the MSO. To make that possible, the modules must be discoverable by an inventory system that can attribute the module to a particular customer site. This may be accomplished by relating the module to the MAC or IP address of the customer’s SD-WAN appliance, for example.

To come under management control without requiring pre-staging or on-site configuration by a trained technician, the units require a method to ‘discover’ the management environment, have their management IP address defined, have the desired configuration provisioned on the unit, have the module registered in the inventory and potentially have the service turn-up testing start automatically followed by service performance monitoring.

One commonly employed method to bring devices under management involves using Dynamic Host Configuration Protocol (DHCP) with options 60 & 43 (refer to Figure 7 below):

1. When a vCPE module is connected to the network, it announces itself using a DHCP request with option 60, which communicates the module’s identifier (device type) to the DHCP server.
2. When properly configured, the server assigns a dynamic IP address to the unit, and responds with option 43, providing the module with address of its “inventory node,” responsible for managing the module.
3. Once under management control, a static IP can be assigned (if desired). A Fully Qualified Domain Name (FQDN) can also be assigned to the module. The FQDN must remain in sync with any link-state change (per RFC-2131), typically realized using automated DNS queries by the module inventory node.
4. Once the module is under management control, automation may be used to trigger an immediate or scheduled turn-up test to validate the service, then provision customer/SLA-specific monitoring sessions, etc. to allow customer-level self-install.

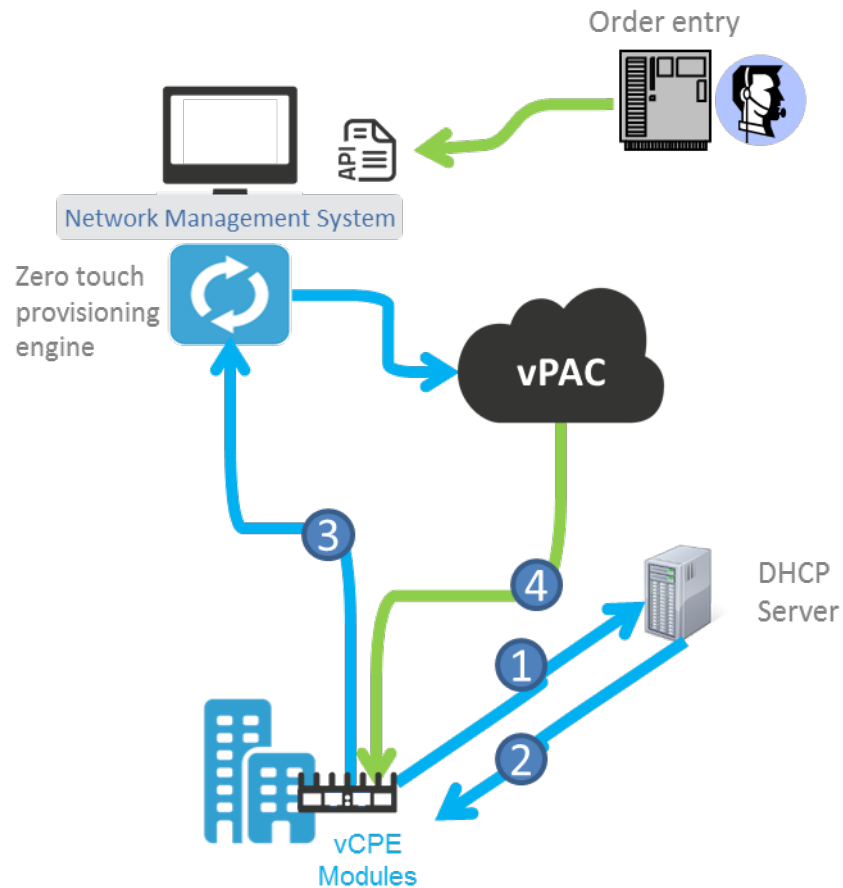


Figure 7 - Low-Touch vCPE Module Provisioning

4.4.2. Stand-alone Software Agents Provisioning

Software agent deployment is very simple and essentially consists of transferring a single executable binary file to the target system.

Software agents should be service-chained in such a way as to be accessible from the WAN link(s). Their configuration is typically very simple and straightforward requiring only a limited number of options to be passed as arguments on the command-line during the agent instantiation. The agents can be configured to reflect performance monitoring probes only on specific system interfaces if desired.

The most basic type of software agent requires no separate management communications channel as it doesn't communicate management information, but merely act as responder to performance monitoring probes sent from a vPAC-type node.

4.5. vCPE Module Management Communication Options

The vCPE module should be capable of adapting to an MSO's particular management and device addressing methodology. This requires the unit to distinguish customer/test traffic from management communication. A variety of methods are commonly used, all implying that the vCPE module must offer support for each of these schemes:

- Layer 2 addressing: separate management and customer MAC addresses. In this case, the vCPE module must support two MAC addresses, one for management traffic, and another for customer, test, CFM, SOAM, and active performance monitoring traffic.
- Layer 2 addressing: separate management and customer VLANs. A single MAC address is used in combination with Q-in-Q VLAN support (C/S tagging, IEEE 802.1ad).
- Layer 3 addressing: similar to the Layer 2 scheme described above; separate management and customer traffic IP addresses may need to be supported by the vCPE module, depending on the method used by the operator. VLAN support, including Q-in-Q (S-VLAN), may also be required to support this scenario.
- Layer 3, transparent IP addressing: an operator may elect not to assign new IP address(s) to the vCPE module, instead using the address of a device located 'behind' the module. This is practical when the operator has a known device at the customer premises (e.g. an SD-WAN appliance, a set-top box, Wi-Fi controller, security gateway, etc.) In this case the vCPE module detects management traffic using a combination of this other device's IP address in combination with other identifiers (such as a management VLAN tag).

5. Industry Proven Methods to Extend Standards-Based Test, Measurement and OAM to Virtualized Environments for SD-WAN and x86 Infrastructure

5.1. SD-WAN Traffic Routing and Performance Monitoring

vCPE Modules are capable of full line-rate test traffic generation, able to create and analyze up to four Layer 2 or Layer 3 unique flows, and run a fully-fledged RFC-2544 or 8-flow Y.1564 SAT suite toward other vCPE modules or third-party endpoints.

Each vPAC is capable of generating up to 4000 performance monitoring flows toward other vPACs, other vCPE modules, or third party-endpoints.

When paired together, the vCPE modules and vPAC will allow service providers to test the multiple service paths at turn-up and re-validate on-demand the capacity of a specific path or service during maintenance windows for troubleshooting.

To direct performance monitoring flows, the SD-WAN controller pushes policies throughout the network, stating that network traffic complying with defined profiles is kept on chosen paths. Then the vCPE modules or the vPAC generate performance monitoring flows that comply with the traffic profile, giving the operator the ability to proactively monitor all SD-WAN paths concurrently and quickly take action whenever a fault is detected.

5.2. SD-WAN Deployment Models and Service Assurance Instrumentation

From a network instrumentation point-of-view, providers can either use a blanket approach to simplify and unify their deployments or use a right-size approach to customize and adapt the instrumentation to the site and its significance in the overall network. The selected service assurance solution should have interoperable solutions that scale from supporting built-in integrated third-party reflectors to fully featured dedicated devices.

As shown in Figure 8, the ability to use a diversity of standards-based service assurance endpoints unlocks the ability to deploy the best tool for the job at each location. Some locations, like aggregations and cores sites, will require the enhanced features provided by a traditional Performance Element NID. Many Enterprise and SMB customers will benefit from the essential features and ease of use offered by vCPE modules, and for those locations where an x86 platform is readily available, a low-touch light-weight software agent can be deployed.

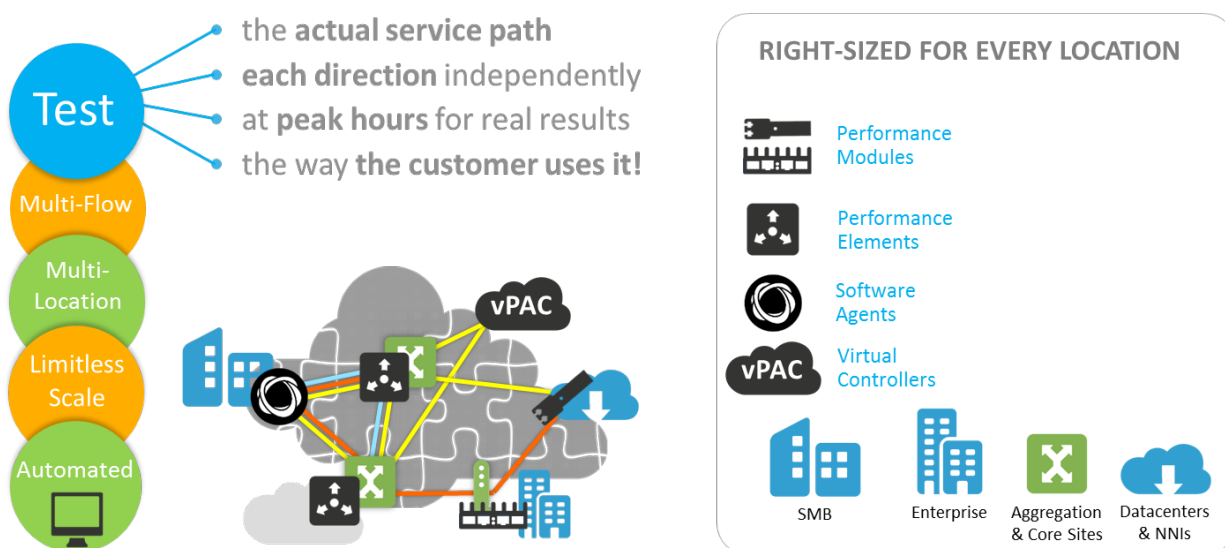


Figure 8 - Performance Monitoring Through Multiple Paths to a Diversity of Endpoints

5.2.1. Centralized Gateway SD-WAN Models

Many SD-WAN vendors offer an architecture based on centralized gateways to act as the virtual hub for any number of remote locations (spokes) as displayed in Figure 9. The connected sites (Branch, Head Office, HQ) need little hardware and a number of network transports (internet links or traditional WAN links) to establish the overlay network needed for the SD-WAN to operate illustrated using the gray lines to the SD-WAN Gateway. The overlay network is built by having each remote site establish encrypted tunnels to the SD-WAN gateway over each provisioned path.




Figure 9 - Centralized Gateway SD-WAN Model

5.2.1.1. With vCPE Modules

In such a model, it is recommended to co-locate a vPAC along with the SD-WAN gateway(s) as compute resources are typically plentiful in the cloud.

When deployed in this manner, the vPAC can be used in large scale hub-spoke and full-mesh topologies to perform active, micro-second accurate, standards-based performance monitoring towards thousands of endpoints continuously.

To ensure the most flexible and featureful performance monitoring solution, the vPAC is supplemented by a vCPE module at each connected site. The remote vCPE modules effectively become remote ports of the centralized vPAC and therefore this deployment model ensures that the full performance monitoring

feature-suite is available for the operator. This also guarantees that the solution will evolve along with the network as the functionality is delivered through NFV capabilities to each remote endpoint.

Further, as displayed in Figure 10, having vCPE modules at each site enables NFV performance monitoring (NFV-PM) and remote SAT generation. Gaining the ability to source the performance monitoring traffic flows from any location and targeting any location ensure that the results will match the user experience 1-to-1.




Figure 10 - Centralized Gateway SD-WAN Model —SAT and Performance Monitoring with vCPE Modules

5.2.1.1. With Software Agents

For locations where the added benefits of vCPE modules, NFV-PM, and remote SAT generation/reflection are not required, a lightweight software agent can be embedded directly into the SD-WAN software stack on-site. The agent can run on bare-metal, in unprivileged mode on a Linux OS, or inside a dedicated container within the NFV infrastructure (NFVi) supporting the SD-WAN deployment.

The benefits of this approach are two-fold: 1) the simplicity of deployment is second-to-none as only a simple software program need to be transferred and started; 2) the site running the software agent can now be the target of highly precise and granular probes, supporting 1-way metrics and return standard-based metrics to the vPAC that will process the results, analyze and monitor the network quality, and report the network state northbound to NMS, analytics, and big data systems.

When software agents are installed on x86 CPE, the overall performance strategy is slightly changed. As shown in Figure 11, the performance monitoring flows are no longer endpoint-to-endpoint (spoke to spoke), but instead generated at the central cloud location (hub) and reflected at the spoke. This falls in line with the visibility provided by the SD-WAN gateway where each path is measured independently. While the model covers the complete network, it does not completely reflect the end-user experience as accurately as the vCPE modules strategy explained previously can.

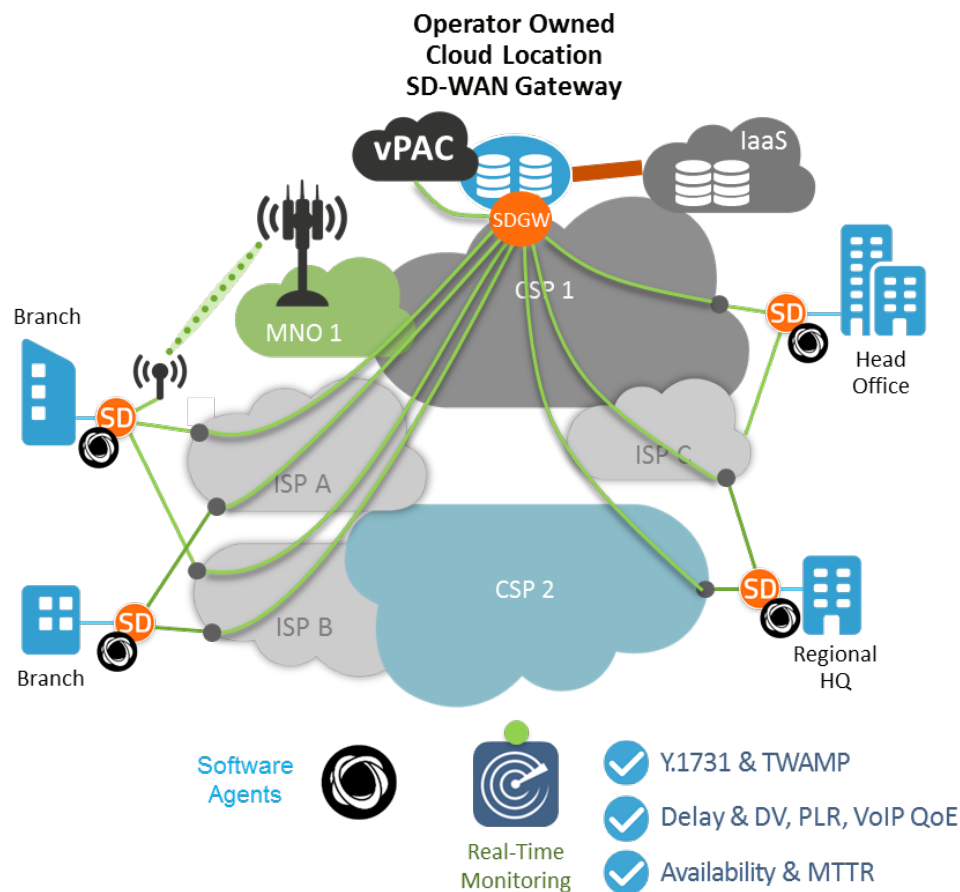


Figure 11 - Hub-to-Site Performance Monitoring from vPAC to Software Agents

5.2.2. Distributed SD-WAN Models

For SD-WAN architectures where each site can directly connect to each other site dynamically and on demand, the need for a centralized SD-WAN gateway is removed. This approach has slightly different pros and cons because the SD-WAN intelligence is distributed into the end-points and the provider can use separate controller (or director) software to centrally configure and control each of the endpoints. Removing the centralized gateway can improve overall resiliency and reduce the required cloud compute resources, but it also removes a centralized location where additional network services (such as the vPAC) could easily be hosted. As shown in Figure 12, in this SD-WAN model, each site is able to establish overlay tunnels over each potential path to any other site dynamically and as a response to underlying network conditions.




Figure 12 - Distributed SD-WAN Model

5.2.2.1. With vCPE Modules

When operating in a distributed SD-WAN architecture, the vPAC can be hosted in a head office or data center location as it requires limited compute resources. In the rare case where no compute resources are available inside the existing network, the vPAC can be hosted in a data center location and then connected to the SD-WAN as a synthetic site.

Distributed SD-WAN architectures are best served when using controller VNFs and NFV-powered hardware modules because this deployment model makes it possible to easily test and monitor each path directly and individually as shown in Figure 13. Having the ability to retrieve all this performance data enables the operations team to easily pinpoint issues and bottlenecks in situations where traffic from one location might hamper the performance when communicating to another site. Given the exponential number of potential paths and the dynamic nature of the traffic, being able to monitor the whole network with a single solution, and easily export the performance data to an open data repository, proves to be instrumental to supporting this new type of environment.

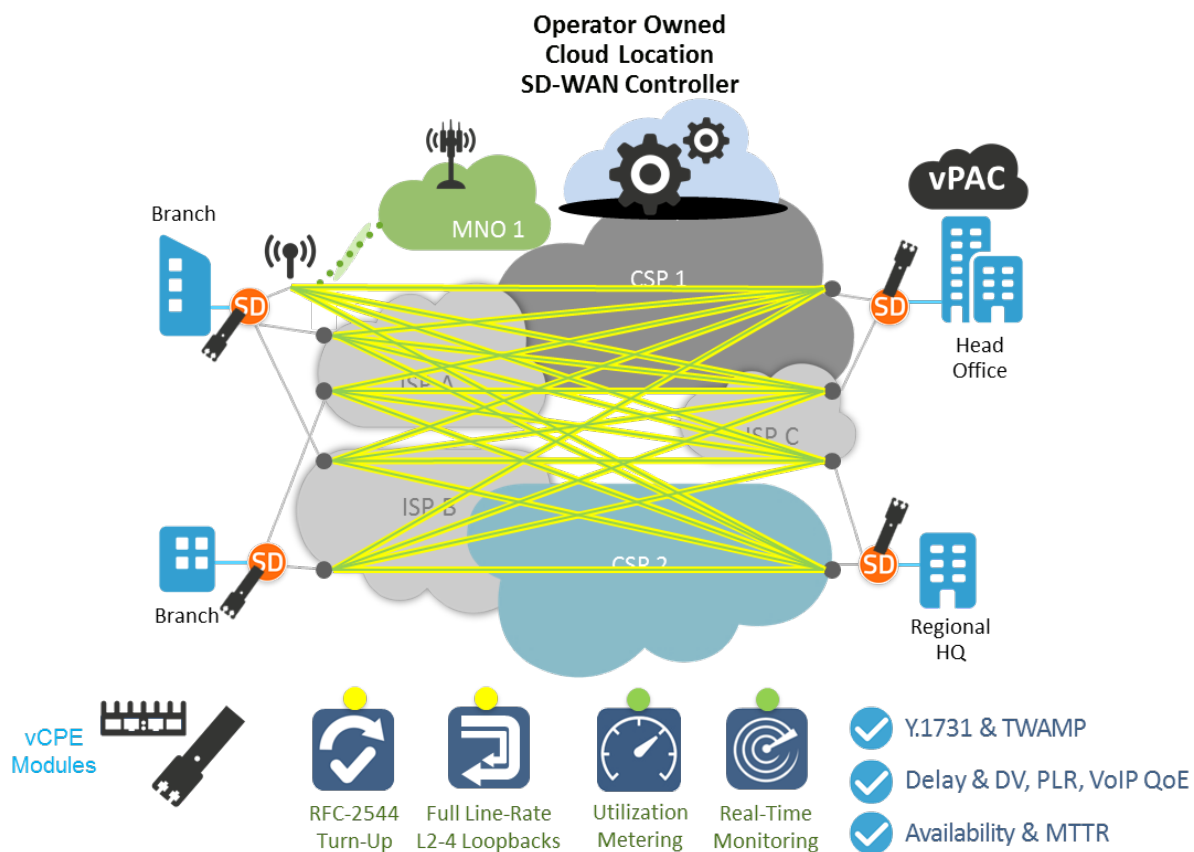


Figure 13 - Distributed SD-WAN Model — SAT and Performance Monitoring with vCPE Modules

5.2.2.2. With Software Agents

When operating in a distributed SD-WAN architecture and using software agents or third-party reflectors, the vPAC should be hosted in locations of greater interest as it will be used to generate performance monitoring traffic flows. When many sites have significant operational importance, each one of these sites should run the vPAC VNF in order to be used as a probe generator.

Because the vPAC is used as the flow generator and the software agent is a reflector, the result of combining a distributed SD-WAN architecture with a 100% software performance monitoring solution results in star-shaped measurements to represent a mesh-like network. Depending on the end-user's use-case and typical network usage, the limited monitoring could be less than ideal and therefore it is recommended to supplement this model with cost effective vCPE modules in key locations to offer the coverage needed for best results.

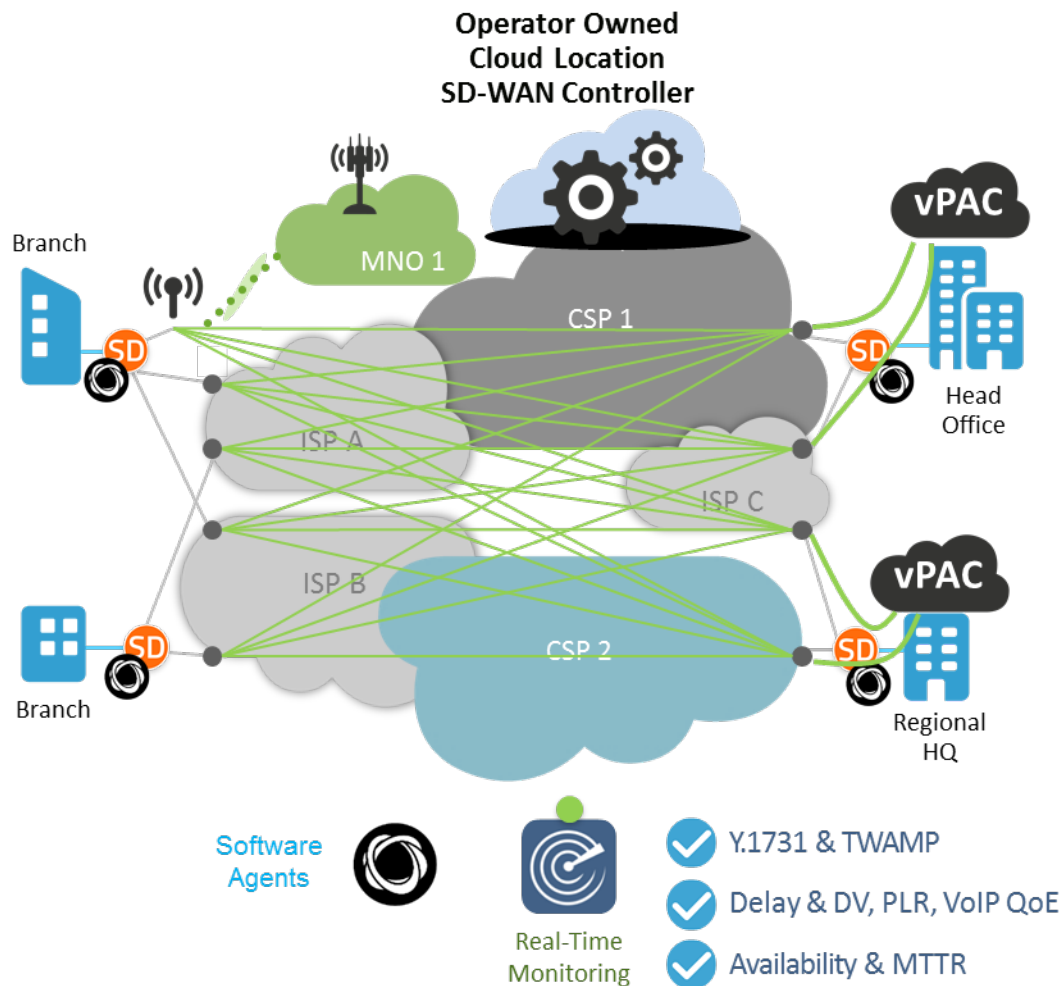
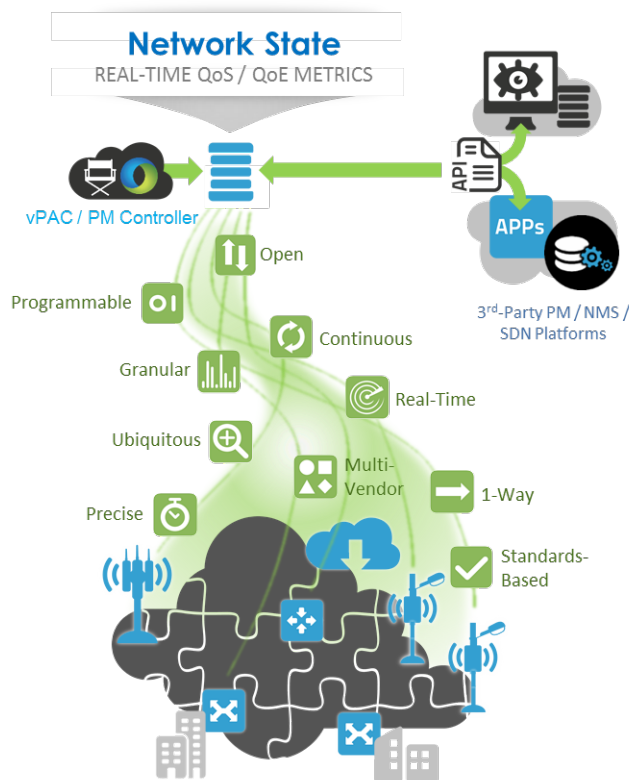


Figure 14 - vPAC to Software Agent Performance Monitoring

6. Methods to Centralize Performance Monitoring Data into Existing Reporting and Fault-Management Systems.

The virtualized NFV-based active and passive monitoring solutions typically provide both fault management (FM) and performance management (PM) metrics and data into an operator's existing FM, PM, and OSS systems. FM data is traditionally provided as threshold-based simple network management protocol (SNMP) traps or extensible markup language (XML) type events, whereas PM data often uses comma separated values (CSV) files or XML/JSON encoded data streams or file transfer.



The metrics data provided to the 3rd party system should fulfil as many as possible of the 10 metrics fundamentals:

1. Open – well documented
2. Continuous 24/7
3. 1-way – path separation
4. Granular – seconds not minutes
5. Programmable – configurable
6. Real-Time – immediately available
7. Ubiquitous – as many sources as possible
8. Multi-vendor – no lock-in
9. Standards based
10. Precise – microsecond where applicable

Figure 15 - Export and Centralize Network State to Northbound Systems

The metrics reported northbound should be tangible, near real-time, and granular. For a modern SD-WAN type deployment the historically de-facto standard of using SNMP-polling every 5 or 15 minutes is far from sufficient to capture the fast transients that occur in today's networks. Active (TWAMP / Y.1731 / UDP) type tests should be able to report at least every 30 seconds, possibly down to every 5 seconds, and use a sampling rate of at least 10 probe packets per second to properly capture any short-lived events. The highly granular data may be aggregated in retrospect, to reduce the amount of metrics stored for long-term reports. It is important, though, that such aggregation does not result in loss of the measured extremes. Use of percentiles (95th, 98th, 99th, etc.) helps preserve a view of the distribution of delay or jitter values, while metrics such as loss burstiness aid in determining length of service interruptions, down to millisecond level with a sufficiently high sample rate.

Ideally, the northbound metrics transport should be streaming based (XML / JSON) for larger installations due to the efficiency of such machine-to-machine (M2M) methods. For smaller, enterprise-type applications, a CSV-based file transfer may be sufficient. The monitoring solution providing the metrics must have methods to filter out any type of monitoring data not desired, and be able to hold the real-time granular data for a limited period of time to enable detailed review of the measurement results for troubleshooting or incident reports. The northbound consumer of the data should keep records for at least one year, possibly rolled up to daily aggregates to reduce data storage requirements.

Reports and dashboards for the monitoring data should be created to suit the SD-WAN applications offered. This may include management-level geographical overview charts, as well as detailed regional views for network operations center (NOC) troubleshooting audiences.

7. Approaches to Optimizing SD-WAN Performance Using End-to-End Monitoring Data

For SDN solutions in general, and SD-WAN in particular, using active and passive end-to-end or hop-by-hop monitoring metrics in a feedback loop towards the SDN/SD-WAN controller provides a means to automatically assist the traffic and service control functions with external quality metrics.

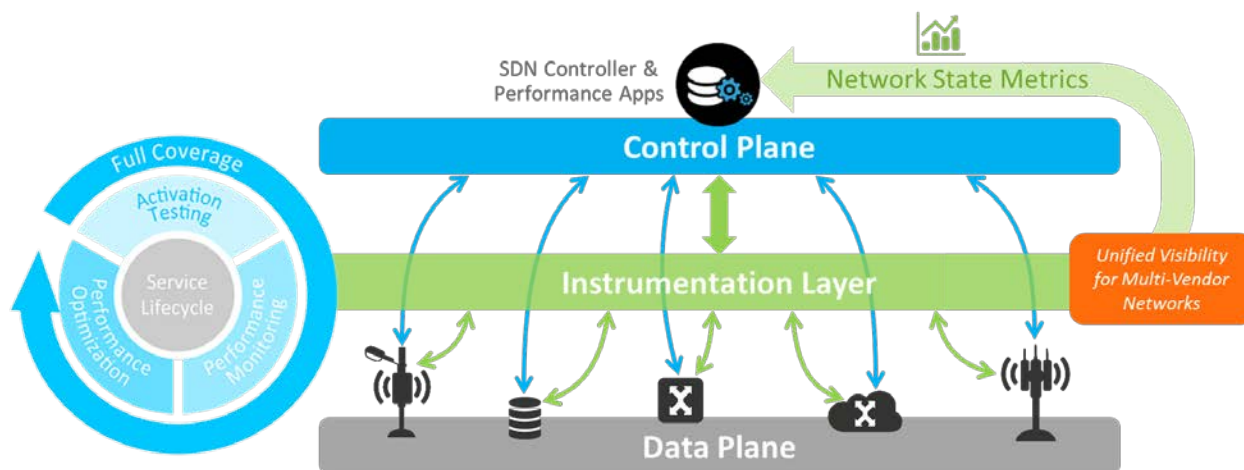


Figure 16 - Unified and Complete Instrumentation Layer Unlocks Value for the Control Plane

Providing the SDN controller with direct network quality metrics allows for immediate feedback on the impact of reroute decisions or path changes, and ultimately allows the controller to assess—via machine learning methodologies—the impact of each decision before it is taken. Automation of the full service lifecycle is thus possible.

This monitoring is not limited to active Layer 2-3 performance data (delay, jitter, loss) but could also be augmented with software-based agent type monitoring of a specific service type or path as well as hardware-based, highly granular utilization type metrics that can alert on links or paths experiencing high frequency of micro-bursts or micro outages.

Conclusion

SD-WAN for enterprises is becoming more relevant all the time, driven by growing use of cloud-based applications. The advantages of SD-WAN—regardless of whether it is deployed using a hybrid model or a full-scale implementation—cannot be ignored. And, as widespread adoption of this technology continues, operators must be able to deliver the same level of quality with their SD-WAN managed services as they do with traditional WAN offerings.

With that goal in mind, all SD-WAN lifecycle phases can benefit from a flexible, NFV-based performance monitoring solution that scales beyond the footprint of the SD-WAN cloud and can send performance flows from any starting location to any destination in the network infrastructure. Such a solution can be used to:

- Cover large scale hub-spoke and full-mesh topologies with active, micro-second accurate, standards-based performance monitoring towards thousands of endpoints continuously.
- Bring standards-based turn-up testing, monitoring, and OAM functions to all SD-WAN endpoints, by adding NFV-enabled vCPE modules or orchestratable lightweight software agents. Since the solution is standards-based, standard networking devices can also act as responders to performance monitoring flows.
- Monitor micro-outages, one-way delay and variation, and SLA compliance by delivering precise and granular metrics.
- Centralize test control and automation, integrated with existing OSS, by pairing vPACs with NMS solutions.
- Deliver a new level of PM workflow automation with results centrally stored for comparison to predefined QoS templates or SLA levels. Tests—conducted one-way or bi-directionally, in an end-to-end or segmented manner—can be scheduled on demand or triggered by service endpoint installation.
- Provide open access to turn-up data and results—including customer-ready reports reflecting their specific SLAs—using the API.

Abbreviations

API	application programming interface
CCM	continuity check message
CoS	class of service
COTS	commercial off-the-shelf
CPE	customer premises equipment
CSV	comma separated values
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data Over Cable Service Interface Specification
DMM/DMR	delay measurement message / delay measurement response
DPI	deep packet inspection
FM	fault management
FQDN	fully qualified domain name
ICMP	Internet Control Message Protocol
JSON	JavaScript object notation
MEF	metro ethernet forum
M2M	machine-to-machine
MSO	multiple systems operator
NFV	network functions virtualization
NFV-PM	network functions virtualization performance monitoring
NFVi	network functions virtualization infrastructure
NID	network interface device
NMS	network management system
NOC	network operations center
OAM	operations and maintenance
OTT	over the top
OSS	operational support systems
PM	performance management
QoE	quality of experience
QoS	quality of service
SAT	service activation test
SaaS	software as-a-service
SD-WAN	software-defined WAN
SDN	software-defined networking
SLA	service level agreement
SNMP	Simple Network Management Protocol
SOAM	service OAM
TCP	transmission connection protocol
UDP	user datagram protocol
vCPE	virtualized customer premises equipment
VLAN	virtual local area network
VNF	virtual network function
vPAC	virtualized performance assurance controller

WAN	wide area network
XML	extensible markup language

Bibliography & References

- MEF Forum*. July 2014. CE 2.0 Service Management Life Cycle White Paper. Retrieved from: http://mef.net/Assets/White_Papers/CE_2_0-Service_Life_Cycle_White_Paper.pdf
- ITU-T*. August 2015. G.8013/Y.1731: OAM functions and mechanisms for Ethernet-based networks. Retrieved from: <https://www.itu.int/rec/T-REC-Y.1731>
- MEF Forum*. January 2012. Implementation Agreement MEF 23.1: Carrier Ethernet Class of Service – Phase 2. Retrieved from: http://dev.mef.net/Assets/Technical_Specifications/PDF/MEF_23.1.pdf
- ITU-T*. February 2016. Y.1564: Ethernet Service Activation Test Methodology. Retrieved from: <https://www.itu.int/rec/T-REC-Y.1564/en>
- Yum, K. IETF*. October 2008. RFC-5357: A Two-Way Active Measurement Protocol (TWAMP) Retrieved from: <https://tools.ietf.org/html/rfc5357>
- Breznick, A. Heavy Reading*, January 2017. How cable can conquer the enterprise market. Retrieved from https://accedian.com/wp-content/uploads/2017/01/HR_Accedian_Cable_Enterprise_WP_1-24-17.pdf
- Sterling, P. Heavy Reading*. February 2017. Operator Success in the New Age of the Software-Defined WAN, Retrieved from <https://resources.ext.nokia.com/asset/201132>
- MEF Forum*. April 2012. Introducing the Specifications of the MEF. MEF 38: Service OAM Fault Management YANG Modules Technical Specification. Retrieved from: <http://slideplayer.com/slide/5687304/>
- MEF Forum*. March 2016. Service Operations Specification MEF 55: Lifecycle Service Orchestration (LSO): Reference Architecture and Framework. Retrieved from: http://dev.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf
- Bradner S., McQuaid J.*, March 1999. RFC2544. Benchmarking Methodology for Network Interconnect Devices. Retrieved from <https://www.ietf.org/rfc/rfc2544.txt>
- IEEE*. December 2007. 802.1ag - Connectivity Fault Management. Retrieved from: <http://www.ieee802.org/1/pages/802.1ag.html>

SD-WAN and Beyond: Delivering Virtual Network Services

A Technical Paper prepared for SCTE•ISBE by

Ralph Santitoro

Head of SDN/NFV/SD-WAN Solutions
Fujitsu Network Communications
(805) 791-0711
ralph.santitoro@us.fujitsu.com

Introduction

Software-defined wide area networks (SD-WANs) have generated much enthusiasm in the industry because they solve real business challenges for both enterprise subscribers and communications service providers (CSPs). SD-WANs leverage software-defined networking (SDN), network functions virtualization (NFV) and lifecycle service orchestration (LSO) technologies making them the ideal foundation for deploying new, on-demand virtual network services. This paper describes the fundamental capabilities of SD-WAN services, their operational and deployment considerations, key benefits to enterprise subscribers and CSPs, and use cases for connecting places and things. The paper also discusses key architectural and deployment considerations required to extend an SD-WAN service via virtual network functions (VNFs) operating on virtual customer premises equipment (vCPE) and virtual private clouds (VPC) to deliver additional virtual network services.

The Journey to Virtual Network Services

SD-WANs are the confluence of several technologies that have developed over the years augmented with newer technologies providing virtualization and centralized management and control resulting in virtual network services. The evolution to SD-WAN as we know it today fundamentally consists of wide area network (WAN) connectivity using the Internet protocol (IP) to create virtual private networks (VPNs) secured typically through IPsec-encrypted tunnels. The IPsec tunnels operate over a physical underlay (transport) network. This enables SD-WANs to operate over multiple WANs types, e.g., dedicated Internet access and MPLS, referred to as hybrid WANs. SD-WANs also support WAN optimization which serves two purposes; increase the amount of usable bandwidth and correct for packet loss over WANs.

SD-WANs can be constructed using the aforementioned technologies. However, SD-WANs didn't become highly popular until automation was added via centralized management and control using SDN, NFV and end-to-end service orchestration. Refer to **Figure 1**.




Figure 1 - The journey to SD-WAN and beyond

What is an SD-WAN and what does it do?

SD-WANs are over-the-top (OTT) virtual overlay networks that operate over any underlay network. This means that with SD-WANs, any network topology can be created over wired or wireless access and core transport networks. This unique property enables SD-WANs to be created over underlay networks using different technologies such as Carrier Ethernet, broadband Internet [digital subscriber line (DSL), Cable, or passive optical network (PON)], WiFi or LTE access networks or IP or MPLS core networks. Also, because SD-WANs create virtual overlay networks, one can create any topology to interconnect sites and connect sites to their public and private clouds, software-as-a-service (SaaS) applications running in the cloud and their data centers.

To simplify operations, reduce costs and enhance agility and security, enterprises are accelerating the migration of applications running on servers on their premises, e.g., Microsoft Exchange Server, to the cloud using a SaaS applications, e.g., Microsoft Office 365. Because of this, SD-WANs will play an even larger role in interconnecting sites more often to the cloud rather than to other sites since information exchange will be done via cloud-centric applications. Refer to *Figure 2*.

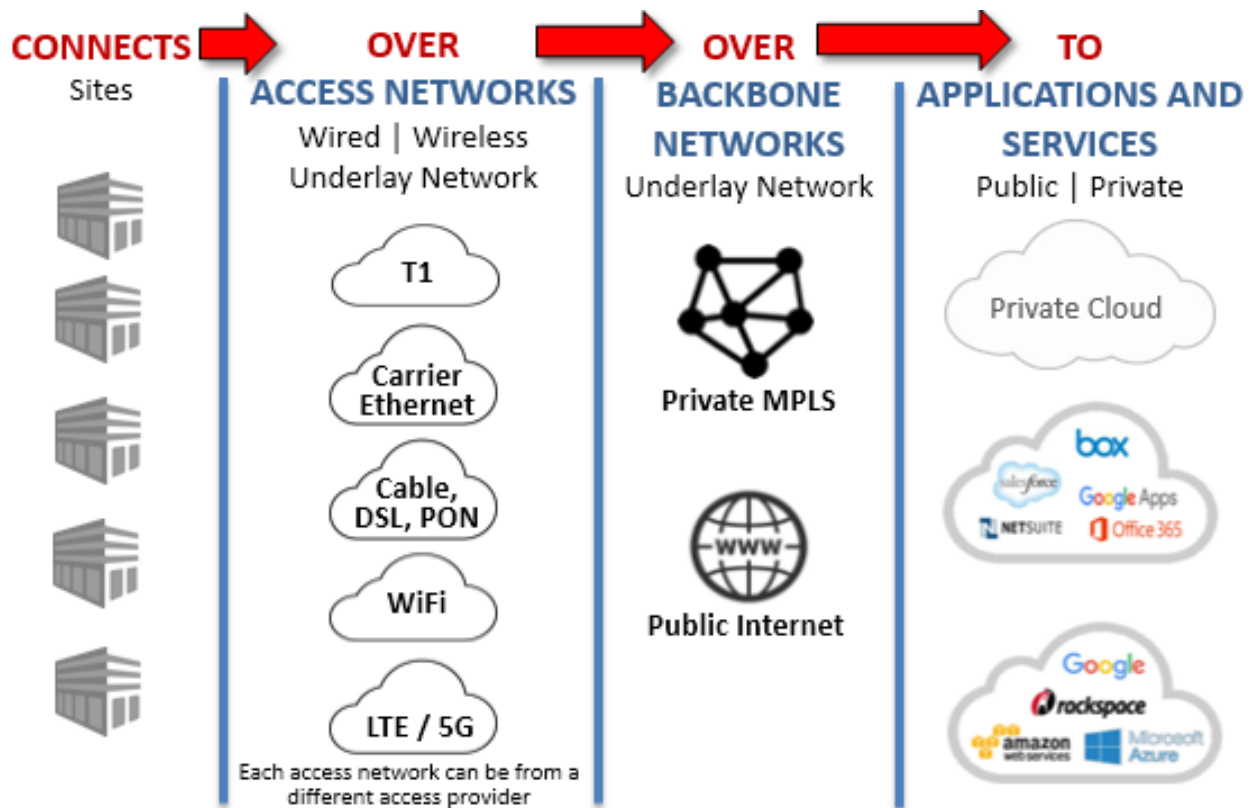


Figure 2 - SD-WAN interconnect sites and application

SD-WAN: Service provider threat or opportunity?

CSPs who currently offer MPLS connectivity services for enterprises, may find SD-WAN services a threat to their existing MPLS business. In some ways this is true given that MPLS service bandwidth typically cost 10-20 times more than broadband Internet. Unlike broadband Internet, MPLS does provide certain quality of service (QoS) performance assurances for some of the bandwidth through different classes of service. However SD-WAN's bandwidth optimization technologies can often compensate for QoS performance limitations of broadband Internet.

This poses an interesting dilemma for CSPs who want to offer or currently offer SD-WAN services. The high demand from SD-WAN services from enterprise subscribers is forcing the issue and CSPs are adapting. MEF Forum sponsored an industry survey on this topic and found that almost half of survey respondents (45%) considered SD-WAN services as a strategic opportunity while only 4% considered them to be a threat. Finally, 37% considered SD-WAN services to be both a threat and opportunity. Refer to *Figure 3*.




Figure 3 - SD-WAN results from MEF survey

CSPs will obtain incremental revenue from SD-WAN services while MPLS service revenue growth will be reduced or remain flat as SD-WAN enables Internet (both broadband and dedicated) to be used to grow bandwidth in addition to the current MPLS bandwidth in use. More importantly, CSPs must view SD-WAN services differently since they provide much more than connectivity. This will become clearer as SD-WAN service capabilities are discussed in the next section.

Fundamental capabilities of SD-WAN services

All SD-WAN services provide some fundamental capabilities. As with any new technology that has become popular and garnered the public's attention, SD-WAN too has had some overzealous marketers associate their products with it. This has been compounded by the lack of any standard service definition for SD-WAN or even the components used to construct an SD-WAN service. MEF Forum has embarked upon addressing this and has created working groups to define the service components, fundamental capabilities, reference architectures and implementations, and market education for SD-WAN services. MEF has also created the technical specification "MEF 55 Lifecycle Service Orchestration (LSO): Reference Architecture and Framework". This will be used for the management and orchestration for MEF-defined SD-WAN services.

Per the MEF work, the following describes the fundamental capabilities for an SD-WAN service. These capabilities are described in more detail in the MEF's paper "[*Understanding SD-WAN Managed Services: Service Components, MEF LSO Reference Architecture and Use Cases*](#)"

SD-WAN services provide encrypted IP tunnels (SD-WAN tunnels) over the underlay transport networks. Since SD-WAN services can be created over the Internet in addition to private networks, e.g., MPLS or Carrier Ethernet, encryption and some firewall functionality is critical to have a viable service. SD-WAN services operate over any wired or wireless underlay transport network. SD-WAN tunnels are built over these underlay networks and do not require any modifications to them.

SD-WAN services take QoS performance measurements (PMs) over each WAN to identify the packet loss, delay (latency) and delay variation (often referred to as jitter). These QoS PMs are used to make application forwarding decisions over SD-WAN tunnels which operate over the different WAN underlay networks.

Unlike other connectivity services, SD-WAN services forward packets based on application type thus making the service much more desirable for enterprise subscribers. SD-WAN services can identify the specific application, e.g., Skype for Business, or grouping by application type, e.g., real-time applications, and decide over which WAN the application should traverse. Other WAN services, such as Carrier Ethernet, MPLS or Internet, focus on forwarding packets at the network layer with no knowledge of the application to which those packets are associated. SD-WAN services forward the packets by application using QoS, security or business priority policies. This capability provides great value to enterprise subscribers because they are more interested in application performance than packet performance. Because policy management is centralized, it is less error prone than having to push policies down to each device individually via a device command line interface (CLI) or scripts.

SD-WAN services use WAN load balancing, diverse WAN and access network providers, and wireline plus wireless WANs to achieve a high availability service. One or more of these techniques may be used in an SD-WAN service deployment.

Unlike other WAN services, SD-WAN services achieve high levels of automation through centralized management, control and orchestration taking advantage of SDN and LSO technologies. This results in new enterprise sites to be turned up literally in minutes. Site configuration information, such as LAN and WAN IP addresses, number of WANs, and WAN types, is collected as part of the site planning and provisioning process. This information can then be prepopulated into a site ‘profile’ which is then used to configure the SD-WAN device. When the device is cabled to an Internet WAN and powered up, it can remotely retrieve its configuration from the site profile in a manner similar to how cable modems remotely retrieve their configuration once they are powered up. This automated configuration is referred to as zero touch provisioning (ZTP).

Finally, WAN optimization is an important part of an SD-WAN service and uses techniques to increase WAN bandwidth utilization for a given amount of WAN bandwidth. WAN optimization can include capabilities to minimize WAN bandwidth using data compression, TCP optimization, data caching, and data de-duplication techniques. These techniques reduce the amount of information that must be transmitted and thus free up WAN bandwidth for other applications. WAN optimization can also reduce packet loss introduced in the underlay network by using forward error correction (FEC). FEC sends additional information enabling the receiving SD-WAN device to reconstruct packets negating the need for the sender to completely retransmit the packet saving WAN bandwidth.

SD-WAN service components

To add clarity to the industry, MEF Forum has defined five service components used in an SD-WAN service as illustrated in **Figure 4**. Some service components are used internally by the service provider while others are located on the customer premises and used by the enterprise subscriber.




Figure 4 - SD-WAN service components defined by MEF

The SD-WAN Edge is a physical or virtual appliance which is placed on the customer premises, data center or cloud. The SD-WAN Edge may consist of a CPE, a VNF running on a vCPE, or a VNF running in a virtual private cloud, e.g., running on a compute instance in Amazon Web Services (AWS). The SD-WAN Edge initiates and terminates the SD-WAN tunnels over WANs plus measures WAN QoS performance. It also enforces application-based QoS, security and business priority policies that are used to steer packets over different SD-WAN tunnels. The SD-WAN Gateway is a special case of an SD-WAN Edge which enables sites interconnected via SD-WAN tunnels to connect to sites without SD-WAN Edges that are interconnected via other private networks such as MPLS or Carrier Ethernet.

The SD-WAN Controller is responsible for the centralized management each SD-WAN Edge and SD-WAN Gateway under its control. This may entail pushing down configuration and policies received from the Service Orchestrator or receiving alerts and alarms which are subsequently sent to the Service Orchestrator. The Service Orchestrator is responsible for lifecycle service management of the SD-WAN

service and may interface with one or more SD-WAN Controllers depending upon the size of the network or geographic placement requirements. Note that some implementations combine Service Orchestrator and SD-WAN Controller. However, these two functions have been separated to facilitate placement into the MEF LSO RA. Finally, the Service Orchestrator is often used to centrally manage other services in addition to SD-WAN services.

The Subscriber Web Portal enables authorized and authenticated enterprise users to modify an SD-WAN service. Such changes may include modifying SD-WAN bandwidth, adding security policies, and adding or removing SD-WAN tunnels (service connectivity) between sites.

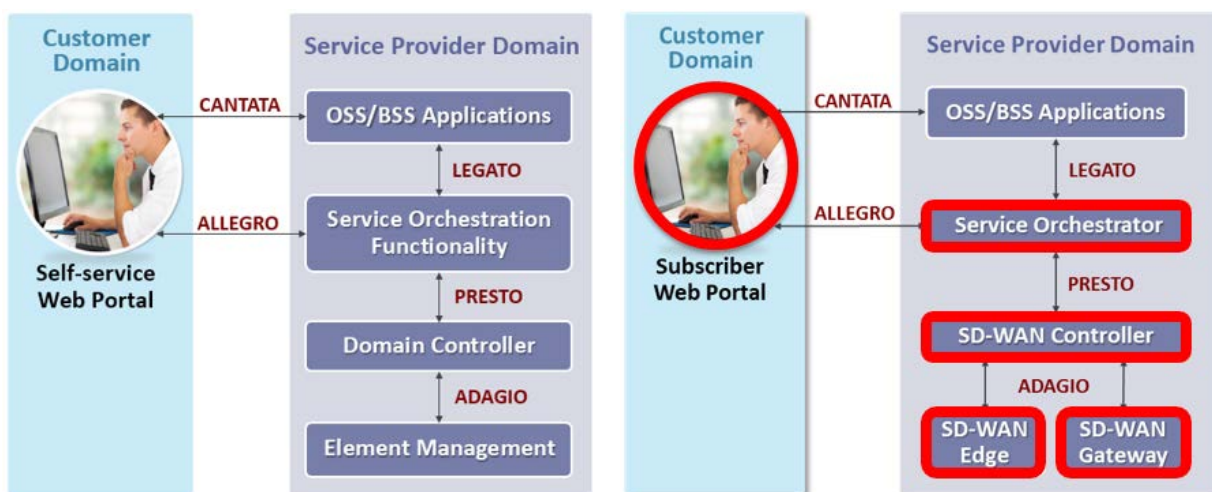


Figure 5 - SD-WAN service components mapped into MEF 55 LSO RA

The diagram on the left in **Figure 5** is the MEF 55 LSO reference architecture (RA) which defines different management interface reference points. The diagram on the right of **Figure 5** illustrates where the aforementioned SD-WAN service components, highlighted in red, are mapped in the MEF LSO RA. MEF members will use this RA to construct reference implementations and developed application programming interfaces (APIs) and data models to facilitate implementations.

SD-WAN use cases

1. SD-WAN Use Case: Hybrid WAN

This use case illustrates how two enterprise sites are interconnected via an MPLS VPN service and how they use the Internet to access public web sites, SaaS applications, cloud service providers, etc. The Internet service provider (ISP) may be different for each site as illustrated by ISP A and ISP B. Since Internet bandwidth often costs 10-20 times less than MPLS VPN bandwidth, the enterprise would like to use both MPLS VPN and a secured Internet to interconnect the sites since their MPLS VPN bandwidth is insufficient. This would enable them to increase inter-site bandwidth without purchasing additional MPLS VPN bandwidth. Refer to the present mode of operation (PMO) in **Figure 6**.




Figure 6 - PMO: Only MPLS VPN used for inter-site connectivity

In the future mode of operation (FMO), the enterprise subscriber uses the SD-WAN service to create SD-WAN tunnels across the Internet and MPLS VPN in effect sharing bandwidth across the different WANs. The SD-WAN tunnels across the Internet are encrypted so the enterprise information is secured. Each site can still connect to the public web sites via local Internet breakouts at each site. Furthermore, since an SD-WAN service provides application-based traffic forwarding, the enterprise can decide which WAN they want the application to traverse as indicated by the two red arrows in **Figure 7**.

The WAN selection for a given application is determined by QoS, security or business priority policies. For example, an enterprise may set a QoS policy to send Skype for Business over any WAN as long as the packet loss is less than 2% and the packet latency is less than 40ms. A retailer may set a business priority policy to send all payment card transactions ahead of any other traffic since these are most important for their business. A financial institution may set a security policy whereby all inter-bank transactions are only sent over the SD-WAN tunnel over the MPLS VPN.




Figure 7 - FMO: MPLS VPN and Internet used for inter-site connectivity

2. SD-WAN Use Case: Secure Connectivity to Virtual Private Cloud

Enterprises are increasingly migrating applications running on site or in their data center to subscribing to a SaaS equivalent, e.g., migrating Microsoft Exchange Server to Office 365 SaaS. Furthermore, enterprises are increasingly renting virtual compute resources like infrastructure-as-a-service (IaaS) rather than purchasing physical servers and operating them on premise. As these applications and workloads migrate to the cloud, enterprises need to provide secure and increasingly higher bandwidth WAN connections to their VPC and SD-WAN services are an efficient and flexible way to support this. **Figure 8** illustrates an SD-WAN service interconnecting two sites over both MPLS and Internet WANs as discussed in the hybrid WAN use case in **Figure 7**. In this case, however, an SD-WAN Edge VNF is instantiated in the cloud compute instance (IaaS) thus extending the SD-WAN to the VPC.




Figure 8 - Use Case for secure connectivity to cloud

3. Use Cases for Placement of SD-WAN Edge and other VNFs

An SD-WAN Edge CPE is the most common implementation since it is the simplest and follows widely established practices deploying a physical appliance at customer premises. Since an SD-WAN Edge can also be implemented as a virtual appliance via a VNF, many interesting possibilities are introduced as to where an SD-WAN service can be extended. As illustrated in **Figure 9**, the SD-WAN Edge VNF could run on a vCPE at the customer premises similar to how an SD-WAN Edge CPE is deployed. However, in this case, the vCPE could be sized to support the SD-WAN VNF and VNFs delivering additional functions and services.

An SD-WAN Edge VNF could also run on a kiosk or automated teller machine (ATM) in a mall, sports stadium or temporary location providing secure connectivity to data centers or the cloud without requiring additional physical equipment and cabling since the VNF is software. Edge computing infrastructure is another use case where compute functionality is provided much closer to the customer premises in addition to traditional centralized, regional data centers (DC).

As network aggregation points near the edge of the network, e.g., cable modem termination system (CMTS) or broadband network gateway (BNG), become virtualized as one or more VNFs, more compute resources become available for other services or functions. SD-WAN Edge or Gateway VNFs, vCMTSs and vBNGs, and other services or functions delivered by VNFs can be added to these edge computing nodes which act as mini edge data centers. Finally, SD-WAN Edge VNFs running in public cloud environments such as AWS, Google Cloud or Microsoft Azure enable enterprises to extend their inter-site secure SD-WANs up to their virtual private cloud applications.




Figure 9 - Interesting possibilities for SD-WAN Edge placement

Conclusion

SD-WAN is the compilation of several networking technologies developed over the years plus newer SDN, NFV and LSO technologies resulting in an agile and flexible virtual network service. Through virtual overlay tunnels, SD-WAN services are decoupled from the underlay transport network. This results in rapid service deployment over any wired or wireless network. SD-WAN services provide much more than connectivity including application-awareness and policy-based packet forwarding. The virtualization of SD-WAN Edges enables SD-WAN services to be delivered beyond traditional 'brick and mortar' buildings, like other WAN connectivity services, and extend secure connectivity to the cloud and other types of devices. Finally, SD-WAN service capabilities, terminology and reference architectures are being defined by MEF Forum which will facilitate and accelerate implementations and service deployments.

Abbreviations

API	application programming interfaces	NFVI	Network Functions Virtualization Infrastructure
ATM	automated teller machine	NOC	network operations center
AWS	Amazon Web Services	OSS	operational support systems
BB	broadband	OTT	over the top
BSS	business support systems	PAYGO	pay as you go
CE	Carrier Ethernet	PM	performance metrics
CORD	Central Office Re-architected Datacenter	PMO	present mode of operation
CLI	command line interface	PON	passive optical network
CPE	customer premises equipment	QoS	quality of service
CSP	communications service provider	RA	reference architecture
DC	data center	SaaS	Software-as-a-Service
DIA	dedicated Internet access	SDN	software-defined networking
DSL	digital subscriber line	SD-WAN	software-defined wide area network
FMO	future mode of operation	TCP	transport control protocol
IaaS	Infrastructure-as-a-Service	vBNG	virtual Border Network Gateway
IP	Internet protocol	vCMTS	virtual Cable Modem Termination System
IoT	Internet of things	vCPE	virtual customer premises equipment
ISP	Internet service provider	VNF	virtual network function
LSO	Lifecycle Service Orchestration	VNS	virtual network services
LTE	long term evolution (4G cellular networks)	VPC	virtual private cloud
MPLS	multi-protocol label switching	VPN	virtual private network
NFV	Network Functions Virtualization	WAN	wide area network

Bibliography & References

[*Understanding SD-WAN Managed Services: Service Components, MEF LSO Reference Architecture and Use Cases*](#); Ralph Santitoro; [MEF Forum](#)

[*SD-WAN Fundamentals, Use Cases and MEF LSO Reference Architecture*](#), Ralph Santitoro; NFV World Congress (May 2017)

[*SD-WAN Managed Services, Terminology, Use Cases and Challenges*](#), Ralph Santitoro and Peter Agnew; BrightTalk Webinar

[*MEF 55 Lifecycle Service Orchestration \(LSO\): Reference Architecture and Framework*](#), MEF Forum

Network Service Descriptors as a Factory

Agile Networking for Differentiating MSO Business Services

A Technical Paper prepared for SCTE•ISBE by

Hans Vanderstraeten

Team Lead Virtual Networks Orchestration
Nokia
Copernicuslaan 50, B2018 Antwerp
323-240-7905
hans.vanderstraeten@nokia.com

Erwin Six, Lead Solution Architect VNO, Nokia

Mihai Fagadar, R&D Team Lead VNO, Nokia

Willem Acke, Lead Solution Architect VNO, Nokia

Introduction

Business Services have been a new ambition for MSOs over the past decade. The Service Level Agreements, among others, associated with connectivity services, imply typically customer delivery lead times of weeks, up to months when the MSO must go off-footprint. By contrast, the same business goes to the Cloud Service Provider, and gets compute and storage infrastructure at his disposal in minutes. These businesses are soon expecting network services to be delivered together with compute and storage, in minutes. SD-WAN delivers on these promises, and drives a need for end-to-end Network Orchestration.

Orchestration has become an overloaded word in the industry. While Network Orchestration and Application Orchestration share a lot of technologies coming from the domains of SDN and NFV, the two approaches serve different objectives. The paper will specifically address Network Orchestration, aiming at launching new revenue generating Business Connectivity services rapidly, with Value-Added Services running in private or public cloud data centers chained-in into the Network Service. MSOs will be able to tap into new revenue streams through the extension of their addressable market going off-footprint, adding new Value-Added Services to their product portfolio, and becoming the trusted middle man in the value chain between Enterprises and public cloud service providers.

The rapid market traction of SD-WAN overlay IP networking, enabled by product suites such as Nuage, fuels the requirements for Network Orchestration: overlay to underlay connectivity, connectivity into Public Cloud, Internet Break-out, flexible onboarding of virtualized Value-Added Services (VAS) such as firewalls, wan optimizers, email filters, and others, all easily customizable for each specific enterprise customer of the MSO.

Various Network Orchestrators, open source and others, have emerged. However, while these all enable writing a Network Service Descriptor (NSD) for a specific service chain use case, including a specific VAS over a purpose-built plug-in, significantly more is needed to deliver Network Services and NSDs rapidly, massively, and reliably. An approach is required that allows to produce NSDs as-a-factory, rather than just ‘executing an NSD’.

The paper will discuss a combination of approaches that together allow to deliver Network Services by MSOs in the most competitive way, and more specifically, how higher automation can be achieved through a ‘best-practice’-based abstraction of enterprise services. A combination of novel Communication and Software technologies are combined to, as examples, automatically generate TOSCA types and plug-ins out of (YAML-based) API specifications; write complex NSDs simply and reliably through substitution and decomposition; and manage a library of NSDs through release upgrades of plug-ins through an extended suite of Continuous Integration and Continuous Testing.

On Orchestration and Virtual Networks

Orchestration is a term frequently used when referring to Software Defined Networking (SDN) and Network Functions Virtualization (NFV) solutions. In essence, Orchestration aims to answer the problem of service agility, service velocity and OPEX control in both network connectivity and application domain. However, different meanings are associated with the word and different types of orchestration are emerging with different domains, roles and responsibilities.

1. ETSI MANO Orchestration

In the communications networks context, including SDN and NFV, orchestration refers to the automation of order fulfillment, service assurance, as well as other business processes in a software-controlled environment. Orchestration of previously manual processes is key to unlocking the OPEX savings expected from SDN and NFV, where moreover each of the SDN and NFV functionalities being orchestrated are autonomous systems on their own. The term alludes to a conductor who orchestrates music for an orchestra and ensures that all the musicians play in tune and rhythm. In turn, a Network Service Designer is the composer determining the notes to be played by each SDN controller, Virtualized Network Function Manager (VNFM), and any other Networking Function, while the Conductor or Orchestrator directs the overall performance of the orchestrated composition.

Further, SDN and NFV do promise for a much more easy consumption of applications and networks, as MSOs and other Communication Service Providers (CSPs) are being challenged by their customers (enterprises and residential, fixed and mobile) to be ‘satisfied instantaneously’ – the audience being the ultimate judge of the performance of the orchestra.

Standardization forums such as TM Forum, ETSI, Metro Ethernet Forum (MEF), the Open Networking Foundation (ONF), and others cover parts of the orchestration spectrum each with a different focus. Some are more specific about network connectivity, some are more specific about Virtual Network Functions (VNFs), and some are more specific about service decomposition.

ETSI proposes the following reference architecture in the NFV domain:

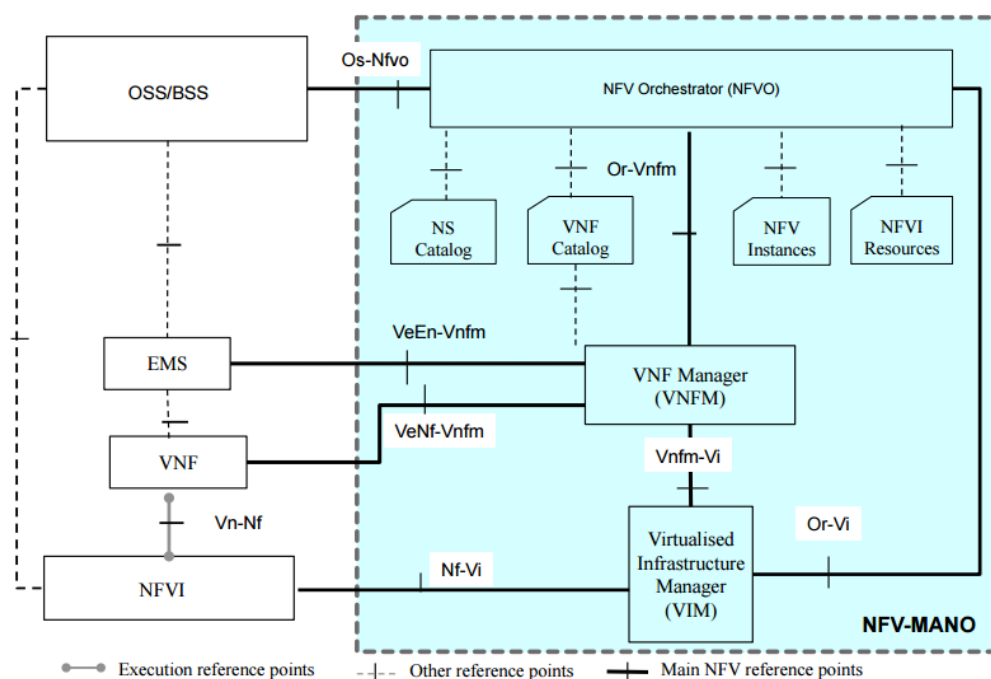


Figure 1 – ETSI MANO Reference Architecture

- NFV Orchestrator (NFVO): Responsible for on-boarding of new Network Services (NS) and Virtual Network Function (VNF) packages; NS lifecycle management; global resource management; validation and authorization of Network Functions Virtualization Infrastructure (NFVI) resource requests;
- VNF Manager (VNFM): Oversees lifecycle management of VNF instances; coordination and adaptation role for configuration and event reporting between NFVI and E/NMS;
- Virtualized Infrastructure Manager (VIM): Controls and manages the NFVI compute, storage, and network resources.

Note that in the ETSI MANO NFV framework, there are essentially two orchestrators: the VNFM as an Application Orchestrator, and the NFVO as a Network Orchestrator. Further, the ETSI framework applies to Virtualized Network Functions, and therefore spans the Datacenter domain only.

Extending upon the ETSI MANO definition, Figure 2 adds the Service Orchestration (SO) layer to the Business Service level which essentially concatenates several Network Services with Physical Network Functions or their management agents, bridging the legacy and virtual domains.

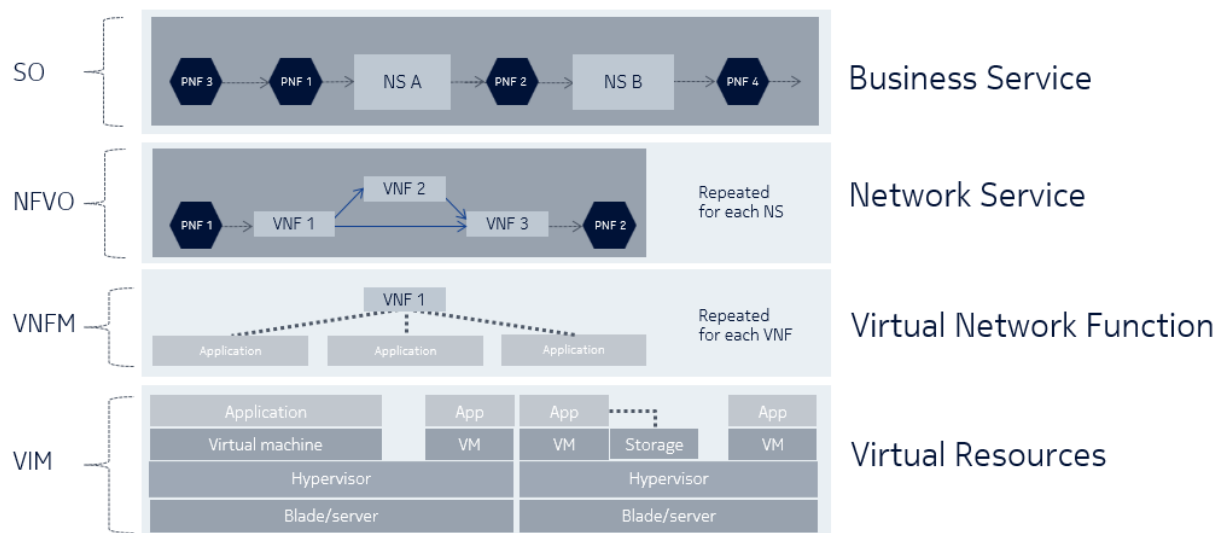


Figure 2 – Service delivery terminology (ETSI-MANO)

Given that orchestration can be defined at different layers and domains, a central concept to come to a meaningful split of responsibilities among the orchestrators is “separation of concern”. A Network Orchestration problem is divided across domains, where each of the domains is defined by the best practices and potentially organizational split of accountability and responsibility. Therefore, the Network Service Orchestration is essentially a well-structured collection of the domain-specific orchestrators, managers, and controllers. Each orchestrator decomposes and automates services, resources, and tasks in their respective domains.

2. Extending Orchestration

In a collaborative effort with the most of the industry actors and vendors, Verizon has brought forward a very specific proposal for the above mentioned “well-structured collection”. Figure 3 is an extract of this public Verizon document:




Figure 3 – Verizon SDN-NFV Reference Architecture

This architecture brings together the best of the “pure” ETSI MANO reference framework, the reality of every operator’s network, and the market pressure to introduce new network services rapidly. It extends the NFVO-VNFM (datacenter) axis, per ETSI-MANO, with two more dimensions:

- Datacenter - CSP networking dimension, through the introduction of multiple domain controllers (DC SDN, WAN SDN, Access SDN, and Domain Specific controllers), therefore extending ETSI MANO outside of the datacenter;
- Virtual-Physical dimension, where Physical Network Functions (outside or inside the Datacenter) also can be orchestrated from the same 'end-to-end orchestration' functional block.

Further, the Verizon reference architecture allows for the extended configuration management as an example, in addition to EMS-led configuration.

3. Network Orchestration versus Application Orchestration

In the previous chapter, Orchestration is explained as a generic, over-arching concept spanning all networks and all applications as they move to Software (Defined Networking) and (Network Function) Virtualization. Given the vast scope of this transition to SDN and NFV, the industry is tackling the problem from two, essentially orthogonal angles: an application (NFV-led) orchestration problem, and a network (SDN-led) orchestration problem.

As an application orchestration example, consider the deployment of a virtualized mobile core. IMS/VoLTE is a very complex application, consisting of several functions that need to be well-orchestrated through the VNFM. The NFVO function will bring the vEPC and vIMS together into a mobile core, and does this in a static, pre-defined way.

Network orchestration orchestrates several networking functions end-to-end, such that an IP packet flows end-to-end according to the rules and policies set out in the network by the orchestrator. This can be a very complex task, as IP underlay and overlay (further defined in Section 4), different domains separated by gateways, physical and virtual appliances and routers may have to be orchestrated to keep the packet flowing. In contrast, the Virtual Network Functions involved in the flow, such as the firewalls, load balancers, WAN optimizers, as applications are typically rather simple “point devices”.

The application and network orchestration approaches currently address very different market demands. Application orchestration focuses on getting the best operational benefits first for otherwise ‘known’ applications such as VoLTE, Network Orchestration focuses on getting the new services enabled by SDN in the market to generate new revenues as soon as possible.

4. Virtual Networks Orchestration and Service Chaining Use Cases

Figure 4 depicts a visual of a Service Chain: a number of Virtual Network Functions, instantiated in the Datacenter as part of the Service Chain through Network Orchestration. In this paper, the authors assume a so-called Overlay Network, VxLAN-based, interconnecting the VNFs through a Software Defined Networking Controller, such as Nuage Networks. The Service Chain further extends the overlay into the WAN domain, where a Network Services Gateway (NSG) connects the Enterprise LAN to the overlay network, transported over the WAN underlay. This Software Defined WAN (SD-WAN) approach allows separating concerns between underlay and overlay, that can support separate administrative domains. The NSG Border Router (NSG-BR) acts as a Datacenter Gateway (DC-GW), bridging underlay and overlay.

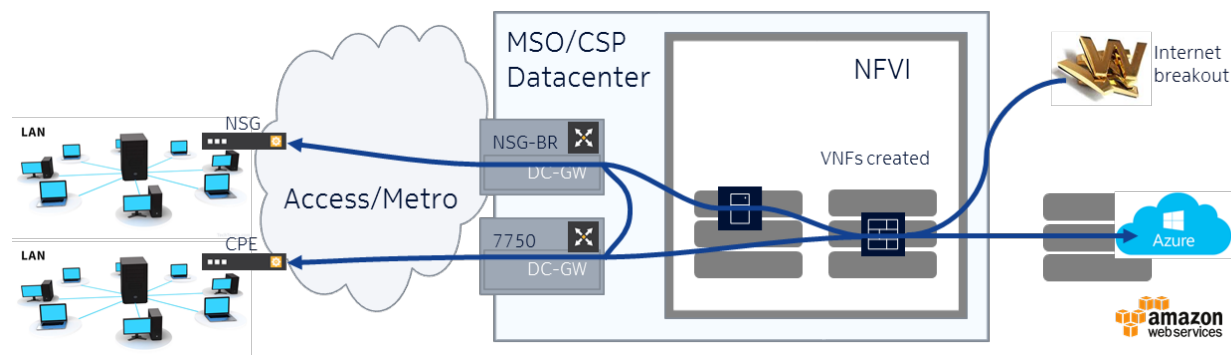


Figure 4 – Enterprise VPN Service Chain example

Service Chains can be further extended into public cloud service providers, such as Amazon Web Services or Microsoft Azure. Interconnection schemes in Virtual Private Clouds (VPCs) are typically proprietary but well-published and accessible through APIs.

The set up and automation of the Service Chain includes three basic elements of scope:

1. End-to-end overlay connectivity across SD-WAN, the Enterprise Cloud, the CSP's Cloud, and the Public Cloud;
2. Connecting the overlay to the underlay where / when needed;
3. Dynamically include IP appliances in the overlay as VNFs.

This approach is denoted as Virtual Networks Orchestration (VNO): as both Network Functions and Connectivity both together start living in the Cloud, a whole new set of opportunities arise, but at the same time, a new set of challenges come with it.

Several Service Chaining use cases have surfaced in the industry:

- Dynamic Enterprise Services: taking L2 and L3 VPN Enterprise connectivity services to the Cloud, with SD-WAN connectivity allowing for on-net and off-net Enterprise VPN 'instantaneous' overlay connectivity, and with VAS service offerings allowing for upsell. The automated connection of enterprises into the Public Cloud Service Providers and their XaaS products will bring the MSOs and CSPs into the value chain as a trusted middleman, opening the opportunity for MSOs to tap into this one hundred billion dollar market. The discussions further in this paper will use Dynamic Enterprise Services use cases as examples.
- Datacenter Consolidation is a common consideration for organizations that plan to reduce the size of a single facility or merge one or more facilities to reduce overall operating cost and reduce IT footprint. Service Chains typically live across datacenters, and will stream-line the network policies across WAN and datacenters.
- GiLAN Service Chaining is contained in ability to steer traffic from the Gi/SGi interface in EPC through different VAS appliances en route to external networks.
- Virtual Residential Gateway (vRGW) transforms the L3 routed Residential Gateway in the home into bridged RGW, by moving its Layer 3 functions into the CSPs Broadband Network Gateway (BNG). Enriching this capability with features such as Home LAN extension does allow for VAS-in-the-Cloud towards the residential market, such as Cloud Network Attached Storage and Parental Control.

5. The multipliers driving the complexity of Service Chains

Referring again to Figure 4, and specifically for the Dynamic Enterprise Services use case, Virtual Networks Orchestration creates complete end-to-end IP networks. Network Architects do realize the potential complexity of designing end-to-end IP networks, and in addition, these networks need to be instantiated on the fly, and can be short-living. Therefore, automation is key. However, more than 'just automation' is needed (Figure 5):

1. Multiple different classes of end-points for the service chains are possible: Thick CPE, denoting CPEs with VNF hosting capabilities in addition to routing capabilities controlled through Openflow; Slim CPE, with only routing capabilities. Even Thin CPEs are possible, bridged devices terminated onto a DC-GW or DC-I router with vCPE capabilities.
2. Different Enterprises have over time been connected onto the MSO/CSP-managed WAN in evolving ways;

3. The MSO/CSP-managed WAN underlay may have different options (IP/MPLS, ...), the MSO/CSP datacenter underbuild may be of different stacks (VMWare, Openstack).
4. At least three different VPC environments will exist in the industry (AWS, Azure, Google Cloud), with a lot of other global and local players (Terramark, OVH, ...). Further, each have several options to interconnect at L2 and L3, and are different between public Cloud SPs.
5. In the MSO/CSP-managed datacenter, VNFs will be on the fly instantiated as part of the Service Chain. Also here, different types will be required (Firewalls, Loadbalancers, WAN optimizers, Access Filters, anti-DDoS, ...), each dominated by two-three different vendors.
6. On top of all previous multipliers, we have for each different versions, stacking up as time evolves.

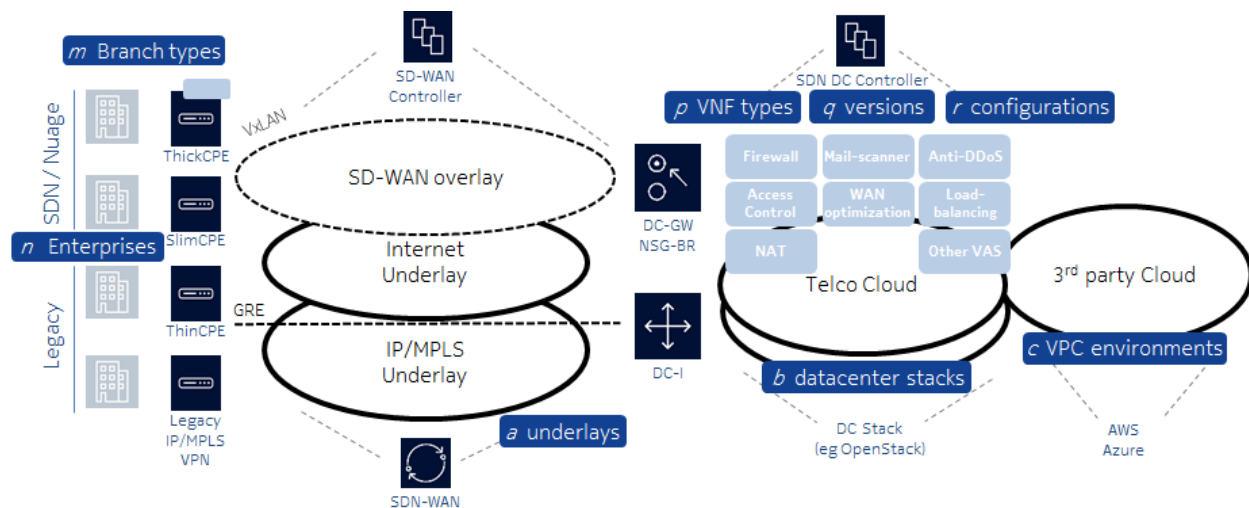


Figure 5 – Enterprise VPN Service Chain in practice: the multipliers

In contrast to these multipliers, enterprises do expect to have connectivity ordered through a simple ordering flow, and delivered at cloud speeds – Network as a Service should not be more difficult than other xaaS services.

The implications are clear: Network Services need to be defined and validated in ‘a day’, for each Enterprise to some extent bespoke, 1000’s of which must be served, with service chains at times living only a few hours.

Therefore, in moving end-to-end IP Network Design from the physical world to the virtual world, all variables governing the problem space change, several by orders of magnitude, and a new approach is required to design, develop, test, sell, deploy Network Services, over and over again, in order words, delivering Network Services as a Factory. The association with the concept of a factory is clear: while delivering new services to the market massively and rapidly, the customer still wants to tune the product to his taste, while the end-product delivered should be reliable and cheap.

Network Service Descriptors

The way the Network Orchestration layer is architected is of crucial importance to cope with new and agile end-customer expectations for their network services, while coping with the multiplier-challenge

introduced in the previous chapter. With Network Service Descriptors, this becomes a reality. Additionally, this orchestration layer needs to facilitate the open design of new network services by telecom vendors, MSOs/CSPs and/or third-party network service designers, specifically keeping in mind that Network Design is a specific technical skill mastered specifically by Network Architects. The Network Orchestration layer needs to have the capabilities to capture these design requirements of the Network Architects and translate them into the right actions of the SDN/NFV IT systems, preferably without too many, if any translation steps from “High Level Design” specifications, to running software.

The Network Service Descriptors (NSDs) are the blueprint of the overlay and overlay-to-underlay network design as defined in Section 4 above and Figure 4, designed and specified by the Network Architect, and which can be interpreted by the workflow engine to orchestrate the correct service fulfillment and later-on the service assurance.

6. Architecture

The architecture of the Networks Orchestration layer (Figure 6) is designed internally around a workflow engine, controlling the Life-Cycle Management of the Network Service. The data-layer of this workflow engine contains a model of the network service, which can be used to monitor the service, and to destroy, upgrade or transition the service.

The model of this network service is called the Network Service Descriptor (NSD). The NSD is a written declaration of the full network service, described in the TOSCA modeling language. Topology and Orchestration Specification for Cloud Applications (TOSCA) is a language that emerged from the Cloud industry and is specified in YAML, and it describes a topology of cloud based web services, their SDN/NFV components, relationships, and the processes that manage them. The language is standardized by the OASIS and adopted by ETSI/NFV.

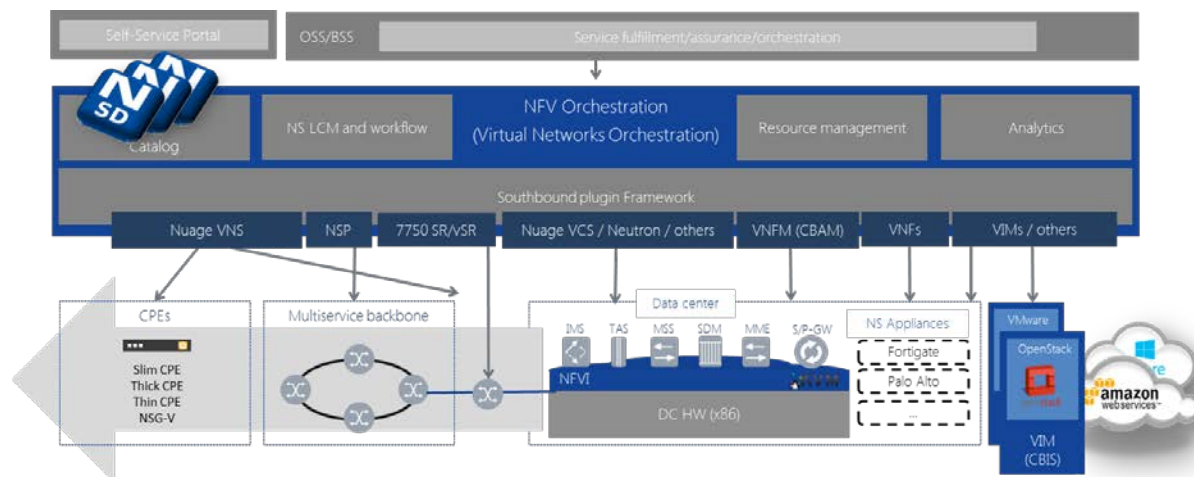


Figure 6 – Architecture of the Virtual Networks Orchestration Layer

Another important component of the architecture is the service catalog. This catalog is a collection of NSDs which are offered, as an example through a self-service portal, as pre-packaged sellable items from which the enterprise customer can select. Through this catalog the provider will be able to offer a large variety of possible network configurations, satisfying the cloud-like service-flexibility requirements of its customers. Additionally, the NSDs themselves will contain a flexible configuration model, which will

allow the enterprise customer to personalize its service the way he needs it, but clearly specified by the operator such that the number of variants/customizations are kept under control.

The Southbound plug-in framework is the other component to keep the multiplier under control. This layer will host various plug-ins towards network and cloud equipment. The plug-ins will translate the TOSCA modeled component to real API calls like RestAPI, Netconf/Yang, CLI, ... As such this plug-in layer creates the important demarcation layer which separates the concerns between the Network service layer and the real implementation specific configuration needed in the underlying systems (CPEs, Routers, Cloud VIMs/VNFs, third-party Cloud services, etc).

Last, this layer contains specific resource management and analytics functionalities to ease Root Cause Analysis (RCA), in case the running Network Services experiences run-time issues.

The orchestration layer has an open Northbound interface, which is controlled both from the Self-service Portal as well as from other OSS/BSS components. For those functional blocks the orchestration layer has the important task to hide the dynamic complexity of the underlying Software Defined components (VNFs, NFV, SDNs) triggered by the multiplier explained above, while keeping more limited and static interactions with the upper-layers.

7. Internals of a Network Service Descriptor

The Network Service Descriptor (NSD) is a deployment template which consists of information used by the NFV Orchestrator (NFVO) or Network Orchestrator for life cycle management of a Network Service. The NSD describes the network service at two different layers of abstraction:

- The logical model (Figure 7): describing standardized components such as Virtual Network Functions (VNFs), Physical Network Functions (PNFs), Virtual Links (VLs), Connection Points (CPs), VNF Forwarding Graphs (VNFFGs), Network Forwarding Paths (NFP);
- The implementation model: describing detailed implementation specific components needed to realize the creation of the logical model. E.g. a logical virtual link could be as simple as an Ethernet LAN, or as complex of an overlay VXLAN network mapped on a IP/MPLS underlay configuration.

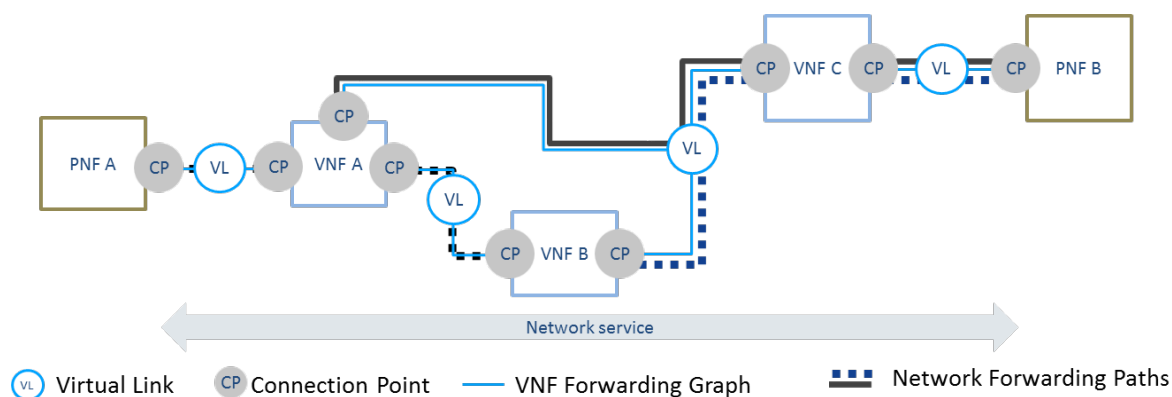


Figure 7 – Logical Model described in the NSD

The components of the logical and implementation models are then written down in YAML notation and packaged together into an archive file, called Cloud Service ARchive (CSAR). This CSAR basically is the packaged network service, which can be onboarded on the service catalog of the orchestration layer.

The YAML description itself contains three major parts: inputs, node_templates, and outputs.

The input part (Figure 8) describes the input variables and represents the flexibility through which the enterprise customer must personalize its networks service: e.g. which IP address ranges to configure, what QoS profile is ordered, whether encryption on the SD-WAN network is needed, ...

```
inputs:
  CPE_uplink_QoS_policy:
    type: string
    description: QoS policy to be applied to the uplink port
    default: Silver
    constraints:
      - valid_values: [ Gold , Silver , Bronze ]

  encryption:
    type: string
    description: Select whether encryption between SD-WAN sites is required
    default: DISABLED
    constraints:
      - valid_values: [ DISABLED , ENABLED ]
```

Figure 8 – Example NSD input parameters

The self-service portal can construct automatically out of the input section of the NSD the input forms for the enterprise customer.

The output part (Figure 9) describes the parameters the NSD provides back to the OSS layer. This can either be information for the enterprise customer (such as its public IP address of its Firewall) or information for the OSS to do Assurance monitoring.

```
outputs:
  fw_mgmt_ip:
    value: { get_attribute: [ fw_vnf, outputs, [ fw_mgmt_ip ] ] }
  fw_public_ip:
    value: { get_attribute: [ fw_vnf, outputs, [ fw_public_ip ] ] }
```

Figure 9 – Example NSD output parameters

The main body of the TOSCA files describes the node_templates. These are a collection of components which describe all the various items which require configuration, together with the properties of this configuration. The important difference with other automatization languages like Ansible and YANG is that complex relations and dependencies can be described between these components. This makes it possible to interpret this description by the workflow engine, in order to make a perfect, well sequenced service fulfillment possible.

```

pub_subnet:
  type: nokia.nuage.nodes.Subnet
  properties:
    name: Public_Internet
    associatedSharedNetworkResourceID: { get_attribute: [ shared_internet , ID ]}
  requirements:
    in_zone:
      type: nokia.nuage.relationships.DefinedInZone
      target: dc_public_zone
    shared_network:
      type: tosca.relationships.DependsOn
      target: shared_internet

```

Figure 10 – Example of a node_template

Given this explanation one could see the parallel between this Network Service Descriptor and the Apps that are installed on smartphones. For the enterprise customer, installing a new network service or upgrading an existing one, will be as simple as instantiating one of the NSDs offered by the ‘webshop’ of the MSO/CSP. The NSD contains all description of the flexibility the service has and at runtime contains all information about that service. This contained modeled information of the network services makes it therefore possible to automatically do advanced life-cycle actions, such as upgrading the service to a new version or transiting the service from the enterprise from one SD-WAN service (eg, a hub and spoke SD-WAN configuration) to a complex other network service (e.g. a full mesh SD-WAN with Firewalls and Intrusion Detection VNFs). By analyzing the dependency graphs, the orchestration layer exactly knows which actions to automatically complete. All of these described features are inherited from the interesting capabilities offered by the TOSCA description language in which the NSDs are written.

Of course, TOSCA has some drawbacks such as limited set of standardized intrinsic functions and relative static topology of the dependency graph. These drawbacks can be overcome by introducing more intelligence in the plug-ins and making a smart NSD decomposition. It is expected that the TOSCA standard will further evolve to onboard these requirements.

When putting all blocks together, one can now fully describe the NSD ‘path’ within the Orchestration layer (Figure 11). The NSD describes all configuration items to be done, together with the relationships between them. The workflow engine in the orchestration layer analyses this dependency graph and uses the plug-ins to do all the configurations required across the various network elements and cloud platforms.

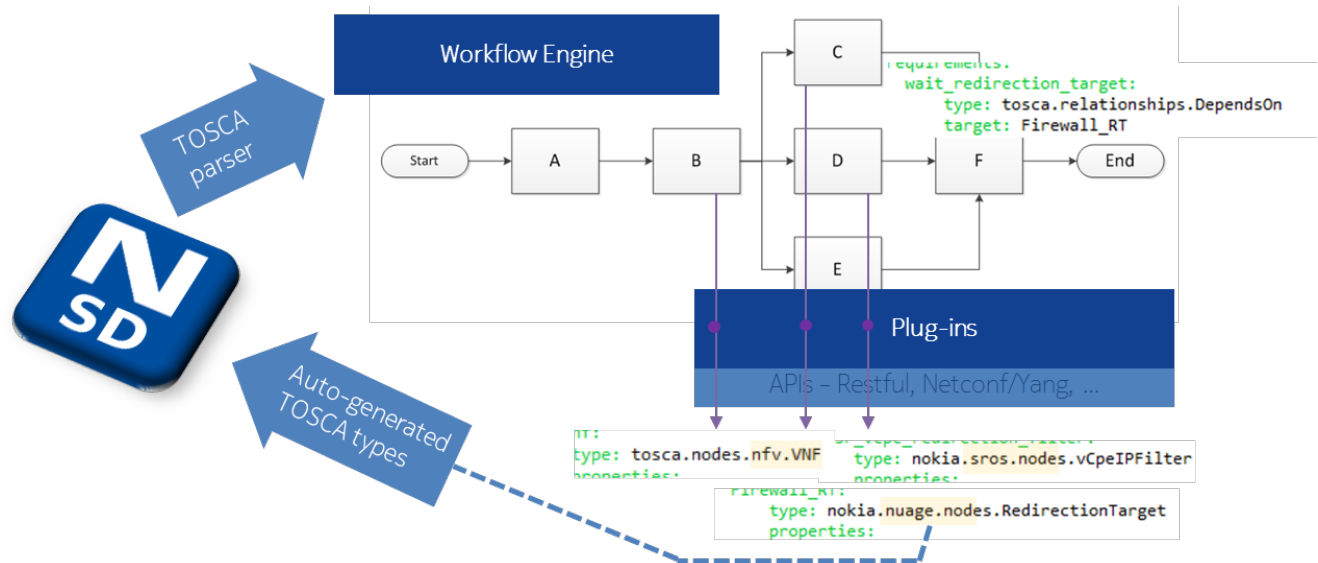


Figure 11 – An NSD into action

The most interesting thing however is that also automation can be realized in the generation of these node_templates and plug-ins. If clearly described interface specifications are available (such as OpenAPI, YANG models, ...) both plug-ins and node-templates can be automatically generated and updated, such that the versioning factor of the earlier described multiplier-challenge is overcome.

‘Best Practice’-based Abstraction

The previous chapter described how one NSD can be written and executed. However, the full end-to-end service, as it is deployed by the Service Orchestration layer, typically consists of a combination of multiple NSDs. As example, an enterprise might initially start from a SD-WAN service with a number of branch-offices. Over a period of time, the enterprise will typically add and delete branch-offices, and might subscribe to new managed services inside the data-center (e.g. firewall, intrusion detection, spam filters,...).

Besides the pure technical skills of writing a NSD, business requirements will highly influence which configuration actions will be grouped together into one ‘sellable item’ – from MSO/CSP to the enterprise customer (adding a ‘branch’ will require the enterprise to pay an additional subscription fee), and which ones will best be decomposed into complementary NSDs (a ‘branch’ hence to become a separate NSD). The following business requirements will highly influence the decomposition of the end-to-end service into NSDs:

- How will the MSO/CSP allow the enterprise-customer to mash-up combinations of service-components in order to personalize its configuration?
- Which migration, upgrade or up-sell scenarios to new functionality would the MSO/CSP like to foresee?

Based on those Best-Practices we currently see already the following collections of separate NSD families (Figure 12), which will together create an end-to-end Network Service.

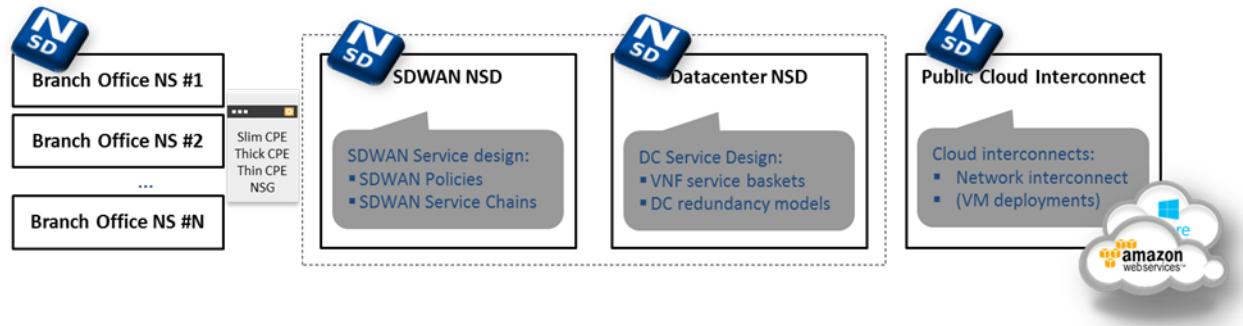


Figure 12 – NSD best-practice decomposition. Example for Dynamic Enterprise Services.

- SD-WAN NSDs will typically contain the model on how an enterprise would like to interconnect its branch offices. This can be a full-mesh scenario, a hub and spoke scenario (where the hub is e.g. the IT headquarter of the company), etc. In addition, it will contain pre-configured configurations of QoS policies, ACLs and Application Aware routing scenarios.
- Branch Office NSDs will typically contain various flavors of branch office interconnectivity. E.g. whether the branch office is connected via single uplink or dual uplink, whether a SDN enabled CPE (so-called ‘SlimCPE’) is used, a ThickCPE - with capabilities to run VNFs distributed between cloud and branch office, or whether a so-called ThinCPE configuration is deployed, which interconnects a non-SDN legacy router to the Overlay SD-WAN.
- Datacenter NSDs will typically contain the needed interconnect and service chaining between the SD-WAN and the Managed Services (VNFs) in the telco-cloud. These NSDs might also contain knowledge of Active-standby configurations when service chains are defined across multiple redundant datacenters.
- Other independent NSDs such as interconnects to various third-party public cloud environments (AWS, Azure,...), interconnects to fixed/wireless infrastructures (like GiLAN, vResidential GWs,...), or complementary service packs (like advanced application monitoring through Application Aware Routing,...)

This best-practice decomposition provides high flexibility to the enterprise customer for defining its required service, and will create a higher-level abstraction towards the Service Orchestration layer, such that the details of the Network definitions are handled in the Network Orchestration layers while the Service Orchestration layer instantiates actions on billable items requested by the end-customer.

Intent-based testing

Testing represents an important aspect of the NSD development process, meant to ensure that a deployed NS performs per the business requirements. This aspect takes center-stage when the development process adheres to modern DevOps or Continuous Delivery methodologies, which is the case of the Network Services as a Factory concept.

Considering the short ‘one-day’ cycles imposed to NS development and validation, it becomes imperative to automatically execute the defined tests every time a change is committed to the NSD implementation. Given the declarative nature of the TOSCA domain-specific language, the most appropriate test methodology is also a declarative one, in which the NS state and behavior is validated against a selected set of inputs.

The test methodology defining the ‘Network Services as a Factory’ approach embraces the notion of behavioral driven testing which checks that the purpose the NSD was designed to fulfill, is successfully achieved. To bridge the gap between business requirements and test use case implementation while retaining full automation capabilities, the authors have selected the Cucumber framework due to its capability of supporting the use of natural language in test definitions.

The domain-specific natural language subset implemented by the chosen test methodology covers the complete creation and deployment lifecycle of a network service according to ETSI standards and is open to further extensions. The below code snippet (Figure 13) shows an example of testing the Internet connectivity established via a data center service chain connected to a full-mesh SD-WAN VPN. An important note is that the best practice related to scenario autonomy in Cucumber was deliberately broken in order to allow (a) the required flexibility in defining large-scale test use cases and (b) avoid the unnecessary repetition of expensive provisioning scenarios.

```
Feature: Connect a DC service chain to an existing Nuage VPN
  In order to provide E2E connectivity
  As a Network Service Designer
  I want to define an DC service chain
  That connects to an existing Nuage VPN
  And provides Internet access to the sites connected by it

Background:
  When I log in to CBND

Scenario: Model the existing VPN as a mesh
  When NSD "../AT_csar/Site-to-site_Mesh_connectivity.zip" is uploaded to CBND with alias "NSD Mesh"
  And NS instance "AT Mesh VPN" of alias "NSD Mesh" is created with parameters:
    | enterprise_name | Enterprise_Test |
    | encryption     | DISABLED       |
    | domain_name     | AT_SD-WAN_Domain |
  And NS "AT Mesh VPN" is deployed within 15 min
  Then NS "AT Mesh VPN" state should be DEPLOYED

...

Scenario: Site-2-site and site-2-outside traffic are enabled
  Given NS "AT Branch 1" state is DEPLOYED
  And NS "AT Branch 2" state is DEPLOYED
  And NS "AT DC SC" state is DEPLOYED
  Then below VMs can pass traffic to each other
    | Name      | Interface |
    | g1lanvm4  | eth2      |
    | g1lanvm5  | eth2      |
  And VM "g1lanvm4" can access URL "www.google.com"
  And VM "g1lanvm5" can access URL "www.google.com"
```

Figure 13 – Intent-based testing – a code snippet.

This solution allows the NS developer or designer to quickly define test use cases that are versioned using a source code management (SCM) system and automatically executed by a continuous integration (CI) server against a specified deployment environment, as shown in Figure 14. The presented Network Service as a Factory implementation relies on the NodeJS implementation of Cucumber - cucumber.js - which runs inside a Jenkins CI server.

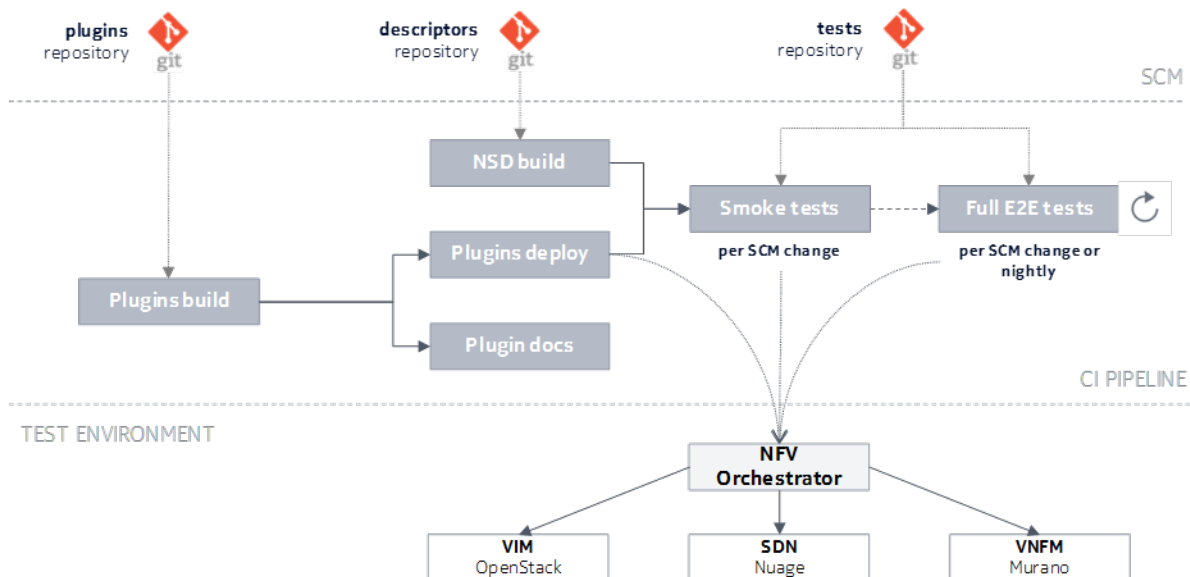


Figure 14 – NSD testing as part of the Continuous Deployment pipeline.

The advantages brought in by the presented NSD testing methodology are multiple. Some of them stem from its alignment to the Continuous Delivery paradigm while others explicitly address the specifics of the NSD design and development process.

- The fully automated nature of the process ensures a continuous, fast and objective assessment of the NSD quality. The network service designer can stay focused on describing the higher-level NSD structure and its intended behavior, knowing that the CI platform takes over the validation efforts. Given that NS designers usually align with the network engineer profile more than with the software developer one, this represents a significant enabler for network engineering teams in transitioning to the Network Service as a Factory approach.
- Every change brought to one element in the toolchain, be it an NFV orchestrator plugin, an NSD or a test definition, will be automatically deployed and validated on the target environment in the context of an end-to-end solution. This keeps quality high, reduces validation times and removes the risk of regressions and their associated high costs. Given NFVO plugin development usually requires technologies different than NSD development, the presented approach bridges the gap between the different teams handling these activities by bringing them under the same validation and acceptance umbrella.
- The use of natural language in test definitions streamlines the customer acceptance process since the test results can be interpreted by all project stakeholders as opposed to just the technical teams. Furthermore, since the platform can be directed at different environments, it is perfectly possible to use it to deploy and validate NSDs on a new customer environment with positive impact on project start-up costs and delivery timelines.
- Besides the abstraction towards the OSS layer brought by the NSD concept as described in the previous chapters, behavioral-driven test definition facilitates the integration of new network services with the OSS by giving OSS teams insight into how the service should be managed and what are its requirements without exposing the complexity of its implementation.

Conclusion

While Network Orchestration and Application Orchestration share technologies coming from the domains of SDN and NFV, the two approaches serve different objectives. The paper has specifically addressed Network Orchestration, aimed at launching new revenue generating Business Connectivity services rapidly, with Value Added services running in private or public cloud data centers chained-in into the Network Service.

The rapid market traction of SD-WAN overlay IP networking, enabled by product suites such as Nuage, fuels the requirements for Network Orchestration: overlay to underlay connectivity, connectivity into Public Cloud, Internet Break-out, flexible onboarding of virtualized Value Added services such as firewalls, wan optimizers, email filters, and others, all easily customizable for that one specific enterprise customer of the MSO/CSP. However, the combination of different Value Added Services vendors and types, SDN domain controllers, cloud stacks, software versions, implies a multiplier in the number of combinations that are possible to deliver just that one single end-to-end service chain over and over again. A different approach is needed, bridging the need for instantaneous Network Service delivery, and network diversity, complexity, and constant evolution.

A combination of technologies, approaches and methodologies, building on abstraction, automation, and decomposition, are brought together, in an approach called Virtual Networks Orchestration. It allows to deliver Network Services and NSDs rapidly, massively, and reliably – Network Services as a Factory.

TOSCA as a YAML-based Domain Specific Language allows to describe Network Services in Network Service Descriptors. More specifically, a best-practice based Network Service decomposition approach allows to abstract out the different areas of complexity in the network, so that, through isolation, a change in the service description in one area does not require to change the full network service. ‘Best-practice’ should be based on the interaction of the market facing units of MSOs/CSPs with its customers (enterprises as examples).

Intent-based testing is explored as a further step towards automation. To bridge the gap between business requirements and test use case implementation while retaining full automation capabilities, the Cucumber framework was highlighted, and its capability of supporting the use of natural language in test definitions. The advantages brought in by the presented NSD testing methodology include the continuous, fast and objective assessment of the NSD quality; the validation of a change to any of the elements in the toolchain, NSDs and plug-ins; the streamlining of the customer acceptance process; and facilitating the integration of new network services with the OSS.

Abbreviations

CI	continuous integration
CLI	command line interface
CSP	communications service provider
DES	dynamic enterprise services
DSL	domain-specific language
EMS	element management system
MANO	management and orchestration
MEF	metro ethernet forum
MSO	multiple systems operators
NFV	network function virtualization
NFVO	NFVorchestrator
NS	network service
NSD	network service descriptor
OASIS	organization for the advancement of structured information standards
ONF	open networking foundation
OSS	operations support systems
SDN	software-defined networking
SD-WAN	software-defined wide-area networking
Tosca	topology and orchestration specification for cloud applications
vePC	virtualized enhanced packet core
vIMS	virtualized IMS
VNF	virtual network function
VNFM	virtual network function manager
VoLTE	voice over LTE
VPC	virtual private cloud
VPN	virtual private networking
xaaS	infrastructure, platform, software as a service

Bibliography & References

http://innovation.verizon.com/content/dam/vic/PDF/Verizon_SDN-NFV_Reference_Architecture.pdf

http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf

http://www.etsi.org/deliver/etsi_gs/NFV-IFA/001_099/014/02.01.01_60/gs_NFV-IFA014v020101p.pdf

<https://www.oasis-open.org/committees/tosca/>

<http://docs.oasis-open.org/tosca/tosca-nfv/v1.0/tosca-nfv-v1.0.pdf>

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/os/TOSCA-Simple-Profile-YAML-v1.0-os.pdf>

<http://yaml.org/>

<http://www.nuagenetworks.net/>

Getting Grounded with DevOps, 2015, <https://www.hpe.com/h20195/V2/getpdf.aspx/4AA4-3696ENW.pdf>

Continuous Delivery: Automating the Deployment Pipeline, 2016,
https://www.microfocus.com/media/white-paper/continuous_delivery_automating_the_deployment_pipeline_wp.pdf

Hallenberg N. and Carlsen P.L., Declarative automated test, Proceedings of the 7th International Workshop on Automation of Software Test, pp. 96-102, 2012

Crowther M., An Introduction to Behavioral Driven Testing, 2009,
http://www.cyreath.co.uk/papers/Cyreath_An_Introduction_to_BDT.pdf

Cucumber reference, 2017, <https://cucumber.io/docs/reference>
Cucumber Best Practices, 2015, <https://github.com/strongqa/howitzer/wiki/Cucumber-Best-Practices>

Real-World Deployment of a Virtual Cable Hub

A Technical Paper prepared for SCTE•ISBE by

Asaf Matatyaou

Vice President, Solutions and Product Management, Cable Edge Business
Harmonic Inc.

4300 North First Street

San Jose, CA 95134

408-490-6834

asaf.matatyaou@harmonicinc.com

1. Introduction

The promise and potential of virtualizing a cable hub has been discussed over the past few years. While the opportunities are limitless, there must be a starting point and manageable steps for operators in migrating to virtualization.

This paper will focus on translating the theory of virtualization into real-world experiences and recommendations based on deployment. Different aspects of a cable hub will be evaluated for virtualization, including using cloud for telemetry and monitoring, configuration and orchestration, operations and back-office elements, as well as a software-based CCAP solution based on a virtual CMTS (vCMTS) implementation. Virtualization implementation approaches will be compared, such as containerization and virtual machines (VM) for different applications and services. In addition to operationalization considerations, support for maintaining legacy services, such as traditional broadcast video, VOD and out-of-band, as well as existing and future IP services will be considered in the transition to a virtual cable hub.

2. The Promise of Virtualization

The promise of virtualization is not new. In fact, it's delivered substantial benefits to many industries, such as data centers, wireless access and Internet of Things (IoT), and has delivered on this promise for decades. The definition of virtualization has broadened over these decades and practical use cases and real-world deployments continue to increase every year.

Virtualization promises to enable change at a pace which meets or exceeds customer demand in the most effective manner. Change can be defined in many ways, such as new services, additional security, elastic storage, more efficient infrastructure and in the case of the cable broadband industry it's all about speed. Virtualization delivers on this promise by separating applications or software from hardware. Key benefits are scalability, sustainability and elastic deployment with the quickness and agility needed to increase business efficiencies and productivity. This separation of software and hardware is the key to quickness and agility of change, as the software can be changed while running on deployed hardware. Upgrading software is not only quicker than replacing hardware, but requires less operational expenditure (OpEx), such as onsite labor and increase in power consumption requirements.

With all the benefits mentioned, why is the cable hub lagging behind in embracing virtualization? In the technical paper, *Transforming the HFC Access Network with a Software-based CCAP*, prepared for SCTE in October, 2015, the considerations for virtualizing the CMTS capabilities in a cable hub were described. Notably, key enablers identified are:

1. CableLabs Remote PHY standard: A standards-based approach separates the physical layer (PHY) from the CMTS Core, which "contains the DOCSIS MAC and the upper layer DOCSIS protocols. This includes all signaling functions, downstream and upstream bandwidth scheduling, and DOCSIS framing."¹
2. Merchant silicon for a full spectrum DOCSIS 3.1 PHY layer: Remote PHY hardware will deliver on a long lifespan, with up to 10 Gbps² of downstream and up to 2 Gbps of upstream bandwidth.

¹ Remote PHY Specification, CM-SP-R-PHY-I07-170524, §5.1, pg. 25

² MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I11-1705, §5.1, pg. 51

3. Performance from Commercial Off-The-Shelf (COTS) Intel-based x86 servers: The performance delivered by today's Intel CPUs exceeds minimum product requirements to serve customers' needs and has the history and potential of continued annual performance gains.

These three key enablers are recent technological advances. Coupled with existing virtualization technologies, these key enablers and associated benefits are now ripe for cable edge deployments.

3. Today's Cable Hub

Today's cable hub is a combination of many legacy and purpose-built hardware-based solutions, which over the past 20 years (DOCSIS turned 20 years in March 2017³) has delivered on data, voice and legacy (MPEG-based) video services. To deliver these fundamental services (other services, such as home security and commercial services, are typically running over the DOCSIS data service as the fundamental service), active equipment found in a typical cable hub includes:

1. Cable Modem Termination Systems (CMTS)
2. EdgeQAMs
3. Routers and switches
4. Out-of-band modulators and return path demodulators
5. FCC and LTE leakage signal generation
6. Provisioning servers (such as DHCP and TFTP)
7. Configuration tools
8. Monitoring tools

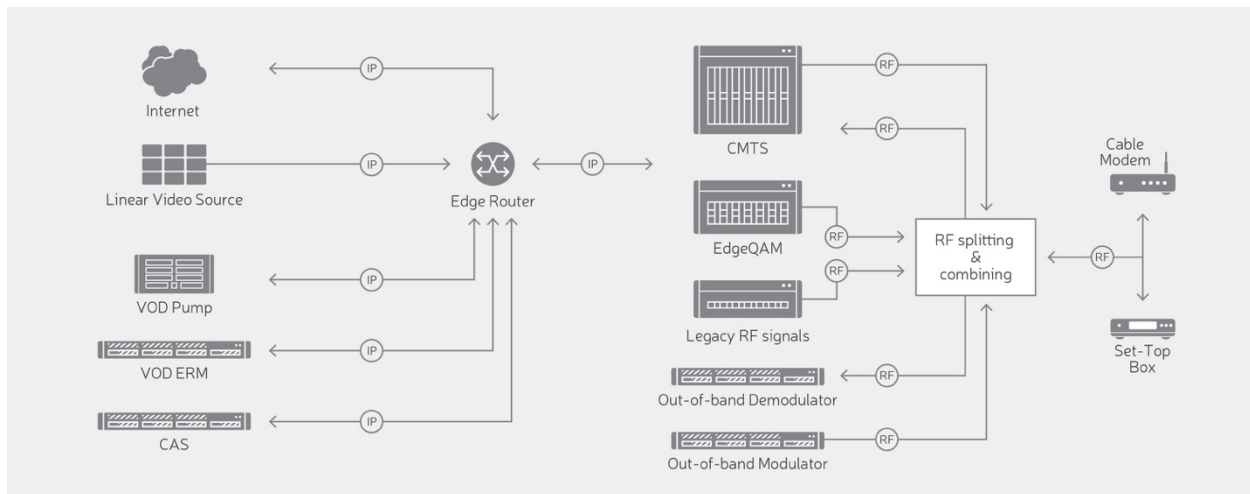


Figure 1 - Today's Cable Hub

Most of the equipment listed above can be categorized as either custom hardware (e.g. CMTS & EdgeQAM) or custom software (e.g. configuration and management tools) and most are purpose-built for the cable edge (with the exception of routers, switches and some of the provisioning server components). The challenges in continuing with a customized cable edge and hardware-based approach as compared to

³ DOCSIS 1.0 specifications Version I01 were released March 26, 1997.

a virtualization approach, as described in *Transforming the HFC Access Network with a Software-based CCAP*, are shown in Table 1 - Virtualization Opportunities in a Cable Hub.

Table 1 - Virtualization Opportunities in a Cable Hub

Hardware-based Challenges	Virtualization Opportunities
Unsustainable space, power and cooling	Dramatically reduced space, power and cooling footprint
Proprietary and custom hardware, with long development cycles	More frequent and shorter development cycles
Complex and infrequent software upgradeability results in slow feature introduction	Agile development with frequent software upgrades delivers fast feature velocity
Limited capacity, scalability and flexibility	Sustainable capacity growth, elastic scalability and increased flexibility
Larger failure domain	Smaller failure domains
High operational expenditure (OpEx)	Improved Total Cost of Ownership (TCO), including reduced operational (OpEx) and capital expenditure (CapEx)

Today's cable hub is most challenged by the equipment which doesn't sustainably scale physically or operationally with service or capacity growth. Sustainable growth will happen when the performance and scale meets or exceeds consumer consumption demands, either driven by actual usage or competition from other access providers, such as Fiber To The Home (FTTH). In other words, virtualizing the cable hub is a high-tech solution for a low-tech problem, specifically, running out of facility space, as well as ever-increasing and recurring electricity expenses.

Another consideration for the ever-evolving cable hub is Remote PHY. While this standard technology enables virtualization in many ways, it also demands many more nodes (or service groups) being deployed, with some estimates exceeding a ten-fold increase in nodes. Today's deployment is manual and labor intensive and doesn't scale operationally, especially when considering the desired deployment rate of Remote PHY nodes. Today's cable hub tools include tried-and-true and very familiar tools such as Command Line Interface (CLI), Simple Network Management Protocol (SNMP) and Internet Protocol Detail Record (IPDR). In many cases, individual cable operators have customized home-grown tools, interfacing to hardware-based equipment over standard protocols (SNMP, IPDR) or proprietary CLI. These configuration and management interfaces are also archaic, slow and manual in many cases.

The equipment and tools which benefit the most when transitioning to a sustainable growth deployment model will be highlighted as components in the cable hub which have been virtualized in a real-world deployment of a virtual cable hub.

4. Virtualization in All Shapes and Sizes

While virtualization is a relatively recent hot topic in the cable industry, it has been present for decades. Since the 1960's the definition of virtualization has broadened over the years. This is good. The power of virtualizing continues to expand, and its reach has delivered benefits to many industries and different types of implementation.

Wikipedia defines virtualization as “the act of creating a virtual (rather than actual) version of something, including virtual computer hardware platforms, storage devices, and computer network resources.” Further, Wikipedia lists many different types of virtualization. To name a few: desktop, operating system-level, application, workspace, service, memory, storage, file system, data, database and network.⁴

When looking at how virtualization applies to the cable edge space, particularly in the cable hub, equipment and tools that are custom hardware-based or benefit from orchestration (for the purpose of eliminating labor intensive human interaction to operate and manage services) will be considered. To be specific, there is a disruptive change coming to cable hubs which is described in Table 2 - Virtualization Changes in the Cable Hub.

Table 2 - Virtualization Changes in the Cable Hub

Old Way	New Way
Application runs on custom hardware	Application runs on COTS hardware
Application is implemented partially or fully on hardware components (e.g. ASIC, FPGA)	Application is implemented in software
Replaced every three to five years	Long lifespan
Application is upgraded infrequently	Application is upgraded regularly
Equipment location is limited	Equipment location is varied
Services turned on are labor intensive	Services are turned on automatically with orchestration
Service monitoring is labor intensive	Service monitoring (telemetry) is performed by software analytics
Service events are limited to local hardware storage	Service events are stored in the cloud

To summarize, virtualization in a virtual cable hub means running virtual applications (for example, DOCSIS, video, OOB) on COTS x86 platforms, which can be located in cable hubs, more centrally in fewer locations (such as data centers) or even in smaller form factors in distributed locations (such as street cabinets).

In this virtual cable hub, telemetry and logging tools perform analysis of streaming data predicting potential impactful events or visualizing historical events in a holistic fashion. In many ways, the idiom “a picture is worth a thousand words” becomes reality. Instead of hours of labor intensive human scrutiny, which may result in a partial analysis and slower resolution of a field issue, a glance of dashboard provides an instantly clear picture showcasing visualized streaming data. The benefit is speed and accuracy in determining a more complete resolution.

⁴ <https://en.wikipedia.org/wiki/Virtualization>

5. Virtualization Considerations

Defining cable edge virtualization as running cable-specific virtual applications in software is a first step in the right direction. This identifies what elements or equipment will be virtualized. However, other virtualization concepts are important to consider, specifically how the software is virtualized. Common concepts include containerization, virtual machines and bare metal approaches. Some of these methods are mutually exclusive, while others are complementary or even dependent on each other. While the methods described are common in different industries, there isn't always definition consensus for each concept. This is due to the wide variety and growing usage of virtualization for different objectives. For this paper, the following definitions will be used for these concepts and are based on the actual usage of virtualizing a cable hub.

- Bare metal: an application is executing on the native operating system (OS), in comparison to executing on a virtual machine or a virtual OS layer. In other words, the application can access “the metal” directly or via a native OS.
- Virtual machine: System virtual machines are capable of virtualizing a full set of hardware resources, including a processor (or processors), memory and storage resources and peripheral devices. A virtual machine monitor (VMM, also called hypervisor) is the piece of software that provides the abstraction of a virtual machine.⁵
- Virtual appliance: “a pre-integrated, self-contained system that is made by combining a software application (e.g., server software) with just enough operating system for it to run optimally on industry standard hardware or a virtual machine (e.g., VMWare, VirtualBox, Xen HVM, KVM).”⁶
- Containerization: “applications can be broken up into manageable, functional components, packaged individually with all of their dependencies, and deployed on irregular architecture easily.”⁷
- Docker: a set of tools to package and deploy containers, which can specify container constraints and access permissions. Additionally, Docker sets up and deploys the container in Linux.
- Cloud native: “cloud native computing uses ... software ... to be containerized, dynamically orchestrated, and microservices oriented.”⁸
- Sandboxing: an isolated computing environment for running applications.
- Single-tenant: a single instance of a single application type running a single physical hardware platform.
- Multi-tenant: multiple instances of one or more application types running on one or more physical hardware platforms.
- Kubernetes: “an open-source system for automating deployment, scaling, and management of containerized applications.”⁹
- Kubernetes-native application: an application which is aware that it's being deployed or managed by Kubernetes.

⁵ https://en.wikipedia.org/wiki/Popek_and_Goldberg_virtualization_requirements

⁶ <https://www.turnkeylinux.org/virtual-appliance>

⁷ <https://www.digitalocean.com/community/tutorials/the-docker-ecosystem-an-overview-of-containerization>

⁸ <https://www.cncf.io/about/faq/>

⁹ <https://kubernetes.io>

- Microservice: “refers to an architectural approach that independent teams use to prioritize the continuous delivery of single-purpose services. The microservices model is the opposite of traditional monolithic software which consists of tightly integrated modules that ship infrequently and have to scale as a single unit.”¹⁰

Clearly, there are many decisions to be made and tradeoffs to consider when evaluating which concept to use and when. The great news is that once custom hardware has been retired and applications are implemented in software on off-the-shelf hardware, there is no decision which can’t be undone and no outcome which can’t be upgraded or improved.

5.1. Virtualization Criteria in a Virtual Cable Hub

The key criteria for deciding between the various virtualization approaches are:

1. Time to market (TTM): the time criticality to deliver a minimum feature set of a virtualized set of applications. Over time, the feature set will grow with periodic software upgrades to the virtual applications.
2. Performance: the minimum application processing required to deliver a cost-effective footprint of COTS x86 servers. Over time, the performance will improve with periodic software upgrades to the virtual applications and the performance per rack unit will increase or the number of rack units will diminish to deliver the same performance.
3. Scale of deployment: the minimum quantity of consumers supported by a virtual cable hub. Over time, the scale of deployment per rack unit will increase or the number of rack units will diminish to support the same quantity of consumers.
4. Application flexibility: the minimum set of application types and elasticity to execute different instances of different applications on a single physical server. Over time, the ability to execute many and different instances of different application types on a variable set of physical servers will be possible with periodic software upgrades.

The common theme for all four criteria (TTM, performance, scale and application flexibility) is that there is a minimum or “good enough” starting point and that future software upgrades improve the virtual cable hub capabilities in different dimensions.

5.2. Crawl, Walk, Run, Fly Approach to a Virtual Cable Hub

The most common assumption when discussing virtualizing a cable hub is the notion that using virtual machines (VMs) is required. It’s a possibility, but not a necessity. Let’s compare bare metal, containerization with Docker and virtual machine approaches to a virtual cable hub, taking into account the four key virtualization criteria.

Two possible approaches include:

1. Single-tenant application running on bare metal or a virtual machine
2. Multi-tenant containerized application instances packaged and deployed by Docker on bare metal or virtual machine

¹⁰ <https://pivotal.io/microservices>

5.2.1. Hit the ground running

The most complex and performance-intensive application in a virtual cable hub is the virtual CMTS (vCMTS) component. As defined by the CableLabs Remote PHY standard, the DOCSIS physical layer is separated from all the upper layers, via the standards-based protocols DEPI, UEPI and GCP.¹¹ When referring to the vCMTS component in the context of virtualization, the CMTS Core functionality (as defined in the CableLabs Remote PHY standards) is implemented as a virtual application.

An incremental approach when virtualizing the cable hub is a crawl, walk, run, fly approach¹². With a virtualized approach rooted in software, maturing from crawling to flying is entirely performed by software upgrades along the way, and the penalty of tripping over oneself is limited to a software release iteration with no need to replace hardware.

One metric of vCMTS performance is the packet processing rate, which in turn results in the bandwidth or throughput capabilities of a vCMTS. Meeting real-world performance requirements dictates that a single-tenant vCMTS application running on bare metal has the quickest TTM, while still delivering on many virtualization benefits. Effectively, this is a virtual appliance approach, which can grow and scale accordingly by adding more servers, each running a single instance of a vCMTS virtual application. Concluding quickly on whether the simplest virtualization approach meets the performance and scale requirements to deliver on the stated benefits of virtualization while meeting or exceeding functional requirements of traditional hardware-based CMTS approaches is vital.

In practice, running a single instance of a vCMTS application on x86 COTS servers delivers tens of Gbps of packet processing performance per x86 server rack unit to dozens of service groups, while reducing the space, power and cooling footprint by up to 90% relative to existing hardware-based integrated CMTSes.¹³

The result of this first phase of a virtual cable hub is already delivering substantial capital and operating expenditure benefits to cable operators, and it can be debated as to whether the benefits already justify stating that this approach is “running.”

5.2.2. How fast do you want to fly?

Continuing with the crawl, walk, run, fly analogy, let’s shift gears and see how fast a virtual cable hub can fly. The next set of critical benefits to a cable operator when looking at a cable hub are still covered by the four criteria (TTM, performance, scale and application flexibility):

1. TTM: the speed to turn on consumer services to a single set of consumers
2. Performance: improving uptime by limiting the scope of service outages
3. Scale: the quantity of consumers which can be supported in a given footprint
4. Application flexibility: the set of virtual applications which are required for a virtual cable hub

The approach of multi-tenant containerized application instances packaged and deployed by Docker on bare metal provides many of these benefits. Let’s identify which virtualization concept delivers on these benefits.

¹¹ Modular Headend Architecture v2 Technical Report, CM-TR-MHAv2-V01-150615

¹² Paraphrasing Martin Luther King Jr from his April 26, 1967 speech in Cleveland

¹³ Based on Harmonic Inc’s analysis of real-world case studies

Kubernetes provides the orchestration to deploy pods (a group of one or more containers), in conjunction with Docker as the tool to package the pods. “A pod models an application-specific ‘logical host’ - it contains one or more application containers which are relatively tightly coupled — in a pre-container world, they would have executed on the same physical or virtual machine. While Kubernetes supports more container runtimes than just Docker, Docker is the most commonly known runtime, and it helps to describe pods in Docker terms. Pods serve as unit of deployment, horizontal scaling, and replication.”¹⁴

The speed to turn on consumer services is increased significantly when using a combination of Kubernetes and a Docker approach to pods deployment. This speed increase is gained when shifting from a human interaction to configure manually each unit of consumer deployment to an orchestrated and automated process. Additionally, reliability increases as the more error-prone manual method of configuring new consumer services is reduced or eliminated.

Improving uptime is a function of reducing failure domain size, which is a major benefit of a containerized approach. Determining the failure domain size of each pod or container provides the knob when determining the tradeoff between potentially more shared resources in a single pod (increasing CPU utilization percentage) and limiting the pod to a single consumer group, such as a service group. Software failures will happen and uptime is improved as failure domains are reduced. By any measure, the performance metrics of uptime percentage and the number of service calls received over a period of time is tracked by cable operators.

Containerization also has the benefit of horizontal scaling (scale-out) in comparison to vertical scaling (scale-up). Vertical scaling increases a single application instance’s set of specifications, such as bandwidth or subscriber count, by increasing the number of x86 CPU cores, storage or NIC speed. A virtual cable hub uses vertical scaling to grow capacity when a single application instance can do more. Horizontal scaling increases a virtual cable hub’s scale with a virtual set of application instances (pods), each with a specific and purposefully bounded scale specification. When the limitation of a single pod is reached, more pods are deployed.

¹⁴ <https://kubernetes.io/docs/concepts/workloads/pods/pod/>



Figure 2 - Vertical Scaling




Figure 3 - Horizontal Scaling

Figure 2 and Figure 3 visualize the difference in approaches between vertical and horizontal scaling. The benefits of horizontal scaling through containerization becomes apparent: smaller failure domain, simpler to develop, simpler to test and more cost-effective COTS x86 server hardware requirements. Development and testing simplicity is explained when considering smaller data sets and test case parameters, such as developing to and testing to a scale of 200 as compared to 200,000 subscribers.

Lastly, multi-tenant in the context of a virtual cable hub, is having multiple instances of a single or multiple application types. Each pod of containers may be a different application type, and potential applications include CMTS Core, Out-of-band Core, Video Core, proactive network maintenance (PNM), orchestration and telemetry. Over time, different and new applications will be considered for virtualization. The benefit is leveraging the same COTS x86 servers for different application types, which delivers improved sparing (for hardware failures and replacement) and economy of scale benefits when procuring many of the same platform. When different applications are able to execute on the same COTS x86 server, dynamic and elastic utilization of CPU resources can be determined during run-time, based on usage metrics and analytics to shift resources from lightly loaded applications to applications running hot.

5.2.3. *Microservices*

The combination of the approaches described and benefits can be also referred to as microservices, which was previously defined as a “continuous delivery of single-purpose services.”¹⁵ It’s evident that shifting from hardware-based to software-based implementations, as well as a shift in software development methodology from waterfall to agile, results in dramatically improved feature velocity. Shifting from traditional monolithic software to a microservices approach is another lever which improves feature velocity and delivers more frequent and higher quality software upgrades.

¹⁵ <https://pivotal.io/microservices>

Monolithic software is released as a single unit, and due to size and complexity, has more software defects as well as a longer regression test cycle time. Microservices are lighter-weight modular units of software, which can be defined with a limited set of capabilities with published interfaces. With monolithic software, it's all or nothing. With microservices, each upgraded service software can be tested with lighter-weight automation and shorter regression test cycle time.

A notable benefit is improved software upgradeability, in terms of total time and service outage potential. Microservices, by definition, have smaller code size than a single monolithic software image. This results in less time to download software images, less time to upgrade software to a limited set of one or more microservices and reducing the minimal amount of code changes when correcting defective software. When performing hitless or in-service-software-upgrades (ISSU), redundancy or protection mechanisms are typically employed to activate a protected unit of software while the originally active software unit is upgraded without impacting service. However, during this time, overall system protection is diminished. With microservices, the total amount of unprotected time during ISSU is reduced, improving overall high-availability of the virtual cable hub.

6. The Virtual Cable Hub

Let's revisit the cable hub previously shown in Figure 1. The CMTS and EdgeQAM, which are historically implemented in big-iron hardware-based chassis, as well as legacy RF signal and out-of-band signal generation performed in the cable hub are shown as example applications which are virtualized as multi-tenant containerized applications turned on with automated orchestration on COTS x86 servers (shown in blue in Figure 4). In this future vision of a virtual cable hub, all data, video and voice services for residential and business customers are IP-based.

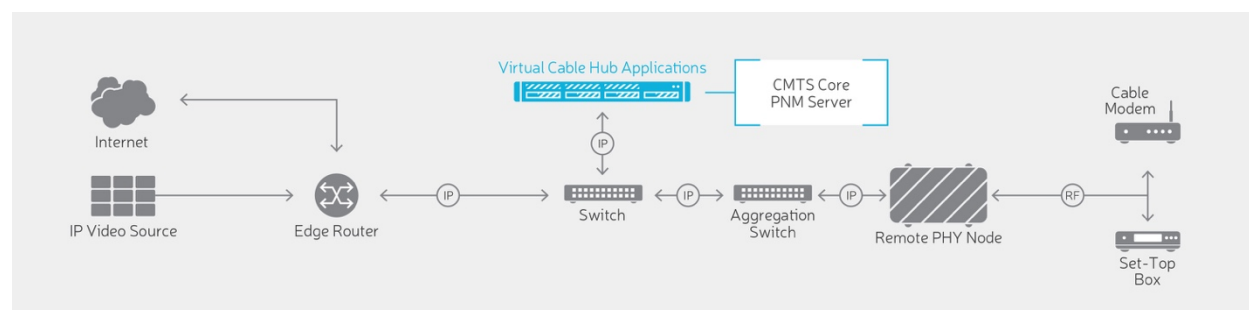


Figure 4 - Virtual Cable Hub

6.1. Moving Forward While Maintaining Legacy Services

The virtual cable hub depicted in Figure 4 is a bit too idealistic as a starting point, with the noted legacy services of linear/broadcast MPEG video, video-on-demand (VOD), switched digital video (SDV), as well as out-of-band signaling for set-top box (STB) and legacy RF signal generation such as HMS and FM. With millions of legacy consumer devices at subscribers' homes, these legacy services will be reduced over many years, eventually being replaced by pure IP-based services. Until that moment, the virtual cable hub, deployed in a DAA architecture such as Remote PHY, will need to support IP transport of the RF signals at the virtual cable hub, with the IP transport converted and modulated to RF signals at the Remote PHY node.

The applications to encapsulate the IP transport of the legacy RF signals don't require much processing and are scalable. With a virtual cable hub, instead of requiring single-purpose custom hardware-based solutions for each and every specialized legacy function, a virtual application can be deployed on available server resources in a cluster of servers to deliver the necessary capabilities. Figure 5 shows a virtual cable hub with legacy services supported.

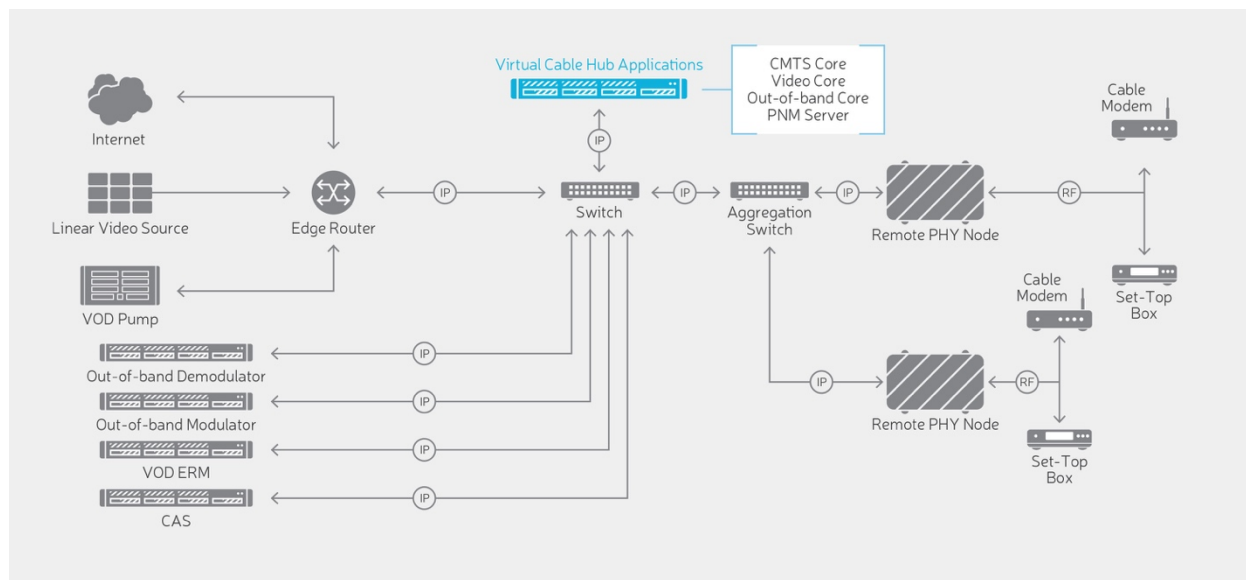


Figure 5 - Virtual Cable Hub with Legacy Services

6.2. The Cloud for the Virtual Cable Hub

The virtual cable hub described in this paper can be considered to be “cloud native” as it meets the conditions defined: containerized, dynamically orchestrated and microservices oriented. However, “cloud native” can also refer to applications executing in the cloud (private/on-premises or public). An example of a public cloud service is Amazon Web Services (AWS), which claims “on-demand delivery of compute power, database storage, applications, and other IT resources through a cloud services platform.”¹⁶

The virtual cable hub is advancing quickly in real-world deployments and delivering a full set of services, meeting demanding performance requirements for speed, latency and jitter. With vital residential and commercial services being delivered, it's important to intelligently weigh the tradeoffs if services might be hindered in a measurable way. Even in the early days of a virtual cable hub, some configuration and monitoring may be serviced in the cloud, public or private. In particular, logging and monitoring is well suited for the cloud, with on-demand increase in database storage and easy accessibility to telemetry and analytics.

Consider the limitations of a traditional hardware-based solution, with finite on-board storage for a small service area. Once the storage is exhausted, the older data is replaced with fresh data, which leads to less samples of data being stored or a short time span of data analytics or logging. Additionally, correlating data is a labor-intensive activity, with a person logging into each platform remotely. In a virtual cable

¹⁶ <https://aws.amazon.com/what-is-cloud-computing/>

hub, a continuous stream of data is sent to the cloud, with machine analytics performed on a much larger service area. The cloud service expands, as necessary, to support longer time spans. Moreover, instead of data taken at a few instants in time, the data is nearly continuous and provides a complete picture of the virtual cable hub health.

7. Conclusion

As was predicted over the past few years, virtualization has arrived in the cable edge. The virtual cable hub will leverage the numerous and substantial virtualization benefits from other industries and apply these tried-and-true virtualization concepts for the first time. These concepts are the keys to unlocking the path to sustainably growing capacity, adapting quickly to customer demands, and a solution which is flexible and elastic enough to dynamically augment and shift resources to the most in-demand applications.

Abbreviations

AAA	Authentication, Authorization and Accounting
ASIC	Application-Specific Integrated Circuit
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
FTTH	Fiber To The Home
FPGA	Field-Programmable Gate Array
Gbps	Gigabits Per Second
GCP	Generic Control Plane
HFC	Hybrid Fiber-Coaxial
HW	Hardware
I/O	Input/Output
IoT	Internet Of Things
ISSU	In Service Software upgrade
MAC	Media Access Control
NFV	Network Function Virtualization
NIC	Network Interface Controller
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
SDV	Switched Digital Video
SLA	Service Level Agreement
STB	Set-Top Box
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
VPN	Virtual Private Network

Bibliography & References

Branson, Tony. "Database Scalability: Vertical Scaling vs Horizontal Scaling." *vcloudnews.com*. VCloud News, 5 Dec. 2016. Web. 31 July 2017. <<http://www.vcloudnews.com/database-scalability-vertical-scaling-vs-horizontal-scaling/>>

"Cloud Computing with Amazon Web Services." *aws.amazon.com*. Amazon Web Services, Inc, 2017. Web. 31 July 2017. <<https://aws.amazon.com/what-is-aws/>>

DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I11-170510, May 10, 2017, Cable Television Laboratories, Inc.

Dwivedi, Vijay. "Container Orchestration Tools: Compare Kubernetes vs Mesos." *platform9.com*. Platform9, 11 Aug 2016. Web. 31 July 2017. <<https://platform9.com/blog/compare-kubernetes-vs-mesos/>>

Ellingwood, Justin. "The Docker Ecosystem: An Overview of Containerization." *digitalocean.com*. DigitalOcean Inc., 1 Feb. 2015. Web. 31 July 2017. <<https://www.digitalocean.com/community/tutorials/the-docker-ecosystem-an-overview-of-containerization>>

Luther King Jr., Martin.: April 26, 1967, Cleveland speech

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.

Popek, G. J.; Goldberg, R. P. (July 1974). "Formal requirements for virtualizable third generation architectures". *Communications of the ACM*. 17 (7): 412–421. doi:10.1145/361011.361073.

"Popek and Goldberg virtualization requirements." *en.wikipedia.org*. Wikipedia Foundation, Inc., 25 June 2017. Web. 31 July 2017. <https://en.wikipedia.org/wiki/Popek_and_Goldberg_virtualization_requirements>

"Production-Grade Container Orchestration." *kubernetes.io*. The Kubernetes Authors, 2017. Web. 31 July 2017. <<https://kubernetes.io>>

Radio Frequency Interface Specifications, SP-RFII01-970326, Mar. 26, 1997, Cable Television Laboratories, Inc.

Remote Downstream External PHY Interface Specification, CM-SP-R-DEPI-I07-170524, May 24, 2017, Cable Television Laboratories, Inc.

Remote PHY Specification, CM-SP-R-PHY-I07-170524, May 24, 2017, Cable Television Laboratories, Inc.

Tholeti, Bhanu P. “Learn about hypervisors, system virtualization, and how it works in a cloud environment.” *ibm.com*. IBM, 23 Sept. 2011. Web. 31 July 2017.

<<https://www.ibm.com/developerworks/cloud/library/cl-hypervisorcompare/index.html>>

“Virtualization.” *en.wikipedia.org*. Wikipedia Foundation, Inc., 11 July 2017. Web. 31 July 2017.

<<https://en.wikipedia.org/wiki/Virtualization>>

“Virtualization and the Internet of Things.” *windriver.com*. Wind River Systems, Inc., Jan. 2016. Web. 31 July 2017. <<https://www.windriver.com/whitepapers/iot-virtualization/1436-IoT-Virtualization-White-Paper.pdf>>

“Virtualization: The Promise vs. The Reality.” *westconference.org*. Tintri, Apr. 9 2015. Web. 31 July 2017. <http://westconference.org/WEST17/CUSTOM/pdf/WEST_Tintri_Thoughtleadership.pdf>

“What are microservices?” *opensource.com*. Opensource.com, n.d. Web. 31 July 2017.

<<https://opensource.com/resources/what-are-microservices>>

“What are Microservices?” *pivotal.io*. Pivotal Software, Inc., 2017. Web. 31 July 2017.

<<https://pivotal.io/microservices>>

“What is Cloud Native?” *cncf.io*. The Linux Foundation, 2017. Web. 31 July 2017.

<<https://www.cncf.io/>>

“What is a virtual appliance?” *turnkeylinux.org*. Turnkey Linux, n.d. Web. 31 July 2017.

<<https://www.turnkeylinux.org/virtual-appliance>>

Yadav, Rishi. “What real cloud-native apps will look like.” *techcrunch.com*. TechCrunch, 3 Aug 2016. Web. 31 July 2017. <<https://techcrunch.com/2016/08/03/what-real-cloud-native-apps-will-look-like/>>

Getting Real Performance from a Virtualized CCAP

A Technical Paper prepared for SCTE•ISBE by

Mark Szczesniak
Software Architect
Casa Systems, Inc.
100 Old River Road
Andover, MA, 01810
978-688-6706
mark.szczesniak@casa-systems.com

Introduction

Virtualization of network functions and software defined networking (SDN) control of those functions promises service providers tantalizing benefits including faster time to market for new services, lower costs, and higher customer satisfaction. These benefits are particularly important as competition and user demand continue to rise year after year. But, virtualized solutions need to not only exist but also perform at least as well as their legacy counterparts. This is the challenge service providers are coming up against as virtualization initiatives try to make their way out of the lab and into the field.

The good news is that there are answers for the performance challenge. In the technical paper “Getting Real Performance from a Virtualized CCAP”, Casa Systems will present the underlying performance challenges inherent in implementing converged cable access platform (CCAP) functionality (video, IP voice and data) on the current generation of commercial off-the-shelf (COTS) x86 servers.

These challenges include:

- Guaranteeing performance of shared network function virtualization (NFV) infrastructure
- Designing for security in the face of cryptographic performance constraints in virtualized environments
- Providing maximum packet throughput for virtual network functions (VNFs)

This paper will explore each of these challenges and present the underlying factors driving these issues. Further, this paper will explore alternative solutions available today and options service providers should consider as they introduce virtualized CCAP in their networks. Some of the solution aspects to be explored include:

- Optimum server performance characteristics
- Optimum NFV infrastructure and configuration
- Using software solutions like Linux new API (NAPI) or data plane development kit (DPDK) to enhance server performance
- Enhancing VNFs for maximum performance and throughput

By analyzing the underlying challenges inherent in virtualizing full CCAP functions, and understanding the options available today to help overcome those challenges, this paper seeks to aid the advancement of virtual CCAP in the field.

Content

1. Background

Before virtualization, CCAP vendors provided physical network functions (PNFs), which included all the hardware and software needed for the full CCAP solution. With virtualization, CCAP vendors only provide the VNFs, which run on generic server hardware. A virtualized infrastructure manager (VIM) is

used to manage and coordinate the physical resources needed by the VNFs. This means that service providers not only need to select the CCAP VNF; they also need to select a VIM and the servers on which everything will run.

Most virtualization infrastructure is tailored to cloud and enterprise applications, which were focused on allocating central processing unit (CPU), memory and disk to applications. Networking was expected as a given, as networking is always shared among applications. With virtualization of service provider network functions, networking performance becomes vitally important and leads to new challenges for virtualization infrastructure.

CCAP virtualization requires the physical / coax layer to be separate from the server. This adds an additional layer of complexity from the need to address security between the CCAP core and the device where the packet is translated to analog. For remote physical (PHY) layer designs, data packets are already encrypted in the core by the data over cable service interface specification (DOCSIS) media access control (MAC) layer, which uses baseline privacy interface (BPI+) encryption. For remote MAC-PHY designs, internet protocol (IP) encryption may be needed to protect the packets between the core and the remote MAC-PHY device.

To achieve the performance needed from virtualized service provider network applications, including virtualized CCAP (vCCAP), solutions need to move packets as quickly as possible in and out of the virtualization platform as well as rapidly encrypt and decrypt packets.

2. Hardware

There are many things to consider when buying a server for NFV. In addition to aspects like disk space, memory and CPU that one normally considers, network interface controller (NIC) bandwidth and non-uniform memory access (NUMA) are important for NFV. The space and number of available slots in the server may limit the number and type of NICs available. A multiple CPU server will have multiple NUMA regions with limited bandwidth between these regions. One may also want to consider adding hardware assist as a valid tradeoff to adding more servers. Hardware assist can be provided via accelerators integrated within the CPU or platform controller hub (PCH) of the server itself, or through specialized silicon (i.e., field programmable gate array (FPGA), application-specific standard product (ASSP), or application-specific integration circuit (ASIC)) added via peripheral component interconnect express (PCIe) add-in cards. These solutions can be used for CPU intensive tasks such as encrypting/decrypting packets or computing cyclic redundancy check (CRC) / checksum where the full packet needs to be traversed by the CPU.

Disk space is not normally an issue for NFV. Most VNFs will need some disk space for the VNF virtual machines (VMs) as well as some configuration storage, but neither consumes very much disk space. If analytics is part of the VNF, disk space may need further consideration, as analytics normally requires a large amount of data.

Memory is important for the control plane to store session/connection information. Some memory is also needed for packet buffers for forwarding, but buffer count should be limited to reduce packet jitter and delay. Even though the control plane may need a lot of memory, the amount of memory a server can have is normally 1-2 orders of magnitude greater than an embedded system, so running out of memory should not be an issue for VNFs.

The remaining choices which are relevant to server selection (e.g., CPU, NICs, and hardware assist) all impact packet throughput.

CPU Key Attributes for vCCAP	
# Cores	Determine number of simultaneous packets that can be processed
Processor Speed	Higher speeds enable faster per packet processing
Memory Bandwidth	Limits transfer of packets between NICs and Hardware Assist and memory
Connectivity	Limits packet throughput speeds

Figure 1 - CPU Key Attributes for vCCAP

When choosing the CPU, one should consider the number of cores, the speed, the memory bandwidth and connectivity, as shown in Figure 1. More cores give the ability to process more packets simultaneously. High speeds will process each packet faster. Memory bandwidth can limit the transfer of data between NICs and hardware assist and memory and thus limit the ability for the core to process packets. Most CPUs use PCIe for connectivity, though some have built-in NICs. Servers have PCIe slots to allow users to add extra capabilities to the servers. These slots are normally used for connectivity (NIC PCIe cards), but can also be used for other types of expansion such as hardware assist. PCIe is the current standard for expandability. Most servers have a mix between PCIe 2.0, which can transmit up to up to 4 Gbps per lane, and PCIe 3.0 which can transmit up to 7.877Gbps per lane. PCIe 4.0 will double PCIe 3.0 speeds when it is available.

There are CPUs on the market, which have 40 PCIe 3.0 lanes, and the expectation is that others will soon come to market with 48 PCIe 3.0 lanes. Some of those lanes may be used for standard interfaces such as universal serial bus (USB) and disk connectivity. Others are made available via LAN on motherboard (LOM) / daughter card and PCIe slots. LOMs / daughter cards are different terminologies to describe a server specific card that provides some NIC functionality. Most PCIe slots are either 8 or 16 lanes. An 8 lane PCIe 3.0 slot would be able to transmit 63Gbps, which could effectively handle a 50Gbps interface. A 16 lane PCIe 3.0 slot would be able to transmit 126Gbps, which could effectively handle a 100Gbps interface.

Some CPUs have a special bus for multi-socket designs. Most servers support 1 or 2 sockets, though some will support 4 sockets. Most servers with multiple CPUs use NUMA and thus need to use the inter-CPU bus for some memory accesses. Because the bus between the sockets is limited, it is better not to move data between sockets too frequently.

Most servers come with a few 1G NICs by default and have a couple of options for adding or upgrading NICs. They normally either have a LOM / daughter card as the first option for NICs or have built-in NICs. If more NICs are needed, the PCIe slots can be populated with NIC cards. For NFV, PCIe slots will be needed to maximize packet throughput. The number of PCIe slots used will depend on the number and size of slots available and whether hardware assist is used, as hardware assist will also require PCIe slots.

In a 1-rack unit (RU) server, there are normally either 2 or 3 slots. In the 3 slot models, all the slots are normally low-profile (LP) or half-height (HH) which leaves less space for ports. In addition, some of the slots are x8 lanes and some are x16 lanes. For a 2 RU server, as illustrated in Figure 2, there can be up to 7 slots, most of which are normally full height, but most are also x8 lanes because the CPUs do not provide enough lanes for 7 slots of PCIe connectivity. Because the standard interfaces normally take about 16 lanes, a 2 CPU server may have 64 lanes for connectivity.

Since most servers that support 2 CPUs are sold in both 1 CPU and 2 CPU configurations, the standard interfaces are normally all connected to the first socket, thus limiting the PCIe lanes available for NICs (and for Hardware Assist) to 24. The second socket will normally be able to use all 40 lanes for NICs (and Hardware Assist). In addition, a 1RU server also may not have the slot connections designed for uniformity in that the first socket is attached to 2 slots and the LOM / daughter card and the second socket is attached to 1 slot. This means that the first socket will have 8 lanes per slot and 8 lanes for the LOM / daughter card and the second socket will have 16 lanes for one slot and 24 lanes unused. 2RU servers do not have this second issue as they have space for more slots to use up to all 40 lanes.

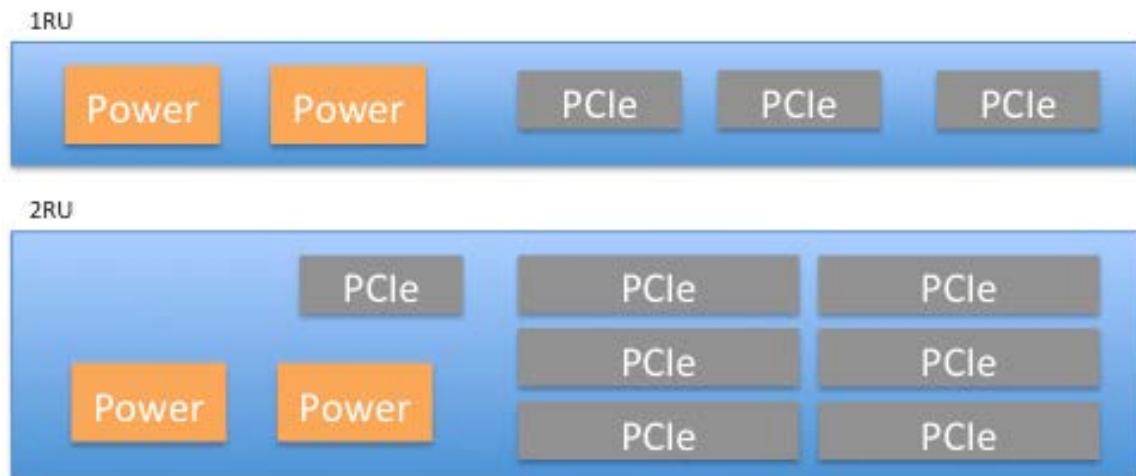


Figure 2 - 1RU vs 2RU PCIe Slots

For maximum throughput, packets should ingress, be processed and egress via the same socket and PCIe lanes attached to the socket. As the CPUs are normally the same across the sockets, it is optimal if the slots / PCIe lanes are allocated the same across sockets (i.e. first socket with 8 lanes for standard interfaces, 16 lanes for daughter card / LOM and 16 lanes for 1 slot and the second socket with 8 lanes for standard interfaces and 16 lanes for each of 2 slots). This symmetric design increases the bandwidth from 250Gbps to 400Gbps for the server. Though 200Gbps is probably a lot more than a high end CPU could process, it allows for half of the bandwidth to be taken by hardware assist (for crypto). Future servers are expected to address the importance of input / output (I/O) balance by providing PCIe connections split equally across both processors.

Additionally, it is important to consider the PCIe generation supported by the slot as some servers have both Gen2 and Gen 3 slots. There can also be a difference between the physical width of a PCIe slot and the width supported electrically. Some slots can physically take an x16 PCIe card, but only 8 PCI lanes are electrically connected to the PCIe controller. The connectivity of the PCIe slots to the PCIe controller

can also impact throughput, as some slots are connected to PCIe controllers integrated inside the CPU as some are connected to ones inside the PCH. The latter creates a longer path to the CPU and memory.

Another limitation on the number of NICs is the size of the slot. The LOM / daughter cards and full height slots can normally support up to 4 ports. Low profile / half height slots can normally only support up to 2 ports. So, a 1RU server could support up to 10 ports. A 2RU server could support up to 30 ports. When you put the size restriction and lane restriction together, the 1RU server could support a max of 100Gbps via (10) 10Gbps ports or 400Gbps with (4) 100Gbps ports in a symmetric server. An asymmetric 1RU server would be limited to (3) 50Gbps (limited by the x8 lanes) and (1) 100Gbps ports. The 2RU server would be limited to (30) 10Gbps ports or (8) 50Gbps ports (limited by the x8 lanes). Since max NIC bandwidth for both 1RU and 2RU servers is 400Gbps, a rack of 1RU servers could provide twice the bandwidth of a rack of 2RU servers if the NICs were the limiting factor.

The NIC bandwidth may also be limited by the use of hardware assist. For most applications, if hardware assist is needed, the hardware assist bandwidth should match the NIC bandwidth and thus NIC bandwidth is cut in half in a symmetric server and is reduced even more in an asymmetric server. The asymmetric 1RU server only has 1 slot and thus cannot handle both a NIC card and a hardware assist card connected to the second socket. This means to use the second socket with hardware assist, all packets would need to traffic the quick path interconnect (QPI) bus, which would limit throughput, as shown in Figure 3.

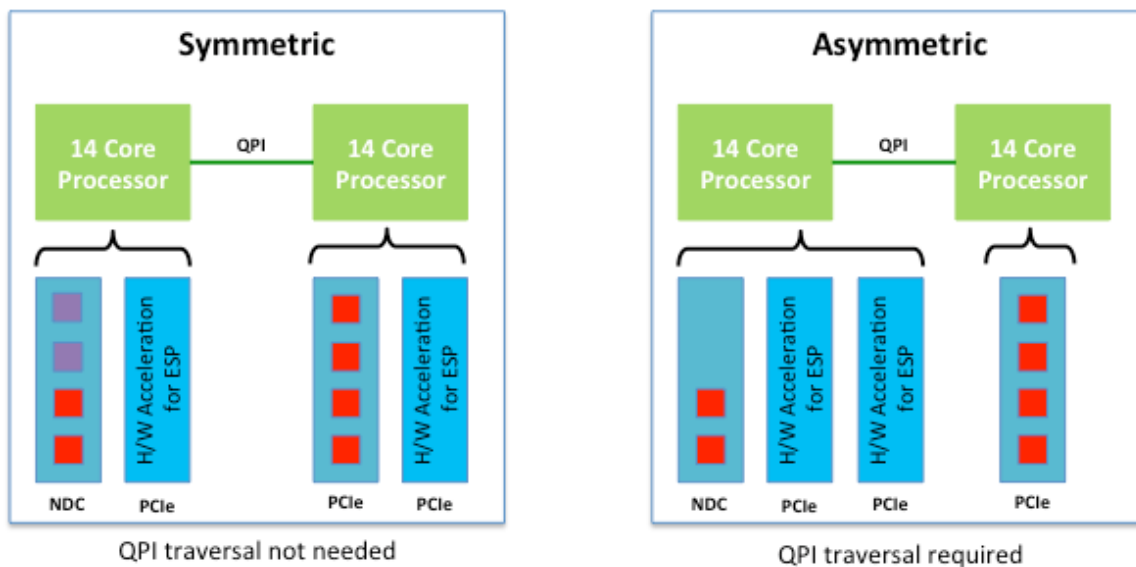


Figure 3 - Symmetric vs Asymmetric QPI Traversal

Another limiting factor is the number of memory channels per CPU. The high-end servers usually provide all the memory channels that the CPU can support. Current server models support 4 – 6 channels. The 6-channel servers have a 50% boost in memory bandwidth, as shown in Figure 4. Even if the server supports the full number of channels, there may or may not be memory in each channel.

Sockets	Channels per Socket	Max Bandwidth*
1	2	230 Gbps
1	4	460 Gbps
1	6	690 Gbps
2	2	460 Gbps
2	4	920 Gbps
2	6	1380 Gbps

*Based on 2400 Mt/s DIMM and 75% efficiency

Figure 4 - Channels per socket vs maximum bandwidth

When purchasing memory for the server, there is a choice between more modules with less memory each or fewer modules with more memory each. The modules with more memory are usually more economical, but they can reduce memory performance. For example, if you want to buy 64G of memory, you could buy (1) 64G dual in-line memory module (DIMM), (2) 32G DIMMs, (4) 16G DIMMs, (8) 8G DIMMs, or (16) 4G DIMMs. Purchasing (8) 8G DIMMs would allow for 1 DIMM per memory channel, thus giving the best memory bandwidth, but it also means that if you need to upgrade, you may need to replace your DIMM as opposed to just adding more.

The information above is all based on rack servers, but there are also chassis / blade systems. These systems can be denser in terms of CPU, but tend not to have the expansion capabilities to add NICs nor hardware assist, thus limiting throughput. In a normal datacenter, CPU horsepower is more important and thus the chassis / blade servers seem to be the better choice. But for network function virtualization, throughput is more important and thus separate servers seem to be the better choice.

Another thing to consider is how the servers are connected. A rack of 30 servers with (10) 10G ports would need a TOR (Top of Rack) switch with (300) 10G ports, whereas, if each server had (1) 100G port, the TOR switch would only need (30) 100G ports. The 10G TOR would be at least 6RU to support the same bandwidth as a 1RU 100G switch for inter-server connections. Either TOR switch would also need uplinks, which would add another 1-2RU to either switch depending on the ratio of uplink traffic to inter-server traffic. The extra size for uplink traffic is due to the fact that the current selections of lower-end 100G switches do not have higher speed uplinks. A 48 port 10G switch can have (5) 100G uplinks and not need to oversubscribe the (48) 10G ports used for server connections. A 32 port 100G switch would only be able to provide 16 ports for servers and 16 ports for the uplink without oversubscription. The 100G switch still has higher overall bandwidth of 1.6Tbps vs. 480Gbps and less cabling (32 vs. 53).

3. Virtualization Infrastructure

There are many different choices for a VIM, with VMware's vSphere and OpenStack being the most prevalent. Many vendors implement their own version of OpenStack with proprietary extensions. Some of these extensions are designed specifically for NFV to help VNFs scale and perform better than just using standard open-sourced OpenStack.

VMs need to be placed where they can access memory and NICs attached to the same NUMA region as the CPU. If not, the QPI bus between CPUs will need to be used, which is limited. This bus is limited because it is expected that most accesses would not need to traverse between NUMA regions. A core in one NUMA region that requires access to memory of the other NUMA region will be required to access the memory via the QPI bus. If this were an update, there would be 2 accesses to the bus (1 for the read and 1 for the write). This can be accomplished either by only giving the VM access to devices in 1 NUMA region or by enhancing the processes in the VM to understand NUMA and group its resources by NUMA region.

VMs can take a bit of time to scale up, so it is worthwhile to spin up an extra VM before it is needed, such that it will be ready when it is needed.

When scaling VNFs that have flow state to multiple servers, a load-balancer is needed to direct packets to the correct VM. The TOR can do this if it is capable of the hash function that will work for the specific VNF. If not, the TOR will need to spray packets across multiple load-balance VMs that can then hash and send the packets to the correct VM to process the flow.

4. Software

Forwarding performance is not just affected by the hardware design of the server as described above. The software running on the server also affects it. Without virtualization, by default, all packets pass through the Linux kernel, which can be very slow, thus limiting packet throughput to under 10Gbps in a server even though 100Gbps NICs are available for servers. This is because the Linux kernel has not been optimized for forwarding. Also, packets that need to get to user space need to have 2 full packet copies when transitioning from kernel to user and then back to the kernel. These full packet copies take up a lot of time for the CPU. Additionally, the kernel uses interrupts to handle incoming packets, which have context switch overhead for each and every packet.

Data plane development kit (DPDK) technology has solved these issues with poll-mode drivers that run in user-space and with optimized libraries for memory and packet management. Packets are no longer handled in the kernel when using DPDK, as shown in Figure 5.

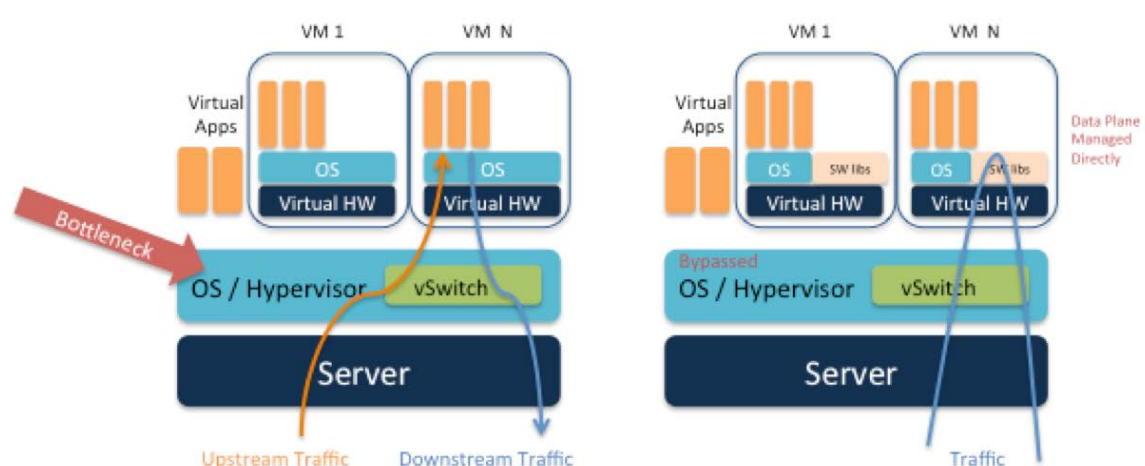


Figure 5 - Direct Management of Data Plane

They are put into a buffer that the user-mode process is polling to process packets. No interrupts will be generated and thus no context switch on a received packet. When virtualization is added, the vSwitch can become the bottleneck. By default, the vSwitch normally runs in the host OS's kernel, including Open vSwitch (OVS), which is used by OpenStack. This can limit server throughput to less than 10Gbps.

There are a few different options for fixing this issue. One is to get an accelerated vSwitch to speed it up. There are also DPDK extensions that will speed it up as well. Another option is to bypass the vSwitch by using single root input / output virtualization (SR-IOV) or Ironic. SR-IOV bypasses the vSwitch and gives packets directly to the VM. Ironic runs the machine natively on the server instead of within a VM. Ironic is OpenStack's terminology for setting up a bare metal server with a specified image and networking. It is limited in that the NICs are determined by the physical hardware and not by the VM setup.

The type of actions performed on the packet can also have an affect on performance. For normal bridging and routing, only the header of the packet needs to be viewed, so only the header actually needs to be pulled into memory. But for operations like crypto and CRC, the whole packet needs to be read and thus needs to be transferred to memory, as shown in Figure 6.




Figure 6 - Bytes Processed per Function

Although CPU can perform the crypto and CRC algorithms, the time it takes to transfer the whole packet into memory and traverse the packet will take a lot of cycles. Adding more cores can compensate for the need for extra cycles, but adding hardware assist is more economical than adding cores. Adding crypto cards to existing servers is much less expensive than adding more servers, both in terms of capital expenditure (CAPEX) and operating expenses (OPEX). The amount of savings will depend on the ratio of cores needed for crypto to the cores needed for the VNF processing of packets.

For the different VNFs required by a vCCAP, the crypto core overhead is between 50% and 250%. There are some optimized crypto libraries than can help increase crypto performance, but they still do not approach the performance of dedicated crypto hardware. Another option for increasing CRC performance is to perform incremental CRC, but this requires access to the input CRC, which is not always available to VMs.

Another consideration is how packets traverse infrastructure and VMs to provide the functionality of the VNF. If a VNF were comprised of 1 VM that processes the packets in 1 pass, a packet would go from TOR to Server to TOR, as shown in Figure 7. A 1 Gbps stream of packets would require a 1Gbps connection from the TOR to the server and from the server back to the TOR. If however, the VNF required 3 VMs to process each packet, the 1Gbps stream could require up to 3Gbps of connection from TOR to server and server to TOR. If all 3 VMs were in the same server, the traffic could be optimized to stay in the server, but if the second VM were on a different server than the other 2, all 3Gbps would be needed. This may not be an issue for low bandwidth applications, but when the application scales up to 100 Gbps streams, requiring 300Gbps to handle the packet within the rack can be costly. The 100/300Gbps is the aggregate throughput of all flows through the application (and thus through all VMs).

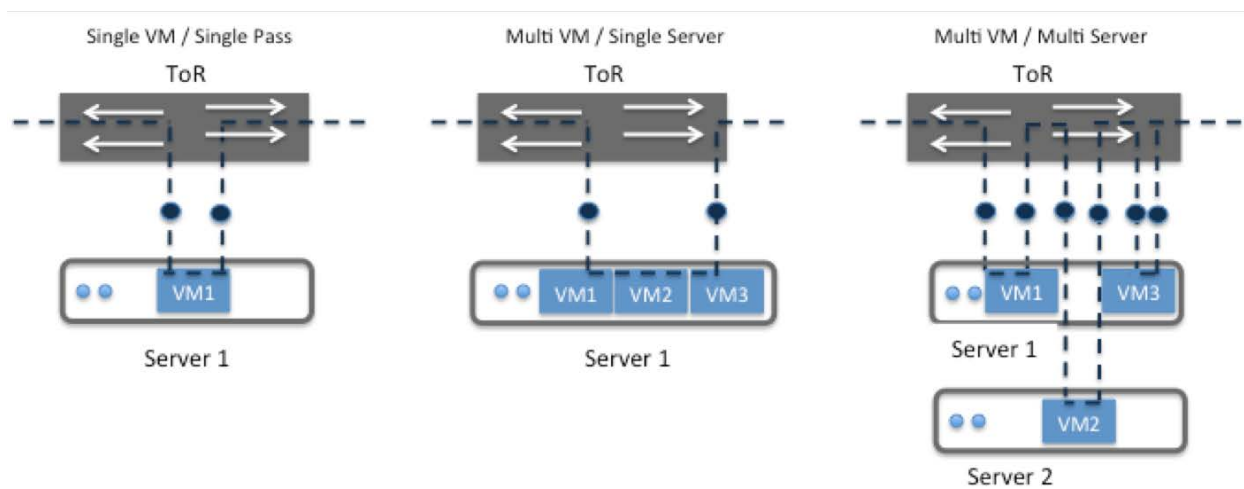


Figure 7 - VM and Server Impact on TOR Traversal

Also, when scaling to high bandwidth streams, there needs to be a way to direct packets to the correct cores within the correct VM from outside the server. A single core will only be able to direct up to 20-30Gbps of traffic to other cores before its queues overflow and packets are dropped, so sending it a 100Gbps stream means dropping 70-80% of the traffic. Most higher speed NICs support Receive Side Scaling (RSS) receive queues and use a hash to select to which VF (Virtual Function) to send packets. This also means VMs will need to use VFs and not the physical interfaces.

There are other new challenges arising from virtualization. Applications may need to evaluate session setup and modification rates as well as messaging speeds. In embedded systems, some applications use shared memory to speed up IPC (Inter-Process Communication) message handling. In a virtualized environment, applications may grow to more than one virtual machine and thus can no longer take advantage of shared memory.

Applications may need to rework their interactions to send larger amounts of data less often. They should also move to asynchronous communication to allow other processing to occur while waiting for replies. These are good practices for any multi-threaded or multi-process application, but were not always necessary in an embedded software environment. Certain VNFs, such as vCCAP may need clock synchronization for functionality such as upstream scheduling. Servers do not usually have a sufficient clock for driving synchronization protocols such as IEEE 1588 such that the master clock will need to come from outside the server. The vCCAP will need to be a slave to that master clock.

5. Wrap-up:

When looking at a server for forwarding performance, like that required by a vCCAP, NIC bandwidth, memory bandwidth and CPU core bandwidth need to be considered. If the application needs 1 core to process 10Gbps of traffic, then each core can match a 10G NIC and only 10 cores are needed in a 1RU server filled with 10G NICs. This would be reasonable for a low-end server with low-end CPUs (6 cores per CPU). Or the 10 cores could be used to fill a 100G NIC per socket, thus (2) 10-core CPUs would be needed. Current top of line servers with 22 cores can support (2) 100G NICs per socket. As long as 4 memory channels per socket have memory, the memory bandwidth is enough for the 200Gbps of traffic per socket. If the application requires 4 cores to process 10G, the top of line server would be needed to support just 100G of traffic (10 10G NICs or 2 50G NICs). The choice for 10G NIC will only work if hardware assist is not needed (as it would take away 2 PCIe slots leaving only space for 6 NICs).

Conclusion

The control plane pieces of VNFs should be able to run well on any server that meets the memory and disk requirements, but the data plane pieces of VNFs need the networking piece to work well. This requirement greatly reduces the number of servers from which to choose. The control plane and data plane parts of the VNF should therefore be split so each can be scaled independently. A VNF will get the most throughput per RU in a 1RU server with 2 (hardware assist needed) or 4 (no hardware assist needed) 100G ports with symmetric allocation of PCIe slots to CPU sockets. Either the VIM will need to allocate VMs for a VNF with all resources from the same NUMA region or the VNF will need to understand NUMA regions and be able to group resources per NUMA region. Data plane VMs either need to be connected via an accelerated vSwitch or use SR-IOV to bypass the vSwitch. They should use DPDK to bypass the kernel.

Abbreviations

AES-NI	Advanced encryption standard instruction set – new instructions
API	Application program interface
ASIC	Application-specific integration circuit
ASSP	Application-specific standard product
BPI	Baseline privacy interface
CAPEX	Capital expenditure
CCAP	Converged cable access platform
COTS	Commercial off-the-shelf
CPU	Central processing unit
CRC	Cyclic redundancy check
DIMM	Dual in-line memory module
DOCSIS	Data over cable service interface specification
DPDK	Data plane development kit
FPGA	Field programmable gate array
Gbps	Gigabits per second
HH	Half height
I/O	Input / output

IP	Internet protocol
IPC	Inter-process communication
LAN	Local area network
LOM	LAN on motherboard
LP	Low profile
MAC	Media access control layer
NAPI	New API
NFV	Network function virtualization
NIC	Network interface controller
NUMA	Non-uniform memory access
OPEX	Operating expenditure
OS	Operating system
OVS	Open vSwitch
PCIe	Peripheral component interconnect express
PCH	Platform controller hub
PHY	Physical layer
PNF	Physical network function
QPI	Quick path interconnect
RSS	Receive side scaling
RU	Rack unit
SCTE	Society of Cable Telecommunications Engineers
SDN	Software defined networking
SR-IOV	Single root input / output virtualization
Tbps	Terabits per second
TOR	Top of rack
USB	Universal serial bus
VF	Virtual function
vCCAP	Virtualized CCAP
VIM	Virtualized infrastructure manager
VM	Virtual machine
VNF	Virtual network function

Cable Access Network Virtualization

Headend Re-architected as a Data Center

A Technical Paper prepared for SCTE•ISBE by

Ruobin Zheng

Cloud Networking Chief Researcher
Huawei Technologies Co.,Ltd.
Huawei Headquarter, Bantian, Longgang District, Shenzhen, China
(86) 075528978020
zhengrubin@huawei.com

Wenle Yang

Senior Engineer
Huawei Technologies Co.,Ltd.
Huawei Headquarter, Bantian, Longgang District, Shenzhen, China
(86) 075528977366
yangwenle@huawei.com

Introduction

A significant rise in the number of diversified services and applications has promoted the rapid development of the broadband industry. With the continuous rollout of new services, such as 8K video and virtual reality (VR), the access bandwidth starts to move from megabit to gigabit. Multiple system operator (MSO) networks face lots of challenges as they try to match the diverse set of service requirements and service characteristics.

1. Management Complexity for Multi-service and Multi-access

MSO networks are migrating to support full services that cover residents, enterprises, mobile backhaul, and wholesale services. With the rise of new services (e.g. 4k, 8k, VR) and the trend of Internet of Things (IoT) and Fifth-generation (5G), service requirements may further include extreme broadband, ultra-low latency, massive connections and ultra-high reliability, which impose strong demands on the access network.

Besides, the new technologies further introduce massive remote nodes. While the cable access network architecture migrates to the distributed mode, the number of standalone remote nodes will scale up greatly. This growth in standalone remote nodes and the need to separately manage them will lead to a longer Time to Market (TTM) and complex Operations and Maintenance (O&M).

2. Scalability and Energy Efficiency in Remote Nodes

In the IoT and 5G deployment scenario, many access elements are widely distributed in outdoor facilities. One of the performance targets is ultra-low energy consumption. These new low-power remote nodes need to be as simple as possible especially as it relates to Operations, Administration and Maintenance (OAM) and a long service life.

Considering the complexity of the IoT network and the heterogeneous network, the remote nodes will need to handle various protocols. With all of these functions, the remote nodes will become very complex, making it hard to meet the requirements mentioned above.

3. Big Gap between Slow Network Evolution and Rapid Service Innovation Requirements

Regarding the traditional way for new service provisioning, a variety of access devices need to be separately configured and may even require hardware changes. In addition, network topology, vendor switch model, and software version all must be taken into account. Especially when many access elements are widely distributed in outdoor facilities, it is not easy to touch all of them. That means the current “rigid” network is extremely hard to evolve to the next generation and will lead to a fairly long time to market for new services.

4. Sharing Difficulty of Access Network

The current methods for sharing the access network (e.g. bitstream), where service packages are only differentiated by bandwidth and few configurable options, limits the ability to provide more advanced features which in turn limits richer service differentiation.

Traditionally, retailers and service providers have only had a limited capability to monitor, control and manage access resources. As access technologies advance, resources that can be accessed by a third party become increasingly more programmable and feature rich, which brings advantages but also introduces an increased risk that network harm could occur if controls are not exercised properly. It can also be time-consuming for retailers and service providers to do service provisioning. The infrastructure provider usually has to deal with a large range of varying requirements from retailers and service providers, which makes it difficult to do network planning or perform technology changes in a timely manner.

To meet the rapid growth of traffic and connections, access to the Headend requires fiber, and the network needs to be more flat. With respect to the challenges and changes in the access network, this paper explores cable access network virtualization with Software Defined Networks (SDN) and Network Functions Virtualization (NFV) technologies, which aim to provide a more flexible and future-proof access network. The network will be transformed into a data center-based architecture, and network functions and services will run in the cloud. Considering some strong demands such as the low latency requirements for 4K / 8K / VR video and the Internet of things (IoT), the data center is expected to move from the clouds down to the Headend where it is closer to the end devices.

With the concept of re-architecting the Headend as a data center, a Programmable Virtual Converged Cable Access Platform (VCCAP) is introduced as a cloud-oriented access network solution. The following content will give a detailed description of this solution, which includes a high-level virtualization architecture and its important components, with the benefits embodied.

High-level Architecture of Cable Access Network Virtualization

By introducing SDN and NFV, the proposed programmable virtual CCAP is based on the high-level architecture of cable access network virtualization. This is depicted in Figure 1. IoT, 5G, Fiber to the Distribution Point (FTTDp) and distributed CCAP are displayed as the example access modes.




Figure 1 - High-level Architecture of Access Network Virtualization

The main concept is that of a centralized controller for a virtual access network. Based on the separation of control plane and forwarding plane, the control plane of the remote node is relocated and centralized in a virtual CCAP (vCCAP) controller. The vCCAP controller can also act as a vCCAP Hypervisor implementing access-network abstraction and slicing to make multi-tenant and multi-service operation in one physical network. The vCCAP controller can be located in a Headend/Hub which can be deployed as a cloud platform or in other words be re-architected as a data center.

The vCCAP application running on top of the vCCAP controller will be discussed in details in the following section. Customer premises equipment (CPE), Content Delivery Network (CDN), Router, and Quadrature Amplitude Modulation (QAM) can also be virtualized in the data center at the Headend. The vCCAP controller talks to the remote nodes through a southbound interface, and also provides an open northbound interface, e.g. open Network Application Programming Interface (API).

With the control and data planes decoupled, the remote nodes are able to be decoupled from services and applications and become dummy but programmable devices which include a programmable forwarding plane. Thus, long-tail service TTM can be accelerated and the need to upgrade all remote nodes can be eliminated.

Thanks to this centralization of intelligence, remote nodes can become plug & play, and will automatically register to, and be controlled by, the vCCAP controller in the Headend. This virtualized access network architecture can greatly reduce Operating Expenses (OPEX) and facilitate new service innovation.

Virtual CCAP

5. Concept of vCCAP and virtual line

A vCCAP is a logical entity that represents a physical node or a part thereof. Each vCCAP has a group of virtual lines, which represent a group of physical lines connected to the corresponding physical CCAP. Similar to the physical line identification (ID) which identifies the physical line, the virtual lines are identified through a virtual line ID.

Figure 2 shows an example of a vCCAP, which represents one physical remote node (e.g. Cable Media Converter (CMC)). The vCCAP interfaces represent the user-side interfaces of the physical node. The activation of the vCCAP includes:

- When a remote node is up, a vCCAP will be automatically generated in the vCCAP Controller.
- vCCAP will automatically get a management IP address for itself.
- vCCAP initiates self-configuration and self-service provisioning to support remote node plug & play.
- Future new functions or protocol enhancements can be implemented in vCCAP without the need to upgrade the physical remote nodes.
- vCCAP can support protocol conversion between Layer 3 (L3) (e.g. Simple Network Management Protocol (SNMP) / NETCONF / File Transfer Protocol (FTP)/ Common Open

Policy Service Protocol (COPS) / OpenFlow) and Layer 2 (L2) (e.g. ONU management and control interface (OMCI) / Ethernet Operations, Administration, and Maintenance (OAM)).

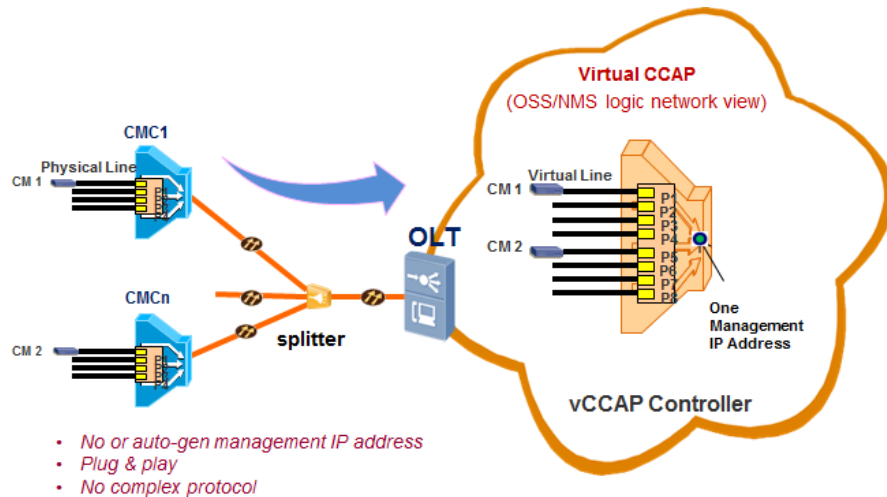


Figure 2 - Overview of vCCAP

The management systems only sees the vCCAP. The vCCAP hides the actual details of the underlying access network infrastructure. A vCCAP allows the physical remote node to be simplified through three aspects:

- The absence of a management IP address or an auto-generated management IP address;
- Plug and play;
- No complex L3 and above protocols, e.g. SNMP/NETCONF/FTP/COPS/Openflow, is possible.

6. Two types of vCCAP

Physical nodes in the cable access network can be combined or segregated to form different vCCAPs. Two types of vCCAP are identified based on different granularity of slicing and abstraction, e.g. per node or per port. In both cases there exists a mapping between physical line IDs and virtual line IDs which are maintained by the vCCAP controller.

Figure 3 describes Type 1 vCCAP where the vCCAP represents one or more than one physical nodes. For example, Remote Node 1 is mapped to vCCAP 1, while Remote Node 2 and Remote Node 3 are mapped to vCCAP 2.

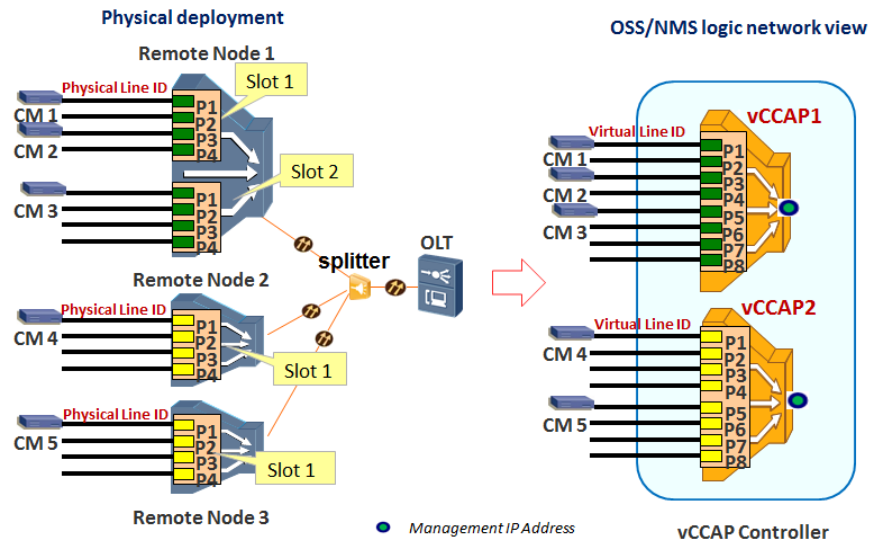


Figure 3 - Type 1 vCCAP

Table 1 shows the mapping between physical line ID and virtual line ID for Type 1 vCCAP described in Figure 3.

Table 1 - Line ID Mapping Table for Type 1 vCCAP

Physical Line ID	Virtual Line ID
Remote node1/slot1/port1	vCCAP1/port1
Remote node1/slot1/port2	vCCAP1/port2
.....
Remote node1/slot2/port1	vCCAP1/port5
.....
Remote node1/slot2/port4	vCCAP1/port8
Remote node3/slot1/port1	vCCAP2/port1
Remote node3/slot1/port2	vCCAP2/port2
.....
Remote node2/slot1/port1	vCCAP2/port5
.....
Remote node2/slot1/port4	vCCAP2/port8

Figure 4 describes Type 2 vCCAP where the vCCAP represents more than one physical interface on more than one physical node. The interfaces of one remote node are assigned to different vCCAPs, and some interfaces may not be mapped into any of the vCCAPs.

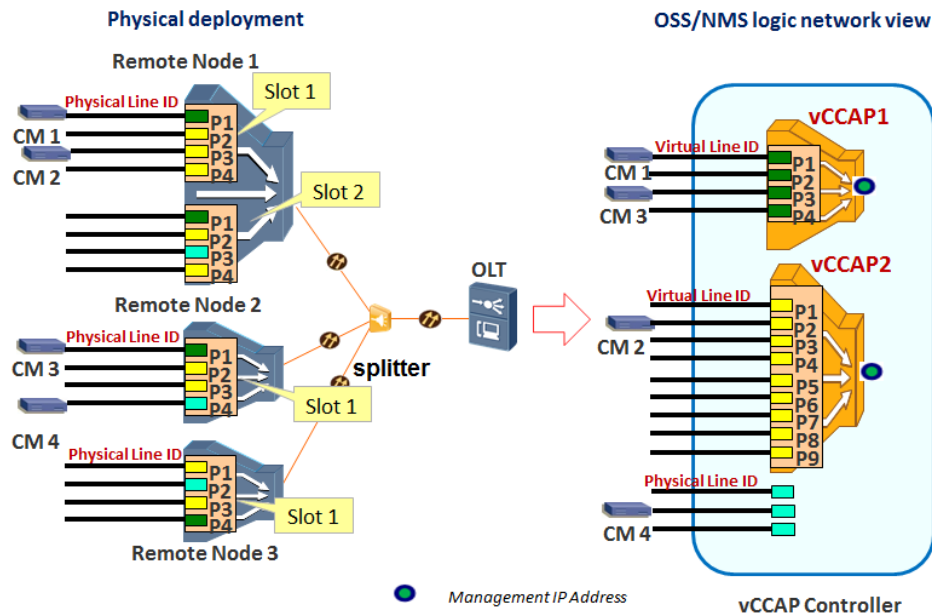


Figure 4 - Type 2 vCCAP

Table 2 shows the mapping between physical line ID and virtual line ID for Type 2 vCCAP described in Figure 4.

Table 2 - Line ID Mapping Table for Type 2 vCCAP

Line ID Mapping Table		Physical Line ID
Physical Line ID	Virtual Line ID	
Remote node1/slot1/port1	vCCAP1/ port1	Remote node1/slot2/port3
Remote node1/slot2/port1	vCCAP1/port2	
Remote node2/slot1/port1	vCCAP1/port3	
Remote node3/slot1/port4	vCCAP1/port4	Remote node2/slot1/port4
Remote node1/slot1/port2	vCCAP2/port1	
Remote node1/slot1/port3	vCCAP2/port2	
Remote node1/slot1/port4	vCCAP2/port3	Remote node3/slot1/port2
Remote node1/slot2/port2	vCCAP2/port4	
Remote node1/slot2/port4	vCCAP2/port5	
Remote node2/slot1/port2	vCCAP2/port6	
Remote node2/slot1/port3	vCCAP2/port7	
Remote node3/slot1/port1	vCCAP2/port8	
Remote node3/slot1/port3	vCCAP2/port9	

Unmapped Physical Ports

Physical Line ID
Remote node1/slot2/port3
Remote node2/slot1/port4
Remote node3/slot1/port2

By slicing and abstracting the physical cable access-network into multiple vCCAPs, the infrastructure provider can support network-sharing for other operators (e.g. retailer) or implement multiple services in one physical access network. The vCCAP and its virtual lines can be controlled and managed by a service provider or a retailer without the need to be aware of changes to physical nodes or physical lines in the

infrastructure network. Security in the infrastructure network can be enhanced, and an easier and more efficient way for the service provider or the retailer to manage their user ports is possible.

Type 1 vCCAP has natural isolation of both forwarding plane and control plane between physical nodes. Only a software upgrade is required during the migration to multi-operator or multi-service sharing in one physical device. It can be deployed in both brown field and green field deployments.

For Type 2 vCCAP, both forwarding resource and control resource of each port have to be isolated in a device. A hardware upgrade may be needed in order to migrate to multi-operator sharing in one physical device. Type 2 vCCAP may only be deployed in a green field deployment.

Access Network Function as a Service

As illustrated in Figure 5, each vCCAP can have a subset of network functions to support different service requirements. These functions can be allocated by vCCAP controller based on the full function set of the infrastructure network. The network function set in each vCCAP can be defined by the infrastructure provider based on the service requirements and the capability of the infrastructure network.

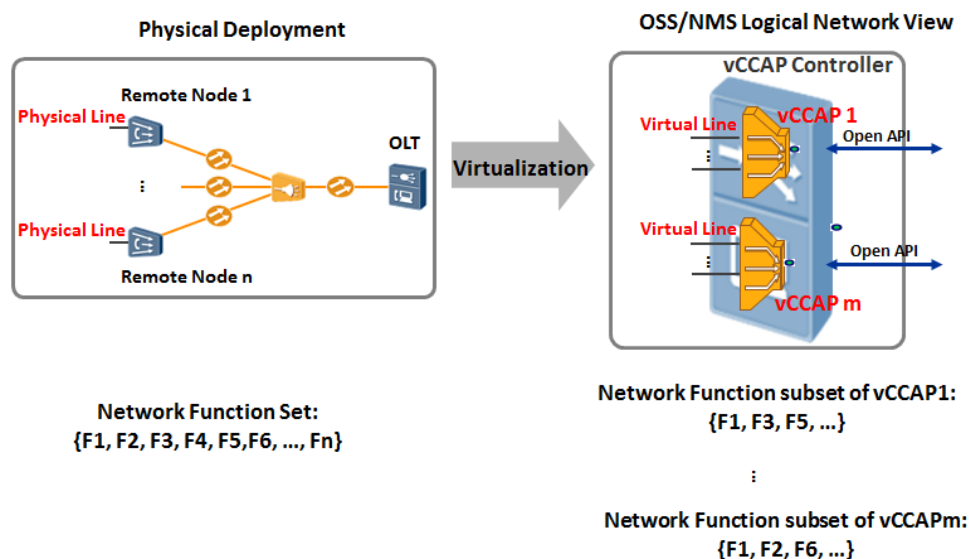


Figure 5 - Access Network Function Virtualization and Allocation

The virtualized access network functions include network functions of the access nodes. They can be functions in the control plane or the data plane. Example functions could be Quality of Service (QoS) policies, filtering, multicast group control, authentication, authorization and accounting, and dynamic address provisioning. The granularity of the network function is flexible so that several functions can be combined into a bigger one and made available as a service.

The access network function as a service can be further illustrated by the following example, where we make the assumption that the infrastructure network has the capability to support the service requirements.

The example assumes that there are requirements for the deployment of mobile backhaul services and residential services, and the infrastructure provider plans to allocate two vCCAPs based on the same physical access network to perform the service deployment. Based on the requirement of mobile backhaul service, the network function subset assigned to one vCCAP can include the clock synchronization function, the Multi-Protocol Label Switching (MPLS) forwarding function, and the IP/MPLS signaling function. Similarly, based on the requirement of residential services, the network function subset assigned to the other vCCAP can include the Authentication, Authorization and Accounting (AAA) authenticator/proxy function, the Dynamic Host Configuration Protocol (DHCP) relay/proxy function, the Internet Group Management Protocol (IGMP) proxy/snooping function and the flow classification and QoS mapping function. Then the two types of services can be provisioned through the two vCCAPs separately yet based on the same infrastructure.

Thus, the network virtualization has the potential to turn traditional access network sharing into Network as a Service (NaaS), which enables:

- Faster time to market and the rollout of new services
- Multi-instance virtualization of access networks

A flexible sharing mode is enabled for the multi-tenant scenario, where the infrastructure operator owns, maintains and virtualizes physical access network resources, while the service operators or third parties can operate, control and manage their assigned vCCAPs and provide differentiated services. This network virtualization is also appropriate for a network operator that wants to slice its access network in order to offer a multi-service solution in one physical network for customers (e.g., for residential, enterprise or mobile backhaul markets) and wants to be able to use a vertical organization structure for aspects related to customers, services and resources.

Programmable vCCAP Reference Architecture

Based on the access network virtualization, Figure 6 shows a detailed reference architecture of the programmable vCCAP. With a vCCAP controller, the programmable vCCAP disaggregates CCAP devices, virtualizes CPE and CCAP control and management functions into the cloud, and distributes Data-Over-Cable Service Interface Specifications (DOCSIS) processing to the remote node.

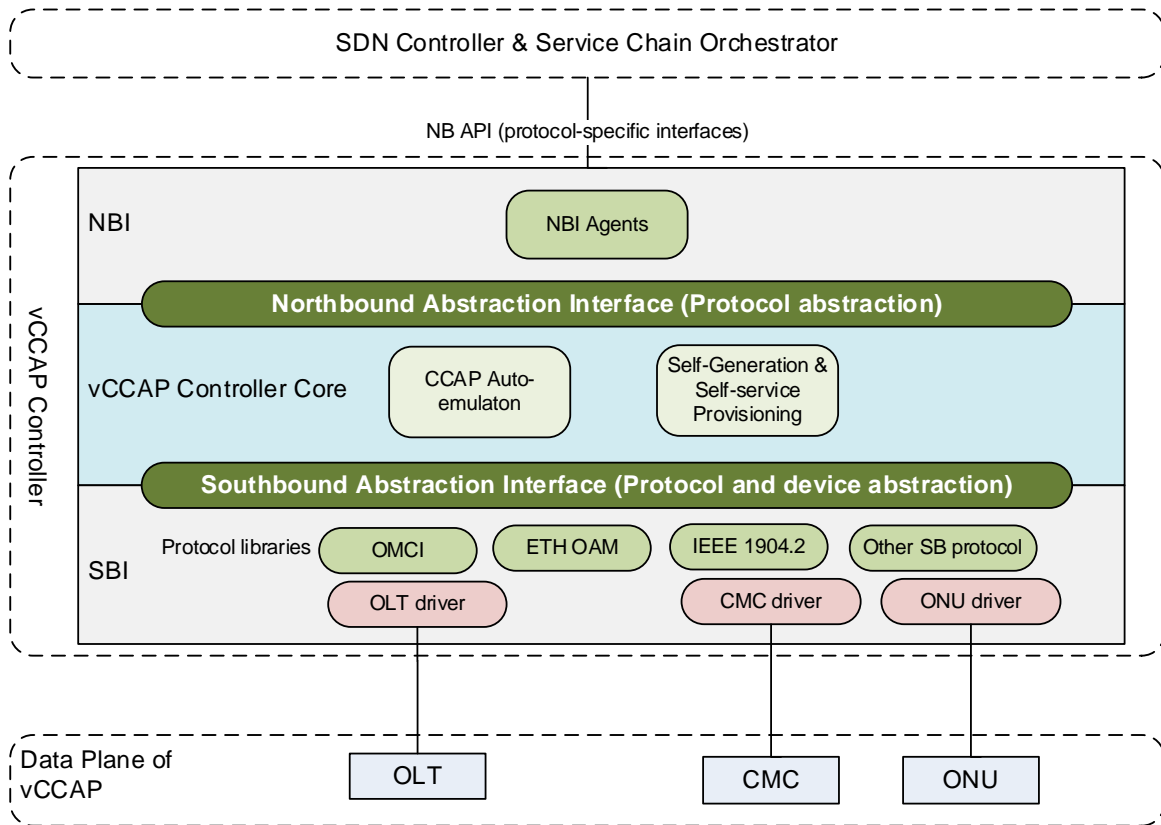


Figure 6 – Programmable vCCAP Reference Architecture

A three-layer architecture is proposed, which consists of an infrastructure layer, a control layer (including vCCAP controller, SDN controller and Service Chain Orchestrator), and an application layer. The vCCAP Controller, which is software running in the cloud, controls and manages the physical devices (e.g. CMC and optical line terminal (OLT)) in the infrastructure layer through the southbound interface (SBI), and the northbound communicates with the SDN controller and Service Chain Orchestrator. The SDN Controller allows applications to tailor network behaviors to suit their needs, and the Service Chain Orchestrator can program services instead of re-architecting the infrastructure layer and the management system for every new service.

Two functions of CCAP Auto-Emulation and vCCAP Self-Generation & Self-Service Provisioning are introduced in the core of vCCAP controller.

The CCAP Auto-Emulation function abstracts and represents the underlying distributed access network elements (e.g. OLT, CMC or Optical Network Unit (ONU)) as one integrated monolithic CCAP, and performs the conversion between this abstract view at the northbound interface (NBI) and the management, configuration, reporting and alarming functions for each of the physical devices at the SBI, which hides the device level details. Different access technologies can be enabled without modifying the existing MSO configuration and management system.

For details, the control and management system configures the forwarding entries for the data plane of the vCCAP without needing to be aware of the access devices' details or the specific access technologies.

Based on some decomposition policies, the CCAP Auto-Emulation function translates the forwarding entries' commands into the forwarding entries of the corresponding access devices (which can be OLTs) and its remote nodes (e.g. CMCs or ONUs), and maps the ingress/egress ports of the vCCAP to the ingress/egress ports on the physical devices.

The vCCAP controller's SBI contains device driver plugins that support communication with the access devices in the network. Device driver plugins, which may be device-specific or generic device drivers, may use southbound protocol libraries provided as common resources, or they can embed their own protocols as needed. Thus, a protocol conversion between layer 3 and layer 2 can be realized by the vCCAP controller, with the purpose of keeping remote nodes as simple as possible (e.g. only layer 2 devices) and leaving the complexity to the vCCAP. By receiving messages with layer 3 protocols such as NETCONF, SNMP, and OpenFlow, the vCCAP controller can use layer 2 protocols in the southbound direction with respect to specific access scenarios such as OMCI for Passive Optical Network (PON), or Institute of Electrical and Electronics Engineers (IEEE) 1904.2 for IoT, 5G, Ethernet 802.3, WiFi etcetera.

With the other module of vCCAP Self-Generation and Self-Service Provisioning in the vCCAP Controller Core, a physical access network can be sliced and programmed into multiple vCCAPs by using machine learning technology. Multi-service or new services can be automatically provisioned through different vCCAPs with different subsets of network or service functions separately under the same physical access network. Thus, it has the ability to turn the traditional business model into NaaS, lower the OPEX and accelerate service innovation and delivery.

Intelligent vCCAP Self-Generator

To create a vCCAP, the main challenges observed are as follows:

- The Operators are usually not aware of what and how many network resources, e.g. the network functions and the service functions are exactly required by the vCCAP, and how to map the virtual resources to the physical resources of the infrastructure. Especially for the service providers who need to lease the virtual resources from the infrastructure provider, it is a waste if their vCCAPs are assigned with spare resources for worst case.
- Facing competition from the OTT providers, the Operators have to speed up the service Time-to-market. Manual upgrade of a vCCAP when service requirements change will lead to a long service deployment time and high OPEX.

With the above considerations, an operation mode of the vCCAP on demand is explored. The Programmable vCCAP system provides an intelligent vCCAP self-generator that allows the vCCAP to be generated automatically based on the abstracted service requirements, as depicted in Figure 7.

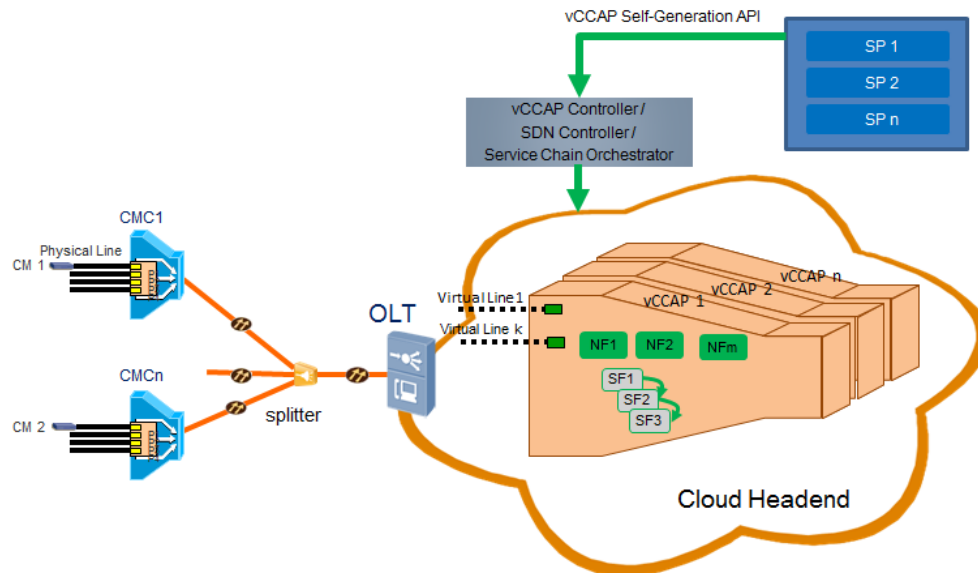


Figure 7 - Intelligent vCCAP Generation

According to business needs, the service provider presents the abstracted requirements through the vCCAP Self-Generation API. For example, a service provider that offers mobile services, requests a vCCAP to provide mobile services at a cost that does not exceed X dollars. Another service provider that offers broadband services, requests a vCCAP to provide broadband services at a cost that does not exceed Y dollars.

When the control system that consists of the vCCAP controller, SDN controller, and Service Chain Orchestrator receives the above requirements, it analyzes these abstracted requirements intelligently, and translates the abstracted requirements into network provision requirements.

In the case where vCCAPs are created for mobile services, the control system may allocate a time synchronization function, protection function and guaranteed bandwidth resources, such that, the total cost of all network resources is no higher than X dollars. On the other hand, for a vCCAP that provides broadband services, it may allocate resources for IGMP multicast network function, best-effort bandwidth resources and parental control service function, such that, the total cost of all network resources is no higher than Y dollars.

The cost of the network resources is determined by a QSN model, which consists of three items, namely, QoS, service function, and network function, as depicted in Figure 8.

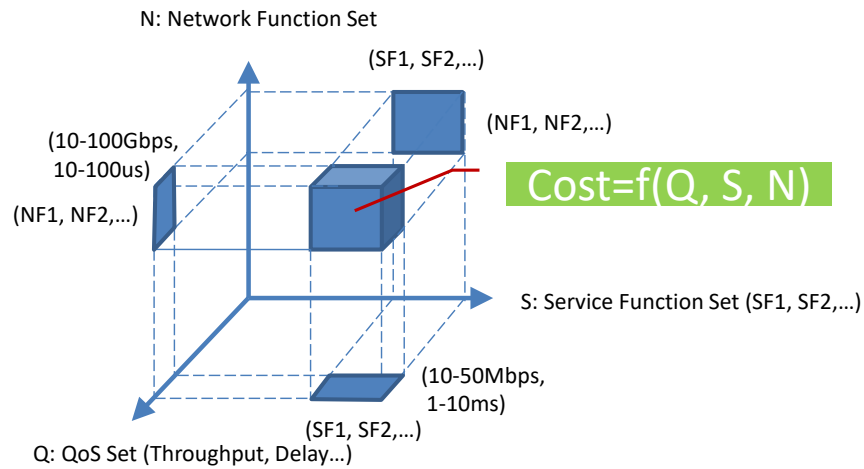


Figure 8 - QSN Model

To be more concrete, Figure 9 shows the principle of the intelligent vCCAP self-generator.

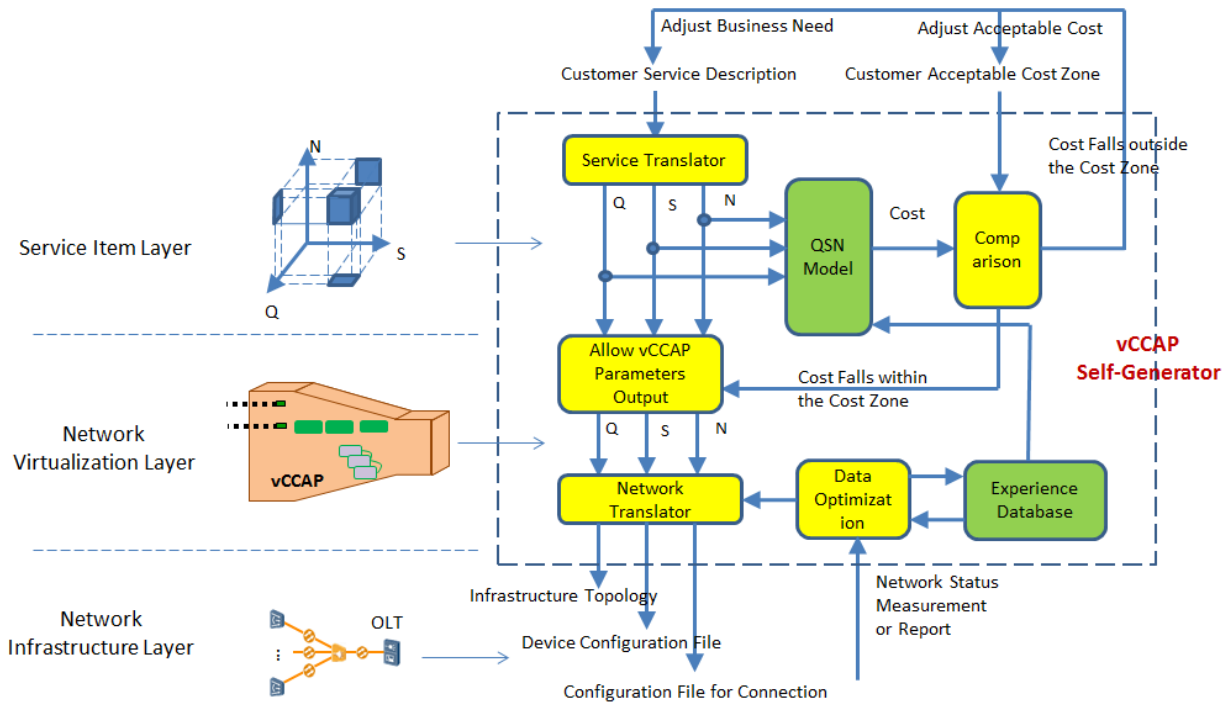


Figure 9 - Intelligent vCCAP Self-Generator Principle

A service translator translates the customer service description into vCCAP QoS (Q) requirements, service functions (S) and network functions (N). Based on the three items, a total cost for the deployment of the vCCAP is obtained by the QSN model, which is then compared with the customer's acceptable cost zone. The service provider can adjust the business needs or the acceptable cost zone to achieve a final agreement if it is necessary. Once the cost is acceptable, a vCCAP with the necessary virtual resources will be generated automatically, and a network translator will translate the QSN parameters for

infrastructure layer deployment. The infrastructure can use pure physical devices, or a mix of physical devices and virtual machines/containers running on a server. In order to enhance the intelligence of the service and network translation, a data optimization module is introduced, which uses machine learning technology and collects data from an experience database, and network status measurements or reports.

With vCCAP intelligent self-generation, the infrastructure provider can be more focused on the network infrastructure and build more functional networks to meet the needs of a richer business, while simplifying the service provider's network knowledge; the service provider's business needs can be automatically converted into network deployment.

The Value Proposition of Programmable vCCAP

The cable access network will benefit a lot with network virtualization. The key values that Programmable vCCAP can bring are summarized as follows.

7. Simple, Green and OPEX reduction

Being re-architected as a data center, the Headend can have a common infrastructure with commodity building blocks. Since the intelligence of remote nodes, e.g. complex control and management plane, is relocated to the vCCAP controller, the remote nodes can be simple and green devices without the need to support complex protocols, especially in the scenario of IoT and 5G.

Furthermore, the operator's configuration and management systems usually only need to maintain the IP address for each vCCAP instead of the IP addresses for all remote nodes. The remote node is either IP address-free or gets an auto-generated management IP address. The IP address configuration and maintenance can be simplified.

Plug and play of remote nodes is also enabled. When a network provider wants to deploy a new remote node, e.g. CMC, in the cable access network, the engineer on site installs a remote node, powers it up and connects its uplink, which comes from Headend or Hub. There is no need for the engineer to configure any parameters on the remote node, which automatically registers and connects to the vCCAP Controller. Therefore, the process of remote node installation and network configuration are simplified, and OPEX is further reduced.

8. Smooth Migration

Future access networks will have to support a whole set of technologies including DOCSIS, PON, active Ethernet, WiFi, Long-Term Evolution (LTE), 5G etcetera. Access network virtualization assures the integrity of multi-service, multi-access, diverse customized modes into a unified access network. This allows the network engineers to ignore the changes in access network topology, leading to a smooth migration from current network to a virtual access network.

Access network virtualization can abstract the common model of DOCSIS management and service provisioning. In this way, the architectural differences of traditional integrated CCAPs and distributed CCAPs can be isolated. MSOs can architect the Headend/Hub with commodity software and hardware building blocks without the need for dedicated hardware at the CCAP Core, and only leave DOCSIS MAC and PHY at remote nodes. Thus, it enables the smooth migration and transformation from Communication Technology (CT) to Internet Technology (IT) solution in cable systems. With access

virtualization technologies, coexistence and smooth migration with respect to the traditional integrated CCAP and distributed CCAP can be enabled, without impacting the MSO's configuration and management systems.

9. Network Innovation Acceleration

Thanks to the access virtualization, the remote nodes are reduced to mere programmable devices and thus decoupled from services, while the complex control and management planes are relocated to a centralized vCCAP. Remote nodes are only responsible for traffic forwarding, which minimizes complex protocol applications and configuration.

The vCCAP controller can help to build a future-proof access network that eliminates the need for hardware and software upgrades in remote nodes. With a programmable forwarding plane, the remote node is flexible and can easily cope with any subscriber session, while meeting requirements for Layer 2/Layer 3/MPLS forwarding, IPv4-to-IPv6 migration or future forwarding mechanisms. New services can be easily introduced into the cable access network, and therefore service innovation is accelerated.

10. Value-added: NaaS

Access network virtualization enables the migration from bitstream mode to a wholesale mode aka "virtual access network as a service". That means a physical cable access network can be virtualized into multiple virtual access networks, and each virtual access network can be wholesaled and be controlled and managed by third party retailers. It enables operators to break free from the simple, low-margin pipe wholesale models of the past through its support of differentiated access network offerings. In addition, MSOs can support the deployment of multi-service offerings in one physical access network, and meet the differentiated needs on the access network from different industries.

Conclusion

The rapid growth and the need for emerging requirements such as multi-access, multi-tenant and multi-service sharing in one physical network place a burden on legacy access network architectures, which find it hard to accommodate all these new technologies. In the cloud era, the access network will be transformed into a data-center based architecture. Network functions and services will run in the cloud. The Programmable vCCAP is a cloud solution with a future-proof access network architecture, which is expected to lead the way in access network evolution.

Abbreviations

5G	Fifth-generation
AAA	Authentication, Authorization and Accounting
API	Application Programming Interface
CCAP	Converged Cable Access Platform
CDN	Content Delivery Network
CMC	Cable Media Converter
COPS	Common Open Policy Service Protocol
CPE	customer premises equipment

CT	Communication Technology
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data-Over-Cable Service Interface Specifications
FTP	File Transfer Protocol
FTTDp	Fiber to the Distribution Point
ID	Identification
IEEE	Institute of Electrical and Electronics Engineers
IGMP	Internet Group Management Protocol
IoT	Internet of Things
IT	Internet Technology
LTE	Long-Term Evolution
MPLS	Multi-Protocol Label Switching
MSO	Multiple System Operator
NaaS	Network as a Service
NFV	Network Functions Virtualization
O&M	Operations and Maintenance
OAM	Operations, Administration, and Maintenance
OLT	optical line terminal
OMCI	ONU management and control interface
ONU	Optical Network Unit
OPEX	Operating Expense
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
SDN	Software Defined Network
SNMP	Simple Network Management Protocol
SP	service provider
TTM	Time to Market
vCCAP	Virtual CCAP
VR	Virtual Reality

Bibliography & References

OpenFlow Switch Specification Version 1.3.5, Open Networking Foundation,
<https://www.opennetworking.org/>.

Future access architecture: Software-defined access networking, Ruobin Zheng; Wenle Yang ; Jun Zhou, Consumer Communications and Networking Conference (CCNC), 2014 IEEE 11th, DOI: 10.1109/CCNC.2014.6940517, 2014 , Page(s): 881 – 886.

H-MPLS: A lightweight NFV-based MPLS solution in access network, Ruobin Zheng ; Wenle Yang, Consumer Communications and Networking Conference (CCNC), 2014 IEEE 11th, DOI: 10.1109/CCNC.2014.6940485, 2014 , Page(s): 887 - 892

G.984, Gigabit-capable Passive Optical Networks (GPON), ITU-T, 2003.

G.988, ONU management and control interface (OMCI) specification, ITU-T, 2010.

Technical Report, SDN Architecture for Cable Access Networks, Cablelabs, June 2015

Media Processing On Cloud

Scalable, Manageable, and Cost Effective

A Technical Paper prepared for SCTE•ISBE by

Srini Akkala

Senior Architect

ARRIS

37 Whitehall Way, Bellingham, MA 02019

301-529-5230

srini.akkala@arris.com

Introduction

Media processing is becoming a complicated problem. The dramatic increase in the number of viewing platforms is creating the need for multiple file renditions for each asset. Add multi digital rights management requirements into the mix and before long VOD libraries can easily expand from a few thousand to over hundreds of thousands of unique assets. This type of scale requires a new approach to business, infrastructure, and networking systems. Due to these added complexities of logic & scale, this new approach requires lot of planning and design in terms of storage, network bandwidth, and choice of tools & technologies.

On-Premises deployment starts with capacity planning (compute capacity, storage capacity, & network capacity). This model is not designed to handle unforeseen processing demands. Even with the best planning and forecasting, an on-premises solution can quickly run out of storage, network bandwidth, and/or processing power with the ever increasing asset library and growing customer base.

Given the ever changing video technology landscape, growth in consumer demand and content availability - setting up a cost effective and “future proof” media processing solution on-premises is often an unrealistic dream. Adapting to cloud based media processing solutions can make this dream a reality.

Solution

VOD processing involves different types of assets from multiple content providers publishing to varying types of screens. A typical media process workflow can be seen below. This paper discusses how each step can be designed to be scalable, manageable, and cost-effective by implementing a cloud-based solution.

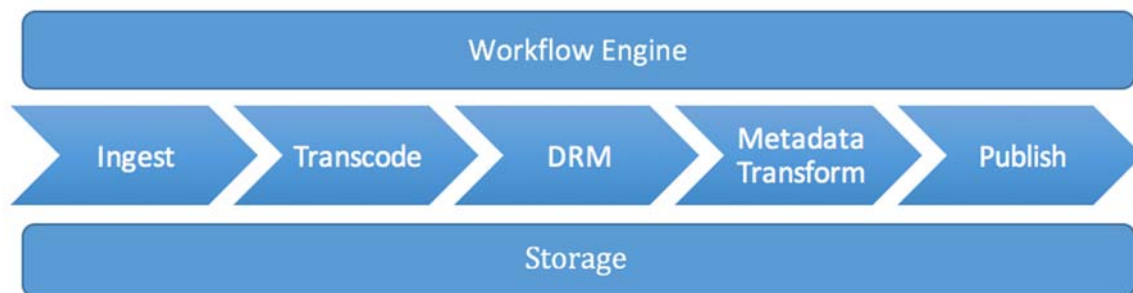


Figure 1 - Content processing workflow steps

The figure below illustrates how the above media processing steps are mapped to Amazon Elastic Compute Cloud (EC2) components.

AWS is taken as an example cloud computing platform, but the solution can easily be implemented in Google Cloud Platform or Microsoft Azure as both have dedicated media processing functionality built-in.

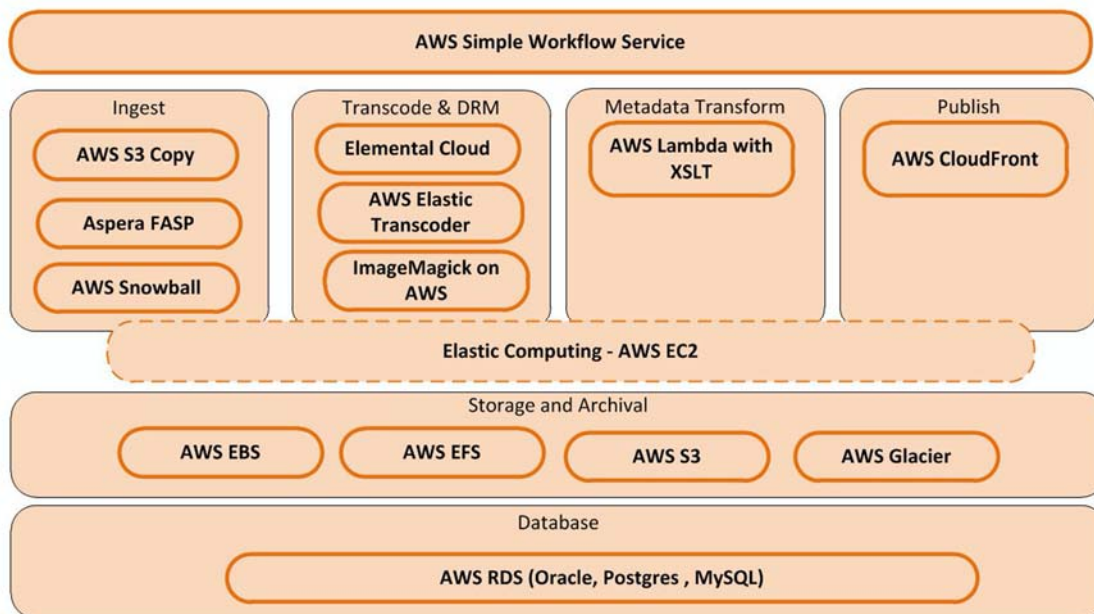


Figure 2 – AWS Components

Elastic Computing

One of the main benefits of cloud computing is elasticity. Elasticity is the ability to instantly grow in scale to meet the resources required during periods of peak or unforeseen load demands. The platform can be scaled up by adding more resources when demand increases and then removed as demand diminishes. Amazon’s EC2 service provides such an elastic computing capability.

One of the primary factors to be considered for compute capacity planning, in the context of VOD media processing, is number of assets required to be processed on a daily basis (e.g.: 200+ per day). Transcoding capacity is determined based on daily loads. On occasion, Cable providers run into situations which require the re-processing of their entire library within a defined period of time. For example, “*Offering 4K content*” requires re-processing of most of the library to generate content in 4k resolution. Accommodating re-processing requirements along with daily loads is not feasible with existing transcoding capacity. In this circumstance, one of the only options is to procure new transcoders to fulfill this kind of re-processing demand, which is not cost effective. Elastic computing enables you to procure (lease) extra transcoding capacity for a certain amount of time and shrink back to the original capacity after the re-process is completed.




Figure 3 – Transcoding Capacity Scaling

Storage and Archival

Media content management generally requires four unique types of storage, each with its own performance requirements.

- High performance local storage for media processing (transcode, segment)
- Shared network storage (e.g. catchers)
- Fast web object storage (e.g. image files, metadata files, file segments for http delivery)
- Archival storage (mezzanine media, backups)

Amazon Web Services offers 4 types of storage -

Amazon EBS provides block-level storage that serves as a virtual hard drive for your Amazon EC2 instance. Amazon EBS is designed for workloads that require persistent storage accessible by single EC2 instances. Content is copied from EFS or S3 to EBS for high speed content processing. Processed content is copied back to EFS or S3.

Amazon S3 is object storage designed to store and access any type of data over the Internet. It is secure, 99.999999999% durable, and scales past tens of trillions of objects. S3 buckets can be configured as catchers to where providers can pitch the content.

Amazon EFS provides simple, scalable file storage for use with Amazon EC2 instances. EFS data can be accessed from multiple AWS instances.

Amazon Glacier is an extremely low-cost and highly durable storage service for long-term backup and archive of any type of data. Source content is archived in Glacier after processing.




Figure 4 – AWS storage tiers

Database

Data storage requirements in the media industry are growing rapidly with each passing day (e.g.: content metadata, service usage data, viewership interaction events data etc.). Databases are getting bigger all the time. Running databases means lots of repetitive work (installation, configuration, administration, disaster recovery etc.). Amazon RDS makes it easy to set up, operate, and scale a database in the cloud. With Amazon RDS, there is no need to buy, rack, & stack hardware and no need to install software. Some of the most important benefits of using Amazon RDS include:

High Availability Amazon RDS is a highly available service which provides a SLA up-time of 99.95%

Scalability Amazon RDS offers two types of scalability features.

Vertical Scalability - Amazon enables push-button vertical scaling. This means that you can scale the size of an RDS instance [memory, CPU, PIOPS etc.] or disk, either up or down, with the click of a button.

Horizontal Scalability – The entire database is distributed across many RDS instances that will work together.

Backups: Amazon RDS offers Automated as well as Point-in-Time snapshot backups.

Available Amazon RDS engines : PostgreSQL, MySQL, Oracle, SQLServer.

Disaster Recovery

One of the main highlights of cloud computing is disaster recovery. With data centers in regions all around the world, AWS provides a set of cloud-based disaster recovery services that enable rapid recovery of infrastructure and data.

Workflow

AWS SWF (Simple WorkFlow service) makes it easy to build media work flows that coordinate operations across work steps involving distributed components. Coordinating tasks across the application involves a great deal of house keeping with respect to the logical flow of the application. Amazon SWF

gives you full control over implementing these tasks and coordinating them without worrying about underlying complexities such as tracking their progress and maintaining their state.

If we take an example of a typical media process flow - large videos are uploaded to Amazon S3 in segments. The upload of the segments has to be monitored. After a segment is uploaded, it is transcoded by downloading it to an Amazon EC2 instance. The encoded segment is stored to another Amazon S3 location. Failures could occur during this process due to one or more segments encountering encoding errors. Such failures need to be detected and handled through Amazon SWF's cloud workflow management.




Figure 5 – Workflow Engine

Some of the most important benefits of using Amazon SWF include:

- Amazon SWF replaces the complexity of custom-coded workflow solutions and process automation software with a fully managed cloud workflow web service.
- Amazon SWF lets you write your application components and coordination logic in virtually any programming language.
- Amazon SWF seamlessly scales with your application's usage.

Ingest

The first step in a media workflow is receiving the content package from a content provider and ingesting it into the system. AWS provides multiple tools for uploading large amounts of data into S3:

S3 Copy: Command line interface tool to transfer data to and from AWS.

Aspera High Speed Transfer Service: Industry leader in high speed file transfers. Direct-to-Cloud integration with S3 APIs ensures data is written directly to S3 storage during the transfer and then immediately made available when the transfer completes. Complete data protection with in-transit & at-rest encryption, data integrity verification for each transmitted block, and automatic retry and resume from point of interruption on failure are some of the key security features.

AWS Snowball: This service is applicable for one-time petabyte-scale data transfer scenarios, for example if you are planning to move an entire library from on-premises to the cloud or a provider is planning to transfer their entire library to the cloud.

Data is transferred into and out of AWS using physical storage appliances. Simply create a job in the AWS Management Console and a Snowball appliance will be automatically shipped to you. Once it arrives, attach the appliance to your local network and copy the library to the appliance. The data is encrypted and transferred to the appliance at high speed. Once the transfer is complete and the appliance is ready to be returned, the shipping label will automatically update and you can track the job status via text messages, or directly in the Console. The data typically appears in your S3 bucket in a couple of days.

Transcode & DRM

Transcoding & DRM can be done on cloud in two ways, by installing any software based transcoders on EC2 instances or by subscribing to cloud based transcoding services. (e.g. Elemental Cloud).

Elemental Cloud is a Platform as a Service (PaaS) built on AWS cloud Infrastructure. The platform automatically provisions and dynamically scales any combination of Elemental video processing, delivery, and storage services within a secure private network. Elemental cloud can automatically scale resources to process and deliver broadcast quality video as demand fluctuates.

Elemental can output multiple container formats include 3GPP, MP4, F4V, MPEG-TS, MOV etc.). Elemental supports adaptive streaming - Adobe's HTTP Dynamic Streaming (HDS), Apple's HTTP Live Streaming (HLS), MPEG-DASH (MP4, ISO, TS), and Microsoft's Smooth Streaming (ISMV). Elemental support content protection using multiple digital rights management (DRM) technologies

Sample Total Cost of Ownership (On-Premises transcoding Vs Elemental cloud) over a five-year period of time is illustrated in Appendix-A.

Other transcoding solutions on AWS include: Amazon Elastic Transcoder, ZenCoder, encoding.com, and harmonic.

Cover art resizing is one of the most important parts of the transcoding process. Content providers usually provide a single high resolution image as part of the content feed which then needs to be resized to fit various client devices. This step is not offered in most off-the-shelf transcoders. ImageMagick is one of most popular open source software products for image resizing. ImageMagick can be installed on an AWS EC2 instance. The Amazon market place also offers multiple SaaS based image re-sizing services which are based on ImageMagick.




Figure 6 – Image Resizing on AWS

Metadata Transformation

Source XML (ADI) metadata needs to be transformed into various formats based on the target device or storefront. Custom software (using XSLT) needs to be developed to achieve these types of transformations. Traditionally, custom code is developed and developers need to spin up VMs (EC2) and install the correct software stack prior to uploading and running code. AWS Lambda can be leveraged for these types of custom code modules.

AWS Lambda is a server-less environment in the AWS cloud. Lambda lets you run code without provisioning or managing servers. With Lambda, you can run code for virtually any type of application or backend service, all with zero administration. Just upload your code and Lambda takes care of virtually everything needed to run and scale your code with high availability. In addition, Lambda automatically scales your application by running code in response to each trigger. Your code runs in parallel and processes each trigger individually, scaling precisely with the size of the workload.

AWS Lambda supports code written in Java, C#, Python, and Node.js




Figure 7 – Metadata Transformation Using AWS Lambda

Publish and Deliver

The final steps of the workflow are publishing and delivery. Processed content is stored in an Amazon S3 object store (origin server). Amazon also offers a global content delivery network called **Amazon CloudFront**. If you store your objects in an Amazon S3 bucket, you can either have your users get your objects directly from S3 or you can configure CloudFront to get your objects from S3 and distribute them to your users.

Using CloudFront can be more cost effective for popular content (most frequently accessed). At higher usage, the price for the CloudFront data transfer is lower than the price for a comparable Amazon S3 data transfer. In addition, downloads are faster with CloudFront than with Amazon S3 alone because your objects are cached closer to your end users (media consumers).

CloudFront has out of the box streaming support for Smooth Streaming, HLS, MPEG-DASH, and RTMP protocols.

Sample content distribution (with offline packaging) is show below:




Figure 8 – Content Delivery on AWS Cloud

Conclusion

Broadcasters, content owners and other media providers are looking to take advantage of the capabilities of cloudbased media processing solutions. The primary intent is to expand/extend their existing platforms and create a unified infrastructure that is more flexible, scalable and future proof in its support for video processing and delivery. The main purpose for writing this paper is to assist these organizations with understanding the options and advantages of these technologies as well as with choosing appropriate implementation architecture.

Appendix A: Transcoding TCO

Appendix A compares the total cost of ownership (TCO) between a hypothetical on-premises solution and an AWS cloud-based solution.

Volume

Average number of assets per day	200
Average size of an asset (TV Show or Movie)	90 mins
Average size of an asset (@ 50 Mbps)	33 GB
Average growth per year	15%

On-Premises Solution

Storage

Total storage Year-1	200,000 GB
Total storage cost (Average \$1/GB) Year-1	\$200,000

Transcoding appliances (Elemental Server)

Number of appliances Year-1	8+8 (redundancy)
-----------------------------	------------------

Average server cost	\$50,000 per server
Average maintenance cost per server	\$4,000 per server
Transcoding cost Year-1	\$864,000

Staff

Number of IT employees for maintenance	1
Average IT employee annual cost	\$80,000

Cloud Solution (AWS)

Storage (S3)

Total storage Year-1	200,000 GB
Total storage cost (Average \$0.35/GB) Year-1	\$77,000

Elemental Cloud Transcoding (Built-in Redundancy)

Average transcoding price per hour	\$1.55
Total number of hours per year	109,500 hours
Transcoding cost Year-1	\$169,725

Staff

Number of IT employees for maintenance	0.5
Average IT employee annual cost	\$40,000

*All amounts are estimates only.

Cost Analysis

On-Premises 5-year total cost	\$2,723,559
Cloud 5-year cost	\$1,933,209




Figure 9 – TCO Analysis

Table 1 – On-Premises vs Cloud Yearly Costs

<i>On-Premises</i>	<i>Year-1</i>	<i>Year-2</i>	<i>Year-3</i>	<i>Year-4</i>	<i>Year-5</i>
<i>CAPEX</i>	\$1,000,000	\$152,400	\$172,860	\$198,789	\$228,607
<i>OPEX</i>	\$144,000	\$165,600	\$190,440	\$219,006	\$251,857
<i>Total</i>	\$1,144,000	\$318,000	\$363,300	\$417,795	\$480,464
<i>Cloud</i>	<i>Year-1</i>	<i>Year-2</i>	<i>Year-3</i>	<i>Year-4</i>	<i>Year-5</i>
<i>CAPEX</i>	\$0	\$0	\$0	\$0	\$0
<i>OPEX</i>	\$286,725	\$329,734	\$379,194	\$436,073	\$501,484
<i>Total</i>	\$286,725	\$329,734	\$379,194	\$436,073	\$501,484

Abbreviations

AWS	Amazon Web Services
DRM	Digital Rights Management
EBS	Elastic Block Storage
EC2	Elastic Compute Cloud
EFS	Elastic File System
HLS	HTTP Live Streaming
PIOPS	Provisioned input/output operations per second
RDS	Relational Database Service
S3	Simple Storage Service
SWF	Simple Workflow Service
TCO	Total Cost Of Ownership
VOD	Video On Demand
XSLT	Extensible Stylesheet Language Transformations

Bibliography & References

Amazon, Documentation for AWS: <https://aws.amazon.com/documentation/>

Cloud Overlay (CLOVER)

Extending the Cloud Virtually

A Technical Paper prepared for SCTE•ISBE by

John Jason Brzozowski
Fellow
Comcast
Philadelphia, PA 19103
484-962-0060
john_brzozowski@comcast.com

Mark Brittingham, Comcast
Chris Luke, Comcast
Zheng Yin, Comcast

Introduction

The term “cloud” is practically a household word today. References to “cloud” have matured, in a relatively short period of time, from what was an abstract concept, to infrastructure and resources used widely by consumers. Adoption of cloud infrastructure has obviously surpassed the specialized usage by large enterprise or service provider adopters. While the use of the cloud has evolved, how adopters access and utilize it has remained largely unchanged. In fact, one could argue that classic, aging networking techniques remain pervasively used today to gain access to third party cloud (TPC) resources and infrastructure. The aggressive adoption of cloud technologies seems to be pushing the limits of traditional techniques, not to mention the associated business and cost models.

The approach described in this paper is one that has been developed to modernize how cloud adopters connect to and utilize modern TPC infrastructures. The objective of this approach, which we call the Cloud Overlay (and hereafter referenced as “CLOVER”), is to marry automation, modern networking techniques, and existing, well-known protocols to help maximize how applications and services are securely deployed to third party clouds. Further, many of the techniques outlined in this paper can be extended and utilized within an enterprise or service provider network to enhance how internal users leverage their own private clouds.

CLOVER sets out to leverage more deliberately the concepts of overlay and underlay networking to provide seamless connectivity to cloud resources that are both on and off network. Today, the line is blurred, perhaps even non-existent, between the concepts of overlay and underlay networking, mainly because they often follow the same layer 3 path. For clarification, an underlay is analogous to how a typical Virtual Private Network (VPN) functions, where the VPN connection is the underlay and the overlay includes corporate email or Intranet communications over the VPN. Further, for CLOVER it is essential to clearly and distinctly differentiate between a service or application interface, and a control interface for a host or collection of hosts that have been deployed into TPC provider infrastructure.

Background

Not all applications or services that run in the cloud are equal. Today’s technology landscape goes far beyond hosting a simple web server in the cloud for an e-commerce offering, or even a personal blog. Most applications and services are quite complex, requiring advanced functionality for storage, performance, security, and network connectivity. In fact, many cloud adopters use TPCs as an extension of their own internal infrastructures -- which means that in many cases, what is hosted “in the cloud” may not even be reachable or usable over the Internet, only internally. Conversely, there are many applications and services that utilize a hybrid model, where they, in fact, are intended to be reachable and usable over the Internet. However, doing so requires access to a complex maze of backend services that are rightfully locked down in secure, on network data centers that are not reachable over the Internet. It is this hybrid model that has introduced complexity, and, in some cases, costs into the TPC adoption equation. This model has simultaneously stressed traditional networking and security technology, while fueling innovation that greatly enhances how TPC infrastructures could be used by the bulk of adopters.

Traditional Cloud Models

The rapid proliferation of TPC offerings, coupled with the desire to minimize the burdens on network data center maintenance, provided much of the early motivation for adopters to migrate toward the use of third party clouds. The attraction of third-party clouds offered the following:

- The promise of a cost-optimized, financially attractive cloud infrastructure, at least in the early stages
- Improved deployment agility, efficiency, and velocity enabled by TPC automation
- The potential to minimize capital investments in new and/or aging on-network, specialized data centers

Organizations that own and operate their own networks, data centers, and cloud infrastructures, and have also embraced TPCs, likely have come to the realization that it's quite a complex juggling act to keep up with demand for the cloud and virtualization, while managing capacity, quenching the thirst for new technology, and minimizing impact and downtime. The same adopters have also likely come to the conclusion that the grass is not green(er) on the other side.

Reference Architecture

There are a multitude of connectivity models available today to enable connectivity to TPCs. Two primary examples will be referenced within this paper: the Full Service TPC, and the Dedicated Infrastructure TPC. Reference to these models will help punctuate some of the challenges associated with their expansion today and in the future, in particular as TPC usage continues to balloon. Specifically, the reference models are:

1.1. Full Service TPC

The Full service TPC is a common model that is most analogous to traditional hosting models, where all resources required by an adopter are deployed on a TPC provider's infrastructure. This includes everything from compute, storage, and possibly application- or service-specific data. The adopter, in this case, effectively outsources most, if not all, infrastructure related activities to the TPC provider.

1.2. Dedicated Infrastructure TPC

Dedicated infrastructure TPC is a model that has grown in popularity. In it, the adopter and the TPC provider establish and maintain dedicated connectivity to the TPC infrastructure. Effectively, the adopter treats the TPC infrastructure as an extension of its own network, data center, or cloud infrastructure. Most TPC providers have an offering of this type.

Observations

Both traditional models carry distinct and immediate benefits. In the Full Service model, the potential to maximize the economies of scale and cost reductions afforded to the TPC provider are by far the most attractive attributes -- providing that the cost savings are at least in part passed on to the adopter or customer of the TPC. Delegating day-to-day operational responsibilities related to managing downtime, releasing upgrades, and augmenting capacity specific to the underlying TPC hardware all yield direct benefits to adopters of this model. Infrastructure delegation does not include responsibilities associated

with the applications or services that are running within the TPC environments, which is deliberate and often viewed as beneficial. All of these responsibilities remain with the application or service owner which is typically the customer of the TPC provider. With these models, adopters are able to focus primarily on application or service excellence and agility. Delegating responsibility related to the underlying infrastructure often allows for the reallocation of time, energy, and budget to application or service innovation and development.

While there are many bona fide benefits, there are also a number of considerations that must still be considered when considering one or more of the TPC models that are available today. The most obvious is relate to “shared fate” and the consequences of the inability to migrate between TPCs, or to use multiple TPC providers simultaneously. Particularly for the Full Service model adopter, if its TPC provider experience issues with their infrastructure (which have been rare to date), those issues typically impact large populations of the TPC customer base. A single TPC adopter is rarely impacted in an isolated manner.

Additionally, many TPC providers have invested heavily in technologies specific to their respective platforms. In most cases, these investments were driven by customer or industry demand, which has truly benefited those adopters. However, those same investments and innovations typically make it difficult (if not impossible) to migrate from one TPC provider to another, or to use multiple TPCs simultaneously.

Further, there are cases -- with the Dedicated Infrastructure model, specifically -- where localization, performance, and redundancy are significantly affected by the resulting network topology. The investment and resources required by adopters of the dedicated TPC model often also requires dedicated access to capacity for a subset of TPC data centers. While offering numerous benefits, this model does have a potentially adverse impact on redundancy and localization -- which, in turn, impacts performance, and ultimately the consumer experience.

Finally, and not insignificantly, the capacity, reliability, and performance of the underlying network between the TPC provider network and the adopter network is critical, specifically for the dedicated infrastructure model. Unpredictability and fluctuations on either side can introduce significant instability and customer impact. Resource requirements in the form of capacity planning and management are essential to the successful use of the Dedicated TPC deployment model. In practice, not all adopters of TPCs have the required resources -- human and financial -- to consider the Dedicated Infrastructure deployment model.

Next-Generation Cloud

The “cloud” has clearly helped to fuel innovation and the deployment of new applications and services. As previously mentioned, this phenomenon helped to push technology to new limits -- and, in some cases, is or soon will be pushing some adjacent technologies past their breaking point. Cloud-based networking and connectivity are near the top of the list of areas where we, as a community of adopters, continue to use traditional techniques that need to be revisited and/or redesigned. CLOVER, our moniker for “Cloud Overlay,” is an outgrowth of those realizations, and aims to offers the following:

- Independence and flexibility across multiple third-party cloud infrastructures
- Automation
- Scale
- Distribution and localization

- Improved redundancy

CLOVER more clearly delineates between the “underlay” and “overlay” aspects of next generation, cloud-oriented applications and services. Today, developers and engineers blur the lines between the underlay and overlay, treating them both equally. While this has worked for years, this approach forces “baggage” to be carried from one chapter to the next, from an infrastructure engineering perspective. Applications and services that have been deployed into the cloud have distinct communication properties -- namely, the application or service interface, and the control interface. To date, all of these communications properties are generally managed as single flows, to and from resources in the cloud. By separating application and services flows from control flows, traffic and connectivity can be separately managed. This allows TPC adopters to maximize and more effectively utilize cloud resources, and the robust infrastructures being built and deployed by TPC providers.

Reference Architecture

CLOVER builds heavily on concepts pertaining to underlay and overlay connectivity and communications. To clarify, underlay connectivity is a term describing the encapsulation or transmission of application, service, or end-user communications using an alternate transport. Overlay communications are typically what is being carried or encapsulated, while the underlay, as expected, is the carrier. CLOVER expands on this by differentiating application and service communications properties from the control communications. The classic example is the common “storefront” web application. Such an application typically consists of the following:

1. One or more web server front-ends
2. One or more backend systems, for example, databases

In the web storefront example, a cluster of web servers lacking customer and product databases does not make for much of an experience or much of a storefront for that matter. Conversely, exposing a customer or product catalog database directly over the Internet is inconceivable without proper security and scale considerations. The two challenges jointly make the case for coupling an HTTP-based application or service interface -- in this case, care of the web servers -- with a control interface to the customer and product databases, via a structured query language (SQL).

The core of the CLOVER target architecture allows for a separation of control communications from application or service communications, with each being managed independently. Control properties often require that communications be secure, and often not available over the Internet-at-Large, leveraging well-known underlay networking concepts. Underlay communications indicate the use of an underlying transport, often Internet Protocol (IP)-based, that can encapsulate or carry others forms of overlay communications -- which are, incidentally, Internet Protocol-based communications as well. With the pervasiveness of Internet Protocol version 6 (IPv6) today, there is often a wide range of underlay and overlay combinations. Overlay communications may include one or more of the following:

- IPv6 only
- IPv4 only
- Dual Stack (IPv4 and IPv6 are both enabled and used simultaneously)

In conjunction with the overlay combinations above, underlay communications can also be deployed and operated in various versions of the Internet Protocol. Unlike overlay communications, underlay

communications need not operate in dual-stack mode, where both IPv4 and IPv6 are enabled and used simultaneously. Underlay communications may, for example, prefer the use of IPv6 with only a fallback to IPv4, in the event that IPv6 is inoperable or is not supported. Notably, a fall back mode of operation is not analogous to dual-stack mode. Beyond the examples listed below, there are additional options for underlay communications, however, addressing these is out of scope for this paper. Further, there are several transport modes or protocols that can be used in conjunction with Internet Protocol to carry overlay Internet Protocol communications. Each of the below offer varying levels of security, including both authorization and encryption. Typically, those that are the most secure often have the greatest overhead, which, in turn, introduces the possibility of performance impacts. Underlay transport modes and protocol examples include:

- IP-in-IP¹
- Generic Route Encapsulation (GRE)²
- Internet Protocol Security (IPsec)³
- Virtual Extensible LAN (VXLAN)⁴

The development documented in this paper focuses on IPsec over IPv4 only and IPv6 only with varying levels of authorization and encryption. For simplicity, during development, pre-shared keys were used for IPsec authorization. However, for a production deployment of CLOVER-based solutions it is recommended that certificates be utilized minimally for IPsec authorization, mainly for simplifications related to automated resource creation and deployment. Several algorithms representing a wide range of encryption levels were used during development, including:

- Null or no encryption⁵ – effectively no payload encryption
- AES-GCM256⁶ – moderate encryption
- AES-256⁶ – better encryption

Several additional algorithms for an IPsec payload encryption are available for use, but for the purposes of CLOVER development, it was of primary importance to determine the performance characteristics of good/better/best levels of encryption. Adopters of CLOVER are likely to select authorization techniques and encryption algorithms that are best aligned with their infrastructures.

1.3. Key Components

The following section outlines the key elements of a CLOVER-based solution. Each component of a CLOVER-based deployment can be implemented and engineered in a manner that best suits the adopter. For the purpose of this work, each component and its function is described autonomously, such that readers can determine independently the best utilization within their own infrastructures.

¹ https://en.wikipedia.org/wiki/IP_in_IP

² https://en.wikipedia.org/wiki/Generic_Routing_Encapsulation

³ <https://en.wikipedia.org/wiki/IPsec>

⁴ https://en.wikipedia.org/wiki/Virtual_Extensible_LAN

⁵ <https://tools.ietf.org/html/rfc2410>

⁶ <https://tools.ietf.org/html/rfc4106>

1.3.1. Virtual Network Appliance (VNA)

A virtual network appliance, or VNA, is, in essence, a virtualized network function where a software module that typically runs on specialized hardware is built to run within a virtual machine, or more specifically as a virtual machine within a private, public, or third party cloud environment. As it relates to CLOVER, the primary function that is being virtualized is one that provides IPsec-authorized and encrypted communications. IPsec VNAs will typically terminate communications against another network element, which is likely to be a physical device that is redundant and fault tolerant. While it is possible, it is not strictly required for VNAs to terminate communications against common implementations, meaning that VNAs and aggregators (see next section) are intended to be fully interoperable. CLOVER aggregators are typically deployed in a centralized manner leveraging a hub and spoke⁷ or star topology⁸. The deployment model for CLOVER aggregation can be duplicated across a large network for increased capacity, improved performance, and localization.

The VNA is a primary communication path for all or a subset of network communications that are sent to and from hosts to where the VNA provides secure, network connectivity. As such, performance and IP transport implications are key considerations for deployment planning.

1.3.2. Aggregator

CLOVER aggregators are typically sets of redundant, high performance network elements that terminate secure network underlay communications for CLOVER virtualized networks functions like VNAs. While it is conceivable that aggregators too can be virtualized network functions, it is common for elements of this type to be dedicated hardware elements. Given the performance of modern computing platforms, commodity hardware is a valid consideration, from a platform point of view. Specifically, commodity computing platforms with high speed network interfaces running open source operating systems, like Linux or BSD derivatives, are legitimate alternatives to specialized, commercial hardware platforms. For example, Vector Packet Processing (VPP) under the Fast Data Plane Project (fd.io) supports packet forwarding on commodity hardware that is comparable to many commercially available platforms.

While there are capital and operational expenditures associated with specialized hardware platforms, open source software and commodity computing platforms are not without their own procurement and maintenance costs. The choice from an aggregator point of view truly boils down to adopter preference and capabilities.

1.3.3. Hosts

CLOVER hosts are the simplest element in the system. Hosts are just that, hosts, either bare-metal or virtualized, that are used to run proprietary or commercial applications and services. CLOVER is simply the mechanism by which network connectivity and communications are provided to them. Hosts can run on any operating system, and can essentially run a wide range of functions with a practically unlimited set of network configurations.

⁷ https://en.wikipedia.org/wiki/Spoke%E2%80%93hub_distribution_paradigm

⁸ https://en.wikipedia.org/wiki/Star_network

The origination of secure, dynamic underlay communications directly from hosts to aggregators, bypassing a VNA, is technically a valid mode of operation for CLOVER hosts, however, is out of scope of this document. This is future work and is in fact a valid construct in the context of a CLOVER system.

1.4. Communication Modes

Building on the definition of the basic components of a CLOVER system, the base communication modes illustrate how the elements of CLOVER can provide a flexible mix of communications paths that offer adopters alternatives compared to traditional cloud connectivity techniques -- private, public, or otherwise. The role of a CLOVER aggregator is to provide termination points for CLOVER VNAs into an area of a private network that may not be reachable over the Internet, or may intentionally be secure. CLOVER VNAs can be placed at multiple entry points in a serving network to provide granular, targeted access to different network segments. The VNA must obviously be able to terminate secure, underlay communications. In this case, the underlay described in this paper was built using IPsec. Operational, virtualized hosts in a third-party cloud environment can in turn utilize the VNA and the active, secure connection back to and through a CLOVER aggregator, over the Internet, for all or a subset of communications. The type, destination, or traffic source are governed by the chosen CLOVER connectivity model. CLOVER provides the flexibility to select a communication model that allows adopters to manage exactly how their applications and services are communicated with, including options to “bring your own” IP addressing.

1.4.1. Converged Model

In a converged communications model for CLOVER, a VNA is used to manage all host-based communications, regardless of IP version and other network communication properties. In a converged model, hosts in the cloud are not reachable through any other paths -- over the Internet or otherwise. From a management perspective, virtualized hosts may, in fact, remain reachable via virtual consoles, however, this is not entirely different than how physical hosts are managed today in a non-CLOVER environment.

The converged model is typically used in cases where an application or service is targeted for a cloud environment and is only required to be reachable from within an enterprise, and not over the Internet. The CLOVER converged model is analogous to a virtual data center connected back to a larger, centralized corporate network using an extension cord. Today, in many cloud environments, especially third party clouds, there are often options to establish direct connectivity to the provider. While these have some attractive properties, they can sometimes be costly, difficult to manage, and inflexible from a resiliency perspective.

In a converged model, all control, application, and service communications flow over the underlay via the VNA, since there are no alternate paths. As such, this does increase the throughput and bandwidth requirements for the VNA. Finally, increased volume of communication through the VNA can potentially decrease the overall quantity of hosts that can be served in this model. This in no way makes the model less usable or less desirable. It simply suggests that adopters must explicitly assess and profile applications and the respective deployment models that are best suited for a converged connectivity model. Access to a CLOVER-based converged cloud infrastructure allows for the expanded or extended use of cloud environments for applications or services that otherwise might be condemned to a life outside of the cloud.

1.4.2. Split Model

Unlike a converged model, a CLOVER “split” model introduces the notion of bifurcated communications. Essentially, control traffic for applications and services running on hosts are routed via the CLOVER VNA over the secure underlay, while actual application or service communications are routed or reachable via the Internet-facing addresses provided by the underlying cloud infrastructure. Again, the underlying cloud infrastructure could be private or public.

Analysis of control traffic for common environments suggests that the majority of communications to and from hosts are to support the currently running services or applications. As such, the use of split mode communications creates a positive imbalance in communications, such that control traffic capacity over the VNA can support a larger deployment of hosts. An assumption of 25-50% of communications to and from virtualized hosts for control communications represents an increase of 4 or 2 times, respectively, for a host’s application or service capacity. In essence, the bifurcation of communications allows for a calculated over-subscription of the CLOVER VNA for network bandwidth and throughput. It is this over-subscription that is a key enabler of increased performance. Further, the use of Internet-facing addressing, provided by the cloud infrastructure, allows significantly enhanced localization and redundancy.

While there are interesting benefits to the use of the CLOVER split communications models, there are also some notable considerations. This model assumes that control information is the minority traffic for network communications. If this is not the case, then many of the gains, from a performance and capacity perspective, will be lost. Additionally, security considerations will need to be closely evaluated specifically from a network and a host perspective. Hosts in a split model may effectively straddle the Internet and secure network segments. This is not unheard of, however, done inadvertently this could compromise security for the adopter.

Observations

Modern day use of the cloud seems to have outpaced network engineering, and in many cases, the underpinnings of network technology. Further, traditional networking in many ways seems to limit the possibilities that the cloud offers to current and future adopters. The concepts explored with CLOVER push network technologies to their limits as they are applied to the cloud. CLOVER advances the cloud and virtualization to truly incorporate the network -- to the extent that data centers (large and small) can now be fully virtualized, in the cloud, while allowing for connectivity back to their parent network to gain secure access to protected resources.

Additionally, the CLOVER architecture more effectively enables adopters to leverage multiple cloud infrastructures simultaneously, which can yield significant cost benefits while maximizing redundancy and reducing the surface area of risk associated with the use of a single cloud provider. Conversely, CLOVER does represent a wholesale paradigm shift around how network and cloud technologies interact. The lines between the two blur, or certainly have the opportunity to blur. How fast, how slow, or if at all are up to the adopter.

Finally, through the use of CLOVER, application and service owners alongside their network brethren will more explicitly dissect network communication. A keen, in-depth understanding of the communications is instrumental in identifying the best CLOVER models for a given application or service set, while optimizing for performance, cost, and efficiency.

Deployment and Operational Considerations

Fully embracing virtualization and the cloud along with the enormous opportunities they represent, also introduces a significant opportunity for other changes. Change in the underlying technology and infrastructure, like CLOVER, can lead to changes in the scale and velocity around how networks, clouds, applications, and services are engineered and deployed.

Automation

CLOVER was built to be fully automated, leveraging the atomic, programmatic building blocks offered by the wide range of cloud technologies and platforms available today. A critical aspect of CLOVER is the automated creation, deployment, and enablement for all of the key elements of the CLOVER system, including:

- Creation of the CLOVER VNAs
- Bi-directional provisioning and licensing CLOVER VNAs with the CLOVER aggregator, including the communication mode (split or converged)
- Automated creation or cloning of virtualized hosts

It is possible, and perhaps desirable, initially, for existing functional areas with an organization to own and manage their respective tasks as noted above. However, it is conceivable that application and service owners or end-users can and will fully automate the creation of cloud-powered virtual infrastructure, end-to-end, including but not limited to the provisioning of the underlying network to enabling users to generate their own configurations. Processes and deployments that currently take days or weeks can now be completed in minutes or seconds. Generally speaking, automation is most often referred to in the context of creation, but decommissioning and resource recovery is equally as important to ensure that antiquated or defective technologies are managed and updated accordingly.

Monitoring and Telemetry

Monitoring and telemetry are critical for physical infrastructure that is largely fixed. The dynamic nature of the cloud and virtualized infrastructure makes this exponentially more important -- especially in a CLOVER-like model, where every element can automatically be created, decommissioned, or moved in a moment's notice. Automated management of CLOVER resources must include dynamic enablement and population of monitoring systems. Further, with such a dynamic virtual environment, it is essential for adopters to re-think their deployment models. Specifically, the opportunity to distribute applications or services more widely into larger quantities of smaller clusters becomes a reality, and in some cases is highly desirable.

Conclusion

CLOVER, or Cloud Overlay, was born out of real operational, technological, and commercial challenges associated with the aggressive use of and deployment into TPC infrastructures. As described, the use of the cloud and virtualization, and the associated business models, will continue to test the limits of the underlying network technologies.

Embracing virtualization to include network functions, coupled with automation, enables infrastructure engineers to keep pace with the evolution of the cloud -- and, in some cases, to fueling innovation. This applies to the use of third party clouds as much as it applies to next-generation private or on premise cloud infrastructures. The virtual landscape that is in front of us now lays the foundation for end-to-end automation and simplification, truly enabling adopters to streamline how the cloud is employed to power their business, products, services, and people.

Cloud utilization will continue to grow, from a sheer scale perspective, as will the related verticals, including the Internet of Things (IoT). This massive horizontal and vertical growth of the cloud will continue to drive innovation, potentially pushing architectures like CLOVER to quickly move beyond concepts like virtual network appliances, to host-based CLOVER models -- where hosts are dynamically and securely negotiating underlay communications paths. The possibility (likelihood) of host-based models thrusts open the doors of innovation, driving CLOVER to evolve to utilize technologies like IPv6 Segment Routing to statelessly and programmatically establish secure, redundant underlay communications over IPv6-only networks.

Abbreviations

GRE	Generic Routing Encapsulation
IETF	Internet Engineering Task Force
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISBE	International Society of Broadband Experts
IPsec	Internet protocol security
RFC	Request for Comment
SCTE	Society of Cable Telecommunications Engineers
SR	Segment Routing
TCP/IP	Transmission Control Protocol/Internet Protocol
TPC	Third Party Cloud

Bibliography & References

Generic Routing Encapsulation (GRE), RFC2784; Internet Engineering Task Force

Security Architecture for the Internet Protocol (IPsec), RFC4301; Internet Engineering Task Force

Internet Protocol, Version 6 (IPv6) Specification (IPv6); RFC8200; Internet Engineering Task Force

IP in IP Tunneling (IP-in-IP); RFC1853; Internet Engineering Task Force

Cloud-DVR Real-Time Splunk-Based Monitoring and Alerting System

A Technical Paper Prepared for SCTE•ISBE by

Shlomo Ovadia, Ph.D.

Director of CPE Engineering
Advanced Engineering, Charter Communications
14810 Grassland Drive, Englewood, CO 80112
720-279-2875
Shlomo.ovadia@charter.com

Jenson Thottian

Splunk/Dev-Ops Architect
Charter Communications
6380 S. Fiddlers Green Cir, Greenwood Village, CO 80111
720-721-2875
c-jenson.thottian@charter.com

Introduction

Cable operators are responding to their subscribers' insatiable appetite for TV programming with new innovative video solutions such as Cloud-DVR (cDVR) and TV Everywhere [1, 2]. cDVR solutions move in-home recording and playout functions to the cloud, and thus enabling remote access of the recorded Linear TV content on a variety of platforms such as laptops, tablets, cellphone, and TVs. cDVR has many advantages such as cost, performance, operation, and business intelligence compared with in-home DVR. For example, no truck-rolls are needed to deploy and fix cDVR issues. The cDVR service offers a virtually unlimited number of tuners, enabling customer to record more than two shows at a time, and providing completely scalable and redundant storage capability. This feature allows customers to easily increase the amount of paid storage without any changes to their home network. Furthermore, cDVR based solutions allow cable operators to deploy faster, and to use more cost-effective CPE devices such as Charter's Worldbox 2.0 with traditional and cloud-based interfaces [3].

However, unlike the in-home DVR use case in which all the time-shifted recorded content is served locally, and thus has no impact on the cable network, cDVR has several infrastructure performance costs to consider. The cDVR use case requires network capacity on the cable access network, and storage capacity for the recorded content. The cDVR solution utilizes unicast video delivery for each subscriber, the total required network capacity is proportional to the number of concurrent cDVR subscribers viewing the time-shifted content. One of the potential cDVR obstacles is the copyright challenge by content providers. In the Private Copy deployment model, the cDVR permits each subscriber to record, store and playback a private and unique copy of the selected content (e.g., Private Copy). This means the cDVR storage capacity is linearly proportional to the number of subscribers. Another deployment option for cDVR is the hybrid storage model where the Private-Copies are maintained for 3 days after the record-time [4]. If the Private Copies are not viewed by then, they can be deleted ("de-duplicated") by the cDVR system and saved only as a Shared-Copies (same recorded content is shared among multiple subscribers). Thus, in this model, the amount of required storage for each subscriber can be significantly reduced compare with the Private Copy model.

Cloud-DVR System Architecture

Figure 1 shows Charter's cDVR system architecture. The main building blocks are the Video Storage and Processing Platform (VSPP) storage nodes, VSPP Manager's nodes, Scheduler servers, Geo-fencing servers, and KUMO servers and software. There are 56 VSPP storage nodes organized in four virtual entities called PODs. Each POD clusters 14 storage nodes using Commercial Off-The-Shelf (COTS) servers interconnected via a LAN topology using 1Gigabit Ethernet (GbE) interface for management, and 10GbE full-meshed inter-connection networks using high-speed switches. All the nodes in the POD contribute their physical resources in terms of storage capacity, CPU power, ingest and streaming throughputs. The VSPP software stack includes virtualized Software-Defined Storage (SDS), which enables high IO performance, high-availability storage solution with seamless fault-tolerance and self-healing operation. The VSPP SDS provides distributed Redundant Array Independent Disk (RAID) 5 storage used for content redundancy as protection against data loss due to disk failures. In addition, each storage node is connected to high-speed

Leaf switch using a 40GbE link, while each Leaf switch is connected to a high-speed Spine switch using a 100GbE link.

All the storage nodes run a Linux-OS and VSPP software stack, and are responsible for the actual data ingest and processing (e.g., transcoding, packaging, Ad-insertion), storage and streaming of content. The VSPP node joins the multicast stream for each linear feed to be recorded. Recording can only be performed for streams that have been requested by at least one user. All the Adaptive Bit Rate (ABR) video streams are delivered to using IP unicast with bi-directional TCP/IP connection between the ABR client device and VSPP storage node. All ABR video profiles for video and audio are stored up to 12 Mbps aggregated throughput. The cDVR peak storage capacity is 13.8PB. The VSPP Manager controls and orchestrates the entire VSPP storage nodes' activities and flows. The Scheduler is responsible for scheduling all the recordings based on the latest received EPG ingest information. KUMO is Charter's abstraction-layer software running on dedicated COTS servers, which receives information logs from client devices and manages and executes the different transactional REST APIs among the Scheduler, client devices, and the other back-office services such as IPVS, security system, NNS, etc.

Since Charter network spans a large portion of the U.S.A, it divides the U.S.A into disjointed geographical areas. Each subscriber has a unique HomeID, and an array of geo-fencing feed tags. A logical channel has different market-level feeds, which are combinations of a unique feed identifier (UUID) and channel name, broadcasted on different geographical regions defined as a list of geographical identifiers (geoID). Not all the logical channels are available over the entire US geo areas. When recording a channel having multiple market-level feeds, the VSPP system must record the market-level feed that applies to the location of the subscriber, which is provided by the Geo-Fence service. There are two Geo-Fencing servers configured in active/passive architecture.

In addition, the VSPP platform includes a Diagnostic server, which is a centralized monitoring and diagnostic back-end server. It continuously collects and aggregates all the VSPP system health and service-level metrics, and sends SNMP traps and Splunk data logs to the cDVR Splunk real-time monitoring tool.

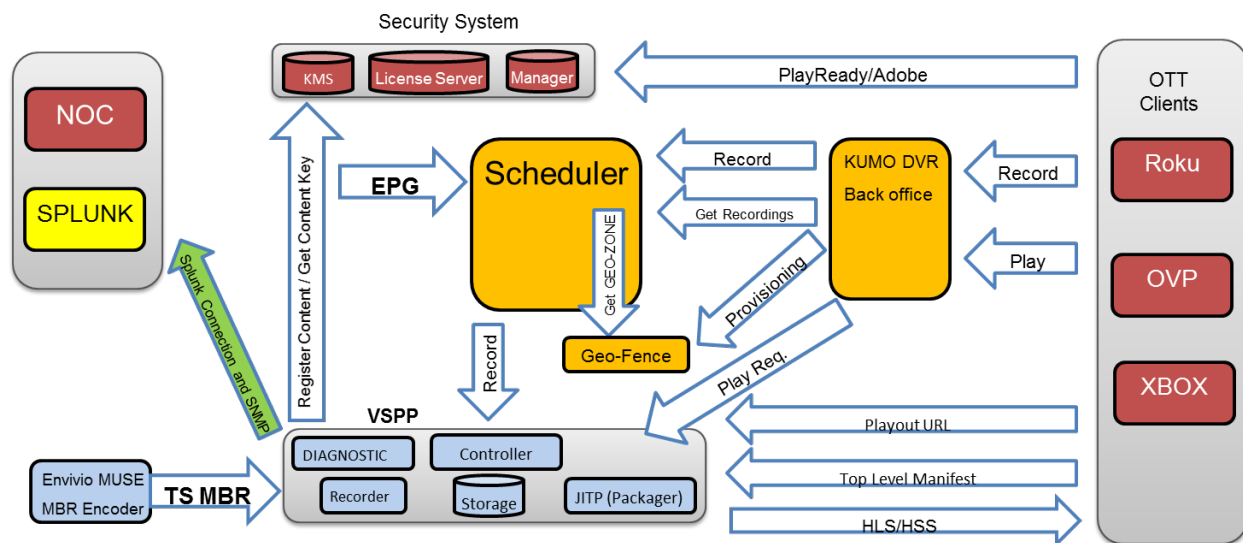


Figure 1 - Cloud-DVR System Architecture Showing the key Building Blocks and Interface to Splunk Environment

Real-Time Splunk-Based Monitoring & Alerting System

1. Monitoring Tool System Architecture

A Splunk-based real-time monitoring and diagnostics tool for cDVR service in production environment was developed. Figure 2 shows, for example, a high-level block diagram of the CDVR main dashboard with all the key components, including the Scheduler and Geo-Fence servers and apps, KUMO servers and apps, the VSPP nodes and apps, VSPP Manager, and Client devices. Two Splunk heavy-forwarders are forwarding VSPP metrics, application logs, Session Data Reports (SDRs), and SNMP traps to the Splunk indexer. In addition, KUMO servers' health metrics and application logs as well as KUMO API analytics, and application logs from Client devices such as OVP and Roku are being forwarded to the Splunk Indexers. All the background jobs are running Splunk ad-hoc searches, and sending all the collected metrics with in-house developed app to Charter's Splunk monitoring tool dashboard. Notice that the number of key components and metrics associated with each component is scalable, depending on the cDVR system complexity, to enable adequate video operation support. Furthermore, the main dashboard includes the following information:

- List of critical issues received in the last one hour
- List of SNMP traps received from the Diagnostics server in the last 4 hours according to their severity level (e.g., the highest-severity traps show-up first)
- Key cDVR plots such as total and used storage status
- cDVR metric definition and threshold levels
- Link to general cDVR analytics with a selectable-time period such as:
 - Number of daily/weekly users as reported by KUMO
 - Number of daily/weekly single and series recordings
 - Number of daily/weekly single and series playback recording
 - Number of daily/weekly recording failures
 - Number of daily/weekly deleted recordings
 - Top 10 watched channels

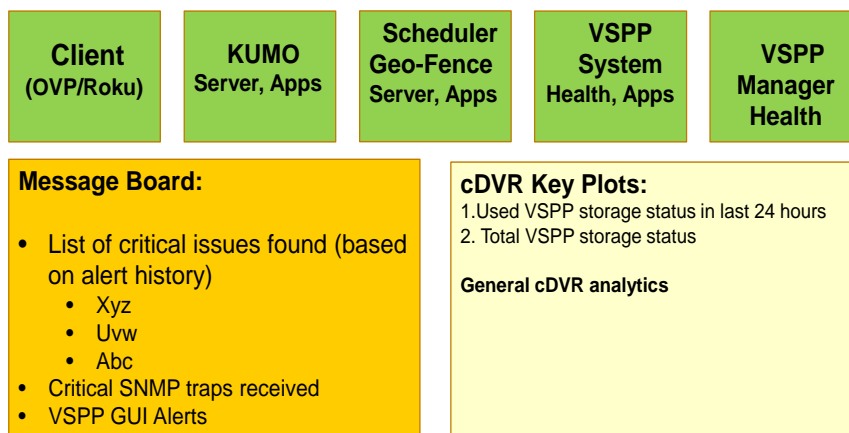


Figure 2: High-Level Block Diagram of Splunk-Based Monitoring Tool Dashboard

Figure 3 shows the main cDVR system health dashboard with its main subsystems, level 1 metrics for each of the subsystems, and their operational status. The listed subsystems and level 1 metrics are shown as an example, and other subsystems such as Arista switches and/or level 1 metrics are planned to be added. KUMO, which is a software abstract layer, receives information logs from client devices and manages and executes the different transactional REST APIs among the Scheduler, Geo-Fencing servers, client devices, and the other back-office services such as IPV5, NNS, etc. The VSPP Manager controls and orchestrates the entire VSPP nodes' activities and flows.

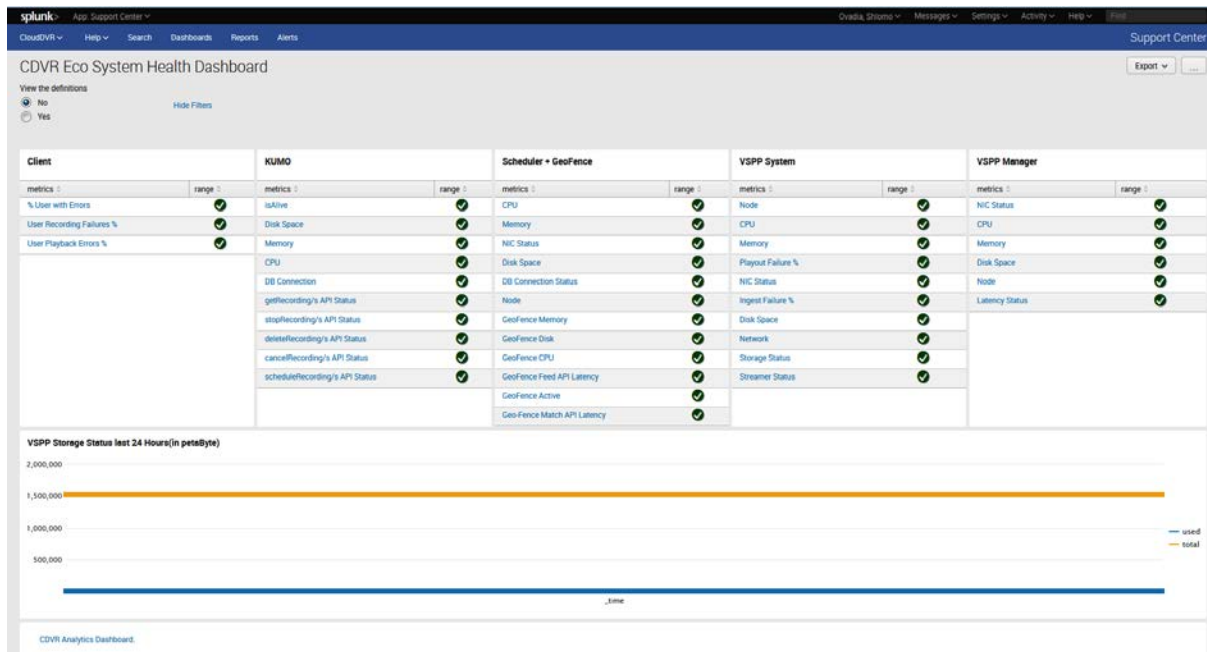


Figure 3: cDVR System Health Dashboard Showing All the Main Components, Metrics, and KPIs

Table 1 - provides a detailed description of the main cDVR dashboard components.

Table 1: cDVR Dashboard Component and Metric Description

cDVR Component	Description
VSPP Manager	<p>Provide real-time level 1 metrics on:</p> <ul style="list-style-type: none"> Health of the VSPP Manager – CPU and memory usage averaged over 15 minute period, most recent available disk capacity, and how long the server has been up and running. App analytics such as round-trip latency between the VSPP Manager and the selected storage node for the recorded sessions

	<ul style="list-style-type: none"> • Maintenance – how many storage nodes are going through maintenance operation such as software updates
VSPP System	<p>Provide real-time level 1 metrics on:</p> <ul style="list-style-type: none"> • Health of each of the 56 storage servers, including CPU and memory usage, most recent available storage disk space, and how long each server has been up and running • Network status – total number of errors while either receiving or transmitting packets • Storage Throughputs – disk write and read throughputs on each VSPP storage node • App analytics – relative performance of various app running on the VSPP system, including live ABR ingest and playout failures (%) in 60s period and round-trip latency of closed recording for each ABR session • Streamer status - provide activity status (active/not active) of the processes running on each of the storage nodes
Scheduler Server	<p>Provide real-time level 1 metrics on the health of the Scheduler servers, including CPU and memory usage averaged over 15 minute period, most recent disk space utilization, and how long the node has been up and running</p>
Geo-Fence Server	<p>Provide real-time metrics on the health of the Geo-Fencing servers, including CPU, memory usage, disk space, NIC status, node status, and the following API analytics:</p> <ul style="list-style-type: none"> • Geo-Fence Match API latency – round-trip latency between the Geo-fence server and its REDIS database • Geo-Fence Active – identify which Geo-fencing server is currently active • Geo-Fence Feed API latency – round-trip latency between Geo-fence server the KUMO server
KUMO Server	<p>Provide real-time level 1 metrics on:</p> <ul style="list-style-type: none"> • Health of the KUMO servers - CPU and memory usage averaged over the last 15 minute period, available disk capacity, and how long the server has been up and running. • DB and VSPP Connectivity – are the KUMO servers connected to the VSPP system and its own database • isAlive – Check if the KUMO app (Java) is alive and running • KUMO API analytics – obtain analytics such as unsuccessful transaction (%), error rate, transaction duration for various APIs: <ul style="list-style-type: none"> ○ getRecording(s) ○ stopRecording(s) ○ cancelRecording(s) ○ scheduleRecording(s) ○ deleteRecording(s)
Client	<p>Provide real-time level 1 metrics on user recording and playback errors (%) and users with errors (%) for various client devices such as OVP and Roku.</p>

2. Dashboard Features and Health Metrics & KPIs

The cDVR system dashboard health metrics and Key Performance Indicators (KPIs) are organized in a hierarchal fashion to facilitate the consumption of the vast amount of available information. When each of the level 1 health metrics and/or KPIs status changes to either a warning or critical condition (orange or red symbol), an operational engineer is able to obtain further detailed information as shown in Table 2.

Table 2: Hierarchal Organization of cDVR System Health Dashboard Metrics

Metric Level	Description
1	Metric status information based on pre-defined threshold levels for each CDVR subsystem
2	Metric status information such as the hostname of the server, its health status based on predefined threshold levels, name of transactional APIs and their status based on error count
3	Time-based behavior of the selected Level 2 metric
4	Event-based result showing the detailed information of the selected level 3 metric at a specific time.

Figure 4 shows an example of VSPP system health dashboard with level 2 metrics such as CPU and memory usage average over 15 minutes period, and disk space utilization on each of the displayed storage nodes. For each level 2 metric, the dashboard shows the hostname of the node, which POD it belongs to, and the corresponding metric health status based on pre-defined threshold levels. In addition, the storage disk utilization (in %) for each POD is displayed. If a storage node undergoes a maintenance operation or has failed, it will be indicated in the level 2 Node Status metric. The number of failed storage nodes is simply equal to 56 – (number of active nodes) – (number of nodes undergo maintenance).

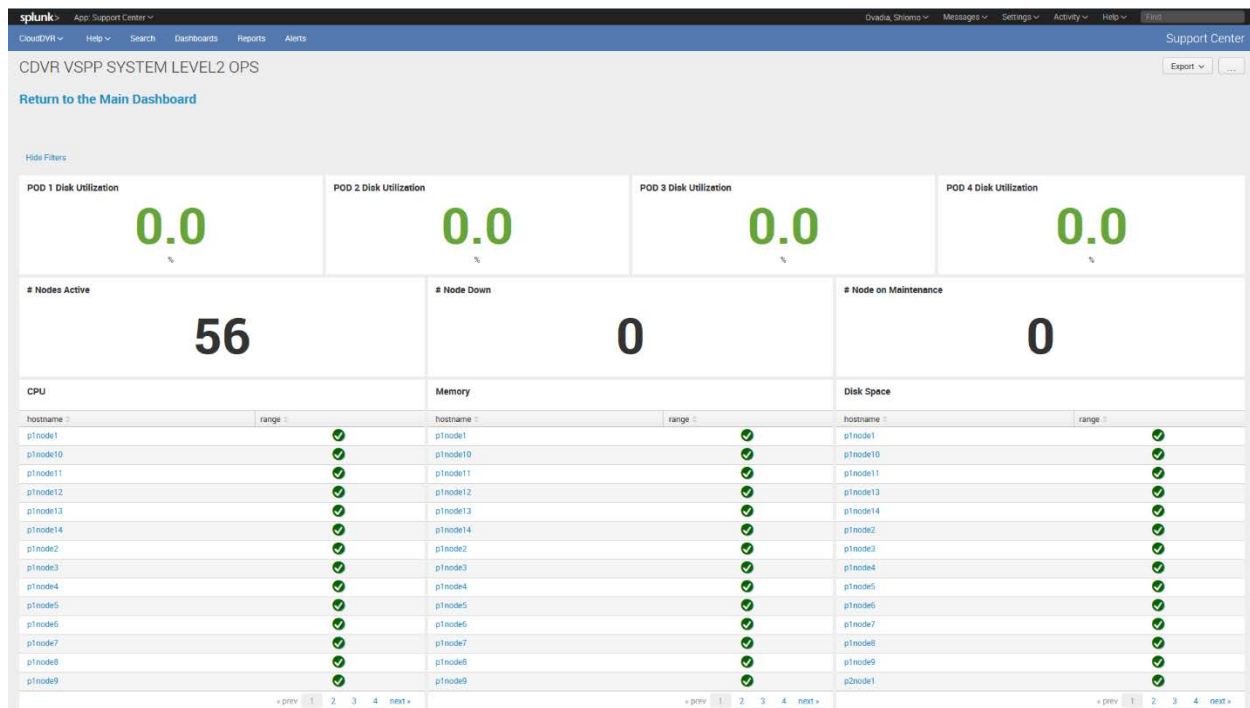


Figure 4: VSPP System Health showing POD Disk Utilization, and Level 2 Metrics such as Disk Space, CPU and Memory Usage for each Storage Node

Level 3 metrics provide the operational engineer with the level 2 metrics behavior versus time. Each level 2 metric can be filtered and displayed as follows:

- **Time-base filtering:** operational engineer can select the specific-time duration of interest to view the selected metric behavior
- **Hostname filtering:** operational engineer can select the specific hostnames to be displayed on the dashboard in the selected time frame.

Figure 5 shows, for example, the CPU and memory usage (%) of p1node12, p2node12, and p3node12 in the last 24 hours. This type of dashboard display is particularly useful for a time-dependent metric comparison among different storage nodes. If there are any storage nodes in maintenance, then the top part of the level 3 metric dashboard will show the hostname of each storage node in maintenance. Splunk query information is available by clicking on the observed trace for a specific metric on the selected hostname and time duration. Each of the dashboard for level 1, 2, and 3 metrics is updated every 15 minutes. Since there are so many concurrent background jobs, it takes two 15 minutes periods to update every dashboard metric. The 15 minute period was selected as a design trade-off between the number of concurrent background jobs for a dashboard update and the rate in which changes occur in the cDVR system.



Figure 5: CPU and Memory Usage (%) on p1node12, p2node12, and p3node 12 in the last 24 Hours

3. Server API Monitoring and Alerting

Various cDVR server transactional APIs such as from KUMO or Geo-Fencing servers are being monitored on the dashboard. Since the application logs are ingested within the Splunk environment, it allows the user to obtain valuable analytics about the monitored APIs. This includes detailed information about various transactions occurring in the VSPP system such as:

- Transaction count and duration in the specific time period
- Average transaction duration
- Transaction error rate (%) in the selected time period
- Percentage of unsuccessful transactions in the selected time period

Figure 6 shows, for example, the getRecordings API transaction status in the last 30 days. The getRecordings API transaction represents all series recordings requests sent by KUMO to the Scheduler. Table 3 summarizes the various reported metrics for getRecordings API. Notice that the getRecordings API transaction duration has large time variations, and 4.267 % of all the transactions were unsuccessful since these transactions timed out. The API transactions were timed out since their duration exceeded the one second threshold level. This programmable threshold level allows the cable operator to make sure that the cDVR system performance is within its design boundaries.

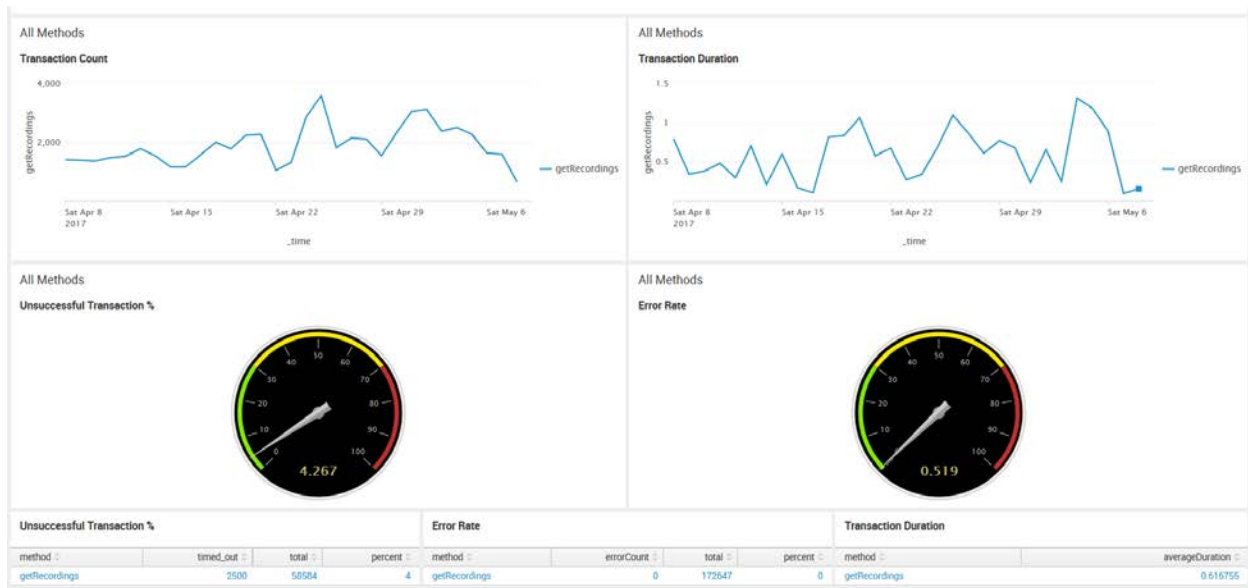


Figure 6: KUMO getRecordings API Status in the Last 30 Days

Table 3: Reported getRecordings API Transaction Metrics

Reported Metric	Reported Value
Peak API Transaction Count	3593
Average Transaction Duration	616.755ms
Transaction Error Rate	0.519%
Unsuccessful Transactions	4.267%

A key feature of the Splunk-based monitoring dashboard is the ability to change the threshold levels for each metric as needed. Programmable threshold levels are defined for each metric to indicate a healthy node condition, a warning or a critical condition. When the status of each metric changes to either a warning or critical condition, the operation engineer is able to obtain more detailed information via level 2 or level 3 metrics. This feature enables threshold-level adaptation of each metric based on cDVR system behavior in production environment.

Another key feature of the Splunk-based dashboard is the ability to send alert notifications and customized detailed reports via e-mail to a group of users when each of the KPI or level 1 health metrics crosses a critical threshold level. Specifically, e-mail notifications are sent when the status of any of the level 1 metrics or KPIs is changed from green to red. Furthermore, e-mail notifications will continuously be sent every hour until the red metric status is back to a healthy condition. With the addition of diagnostics capabilities, the operational engineer is able to take the necessary steps to resolve the observed issues.

4. SNMP Traps and Messages

The Splunk-based monitoring tool receives SNMPv3 traps and GUI messages from various cDVR components that are not shown in Figure 2. Table 4 lists other SNMP traps and messages that are received by the cDVR system dashboard. Both the health status of the high-speed network switches and the low-level server's hardware alerts information is available on level 2 of the VSPP system dashboard.

Table 4: Other SNMP Traps and Messages Received by cDVR System Dashboard

cDVR Component Name	Description	Provided Information
Network Switches	High-speed Spine and Leaf network switches connecting the VSPP storage and management servers	<ul style="list-style-type: none">• Deep insight and visualization of Spine and Leaf switches' health metrics• SNMP traps to identify any failures and performance degradation.
Server Hardware Alerts	Low-level storage and management servers' hardware alerts	Critical and major server's hardware fault events such as power supply, memory, storage disk, fan or smart battery array failures as defined via the vendor's SNMP MIBs.
VSPP GUI Messages	VSPP system syslog messages	VSPP syslog messages about various VSPP internal configuration setting changes listed according to their severity level and received time.

5. Diagnostics Features

A real-time list of new critical or warning issues are generated by cDVR system alerts and received SNMP traps based on pre-defined threshold levels for each metric. For rapid and scalable diagnostics of the observed issues, the monitoring tool needs to guide the operational engineer with an action plan in order to resolve the observed critical issues. Two primary diagnostics capabilities are proposed as follows:

A. Metric Diagnostics and Escalation Path

When the status of any level 1 metrics changes to either warning or critical, a new diagnostics screen becomes available. This new diagnostics screen is in addition to the level 2 metrics screen, and occurs when the operational engineer clicks on the selected level 1 metrics. The diagnostics screen provides the following capabilities:

- Suggested list of steps based on all the received system alerts and SNMP traps for the operational engineer to check in order to address the observed warning or critical level 1 metrics.
- Escalation path with contact info (e-mail address and phone #) of advanced engineers for further debugging of the observed issues.

B. Knowledge Based Diagnostics

Based on the video operation experience of the cDVR system, it is expected that some of the observed failures or critical issues may not be immediately solved by the operational engineers. A potential solution may be found based on the generated application error codes and their severity. When an observed failure, warning or critical issue is resolved, a detailed report may be generated with the suggested best engineering practices how to address such a future failure or critical issue. Depending on the Root-Cause Analysis (RCA) of the observed failures and the number of impacted customers, the threshold levels of some level 1 metrics may be changed, and/or new metrics maybe added. For example, after the Scheduler software was recently upgraded, it was found that the Scheduler had a stale EPG data, resulting in many failed recordings. Based on the vendor RCA, a new EPG ingest status metric that checks if the Scheduler has the latest EPG data is being added to the Scheduler dashboard.

Such knowledge-based reports are converted into a list of steps for the operational engineers to check before escalating to Tier 2 or 3 support engineers. Thus, the integration of this type of knowledge-based diagnostics capabilities into the monitoring tool can significantly reduce the time it takes the operational engineers to resolve new failures, warning or critical issues.

6. Self-Learning Monitoring Features

Self-learning monitoring capabilities are essential in order to continuously improve the monitoring tool. Three different types of self-learning capabilities are identified as follows:

1. **Tool health Check:**

The tool has a shell test script that periodically runs and checks if any background jobs are not running and reporting the assigned metrics. If the test script finds one or more such background jobs that didn't run, the test script performs the following steps:

- a. Check the status of all the background jobs that were run in a last period of time (i.e., 15 min.) according to their priority, and
- b. Rerun all the jobs whose test result was inactive (e.g., these jobs didn't run)
 - i. Check if this issue previously reoccurred with specific background jobs. If yes, the test script assigns these jobs a highest priority when executed. Test script monitors if no issue reoccurred, for example, in a four-week period, then the test script reduces the specific background job priority.

This feature allows the operational engineer to monitor that the health of the monitoring tool.

2. **Self-Optimization of Metric's Threshold Levels:**

By accessing historical Splunk logs and observing time-dependent behavior patterns of various metrics (level 3) in different dashboard components, the monitoring tool automatically reprogram the pre-defined initial threshold levels of the metrics, and provides an updated report on the updated threshold levels for each metrics. The tool provides a gradual incremental change in the pre-defined threshold to avoid false reporting of system health. In addition, this allows the operational engineer to reject some of the tool's reprogram threshold levels, and enter new threshold levels for the specific metrics and/or KPIs.

3. Operation Intelligence:

The VSPP Manager, which control and orchestrate all the activities in the cDVR system, monitors the health of each storage node as well as the apps that are running on the node. If, for example, a storage node starts to exhibit hardware failures or an abnormal behavior such as:

- One or more storage disk failures as received by an SNMP trap
- High-temperature inside a storage node as received by the iLO alerts
- Very high CPU or memory utilization (> 90%)

In this case, the monitoring tool sends a REST API request command to the VSPP Manager to take this storage node in a maintenance mode. The VSPP Manager put the specified storage node in a maintenance mode, and sends REST API acknowledgement to the tool. This allows the operational engineer to further debug the issue and take the appropriate action such as:

- Replace the failed storage disk
- Replace other parts within the server
- Reboot the node to check if the node is reporting healthy behavior, or
- Perform a scheduled software update/upgrade

After the hardware failure is fixed or the software upgrade is completed, the monitoring tool performs the following tasks:

- Checks the status of the following metrics and alerts:
 - All the reported level 2 metrics are healthy from this storage node
 - There are no SNMP alerts from the Diagnostics server
- If yes, it sends a REST API request to the VSPP Manager to take off the storage node from maintenance mode
- Generate a detailed report of the incident

The VSPP Manager takes the node off the maintenance mode, and completes the request by sending a REST API acknowledgement to the monitoring tool. By logging these cases, the monitoring tool can take similar actions if hardware or software failures re-occurred in other nodes.

Comparison with Other Monitoring Tools

There are other enterprise-level monitoring tools that can be used to monitor the cDVR system. Table 5 shows a high-level comparison between Splunk Enterprise [5], Graphite/Grafana [6], Nagios XI [7], and ELK [8] monitoring tools. The Splunk Enterprise is a flexible and scalable platform that makes it simple to collect and analyze vast amount of machine data, and act upon the received system alerts and SNMP traps. However, this is a proprietary monitoring solution that requires customer subscription. Fortunately, there are many vendors that already developed Splunk-based telemetry applications for their systems, which simplifies the integration of these applications into a Splunk-based dashboard.

Graphite is an open-source app for collecting, analyzing, and providing real-time monitoring of server health and application metrics for enterprise platforms. In addition, it offers a user-friendly graphical presentation of the data via the Grafana dashboard. However, the Grafana dashboard doesn't receive SNMP traps, and doesn't send e-mail notification or has an ability to set-up threshold levels for various metrics, which limits its usefulness in a large-scale video operation.

Table 5: Comparison between Splunk, Graphite, and Nagios XI Monitoring Tools

Monitoring Tool	Pros	Cons
SPLUNK>	<ul style="list-style-type: none"> Enterprise-class high availability Scalability Customized dashboard Interactive graphs Server and App metrics Ability to receive SNMP traps Ability to send e-mail notifications Generate customized technical reports Data logs retention & reporting Provide analytics for Nagios XI Specialized modules are available for security, IT services, and user behavior Easier to use logs for troubleshooting 	<ul style="list-style-type: none"> Proprietary Required customer subscription Potentially expensive
Graphite/ Grafana	<ul style="list-style-type: none"> Open-source tool Interactive graphs Scalability Low-cost subscription 	<ul style="list-style-type: none"> Doesn't receive SNMP traps No ability to set threshold-levels to dashboard metrics No alerts or e-mail notifications No or limited technical support Limited data logs retention No support for string metric values
Nagios XI	<ul style="list-style-type: none"> Open-source tool Customized dashboard Scalability Server and App metrics Ability to send e-mail notifications Generate technical report Data logs retention and reporting 	<ul style="list-style-type: none"> No ability to set threshold-levels to dashboard metrics No ability to receive SNMP traps GUI lacks user-friendliness Requires subscription for Enterprise-level tool (potentially expensive)
ELK	<ul style="list-style-type: none"> Open-source tool Scalability Customized dashboard Server and App metrics Data logs retention and reporting Alerts or e-mail notifications (with X-Pack) Generate technical report 	<ul style="list-style-type: none"> Missing user management features (in basic ELK) No SNMP traps (w/o using external modules) ELK cluster deployment requires more time & resources than Splunk Data onboarding is harder than Splunk Feature-poor UI compared with Splunk Only accept JSON-formatted data No specialized modules are available for security, IT services, etc.

Nagios XI provides monitoring of all mission-critical infrastructure components including applications, services, operating systems, network protocols, systems metrics, and network infrastructure. Hundreds of third-party add-ons provide for monitoring of virtually all in-house applications, services, and systems. Although Nagios is the best known free monitoring tool, its open-source version is limited in terms of dashboard features and capabilities, there is a steep and costly learning curve, and the GUI lacks user-friendliness. This is one of main reasons that the Nagios XI tool is losing its appeal among corporate customers.

ELK is an open-source monitoring tool that is gaining popularity among cooperate users similar to Splunk. It uses Elasticsearch for ingesting data logs, and Kibana as the visual UI for displaying customized dashboards. Enterprise users can purchase X-Pack, which is an Elastic Stack extension that bundles security, alerting, monitoring, reporting, and graph capabilities. X-Pack components are designed to work together seamlessly, allowing the user to enable or disable the desired features as needed. However, a larger ELK development effort for a customized monitoring solution than in Splunk may be needed, depending on the growth rate and complexity of deployment use cases. Another main difference is the way the data is parsed. ELK requires you to identify the data fields before it's shipped to Elasticsearch, while with Splunk, you can do that after the data is already in the system. This makes data onboarding easier by separating shipping and data classification/field labeling.

Summary

In this paper, the cDVR system architecture with its key hardware and software components was reviewed first. Then, the real-time Splunk-based cDVR monitoring system architecture, main features, health metrics, and KPIs were explained. This includes four-level hierarchal organization of server health metrics from VSPP storage nodes, VSPP Manager, Scheduler, Geo-Fencing and KUMO servers as well as all the transactional APIs' analytics from KUMO, Geo-Fencing, and VSPP Manager. The cDVR system dashboard includes the status of each health metric based on pre-defined programmable threshold levels for healthy (green), warning (amber), and critical (red) condition. When the status of each metric is changed from green to red, e-mail notifications and customized health reports are sent to the operational engineers to take the necessary actions to resolve the issues. A complete end-to-end system health report is achieved via the received SNMP traps by the cDVR dashboard from the high-speed network switches, low-level storage and management servers' hardware, and VSPP system syslog messages listed according to their severity level.

Another novel feature of the cDVR monitoring and alerting system is the self-learning capabilities such as the monitoring its own health. It periodically runs and checks if one or more background jobs are not running and reporting the assigned metrics. Furthermore, the monitoring system provides self-optimization of the threshold-levels for each of the monitored metrics and KPIs based on historically system behavior and observed failures. Another key aspect of the self-learning capabilities is the operation intelligence to identify low-level hardware failures such as a disk, fan, memory or smart array battery failures in a VSPP storage node, put the node in maintenance for further user diagnostics, and take the node off maintenance mode after either hardware repairs or software update tasks are completed.

Although the monitoring and alerting system cost is an important factor when comparing Splunk-based monitoring tool with other enterprise-class monitoring tools, the available features set, data logs ingest, scalability, user-friendly graphical interface, and time-to-market to develop a real-time operation-ready monitoring tool are also important considerations. When comparing the pros and cons of all these considerations (e.g., Table 5), the Splunk-based monitoring tool appears to be the most suitable for our

cDVR system. Furthermore, the Splunk Enterprise environment is conducive for integration with other Charter’s back-office applications such as IPVS, NNS, etc., resulting in faster video operation readiness.

Abbreviations

Table 6: Table of Abbreviations

Abbreviation	Stand For
ABR	Adaptive Bit Rate
API	Application Programming Interface
cDVR	Cloud Digital Video Recorder
CPE	Customer Premise Equipment
COTS	Commercial-Of-The-Shelf
CPU	Central Processing Unit
EPG	Electronic Program Guide
GbE	Gigabit Ethernet
GUI	Graphical User Interface
iLO	Integrated Lights Out
IPVS	IP Video Systems
KPI	Key Performance Indicator
MIB	Management Information Base
NNS	National Navigation Services
NOC	National Operation Center
OVP	Online Video Platform
RAID	Redundant Array of Independent Disks
RCA	Root Cause Analysis
REST	Representational State Transfer
RS-DVR	Remote Storage DVR
SDR	Session Data Report
SDS	Software-Defined Storage
SNMP	Simple Network Management Protocol
VSPP	Video Storage and Processing Platform

Bibliography & References

- [1] Carol Ansley and John Ulm, “The Dawn of Cloud-Based DVR Services”, SCTE Cable-Tec Expo, October (2013).
- [2] John Horrobin and Yoav Schreiber, “Unicast or Multicast for IP Video? Yes!", SCTE Cable-Tec Expo, October (2014).

- [3] <http://www.multichannel.com/news/content/charter-taps-arris-key-development-partner-worldbox-20/408379>
- [4] I. Tomer, Hybrid Solution for Cloud DVR: Meeting the Needs of Converging Legacy and OTT Platform, BEC proceedings (2016).
- [5] https://www.splunk.com/en_us/products/splunk-enterprise.html
- [6] Overview of Graphite tool can be found at <https://graphiteapp.org/#overview>
- [7] <https://www.nagios.com/products/nagios-xi/>
- [8] <https://www.elastic.co/webinars/introduction-elk-stack>

Interference Group Discovery for 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
tonliu@cisco.com

Introduction

In legacy DOCSIS, data can only be transmitted in one direction across any part of the spectrum. Compared to the passive optical networks (PONs), a cable access network is severely limited in the maximum symmetrical data speed due to the upstream RF spectrum scarcity. Since bringing fiber to the home is extremely expensive, cable operators have searched for an alternative to deliver the multi-gigabit services promised. This need together with recent trends in the cable industry (i.e. the deployment with DOCSIS 3.1 Orthogonal Frequency Division Multiplexing (OFDM); the deep fiber migration; and the remote PHY network architecture) has resulted in the rapid development and standardization of the full duplex (FDX) DOCSIS technology. With FDX DOCSIS, the RF spectrum can be used simultaneously in both the upstream (US) and downstream (DS) directions, allowing up to 5 Gbps US service and 10 Gbps DS service over the cable access network.

In FDX communications, a system supports simultaneous bi-directional transmissions across the same spectrum. Interferences between the bi-directional transmissions therefore must be mitigated for the intended signals to be properly received. DOCSIS is a point to multi-point system, where multiple cable modems (CMs) are connected to the same Cable Modem Termination System (CMTS) port via a coax distribution line. 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. Since the source of the interference is unknown to the receiving CM, PHY layer echo cancellation cannot be used. FDX DOCSIS address this issue by grouping CMs that interfere with each other into an Interference Group (IG). CMs in the same IG must transmit or receive along the same direction at any given frequency and time. CMs from different IGs have enough RF isolations to allow simultaneous US and DS transmissions at the same frequency.

In this paper, we will discuss IG discovery, a new process introduced in FDX DOCSIS to determine the IGs based on the CM to CM interference measurement obtained via sounding. We will start by introducing the basic IG concept and the operational principles to conduct sounding. We will examine the system overhead in terms of the spectrum cost and the time to converge for sounding among a given number of CMs at the desired frequency granularity. We will then propose a set of optimization techniques to improve sounding efficiency. We further extend the solution space by incorporating an iterative IG Discovery model to allow the system to automatically adapt to the changing network environment for optimized system performance.

IG Discovery Overview

1. Interference Groups

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 in a full duplex mode. This occurs when the co-channel interference (CCI) levels at the receiving CMs are above a design threshold when a CM is transmitting simultaneously over the same FDX spectrum.

FDX DOCSIS uses a sounding procedure to measure the CM to CM CCI. During Sounding, the CMTS selects one or more FDX capable CMs as test CMs to transmit test signals on designated subcarriers, while directing other FDX capable 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.

The measured CCI, in the form of the RxMERs collected from the measurer CMs, can then be used to sort CMs into IGs. Quantitatively, given a set of CMs, cm_1, cm_2, \dots, cm_N in a service group, cm_i 's IG group, $IG(cm_i)$, can be determined, such that,

$$\text{for any transmitting } cm_j \in IG(cm_i), RxMER_{ji} < \overline{MER} \quad (1)$$

or,

$$\text{for any transmitting } cm_j \notin IG(cm_i) \quad RxMER_{ji} > \overline{MER}; \quad (2)$$

Where, $RxMER_{ji}$ is the RxMER obtained at cm_i when cm_j is transmitting test signals. \overline{MER} is the threshold designed for $IG(cm_i)$, for its member CMs to properly demodulate a target modulation scheme.

Since the path loss of the interfering signal is reciprocal in a passive coax plant, symmetrical CCI is expected between a pair of CMs, therefore,

$$\text{if } cm_j \in IG(cm_i), \text{ then } cm_i \in IG(cm_j); \quad (3)$$

However, as the RxMERs are also impacted by the noise sourced internal to a CM, the RxMER level may not be the same. Sounding is thus required at both cm_i and cm_j to accurately detect the interference.

Figure 1 shows an IG Discovery example using the RxMER measurement data listed in Table 1. The shaded cells mark out the three IGs after applying a 35dB MER (or 10 bits/subcarrier) threshold, namely IG1 for CMs under Tap1, IG2 for CMs under Tap2, and IG3 for CMs under Tap3, Tap4 and Tap5.

From the example, we can observe the following:

1. Low MERs for CMs under the same tap; for example, the MER is 6dB for CMs under Tap1, as the RF path loss between the CMs under the same tap is much less compared to the inter-tap case.
2. Low MERs for CMs under the taps close to the end of distribution line; for example, CMs under Tap 3 through Tap5 all have MER below 35dB, due to the poor coupling loss of the lower-value taps.
3. Symmetrical CM-to-CM interference indicating reciprocal path loss of the passive plant.

Table 1 - CM-to-CM Interference and IG Formation

R E C E I V E	MER (dB)	TRANSMIT				
		Tap1	Tap2	Tap3	Tap4	Tap5
	Tap1	6	39.9	39.9	39.9	39.9
	Tap2	39.9	9.8	37.5	37.5	37.5
	Tap3	39.9	37.5	13.2	34.5	34.5
	Tap4	39.9	37.5	34.5	15.8	31
	Tap5	39.9	37.5	34.5	31	18.2

R E C E I V E	Mod Order	TRANSMIT				
		Tap1	Tap2	Tap3	Tap4	Tap5
	Tap1	-	11	11	11	11
	Tap2	11	-	10	10	10
	Tap3	11	10	-	9	9
	Tap4	11	10	9	-	8
	Tap5	11	10	9	8	-

DS output power: 39 dBmV/6 MHz
US Input power: 8 dBmV/6.4 MHz

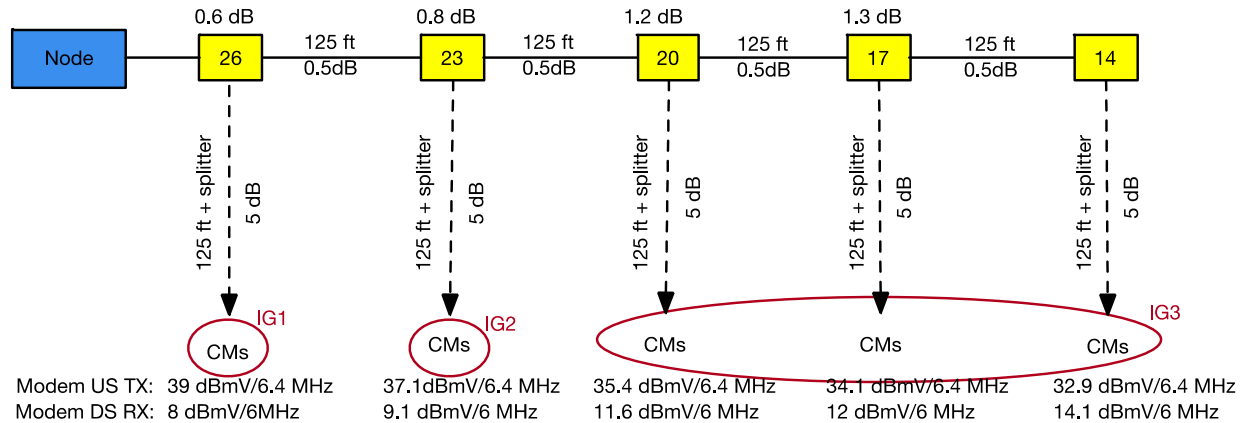


Figure 1 - CM Interference Groups over a Passive Coax Distribution Line

2. Sounding Techniques

There are two sounding methods proposed in FDX DOCSIS [3][4].

1. Sounding with OFDMA Upstream Data Profile (OUDP) test bursts
2. Sounding with continuous wave (CW) test signals

The OUDP method is intended for the deployment scenario where the legacy high-split DOCSIS 3.1 CMs, after necessary software upgrade, can share the US spectrum between 108 to 204 MHz with the FDX CMs. Since the DOCSIS 3.1 CMs cannot generate a multiplicity of CW tones as required in the CW sounding method, the DOCSIS 3.1 OUDP test bursts must be used instead as the test signals. When the OUDP test bursts are being transmitted by a test CM, other CMs that are capable to receive in this frequency band measure the RxMERs in the time and frequency encompassed by the continuous OUDP bursts. The OUDP test burst is intended to cover all DS subcarrier frequency locations by taking advantage of a faster RxMER measurement scheme to be implemented on the new FDX CMs.

The CW method is intended for the deployment scenario where the DOCSIS 3.1 CMs, after necessary software upgrade, can share the DS spectrum with FDX CMs. For example, a low-split or mid-split DOCSIS 3.1 CM can share the DS spectrum between 108 to 684 MHz, and a high-split DOCSIS 3.1 CM can share the DS spectrum between 258 to 684 MHz. During CW sounding, one or multiple FDX test CMs send CW test signals at selected DS subcarrier frequency locations, while the rest of CMs, including both legacy D3.1 CMs and FDX CMs measure the MER using the DOCSIS 3.1 RxMER measurement method.

3. Spectrum Overhead

A sounding test opportunity requires spectrum resource in time and frequency for both the US and DS directions. As shown in Figure 2, in the US direction, a test signal transmission opportunity is required for a test CM to send the test signals. In the DS direction, a test signal interference region is required to carry zero-bit-loaded symbols, to avoid any packet caused by the interference from the test signals.

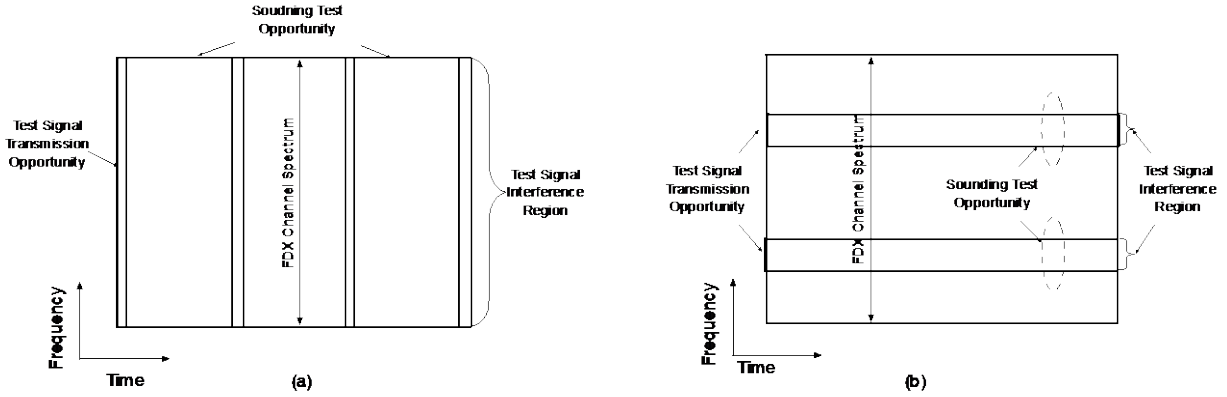


Figure 2 - Sounding Test Opportunities (a) OUDP Test Opportunities (b) CW Test Opportunities

For the OUDP sounding, a sounding test opportunity covers the entire FDX channel width in frequency and lasts about 20 to 60 milliseconds in time [4]. Thus, no spectrum can be used for traffic when the OUDP sounding burst is present on the FDX channel under test.

For the CW sounding, a sounding test opportunity includes a single CW subcarrier and a few guard subcarriers on both sides, to prevent inter-symbol interference at adjacent data subcarriers. Comparing to the OUDP sounding, a CW test opportunity occupies much narrower spectrum however lasts longer in time. It typically takes around 200 to 300 milliseconds for DOCSIS 3.1 RxMER measurement scheme to converge.

With the CW sounding, the CMTS has the option to limit the number of sounding test opportunities, so traffic can be sent using the data subcarriers outside the CW interference regions, particularly, the DS traffic to the measurer CMs, and the US traffic from a test CM if the test CM's IGs have been identified through previous sounding.

The spectrum overhead S_{avg} spent on sounding can thus be expressed as the percentage of the sounding dwell time multiplied by the percentage of the number of subcarriers budgeted for sounding,

$$S_{avg} = (Sb_{sounding}/Sb_{total}) * (T_{sounding_cycle}/T_{sounding_interval}) \quad (4)$$

where,

$Sb_{sounding}$: total number of subcarriers in all concurrent sounding test opportunities

Sb_{total} : total number of subcarriers on a given FDX channel under test,

for OUDP sounding, $Sb_{sounding} = Sb_{total}$;

for CW sounding, $Sb_{sounding} < Sb_{total}$;

$T_{sounding_cycle}$: duration of a sounding cycle to sound all intended Test CMs on a given FDX channel

$T_{sounding_interval}$: the average time interval between subsequent sounding cycles.

4. Sounding Cycle

As mentioned in the previous section, a sounding cycle includes all the necessary operational steps to identify the interference relationships among all CMs that may transmit and/or receive on a given FDX channel. As shown in Figure 3, a sounding cycle includes preparation, interference test and recovery three phases:

- Preparation Phase

To prepare for sounding, the CMTS has to ensure the FDX channel operates in the DS direction from the measurer CMs' point of view. If the FDX channel has been operating in the US direction in regarding to the measurer CMs, CMTS must switch it to the DS direction and wait for the measurer CMs to acquire the DS channel prior to sounding starts.

- Interference Test Phase

The interference test phase consists of one or more test windows. Each test window marks the time span of one or more parallel test opportunities as shown in Figure 3. In case of OUDP sounding, a single test opportunity covers the entire FDX channel width, hence the number of test windows required is equivalent to the number of test CMs. In case of CW sounding, a test window may contain multiple concurrent test opportunities arranged at difference frequency locations. These test opportunities can be assigned to one test CM or a group of test CMs to sound in parallel. The number of test windows required therefore equals to the number of parallel test groups that can be arranged among the test CMs. Parallel sounding is an optimization technique to shorten the sounding cycle.

- Recovery Phase

After the interference test is done, a recovery phase is required for the CMTS and the CM to resume regular operations. The recovery phase may include channel direction change to recover the traffic throughput prior to sounding.

The sounding cycle duration can be simply expressed as,

$$T_{\text{sounding_cycle}} = T_{\text{prepare}} + N * T_{\text{test_window}} + T_{\text{resume}} \quad (5)$$

Where,

T_{prepare} : sounding preparation time

N : the number of sounding test windows

$T_{\text{test_window}}$: duration of a sounding test window

T_{resume} : recovery time to resume FDX traffic operation post sounding.

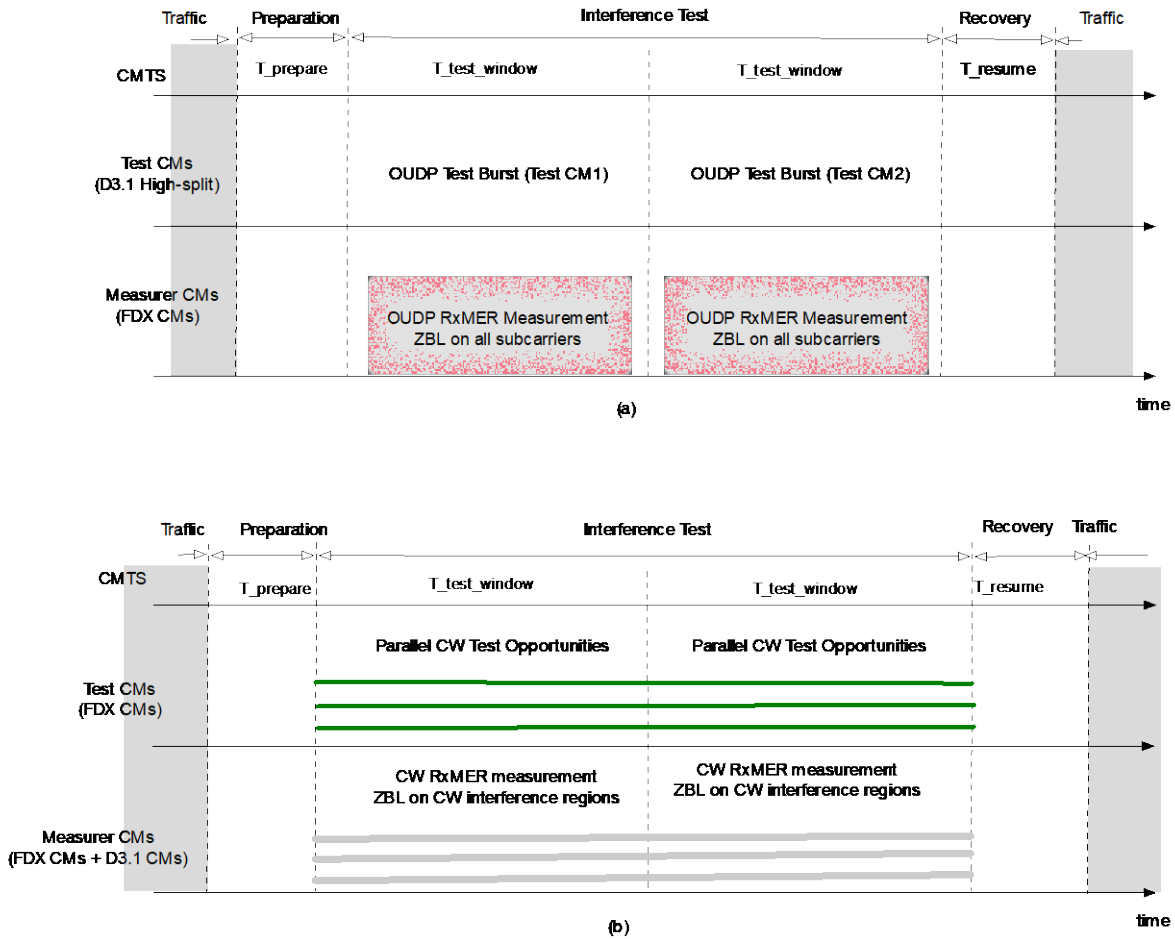


Figure 3 - Sounding Cycle (a) OUDP Sounding (b) CWT Sounding

The sounding cycle duration is a performance benchmark from the FDX operation point of view. It quantifies the FDX bandwidth access time when a new FDX CM is coming online and the traffic interruption time when there are active FDX CMs already operating on the given FDX channel prior to start of sounding.

For CW sounding, the sounding cycle duration is inversely proportional to the number sounding subcarriers at a given spectrum overhead level, as shown in equation (4). It is also impacted by the number of concurrent CW test signals that a CM can send. Figure 4 shows the sounding cycle duration in relation with the sounding subcarrier percentage and the number of CW test signals per CM.

From the chart, we can observe that at given sounding frequency granularity:

- The CW sounding cycle duration decreases as the number of sounding subcarriers increases.
- The sounding cycle duration remains the same if the number of sounding subcarriers allocated results in the same number of test windows.

- The number of CWs a CM needs to generate is bounded by the available number of sounding subcarriers. For example, there is no time advantage for a test CM to generate more than 255 CW tones if only 5% of the subcarriers can be used for sounding at any given time.
- At given frequency granularity, spectrum budget and number of CMs to sound, an optimum number of concurrent CW tests per CM exists that can result in the shortest sounding cycle duration.




Figure 4 - CW sounding duration at different sounding subcarrier percentage and number of CWs per CM

Table 2 - Assumed Parameters for the CW Sounding Duration Calculation

Number of test CMs	60
Number of subcarrier a CM need to sound	3976
Number of subcarriers in a CW interference region	7
Time to prepare for CW test	200 ms
CW Test window duration	500ms
Time to resume FDX operation	100ms

IG Discovery Optimizations

This section looks at a set of optimization techniques for IG Discovery based on the following realizations:

1. Since a CM cannot be both a transmitting CM and receiving CM on a given FDX channel at the same time, sounding can be decomposed into two directional tests, namely, a transmitting test and a receiving test that can be conducted independently. This leads to the Partial Sounding technique.
2. The frequency granularity required for sounding is bound by the MER margin acceptable to a modulation order and the corresponding correlation bandwidth in plant's frequency response. This leads to the MER sub-sampling technique.
3. IG discovery accuracy is relative to the DS spectrum efficiency. Errors in interference measurement and estimations can be compensated with lower modulation orders. IG Discovery may never complete as the interference environment keeps changing. This realization leads to the iterative IG Discovery technique.

The following subsections describe each technique in detail.

1. Full Mesh Sounding vs. Partial Sounding

Full mesh sounding is intended to proactively test all pairing permutations between the transmitting CMs and the receiving CMs. To perform full mesh sounding, the FDX channel under test must be changed to the DS direction for all potential measurer CMs. Consequently, full mesh sounding lasts longer in time and causes longer traffic interruptions. Full mesh sounding may not be desirable if the traffic condition does not permit the necessary time and spectrum required.

Partial sounding attempts to minimize the traffic impact by opportunistically pairing the test CM and Measurer CMs based on the channel direction in use. Partial sounding can be either a transmitting test or a receiving test as shown in Figure 5. The transmitting test allows the CMTS to evaluate if a new CM can transmit upstream on a FDX channel when a specific set of CMs are receiving over the same spectrum. The receiving test allows the CMTS to evaluate if a new CM can receive on a FDX channel when a specific set of CMs are transmitting upstream over the same spectrum. Based on the partial sounding, the CMTS can conditionally enable a CM's FDX service if the operation conditions match the tested scenarios




Figure 5 - Partial Sounding, (a) Transmitting Test; (b) Receiving Test.

Full mesh sounding and partial sounding can be combined to provide an optimum system solution, for example applying full mesh sounding upon boot up to acquire the interference relationship base line, and applying partial sounding repetitively when a new interference condition is present.

2. Sequential Sounding vs. Parallel Sounding

Parallel sounding is used to reduce the sounding cycle duration. Parallel sounding is possible when the number of sounding test opportunities is greater than the number of test signals a CM needs to generate at a time.

The following is an example to exam the timing advantages of the parallel sounding. Figure 6 shows a service group with N (64 in this example) FDX CMs that are capable to transmit and receive on a FDX channel. The time to conduct full mesh sounding requires N CW sounding test cycles, if sounding is performed sequentially with only one CM transmitting in each test window.

Figure 7 shows a parallel sounding algorithm that sounds 8 CMs at a time. First horizontally by arranging each column of 8 CMs transmitting on different subcarrier locations while the rest of CMs in the service group measuring MER on all DS subcarriers. After this step, the only unknown interference is between different rows, so the second step is to sound vertically by arranging each row of CMs to send test signals in parallel while the rest of the CMs measure. The total number of CW test cycles with this approach is 16. Assuming each CW test cycle takes 800ms, parallel sounding in this example only takes 12.8 seconds, while the sequential sounding method would take 51.2 seconds.

Compared to sequential sounding, parallel sounding takes less time but a cost of frequency granularity. Parallel sounding is suitable to identify interferences at restricted frequency locations or form coarsely grained IGs to speed up FDX service access.




Figure 6 - Sequential sounding example




Figure 7 - Parallel Sounding Example

3. Complete Sampling vs. Sub-sampling

Complete sampling refers to the type of sounding in which sounding is attempted on all subcarriers of a given FDX channel. In the case of CW sounding, complete sampling can only be achieved with incremental subsampling, which may take an extended period of time to complete.

The complete sampling is generally not necessary for FDX operation. The frequency granularity required for sounding is bound by the MER margin acceptable to a given modulation order and the corresponding correlation bandwidth at a given frequency. Results from the subsampling can be directly used for IG discovery. The CCI level on the unsounded subcarriers can be interpolated with a maximum likelihood estimation with certain error margins.

Figure 8 shows a subsampling example with the measured MERs scattered across a few subcarriers. Figure 9 shows the MER interpolations in between the sparsely spaced measurement samples. For each estimated MER value, a variation range is incorporated to bound the worst-case estimations. As time progresses and more subcarriers are sounded, the cumulative subsampling approaches the full sampling with less estimation errors as shown in Figure 10.

Subsampling allows the CMTS to quickly enable the FDX operations with coarsely grained initial IGs, and incrementally refine the IG formations with continuous subsampling.




Figure 8 - Sub-Sampling At Selected Subcarrier Locations




Figure 9 - Subsampling With Interpolated MER (Inner) Estimations




Figure 10 - Cumulative subsampling over time

4. Iterative IG Discovery

The iterative aspect of IG Discovery is important. As the interference environment changes, either triggered by a new CM coming online, channel allocation change or temperature fluctuations, the system must be able to adapt, using previous computations together with any new sounding data to produce reliable IG decisions.

The iterative IG Discovery process can be modeled as a multi-stage feedback loop that constantly refines the IG decisions based on the new measurement data and the feedback for positive and negative outcomes. As shown in Figure 11, the iterative IG Discovery process includes the following four steps:

- Sounding

This is for measuring the interference between the specific transmitting and receiving CM pairs at given frequency locations. The measurement data obtained will be used for IG formation.

- IG Formation

The new measurement provided by sounding, together with previous computation results, is used to form IGs to enable FDX operation with acceptable error margins.

- FDX Operation

The FDX operation is constantly monitored. Events and statistics, such as CM population, traffic condition and signal quality are collected for IG evaluation.

- IG Evaluation

IG decisions are re-evaluated based on the operation events and statistics. The evaluation results in a new set of transmitting and receiving CM pairs and specific frequencies targeted for the next round of sounding.

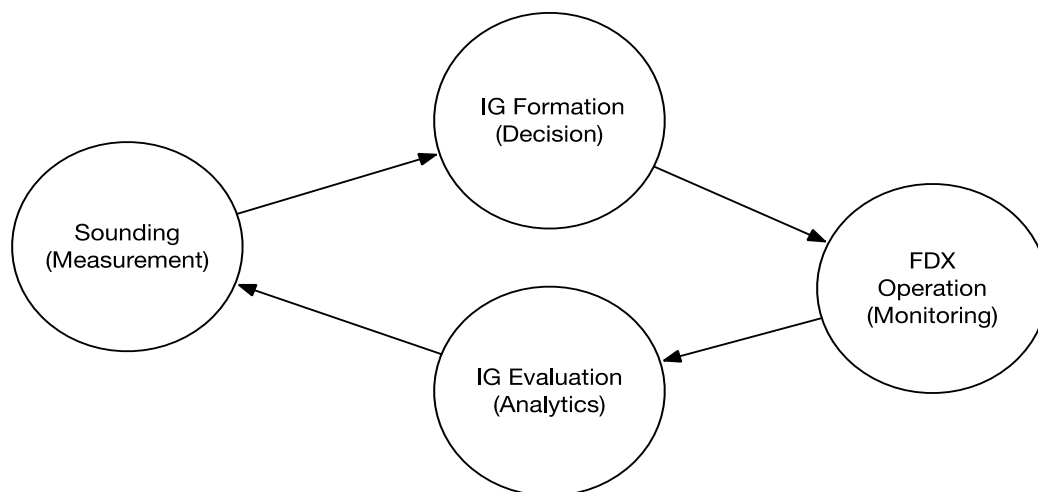


Figure 11 - Iterative IG discovery process

Conclusion

The operational requirements for IG Discovery results in conflicting design considerations, in terms of spectrum budgeting, time to convergence and the interference detection accuracies. In search for a balanced, optimization solution, a system approach is used to identify the key performance impacting elements and their tradeoff relations. Based on this, a set of optimization techniques are described including:

- partial sounding
- parallel sounding
- interference subsampling with interpolations

The solution space is further extended by incorporating an iterative process that follows a measurement – decision – monitoring – analysis feedback loop, to allow the IG Discovery to be constantly refined and adaptive to the changing interference environment.

Abbreviations

CM	Cable Modem
CMTS	Cable Modem Termination System
CW	Continuous waveform
DS	downstream
FDX	Full Duplex
IG	Interference Group
HFC	hybrid fiber-coax
MER	Modulation Error Ratio
Hz	hertz
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing with Multiple Access
OU DP	OFDMA Upstream Data Profile
PON	Passive Optical Network
US	upstream

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] FDX MAC EC: MULPIv3.1-x-17.1764-1
- [4] CM-SP-PHYv3.1 Annex F

Bridging the Gap Between ETSI-NFV and Cloud Native Architecture

A Technical Paper prepared for SCTE•ISBE by

YuLing Chen

Senior Technical Leader
Cisco Systems Inc.
375 East Tasman Drive
San Jose CA 95134
408-393-5606
yulingch@cisco.com

Alon Bernstein

Distinguished Engineer
Cisco Systems Inc.
375 East Tasman Drive
San Jose, CA 95134
alonb@cisco.com

Introduction

In recent years, Network Function Virtualization (NFV) has been introduced into the Telecom industry to deliver reliable and efficient commercial networking services in programmable standard hardware systems, called Virtualized Network Functions (VNFs). NFV promises benefits in the savings of operational and capital expenditure (OpEx and CapEx), as well as the increased automation, operations simplification, business agility, and faster time to market.

The cloud native microservices container architecture was originated from the webscale providers such as Amazon, Google, and Netflix. The approach of cloud native is to break down a monolithic application into small microservices and deploy as containers in the cloud. One of the attractions of this approach is that applications can be tested in an iterative and distributed model, without taking applications offline. In the cloud world, large scale applications have been developed, tested, and deployed with more agility using this distributed model.

Since 2016, several large service providers have publicly embraced the move to a microservices architecture in the telco cloud. [1] There have been announcements from major service providers to use containers to build out their network function virtualization infrastructure. Some key telecommunications equipment suppliers are using microservices to implement some of their software. Open-source initiatives are moving towards microservices and containers. In NFV space, there is a trend of moving from the virtual appliance based solutions to the cloud native approach, which is referred to as the Cloud Native NFV.

The NFV world has been following ETSI NFV references. However, most of the ETSI published documents were based on case studies and Proof of Concepts built on virtual appliances. There is a gap between ETSI NFV and the cloud native approach. With more and more cloud native solutions appear in NFV, there is a need to augment the existing ETSI NFV specifications so as to continue guiding the NFV world towards interoperability and standardization.

To support this effort, this paper identifies the elements in the ETSI NFV Management and Orchestration (MANO) reference architecture that need to be adjusted when applying the cloud native approach in NFV. We also propose a pragmatic software architecture that realizes the NFV MANO functionality using the cloud native approach. With the focus on the network service design and deployment, which is the core functionality of the NFV Management and Orchestration systems, we exercise the TOSCA language for the service modeling in the cloud native environment.

The Cloud Native Trend in NFV

Since 2012, driven by leading telecoms network operators, the European Telecommunication Standards Institute (ETSI) has been working on NFV requirement prioritization, high level architectural framework definition, development guideline specification, and Proof of Concept organization. ETSI has published a series of documentation and specifications in these related efforts. The documentation has been widely referenced and adopted in the NFV space.

What ETSI NFV advocates has been the moving of network functions from specialized proprietary hardware to virtualized software that can be deployed on standard hardware equipment. [2] Now, with the

cloud native adoption in NFV, we observed that the network functions together with the MANO systems are moving into the cloud as a form of microservices containers.

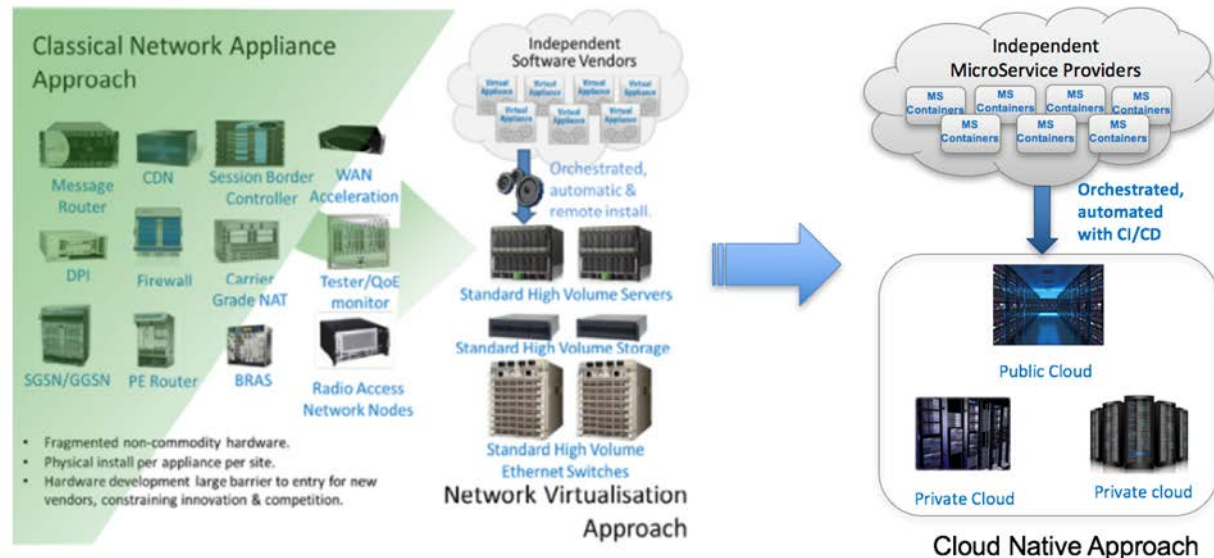


Figure 1 - The Trend of the Cloud Native NFV

As described in Figure 1, with the Network Virtualization Approach, which ETSI-NFV has been focusing on, the classical network appliances with non-commodity hardware move to the software based virtual appliances deployed in the standard equipment. With the cloud native approach, the network functions, which were implemented as monolithic applications, now are broken down into smaller microservices, and deployed as containers in both the public and private clouds. Leveraging Continuous Integration and Deployment (CI/CD), these microservices containers are orchestrated and deployed with automation. The independent software vendors who used to produce full-fledged network functions now become the vendors of smaller microservices.

More specifically, we observed the adoption of the cloud native solutions in the following NFV areas:

- VNFs
 - More and more VNFs are broken down into smaller microservices containers.
 - More and more NFV applications are packaged as microservices containers and deployed in the cloud native environment.
- MANOs
 - More and more NFV Management and Orchestration systems are deployed in the microservice container environment
- VIMs
 - Container orchestrators such as Kubernetes and Docker Swarm appear in NFV applications

ETSI NFV Adaptation to the Cloud Native Architecture

ETSI NFV ISG has published a series of specifications including the ETSI MANO GS (Group Specification) *Network Functions Virtualisation (NFV); Management and Orchestration*. [3] The specification lays out the NFV MANO objectives and concepts, defines the high level reference architectural framework, and specifies the information elements in an NFV MANO system. ETSI Management and Orchestration Architectural Framework

ETSI NFV reference architectural framework defines three functional blocks in the NFV-MANO domain: NFV Orchestrator (NFVO), VNF Manager(s) (VNFM(s)), and Virtualized Infrastructure Manager(s) (VIM(s)).

- **NFV Orchestrator (NFVO)**

NFVO is responsible for the on-boarding of a new Network Service (NS) composed of multiple VNFs, VNF forwarding graph, Virtual Links, and, as an option, Physical Network Functions (PNFs). The orchestrator also controls the life cycle of the Network Service, validates and authorizes NFVI resource requests, manages global resources, as well as the policy of the Network Service instances.

- **VNF Manager(s) (VNFM(s))**

The VNFM focuses on the life cycle management of individual VNF instances. A VNF manager takes the responsibility of the management of a single VNF instance, or the management of multiple VNF instances of the same type. VNFM also serves as an overall coordination and adaptation role for configuration and event reporting between the VIM and the EM systems of traditional operator architectures.

- **Virtualized Infrastructure Manager(s) (VIM(s))**

The VIM is responsible for controlling and managing the NFVI compute, storage and network resources. At the same time, it collects performance measurements in the infrastructure and makes the data available from other functional blocks for monitoring purposes.

The NFV-MANO architectural framework also identifies main reference points for the exchange of data between the corresponding defined functional blocks.

Proposed ETSI MANO Reference Architecture Augmentation for Cloud Native NFV

As a standard specification, ETSI focuses on high level architecture, development guidelines, and interoperability enabled by open interfaces. Most of the specifications in ETSI MANO GS continue serving the purpose when applying to the Cloud Native NFV. Nevertheless, augmentation is needed in some areas because of the differences between the VM based and cloud native solutions.

In the cloud native architecture, the network functions are deployed in the cloud as microservices containers. The granularity of the deployed instances is much smaller based on microservices design and implementation. A VNF in ETSI context would contain multiple microservice containers working together in the cloud native context. This adds complexity to the management and orchestration of the system. The fully distributed architecture, coordination and communication among the microservices, fault monitoring and recovery in smaller but more specific portion of the software, all contribute to the complexity of the MANO system.

On the other hand, the cloud native architecture brings advantages to NFV MANO systems including some critical pain points. One pain point in the VM based solutions is to provide high availability to the service providers by spawning new instances in the cases of faults, failures or scaling out to handle larger workloads. The time needed for spinning up a new VM has been the bottleneck in the VM based solutions. Using the cloud native approach, because the granularity of the independent unit for recovering or scaling out is much smaller, and the container start/stop is much faster, the latency of spawning a new network function composed of a set of microservice container instances is much smaller than that in the VM based solutions.

To highlight the differences between the VM based solution and the cloud native based approach in NFV, we propose the Cloud Native NFV high level reference architecture with the revised terminologies as an augmentation to the original ETSI NFV MANO reference architecture, which is illustrated in Figure 2.

As described in the diagram, The Network Function Cloud Infrastructure (NFCI) contains the public cloud and the private cloud(s) with containerization layer(s) to provide the infrastructure for the network services deployed as containers in the Cloud Native architecture. On top of the NFCI, a set of microservice containers work together to realize the functionality provided by a Cloud Network Function (CNF). Each microservice in a CNF is called a Cloud Network Function MicroService (CNFMS). A list of CNFs chain together with traffic flowing through the network functions becomes a CNFFG.

NFV Orchestrator (NFVO), CNF Manager (CNFM), and Cloud Infrastructure Manager (CIM) are the functional blocks of NFV MANO systems. The NFV MANO system communicates with OSS/BSS, CNFs in CNFFG, and NFCI through interfaces to manage and orchestrate the network services provided by the network functions deployed as microservices in the NFCI.

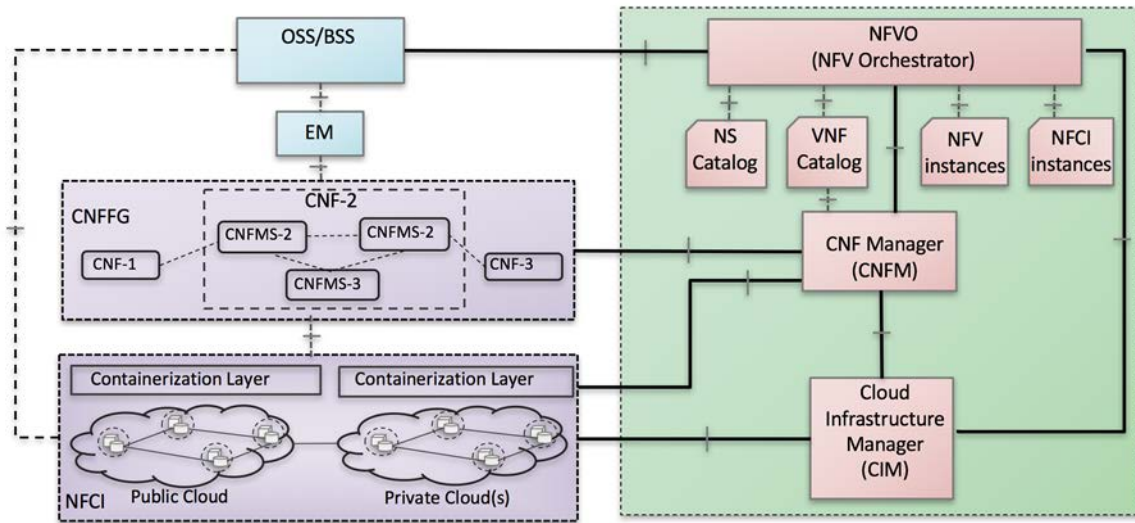


Figure 2 - Proposed ETSI MANO Reference Architecture Augmentation for Cloud Native NFV

The detailed description of each functional block and component in Figure 2 is as follows:

- **Cloud Network Function MicroService (CNFMS)**

CNFMSs are the microservice containers from which a Cloud Network Function is composed.

- **Cloud Network Function (CNF)**

CNFs are the network functions deployed in the cloud as microservices, usually in container format.

- **Network Function Cloud Infrastructure (NFCI)**

NFCI provides the underline physical infrastructure for the network functions. This includes the hardware equipment for the computer, networking, storage, as well as the containerization layer on top of the hardware platform. CNFs are deployed on top of the NFCI.

- **Cloud Infrastructure Manager (CIM)**

CIM is responsible for controlling and managing the NFCI compute, storage and network resources, as well as scheduling the microservice containers in the cloud. It manages the lifecycle of the containers in the cloud.

CIM also collects performance measurements in the infrastructure including container level, and makes the data available for other functional blocks for monitoring purposes.

Other responsibilities of CIM include virtual networking control and management, as well as the southbound integration with various network controllers to achieve the physical network control and management capabilities.

Examples of the CIMs available in the market are Kubernetes, AWS ECS, and Docker Swarm.

- **Cloud Network Function Manager (CNFM)**

CNFM focuses on the life cycle management of individual CNF instances. In the cloud native architecture, a CNF is usually composed of a set of containers that implement a network function. A CNF manager takes the responsibility of the management of the multiple container instances of the same network function. To control the lifecycle of the CNFs, CNFM works closely with CIM, which manages the lifecycle of the individual container of the CNF.

CNFM also serves as an overall coordination and adaptation role for configuration and event reporting between the CIM and the EM systems of traditional operator architectures.

- **NFV Orchestrator (NFVO)**

The NFVO continues serving the responsibility of on-boarding a new Network Service (NS) composed of multiple CNFs, CNF forwarding graph, Virtual Links, and, as an option, Physical Network Functions (PNFs). The orchestrator also controls the life cycle of the Network Service including instantiation, scale-in/out or up/down, performance measurements, event correlation and termination. Further key operational functions are global resource management, validation and authorization of NFVI resource request, as well as policy management of Network Service instances.

- **Cloud Network Function Forwarding Graph (CNFFG)**

The CNFFG contains a list of CNFs and the virtual links among the CNFs and the physical endpoints.

Cloud Native NFV MANO Software Architecture

As a standard specification, ETSI focuses on high level architecture, development guidelines, and interoperability among systems produced from different vendors. In order for the industry to generate real products in microservices container architecture, we further refine and develop the functional blocks into micro services, helping to implement and realize the functionality mentioned in the previous section. We also realized that a real NFV MANO product needs to address more operational issues and challenges than what ETSI has specified. This includes the design phase support, network control in the cloud native environment, data collection, monitoring, and analytics. [4]

Figure 3 illustrates a pragmatic NFV MANO architecture using the cloud native approach.

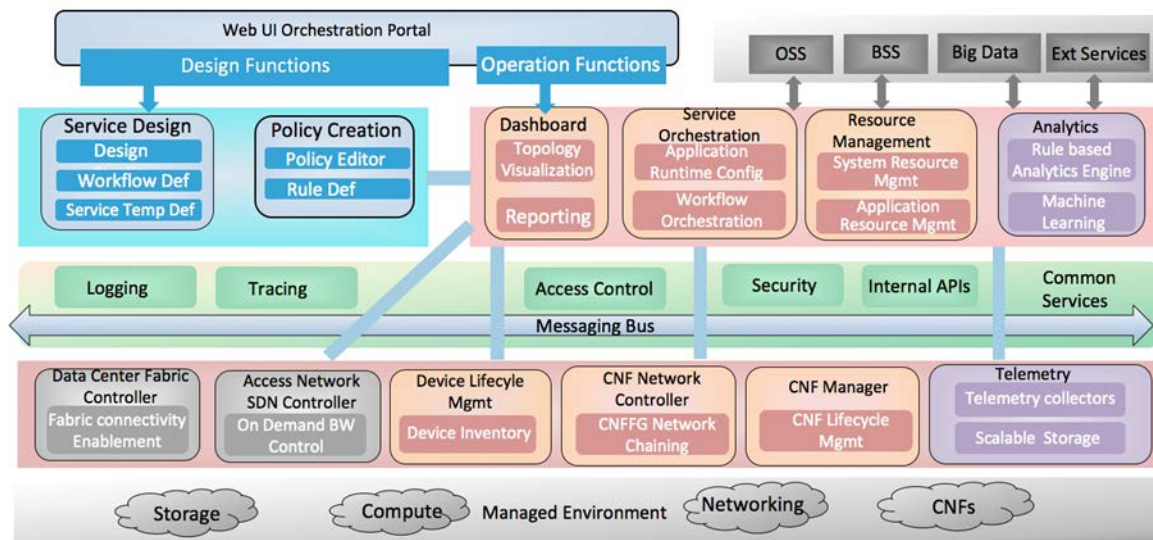


Figure 3 - A pragmatic NFV Software Architecture in the Cloud Native Environment

Figure 3 contains the microservices for the NFV Management and Orchestration in both the Design Phase and the Runtime Execution Phase. The design phase generates the network services model to be deployed onto the target private or public cloud(s). The Runtime execution phase contains the microservices for the network service deployment, orchestration, lifecycle management, network control, monitoring, and analytics. Besides the functional services, there are a set of common infrastructure services for all the microservices containers in the cloud native architecture environment.

1. Network Service Design Phase

One of the key functionality of NFV MANO is the Network Service Orchestration. Before the MANO can orchestrate the virtual and physical network services, we first need to design the services with rules and policies to indicate how the services are being deployed at run-time.

The Service Design microservice provides interfaces, usually in a visual studio way, for the user to design and model the network services and store them into a network service catalog (NS Catalog). The definition of the services need to specify how and when the CNFs are realized in a target environment. In particular, the definition would need to include the logical catalog items, together with selected workflows and instance configuration data, completely defines how the deployment, activation, and life-cycle management of CNFs are accomplished.

To facilitate the portability of the service design, it is desired to use a standard modeling language to describe the service definition. With such definition, any NFV orchestrator that is compliant with the standard modeling language can take the service specification and deploy into the target cloud environment.

2. Run-time Execution Phase

Various microservices work together during run-time to realize the functionality of the network service deployment, configuration, monitoring, and data analytics.

Taking the generated service definition as the Network Service Descriptor (NSD) and Cloud Network Function Descriptor (CNFD), the Service Orchestrator executes the workflow specified in NSD and works with other microservices to control the lifecycle of the network service; the Dashboard visualizes the topology of the network with basic monitoring of the health of the system; the Global Resource Manager allocates both the system resources and the network application resources; the CNF Manager deploys and manages the lifecycle of the Cloud Network Functions; the CNF Network Controller focuses on the control and management of the virtual networking among the CNF instances; and the Device Life Cycle Manager manages the physical life cycle of the devices.

Container networking in the cloud native environment is important to NFV applications. Since the CIM is realized by third party container orchestration tools, and most of the tools lack the full-fledged networking feature including policies and security support across different compute nodes, we would need to plug in various types of Network Controllers to augment the capability needed for the network functions deployed in the cloud.

The CNF Network Controller stitches the network connectivity between microservice containers. Usually there are two types of network traffic going through between CNFs: the control traffic and the data traffic. The control traffic usually requires lower bandwidth with relatively longer latency. The data traffic requires higher bandwidth with lower latency. Most container orchestration tools contain sufficient support for the control traffic. However, for data traffic, specialized CNF controllers will work with high speed virtual routers to realize the data plane acceleration in the cloud native environment.

The Data Center Network Controller and Access Network Controller are southbound plugins to the CIM to provide the physical networking to virtual networking mapping in the NFV system.

Telemetry and Analytics microservices work together to realize the collection, streaming, storage, and analytics of the operational data collected from the network.

Besides the functional services to achieve the NFV MANO capabilities, there are a set of common services that are particularly important in microservice container architecture.

The messaging bus enables the loosely coupled integration architecture in microservice container environment using publish/subscribe way of the inter container communication. Another important role of the messaging bus is to support large amount of telemetry data pushed from the network functions. The scalability and performance of the messaging bus is critical to the success of the cloud native architecture. Currently Kafka is one of the most popular messaging tools that are widely used in the microservices container environment.

Logging and tracing help with efficient and effective troubleshooting across distributed microservices. Open source tools such as fluentd, ELK, open tracing, and zipkin are popular ones to enable centralized logging and cross service tracing capabilities.

Service Design and Deployment using TOSCA Modeling Language

Network Service design and deployment is fundamental to NFV MANO functionalities. As the result from the service design, network service descriptor is generated. At run time, the NFV Orchestrator deploys the network service as CNF instances in the cloud native environment.

1. Service Design and Deployment Process Flow

The process flow of the Network Service Design and Deployment is described in Figure 4. The first step is to design the service using visualization tools in the Design Studio. The Design Studio is a graphical interface for the user to define the network services that contain the network nodes and the relations among them as links. A typical graphical Design Studio provides a set of drag and drop tools to help make the modeling process easy and intuitive to the user. During this process, the user enters the workflow definition, service template definition, and policy description as the input. As the result from this step, a Network Service Descriptor (NSD) will be generated and stored in a database called NS Catalog. At run time, as part of the NFVO, a NS Deployer would read the NSD from the catalog. The NS Deployer decomposes, translates, and converts the NSD into CNFD and stores it into the database called CNFD Catalog. The CNFD contains the specifications of the Cloud Network Functions that are ready for being deployed in the cloud. After that, the CNF Deployer reads the CNFD from the catalog and converts it into the executable artifacts accepted by the target cloud provider. Finally, the CNF Deployer communicates with the CIM to deploy the workload into the cloud.

Figure 4 describes the process flow of the Network Service Design and Deployment.

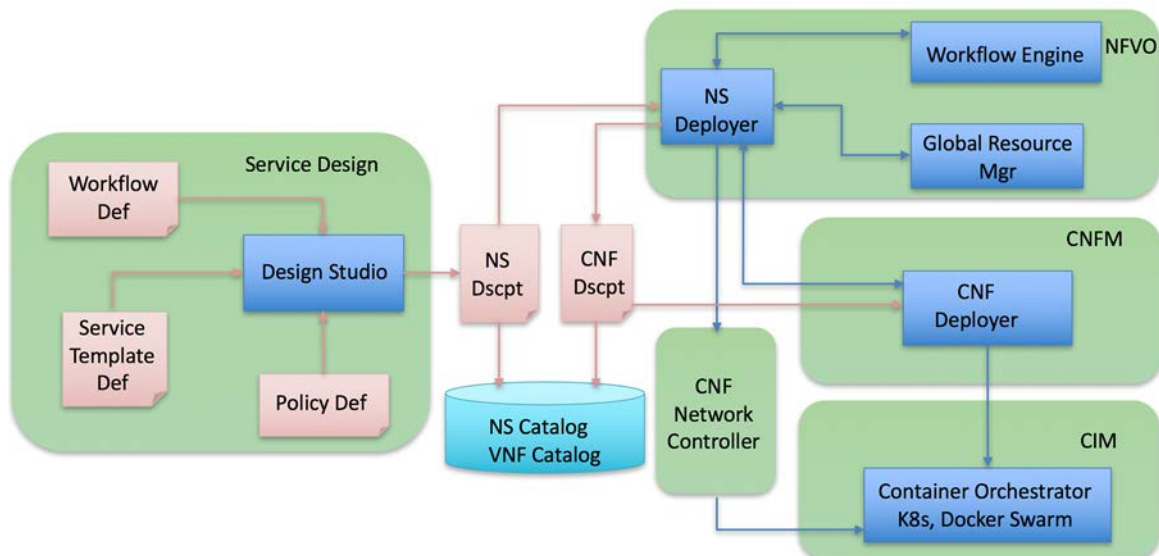


Figure 4 - Service Design and Deployment Process Flow

In the process described above, if the descriptor of the Network Service (NSD) and Cloud Network Function (CNFD) is specified in a standard modeling language, any NFV MANO that is compliant with

such standard language will be able to deploy the workload into the cloud. This has been the goal of TOSCA modeling language.

2. Network Service Modeling using TOSCA

Topology and Orchestration Specification for Cloud Applications (TOSCA) is a standard modeling language managed by industry group OASIS that can be used to orchestrate NFV services and applications. TOSCA delivers a declarative description of the application topology for a network or cloud environment that includes all its components, which may include the need for load balancing, networking, computing resources, and other software. It can also be used to define the workflows that need to be automated in the cloud.

The TOSCA modeling language includes concepts such as nodes and relationships, whereby a node is an infrastructure such as network, subnet, or a server software component. TOSCA can help define how these nodes and services work together. TOSCA uses templates to automate the configuration of these relationships.

TOSCA facilitates high levels of service portability, making services portable to any cloud or application that is TOSCA compatible. The data model also enables easier migration of applications. It is inherently infrastructure-agnostic, and thus is extensible to enable the automation of software-defined networks, in combination with NFV and clouds, to simplify end-to-end service orchestration for cloud and telco operators.

Figure 5 illustrates how to use TOSCA modeling language to model NFV Network Services and achieve the portability of deploying the same services in different cloud providers. [6] As described in the diagram, a CNFFG defined in TOSCA language can be deployed by a TOSCA compliant orchestration engine to different types of Cloud Providers such as AWS, Kubernetes, or Docker Swarm. The TOSCA compliant orchestration engine applies the necessary automatic matching, translation and optimization between the application requirements and the NFV infrastructure capabilities provided by the target cloud providers to achieve the portability.

TOSCA supports XML, JSON, and YAML implementation of the data model. With the industry trend moving to microservices container architecture, more and more TOSCA implementation products started the effort in deploying TOSCA defined services as microservice containers in the cloud native environment.

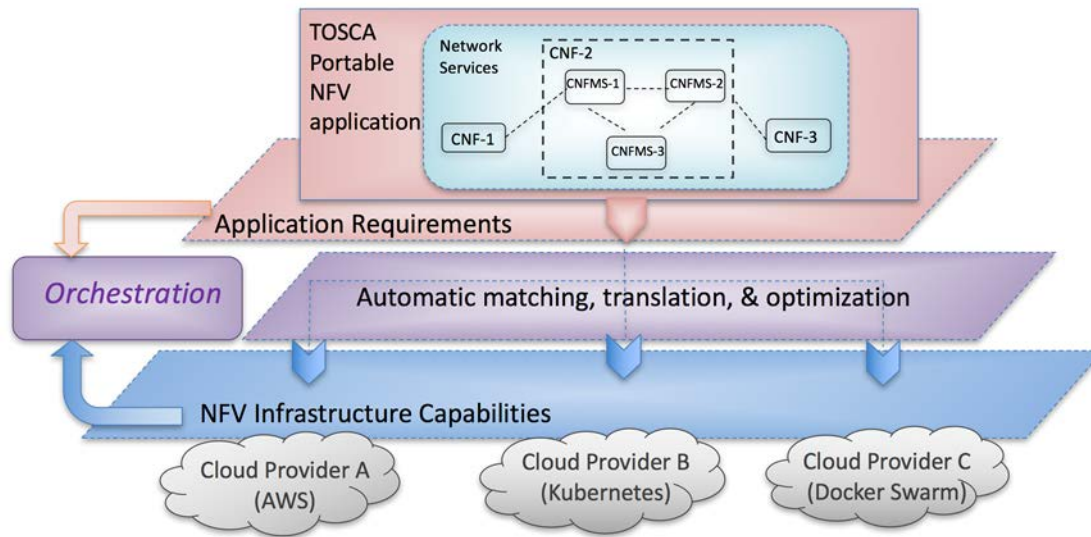


Figure 5 - TOSCA to support portable NFV applications[6]

3. TOSCA and YANG

YANG is a data modeling language used to describe configuration and state information. It was published by Internet Engineering Task Force (IETF) in 2010. YANG has been used to model networking devices and services – i.e., an object and its attributes. YANG defines the data models that are manipulated through the NETCONF protocol.

There has been a battle between using TOSCA or YANG modeling languages in the NFV context. These two modeling languages are not competitors but complementary to each other in different perspectives. TOSCA focuses more on the network topology, cloud workload, workflow representation, and deployment artifacts specification of the network services. The goal of TOSCA is to prepare a declarative specification for the workload being deployed in the cloud. YANG focuses more on the configuration of the network functions in the cloud. YANG provides the ability to easily configure network devices in a human readable fashion.

Because of YANG's strength is in configuring networking devices while TOSCA's strength is orchestration, we suggest using TOSCA in NFVO and CNFM, while using YANG for the configuration of CNFs and PNFs in NFV MANO architectural framework. [5] Figure 6 illustrates the idea of using both TOSCA and YANG in different functional blocks of the ETSI NFV MANO reference architecture.

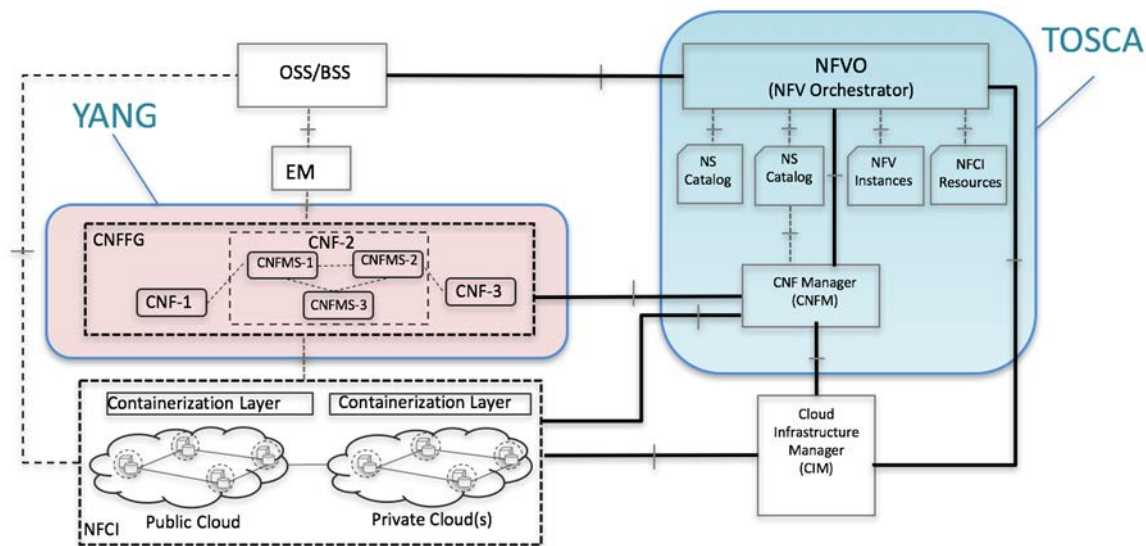


Figure 6 - Using TOSCA and YANG in different perspectives of NFV applications[5]

4. TOSCA for Cloud Native NFV

The TOSCA metamodel uses the concept of service templates to describe cloud workloads as a topology template, which is a graph of node templates modeling the components that the workload is made up of, and as relationship templates modeling the relations between those components. TOSCA further provides a type system of node types to describe the possible building blocks for constructing a service template, as well as relationship types to describe possible kinds of relations. Both node and relationship types may define lifecycle operations to implement the behavior an orchestration engine can invoke when instantiating a service template.

An orchestration engine processing a TOSCA service template uses the mentioned lifecycle operations to instantiate single components at runtime, and it uses the relationship between components to derive the order of component instantiation. Network Service workflow can be modeled using TOSCA leveraging the relationship modeling capabilities from TOSCA.

The TOSCA simple profile defines a number of base node types and relationship types to be supported by each compliant environment. Furthermore, it is envisioned that a large number of additional types for use in service templates will be defined by a community over time. At the same time, TOSCA is highly customizable with type inheritance capabilities built in the language. Specialized TOSCA engines can build the support for customized node types and relationship types to satisfy the needs.

TOSCA has been popular in the network service modeling in NFV. OASIS published the TOSCA Simple Profile for Network Functions Virtualization (NFV) Version 1.0 in May 2017. [8] However, the specification is VM based without the microservices container support. There is a need for extending the TOSCA specifications to support the Cloud Native NFV. With the strong growth of the Cloud Native NFV in the Telco space, this could be the next step from OASIS in the near future.

Before we have a set of TOSCA base node types and base relationships types defined by OASIS for the Cloud Native NFV, we can leverage the generic TOSCA base types with customization. Recently, there

has been research looking into how to use customized types in TOSCA to model generic applications deployed as Docker containers in the cloud [9][10]. We can take the same approach while adding additional custom types needed in the Cloud Native NFV context.

5. Example of using TOSCA Modeling Language

In this section, we will exercise an example using TOSCA modeling language to design and deploy a network service in a CMTS orchestration system. The goal is to illustrate the approach of using TOSCA to model the Network Service that contains sufficient information for the NFVO to deploy in a cloud native environment.

In this example, we will model a Physical Network Function called Remote PHY Device (RPD), which connects to a Cloud Network Function called Cloud Cable Modem Termination Service (CCMTS) to consume the CMTS services. To make it simple, we assume the CCMTS contains only one microservice container. We also assume that the CCMTS will need to be deployed as a Docker container in the cloud. The lifecycle of CCMTS and RPD needs to be specified in the service model.

To achieve the above goal, after analysis, we decided to add four node types and two artifact types to the existing TOSCA base type definition included in the TOSCA Simple Profile in YAML Verion 1.0 [7]. The current TOSCA relationship types defined in the referred document are sufficient to support the relation definition between the nodes.

Table 1 describes the custom types that we need to add to the TOSCA base type definition.

Table 1 - TOSCA custom types to support network services modeling in an example CMTS System Ochestraor

Node Types	Extends	Artifact Types	
Container	tosca.nodes.Root	Image	tosca.artifacts.Root
Container.Executable	cabu.nodes.Container	Dockerfile	tosca.artifacts.Root
Software	tosca.nodes.Root		
Volume	tosca.nodes.Root		
Phyendpoint	tosca.nodes.Root		

In the table, all the custom types extend the root Data Type in the standard TOSCA specification. The detailed definition of the toscanodes.Root can be found in [7]. The Container data type defines the basic meta data of the Docker container; The Docker executable container defines more data when the container is deployed in an actual cloud provider environment. The Software data type defines the actual software that runs inside a container; The Volume specifies the storage attached to the container; The Dockerfile and the Image are the two artifacts that the container uses for the software packaging and deployment; The Phyendpoint is the data type used for modeling a physical device.

Figure 7 describes the attributes of each Node Type and Artifact Type in the above table.




Figure 7 - Visualized TOSCA custom types for CMTS System Orchestrator

The above custom types help with the modeling of the CCMTS instances in the cloud. The Docker container and Docker executable container provide the attributes for the user to specify the ports, environment variables, and commands need to be executed when starting the container. Since the container is transient and the data needs to be stored in a Volume, the storage property of the Docker container and Docker executable container will allow the user to specify the Volume to be attached to the container. The standard lifecycle operation interfaces allow the user to plugin customized scripts to run during the lifecycle of the containers.

The Physical endpoint helps with the modeling of RPDs. Although it is not part of the workload being deployed in the cloud, we need this entity in the data model to specify the relationship between the RPDs and the CCMTS instances. The attributes and the lifecycle operations of the Physical endpoint will allow the user to uniquely identify this physical network function, and to insert customized workflow operations during the deployment of the RPD. Appendix 2 specifies the detailed TOSCA YAML definition of these custom types.

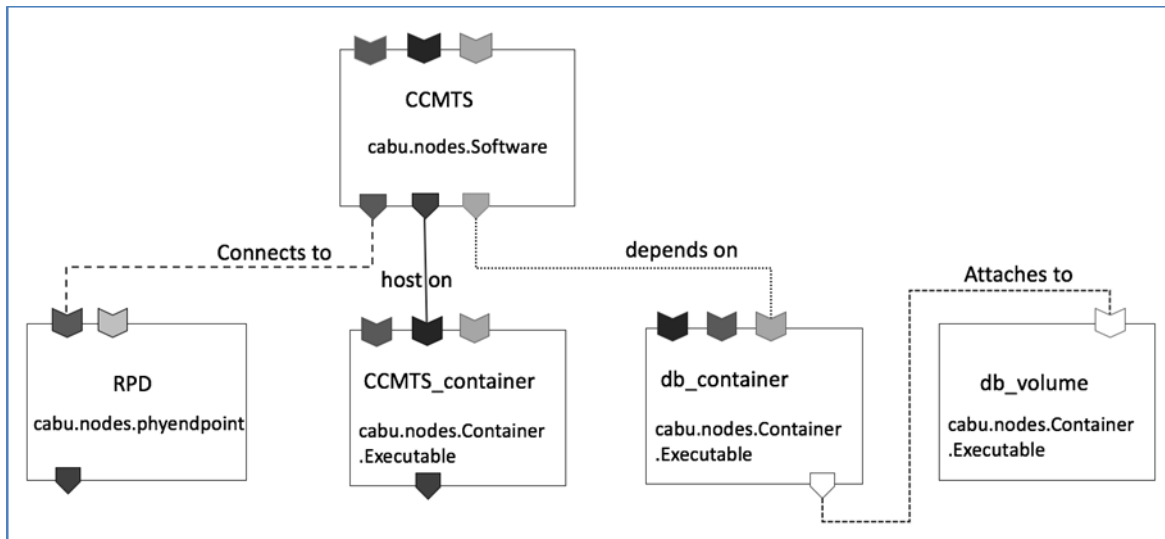


Figure 8 - The TOSCA model of an example CMTS System Orchestrator

Using the custom types, we can model the relationship between the RPDs and CCMTS instances, which is illustrated in Figure 8. In the diagram, CCMTS instance is a type of software node hosted on a CCMTS_container node with the sufficient information to spin up a microservice container in the cloud. This includes the Docker image file location for downloading during the deployment time, the runtime parameters passed to the container, the volume to store the data, which are modeled as different node and artifact instances using the custom types defined earlier. Moreover, the above data model specifies the relation between the nodes using the standard TOSCA relationship types.

The following YAML snippets give an example of the TOSCA definition of the above model. For the complete definition, please refer to Appendix 1 for more details.

```

node_templates:
  ccmts:
    type: cabu.nodes.software
    requirements:
      - host: ccmts_container
      - connection: RPD1
    interfaces:
      standard:
        create:
          implementation: scripts/api/install.sh
          inputs:
            repo: <git_repo_of_the_scripts>
            branch: {get_input: api_branch}
        configure:
          implementation: scripts/api/configure.sh
        start:
          implementation: scripts/api/start.sh
        delete:
          implementation: scripts/api/uninstall.sh

  ccmts_container:
    type: cabu.nodes.Container.Executable
    artifacts:
      ccmts_image:
        file: ccmts:1.0
        type: cabu.artifacts.Image
        repository: <cabu_docker_hub_url>
      requirements:
        - storage:
            node: ccmts_volume
            relationship:
              type: tosca.relationships.AttachesTo
              properties:
                location: /data/ccmts/db

  ccmts_volume:
    type: cabu.nodes.Volume

```

Figure 9 - TOSCA definition of the RPD to the CCMTS Connection in the Cloud Native Environment

After we generate the TOSCA definition of the CMTS service, we can use a TOSCA compliant orchestration engine to deploy the service into the microservices container environment. For example, if the target deployment cloud is Kubernetes managed microservices container environment, the TOSCA engine converts the service definition into Kubernetes deployment scripts and instruct Kubernetes to deploy the specified network services into the cloud.

The example described in this section is extremely simple. A real world NFV service is much more complex. For example, a CCMTS instance in the cloud will contain a set of microservice containers to realize the CNF functionality.

In this case, there are two solutions for the deployment. The first one is to model the service at CNF level without getting into the details of the microservice containers that CNF is composed of. In the data model, we rely on the life cycle operations of the CCMTS to run a script that spin up a set of microservice containers required by the CCMTS instance in the cloud.

Another solution is to rely on NFVO to convert the high level network service data model into the detailed model that contains all the microservice container definitions deployable in the cloud. Then the NFVO sends the generated detailed data model to the CNF Manager, which converts the service definition of the CCMTS into the deployment script that the target cloud understands. Then CNFM instructs the CIM to deploy the CCMTS into the cloud.

In the CNF Manager, we can either incorporate a third party TOSCA compliant orchestration engine or develop a TOSCA orchestration engine from scratch to parse the TOSCA definition based on the predefined node types and relationship types. There are open source TOSCA engines and TOSCA parsers available in the market including Cloudify, Tacker, and Open-O. [11] [12] [13]

Conclusion

Cloud Native NFV is the next wave in the telecommunication and network function virtualization space. At this early stage of Cloud Native NFV, we observe the gaps between the ETSI NFV references and this newly introduced approach in NFV. This includes:

1. ETSI NFV specification focuses on virtual appliances based solutions, and lacks the information and guidelines to support the cloud native architecture and environment.
2. ETSI NFV specification focuses on the service deployment and orchestration. A pragmatic NFV application needs to address other perspectives including service design, modeling, monitoring, and analytics.
3. Network Service design and modeling needs a standard modeling language. Current ETSI NFV compliant service design tool TOSCA is a good candidate but lacks the support for Cloud Native NFV.

This paper helps to bridge the gap between ETSI NFV and the Cloud Native NFV by identifying elements in the ETSI specification that needs to be augmented to support the microservices container environment. We also proposed a pragmatic software architecture to illustrate a Cloud Native NFV MANO system. Because using a standard modeling language for Network Service design is critical to the portability and interoperability of NFV MANO systems, we proposed using TOSCA modeling language and explained how to use its customization capabilities to support the Cloud Native NFV.

Cloud Native NFV is still in its infant stage. The purpose of this paper is to share the observations, practices, and examples in the related areas to help with the building of actual products using this approach. We believe with more and more NFV applications using the cloud native approach, more and more tools to support the NFV applications in the related areas will appear. More and more best practices discussions, guidelines, and specifications will be generated towards the eventual interoperability and standardization in the Cloud Native NFV.

Abbreviations

CIM	Cloud Infrastructure Manager
CCMTS	Cloud Cable Modem Termination System
CMTS	Cable Modem Termination System
CNF	Cloud Network Function
CNFMS	Cloud Network Function MicroService
CNFM	Cloud Network Function Manager
CNFFG	Cloud Network Function Forwarding Graph
EM	Element Manager
ETSI	European Telecommunications Standards Institute
ETSI ISG	ETSI Industry Specification Group
ETSI ISG GS	ETSI ISG Group Specification
IETF	Internet Engineering Task Force
MANO	Management and Orchestration
NFCI	Network Function Cloud Infrastructure
NFV	Network Function Virtualization
NFVI	Network Function Virtualization Infrastructure
NFVO	Network Function Virtualization Orchestrator
NS	Network Service
NSD	Network Service Descriptor
PNF	Physical Network Function
RPD	Remote PHY Device
TOSCA	Topology and Orchestration Specification for Cloud Applications
VIM	Virtualized Infrastructure Manager
VM	Virtual Machine
VNF	Virtualized Network Function
VNFD	Virtualized Network Function Descriptor
VNFM	Virtualized Network Function Manager
VNFFG	Virtualized Network Function Forwarding Graph
AP	access point
bps	bits per second
FEC	forward error correction
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[1] *Microservices architecture in the Telco Cloud*, SDX Central; Available from <https://www.sdxcentral.com/nfv/definitions/microservices-architecture-telco-cloud/>

[2] Network Function Virtualization – Introductory White Paper; Available at https://portal.etsi.org/nfv/nfv_white_paper.pdf

- [3] *Network Function Virtualization Management and Orchestration Group Specification*; European Telecommunications Standards Institute (ETSI) Group Specification; 2014; Available from http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf
- [4] *OpenNetwork Automation Platform*, ONAP Wiki; Available from <https://wiki.onap.org/display/DW/Architecture>
- [5] *TOSCA vs. Netconf – a Comparison*, SDX Central; Available from <https://www.sdxcentral.com/nfv/definitions/tosca-vs-netconf-comparison/>
- [6] *Making TOSCA Truly Portable*, Nati Shalom, May 12, 2016; Available from <http://cloudify.co/2016/05/12/making-tosca-truly-portable-openstack-cloud-nfv-open-source-orchestration.html>
- [7] *TOSCA Simple Profile in YAML Version 1.0*; Available from <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csprd02/TOSCA-Simple-Profile-YAML-v1.0-csprd02.html>
- [8] *TOSCA Simple Profile for Network Functions Virtualization (NFV) version 1.0*; Available from <http://docs.oasis-open.org/tosca/tosca-nfv/v1.0/tosca-nfv-v1.0.html>
- [9] *Orchestrating applications with TOSCA and Docker*, Luca Rinaldi; Available from <https://core.ac.uk/download/pdf/79623650.pdf>
- [10] *Docker.io*; Available from <https://www.docker.com/what-docker>
- [11] *Cloudify.co*; Available from <http://cloudify.co/>
- [12] *Tacker – OpenStack NFV Orchestration*; Available at <https://wiki.openstack.org/wiki/Tacker>
- [13] *OpenO.org*; <https://www.open-o.org/>

Appendix 1. Cloud Native CMTS System Orchestrator TOSCA Model

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: TOSCA description of the CCMTS and RPD Orchestration
application.

repositories:
  docker_hub: <cabu docker hub url>

imports:
  - cabu_tosca:<cabu_tosca_type_def_url>

topology_template:
  ccmts_port:
    type: integer
    default: 8080
    description: REST port

node_templates:
  ccmts:
    type: cabu.nodes.software
    requirements:
      - host: ccmts_container
      - connection: RPD1
    interfaces:
      standard:
        create:
          implementation: scripts/api/install.sh
          inputs:
            repo: <git_repo_of_the_scripts>
            branch: {get_input: api_branch}
        configure:
          implementation: scripts/api/configure.sh
        start:
          implementation: scripts/api/start.sh
        delete:
          implementation: scripts/api/uninstall.sh
```

```

ccmts_container:
  type: cabu.nodes.Container.Executable
  artifacts:
    ccmts_image:
      file: ccmts:1.0
      type: cabu.artifacts.Image
      repository: <cabu_docker_hub_url>
  requirements:
    - storage:
        node: ccmts_volume
        relationship:
          type: tosca.relationships.AttachesTo
          properties:
            location: /data/ccmts/db

ccmts_volume:
  type: cabu.nodes.Volume

rpd1:
  type: cabu.nodes.phyendpoint
  requirements:
    -connection: ccmts
  properties:
    -macAddress: 40:00:00:00:00:04
  interfaces:
    Phyendpoint:
      initialize:
        implementation: scripts/api/initialize.sh
        inputs:
          repo: <git_repo_of_the_scripts>
          branch: {get_input: api_branch}
      configure:
        implementation: scripts/api/configure.sh
      connect:
        implementation: scripts/api/connect.sh
      disconnect:
        implementation: scripts/api/disconnect.sh

```

Appendix 2. Cloud Native CMTS System Orchestrator TOSCA Custom Types

```
Tosca_definitions_version: tosca_simple_yaml_1_0

Description: Definition of the custom types of cmts orchestrator

Node_types:
  cabu.nodes.Container:
    Derived_from: tosca.nodes.Root

    Attributes:
      Id:
        Type: string
      Private_address:
        Type: string
      Public_address:
        Type: string
      Ports:
        Type: map

    Properties:
      Ports:
        Type: map
        Required: false

    Requirements:
      -storage:
        capability: tosca.capabilities.Attachment
        occurrences: [0, UNBOUNDED]
        node: cabu.nodes.Volume
        relationship: tosca.relationships.AttachesTo

    capabilities:
      host:
        type: tosca.capabilities.Container
        valid_source_types: [cabu.nodes.Software]
        occurrences: [0, UNBOUNDED]
```

```
cabu.nodes.phyendpoint:
  derived_from: tosca.nodes.Root

  attributes:
    mac_address:
      type: string
    ip_address::
      type: string

  properties:
    mac_address:
      type: string
      required: yes
    ip_address:
      type: string
      required: false

  requirements:
    -connection:
      capability: tosca.capabilities.Endpoint
      occurences: [0, UNBOUNDED]
      node: tosca.nodes.Root
      relationship: tosca.relationships.ConnectsTo

  capabilities:
```

```

cabu.nodes.phyendpoint:
  derived_from: tosca.nodes.Root

  attributes:
    mac_address:
      type: string
    ip_address::
      type: string

  properties:
    mac_address:
      type: string
      required: yes
    ip_address:
      type: string
      required: false

  requirements:
    -connection:
      capability: tosca.capabilities.Endpoint
      occurrences: [0, UNBOUNDED]
      node: tosca.nodes.Root
      relationship: tosca.relationships.ConnectsTo

  capabilities:
    endpoint:
      type: tosca.capabilities.Endpoint
      valid_source_types: [cabu.nodes.Software,
cabu.nodes.Container.Executable]]
      occurrences: [0, UNBOUNDED]

  interfaces:
    cabu.interfaces.node.lifecycle.phyendpoint:
      description:
        this interface defines the lifecycle operations related to
the physical endpoints
      initialize:
        description: lifecycle initialization operation for
physical endpoints
      configure:
        description: lifecycle configuration operation for
physical endpoints
      connect:
        description: lifecycle connect operation for physical
endpoints
      disconnect:
        description: lifecycle disconnect operation for physical
endpoints

```


Fungible Virtualization Stacks

Refocusing on Optimization of Underlying Resources

A Technical Paper prepared for SCTE•ISBE by

Keith Alan Rothschild

Principal
kar@cox.com

Guy Meador III

Senior Solutions Architect
Guy.Meador@cox.com
Cox Communications
Technology Solutions Engineering,
6305B Peachtree Dunwoody Road
Atlanta, GA 30328

Brian Kahn

VP, Solutions Architecture
Sea Street Technologies
401 Edgewater Place, Suite 570
Wakefield, MA 01880
bkahn@seastreet.com

Summary

Industry focus on stack providers such as VMware and the various OpenStack implementations has prevented operators from focusing on the core problem: optimization of the underlying resources. Leveraging an integrated Policy-Driven Model-Based Service-Orchestration and Resource-Automation framework, we developed a solution where we can optimize the underlying resources, and where the span of control for any given virtualization stack is what is dynamically managed. This helps us address several of our critical use cases in addition to optimizing raw underlying resources including: (a) power management - turning off un-utilized resources, (b) services of bare-metal requirements for CDN, some real-time processing applications, etc., and (c) managing the needs of specific virtualization stacks (including isolation into secure network zones) required for specific VNF and/or IT domains.

Introduction

The systems we are designing now, and will be implementing over the next few years, need to support innovation over the next decade. It is useful to consider potential future scenarios when positing the span of control of each system. Consider two system environments: the consumer (home or business) and the data center.




Figure 1 - Universal CPE

There are competing formulations of what makes CPE “universal”, with many aspects depending on viewpoint. Consumers may want it to be a retail-available next-generation version of the cable modem, bringing WAN and LAN connectivity, and to allow for any number of additional features such as connected storage, compute, and/or to be converged with their home automation hub. The operator wants to minimize the number of devices deployed at a given customer location to support all of the operator’s products at that location – even to the point of deploying a single device - so that the operator neither maintains multiple versions of CPE to support any given customer, nor maintains multiple profiles of CPE across broad customer segments. When determining the functions to be performed on the CPE vs. elsewhere in the operator’s span of control, issues related to latency and the importance of continuity of service during loss of WAN connectivity may be important.

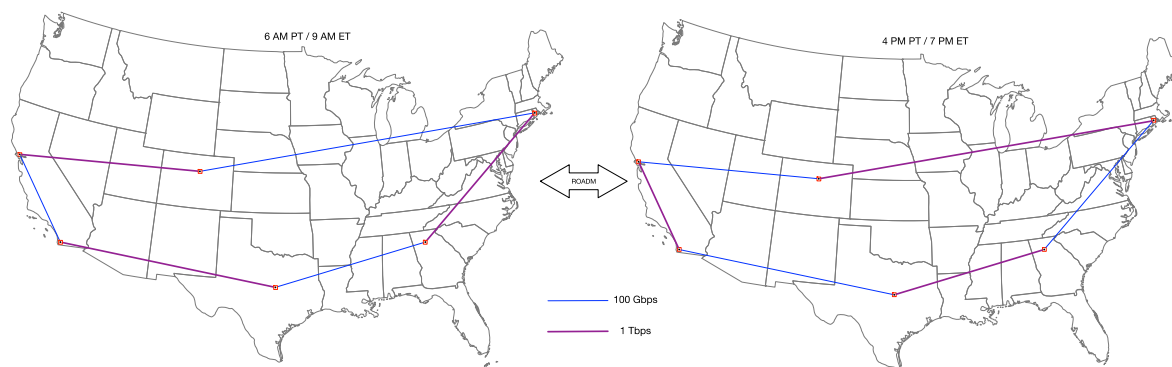


Figure 2 - Time-of-Day Network Reconfiguration

For the data-center environment, the ability to get the most benefit at the lowest cost is highly desirable. This may mean rebalancing optical links between data centers using Reconfigurable Optical Add-Drop Multiplexers (ROADM) to minimize the quantity of the most expensive optics required. To the extent that demand drives electricity and cooling costs, managing demand at operator locations throughout the day may become an important factor in reducing overall operating costs. Demand-shifting strategies may become important, such as making processing payloads “follow-the-sun” (to make use of solar/photovoltaic energy while addressing peak demands). Another use case is reduction of number of compute centers, for example, using excess capacity in the West or East to serve Central needs based on differences in consumption during different times of day, as shown in Figure 2.

With these scenarios in-mind, we need to identify the most important design principles, the universe of resources involved, and the optimal way to manage those resources over the long-term.

1. First Principles

1.1. Service Exposure Pattern and Approach

It is recommended to adopt a top-down, service-centric approach. Specifically, do not define offerings by the way they are technologically constructed and deployed (ex.: resources, stacks or physical location), but by service functionality, characteristics, interfaces, and consumption model. The service exposure pattern has usually been the domain of “cloud” orchestrators and software-only services. It is not identical to SOA, but shares many related concepts applied in a more generalized way, including constituent hiding, functional service interfaces, well defined attributes, and interfaces to control and automate services. It is applicable not only to software-based services, but should include hardware, network, and software. The key result: all services exposed through the pattern having the same functional interfaces and essential service attribute values are equivalent regardless of composition.

1.2. Composites and Composability Across Domains

A *Composite* is defined as an entity that contains a number of parts, called *constituents*, that are used by the entity as a whole. At any particular time, a constituent may be integral and dedicated for use by the whole or, in some cases, may be shared between composites. For this paper, composites and constituents are technological resources or services within, or across, one or more domains.

Composites can be assembled with one of two high-level approaches: static assembly or dynamic assembly. Static assembly implies a tight integration and association between the composite and constituents in a manner that is not intended to be changed over time, except as a result of undesirable events (ex. breakage) or design changes. Dynamic assembly may, indeed, have tight integration and association to the composite, but that association is more changeable in nature and occurs at any time in the normal course of composite operation (not only, for example, in response to constituent breakage).

Virtualized and software-controlled environments make dynamic assembly a more viable strategy as compared to a hardware-centric environment. Moreover, dynamic assembly of composites is essential to flexible, responsive, and agile service delivery. This gives rise to the Principle of Dynamic Composability: The aggregation relationship between a composite and its constituents should be changeable at any time in the composite's life cycle.

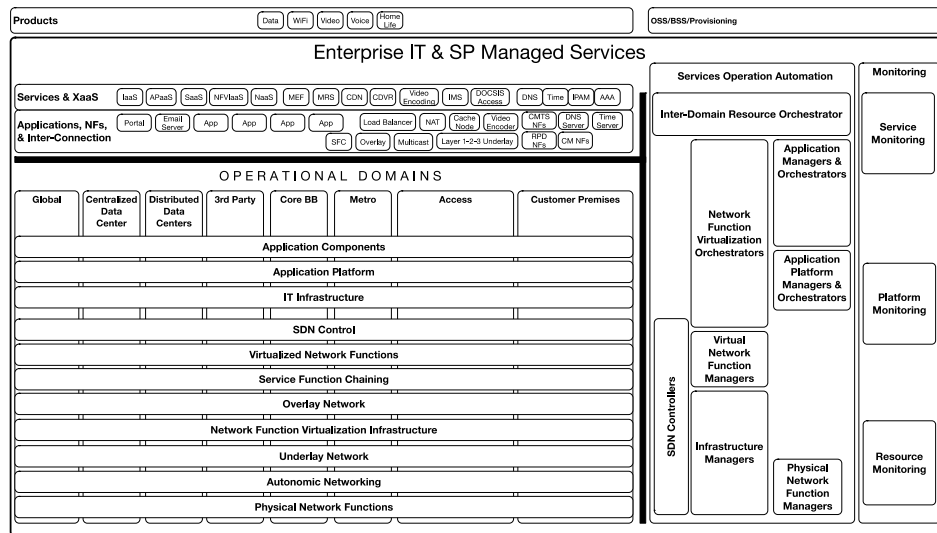


Figure 3 - Domain Diagram

Each domain and the services it provides should be dynamically composed of resources (and, possibly, other services), with such composition being created and changed over time under the control of a higher-level (sovereign) system, such as the Inter-Domain Resource Orchestrator (IDRO) proposed by BT:

“The development of rich NFV business cases depends on agreeing [on] each layer's responsibilities and the application programming interfaces (APIs) and service-level agreements (SLAs) through which each exposes its functionality to others. Otherwise, BT points out, nonsensical scenarios will arise, such as an NFV IaaS provider ceding resource allocation control in its own infrastructure to its customer's NFV Orchestrator (NFV-O). [...] Global resource management is parked in the NFV-O alongside service orchestration, but BT and others argue that it should be separated out into a fourth layer: the inter-domain resource orchestration layer.” (Chappel, 2015, p.3)

This arrangement enables global, end-to-end, policy-controlled decisions to be centrally determined and to take effect uniformly across and between the domains. Indeed, resources could be provided to one domain for a time and then reassigned to another domain, driven by global decisions and policies arising

from, for example, business needs or operational considerations that exceed the scope of any one of the domains.

1.3. Resource Abstraction

Expect resources to change over time as services, technologies and costs change. Focus on a stable, extensible approach to achieve thorough implementation and consistent reuse of policy and business logic. Resources will converge. Policy and business logic will be used to create business differentiation.

Incorporate abstraction and software infrastructure convergence so resource changes are seamless and straightforward. Make certain that any service can call upon, consume and control resources from any-and-all underlying infrastructure or cloud systems (i.e., compute, storage, networking, and future systems) so evolution is modular and straightforward.

Virtualization Stacks should be Composites supporting the Dynamic Composability Principle and Service Exposure Pattern, and should be managed with a layered operations automation approach that composes the resources and services and exposes the desired interfaces and functionality without revealing their composition.

1.4. Optimize for the Business First

Optimize for the business first and the infrastructure second. Expressly enable optimized recurring configurations, even if they are different in different organizations. Optimize infrastructure at the raw resource level (i.e., compute, storage, network links, PNFs/VNFs, stack licenses, etc.), not at the finished configuration level. Enable the business first. Execute, control and assure the full operational lifecycle of services for maximum efficiency, reliability and value. Do not stop at fulfillment.

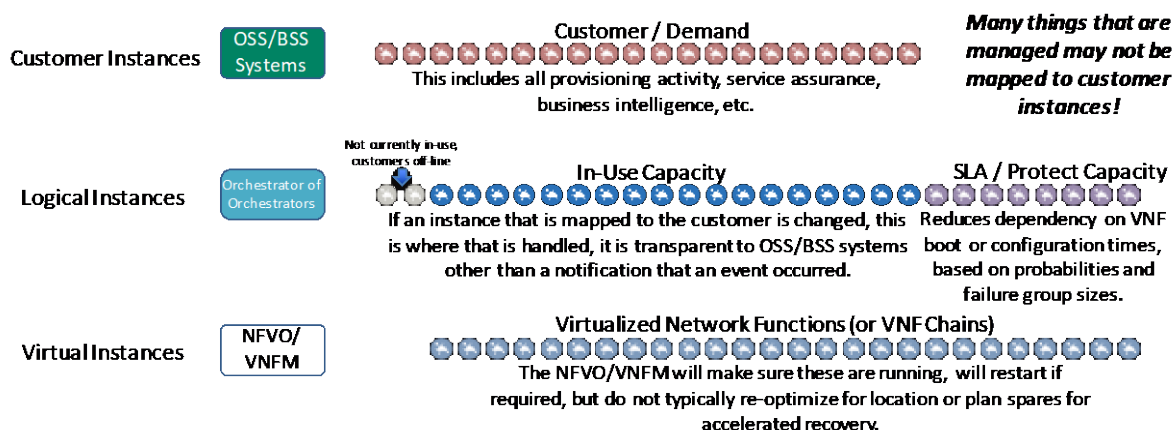


Figure 4 - The Relationship Between Virtual Instances and Customer Instances

1.5. Model Based Approach

Use models to operate the lifecycle of services and the resources. Manage the models to manage the service or resource. Encapsulate service lifecycle operations using a consistent, model-based service

induction methodology so operations logic is portable and can be understood and simulated alone or in context.

Control operations via design-able, reusable policy models. Encapsulate business logic in policies. Design policies once and reuse them across multiple services, offers and applications as needed. Policies scale in a way that people cannot.

Expect hybrid cloud, inter-cloud, NFVI, and hyper-converged infrastructure models to present and reinvent themselves over time and be ready to include them. Choose a methodology and a platform to allow modular flexibility to progressively add innovative resource types as required.

Perform service level cost modeling before and after implementation to identify and prioritize business, operational, financial and compliance factors. Let the numbers drive the resource and operational decisions. Remember, the cost to automate should be compared against the cost of not automating, including the development and support of M&Ps which will often have a comparable cost to automation independent of how often the M&P might be utilized. An M&P that is rarely used is more likely to introduce uncertainty than the corresponding automation.

2. Resources

For the purpose of this paper¹, we will address two major classes of generalized telecommunications resources, separating them into compute (processing, storage, etc.) and links (networks). Improvements in hardware capabilities have resulted in general purpose hardware being able to displace specialized hardware needed for supporting network functions. The ability to deploy virtualized network functions (VNFs) and the separation of the data plane and control plane in networking enables the creation of physical (underlay) networks and the use of software to control the realized (overlay) networks in a manner that can be flexibly reconfigured as needed. These trends are both in-line with what we would expect from general technology evolution - but are far from where this evolution will end.

Responsible architectural analysis and design will contemplate both the implications of these evolutionary steps, predict the likely vectors of subsequent evolution, and determine if it is worthwhile to predispose solutions to support any specific vector. Optimal placement of compute and link capacity is moving from highly specialized fixed components to components that are generalized and more flexibly configurable, creating a rather complex resource utilization problem.

Placement of general purpose (compute) resources should be contemplated for these variants: locally on customer premises (CoP, Compute on Premises), regionally (as it relates to the provider network), centrally (as it related to the provider), and remotely (third-party). Similarly, (network) links should be

¹ This differs from normal treatment in that what is normally referred to as compute is described as processing, and the term compute is used to describe both processing and storage. Similarly, a third class is often referred to as networks, whereas here, the term links is used instead. Components such as switches or routers are often classified as “network” rather than as compute, however, they are packet processors, and as they become virtualized, may be replaced by generalized compute. As such, the paradigm of “compute, storage, and network” is replaced with “compute and links”, which is more germane to this analysis.

contemplated as they connect the customer premises (access), connect resources (data-center), connect regional data-centers (metro), connect centralized data centers (backbone), and provide for connection to third-parties (interconnects).

Service Provider facilities encompass a variety of configurations and locations into which generalized compute resources are placed. These facilities range from large, centralized data centers to regional data centers, to edge facilities. The evolution of the kinds of generalized compute resources that are placed at each of these facilities, and the mix and scale of compute resources within them will be dynamic and ever-changing (ex.: hybrid cloud, inter-cloud, NFVI, and hyper-converged infrastructure, etc.). In tension with these trends is the goal to select and manage compute resources in a manner that makes them usable by the largest set of services over time without the need to perform manual/physical changes; once put in place, the same compute resources should be usable by any service's software elements as long as the resource meets the minimum capability profile for the software element.

At present, CoP devices are likely to be x86-based, but, at scale, more cost-effective solutions are expected to emerge. Over time, device capability will improve, with the potential for device availability through retail channels and device upgradability (ex. compute capacity). These dynamics will further fragment the uniformity of capabilities across the device population. Independent of capability, these devices will be expected to support a combination of services, as subscribed to by the customer.

As denser alternatives to x86 CoP devices become popular, the desire to benefit from this in the datacenter will begin to gain momentum, and we will see a combination of hybrid virtualization stacks and virtualization stacks dedicated to dense packet processing.

More capable CoP devices may be able to reduce the traffic demands on the access network, and depending on business models and who provides the CoP device, the bandwidth allocation on the access network may need to be dynamically adjusted to accommodate placement of network functions on the CoP or remotely.

The concept of Universal CPE (uCPE) is closely tied to Compute on Premises (CoP). The understanding of what constitutes uCPE may lay within the perspective of the stakeholder, and even then, could be differentiated by the role. For Service Provider (SP) product managers, the universality of uCPE could mean that it will support any combination of services that SP offers, for instance, some combination of video, voice, internet connectivity, home security and/or home automation services. Supply Chain or Field Service stakeholders might also include flexible (possibly modular) support of a multitude of WAN connectivity options. Additionally, we can expect the desire for this to be equipment provided by the customer and available in retail outlets, in order to reduce expected capital outlays. For retail-available components, the relationship between the connectivity aspect and CoP aspect of CPE may be severed, or at least may include the ability to modularly supplement CoP with devices such as NAS and Home Automation Hubs.

The desire for CoP, from the consumer's perspective, goes beyond simple optimization of compute and links, but includes capabilities such as caching of IoT data or the desire to improve performance of automation capabilities by having some of the decision processing occurring on-site. Retail providers of CoP will want to optimize cost at volume, and even if x86-based devices were the preference of the SP community, the OTT and IoT communities would likely move to a more cost effective dense packet processor. As VNFs evolve towards a micro-service orientation, the CoP domain is likely to be a hybrid of SP and customer-provided compute resources; therefore, convergence is expected across SP-services,

OTT-services, and IoT device capabilities. The question isn't whether the CoP will evolve to support multiple processing paradigms, but whether the edge-compute and centralized-compute environments will also embrace supporting a hybrid of x86-based and dense-packet-processing-based paradigms. Similarly, we will have to consider management based on the availability of equivalent VNFs for each of the compute paradigms. The need to satisfy multiple compute domains across multiple service offerings that are provided and managed by multiple parties will greatly challenge the practice of treating compute as a resource under the control of the SP NFV architecture.

An implication of the convergence of CoP usage is the provider/consumption model for CoP itself. Consider that businesses and residences are inherently dependent on the managed delivery of services such as power, water, and sewer. In fact, it is the rare circumstance that these services are offered and consumed as anything other than fully-managed services. It may be that within a short amount of time businesses and residences will also become inherently dependent on fully-managed CoP, used as assumed infrastructure. The main aim of such a managed CoP service is to provide available infrastructure positioned between the home network and external networks for use by multiple services and parties.

3. Resource Consumption

Many applications utilized by enterprises, which may be the most common class of applications supported in the IT environment, require a highly available infrastructure. A highly-available infrastructure drives the need for expensive components. Virtualization stacks that offer high-availability and nearly transparent VM migration are likely best suited for such enterprise applications, but, this approach, generally, comes at a cost premium as compared to an approach where the virtualization stack does not offer those capabilities and the application is designed accordingly.




Figure 5 - Enterprise Application Model vs. Cloud Application Model

The Cloud Model, which is becoming more prevalent, especially in the service provider (SP) space, shifts the burden from the infrastructure to the application (depicted above). The application is designed to account for the fact that a fault may occur in the infrastructure and it must handle these additional failure-case scenarios itself. The virtualization stack in the cloud-model is responsible for making sure that the elastic infrastructure demands of the application, especially those used for self-healing, are handled transparently, and that seemingly immutable infrastructure is supplied to the application on-demand.

Significant application re-design is required to move an application from the Enterprise Model to the Cloud Model, and the expense, which may or may not be capitalized as NRE, may not be warranted for

many applications. For those situations, continued use of high-availability virtualization stacks becomes appropriate.

On the other hand, the nature of the Enterprise Model may limit the availability potential of the application to roughly 99.9%, which meets the requirement of many IT applications, but doesn't meet the requirement of many SP applications, which require either three-nines-five (99.95%) or five-nines (99.999%) availability.

In keeping with the 'Business First, Resources Second' optimization principle, we believe it is important to support a reasonable number of environments such as the ones above so that the various constituents gain repeatable, optimized environments that suit their specific needs.




Figure 6 - Desired Relationship of Stacks to Resource Pools

Many virtualization stacks support multi-tenancy, and this is something that should be employed with great caution, especially with service provider applications that leverage and depend on the SDN capabilities integral to these virtualization platforms. We must avoid putting operations in the position where one tenant application requires a specific version of a virtualization stack that conflicts with another tenant application's requirements due to known compatibility issues. Additionally, secure isolation enforced at the network level may be more realistic than performing (or relying on third-parties to perform) security audits on the virtualization stack.

In this approach, virtualization stacks should be treated as resources, just like compute and storage, that are consumed based on policy and the actual requirements of the service being placed. OpenStack, VMWare, containers, bare metal and future stacks should all be enabled to the degree they are required by the services. Resource optimization will still occur, just at the layer underneath stacks, through the use of common processing and link resources to the degree permitted by the services.

From an equipment standpoint, the goals are to have the smallest possible footprint required to meet the demands, and to maximize utilization and reusability while minimizing 'flavors' of equipment, and, therefore, to minimize CapEx. A related goal is to ensure services are portable across different infrastructure platforms and providers so the SP is free to take advantage of lower cost options as they are available.

4. Automation

Automation, Workflow and Orchestration are utilities that are manually operated or reactively triggered to achieve a specific result. They create leverage for human operators and significantly reduce the manual

steps required in a given process. Today there are common tools for deployment, also called fulfillment, of cloud services.

While their names might lead one to believe otherwise, these utilities do not cover the full operational lifecycle of a service. Today, their actions are limited to deployment, single-app scaling and application-based fault tolerance at best. This is not to say they are not useful - only that they aren't as broad as they may seem.

A comparison of scripts, automation, workflow, orchestration and sovereign systems is shown below in Figure 7. The top line in the diagram shows these utilities compared to one another in terms of suitability for single vs. multiple tasks, synchronous (potentially single-threaded) vs. asynchronous (potentially multi-threaded) operations, and their abilities to incorporate rules and policies. The lines below address the levels of information with which they can interact. Finally, the bottom line shows the level of intelligence they can encapsulate.

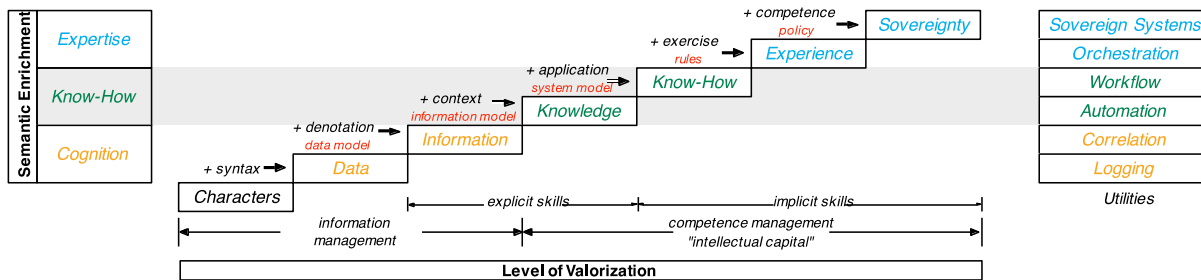


Figure 7 - Comparison of Scripts, Automation, Workflow, Orchestration, and Sovereign Systems

The hierarchy of automation that begins with scripting. Scripting is a series of commands run in a specific sequence, and is kicked off by calling that script. Automation builds upon scripting by explicitly removing the manual need to invoke the script. Automation can be invoked based on some scheduling system or trigger. Scripting and Automation are typically focused on accomplishing a specific task (or closely related set of tasks) in a single- threaded or synchronous manner.

Automation is a utility that exists at many levels in the software and hardware of a cloud. It provides rule-based actions and reactions for a single task. Automation is very useful, but it is generally limited to a set of known actions. Automation can perform poorly in the face of previously unseen conditions (exceptions). Automation systems cannot be used effectively for integration because they only focus on a single task at a time.

Workflow builds upon scripting and automation by executing a series of loosely related tasks, retaining the single-threaded/synchronous nature. It is a utility that provides rule-based actions across multiple tasks. Workflows are commonly used to execute processes. They are single threaded per task. Like automation, workflow systems can perform poorly in the face of exceptions.

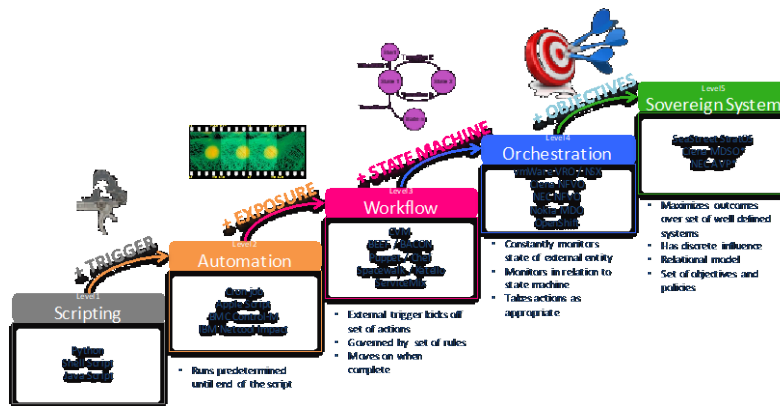


Figure 8 - Evolution between Scripts, Automation, Workflow, Orchestration, and Sovereign Systems

Orchestration deals with the relationship between multiple workflows in a multi-threaded asynchronous manner, and can deal with invocation at multiple points of a process based on rules.

Orchestration systems are commonly used for deployment of applications across cloud resources. They might spawn a thread to acquire compute resources, and deploy and configure software on those resources, while another thread acquired storage and configures it, all for the same application. While orchestration systems typically terminate an activity when fulfillment is complete, there are some that can be called repeatedly to affect scaling. Beyond deployment, initial configuration, scale, and decommissioning, orchestration systems do not effectively address other aspects of the operational lifecycle of a service.

4.1. Resource Management

Model-Based Systems (MBS) are being deployed to close the gaps in automating fragmented manual operations. MBS are replacing Run Book Automation (RBA) systems, which traditionally are fragmented and costly because they rely on manual scripts and templates that represent legacy software. Furthermore, RBA is constrained to only those operations for provisioning, which are often specific to one infrastructure domain (i.e., Compute or Network or Storage and their related configurations). The one-to-one relationship of RBA binds automation to a vendor's technology implementation. Not only is the SP wedded to that vendor, RBA also limits automation to simply those provisioning transactions for a specific infrastructure technology. RBA requires the vendor's customer to maintain multiple teams of operating experts to administer and manage its various systems. These teams are repeated in each technology domain. MBS solve the RBA problem through advanced abstraction by de-coupling and converging automated operations. Through automation, MBS significantly reduces OpEx and improves resiliency, control and availability of those services consuming the underlying infrastructure.

It is typical in the NFV and IaaS/PaaS spaces for solutions to provide an end-to-end ecosystem that can manage almost any task required. The question becomes, even if they can, should they? What is the appropriate set of criteria to make those decisions? One could make the argument that each link-domain and each compute-domain should be managed by an independent domain-specific controller (orchestrator) and that coordination between these domains should occur using a higher-level (sovereign) system. This does not mean that these domain specific systems don't communicate with each other, but there would be no static configuration between them and no clear master among them.

As resources increasingly move to orchestrated virtualization, operators will see the number of resource-oriented orchestration systems increase in their network. For most domains, there are orchestrators that handle that specific domain; but, there is also a need for a higher-level orchestrator which is orchestrating the end-to-end as a whole. Similar to the situation for early automation and workflow-based control systems, directly integrating third party components tends to increase capital burden over time (left-side of diagram below). That approach becomes an integration nightmare and locks in domain orchestrators. By using a higher-level cross-domain orchestrator, policies can be defined generically and then applied down into the domain orchestrators. This preferred approach enables policies to span domain orchestrators, allowing for quicker introduction of new domain orchestrators over time with reduced integration overhead and minimal redesign.



Figure 9 - Integration Platform

The desirable approach is to have an integration platform intermediate between resource orchestrators (right-side of diagram above). It is necessary to avoid use of a simple automation or workflow-based integration platform for this purpose; use, at a minimum, an orchestration-capable platform to act as this integration layer.

As offers become more complicated, a cross-domain orchestrator is needed. Offers like in-home services require a handful of additional service provider services in order for them to work properly; examples include physical network, overlay network, firewall, CoP or CPE device, etc. The ability to coordinate the service across domain orchestrators is critical to understanding the overall SLA for the end customer offer.

Service deployments are becoming more complex and need to be deployed faster in order to keep up with demand. The ability to abstract common functionality into reusable and highly-reliable service models enables deployment to be faster and creates a better service experience. These models should be abstracted from the underlying stacks to make them portable across infrastructure and component services.

The ability to dynamically adjust to real-time telemetry to ensure the SLA and service needs being met at all times is going to become very critical. The number of services that will need to be supported by service providers is increasing quickly and the budget for operations is only going to shrink. Autonomous Operations enables the ability to automate known tasks, scale up repetitive tasks and execute them at computer speed with customer accuracy. The days of having people manually monitor and operate services is going to come to an end given the size and complexity of what is coming.

4.2. Sovereign Systems

Sovereign systems are top-level controllers that operate multiple services together in context with each other and across varied sets of infrastructure. Such systems are almost always model-based and use models to analyze, predict and manage the operation of services on top of ‘real managed things.’

Sovereign systems combine four key elements: (1) stateful awareness, (2) converged declarative telemetry, (3) late-binding policy, and (4) abstraction. A runtime, state-aware process that understands the needs of all of the services, capabilities and capacities of all the resources, and the requirements of the policies simultaneously. Converged and declarative telemetry related to each service. (Note: declarative telemetry means the service declares the telemetry that is important and relative to it and consumes that data specifically. Contrast this with the more typical raw event data that has missing context and requires correlation.) Policy control is applied continuously to manage the automated processes across the service or resource lifecycle. An abstraction layer separates the traditional infrastructure systems and controllers from models that define the service’s operational lifecycle.

Sovereign systems sit above traditional infrastructure systems and controllers (e.g.: scripts, automation, workflow and orchestration) and below traditional OSS/BSS systems and portals. They do not replace these systems. Sovereign systems typically receive an order from an OSS/BSS system or portal for a service and then work through an API fabric to operate infrastructure systems and controllers, to gather telemetry, and to exert control on services and resources to fulfill and lifecycle operate the ordered service, enabling these systems to operate autonomously.

Key differentiators of Sovereign Systems include:

1. They operate continuously, covering the full service-lifecycle.
2. They run many services simultaneously, in context with each other.
3. They operate under policy control and make policies real in the services and infrastructure
4. They are stateful and understand the meaningful condition of each service
5. They understand resources and capacity and know how much capacity remains
6. They contain a real-time ‘as-built’ graph of the cloud components they are responsible for and a real-time dependency map.
7. They understand service level health, and can remediate, migrate, scale, and move services
8. They can handle complex multi-stage activities like upgrade
9. Sovereign systems can be programmed to learn from experience, and services operating under a sovereign system can learn from each other.

Sovereign systems make use of automation and orchestration utilities that exist southbound where appropriate. For instance, it is common for a Sovereign system to use automation and orchestration utilities presented by stacks like VMWare or OpenStack.

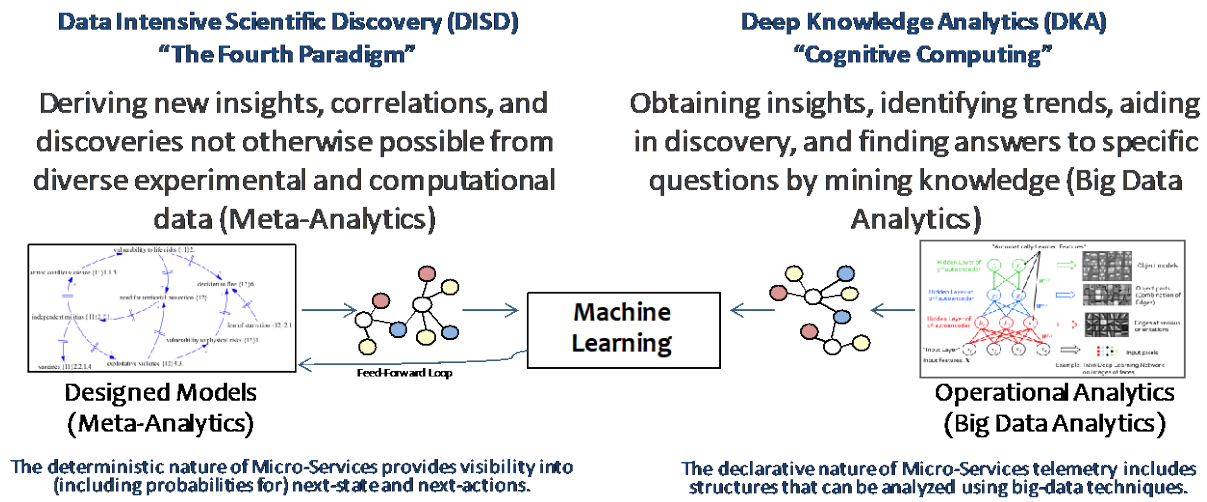


Figure 10 - Dimensions of Machine Learning

Through modeling, sovereign systems allow methods and procedures (M&Ps) to be written in software rather than on paper, and they allow these M&Ps to be executed with computer speed and reliability.

Where historic cloud approaches have previously attempted to address deployment and configuration only, a sovereign system addresses the full lifecycle of a service: order, deploy, assure health, assure compliance, remediate faults, failover if required, upgrade, move, port, and destroy/delete. Because they address the full lifecycle, Sovereign systems have the potential not only to mitigate the additional OpEx that comes with virtualization and cloud, but also to change fundamentally the cost of business by eliminating manual work.

Sovereign systems also need to address very specific concerns related to the potential to accelerate and amplify the effect of failures. Domain orchestrators may not act as expected, either due to malfunction or issues in the way they were modeled. The sovereign system must be able to predict the degree of impact that is expected, monitor the implementation, detect when such an anomaly occurs, and accommodate the anomaly as it can. It should employ circuit-breaker logic with corresponding alarming and notification, rather than enforcing changes that could spiral out of control, or that might result in excessive reconfiguration and possibly reconfiguration loops.

A point of clarification is that Sovereign Systems should not be confused with the concept of data sovereignty. Data Sovereignty is a policy that states where data can be stored or routed or used. For instance, the Australian National Healthcare System has a data sovereignty requirement that the medical information for Australian citizens can only be stored, routed and used within the nation of Australia. You may use a Sovereign System as one way to accomplish this (there are others), but otherwise the concepts are not connected.

5. Proof of Concept

We created 3 different proof-of-concepts (PoCs) with a Sovereign system from Sea Street called StratOS. These PoCs were designed to prove out the concepts described in this paper across different types

services, which includes service provider services, customer-facing offerings and implementing/integrating multiple stacks together across domains in the IT operations domain.

For the service provider service we used video encoding. This PoC deployed, configured and managed the operational lifecycle of encoders deployed on bare-metal blade servers. The StratOS system modeled the behaviors of the encoders, video streams, video quality monitoring devices and the infrastructure deployed on. StratOS models each component in the system and aggregates these components into higher-level models that provide a combined higher-level view of the end-to-end server. This means that the StratOS system has a complete and fully stateful awareness of the entire encoder system. This includes the encoder location, configuration, mapping of the health of each video stream both upstream and downstream of the encoder. Once deployed, the system collects telemetry from the infrastructure, encoders and video quality devices and converges this data to determine the health of each individual stream, the infrastructure, encoder application, network, etc. Policies are then applied to the converged data to determine what action, if any, is required to be taken. For example, if the output video quality device reports “no video”, the source quality device would be checked to see if the issue is an encoder issue or upstream from the encoder. Given these data points, StratOS would then take actions, as needed, to remediation to solution such as failover to a standby encoder.

For the customer-facing offering we focused on building a managed router solution with virtualized network functions (VNFs). This PoC created the overlay network for a managed router customer utilizing VNFs (i.e. firewall & content filter) from multiple vendors. The StratOS model for this deployment created each VNF individually and provided the required configuration to connect the service chain together. As with the encoder PoC, the StratOS model contains higher-level models that aggregate lower-level models together enabling StratOS to operate the operational lifecycle at both an end-to-end offer level and component level. This late-binding of network services proved the ability for a Sovereign system to be able to create dynamic offers for customers based on individual needs. In the PoC we created multiple types of SLAs that can be applies to each customer to determine the HA configuration of the service chain, determined placement of services and types of VNFs to be used. Once deployed, StratOS collected telemetry from each of the VNFs, network and the CPE devices to determine the health of the end-to-end offer. When a failure in the offer is detected, StratOS was able to failover the customer to a new service chain within seconds and automate the remediation of failed service chain.

For implementing/integrating multiple stacks together we looked at the IT Operations challenges where private local networks use one type network and the operations backbone uses another. For this PoC StratOS modeled the generic requirements of an IT operations system; a set of VMs and network access to NTP, DNS, the Internet and other private networks. StratOS then determine the correct configuration required within in each environment in order to create the proper routes in the multiple networks showing that a Sovereign system can work across and connect multiple environments.

In each of these cases, the StratOS models are designed to be portable across infrastructure through their abstraction layer. In some cases, we ran the same PoC across different infrastructure stacks and vendors without changing the models or policies defined in StratOS.

During the proof-of-concept work we did, we examined the cost savings for two of services: (1) managed router and (2) video encoding. The result of these analyses showed additive operating savings beyond the benefits of virtualization and standardization of 41% and 31%, respectively. These projected savings were gained through the efficiencies of automated failure detection and remediation, fewer human errors and consistent operations engendered from lifecycle automation.

Services that don't scale often and are stable will see closer to a 30% cost savings through automation, while services which are very dynamic, tend to be less stable, or are scaled often will have closer to a 60% cost savings.




Figure 11 - Factors that Impact Potential for Operational Savings

Automation of the Operational Lifecycle of Services reduces operating costs by:

- Automating manual tasks typically taken on by tier 1 and 2 operations staff, these include typical daily tasks such as monitoring and telemetry convergence, and less frequent tasks such as failure detection, troubleshooting, cause determination and remediation.
- Automating deployment of applications and services, including upgrades and mass configuration changes (e.g., re-IPing a service), through common models that contain instance specific properties.
- Automating security requirements/audits, capacity management, service/datacenter load balancing and other maintenance tasks performed by operations.
- Automating the recovery of a failed offer, application or service by whatever means necessary to get the offer or service up and running again. This includes changes to compute instances, network configuration and more.
- Automating the data flow between the offers, OSS and BSS systems reducing the need for operational staff to enter or translate data to multiple systems.
- Automating the collection of log and other data needed to analyze failure and determine what actions need to be taken to prevent similar failures in the future.

An interesting phenomenon we ran into when evaluating these cost savings relates the relative fallacy of contemplating these as true savings; they may be more accurately described as cost avoidance. Companies typically have a finite pool of resources that they attempt to spend in an optimal fashion, with far more demands than the budget can accommodate. Reducing the cost per action will often allow the company to do more using the same budget, thus a “savings” or “avoidance” of 33% might be more realistically viewed as increasing the operational capacity of the organization by 50% (you can do half again as much as before) and a cost savings of 67% can be viewed as tripling the operational capacity of the organization!




Figure 12 - Factors that Impact Potential for Operational Expense Cost Avoidance

While claims of 33%-75% may be hard to believe. Being able to operate 2x – 10x the amount of equipment (the equivalent of 50%-90% “savings”) is easier for people to understand.




Figure 13 - Using Virtualization-Enabled Automation to Empower the Workforce

The reality may lay someone in between, for example, leveraging virtualization-enabled automation may enable a smaller workforce to handle 5x the amount of equipment with half the expense rather than achieving a 90% cost savings.

Conclusion

Although virtualization increases the complexity of solutions, it enables highly flexible automation. Initially, only a few things may leverage this virtualization-enabled automation, however, it will become a pervasive technology paradigm. To prevent virtualization for the sake of virtualization and over-integration of technologies from disparate domains, it is important to have a future-state vision that can be used to guide decisions. Although it is tempting to describe the benefits in terms of cost savings or cost avoidance alone, the theoretical costs would likely have been prohibitive, and an expression of operational efficiencies (do-more-with-less) may be more realistic.

Abbreviations

API	Application Program Interface
BSS	Business Support Systems
CDN	Content Distribution Network
CoP	Compute on Premises
CPE	Customer Premises Equipment
IaaS	Infrastructure as a Service
IDRO	Inter-Domain Resource Orchestrator
IoT	Internet of Things
IT	Information Technology
LAN	Local Area Network
M&P	Method & Procedure
MBS	Model Based System
MBSE	Model Based Systems Engineering
NAS	Network Attached Storage
NFVI	Network Function Virtualization Infrastructure
NFVO	Network Function Virtualization Orchestrator

OSS	Operations Support Systems
OTT	Over-The-Top
PaaS	Platform as a Service
PNF	Physical Network Function
PoC	Proof of Concept
RBA	Run Book Automation
ROADM	Reconfigurable Optical Add-Drop Multiplexer
SLA	Service Level Agreement
SOA	Service Oriented Architecture
SP	Service Provider
uCPE	Universal Customer Premises Equipment
VM	Virtual Machine
VNF	Virtual Network Function
VNFM	Virtual Network Function Manager
WAN	Wide Area Network

Bibliography & References

Chappel, Caroline. (2015) NFV MANO: What's Wrong and How to Fix It. Heavy Reading 13(2).

Chen, Y., Qin, Y., Lambe, M., & Chu, W. (2015, November). Realizing network function virtualization management and orchestration with model based open architecture. In Network and Service Management (CNSM), 2015 11th International Conference on (pp. 410-418). IEEE.

Garcia-Gomez, S., Jimenez-Ganan, M., Taher, Y., Momm, C., Junker, F., Biro, J., ... & Strauch, S. (2012). Challenges for the comprehensive management of Cloud Services in a PaaS framework. Scalable Computing: Practice and Experience, 13(3), 201-214.

Gevorgyan, A., Krob, D., & Spencer, P. (2016, July). Functional Analysis and Design Approach for an Optimal Virtual IP Multimedia Subsystem (IMS) Architecture. In INCOSE International Symposium (Vol. 26, No. 1, pp. 1463-1476).

Ivezic, N., & Srinivasan, V. (2016). On architecting and composing engineering information services to enable smart manufacturing. Journal of computing and information science in engineering, 16(3), 031002.

Ortiz, A. M., Rios, E., Mallouli, W., Iturbe, E., & de Oca, E. M. (2015, September). Self-protecting multi-cloud applications. In Communications and Network Security (CNS), 2015 IEEE Conference on (pp. 643-647). IEEE.

A Practical Approach to Virtualizing DOCSIS 3.1 Network Functions

A Technical Paper prepared for SCTE•ISBE by:

David S. Early
Data Scientist
Applied Broadband, Inc.
2741 Mapleton Ave
Boulder, CO 80304
720-470-7460
david@appliedbroadband.com

Paul E. Schauer
Distinguished Engineer
Comcast
183 Inverness Dr W,
Englewood, CO 80112
303-372-1215
paul_schauer@comcast.com

Jason K. Schnitzer
Founder
Applied Broadband, Inc.
2741 Mapleton Ave
Boulder, CO 80304
720-838-4465
jason@appliedbroadband.com

Introduction

There is broad consensus amongst the networking community that programmable network architectures (including SDN and NFV) represent the next stage of connected infrastructure evolution. Benefits for providers are many, stemming from innovative new approaches that re-factor the development and deployment of networks and services.

The service provider's ability to rapidly and continuously develop and deploy new software and tools necessary to realize these objectives is central to the success of this model. Not only will the network see technical change with the introduction of virtualization, but the service provider's organization will require significant changes as well.

This paper presents a structured approach to the evolution of virtualization within the service provider's practice. We establish criteria for the virtualization of network functions within the broadband access network. In doing so, we provide system considerations for the tradespace between the use of centralized and distributed deployment architectures. We provide an overview of Service Provider DevOps (SP-DevOps) and how it can be applied within the cable service provider environment for the continuous deployment of virtualized networks and services.

We present a practical example to illustrate key concepts, developing a simple microservice that provides an implementation of the CableLabs DOCSIS 3.1 Common Collection Framework (DCCF) software system. Use of this this DOCSIS 3.1 microservice will be shown using a container architecture within a cable operator provider's cloud infrastructure.

Service Provider DevOps

In contemporary software development, it's almost impossible to ignore the term DevOps, a portmanteau of "Development" and "Operations." DevOps represents a practice within Information Technology (IT) that defines an organizational model for collaboration between software development and software operations.

Though DevOps is a popular topic in software engineering research, it currently lacks a widely accepted standard definition. The following is a common definition of DevOps based on a novel analysis of peer-reviewed articles:

"DevOps is a development methodology aimed at bridging the gap between Development and Operations, emphasizing communication and collaboration, continuous integration, quality assurance, and delivery, with automated deployment utilizing a set of development practices". [1]

The goal of DevOps is to "bridge the gap" between the internal functions responsible for producing software and the functions charged with running it. Major Internet software, services, and cloud provider companies have adopted DevOps practices to improve the velocity and quality of software development and delivery within their practice [2][3].

With success as a model for software development in large internet enterprises, it has been proposed that DevOps may be useful within the service provider environment as well (including DOCSIS access networks). Though many of the DevOps concepts readily apply, it is worth examining some key differences between target infrastructures. Broadband access service providers have the following distinct

requirements, relative to most other organizations implementing DevOps principles. Notable differences between service provider access network environments and Internet enterprise networks include:

- *A high cost of operation in terms of time and human/financial resources due to the physical management of distributed network nodes. Access Network service providers supervise a management domain that extends beyond a centralized facility (e.g. Data Center), to intelligent devices at the edge of the network with execution environments extending through the cable plants and into each broadband subscriber's home.*
- *Limited visibility of distributed network and service states that make it difficult to assure the Quality of Experience (QoE).*
- *Difficulties in pinpointing the cause and location of problems (troubleshooting) and debugging.*
- *Difficulties in deploying services quickly and frequently, e.g. due to the validation of service integration or regression testing.*

This short list of maladies, shared within large scale Information Systems (ISs) development and operations, is not unlike those faced by service providers when capturing the limitations born of current network management systems and organizational practices. In recognizing this, service providers and network engineering researchers have begun to evaluate the applicability of DevOps practices to analogous challenges of scale and complexity inherent in both traditional (hardware physical element centric) and more contemporary (software virtual element centric network architectures).

It has been noted that though similar management automation goals exist within the programmable network infrastructure domain, service providers face challenges unique to access network systems. Where DevOps was formed from IT organizations working within concentrated data center environments, large service providers operate an inherently more geographically distributed system. Project UNIFY [4] proposes a list of four key characteristics of telecommunications networks making them different from IT organization data centers. These include:

- *Higher spatial distribution with lower levels of path and equipment redundancy;*
- *High availability;*
- *Strictly controlled latency;*
- *Larger number of distributed datacenters.*

However different, greater similarities have motivated new research into the applicability of DevOps principles in a service provider environment. When applying key concepts of DevOps to the Service Provider domain, the integration of these models took the form of Service Provider DevOps (SP-DevOps). The adjusted model proposes the following aspirations:

- *Iterative Development / Incremental feature content*
- *Continuous deployment*
- *Automated processes*
- *Holistic/Systemic views of development and deployment/operation.*

The UNIFY project represents a first attempt to codify SP-DevOps practices in the form of a standards-based approach. The initiative aims at applying DevOps concepts to telecom operator networks and supporting the idea of fast network reconfiguration.

Microservices:

Building Modular Distributed Applications

Microservices are small, purposeful modules of software executing in a distributed environment. Microservices can be used to build larger applications (northbound) that implement specific business solutions through interactions with the network operator's back-office infrastructure. The microservices architectural model is heralded to be a more agile approach to complex system development while being better aligned with the goals of SP-DevOps and a continuous deployment model.

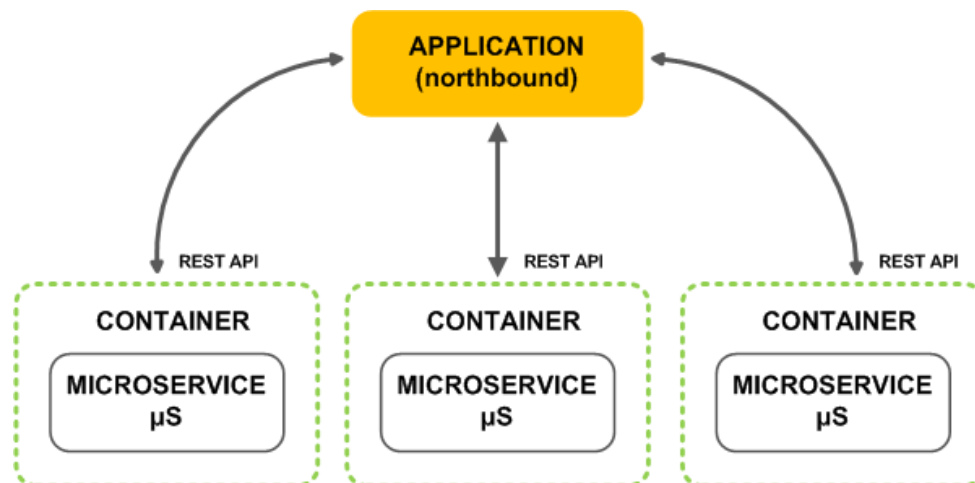


Figure 1 - Building Networked Applications from Microservices and Containers

Containers:

Practical Architecture for Network Virtualization

Containerization, also called container-based virtualization, is an OS-level virtualization method for deploying and running distributed applications without launching an entire Virtual Machine (VM) for each application. Instead, multiple isolated systems, called containers, are run on a single control host and access a single kernel.

Microservices are often executed in containers. Containers offer a new form of system virtualization that significantly reduces the resource cost and complexities of their VM predecessors. Where VMs embody a complete executing environment and copy of the OS, containers perform execution isolation (virtualization)

at the OS level. In this way, a single OS instance can support multiple containers, each with a microservice executing inside [5].

The Rise of REST

The REpresentational State Transfer (REST) protocol has risen to become the de facto choice for web development with the evolution of Inter-Process Communication (IPC) to make use of text-based serialization formats, like XML and JSON. Protocols such as SOAP allowed IPC across HTTP, and soon web developers were not just building web applications that served content to browsers, but web services that performed actions and delivered data to other programs. This services-based architecture proved to be very powerful, as it eliminated dependencies on shared code libraries, and allowed application developers to further decouple their application components. The SOAP protocol and the related WS-* standards soon became increasingly complex and heavily dependent on specific implementations in application servers, so developers migrated to the much lighter-weight REST protocol. As the use of mobile devices exploded, and as web interface development switched to AJAX and JavaScript frameworks, application developers started to make extensive use of REST for transmitting data between the client devices and the web servers.

In popular Software Defined Network (SDN) architectures such as OpenDaylight [6] and ONOS [7], REST has become a critical infrastructure component, extended by the IETF to support configuration management concepts by offering a lighter-weight transport to NETConf's more cumbersome transactional model [8]. As we will see in the example provided, the use of REST lends itself well to microservices based architectures and to key SP-DevOps development and delivery principles.

DOCSIS 3.1 Management Data

1. Proactive Network Management (PNM)

The DOCSIS 3.1 Physical Layer (PHY) specification introduces a new network management technique for gathering highly valuable operational data from CCAP and CM devices [9]. The Proactive Network Maintenance (PNM) data can be sourced from the CCAP, the CM or as a collaboration of both CCAP and CM and for larger data sets delivers bulk data faster than traditional data collection methods (e.g. SNMP) via TFTP file transfers. **Error! Reference source not found.**² illustrates the various test points and test functions defined in the specification (sourced from [9]).

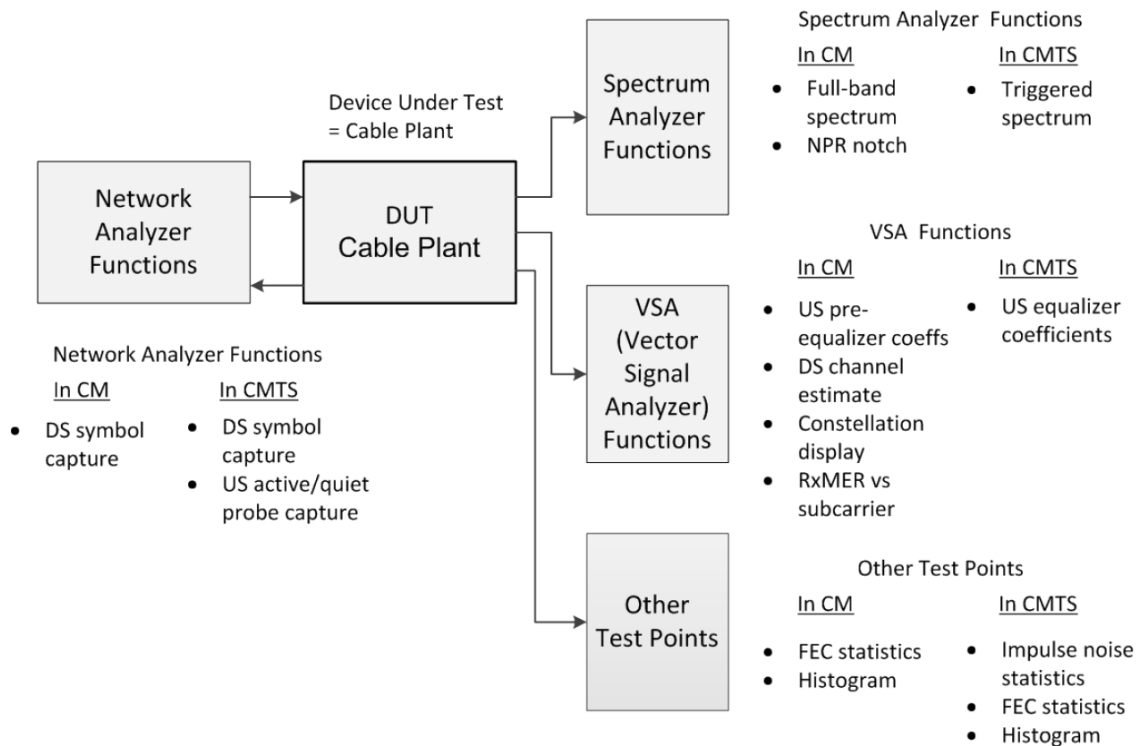


Figure 2 - CM and CCAP test points as illustrated in [1] (Figure 9.1)

These features provide capabilities similar to test and measurement equipment without the limitations associated with this type of testing (e.g. cost, limited deployment, limited access). By placing the test functions in the system to be tested and providing easily accessible methods of data acquisition, PNM data represents a valuable source of operational information for the MSO to identify physical layer anomalies and proactively address issues that affect access network performance and service quality.

1.1. DOCSIS 3.1 PNM Data Collection Workflow

Collection of DOCSIS 3.1 PNM data from a DCOSIS 3.1 CM device involves the following steps:

1. SET via SNMP any configuration parameters for the test:
 - a. Configuration parameters may include offsets, windows, timeouts, etc.
 - b. For tests that return a file, set a filename and TFTP service IP.
2. Trigger the test via SNMP set to the CM device:
 - a. All tests include a “TriggerEnable” field to initiate the test.
3. Collect the data:
 - a. For file-based data, part of the SET step is setting a TFTP destination. The data, when complete, will be forwarded to this TFTP service
 - b. For SNMP based data, the data must be retrieved from the device via SNMP.

These steps are non-trivial, order-dependent steps for obtaining PNM data – making this a prime candidate for a distributable microservice providing a uniform user interface for gathering and using PNM data.

1.2. D3.1 CM PNM File Data

Note that the specific definitions of these data sources can be found in [10] and [11]. Please refer to the most recent version of this document for a full definition of the test, its data sources, and content.

Table 1 - Defined CM PNM Tests

PNM Test	Source	Note ¹
CM Symbol Capture	PNM File	Analyze the response of the cable plant from the CM's perspective based on a sample symbol captured at the CCAP and CM. Paired with the CCAP equivalent below.
CM Channel Coefficient Estimates	PNM File	CM estimate of the downstream channel response coefficients, typically used for the CM's downstream equalizer.
CM Ds Constellation Display	PNM File	CM downstream constellation display providing received QAM constellation points.
CM Ds OFDM Rx MER	PNM File	Measurements of the receive modulation error ratio (Rx MER) for each subcarrier
CM Ds Histogram	PNM File	Measurement of nonlinear effects in the channel such as amplifier compression and laser clipping. CM captures the histogram of time domain samples at the wideband front end of the receiver (full downstream band).
CM Pre-equalizer Coefficients	PNM File	CM upstream pre-equalizer coefficients. The CM pre-equalizer coefficients and the CMTS upstream adaptive equalizer coefficient update values, when taken together describe the linear response of the upstream cable plant for a given CM.
CM FEC Summary	PNM File	A series of codeword error rate measurements on a per profile basis over a set period of time (10min or 24hr).
CM Spectrum Analysis	PNM File	CM downstream spectrum analysis function, each measurement is a data collection event that provides the energy content of the signal at each frequency within a specified range.
CM OFDM MER Margin	SNMP	An estimate of the MER margin available on the downstream data channel with respect to a candidate modulation profile. This is similar to the MER Margin reported in the OPT-RSP Message [MULPIv3.1].
CM OFDM Required QAM MER	SNMP	Calculated Required Average MER based on the bit loading for the profile and the Required MER per Modulation Order provided in the CmDsOfdmRequiredQamMer Table.

¹ Paraphrased from [7]

2. DOCSIS Common Collection Framework (DCCF)

Introduced early in 2017, the CableLabs DOCSIS 3.1 Common Collection Framework (DCCF) was introduced to abstract the complexity of low-level data collection in DOCSIS 3.1 network deployments.

CableLabs Proactive Network Maintenance (PNM) [PNM] program has been met with measurable success by Cable operators working to enhance best practices for DOCSIS 3.0 network operations. With the introduction of D3.1, it is anticipated that PNM adoption will increase while network data complexity and volumes will rise.

In addition, new platforms that enhance visibility into other critical segments within the service provider access network infrastructure, both wireless and wireline, will also be addressed by companion PNM initiatives. In doing so, the Cable operator will be enabled with a common platform and methodology upon which to build enhanced applications to support current and future operational use cases. Potential network infrastructures within the operator's management domain that would benefit from common PNM practices could include Wi-Fi, MOCA, R-PHY, and optical.

It is expected that multiple network technologies will be under PNM management within an operator's infrastructure at the same time. It is also understood that different network types will expose different forms of operational instrumentation based on information models inherent to their design and deployment disposition. In addition, over the course of broadband network evolution, a number of different network management protocols have been adopted to manage D3.1 networks.

Though data is collected from the same network, it is often gathered by multiple protocols (SNMP, TFTP, SYSLOG) embedded in disparate and closed network management systems. This adds to the burden of network data collection, making holistic visibility unattainable.

In order to motivate wider adoption of PNM practices across current and emerging network technologies, it has been proposed that a structured approach to network data collection would accelerate development and deployment by abstracting the complexities of multi-network visibility through the support of standard network information models and protocols. The Common Collection Framework (CCF) provides a structured approach to the collection of data from standards-based network deployments.




Figure 3 - CableLabs Common Collection Framework Architecture

As the name indicates, the DCCF is concerned with connecting D3.1 management applications with data instrumented by the underlying D3.1 network plane. This restricts the DCCF to the following subset of logical components described in the CCF architecture:

- D3.1 PMA and D3.1 PNM applications.
- D3.1 network data service.
- SNMP, TFTP, and SYSLOG protocol modules.
- D3.1 CCAP and CM devices.

The overall DCCF architecture follows the model described for CCF in the preceding sections.

An Example DOCSIS 3.1 Virtualized Microservice

In this section, we bring together the concepts introduced earlier in the form of a basic DOCSIS 3.1 virtualized microservice using the DCCF software system.

3. DCCF Microservices Architecture

3.1. Software Modules

The DCCF design is well-suited for deployment as a scalable microservice. DCCF consists of three (3) primary modules:

1. **DCCF REST API (RA)** - All user REST requests are directed to the DCCF RA. The RA acts as a request router, forwarding any incoming requests to the correct destination. Requests are forwarded to the Workflow Controller via an internal REST interface.
2. **DCCF Workflow Controller (WC)** - The WC manages the execution of individual tasks against network devices or external data sources.
 - a. The WC contains “Drivers” for different network operations. The first driver is an SNMP driver. Future drivers might include:
 - i. Interfaces to provisioning systems
 - ii. Data retrieval from other data sources (DB, other collection systems)
 - iii. IPDR collection (most likely through file import)
 - b. Each driver is made up of one or more driver modules that performs an action. These are interchangeable as long as the input and output formats remain the same. This means that in most cases, a single action/function can be updated without resetting the system.
3. **DCCF TFTP Service (TFTP)** - This is a TFTP service customized for specific file management functions required by DCCF. It is a service for requests to PUT files from remote devices such as CMs, and for GET requests from the DCCF for retrieving remove file.
4. **DCCF Disk Cache (DC)** - This is the local data storage associated with each WC. The RA does not have a DC. The DC is not meant to be long term storage, and any long-term storage needs to be done external to the DCCF. The DC has no software component.

Each module of the DCCF can be updated independent of the others, as long as the interface characteristics (parameters in and out) remain the same.

4. DCCF Deployment and Scale

Error! Reference source not found.4 shows a simple standalone DCCF installation, with all modules located in a single container. An obvious deployment case for testing, it also represents the most atomic deployment possible with an assumed 1:1 relationship with a CCAP. Using an external proxy router for incoming requests, an entire production environment could be made of small, more atomic installations.

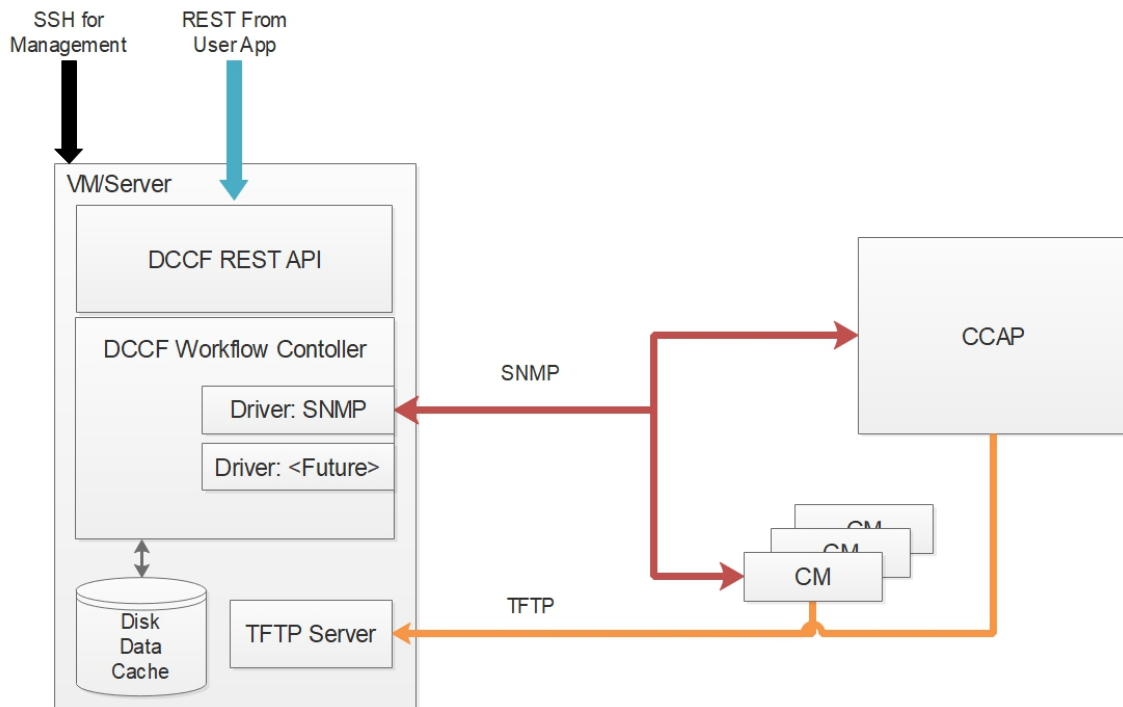


Figure 4 - Basic DCCF installation

As each module is a standalone entity, they may be distributed in different containers. **Error! Reference source not found.** shows a distribution of the RESTful, WC, and TFTP services.

This configuration introduces the concept of a remote TFTP service which requires a new Driver module, a “GetTFTP” module, to retrieve files from the remote TFTP service. When a file arrives from a remote device, the TFTP service sends a REST alert to the WC which triggers a GetTFTP action to retrieve the file and store it in the local cache.

Configuring a remote TFTP service requires some minor changes to the DCCF configuration file. Also note this example maintains a 1:1 relationship between RA and WC: One RA serves one WC.

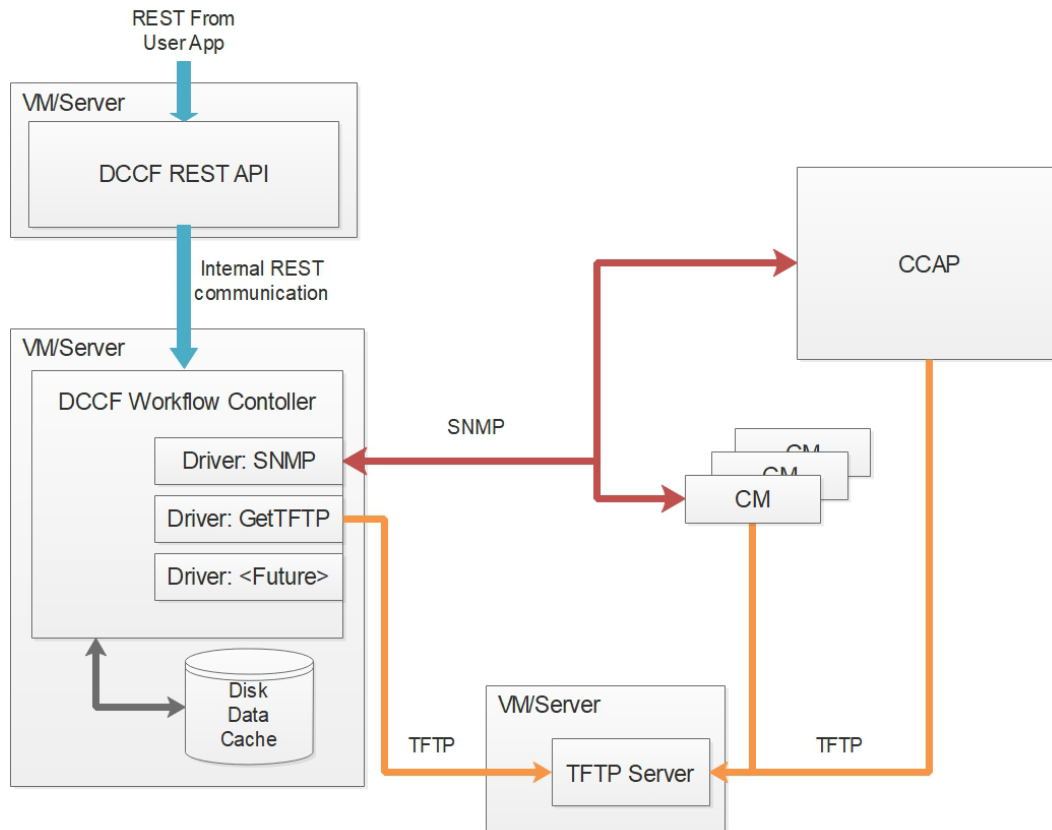


Figure 5 - Distributing the DCCF modules

Finally, a fully distributed deployment is illustrated in **Error! Reference source not found.**. This is not possible with the current version 1.0 of DCCF, due to an immature routing function in the RA, but is totally supported with the current architecture:

1. Multiple RAs, mostly likely served by a simple commercial load balancer, take incoming user REST requests.
2. The routing in the RA distributes the requests so the correct WC.
3. The WC performs the actions requested, storing and making available data in the local DC.
4. TFTP activity is managed through one or more remote TFTP's (each WC configured to use an appropriate TFTP service).

By controlling the resources associated with the virtualized WC, this deployment strategy gives excellent horizontal scaling options as well as offering multiple options for physical distribution in the operations network and container environment.

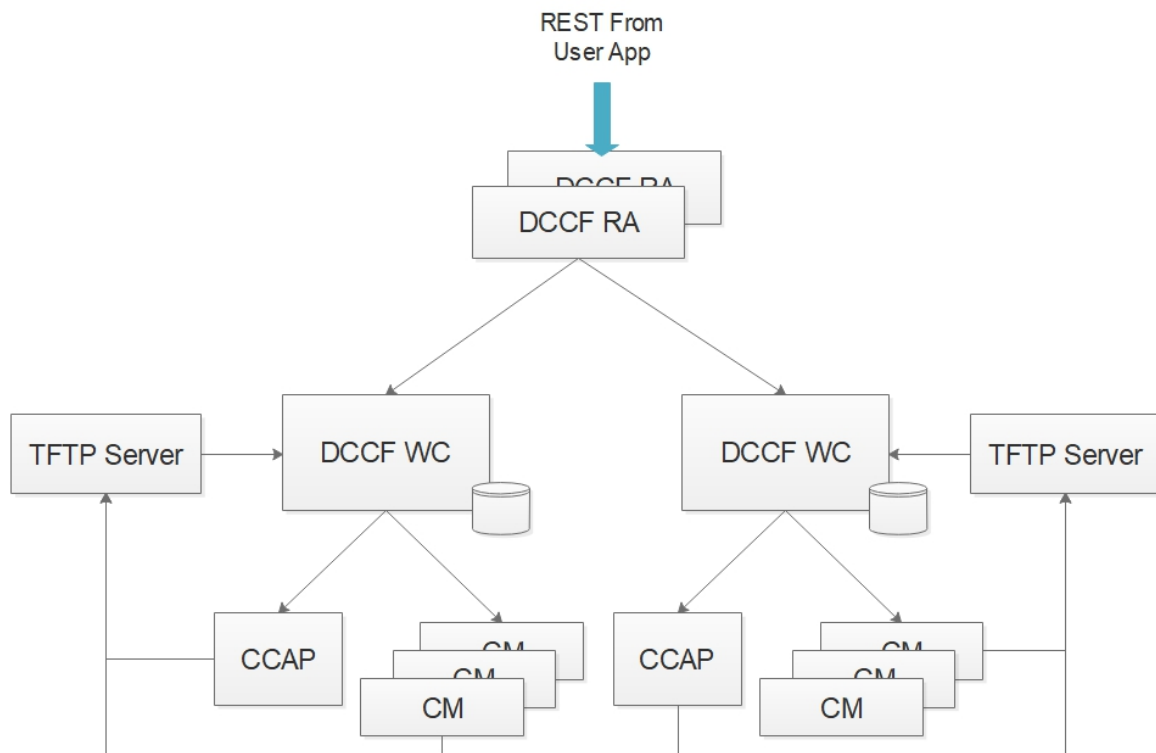


Figure 6 - Distributed DCCF with Multiple TFTP Service containers and Load Balancing

5. DCCF Features and Functions

The DCCF REST API implements a growing number of useful commands for the exploration and management of DOCSIS 3.1 networks. Table 2 shows a summary of available DCCF operations available as of release 1.0.

Table 2 - DCCF v1.0 REST API commands

Operation	Description
POST /dccf/ccaps/(CCAP)/cmPNMDsRxMer	Initiate collection of cmPNMDsRxMer data for CMs devices listed in attached JSON file.
POST /dccf/ccaps/(CCAP)/cmPNMFecSum	Initiate collection of cmPNMFecSum data for CMs listed in attached JSON file.
POST /dccf/ccaps/(CCAP)/initialize	Add new CCAP to DCCF.
POST /dccf/ccaps/(CCAP)/registered31cms	Create/update list of D3.1 CMs on CCAP.
POST /dccf/ccaps/(CCAP)/topology	Create/update CCAP topology information (Multiple GET options).
POST /dccf/ccaps/(CCAP)/cms/(CMMAC)/cmDeltaDsFecStats	Poll the latest cmDeltaDsFecStats (SNMP) data (GET retrieves data).
POST /dccf/ccaps/(CCAP)/cmPNMDsRxMer	Initiate collection of cmPNMDsRxMer data for CMs listed in attached JSON file (GET retrieves data).
POST /dccf/ccaps/(CCAP)/cms/(CMMAC)/cmPNMDsRxMer	Initiate collection of cmPNMDsRxMer data from specified CM (GET retrieves data).
POST /dccf/ccaps/(CCAP)/cmPNMFecSum	Initiate collection of cmPNMFecSum data for CMs listed in attached JSON file (GET retrieves data).
POST /dccf/ccaps/(CCAP)/cms/(CMMAC)/cmPNMFecSum	Initiate collection of cmPNMFecSum data from specified CM (GET retrieves data)
GET /dccf/jobs/(JOBID)	Return status information for specified JOBID.

DCCF Microservice Example

With the concepts of D3.1, PNM, DCCF, and microservices in hand, we can now proceed with a simple example using the DCCF software (release 1.0) running on a virtualized host within a container environment.

The DCCF is running on a host (DCCF_HOST) with SNMP and TFTP access to a DOCSIS 3.1 network. This first query (figure x) is executed in the DCCF client terminal using the **curl** (<https://curl.haxx.se>) command utility to generate the REST API CM topology discovery command over HTTP. In this example, the DCCF returns a list of all D3.1 CM devices registered on the CCAP (CCAP_IP) is returned describing each by MAC address.

Figure 7 shows a request for a single CM's Downstream Receive MER report. In this example, a client executes a curl command which sends the REST API request to request CM PNM data for the device. The DCCF returns an acknowledgement that the command has been received and is in process.

To retrieve the CM PNM MER measurement data requested, a final REST command is sent via curl from the client's terminal. The command is sent to the DCCF which returns the current data of the CM's Downstream RX MER in an efficiently compressed and archived format.

```
[dccf_host]$ curl -X GET
http://${DCCF_HOST}:8888/dccf/ccaps/${CCAP_IP}/registered3lcms
{
  "function": "wc_get_ccaps_registered3lcms",
  "json_data": [
    "AC202E772B70",
    "F8A097EF242C",
    "1CABC0B999E4",
    "F8A097EF24B3",
    "64777D90D8C0",
    "64777DE45830",
    "A84E3FCA5B50",
    "1CABC0B99AC6",
    "AC202E772D60",
    "1CABC0B99AF0",
    "1CABC0B99ADC",
    "1056118A0B9E",
    "AC202E7727E0",
    "64777DE45890",
    "64777D5EC500"
  ],
  "message": "wc_get_ccaps_registered3lcms: Completed on CCAP Status code: 200",
  "results_in": [
    "json_data"
  ],
  "status": "OK",
  "status_code": 200
}
```

Figure 7 - Retrieving CM Topology Information for a CCAP from DCCF

```
[dccf_host]$ curl -X POST
http://${DCCF_HOST}:8888/dccf/ccaps/${CCAP_IP}/cms/${CM_MAC}/cmPNMDsRxMer
{
  "function": "ra_post_ccaps_cms_pnm",
  "json_data": {
    "function": "wc_route_handler",
    "json_data": "{\"name\": \"cmPnmFile\", \"initialTime\": \"2017-07-20
20:36:20\", \"currentState\": \"ACCEPTED\", \"action\": \"cmPnmFile\", \"jobid\":
\"20170721003620_4b70b0bd_010010010001\", \"ccap\": \"010010010001\",
\"updateTime\": \"2017-07-20 20:36:20\"}",
    "message": "wc_route_handler: Completed cmPnmFile on CCAP",
    "results_in": [
      "json_data"
    ],
    "status": "OK",
    "status_code": 200
  },
  "message": "ra_post_ccaps_cms_pnm: POST
http://${DCCF_HOST}:8888/dccf/ccaps/${CCAP_IP}/cms/1CABC0B99AF0/cmPnmFile/4
workflow_controller response status code: 200",
  "results_in": [
    "json_data"
  ],
  "status": "OK",
  "status_code": 200
}
```

Figure 8 - Requesting CM DS Rx MER Measurement Data From DCCF

```
[dccf_host]$ curl -vv GET
http://${DCCF_HOST}:8888/dccf/ccaps/${CCAP_IP}/cms/${CM_MAC}/cmPNMDsRxMer >
/tmp/pnm_dsrxmer.tar.gz
* Trying ${DCCF_HOST}...
  % Total    % Received % Xferd  Average Speed   Time    Time       Time  Current
                                 Dload  Upload   Total   Spent    Left   Speed
  0     0    0     0    0     0      0      0  --:--:--  --:--:--  --:--:--    0*
Connected to ${DCCF_HOST} (${DCCF_HOST}) port 8888 (#0)
> GET /dccf/ccaps/${CCAP_IP}/cms/1CABC0B99AF0/cmPNMDsRxMer HTTP/1.1
> Host: ${DCCF_HOST}:8888
> User-Agent: curl/7.49.0
> Accept: */*
>
* HTTP 1.0, assume close after body
< HTTP/1.0 200 OK
< Content-Type: application/x-tar
< Content-Disposition: attachment; filename=wc_get_ccaps_cms_pnm-cmPNMDsRxMer-
1CABC0B99AF0_MOST_RECENT_None_None.tar.gz
< Last-Modified: Fri, 21 Jul 2017 00:37:33 GMT
< Expires: Fri, 21 Jul 2017 12:37:33 GMT
< Content-Length: 2145
< Date: Fri, 21 Jul 2017 00:37:33 GMT
< ETag: "1500597453.866377-2145-3238667453"
< Cache-Control: max-age=43200, public
< Server: Werkzeug/0.11.11 Python/3.5.2
<
{ [1024 bytes data]
100 2145 100 2145    0     0 81342      0  --:--:--  --:--:--  --:--:-- 85800
* Closing connection 0
```

Figure 9 - Retrieving CM DS MER Measurement Data From DCCF

In this way, we have demonstrated remote interaction with a virtualized DOCSIS 3.1 microservice that provides visibility into the network while abstracting the complexity of low level data collection and topology discovery. DOCSIS network functions and applications can now be developed without requiring low-level access network data collection capabilities. The remainder of this example will present a simple DOCSIS 3.1 application that displays the OFDM DS MER data returned by the DCCF microservice.

To illustrate the CM MER data content, Figure 10 shows visualizations for three CM devices created with a simple Python script that retrieves data from the DCCF REST interface and generates a graph using an open source visualization library.




Figure 10 - Example OFDM MER for three different CM devices

Conclusion

We've presented a practical example to illustrate key concepts of virtualizing DOCSIS 3.1 network functions. We presented a simple microservice that provides an implementation of the CableLabs DOCSIS 3.1 Common Collection Framework (DCCF) software system. Use of this DOCSIS 3.1 microservice was shown using a container architecture within a cable operator's cloud infrastructure.

Abbreviations

PNM	Proactive Network Management
CM	Cable Modem
CCAP	Converged Cable Access Platform
DUT	Device Under Test
CMTS	Cable Modem Termination System
SNMP	Simple Network Management Protocol
MSO	Multiple System Operator
DCCF	DOCSIS Common Collection Framework
CCF	Common Collection Framework
MSA	Microservice Architecture
PO	Profile Optimizer
POC	Proof of Concept

Bibliography & References

[1] Teaching Agile Development with DevOps in a Software Engineering and Database Technologies Practicum, 3rd International Conference on Higher Education Advances, HEAd'17 Universitat Politècnica de Valencia, Valencia, 2017, Mason, Robert T., Masters, William and Stark, Alan

[2] From Virtual Machines to Containers and Micro-Services: The Next Generation of Virtualization

[3] <http://about.att.com/innovationblog/08252015nextgenerati>, August 25, 2015, By Andre Fuetsch

[4] http://www.fp7-unify.eu/files/fp7-unify-eu-docs/Results/Deliverables/UNIFY-WP4%20M4_1%20SP-DevOps%20concept%20evolution%20and%20initial%20plans%20for%20prototyping.pdf

[5] <https://www.opencontainers.org/>

[6] <https://www.opendaylight.org/>

[7] <http://onosproject.org/>

[8] <https://tools.ietf.org/html/rfc6241>

[9] *Physical Layer Specification*, CM-SP-PHYv3.1-I11-170510, May 10, 2017, Cable Television Laboratories, Inc.

[10] *Cable Modem Operations Support System Interface Specification*, CM-SP-CM-OSSv3.1-I06-151210, May 10, 2017, Cable Television Laboratories, Inc.

[11] *CCAP™ Operations Support System Interface Specification*, CM-SP-CCAP-OSSv3.1-I09-170510, May 10, 2017, Cable Television Laboratories, Inc.

Leveraging Machine Intelligence and Operations Analytics to Assure Virtualized Networks and Services

A Technical Paper prepared for SCTE•ISBE by

Andrew Sundelin

Director, Product Management -- Cable Innovation
Guavus

1800 Gateway Drive, Suite 160

San Mateo, CA 94404

303-883-1226

Andrew.sundelin@guavus.com

Abstract

Operators will benefit from the flexibility that comes with the ability to apply IT virtualization technology to CPE and network functions. However, virtualization introduces additional complexity, creating the need for an orchestration layer with more sophisticated assurance capabilities, including data center and network analysis capabilities, encompassing physical and virtual resources in real time.

This paper focuses on a valuable use-case, which leverages operations analytics (OA) and machine intelligence (MI) to drive a variety of resource allocations within virtualized networks. This resource allocation applies both within the “cloud” and in the access network.

Within the cloud, OA and MI can be utilized to predict when additional capacity is needed for services such as additional compute to maintain quality of experience (QoE) for a cloud-based guide or additional storage for a cloud-based DVR system. These technologies can also be utilized to predict unexpectedly popular content and pre-position it optimally within content distribution networks. Within the access network, OA and MI can be utilized for congestion prediction and service optimization. For example, trends in PHY performance parameters can be leveraged to predict the need to tweak OFDM/OFDMA profiles for maximal efficiency.

This paper also explores an emerging concept: “Software-Defined Operations (SDO),” which applies Software-Defined Networking’s (SDN’s) separation of data and control planes to key pieces of modern care and operations equipment. For example, if one views technical support calls as the data then the Interactive voice response (IVR) system is the equivalent of a router (and, thus, part of the data plane) and programming that IVR system is, thus, part of the control plane. With that concept established, the paper looks at the power SDO can have when combined with Operations Analytics.

Virtualization creates benefits and challenges, both of which are driven by the new dynamicity of network services, topology, inventory and hardware resources. The combined application of OA technologies and MI can help operators to assure virtualization by providing an automated, real-time approach that learns from data and autonomously adapts to new information, intuiting connections and relationships to proactively detect anomalies and prescribe solutions.

Introduction

Networks are once again in transition. First there is ongoing development of existing cable technologies. Then there is the adoption of innovation from outside the industry. This has happened in the past, for instance, with fiber and high-speed data. It is beginning to happen with cloud computing and virtualization. MI is another promising and timely import.

Operators will benefit from a new level of flexibility that comes with the ability to apply IT virtualization technology to customer premises equipment (CPE) and network functions. However, virtualization introduces additional complexity, creating the need for an orchestration layer with more sophisticated assurance capabilities, including data center and network analysis, encompassing physical and virtual resources in real time.

Combined with OA, MI can play a key role in this emerging era. After a brief background discussion, this paper addresses the applicability of MI/OA through several use cases. Of special relevance are solutions to resource allocation problems within the cloud (nDVR, guide) and access (DOCSIS) portions of the cable network. The emerging category of “Software Defined Operations” holds additional promise.

Advanced analytics work apart from virtualization. But OA technologies and MI are especially well suited to help operators assure virtualized networks and services by providing an automated, real-time approach that learns from data and autonomously adapts to new information, intuiting connections and relationships to proactively detect anomalies and prescribe solutions.

Clouds and Virtualization

Unveiled about five years ago, notably at a SDN gathering, and directed toward telecom service providers, network functions virtualization (NFV) was aimed at a number of goals, including cost reduction, speed, agility, innovation, and improved services. [1] While having appeared first, SDN complements NFV. The separation of control and data planes that SDN delivers can enhance the infrastructure enabled by NFV.

At NFV’s debut, MSOs already were assessing the relevance of SDN. [2] They were engaged in cloud initiatives, such as network DVR and cloud-based guides. Industry leaders also agreed that web technologies could apply to networking. [3] In terms of deployment, however, neither SDN nor NFV are in wide-scale production. For many operators, virtual CCAP (vCCAP), among other technologies, fits that description. At best, virtualization falls in what one industry leader calls the upcoming “second wave.” [4]

To date, there have been limited implementations of SDN-style control planes in DOCSIS networks, [5] but industry leaders are aligning with the broader NFV initiative. Comcast, for instance, serves on the board of the Open Network Operating System (ONOS) project, according to which central offices (and, by inference, head-ends) are being ‘re-architected’ as data centers. [See the related Central Office Re-architected as a Datacenter (CORD) and Headends Re-architected as Datacenters (HERD) concepts.] [6] [7]. Where data centers go, clouds are sure to follow.

Machine Learning, Machine Intelligence & Operations Analytics

Figure 1 shows the evolution of systems from legacy database systems on the left which are generally fairly siloed in the data they contain (e.g. operated by a single department and only containing that department’s data) to big data systems (for example, data lakes) which federate disparate data and often provide wider-scale to the organization via cloud-based access. However, one common complaint about data lakes is that operators don’t know how to derive actionable insights from all the data that they have centralized – this is where big data analytics come into play. The right side of the figure illustrates real-time streaming analytics – leveraging machine intelligence and integrated with an orchestration layer – more realizing the potential of the data lake (for data at rest) while both accelerating and automating operations decision making incorporating data in motion and machine intelligent algorithms.

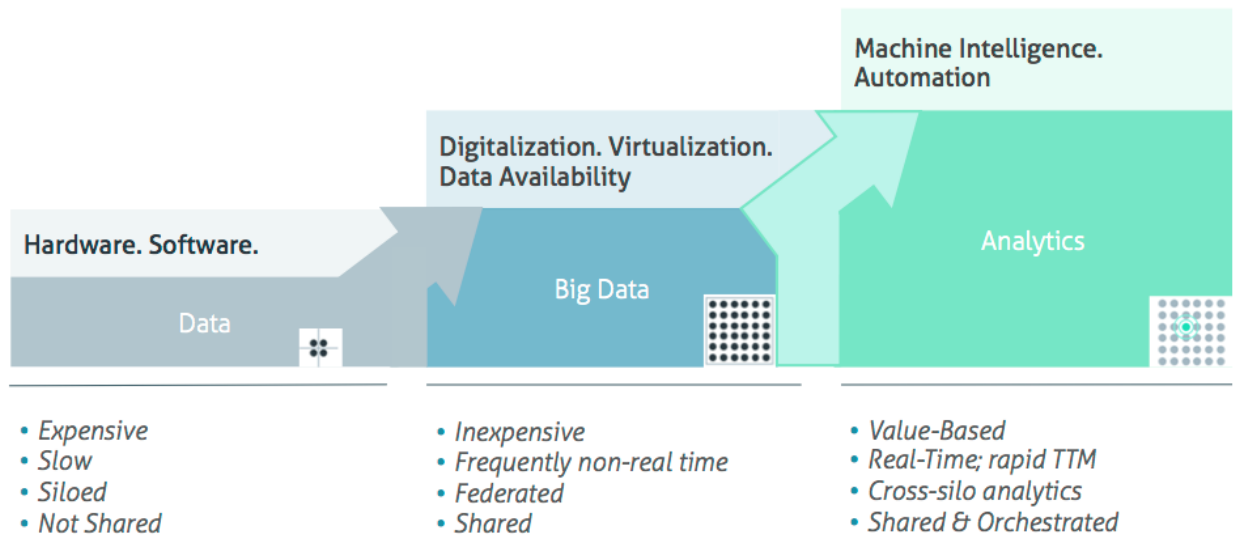


Figure 1 - The Evolution of Analytics

Machine learning (ML) is a relatively mature branch of analytics in which machines improve their ability to recognize patterns as they continue to be trained with additional examples – without having to be programmed to handle each new example or pattern. MI is a somewhat newer field which improves upon ML by adding the ability to reason. Machine intelligent systems are capable of forming hypotheses from raw, disparate data to develop new, valid information that is not a direct result of data in the original data set.

MI is a branch of the broader field of analytics, which involves studying past historical data to identify and interpret patterns. OA applies analytics techniques to the operational realm. More specifically, OA relates to historical and real-time business processes, including resource planning, service monitoring and service diagnostics.

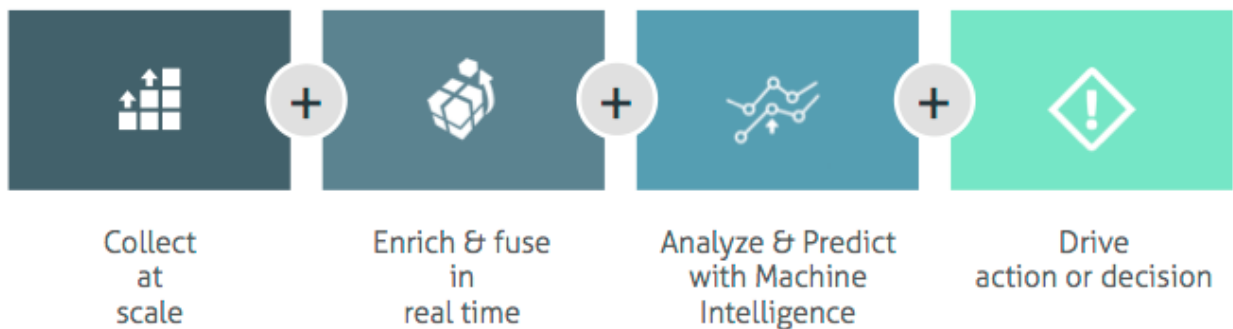


Figure 2 - Operations Analytics Functions

This paper looks at Operations Analytics systems which can address virtualization use cases in real time by leveraging a number of broad functions:

- **Collection at Scale:** collect event data in real-time and at massive scale from a variety of sources leveraging a big data engine;

- Real-time Enrichment: enrich and fuse cross-functional data in real-time with other events and reference data – combining data in motion with data at rest;
- Analyze & Predict with MI: monitor millions of event time-series, apply machine learning for baselining, anomaly detection and real-time prediction with models built via machine-intelligent algorithms for actionable intelligence; and
- Drive Action or Decision: prescribe actions and integrate with downstream systems to perform a network or operations function.

OA can have numerous applications to the service provider domain, but this paper focuses on those applications which relate to virtualization.

How does MI/OA Apply to Virtualization?

From the start, NFV was envisioned as interacting with legacy OSS and network management system (NMS), especially given the requirements of service provisioning and management, as well as with new elements, such as network controllers and cloud managers. [1] [8]. Within an SDN/NFV framework, the orchestration layer handles those complex tasks bridging physical and virtual resources.

For SDN/NFV to reach its potential, automation became a clear requirement. The logic is compelling: Programmable networks generate considerable amounts of data, which creates the potential for more intelligent, closed-loop decision-making. And as services can be spun up (or torn down) in an NFV environment, what was once challenging enough to manage becomes even more so. Operations analytics – which adapt to circumstances and draw new connections and insights through real-time data processing – therefore become integral to orchestration and next-gen OSS. [9]

The link between analytics and the programmable or virtualized networks does not occur in a vacuum. In practice, hybrid implementations draw upon legacy and new network elements for their purposes, such as anomaly detection and prescriptive decisions. [10] Of special value, as noted in some of the following use cases, is OA's ability to be more accurate and timely than legacy systems. Typically involving long-term and quickly outdated studies, the status quo approach yields static metrics that are applied universally, often misaligned with particular circumstances.

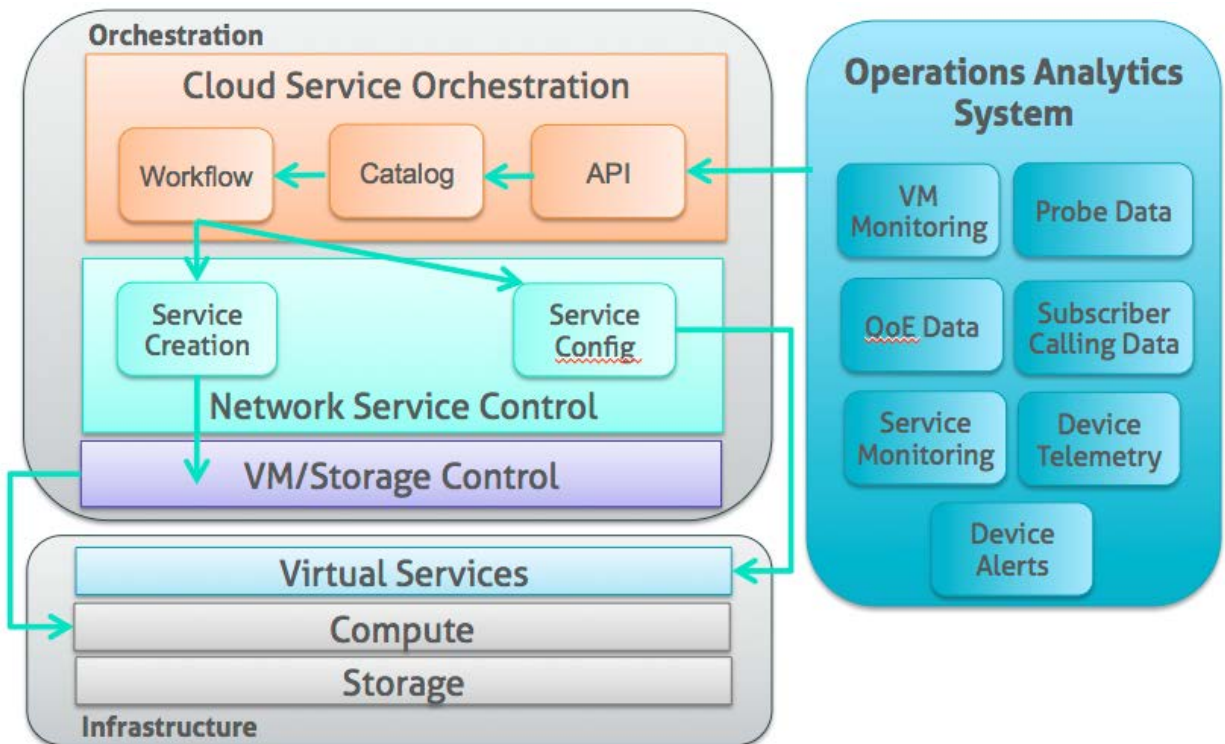


Figure 3 - Operations Analytics and Network Function Virtualization

Operations Analytics Use Cases

These first three use cases apply to cloud computing; the next two to the access network. The final case is an emerging category. They all drive a variety of resource allocation optimizations in networks that are or could be virtualized. Each case reviews the service or scenario, considers the status quo approach, and indicates what the application of OA could deliver.

1. Cloud Guide

A cloud-based, media-rich user interface no longer depends upon limited compute resources of the set-top box. Yet cloud-based guides have limits, too. A system overwhelmed by requests will generate denials for access. The worst approach is simply to wait until those denials are issued and customer complaints begin hitting the care staff. A less naive approach involves fixed thresholds, such as waiting until the resources hit a certain percentage, and then adding CPU capacity. But how can it be assured that the threshold chosen is the correct one? Or that it is correct for all parts of the network?

The ideal approach would be right-sizing a dynamic threshold for given portions of the network with the goal of optimizing resource allocation and only adding virtual machine (VM) capacity when it was truly needed, but also performing this optimization with an eye towards real subscriber QoE measures. Thus, meeting the dual goals of maximizing resource utilization and QoE at the same time. One way to achieve

these goals is to leverage automatic feature selection and develop a predictive model for a given target variable. When building a predictive model, different “features” (or data attributes) are evaluated for their level of correlation with a target variable. For example, one might use a target variable of technical support calls about video issues from cloud-based guide users and use CPU utilization as a feature; with this approach one can build a predictive model to determine at what level of CPU utilization do video-related calls from this population begin to deviate significantly from their baseline rate (i.e. become anomalous). Similarly, one might use a target variable of set-top box retry messages and CPU utilization as a feature to determine at what level of CPU utilization does the rate of retry messages begin to be anomalously high.

By performing anomaly detection on the rate of set-top retry messages, an operations analytics engine could determine the level of CPU utilization when retries on the guide begin, or perhaps when retries reach a certain percentage of all requests within a given market. Sensitivity to slow-downs or retries in different regions of the network could vary, depending upon various demographic factors – and could vary over time as STB software or end-user expectations change. Thus, by utilizing a machine-intelligence-based approach to determine when to add additional capacity to a cloud-based guide system an operator creates a closed-loop system that provides optimized, just-in-time capacity all the while maintaining QoE and, yet, adapting to both localized and time-varying user perception.

Figure 3 shows a variety of data which might be applied to determining when to add VMs to a cloud-based guide service (e.g. probe data, QoE data, subscriber technical support calls related to cloud-based guide, device alerts [such as retry messages], service monitoring data, etc.). It also illustrates how an operations analytics system can utilize an orchestration layer to spin up/down VM instances as indicated by the feature-driven model.

2. Network DVR

Leveraging the popularity of CPE-based DVRs, operators have deployed cloud or network DVR service to reduce costs and provide more value to subscribers related to greater mobility, tuning options and storage capacity. Business models vary. In a standard scenario, there may be a static threshold, at which point a subscriber receives more space or an offer to purchase more. Or the business rules may call for deleting content after a certain number of days.

A more dynamic, MI-driven approach, however, will look at the trends in the recordings made by subscribers. This kind of system assesses the need for additional capacity on the basis of not only virtual disk capacity, but also variables such as file-size distribution, delete history and recording frequency. At the macro level, operations analytics could also guide operators on the allocation of resources within clusters of VMs. A large network storage array, for instance, is likely serving not only the network DVR, but also the guide, billing and other needs. Operations analytics could determine which services get how many CPUs allocated to them, and at what times of the day, week, month or year this allocation takes place.

3. CDN Placements

MSOs first engaged with streaming servers and content libraries over QAM VOD. With the introduction of IP VOD and the migration toward all IP, many operators now run their own content delivery network (CDNs). In any case, for efficient transport and better user experience, placing content as close as possible to end-users (“content at the edge”) is a best practice. But methods differ. Whereas the status quo

approach may use Nielsen ratings or a human expert to make a determination of what is normally popular in a given region, MI/OA will detect more granular patterns, which allows for better QoE while minimizing resource consumption.

With a real-time understanding of actual user behavior in some regions, an operator with a closed-loop operations analytics platform can make timely predictions about user behavior in other regions. For example, as a baseline such a system would prepopulate content based on historic local popularity across all time zones. At the same time, real-time OA would use MI to predict what is likely to become popular – in specific locations within later/western time zones – using actual viewership and demographic data from earlier/Eastern time zones. This approach minimizes latency and maximizes QoE while also minimizing WAN capacity requirements (and avoiding potential WAN congestion) ensuring that both normally popular content and surprisingly popular content are optimally positioned within the CDN.

4. DOCSIS Channel Licenses

With the emergence of the CCAP and massive channel densities, most CCAP vendors have moved from a pure hardware-based model for selling DOCSIS channels (e.g. a 5x20 card where you get 5 DS QAMs and 20 US channels solely by buying the card) to a license-based approach where the hardware itself supports hundreds, if not thousands of channels, but these channels are only usable with a valid license. CCAP vendors have also allowed operators to pool licenses across their networks, decoupling them from a single physical device. This allows operators flexibility in deployment of DOCSIS channel licenses, but with that flexibility comes an optimization challenge – where are these licenses best deployed?

This situation presents an opportunity for MI/OA and, essentially, license-based SDN within the access network. By way of smart license allocation, an operator can optimize license allocations such that licenses from under-utilized portions of the network can be reallocated to “hot spots” in other portions of the network.

Part of the challenge here is essentially a classic capacity planning challenge – but applied in an SDN-oriented fashion. This MI-based approach allows operators to transcend simple, static “one size fits all” notions of capacity thresholding, where rules of thumb might say that any DOCSIS channels are congested at, say, 75 percent utilization. By combining QoE measures such as speed testing probes, customer calls about “slow speeds”, etc., MI and OA can be applied to capacity planning, empowering operators to right-size capacity based on the specific sensitivities of the local subscriber population. For example, some populations may be very sensitive to congestion and start to “feel” the impact of congestion at, say, 72 percent utilization. Other populations may be less sensitive to congestion (either due to different expectations or different traffic types) and not “feel” congestion until 80 percent. This variation presents an opportunity to minimize CAPEX by optimally assigning DOCSIS channel licenses.

Whereas a status quo approach to making these licensing decisions might rely upon one-time historic generalizations, operations analytics can identify real-time variations from narrower sub-populations and prescribe precise and ongoing resource reallocation solutions with much greater efficiency.

5. DOCSIS 3.1 Profiles

As the deployment of DOCSIS 3.1 gains speed, it is useful to recall how its use of Orthogonal Frequency Division Multiplexing (OFDM) enables thousands of QAM sub-carriers per channel, each with its own profile, or modulation value and amplitude. This is the property that enables OFDM to “fit a downstream

transmission path like a glove.” [11] However, there are only 16 OFDM profiles per downstream channel and each CM has its own unique RF characteristics. Thus, optimizing this limited set of profiles for a large number of CMs per channel can be a difficult task. Configuration of CCAP modulation profiles is relatively manual in early deployments, but over time this process will become more optimized. While embedded processing within the CCAP itself could perform this optimization, this computationally intensive task is better handled externally.

Using the SDN strategy of separating control from data planes, a CableLabs working group proposed and trialed a remote application to manage these profiles, as reported on last year. [12]. The overall effort to use SDN to optimize DOCSIS 3.1 turns on the applicability of k-means clustering, a type of unsupervised machine learning algorithm. [13]. This effort remains a leading example of the application of ML/OA within a software-defined framework for cable. These applications take advantage of the lower price and higher performance of standard computing platforms as well as much longer data retention capabilities of standard computing platforms. Bringing ML/OA to the table allows for not only looking at current conditions when optimizing OFDM profiles, but also predicting future conditions. Operators can do so by leveraging current and past performance and RF conditions. By looking back at a longer history, they can then incorporate seasonality and other longer-term factors in the profile optimization.

6. An Emerging Concept: Software-Defined Operations

A new case involves the extension of SDN into the OSS arena. As the “data plane”, for instance, an IVR system is being considered. In effect, it is a network element making decisions about routing calls. The external systems that program the IVR represent the “control plane.” In this framework, it can be argued that SDN concepts can be applied much more broadly within an operator’s infrastructure than for just intelligently routing data and optimizing network performance. For lack of a better term, this emerging trend could be called “Software-Defined Operations.” Many of the leading MSOs already are applying SDN concepts more broadly within their operations.

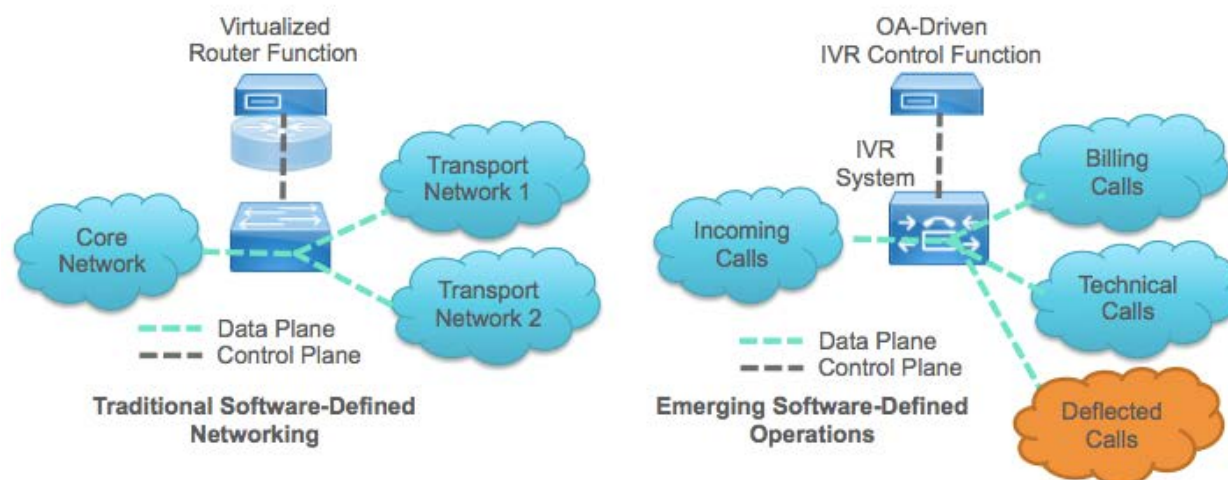


Figure 4 - Software-Defined Networking & Software-Defined Operations

ML/OA is a natural fit in this case, and there is a growing base of operator experience with this approach. One MSO’s customer experience analytics (CEA) engine was built to detect network anomalies and automatically trigger a call deflection within the IVR to inform subscribers of a known or likely outage.

The results were impressive, leading to \$6.7 million in savings from call deflections and elimination of unnecessary truck rolls. [14] While any outage inflicts damage, proactive and accurate customer engagement driven by “Software-Defined Operations” can not only reduce operations costs but also limit the impact of outages and other service issues on a net promoter score (NPS). With customer satisfaction levels continuing to fall, there is some urgency to finding applications in that domain. [15]

Other potential applications of Software-Defined Operations include:

- Auto-generating tickets in trouble-ticketing systems using MI/OA on network alarms that are predicted to lead to incidents/outages;
- Auto-cancelling booked truck rolls for subscribers which MI/OA has determined are unnecessary; in these cases, the true root issue of their problem is an outside-plant (or other multi-user issue) and not a single subscriber issue, as originally diagnosed by the care agent.

Conclusion

The application of MI and OA is not dependent upon virtualization. As discussed in a paper presented at last year’s SCTE Cable-Tec Expo 2016, MSOs are deploying advanced analytics today, to good effect. [16] With other results showing that analytics can reduce truck rolls by 30 percent, the potential for savings in customer care and network maintenance is tremendous.

MSOs can realize those benefits now, apart from SDN or virtualized network functions. Indeed, now is a good time to explore such applications of MI and OA. But the time is coming – possibly in the next technology “wave” – when deploying these smart technologies to assure more fully-fledged virtualized platforms will become critical. Already cases can be sketched out where the application of MI/OA is a good fit for both cloud environments (guide, nDVR) and access networks (DOCSIS). Another promising area for active exploration and experimentation is the emerging category of “Software-Defined Operations.”

Abbreviations

CCAP	Converged cable access platform
CEA	Customer experience analytics
CORD	Central office re-architected as a datacenter
CDN	Content delivery network
CPE	Customer premises equipment
DOCSIS	Data over cable service interface specification
HERD	Headends re-architected as a datacenter
IVR	Interactive voice response
MI	Machine intelligence
ML	Machine learning
MTTU	Meant time to understand
nDVR	Network personal video recorder
NFV	Network functions virtualization
NMS	Network management system
NPS	Net promoter score

OA	Operations analytics
OFDM	Orthogonal frequency division multiplexing
ONOS	Open network operating system
OSS	Operations support system
PCMM	PacketCable multimedia
QoE	Quality of experience
SDN	Software-defined networking
SDO	Software-defined operations
vCCAP	Virtual CCAP
VM	Virtual machines

Bibliography & References

- Network Functions Virtualization: An Introduction, Benefits, Enablers, Challenges & Call for Action. SDN and OpenFlow World Congress, 2012.
- G. White. Can DOCSIS Networks Leverage SDN? NCTA, 2013.
- B. Field. Applying Web Principles to the Network. NTCA, 2014.
- R. Howald. Aboard the Technology Wave: Surf Conditions Report. Cable-Tec Expo, 2016.
- D. Early. A Practical Guide to Implementing Software Defined DOCSIS. Cable-Tec Expo, 2016.
- N. Nandiraju. Distributed Access Architecture - Goals and Methods of Virtualizing Cable Access. Cable-Tec Expo, 2016.
- S. Chatterjee. Headend Re-architected as a Data Center (HERD). Cable-Tec Expo, 2016.
- Telcos eye servers & software to meet networking needs. (Interview with Don Clark, technical manager of NFV industry specification group.) Gazettabyte, April 1, 2013.
- CEO Chat with Anukool Lakhina, Guavas, LightReading, July 20, 2015.
- B. Lynch. Anukool Lakhina. Implement Closed-Loop Network Decisioning Now with Big Data Analytics and Fuel Future-State SDN Use Cases Through a Common Platform Deployment. NCTA, 2014.
- J. Chapman. The Power of DOCSIS 3.1 Downstream Profiles. NCTA, 2013.
- S. Rahman. J. Solomon. J. Schnitzer. D. Early. DOCSIS 3.1 Overdrive: Dynamic Optimization Using a Programmable Physical Layer. NCTA, 2016.
- J. Schnitzer. DOCSIS 3.1 Profile Optimization Using SDN. Cable-Tec Expo 2016.
- A. Sundelin. J.P. Goyet. Using Analytics to Extract Operational Intelligence and Redefine Customer Experience Management. Cable-Tec Expo, 2016.
- Wireless Competition Boosts Customer Satisfaction, While Subscription TV and ISPs Face Problems. ACSI Telecommunications Report 2017.
- C. Menier. Contextualizing Data to Gain a Real-Time Integrated View of the Network and Improve Customer Experience. Cable-Tec Expo, 2016.

Access Network Operations Savings Through Extending Automation and Orchestration Beyond Remote PHY

A Technical Paper prepared for SCTE•ISBE by

John Holobinko
Director Access Networks Strategy
Cisco Systems, Inc.
Lawrenceville, GA
Phone: 770-236-1123
jholobin@cisco.com

Ron Zimmerman
Director, Technology
Cox Communications, Inc.
ron.zimmerman@cox.com

Todd Greene
Product Manager, Marketing
Cisco Systems, Inc.
greenet@cisco.com

Introduction

The basic design of optical nodes has gone unchanged for over a decade. This paper explores the notion that rethinking the design and function of the optical node, specifically by incorporating intelligence and leveraging automation, can potentially have a material effect on cable operations expense (OpEx) and network availability.

Content

1. Problem Statement

Cable access networks are built with a combination of optical nodes and amplifiers, coaxial cables, taps, connectors and drop cables. By their very nature, cable access networks are more OpEx intensive than passive optical fiber networks (PONs). For example, the passive elements of a coaxial distribution network include many separate pieces of cable, taps and connectors that join these together. Connectors may loosen; water seals may fail; cables may be dented resulting in impedance changes. All of these deleteriously impact network performance and require corrective maintenance. While not immune, fiber plant is less susceptible to such issues.

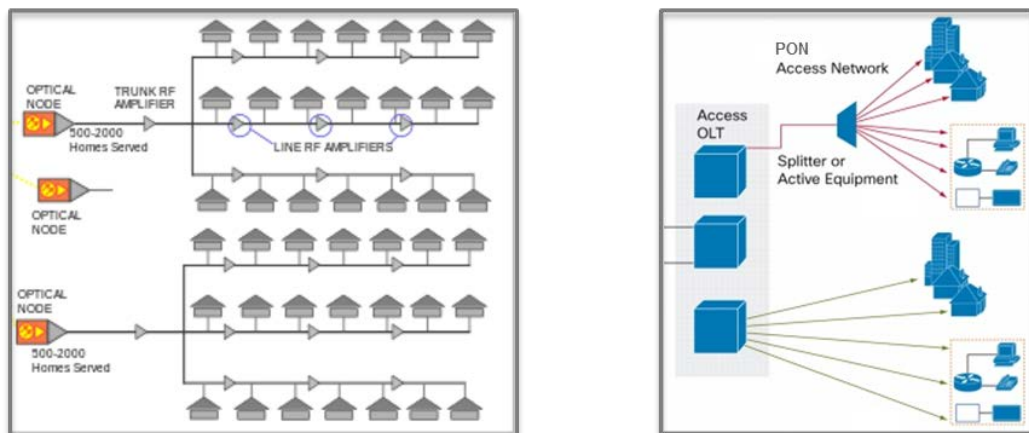


Figure 1 – Cable Access Network versus PON Network

Even small cascade/no cascade cable access networks remain more costly to maintain than PON networks.

In a cable access network there are many “cable specific” potential failure points including:

- Node (N+0) or Node + Amplifiers (N+M)
- Taps – water ingress, etc.
- Hardline cable – dent, rodent damage, shield issues
- Drop cable – cut, damaged, improper connections
- RF connectors – loose, water, damaged

In a PON, the number of potential failure/performance impacting mechanisms is significantly less than current HFC networks. The major failure mechanisms in PON are:

- Optical network terminals and optical line terminals (ONTs + OLTs)
- Fiber cable cut
- Optical connectors

However, PONs are much more expensive for cable operators to implement for serving existing subscriber areas than cable network upgrades to N+1 or N+0 architectures. Put into perspective, in PON there is a dramatic trade off of much higher CapEx in exchange for lower OpEx.

As discussed below, given new advances in bandwidth achievable by implementing DOCSIS 3.1 and 1.2GHz bandwidth, cable systems can potentially deliver 10Gbps downstream. With emerging full duplex DOCSIS (FDX) technology, the upstream is potentially capable of 5 Gbps. Therefore, the cable system challenge vis a vis PON is not as much bandwidth as it is to shrink the operational cost of the cable access network to approach the operational cost of the PON network.




Figure 2 – The Chasm: Cable Access OpEx vs. PON OpEx

2. Postulating a Change in Node Thinking

We began our study with a question: Can automation and additional intelligence be incorporated within the cable access network in such a way as to reduce the gap between the cable access network and PON relative to outside plant installation and maintenance?

One of the current means of reducing maintenance is a radical reduction in the number of active devices between the cable headend and the subscriber. Reducing amplifier cascades to N+2 will typically provide sufficient bandwidth per subscriber, while N+0 networks have the additional benefit of having only a single active device between the headend and the subscriber, albeit for a much higher capital investment. As evidenced by announcements for N+0 network plans over the last year, more cable operators appear willing to trade off this higher capital expense (CapEx) to obtain gains achieved in OpEx reduction within a low/no cascade cable access network.

In addition to N+0 benefits, significant technology advances are coming to the cable access network. In the near future, a cable network implementing Remote Phy and Full Duplex DOCSIS (FDX) in N+0 architecture across 1.2GHz of spectrum will provide more bandwidth to a given group of subscribers (a

service group) than a classic gigabit PON network (GPON) serving the same subscribers. However, even as these advances will deliver a dramatic technical advance in bi-directional bandwidth capacity, the operations costs of an HFC network still remain significantly higher than most of today's installed legacy PON networks, which are 1G EPON or 2.5G x 1G GPON.




Figure 3 - Various PON Effective Throughputs versus DOCSIS 3.1 with FDX in 1.2GHz System

The performance and reliability of a cable access network remains highly dependent on the skills of the RF technicians who are tasked with installing and maintaining the nodes (and amplifiers where applicable), passives, taps and coaxial cables. Even with proactive network management (PNM) tools, the overall installation, plant operation and problem determination of the cable network remains dependent on the skill of the cable technician. Remote PHY device (RPD) technology can mitigate some of the issues previously associated with the analog fiber portion of the cable plant, but it also brings operational complications. Ironically, an N+0 network employing 10GE optics connected to RPDs is in some ways more of a challenge to maintain than a traditional hybrid fiber coax (HFC) network employing analog optics, because the use of some network diagnostic tools are either not supported or simply become too expensive. (The reason for the latter is that in a rebuild to N+0, the number of nodes in the network increases by an order of magnitude.)

In a low/no cascade cable access network, the node becomes the central control point for the service group. The design and function of HFC optical nodes has remained relatively the same since their introduction into the network outside plant over twenty years ago. A cursory look at node performance and failures would conclude that nodes are very reliable components of the network, yielding little potential for improving cable access network costs and availability. However, what that does not consider is the node as a remote information and remote control point within the cable RF distribution plant. Therefore we posed the following question:

3. Can a node be redefined to positively impact cable network OpEx

Can the node be designed in such a way as to have a material effect on plant power consumption? What if the concept of the node could be redefined in such a way as to enable more control and problem determination of key plant issues and failure mechanisms? Based on these questions, we set out to

determine whether a rethinking of the node and its functions could have a material effect on operations costs and network availability.

For many years, the cost of automating nodes and providing a smattering of additional information remotely from nodes to hubs was so expensive as to make wide spread deployment impractical in most systems. In North America and Europe, only a few operators have even implemented the ability to report back basic information such as optics status, and the ability to attenuate a return path leg (also known as a “wink” switch) as a means of isolating some cable system problems. But, what if by rethinking the node, automation technology and intelligence could be directly incorporated into the node design, offsetting certain costs and making it possible to build an advanced node at approximate cost parity with traditional “dumb” nodes?




Figure 4 – Traditional node with DOCSIS transponder vs. node built with integrated intelligence

4. Potential Cost Savings of An Intelligent “No-touch” Node

Importantly, to determine if enhanced node functionality can have an impact on network reliability and services availability requires examining the major or areas of OpEx cost for the network.

1. Power Costs
2. Operations Costs and System Availability
3. Lost Subscriber Costs (i.e. customer retention)

4.1. Power savings

For an N+0 architecture to be the most capital cost efficient requires redesign of the system to maximize the number of subscribers able to be reached by a node, and for the node to achieve the highest output power possible on every leg. While the number of households passed (HHP) per node will vary based on home density, a North American average is 40-70 HHP per node. However, homes are not laid out linearly. Densities and topologies vary. In 1.2 GHz N+0 cable access network designs Cisco has performed, we find that not every node will utilize all four outputs. The average number of outputs is approximately 3.4 versus 4.0. In other words, in a 50,000 HHP design with 50 HHP per node, there will be 1000 nodes, and a total of 4000 outputs, of which 600 will not be used. This means that with

traditional nodes, the operator will be paying for substantial power consumption across the N+0 footprint that is unnecessary over the life of the system. When systems are redesigned for N+0, they are mostly designed for 1.2 GHz operation with the goal of 10 year lifetimes before they reach traffic saturation. Operating at 1.2 GHz is substantially more expensive than operating at 860 MHz or 1.0 GHz, yet the operator may not utilize the full bandwidth for a number of years. The difference in power consumption between 860 MHz, 1.0 GHz and 1.2 GHz is also substantial.

With today's traditional high output nodes there is no means to turn off unused RF outputs, nor is there the ability to run an RF output hybrid at less than full power (i.e. full bias current and bias voltage). Therefore these savings are currently unobtainable with traditional nodes. However, what if the node were intelligent and control over all of its internal operating parameters could be done remotely without touching the node? Specifically, if the node were designed to control both "on/off" plus variable bias voltage/ current of each output of each output hybrid amplifier, what are the potential savings?

We calculated the power savings of the unused outputs across the footprint using the ratio of 3.4 active ports per 4 port optical node. For power savings from lower initial spectrum use, we assumed five years (the estimated system half-life as defined by the cable operator) as the point when a system would require full use of 1.2 GHz bandwidth. We then calculated the power in those first years that is wasted based on two scenarios: with no traffic occurring above 860MHz or 1.0 GHz. For systems that pay for metered power versus bulk rate, the potential cost savings of lower power consumption is substantial. Using a cost of \$0.10 per KWH and a total footprint of 1 Million HHP results in the following total potential combined power savings from turning off unused ports and operating the remaining ports at reduced power until they are fully enabled at the five year point:

Table 1 - Computed Power Savings From Automated Control of Output Power Hybrids

10 Year Total Savings @ 860 MHz	\$2,993,000
10 Year Total Savings @ 1.0 GHz	\$2,620,700
10 Year Average Savings/Node 860 MHz	\$179.58
10 Year Average Savings/Node 1.0 GHz	\$157.24

4.2. Operations Costs Savings and System Availability

We address OpEx savings and systems availability simultaneously as they are interrelated. Operations costs savings are derived by determining a reduction in the time and effort required to detect and remediate system issues. Faster detection, remediation, or rectification of issues reduces the time that the system is not operating or services are operating in a degraded state, resulting in higher network availability. Higher network availability is a factor that corresponds to reduced customer churn.

System/services availability is dependent on at least three major factors:

1. The number of hard cable access network failures
2. The number of degraded performance issues in the access network
3. The time to detect and correct hard failures and degraded performance issues times the number of subscribers per event

In today's cable networks, more customer complaints originate from degraded performance than hard network failures. As stated previously, the number of hard failures in an N+0 network is substantially less than in a cascaded amplifier network. Hard failures are limited to node failures and coax/fiber cuts. By far, degraded performance incidents dominate service calls related to the access network.

Cable plant issues that degrade return path performance represent some of the most difficult problems to resolve. They can result in slow internet performance, or intermittent loss of service as cable modems lose registration, re-register, and lose registration repeatedly. Even with next generation Remote Phys installed in nodes, the operator is challenged to determine which leg of the node has the problem. PNM can be used to somewhat isolate the problem to a localized area of the cable plant, but PNM can only narrow down the likely causes of the problem, versus pinpoint the exact cause in every case.

Return Path Affecting Problem Detection, Mitigation and Resolution

Today, usually as a result of a customer triggered trouble ticket, technicians are sent out to the service area. The first course of action is to visit the node in order to get a reference signal. Then the technicians go from the node to the location where the suspected problem exists. A few examples include loose connector, water in tap, rodent damage to cable, dented cable, etc. To confirm that they have located the problem either requires an additional person at the headend or a special diagnostics program. The technician corrects what s/he believes the problem to be, then waits for feedback to see if the problem is corrected. In some cases, they also use a diagnostic program that interfaces to the cable modem termination system (CMTS). This multi-step corrective action takes time and resources.

How could functionality in the node be changed to significantly reduce the effort to address return path problems?

1. *Sense the problem before the customer trouble ticket even occurs.* The node should have the ability to take measurements of return path waveforms for each individual return path at the node RF port, independent of the RPD, and send these to a software package such as an enhance PNM system with big data analytics in order to perform trend line analysis and detect issues before the customer calls. Such software goes beyond a traditional PNM system.
2. *Enable the problem to be mitigated before the technician is dispatched.* Waveforms for each node leg should be remotely visible and recordable and the node should provide a capability to remotely attenuate each individual return path leg, so that in the case of impulse noise, customers on that leg may still operate with reduced performance instead of total loss of service. In the case where the problem is too severe, the node should be able to turn off only the offending leg, so that the other 75% of the homes attached to the node remain at full service.
3. *Eliminate the technician's need to visit the node before going to the suspected problem location.* The node should provide a remote spectrum analysis function and eliminate the need for the technician to calibrate test equipment. Note that spectrum analysis differs from spectrum capture.
4. *Enable the technician to validate the problem and the solution from the point in the plant where the problem originates.* This requires a real time spectrum analysis function at the node on the specific leg affected. It is not sufficient to provide a static "snapshot" of the spectrum. The technician will want to view the waveform live, therefore the node function should be able to

provide multiple images per second with sufficient resolution to confirm both the problem and resolution.

Notably, the measurement capability of the node must enable the technician to be able to obtain all of the information necessary obtained previously at the node test points, without the need to go to the node. Therefore, the intelligent no-touch node needs to provide the primary functions of a technician's hand held spectrum analyzer and to do so remotely. Ideally, a secure application will enable the technician to do remote inquiries of the node. Typical spectrum analyzer functions that should to be provided include the following measurements:

- Full band spectrum capture and display
- Partial band spectrum continuous sweep (at least once every 100 msec)
- Channel power / Total composite power (TCP) Measurement
- Max. hold, Min. hold, average, clear write
- CCDF (complementary cumulative distribution function used for peak to average power)
- MER / BER (modulation error ratio/ bit error rate)

While the number of truck rolls per year per 1000 subscribers is based on a number of factors, and the cost per truck roll varies by cable operator, what we can say is that the time in the field per call will be reduced. We estimate that the time to diagnose and remediate can be improved by at least by at least 25% per truck roll for these types of problems over traditional systems. Given that a truck roll costs between \$100 and \$150 for most operators, the potential savings over the life of the node are potentially substantial.

In addition to truck roll costs, there is the cost of service calls to customers when any service degradation or outage happens. While the cost per call is significantly less, a number of potential calls may result from a single event.

4.2.1. Subscriber Losses Due to System Issues (Customer Retention)

Network outages and degraded performance are key contributors to subscriber losses. Correspondingly, we know that the ability to detect problems sooner and to mitigate any disruption of services to the smallest number of customers and the shortest overall time has a correlation to customer retention. While we do not yet have sufficient operating data to measure the impact of mitigating these issues faster, it is clear that these abilities will positively affect retention. Given that the enterprise value of a cable customer today is between US\$4,000 and \$5,000 it is quite possible that benefits of customer retention may actually be higher than OpEx savings provided by intelligent, touchless nodes.

What about today's nodes outfitted with a remote phy module? Don't they provide the same functionality and the same benefits? They do not. The table at the end of this section summarizes the functions we have identified, versus what is provided by traditional nodes, or a traditional node plus RPhy.

Performance and Reliability Issues Related to RF Technician Actions

Another cause of service problems is ironically the result of RF technician behaviors. A technician may be called to a subscriber residence to address a service issue. When the problem is not easily resolved, in hope of a fix some technicians will go back to the node and raise the downstream output level by removing and replacing the attenuator pads in the node housing. Not only does this change system

behavior, changing the pad value creates a system outage for all customers on the RF leg while the pad change is being made. The new pad value changes the RF output level, overriding the original system design and impacting the RF levels on all of the homes attached to that leg of the node. Over time, these changes usually result in additional customer service calls. In the operator employs PNM to detect level imbalances, a trouble ticket may be issued to send a technician to the node and revert the pad levels to those specified in the network design, and in doing so create another service outage. Because of these outages, the reversions need to be done in the overnight maintenance window, resulting in additional costs. The cost of these truck rolls and any associated additional customer calls are unnecessary, additional OpEx expense. The total unnecessary OpEx cost involves the original time of the technician to modify the RF output level and the time and additional truck roll required to revert the system back to its design specified level. In some systems, this is a major cause of needing to perform periodic system sweeps.

An additional expense of traditional nodes is the cost of maintaining passives inventories for technicians. While the passives themselves are not very expensive, the time and effort to maintain a complete set of pad (and equalizer) values, as these are consumed in the field. This is not a trivial matter, as technicians only tend to ask for new values when they have entirely consumed the value of a pad and need that value for the job. It can result in wasted truck rolls when a technician discovers that s/he does not have the right pad or equalizer value, plus the time spent ordering missing tap values on a rush order.

How could node functionality be enhanced to eliminate this type of technician behaviors, eliminate the pad and equalizer parts inventory issue, and simultaneously reduce OpEx?

1. Eliminate all pads and equalizer plug-in accessories from the node design – the node electronics must be designed to self-align and not require manual adjustment by the technician. Any adjustments must only be able to be done electronically and remotely without the need for any plug-in pads or equalizers.
2. The operator must be given the choice to prevent any changes to node settings by technicians, or at their option, to be able to authorize only certain technicians to modify values on certain nodes while automatically logging and time stamping any change made to node settings. By providing absolute control over node operation, unauthorized changes to levels are eliminated.

By automating set up and control of the node, and eliminating the need for any plug-in accessories, the potential OpEx savings are elimination of accessories management, plus avoidance of truck rolls necessary to rebalance systems to their original design settings. But most importantly, for each unauthorized level change, two outages are eliminated: The first one by the rogue technician and the second one for restoring the RF levels to the correct design values.

Other forms of diagnostics

Some problems are highly difficult to detect because they are intermittent – either occurring in unpredictable time intervals or very infrequently. If a particular node leg is suspect, polling of the suspected leg at very short intervals for long periods of time is not practical today, due to the heavy load it puts on the DOCSIS platform and the impact on data throughput. If the node were able to do full band or partial band capture and its polling rate is adjustable, a large number of samples could be stored in local memory and downloaded to the hub at the operator's convenience. This polling could be done completely independent of the CMTS and without impacting customer DOCSIS performance. But to do

this requires intelligence in the node in the form of a processor, local memory, plus circuitry to measure the spectrum at critical points within the node.

Simplifying Initial Installation and Eliminating Set-Up Errors

As addressed previously, the ability to set levels initially lower in order to save power and then later remotely change levels to support full spectrum operation results in considerable power savings. Importantly, for a system that is going to be rebuilt to N+0, there will be between ten and fifteen times the number of current nodes in the network, and all must be installed at the same time. Typically, a contractor hangs the node, while a second visit is made by a technician to install the node and confirm its proper operation.

If the node functions can be automated using intelligence, then the node installation can be made automatic, eliminating the need for the skilled technician at installation. If the contractor can install the fiber connectors and RF power, the Remote PHY can advertise itself to the network and self-install, followed similarly by the node if it has the requisite intelligence. Geolocation data and serial numbers can be used to identify the specific node, and a look up table used to download the specific RF design values specified for the node location. No previous set up of the node is required and there is no need to deal with any pads or equalizers at installation. This will result in considerable savings at installation as well as insuring that the proper node configuration is made 100% of the time.

Ideally, all of these communications functions should be made controllable using a standards based interface and standards based communications. In nodes that employ DOCSIS Remote PHY and the Open PHY standard this means utilizing Ethernet at the transport layer and using a Netconf/Yang model to provide an open interface to third party software.

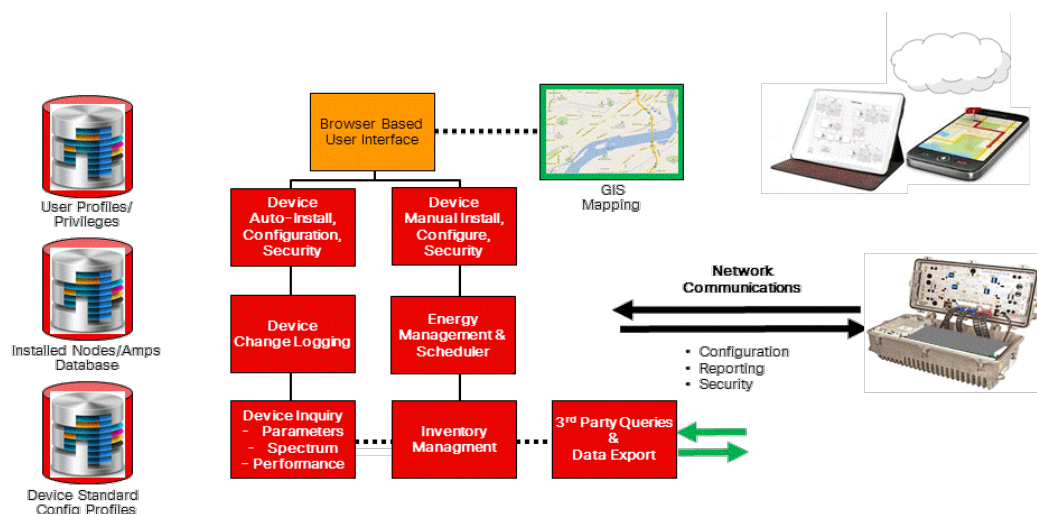


Figure 5 - Touchless Control of Intelligent Node

Follow-on Operation with Advanced Modulation Schemes

No operators are fully utilizing the entire 1.2GHz spectrum in their downstream in N+0 systems today. While it is feasible to calculate required RF levels for full spectrum utilization, in practice it may be possible to lower levels, or do other adjustments in particular areas of the spectrum. As FDX is implemented, there will be an additional learning curve on signal levels and settings. Finally, there may be improvements in customer premises equipment (CPE) that change the definition of what ideal levels are in both the forward and return path of the system.

With traditional nodes, the ability to redefine levels, slopes, and other node operating parameters is severely handicapped by the need to visit each individual node to effectuate any changes, and every change results in a network outage for those customers attached to the node under adjustment. In an N+0 network, the sheer number of nodes makes any such adjustment across the footprint non-feasible.

With an intelligent, no-touch node, the ability to remotely set and adjust node levels and other parameters individually, regionally, or globally provides a sort of futureproofing for future learnings. None of this is possible with traditional nodes.

Table 2 - Comparison of Advanced No-Touch Node vs. Traditional Node Table

	Advanced Node with Intelligence	Traditional Node	Traditional Node with 1 x 1 RPD	Traditional Node with 1 x 2 RPD
Waveforms remotely visible on a per leg basis	YES	NO*	NO, per combined returns only	NO, 2 legs combined at a time
Return path level and turn off control per leg	YES	NO	NO	NO
Remote adjustment of signal levels, slope, while eliminating pads or EQs	YES	NO	NO	NO
Capture of waveform and other data without DOCSIS performance impact	YES	NO	Limited information	Limited information
Remote power savings and power on/off for each RF Power Hybrid	YES	NO	NO	NO
Ability to control access to RF section, log changes, etc.	YES	NO	NO	NO
* Only possible when equipped with internal cable modem and controls at significant additional cost				

5. Next Generation Intelligent No-Touch Nodes in non N+0 networks

While the benefits of complete automation are reduced in a network with amplifier cascades, there remain benefits that are not available with traditional nodes. The ability to see return path levels on each leg enables each return path to be examined independent of the other return paths. This is beneficial because in today's nodes, two to four return paths are typically RF combined before sending to the RPD or burst receivers in the CMTS. Additionally, for integrated CMTS's, very often the return paths from two to four nodes will be combined into a single service group, meaning that the burst receiver sees the simultaneous return from up to 16 return path legs. Being able to see each individual return path signal at the entry point to the node and after amplification gives up to 16 times the resolution as a traditional return path service group solution.

The ability to attenuate one return path to minimize the number of subscribers impacted by a high noise event remains valuable. The ability to bring back additional telemetry for PNM allows better problem determination. For example, in an N+1 system where the amplifier is a single port line extender, spectrum display and measurement of return path inputs and forward path outputs at the node can provide additional information for diagnosing problems at the line extender or even the cable beyond.

Conclusion

We believe that the industry is on the verge of the first major change to node design in years. Next generation intelligent, no-touch nodes will integrate digital control, intelligent signal adjustment and more granular plant measurement capabilities such as rapid sweep spectrum analysis per leg (not to be confused with static spectrum display) within their basic functionality. This will require an ability to develop these nodes via new digital design techniques such that the cost of these nodes to the operator will be on parity with traditional non-intelligent nodes. Complementing these intelligent nodes and the RPDs installed inside them will be automation and orchestration software that will automate installation, control technician access, and thereby remove variables out of the technician equation. These new capabilities promise a significant impact on reducing field problems, outage times, thereby improving OpEx and customer satisfaction and reduce the PON operations advantage.

Abbreviations

BER	bit error rate
CapEx	capital expense
CMTS	cable modem termination system
CPE	customer premises equipment
DOCSIS	data over cable service interface specification
EPON	Ethernet passive optical network
FDX	full duplex DOCSIS
FEC	forward error correction
Gbps	gigabits per second
GHz	Gigahertz
GPON	gigabit passive optical network
HFC	hybrid fiber-coax
HHP	households passed
Hz	Hertz
ISBE	International Society of Broadband Experts
MER	modulation error ratio
MHz	Megahertz
OpEx	operations expense
OLT	optical line terminal (PON network element)
ONT	optical network terminal (PON network element)
PNM	proactive network management
PON	passive optical network
RF	radio frequency
RPD	remote PHY device
SCTE	Society of Cable Telecommunications Engineers
TCP	total composite power

Using Digital Identity to Drive Personalization, User Experience and Monetization

An Operational Practice Prepared for SCTE•ISBE by

Doug Fantuzzi

Vice President – Amdocs Media & Entertainment Solutions
185 Hudson St.
Jersey City, NJ 07311
240-751-5089
Douglas.Fantuzzi@amdocs.com

Hadar Sharon Amdocs Entertainment Product Management, Hadar.Sharon@amdocs.com

Ira Kogan Amdocs Entertainment Global Business Group, Ira.Kogan@amdocs.com

Introduction

Gartner stated: “By 2020, customers will manage 85% of their relationships with enterprises without interaction with a human.” User identity will be key to such digital interactions and the Digital Service Provider (DSP) that manages this identity will also need to stitch together these different interactions to create the personalized user journey. This journey will bring together digital technologies, the digital customer and the digital economy. Figure 1 below depicts the different components and characteristics of the future DSP which will play a key role in realizing Gartner’s prediction.

Digital technologies

- Cloud
- Network – 5G, NFV
- Open APIs
- Service creation
- AI
- Analytics, Big data



Digital customer

- In control
- Always connected
- Channel of choice
- Social
- Self-service
- Apps
- Consumer or enterprise

Digital economy

- OTT
- Content
- IoT
- Financial services
- Advertisement
- Healthcare

Figure 1 - Digital Service Provider

Today's digital providers like Google, Microsoft, Alibaba, and Amazon understand how to utilize user identity via a username and password that enables them to access a user's personal information. The user ID enables them to personalize services, create innovative monetization models, and deliver great user experiences where consumers can purchase and consume digital goods, wherever and whenever they want.

Communications Service Providers (CSPs), which for the purposes of this paper also includes Multi-System Operators (MSOs) and cable operators, typically manage their relationships and interactions through a physical street address without much understanding of individualized usage. As an example for cable operators, a single address would have all or some of the family members watching the same TV in the living room or sharing the same broadband connectivity, with the operator not knowing which family member actually interacted with the service. But with the evolution of digital TV and personalized connected devices, each member of the family can watch TV or connect to digital applications via the internet using their own personal or shared device, such as a set-top box, a game console, a streaming media stick, or a tablet.

Today's viewers expect a more personalized video viewing experience. MSOs and cable operators need to implement capabilities to better understand who the users are in order to provide smarter personalized experiences—for example, knowing the user's digital video recording (DVR) history, favorite channels, preferred genres, favorite actors and actresses, or the next episode they've queued up to watch. All of this contributes to a more satisfied audience.

The transformation from a CSP to a DSP providing personalized digital experiences will require a change in the way CSPs engage their customers as well as their business partners. Every video or broadband interaction, promotion, trailer, advertisement, landing page etc. should be captured and matched to an individual user. Personalization eliminates the guesswork. It “unbundles the bundle” and opens the door to a wide variety of futuristic service offerings and business models. But, it requires a thorough, accurate and integrated digital user identity mechanism. This paper discusses a new digital user-identity approach (both the challenges and monetization opportunities) that will enable CSPs to transform to DSPs by identifying and managing their users as active individuals versus passive members of a household address.

The new digital identity consists of:

- A secure authentication and authorization process enabling an excellent user experience
- An intelligence-powered, personalized customer journey using a single ID
- An integrated experience across owned and partner-enabled digital services
- A simple entitlement management capability
- Dynamic grouping, such as families and ad-hoc communities

Digital Identity

1. Challenges of the Digital Service Provider

1.1. Connecting to the Users

For mobile service providers, it's the SIM that identifies a user. The IMSI number that attaches to a SIM is used to identify the user and the user's entitlements, while billing is implemented based against the user's telephone number.

For cable operators, transactions are completed based on a home address, which was originally used as the basis for feasibility and serviceability. Address-based identification is not a sustainable identification approach for cable operators wanting to compete in a market where over-the-top (OTT) digital companies have redefined the user experience. DSPs of the future will need to center their business around user-specific digital IDs to improve the user experience while providing user-centric personalization, as well as giving users the ability to self-manage their entitlements.




Figure 2 - Digital ID Transformation

1.2. Collecting the User's Journey

DSPs will need to provide services to a wider variety of users with different expectations and needs, including:

- An occasional user (not a subscription customer) who purchases a one-time event
- A wallet user (not a subscription customer) who wants to pay for digital services
- An user of OTT services who pays a subscription fee using a credit card
- A connectivity subscriber (recurring monthly subscription)

For the traditional SP, these different types of users are typically managed in different systems and the same customer may have different user names. If customers are required to self-onboard for multiple services, the SP may be unable to collect necessary information, personalize their experience and manage their journey.

The examples below (see Figure 3) depict a variety of personalized digital experiences. Each experience typically requires a unique registration and authentication. How will the DSP be able to capture the full user identity and manage the user's journey if they cannot associate the distinct interactions with a single user? How can a DSP deliver a true, holistic personalized experience with a disconnected user journey?




Figure 3 -The Disconnected User Journey

1.3. Creating a Trusted Environment

Once the DSP is able to fully capture the user journey, it needs to navigate the delicate association between trust and the likelihood of individuals to leverage and embrace personalization. How do you avoid making personalization intrusive and creepy? DSPs must find the right balance regarding the level

of personalization and privacy, meeting regulatory requirements, and the value and user experience the consumer receives in return.




Figure 4 - User's Trust and Value Balance

On the one hand we have privacy regulations, set by the regulator. In the European Union (EU), General Data Protection Regulation (GDPR) defines the new data regulation while in the US, the FCC is in charge of privacy regulation. For example, Verizon was recently fined more than \$1 million by the FCC for “supercookie” tracking which captured user activity without any notification or “opt out” option.

On the other hand, what the user pays varies by digital provider and the value to the consumer is in the “eye of the beholder.” A variety of Google services are free, but in return for a free service, a consumer is willing to give up much more information and privacy (and perhaps a drained battery). Some people use and rely heavily on Google services while others hardly ever use Google. Consumers pay an annual fee for Amazon Prime, but in return Amazon provides free streaming and free delivery as part of this prime subscription. and a different value proposition for their user community.

Consumers are constantly checking and validating the value they are receiving in return for the data they share. If consumers acknowledge that they are getting benefits from sharing their data, they will be more likely to provide consent for collecting and using their data for personalized purposes. As an example, Telefonica’s “Giving the Data Back” initiative provides an entirely new edge for Telefonica to differentiate themselves from their competitors by giving data back to their customers. One of the use cases is enabling consumers to manage some data functions on their own, and eventually generating royalties from partnerships or selling data to partners with customers’ permission. So Telefonica collects and maintains consumer data but returns it to their customers for monetization purposes.

The amount consumers are willing to pay depends on the value they expect and experience in return. For example, consumers are willing to pay a monthly fee for Netflix for the value of the commercial-free content (on-demand and original) they can access on top of the standard cable and broadband monthly fee. Comcast recently announced the launch of a new service for \$5.99 offering users the ability to

receive an ad-free, premium FX network experience. You can expect to see more SPs experimenting with the user's trust and value balance equation as part of their DSP transformation journey. Finding the correct balance is extremely important for user adoption of true 360 degree personalization.

2. The Opportunities with Digital Identity

Digital identity plays a critical role in creating a new personal TV user experience, building trust between the DSP and the consumer, creating a strong digital brand, and finding new ways to monetize the DSP's services.

2.1. Digital Identity for TV

For cable operators, the most important piece of information they currently have is the home address, and while we can all imagine the traditional picture of the family (or some of the family) watching TV in their living room, unfortunately, the cable operator doesn't know *which* person is actually consuming the service at any particular time.

"Watching TV" no longer means viewing content via a physical TV mounted on the wall. Each family member can now "watch TV" using personal devices like tablets, or shared devices such as a set-top box or a streaming media stick or within a digital application. Since viewing and internet habits differ from person to person, it's become critically important from a service standpoint for MSOs to understand exactly who the user is, as well as the groups they belong to, in order to personalize the services for that specific user and so deliver a user experience to rival or exceed those provided by OTT and other digital service providers.

Leveraging data at the user level will enable cable operators to:

- Provide recommendations about new series and episodes
- Give consumers the option to switch between favorite channels easily
- Give consumers a more personalized DVR experience
- Give consumers the option to consume on multiple devices switching mid-session
- Create a connected user journey providing a more holistic and intimate understanding of the individual

MSOs will also need to personalize at a group or family level in order to enable users to perform so tasks as:

- Managing parental controls, such as the type of content a child is allowed to watch, (and during which hours)
- Managing the amount users are allowed to spend (e.g. video on demand, gaming) from the family or group allowance
- Enabling the user to share specific events with their friends

2.2. Engage with Every User

A single digital identity enables collecting the information on the user through the user lifecycle in the correct sequence and context. Starting with a common sign-up and easy onboarding for new customers, to managing user journeys with different services consumed, there are many opportunities to provide a compelling customer experience while increasing the revenue per user via a single digital identity.

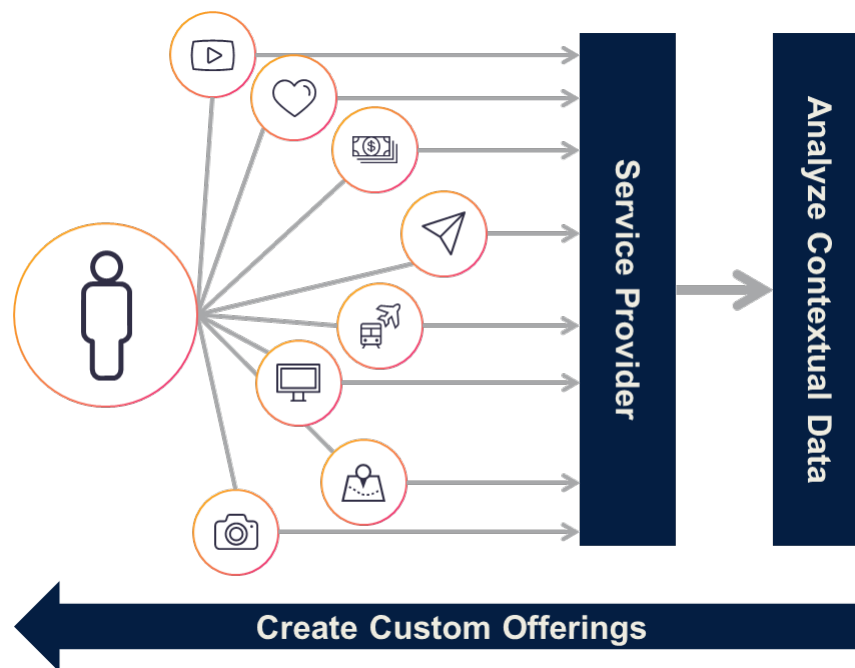


Figure 4 - Manage User Journey; Increase Revenue Per User

Big data and advanced, contextualized analytics are game changers and are absolutely required for extracting the value out of the digital identity and user journey. However, CSPs and DSPs continue to struggle with solving the challenge of managing massive amounts of data collected from all contributors to the user journey in a highly structured way. DSPs will need to leverage big data in order to make consumers feel individually valued at precisely the right moment and the right location.

2.3. Digital Identity – Asset to Generate New Revenue

Digital identity is a key enabler for DSPs to create additional revenue in the following areas:



Figure 5 - Monetization opportunities

2.3.1. Targeted Promotions

With the transition from an address-centric paradigm to a user-centric paradigm, cable operators will be able to engage with users in the household, personalizing their experience and recommending new services – for example offering a parent in the household a gift card for their son’s birthday which could be used for online games, pay-per view movie, etc. Perhaps the cable operator partners with an on-line retailer (another digital service provider) to offer a gift card for the toy that their son has been searching and researching on his tablet over the past couple of weeks using the family’s broadband service. This type of robust targeted promotion offering is only made possible with Business to Business (B2B) partnerships, but most importantly, digital identification and user journey capabilities enabled!

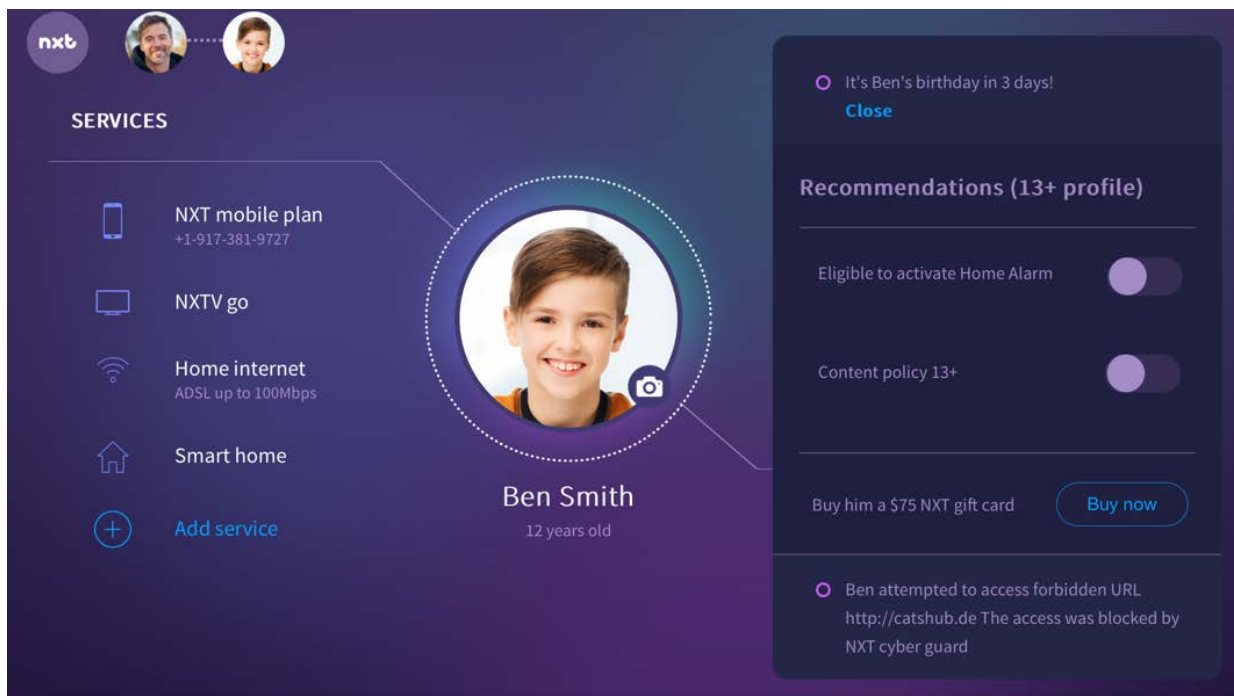


Figure 6 - Example of Targeted Promotion

2.3.2. B2B Partnerships

Becoming a DSP, or digital aggregator, requires working with partners from different industry sectors such as entertainment, retail, hospitality, travel and leisure, smart home, connected cars, connected health, IOT device providers, and many more in order to offer consumers a range of timely and contextually relevant digital services. There are several challenges to creating a meaningful and seamless consumer experience when working with digital partners – these include:

- How will the user purchase the partner’s service? Will they need to register on the partner’s portal?
- How will the user consume the service? Do they need to login to the partner’s portal?
- How will the DSP differentiate between the users?
- How will partners capture user information in order to deliver a personalized experience?
- Will the user be able to manage family/group entitlements?

Russian service provider VimpelCom recently rebranded itself Veon as part of their DSP transformation. “We are doing two things”, explained CEO Jean Yves Charlier. “We think we have to be a great telecoms business providing connectivity, but we also think we have got to do much more than that, and that’s what we are focused on. We want to bring a new digital model to the industry, not just a bricks-and-mortar model.”

The potential acquisition of Time Warner Inc., by AT&T, who once referred to themselves as “The World’s Networking Company,” is yet another step in AT&T’s plan to become a DSP providing premium content, advanced advertising and a variety of OTT capabilities.

Fifty percent of CSPs’ new digital services originate from partnerships and investment. The diagram below shows possible partnerships to fulfil the needs of the digital savvy consumer.

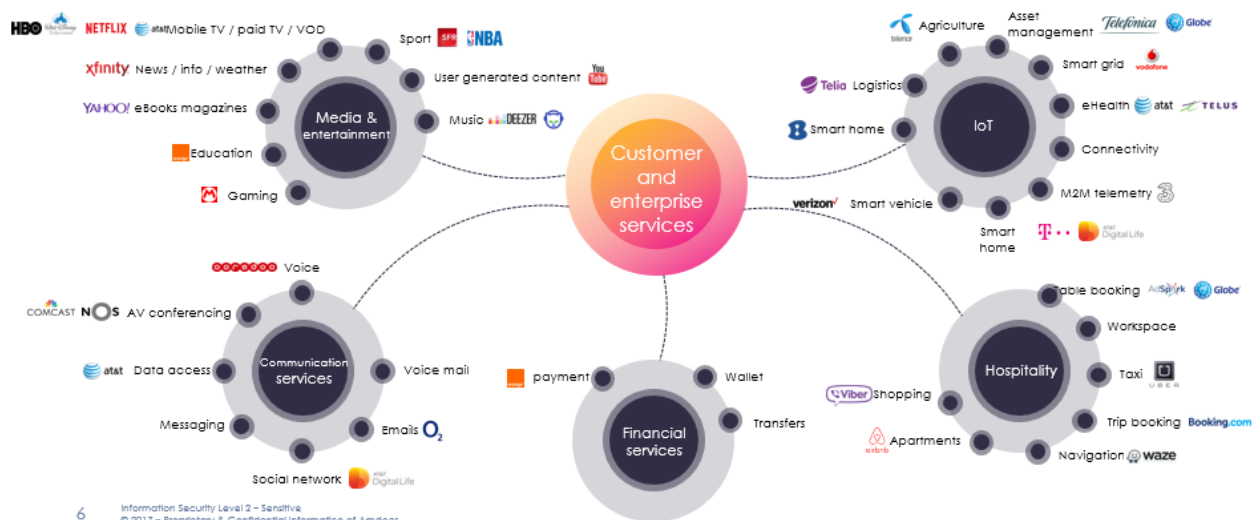


Figure 7 - The Digital B2B Economy

2.3.3. Advanced Advertising

Targeted advertising enables DSPs to reduce advertising ‘clutter’ and improve the overall customer experience. But targeted advertising is challenging and less effective when customer information is scattered. An orchestration function or layer to correlate information from multiple systems is required to get a more comprehensive user profile for effective targeted advertising.

Personal profile data is today’s advertising currency reaching across channels like mobile, web and TV as it increases the targeting capabilities and correspondingly the advertising yields. Digital advertising is at a point where advertisers are willing to pay a premium to service providers for subscribers’ first-party data. Generally defined, first-party data is specific data collected by you about your audience (i.e., the first party is “you”) as compared to third-party data which is acquired from outside resources (i.e., a third party) and provides generic, segmented information. This creates an immediate opportunity for service providers to monetize new channels by using an enriched customer profile that combines relevant first and third-party data for use with interactive targeted ads, which increase subscribers’ engagement and affinity with the advertised product.

First-party data or customer profile data available to SPs is unique. It combines a variety of location, usage and spending behavior that provides good indications about the customer’s interests and needs. Data can be gathered ongoing from multiple platforms and interactions, and then packaged in a data model that enables service providers to create audiences. These audiences, based on rich customer insights, are highly sought after by brands and advertisers in the advertising ecosystem.

DSPs with properly implemented digital identity and B2B relationships will be able to go beyond regular behavioral targeting. They will be able to access enhanced customer profiles, consisting of content, service and user data updated in real time by multiple subsystems in a flexible, efficient, and cost-effective way. By leveraging this information, service providers can create a high-value advertisement inventory for themselves as well as for advertisers.

Targeted advertising is not just about offering the right product to the right person at the right time. It is about “closing the loop” and encouraging the consumer to take action. For example – a car advertisement which also includes directions to the nearest car dealer to the customer’s location, and which can be accessed by the customer across multiple channels, touch points and devices.

By leveraging the digital identity of their customers, and the data collected from their activities, DSPs can:

- Improve user experience by providing relevant ads to their customers
- Increase the value of multi-screen advertising
- Deploy a targeted ad platform with innovative third parties – the operator can license, partner, or sell ads to the content owners and television networks
- Maximize revenues from digital advertising

2.3.4. Beyond the Quad-play

Some OTT providers are already monetizing their service with different paradigms – examples include:

- *Family Sharing with Apple* – family sharing makes it easy for up to six people to share apps, iBooks, iCloud, music and more
- *Netflix's screens pricing model* – pricing determines how many users can access a Netflix account at the same time

The new bundle will definitely embrace personalized and targeted content, and for these new innovative monetization models, the service provider needs to have information about individual users and devices.

2.4. The Trust Factor

It is not surprising that in today's digital world and the rise of hacking, consumers are paying more and more attention to their personalized digital data and how it is used. In a recent global consumer survey almost half the respondents emphasized how critical it is for them to know that their provider keeps consumer personal data in a place that is extremely private and secure. CSPs and pay TV providers are relatively well trusted. They can leverage this trust as they transition to become DSPs, or digital aggregators, where they offer a more personalized experience and sell more services. In order to meet privacy requirements, DSPs will need to identify who the user is and make sure they have appropriate consent for any action that has to do with use of personal data.



Figure 8 - Level of Trust Cross-Industries

2.5. Contact Center of the Future

Customer support is changing and DSPs will have to engage with their customers across a range of channels – traditional channels such as call centers and interactive voice response (IVR), and increasingly online (online chat, email, etc.) and social channels (e.g. Facebook, Twitter). Service providers need to keep track of their customers digital journey across all channels for all services. By having a complete

view of their customers, and by analyzing the interactions with these customers, service providers are able to move from reactive to proactive in their support and sales efforts.

3. The Digital Identity Solution

3.1. Digital Identity Ecosystem

DSPs need to create a digital identity ecosystem that maintains trust with their customers regarding use of personal information and fosters collaboration with their digital partners. The primary component of such an ecosystem is a digital identity solution that eventually empowers a digital locker while enabling new, flexible monetization models.

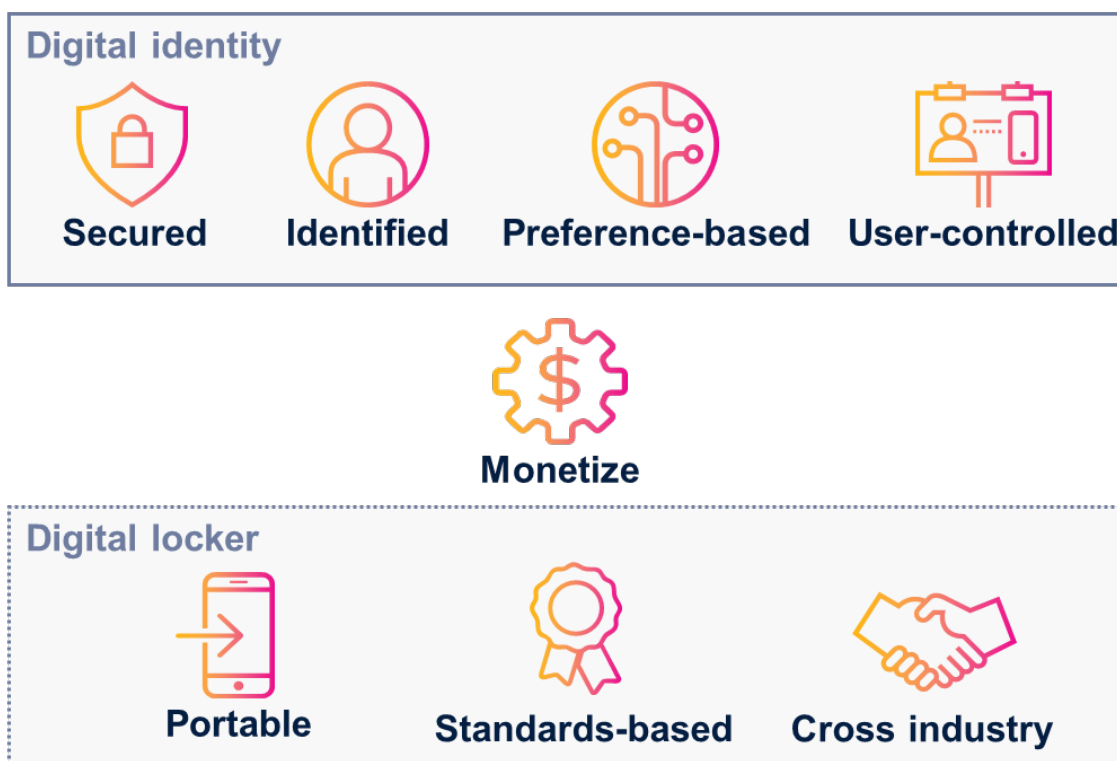


Figure 9 - Digital Identity Ecosystem

The digital identity component of this ecosystem consists of the following attributes:

- Highly secured user identity information
- Individual user identification for all services on any device
- A collection of user preferences based on user consent
- User control determining who can see and access their data

In order to facilitate usage, sharing and, eventually, advanced monetization models, the user's digital identity should be stored in a digital locker consisting of the following attributes:

- Users should be able to port their identity using it wherever and with whoever they choose
- User preference and API standards need to be defined to support porting
- Adoption between many partners, across different industries

The digital identity component is the foundational enabler for expanding monetization opportunities. It needs to be fully integrated with partner contract management and billing systems and enhance the DSP household level billing to individual user billing.

3.2. Digital Identity Solution Functions

To implement this digital identity ecosystem, CSPs looking to transform to DSPs need to ensure they have the following areas addressed:

3.2.1. Identity Management

Identity management should maintain all of the consumer's identities and associated information. The DSP will need best practices defining how to manage the user lifecycle: onboarding users, converting anonymous users into subscribers, how to identify a suspended user, etc. Identity management also needs to enable a user to manage both their own and related identities. Some examples of required functionality include:

- Self-registering for services
- Granting permissions to other users within an account (e.g. parent enabling service for child)
- Registering a new device and associating it with a user on the account

3.2.2. Access Management

Access management to provide user authentication and authorization can be activated in various ways (e.g., username/email and password, biometric authentication, two-factor authentication, mobile authentication). For effective authentication, the DSP will need security policies such as strength of password or token expiration period.

Access management in many complex organizations is determined by a single sign-on (SSO) property.

3.2.3. Authorization

Once users have been authenticated by the access management system, they can use the services that are enabled by the provider based on their authorization level. The authorization can be based either on the role of the user or according to their individual entitlements. Roles can be provided to a group of people and dictate their permission to access specific activities, and/or, information.

3.2.4. Single Sign-on and Easy Integrations

While the DSP will rely on partners to provide different digital services, the desired user experience is that the entire registration process will be via the DSP's portal. The user should not be required to register on the partner's portal or obtain an additional username and password – a single name and password should enable the user to access both the digital service provider and their partners' services.

The DSP's access management should support single sign-on standards like SAML or OpenID Connect. If the partner does not support these standards, then integration is required and integration tools must be put in place.

3.2.5. User Entitlements

User entitlements cannot be based solely on the user's role. Entitlement models are complex and need to take into account such actions as:

- Did the user pay for a specific service?
- What is that user's position in the account hierarchy?
- How many users are already signed on to the service?
- Maintaining previous entitlements

Identity management should support the data model of the entitlements, and access management should check when a user is trying to access a resource whether that user is entitled to such an activity.

3.2.6. Device Registration, Entitlements, Consent Management

When it comes to the Internet of Things (IoT), it is very important to enable the user to register new devices – even if they haven't been purchased through the DSP's channels – and allow them to manage entitlements for specific users on these devices. For example, a person could configure the system so that only the household adults have access to adjust thermostats in the home.

As part of the new EU security regulations (i.e. GDPR in 2018) and in the USA, users own their personal data, and as such, the user must give their explicit consent for their data to be used for a specific service.

A consent management system is therefore required in order to:

- Enable a consent lifecycle (i.e. the ability to consent, revoke and freeze data consent)
- Enable parental consent management
- To link between consent management and the user's personal data
- Allow user consent to be distributed to partners.

The consent information is stored in the user identity as part of the user profile. Access management should restrict access to personal information according to the user's consent.

4. Integrating Digital Identity E2E Solution

Realization of the digital identity ecosystem requires integration into the existing business and operational environment. The comprehensive, trusted digital identity becomes the cornerstone of the end-to-end operational flow. Below is a sample solution of the digital identity ecosystem integrated into a complex end-to-end DSP environment.

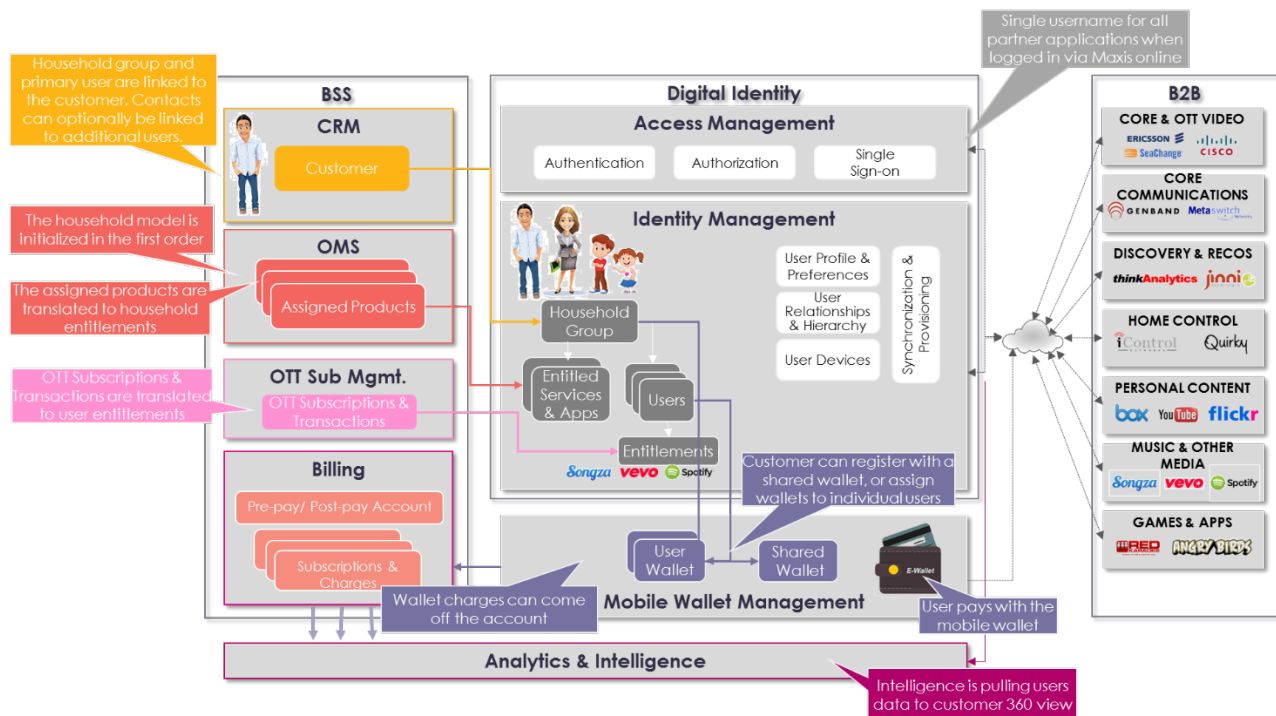


Figure 10 - Integrated Digital Identity Ecosystem

The user ID serves as an identifier to many of the surrounding systems and connects to the mobile wallet, BSS, partner B2B and data analytics systems.

User identity is foundational to any BSS environment including global billing and monetization engine(s). It enables the introduction of fast time to market digital content and services, including real-time personalized promotions and merchandising for both existing and anonymous customers. In the pre-paid system, the user ID identifies a specific balance. In the wallet system, the user and group's identification will authorize who within the family is permitted to pay using the wallet. In the video system, the user identity validates the customer's entitlement for a content.

The user identity stores previous customer actions and connects to the analytics engine. This enables the delivery of intelligent, contextual and personal interactions across the customer journeys with up-sell and cross-sell opportunities by collecting and analyzing the variety of customer data available through the provider's own platforms as well as partner-originated data, turning it into action by predicting and automating any customer engagement across the lifecycle. Digital identity drives higher revenues, service uptake, loyalty, satisfaction and engagement.

Digital identity powers efficient connection to partner B2B systems and manages all the digital services from one place.

The digital identity in the above solution is comprised of two main components: access management and identity management.

The user management system stores the users, groups, locations, devices, the customer entitlements and the user entitlements information. When a customer purchases a new service, that service is distributed to the user management system along with the customer entitlement, devices and location information specific to that user. The user can manage these entitlements for the family, along with the location and devices data, by using a graphical user interface (GUI).

When a user logs in to the system, the access management GUI will invoke a series of authentication steps which are processed according to its configuration. One of the authentication steps is to validate the user entitlements. Once the user is authenticated, the SSO component in the access management system uses an access token to login to any service that the user is registered, including operator services and partner services.

Conclusion

In today's all-encompassing digital world, where many of our day-to-day business and social interactions are anonymous, customers' are demanding a well-shaped, personalized experience. Spotify and other digital services, such as Netflix and YouTube have given birth to the personalized music and video DJ, supplying us with targeted content at the exact moment we want to consume it.

But today's TV viewers have an overabundance of content to choose from, and have come to expect a timely, purposeful, customized experience. It will not be long before traditional TV, and for that matter all connected digital services will be provided by our own personalized content DJ.

Transforming to a DSP requires adopting a more holistic approach to user identity, which takes into consideration the way today's viewers are consuming their entertainment: anytime, anywhere, on multiple devices. In order to put the end user in the personalized driver's seat, individual consumers will need to be able to create and manage their own profiles. But this movement from household to individual is not sufficient as it only addresses one of the many dimensions of a total personalized experience.

A true 360-degree personalized experience will require the DSP to deliver highly relevant experiences across every connected device and digital channel. There are large amounts of customer digital touch points and a wealth of data available, but it's scattered, often non-standard and disjointed. To really understand customers and deliver the most relevant experiences, a DSP needs to be able to stitch together all the digital interactions and associated data sources under one digital ID in order to track the user journey.

To enable a well-connected and contextualized user journey, DSPs must demonstrate they can be a trusted partner. Customers are only willing to share their information if they know (1) it is being using to help them get what they want, when they want it, (2) the information is securely maintained, and (3) consumers can control how their data is used.

Once the digital identity ecosystem is established and embedded within the end-to-end operational environment and the CSP has transformed to a trusted DSP, the monetization door will swing wide open, providing access to a variety of new, innovative B2B and B2C monetization possibilities. If executed properly with the right balance between trust and value, DSPs can enable and support blended, contextualized, and customized experiences bridging and integrating our different day-to-day digital interactions along with our entertainment and social experiences.

For the consumer, it will feel like they are carrying around their identity inside a digital locker which can be plugged into their daily activities and subsequently leveraged to drive an enjoyable, finely tuned experience. Customers will then become comfortable sharing the contents of their digital locker across different industries to support a contextually relevant experience that reaches from their living room, to their car, to the store while traveling with them to and while on vacation.

If you are a CSP today, are you ready to transform and become a DSP, assuming the responsibility of managing users' identities as a trusted partner, creating new B2B partnerships in order to enable great personalized experiences for your users with a rich portfolio of digital services?

We are all consumers of digital services and content. As a consumer of a variety of digital services, are you ready to move from account centricity to individual identities for your family? Are you willing to have your digital locker full of personalized information? Are you ready to trust your identity with your DSP? But most importantly, are you comfortable sharing your personalized journey in exchanged for unique, effortless, meaningful digital experiences?

Abbreviations

B2B	Business to Business
B2C	Business to Consumer
CSP	Communications Service Provider
DSP	Digital Service Provider
DVR	Digital Video Recorder
E2E	End to End
EU	European Union
FCC	US Federal Communications Commission
GDPR	General Data Protection Regulation
GUI	Graphical User Interface
IMSI	International Mobile Subscriber Identity
IoT	Internet of Things
IVR	Interactive Voice Response
MSO	Multi-System Operators
OTT	Over The Top content
SAML	Security Assertion Markup Language
SIM	Subscriber Identity Module
SP	Service provider
SSO	Single Sign On

Bibliography & References

“Level of trust cross-industries – driving the roadmap to digital success”, lecture by Rob Van Dem Dam, Global Telecommunications Industry Leader for the IBM Institute for Business Value, TM Forum Live!, May 2017

[Telefónica to create personal data bank for customers, expose “unfair” apps](#)

[Telefonica Plans to Give Customers More Control Over Their Data](#)

[VimpelCom ditches bricks and mortar for digital model](#)

Smart Entertainment in the Smart Home

Reducing Friction in Content and Service, Discovery and Consumption, Across Devices at Home

A Technical Paper prepared for SCTE•ISBE by

Arsham Hatambeiki

Vice President, Corporate Product Strategy

Universal Electronics Inc.

201 E. Sandpointe Ave., Santa Ana, CA 92707

(925) 567-3121

arsham@uei.com

Introduction

The connected home has evolved in many exciting ways. Expedited growth of more affordable and compact computing power, and expanded connectivity have enabled rapid growth of “smart” devices capable of offering new services and creating new ecosystems within the home. Many entertainment devices take advantage of this connectivity to create a cooperative ecosystem of apps and features covering different forms of content delivery as well as smart home services, resulting in a series of new compatibility requirements. **Defying popular belief, new does not always replace the old and quite often can be an addition.** Most homes contain a wide array of traditional devices that are not connected and are now mixed in with the newer connected devices with enhanced capabilities such as voice control. Based on a QuickSet Data Insight study, more than half of the television (TV) installed base in the United States are five years or older. This varying level of “smart” and interactivity causes discontinuity and confusion for consumers in accessing and controlling each device to enjoy their content.¹ This obvious gap in a common approach to interactivity has created an area of opportunity for device manufacturers and service providers. Consumers need a simple and unified approach to interacting with their devices at home, new and old, connected or not, offering a reliable experience with all.

The delivery of entertainment content has been particularly affected by continued improvements in Internet connectivity speeds and processing power in devices. Innovative startups and well-known names in entertainment alike are looking for a winning edge to lead consumers to their content and their services. Content providers are offering a dizzying array of almost unlimited content, available through multiple channels, and some exclusives to a single channel. The starting point in content discovery for consumers now includes necessary decisions such as which app and on which device, and finally, the main question-- what to watch!

Today consumers are expected to decipher where and how to access a specific service or content, **which app on which device through which control interface, a constant context switching from one content source to another, adding friction to daily TV watching activities.** Friction in the user experience will cause fatigue, and users will then quickly stop using a device or service and revert back to the previous methods they understand; after all, watching TV is supposed to be fun. A consumer will naturally be attracted to devices that offer a consistent and reliable experience in finding content and services. This has created an opportunity for service providers and consumer electronics manufacturers to provide an elegant solution to this daily problem.

Successful solutions, with wide adoption, have taken the realities of the home ecosystem into consideration, providing a reliable experience across the installed base of old and new devices in consumer’s home. As an example, while in the first generation Apple TV attempted to address the user experience and TV compatibility with a single protocol such as High-Definition Multimedia Interface Consumer Electronics Control (HDMI-CEC), the second generation now includes multi-protocol support including limited infrared, for TV control; an improvement, but still not the ideal experience for a mainstream product.

Trend-setting set-top boxes from Dish Network and Comcast have employed techniques to automate the set up experience and make their **content reachable with a single touch**, while smart TVs have offered a **unified discovery and control experience across different content sources** and services to manage the daily activities. Some well-known brands also offer universal content search solutions that focus on the experience on a single device, searching across a catalog of content offered by installed apps and

channels, while others strive to unify the experience across devices and services through a combination of content recognition techniques and pre-defined indexed catalogs.




Figure 1 - A Day In A Life

Leading consumer electronics companies are already benefiting from providing a unified interface to interact with different devices and services, taking hold of the **starting point in the users' daily content consumption experience**. This unique position has enabled new business models and revenue opportunities. Solutions such as dynamic media recognition offered by Gracenote and Shazam, or device-centric whole-home discovery and control capabilities have enabled consumers to automatically discover and interact with all points of access to content and applications in the home through a single point of control (a voice assistant; a touchscreen device, or a remote control). These capabilities are offered through set-top boxes, TVs, game consoles, smart home gateways or any connected device offering a service.

Media recognition systems rely on a database of content to dynamically identify matching signatures against live or recorded content. A successful solution, as a device-centric approach, will complement this approach through device specific knowledge, and provide a unified discovery and control interface for devices in the home, through different communication mediums, including High-Definition Multimedia Interface (HDMI), Internet Protocol (IP), and different wireless protocols such as 802.15.4, Bluetooth, and widely used Infrared.

At its core, the ideal solution would enable anyone to **enter a room, seamlessly discover all nearby devices and services offered** by these devices, and interact with these services through an intelligent, simple and natural interface. It should operate independently of any framework tied to specific devices, access protocols or proprietary software ecosystems.




Figure 2 - Device Model in Knowledge Graph

A complete model of a device is the most basic, yet necessary, building block of a **knowledge graph representing devices and services in a home that can add context to all user actions**. Viable platforms rely on algorithms that utilize a knowledge graph of devices with varying control capabilities, communication interfaces and protocols, each identifiable with unique fingerprints.

At the most fundamental level, a capable solution automatically discovers nearby devices through different communication mediums, generates unique fingerprints and matches them to the knowledge graph, to serve up a full range of capabilities. Beyond control, this knowledge graph also adds the much-needed context to all user commands and actions, making dynamic capability discovery of nearby devices possible.

This highly versatile approach to supporting a unified customer experience lends itself to a wide range of brand strategies. Companies such as Comcast, Dish Networks, AT&T, Microsoft and LG Electronics have widely divergent approaches to offering services, but they are all trying to reduce friction in consumption of their services and enhance the customer experience.

In the discussion that follows, we begin by exploring market conditions driving the need for a holistic, automated approach to connected-home device management, followed by a look at how service providers and original equipment manufacturers (OEMs) are meeting that need. In later sections we'll review the enhanced capabilities required to power content discovery across devices, and better smart home services.

Market Overview: More Content, More Apps, More Devices

The ability to provide a consistent experience across devices, hence the need for a universal connected-home control solution, has gained increased importance with cable and satellite operators, over-the-top (OTT) providers, and consumer electronics (CE) OEMs competing for consumption of content and newer services in Internet of Things (IoT) markets. As the increasing volume of gadgets in the home promise a better home, the resulting friction in the user experience imposes an even greater inconvenience to consumers.




Figure 3 - U.S. Pay-TV and OTT Penetration

This presents both a challenge and an opportunity to provide an elegant and seamless experience in how users discover and access their content and services, thus improving customer satisfaction, lowering attrition.

The increasingly urgent need to simplify device control and navigation is affecting not only traditional cable/satellite providers but also makers of smart TVs and streaming services. According to Parks Associates, **63 percent of U.S. broadband households subscribe to at least one OTT video service and 31 percent subscribe to two or more**ⁱⁱ underlining the increasing friction due to context and input switching between devices and services.

In another report, Parks found that 52 percent of households subscribing to at least one OTT service also subscribe to a traditional pay TV service.ⁱⁱⁱ Sixty-one percent of households headed by people from the millennial generation are in this category, suggesting the multiple subscription phenomenon is here to stay.

Consumers use various combinations of devices to get the content they want, including cable/satellite-supplied set-top boxes, smart TVs, media streamers, game consoles, and personal devices such as smartphones and tablets. Measuring actual network usage in 2016, network traffic tracker Sandvine found that the average North American household had at least **seven active IP devices in use daily, with 65 percent of the usage dedicated to video streaming.**^{iv} By 2019, the average number of connected media-enabled devices per North American household will climb to ten, according to IHS projections.^v

These trends are replicated worldwide. U.K. analyst Ovum in a recent report predicted online video subscriptions, which topped 100 million worldwide at the end of 2015, will increase to 177 million by 2019.^{vi} In another report, Digital TV Research projected that by 2020 online subscription revenues will reach \$21.6 billion, three times the 2014 total, with penetration exceeding 33 percent in ten countries.^{vii}

All these statistics point to the fact that streaming service subscribers, including those who subscribe to cable/satellite services, will be **consuming content through multiple channels and on multiple devices, expanding the ecosystem of devices connected to their TV sets** and adding to the complexity that comes with having to switch TV inputs, applications, and remotes from one service to another. The ability to simplify accessing a provider's services by eliminating this difficulty will only grow more significant to their overall value proposition over time.

In another survey involving 1,812 participants in the UK, France and Germany, researcher Trendbox found that, with an average of 3.3 remote control devices per household, consumers are hungry for a solution that overcomes their usage hassles.^{viii} **80 percent said they want a control device that's easy to set up and use** for accessing all their content on their TV sets. In an acknowledgement of the confusion people experience navigating across TVs, set-top boxes, game consoles and other content sources, **34 percent said they have difficulty getting the content they want onto their viewing screens.** Also important to note 41 percent stressed the importance of eliminating the button clutter common to most remote controls, which highlights the importance of a simple and automated solution.

Satisfying these demands not only adds to the appeal of a given provider's services, it also saves money by reducing customer support calls. 10 percent of the respondents in the Trendbox survey said **they contact customer support to get help when they have trouble navigating among content sources.** This is an especially big headache for cable/satellite providers, who are often held accountable for consumers' OTT as well as legacy pay TV usage issues. Such calls, owing to the complexities of the issues, can potentially consume more time than other types of calls.

The multi-device usage issues surrounding TV services and applications are becoming equally problematic in the smart home realm as consumers engage with multiple applications from different sources. Gartner predicts the number of installed IoT-related devices in homes and businesses worldwide, not counting smartphones, tablets or PCs, will grow to 26 billion by 2020.ⁱⁱⁱ Gartner says that by then IoT component and connectivity costs will be so low that just about everything – light fixtures, windows, doors, appliances, toys and sporting equipment – will be equipped for IoT connectivity to support control, monitoring and sensing.

As noted by Parks Associates, the adoption rate is especially high among millennials. As IDC noted in a recent report, by 2018 millennials, who are the leading adopters of IoT technology, will comprise 16 percent of the world's population, which means IoT adoption will be far more commonplace than it has been so far.^{iv} Consequently, there's a growing opportunity for IoT service providers who can aggregate whatever specific applications consumers choose into a unified experience under the control of a single device.

Case Studies: Power & Versatility in Deployments

Every device manufacturer and service provider has different needs when it comes to their brand experience. TV manufacturers have different goals from streaming service providers, which have different objectives from cable/satellite providers.

Smart TV brands tend to be content-agnostic, and **newcomers as service providers**. Although in recent years, more and more TV brands have begun to serve content from sources they promote, many consider their platform as a common interface for all content. TV manufacturers are in a very competitive marketplace, striving to add value and gain recognition among many competitors. An excellent method to reaching these goals is by providing an easy and satisfying out-of-box experience that can quickly and easily discover and connect to content sources - a consistent and reliable starting point to the daily journey of content discovery and control across different sources.

Streaming service providers and more popular cable and satellite service providers are looking to make their service as easy as possible for subscribers to find and consume. This is necessary in both initial set up as well as daily use where consumers increasingly can be **jumping between content sources**. Getting back to the service provider's content easily is paramount, and the best solution is a single touch or a single "Watch TV" voice command.

1. Simple Content Access on Set-Top Box: Comcast Xfinity

Incumbent service providers are faced with an onslaught of competition from OTT, TV, and other streaming services that are vying for content viewing from the consumer. Traditional grid-based menus and scrolling through long lists of content are now a thing of the past for a competitive service provider. Since content is coming at consumers from multiple devices and from multiple service providers, cable operators like Comcast have taken steps to further refine their experience at finding and switching to content.

Comcast, one of the largest providers in the world and a leading innovator in the segment, has recognized and made considerable advances in further refining their subscriber experience in accessing and enjoying live content and video on demand. Building upon their unique infrastructure capabilities, the recent migration to the cloud based X1 platform delivering Xfinity services has enabled a much more agile approach to delivering better services. This approach has shown great promise for Comcast in maintaining subscribers and providing improved video on demand (VOD) content consumption.

Switching inputs, changing remote controls, and keeping it all straight was a threat to continuity for subscribers and was not a pleasant experience in finding entertainment content. Comcast was able to address the complexities of providing new and improved personalized services in shorter refresh cycles, and adapting to a more complex home environment with increasing friction from context switching between content sources. This **reduced initial install costs** through a combination of self-install kits and shorter professional install times, enabled by a better out-of-the-box experience.



Figure 4 - Xfinity X1 and Voice Remote Provide an Easy Viewing Experience

The solution’s ability to discover, recognize, and automate control of multiple devices brought the X1 platform one step closer to a truly automated experience in connecting consumers to their desired content. “What excites us about this feature is that it brings our customers closer to the experiences they love,” notes Jonathan Palmatier, vice president of product development and consumer devices at Comcast. “Instead of fumbling with remote codes and instructions, they can just unbox a new remote and start discovering great content right away.”



Figure 5 - Xfinity X1 Cloud Services for Device Discovery and Control

Through a scalable cloud-to-cloud integration model, the Xfinity X1 platform has been able to roll-out cloud-enabled features across their nationwide footprint quickly and efficiently. The X1 platform is now capable of automatically discovering televisions and audio devices connected to a set-top box, and retrieving a complete set of control capabilities for these devices. **This dynamic and real-time configuration capability is available throughout the lifecycle of the product**, and automatically reconfigures to address any changes in the consumer entertainment system, such as a purchase of new equipment, or rewiring of the audio/visual (AV) system.

As an example, a simplistic implementation for volume control based on a single protocol such as HDMI-CEC would limit the compatibility of such a rollout to less than 8% of households in North America, where Comcast provides compatibility with over 98% of the install base in the same market.

2. Unified Content Interface on Smart TV

For TV brands, simple internet connectivity is no longer enough to generate consumer excitement. A TV, as the single screen shared by all entertainment devices and applications, is in desperate need of harmony and unification between the available content sources, and services in the home.

In growing from a device manufacturer towards also a reliable service provider, today's TV platforms must gain consumer trust to compete with incumbent sources of content and services. **TVs are in control of the starting point in our daily content consumption journey.** The first step is turning on the display, which makes it one of the more natural places to offer a unifying interface across devices and services. Smart TV content dashboards and applications can fall short in providing continuity in the experience if they simply focus on native content, and overlook dominant sources of content in the home today like cable and satellite set-top boxes and streaming boxes.



Figure 6 - Automatic Discovery and Setup of Connected Devices

In the OEM domain, TV manufacturer Samsung caused a stir in early 2016 with the introduction of its 9500 series Smart TVs, marking an industry-leading shift to a novel approach in consumers' experience across content sources and devices. They turned the display into **an intelligent hub for what consumers truly care about: content and services**. Leveraging a new solution and an external HDMI hub concept, Samsung launched a television platform and remote that discovers and automatically controls entertainment devices connected to it. The automatic set up allows consumers to immediately switch seamlessly between linear and on-demand programming delivered via cable/satellite operators, and streaming content delivered directly via IP.

The manufacturer has won wide praise from reviewers for this innovative approach, putting consumers' needs for multi-brand, multi-control connectivity ahead of single brand compatibility. Personalized content and history is presented in a single interface, regardless of the sources, surfacing all services contained within devices attached to the TV. Users can **access a channel, content or service in a single touch**, and the system then handles all the prerequisite steps and configurations necessary to serve up the desired content on that source.




Figure 7 - No Inputs: Unified Content Dashboard with One Touch Tune

In a video review, Geoffrey Fowler, personal technology columnist for the Wall Street Journal, offered these comments: “[The TV] actually fixed the three-decades old pile-of-remotes problem....I can go from my cable box to my Xbox to Netflix in a snap....I can go to my favorite network without needing to switch inputs or touch that gosh-awful TV programming guide....[The platform is] even smart enough to

pull channel data in over the Internet and make recommendations on the screen....A truly universal TV remote has an immediate impact on my life.”^{ix}



Figure 8 - A Unified Dashboard for Devices and Content with Global Content History

Through an expanded set of discovery and control medium capabilities, the device interface software has further refined the vision of a unified and personalized interface across all devices. The solution is now capable of reaching further into devices, discovering expanded content metadata such as current playing media, installed applications, and services contained within a range of devices. This helps address the issue of app confusion in this day and age where the same app may be offered through different smart devices. **This capability, used in conjunction with traditional media recognition and watermarking techniques, can help build a much more scalable, efficient, and personalized recommendation engine** based on universal history across multiple content sources.

3. Competing for “Input 1” by Streaming Services

Streaming content providers, also known as Over the Top providers (OTT), struggle as the secondary source for content consumption. As such, these providers are particularly vulnerable to context switching issues and resulting friction that prevents subscribers from easily getting to their content.

OTT premium service provider Sling TV deployed the [AirTV player](#) to integrate OTT packages and over-the-air (OTA) broadcast TV options into a complete premium TV viewing experience. Sling TV launched the AirTV at CES 2017 to rave reviews.



Figure 9 - One Touch, and You're Back to Sling

The AirTV Player, is a 4K Android TV streaming device that leverages an advanced device discovery and control solution, uses a Bluetooth Low-Energy (BLE) advanced voice remote to access content through a simpler user interface. The AirTV player set up is completely automated, enabling the consumer to simply plug it into their TV and audio devices to begin the process.

In order to address the exaggerated context switching friction, this implementation includes a feature called **One Touch View**. This allows consumers to immediately **go to the desired content source with a single touch**. In this multi-protocol implementation, the platform uses the fingerprint for the discovered devices to identify the characteristics based on the evolving knowledge graph of devices, and select the optimal combination of commands, across protocols, to take the user back to the desired content.

A single protocol implementation with only coverage according to specifications such as HDMI-CEC will limit an essential feature such as One Touch View to less than 28 percent of households in North America, underlining the necessity for a multi-protocol approach for achieving mainstream viability and 3X compatibility for such features.

4. Voice Assistants and the Rise of AI in Home Control

Artificial Intelligence is enabling many new use cases, predominantly through voice assistants available on smart speakers or voice remotes; home gateways and mobile applications in home control. AI implementations are limited by two important factors: (1) **compatibility with mainstream services and devices** that consumers care about, referred to as “skills,” and (2) **a basic understanding of context** which includes nearby “things” and their current state.

[LG Electronics](#)’ transformation of smartphones into whole-home personal assistants illustrates the vision of a handheld device manufacturer in offering a more complete experience for the consumer.



Figure 10 - LG QRemote as a Personal Home Assistant

In recent years, smart speaker category has shown great promise and major improvements in quick cycles; however, popular implementations remain purely focused on cloud service integration. Such an approach falls short in providing features needed for everyday life, and overlooks daily activities such as content consumption on common AV systems. **A true knowledge graph of devices and capabilities can apply context to user commands, and enable compatibility** with popular everyday devices. The need to understand the context of where content is located and how it can be accessed and controlled, available services on nearby devices and their current state, all provide essential signals necessary in properly deciphering the intent and executing the desired actions. Processing a simple command such as “open Netflix” requires knowledge of nearby devices, and available apps on each, and executing this command requires compatibility to execute this command.

A true home butler would need to provide a consistent experience in discovering and interacting with nearby devices at home, regardless of communication protocol and varying implementations of similar capabilities. This multi-protocol requirement, as well as context knowledge, represents an underlying vision of a connected home. In this vision, infrared and RF compatibility is as essential as IP capabilities, to provide both the necessary coverage as well as the expanded capabilities.



Figure 11 - Discovering Nearby Devices, Apps, and Real-time Status

The ultimate solution adds the necessary intelligence to voice assistant implementations as it is capable of not only discovering nearby devices, but also their applications and status. **Device status discovery addresses the limited contextual understanding of voice assistant systems.** Imagine a simple command such as “pause” or “what’s playing?” instead of an unnatural and verbose version where the user needs to convey context in every command. Simple context inferred from the previous command is relatively limiting in real-life scenarios, content playback can be initiated from different control points, or through different playback devices. Media watermarking and recognition techniques may be employed, but can be limiting in reliability, cost and efficiency. A hybrid implementation which includes a dynamic and device-centric discovery platform would enable a scalable and reliable solution.



Figure 12 - Proactive Notifications based on Nearby Device Status Change

Alternative implementations in the market can achieve this in a closed ecosystem of their own brand or tightly integrated partners, but fail when they meet the realities of a consumer's home where multiple brands and providers need to work in a single environment. So if a voice assistant of the smart home is to be the future, it also needs to address these realities.

Discovery and Interaction with Nearby Devices



Figure 13 - Discovery and Interaction with Nearby Devices

As a starting point, a viable platform should consider varying business and technical needs, hardware and application architectures, and global ecosystem differences. Its mission is to surface the services offered by all devices in home, no matter what wired or wireless communications interfaces they use, and intuitively present them to the consumer, at the right time and place.

The user doesn't have to give any thought to where the content and applications are coming from. They're just there, **hiding the technology behind the experience.**

To accomplish this mission, operators must invest widely to continually update their platform with a wide range of tools and functionalities including software developmental kits (SDK) for simple integration with popular platforms such as Android and Linux, and Web application programming interfaces (APIs) that enable cloud to cloud integrations, as well as resource constrained edge nodes.

5. Device Discovery and Control: Under the Hood

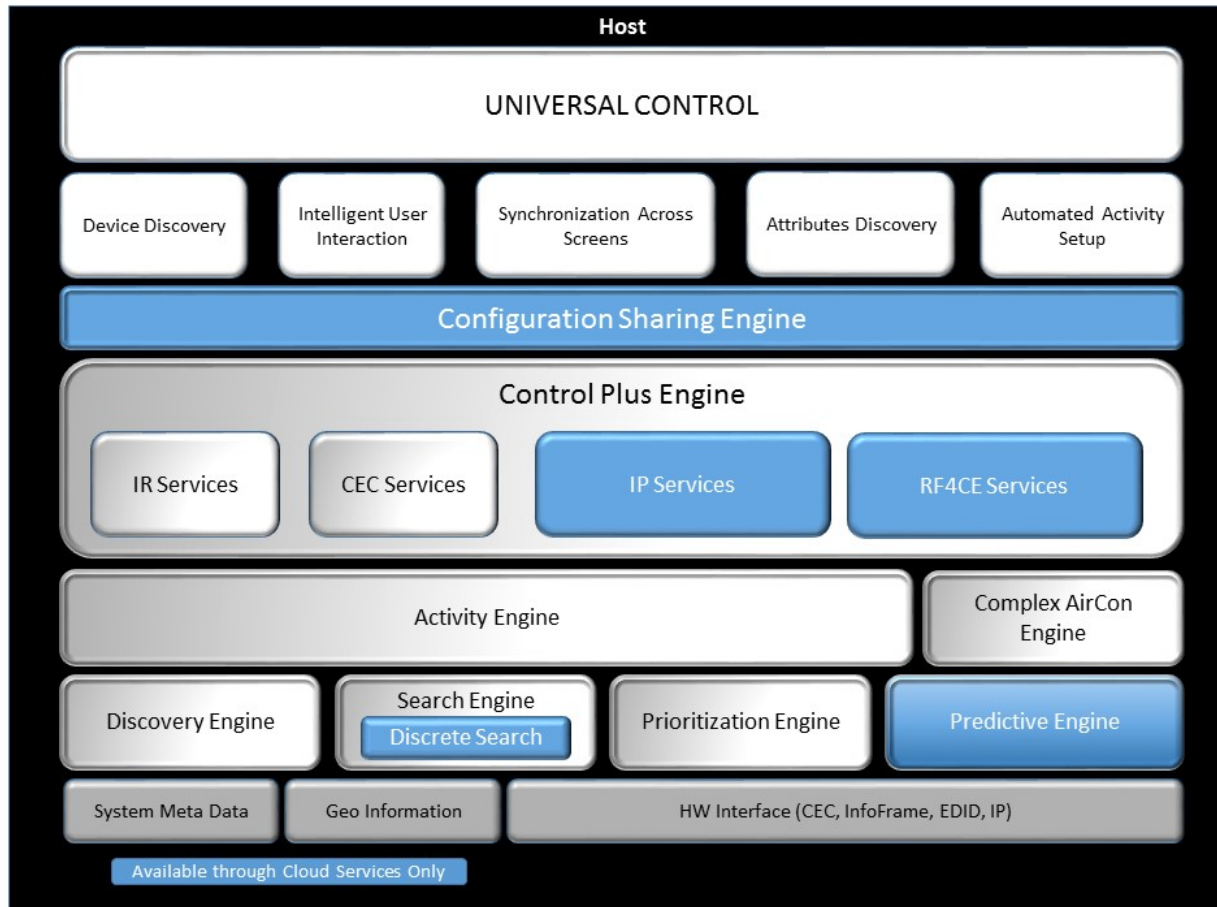


Figure 14 - Device Discovery and Control: Under the Hood

Today, the full range of capabilities embodied in the most advanced solutions comprise the industry's most comprehensive approach to realizing the vision of a unified and consistent user experience across devices in the home. **A Discovery Engine can be built on top of a range of protocols, and millions of device commands and attributes available in the device knowledge graph**, that can be continually updated as new or updated devices come on the market. The Discovery Engine can discover and apply data native to hundreds of thousands of products used throughout the world with no restrictions as to messaging and link protocols used by any given device.

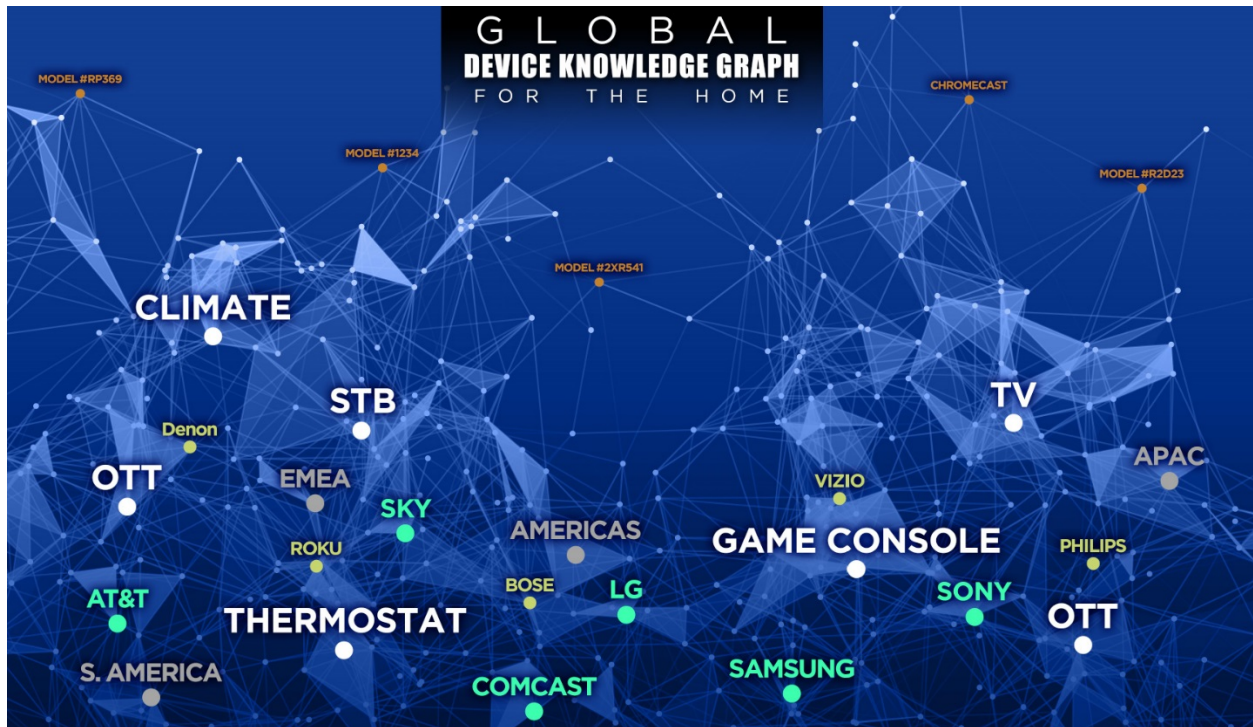


Figure 15 - Global Device Knowledge Graph for the Home

A **multi-protocol method and a unique fingerprinting approach** for identifying devices based on non-structured and non-standard datasets is capable of fingerprinting devices across all communication mediums and physical interfaces in an environment, including a range of characteristics exposed on an IP network, HDMI network and other wireless mediums. This makes it possible to detect the broadest possible range of data, from model and serial numbers to age and region, native video and audio formats, in addition to control capabilities and characteristics.

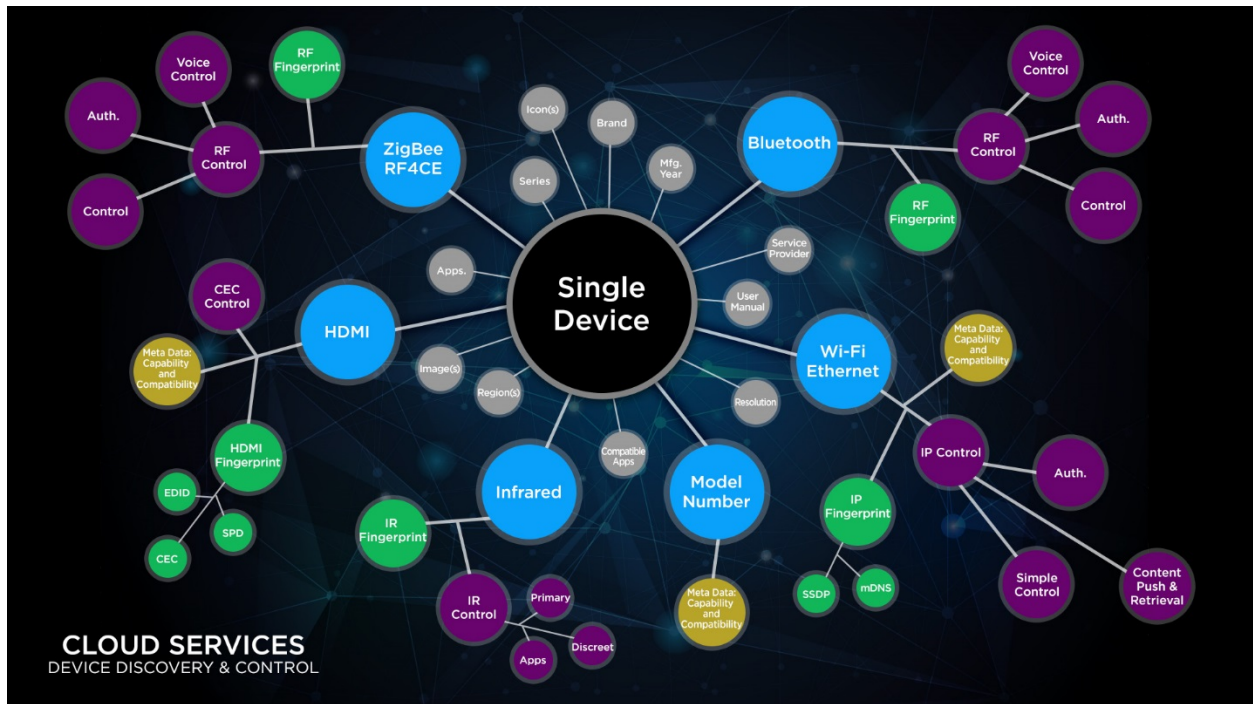


Figure 16 - Multi-Protocol Device Model with Fingerprint

When it comes to control information, multiple protocols must be written and deployed for interacting with multiple devices. This coding must identify the optimal control method(s) for actuating a desired response, which can be on one or multiple control mediums and protocols, including most common wired and wireless protocols in connected home devices.

These capabilities make it possible for providers to fully exploit multi-protocol chipsets and other advances that have been embedded with state-of-the-art smart home gateways, nodes and remote controls.




Figure 17 - Hybrid Architecture: SDK backed by Cloud Services

Today operators are taking advantage of the flexibility, power and intelligence added with cloud-based platforms. The cloud's vast processing power, scalability and updateability to enhance recognition capabilities and new services through state-of-the-art machine learning techniques are generally being appreciated.

A viable cloud-based solution can power the simplest nodes including resource constrained devices in need of better discovery and interaction with nearby devices. Equally important, it can be easily integrated over server-to-server links with any connected host platform through simple Web APIs, making it possible to extend features and benefits to managed networks and ecosystems.

Providers have the flexibility to deliver all these benefits through cloud-based solutions, or in hybrid modes that enable close coordination with offline locally hosted snapshots of the knowledge graph in a compressed and efficient format. Cloud connectivity enables near real-time access to an expanding knowledge graph of devices, and continuous updates on a much more flexible schedule.

In fact, a hybrid model with cloud-first configuration has become a common approach among deployments, including the intelligence to dynamically switch between cloud based APIs and offline fallback to provide the best user experience in a flexible integration model.

In a multi-service and multi-device age, cable/satellite operators and CE OEMs can enable second-screen applications, and cross device configuration synchronization to allow control through different interfaces. Coming from the opposite direction, mobile providers can expand the capabilities beyond the initial AV centric applications to **include capabilities for the smart home**.

6. New Experiences in Homes Powered by Device Knowledge Graphs

In tandem with the expanded modes of device discovery, an expanded range of metadata and processes to support new and better experiences targeted at the connected home are necessary. These capabilities power a range of new possibilities for service providers to offer new and personalized services that can be used for monetization through advanced advertising, e-commerce and marketing applications.

6.1. Device Discovery and Contextual Awareness in the Smart Home

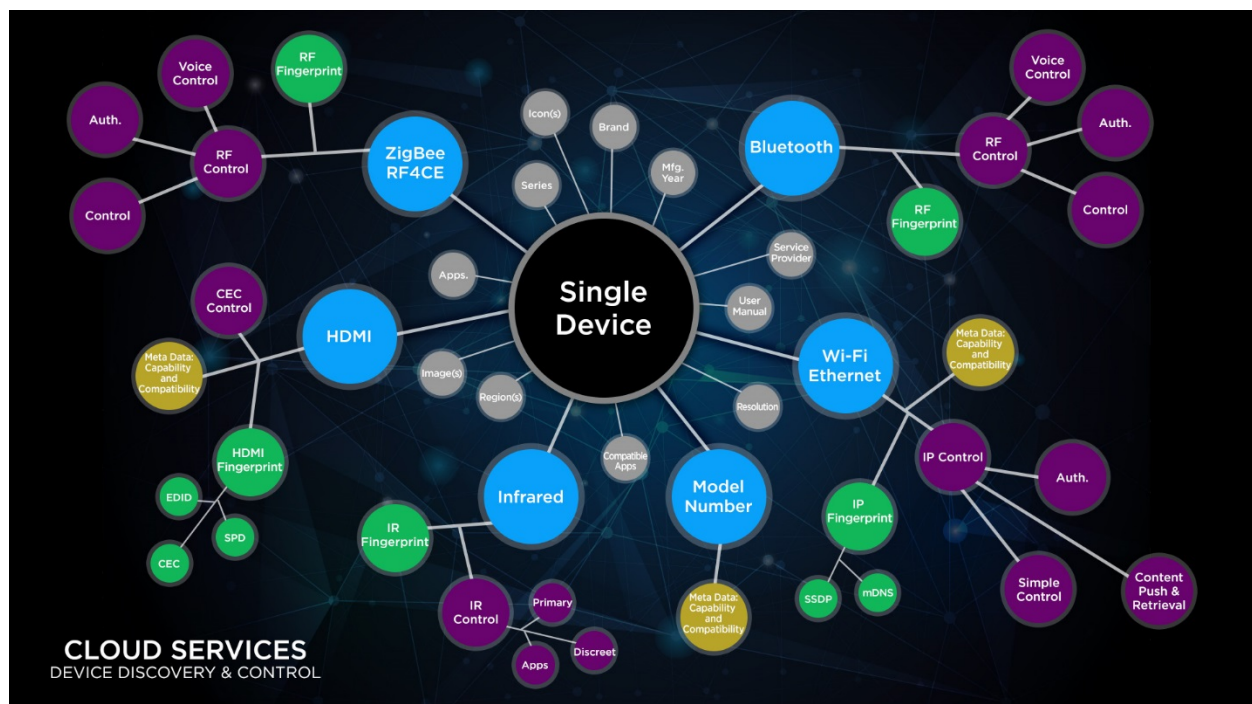


Figure 18 - Expanding Device Model for New Experiences

The **aggregation of additional metadata and device characteristics** are compiled through a combination of advanced data analytics and machine learning processes, combined with a rigorous validation and curation process to assure quality and reliability. The expanding range of metadata can be used by providers to further automate device discovery in the home and to help **improve user experience, or open the door to new business opportunities**.

Available devices, services and content within the home provide the necessary **contextual awareness for better and more personalized services**. In a simple approach, available devices and services are used as signals to trigger a specific persona, such as a household with a specific brand of game console, or subscription to a specific streaming service, or a household in need of a TV system upgrade due to the age of their display. In the world of content, implementations targeting some of these use cases utilize the traditional media recognition techniques which may be limiting and expensive to scale.

These capabilities, when applied to the smart speaker category, can materially improve the experience through better contextual awareness for voice assistants competing to be the primary interface to the

smart home, and a seamless integration with the most common and widely used devices in the household, including new and old entertainment systems and air conditioner units.

6.2. Content Consumption Across Devices and Services

First and foremost, unified discovery and control capabilities satisfy the need among consumers for a way to surface content and applications for easy consumption, regardless of how and where the content is served from.

Upon the completion of initial device discovery on local networks, a solution can apply whatever scripts and protocols are needed to scan and index the content and apps hosted by nearby devices, including any content or apps that might be currently in use. Examples of available insights include which channel is being watched on a cable/satellite service, which application an OTT user is watching or what tune is playing on Spotify. This data is applied in real time to enable aggregation and prioritization of content and applications for display in the provider's **user centric global history and content dashboard**.

There is no set template providers must adhere to in executing this capability. Instead, through the use of the advanced engines, providers have the maximum flexibility to apply metadata and tools to execute all the capabilities in accord with their strategies and user interface (UI) designs.

Critically, providers have the flexibility to provide an intelligent dashboard that adapts to what the end user cares about, dependent on the time of day, and based on their consumption across devices. This global usage history can be used to localize and personalize what appears in the user's UI by referencing the usage profiles captured while adhering to the most stringent privacy and security guidelines.

6.3. Advancing Personalization & Monetization

When the user data is aggregated across devices and services using a unified discovery and control solution, that data can be used to generate **personalized and contextual content recommendations**. A provider can utilize in-house developed or commercially-available recommendation engines to consume data insights to generate content suggestions based on user habits -- not only while consuming the provider's content, but also while interacting with other services.

A platform's support for personalized experiences sets in motion monetization opportunities as well, starting with advertising. Using the metadata aggregated, providers can improve the experience of their ads to be **tuned to specific user or household interests and needs, derived from both user activity as well as the installed base of devices in the home**. Nearby devices in a home can be a strong signal indicating preferences and needs. Such ads might be based on the content a user is drawn to, such as local concert ads for users who frequent a music app; or based on discovered devices such as ads promoting games developed for a specific user's game console.

With assistance from the Predictive Engine, the platform can help OEMs filter their messaging by alerting them to instances where households are using TVs that are at the end of their life spans and would be a strong candidate for upgrades. Similarly, the information identifying types of TV sets in use by any given household can be used by OEMs to avoid running ads promoting new TV models to people who already have them.

For service providers, the opportunities revolve around identification of user behavior that signals potential interest in a service they don't already have. For example, if cable/satellite offers 4K ultra-high-

definition (UHD) service, they will want to alert new buyers of 4K TV sets that the service is available. Or, knowing a user is primarily accessing Netflix for movies, the cable/satellite operator may want to promote its VOD service as a better option.

Conclusion

A simple and enjoyable experience in interacting with devices and content at home has become a key differentiator for service providers and manufacturers at a moment when gadget saturation has become a major drag on consumer satisfaction. **Technology can work behind the scenes to automate discovery of nearby devices and services**, and put content under the fingertips of the user.

Purist approaches have tried to redefine the household devices and protocols, but failed. We believe a flexible approach is capable of scaling across brands and ecosystems to improve daily life for the end users.

These platforms can enable a unified **one-touch** control experience over multiple household platforms, capable of automatic discovery and control over any device regardless of communication protocol. The platform must:

- Aggregate and present content and applications offered across devices in the home to power a unified and personalized dashboard;
- Extend universal control mechanisms to mobile devices and voice assistants;
- Facilitate device and application integration with third-party cloud-based services;
- Support personalization and monetization of the whole-home experience with advances in voice technology, data analytics, and other AI powered services.

It is important to acknowledge the opportunity to improve daily usage experiences for all users through a unified and consistent AI-powered interface, going across device and application boundaries. As a trusted assistant at home in charge of a user's daily journey in finding and consuming new content and services, AI-powered interfaces are now dangerously close to the needs of such an application, only lacking the proper contextual awareness and knowledge of nearby devices and content.

Abbreviations

API	application programming interfaces
AV	audio/visual
BLE	Bluetooth Low-Energy
CE	consumer electronics
HDMI	High-Definition Multimedia Interface
HDMI-CEC	High-Definition Multimedia Interface Consumer Electronics Control
IoT	Internet of Things
IP	Internet Protocol
OEM	Original Equipment Manufacturer
OTT	over-the-top
SDK	software developmental kits
TV	television
UHD	ultra-high-definition
UI	user interface
VOD	video on demand

Bibliography & References

- ⁱ Universal Electronics, [QuickSet Data Insights - U.S. TV Installed-Base](#), November 2016
- ⁱⁱ Parks Associates, [OTT Video Market Tracker](#), December 2016
- ⁱⁱⁱ Parks Associates, [23% of Millennials are OTT-only Broadband Households](#), June 2016
- ^{iv} Sandvine, [Global Internet Phenomena Report](#), August 2016
- ^v Broadband News, [Five Connected Media Devices per Home by 2019](#), September 2015
- ^{vi} Rapid TV News, [“OTT Streaming to Hit 100 MN Subs,”](#) April 2015
- ^{vii} Digital TV Research, [Global OTT TV & Video Forecast](#), June 2015
- ^{viii} Broadband TV News, [Viewers Want Simple Remote Controls](#), March 2016
- ^{ix} Wall Street Journal, [Samsung TV with UEL-Enabled Remote](#), April 2016

We Have Arrived. Our Light Bulbs Finally Have IP Addresses!

Approaches for Proactively Managing Customer Experience and Reducing OPEX in a Cable Operations Environment

An Operational Practice prepared for SCTE•ISBE by

Gary Cunha
Sr. Director, Product Management
ARRIS
900 Chelmsford St.
Lowell, MA 01851 USA
978-614-2900
gary.cunha@arris.com

Introduction

We have arrived. Our light bulbs finally have IP addresses!

Meanwhile, operators must hire buildings full of customer service representatives (CSRs) to answer questions about networks and devices over which they have little control. There's a better way to maintain and grow the Customer eXperience (CX) by leveraging data, technology, and methods that are available today.

Homes that previously hosted a single computer with high speed data service now have numerous Internet of Things (IoT) devices, complex Wi-Fi networks, and subscribers that insist it all work seamlessly. Business subscribers and home-based workers have even higher expectations. Call centers are strained trying to keep up with technology advances and are buried under a growing call volume.

In their *Customers 2020* research, a Walker study finds that CX is overtaking price and product as the key brand differentiator. In fact, over 85% of subscribers would be willing to pay more for a better customer experience [Campbell]. Poor experiences drive consumers to buy less and share their bad experiences more. According to the White House office of consumer affairs, news of a bad customer experience reaches more than twice as many ears as praise for a good service.

For call centers, the days of simple workflow, scripted line-of-questioning tools have passed. Energetic cable operators have attempted to take a next step by aggregating piles of data into one screen, only to find they still are not getting ahead. CSRs are not analysts: there's not enough time on calls and not enough training available for operators to expect them to arrive at a clean situational view of a subscriber's service challenges. While some tech savvy cable operators are performing some form of periodic subscriber churn analysis or leveraging basic data aggregation to guide call center workflows, a more comprehensive approach is needed.

Additionally, with CX impacting all areas of the cable organization, advanced strategies must be up-leveled and made a top priority across the organization, combining resources to establish a systematic and aggressive CX plan. Business survival is at stake, as a growing number of subscribers have alternative service options available and are being marketed to for additional revenue.

This paper discusses CX drivers, reviews metrics, and lessons learned from several CX initiatives. It also provides an approach for addressing CX challenges. A cross-group view is used here as the basis for discussion – in an effort to give the reader a larger vision of the CX opportunity.

The Changing Service Environment

1. Dramatic Changes in the Service Landscape

This section outlines several of the complexities operators must deal with in their pursuit of the happy customer.




Figure 2 – Total “Smart” Devices in Households (Estimated)

CPE: To satisfy connectivity needs for these devices, CPE (including modems and home gateways) is playing a more complex role in the home. Devices are often shipped with dual-band Wi-Fi capabilities, while also managing all IP traffic, high speed data functions, voice service, security, parental controls, and acting as a local router. In the role of a residential gateway, subscribers expect it to handle any flavor of consumer device in the home. Subscribers attempt to be their own IT administrators, dealing with the numerous subscriber-facing configuration options – which allow a complex set of combinations.

1.3. Medical Home-Care

Broadband service has become mission critical for a growing population of subscribers using network-enabled medical devices from within the home. Berg Insight forecasts that the number of patients using connected home medical monitoring devices will grow at a compound annual growth rate (CAGR) of 44.4% to reach 19.1 million in 2018.

In this scenario, CX expectations could be magnified by potential business partnerships between service providers and health organizations, where premiums are paid by health organizations for a higher service level. Here, it's possible auditable records would be required along with faster response times during service impairments. Even after subsidizing connectivity, cost savings in health care costs could be recognized [AHA]




Figure 3 – Connected Home Medical Monitoring Devices, (World 2013-2018)

1.4. Capacity and Architecture Challenges

In addition to the expanding home network, operators battle numerous challenges as they expand and improve service.

Bandwidth: Consumers' insatiable appetite for bandwidth is triggering massive shifts in the Hybrid Fiber Coax (HFC) network architecture. Global consumer Internet traffic has seen a 5x increase in the last 5 years and could increase 3-4x by 2020 [Davidson]. Driven by this, complexity is being magnified by new DOCSIS® standards, updated hardware for headend, plant, and home, all with new management interfaces. This creates a balancing act as operators work to align the timing of their network expansion to meet consumer demands.

Consumers are impacted by the legacy HFC plant, neighbors consuming bandwidth on shared channels, network interference resulting in Codeword Error Ratio (CER) and lower throughput, and home network and device issues.




Figure 4 – Data Transfer Projections to 2020 (source: Bell Labs)

Deployment of higher-level QAM for increased capacity demands tighter Radio Frequency (RF) performance tolerances, further increasing the demand on operators to proactively manage the plant. Forward Error Correction (FEC) and Adaptive Equalization (EQ) techniques compensate for RF/interference issues to a point, yet can't protect subscribers from inevitable service impairments. Growing trends toward Distributed Access Architectures (DAA) and migration of core functions from the headend to the outside plant could exacerbate challenges in aerial networks as weather and other environmental factors wreak havoc.

1.5. A Typical Day: One Operator's Struggle

There are so many things that can interrupt service and whole operator teams are created around manually triaging and resolving service issues.

The day-to-day manifestation of these operational challenges was summed up by an operations director of a large cable operator earlier this year:

"We're still reactive; When a spike in calls hits the call center, we really have no idea what caused it. Another analysis begins."

This director made it clear that it goes far beyond issues in the plant. He continued,

"Despite our efforts gathering data and making it visible to engineers and CSRs, most correlations are done manually. We can't tell if this event was caused by a recent CPE firmware update, a model of consumer device that just received an update, upstream noise in the plant affecting subscribers all over the node, or even some intermittent noise from an outside source in the area. Those intermittent ones are the toughest."




Figure 5 – Ops Director’s Quote on Struggles Triaging Call Center Spikes

2. Taking the Next Step

2.1. Who’s Got The Ball?

In the previous section, the director’s statements identify a lack of visibility across realms and a gap in understanding and sharing data. In organizations across the globe, we see endless amounts of data being gathered, maybe post-processing for post-mortem review. However, insights and correlations are not keeping pace with expectations, they are not leading to actionable, impactful results.

While the scenario above begins in the call center, the larger issue really is organization-wide. Managing CX and related OPEX-heavy activities starts with the very first truck-roll.

A simple sports analogy helps explain the group dynamics involved: In soccer, whose job is it to prevent a goal from being scored by the opposing team? It’s not just the goalie/keeper. Every person on the field has a role. Translation: whose job is it to address a service-affecting event that drives calls to the call center? The whole organization. The goalie (call-center) should be treated as a last resort in preventing loss.




Figure 6 – Entire Team Must Protect the Goal (a.k.a. Call Center)

In business terms, the deeper into the organization a service event reaches before resolution, the more expensive it is and the more impacted/upset the subscriber (lower CX). This is why programs aimed at CX and call center success must be a cross-group effort. After reading this, an expected reaction is “Well,

that's obvious...” Maybe, but here’s the reality in most organizations, as observed during numerous engagements with operators around the globe over many years:

Table 1 – Reality of Inter-team Dynamics & CX Efforts

Examples of Operations Not Targeting Overall Company Targets
Directors protect their own OPEX/P&L hoping for another team to pick up the expense of intelligence projects
Revenue-generating efforts are funded heavily while OPEX-driven programs are constrained, despite the expectation of positive ROI on OPEX-saving investments
Internal teams point to each other as the culprit in service outages
Technicians are lightly armed with tools needed to properly perform Install/Repair (I&R)
Technicians are often pressured/incented to get on to the next job
Network Operations Center (NOC) analysts are not given the time to appropriately study situations, capture metrics, create an audit trail, and root cause; losing a golden opportunity for future pattern recognition. Trouble tickets often go unclosed or not updated with root cause and summary info
CSRs find ways around any Method & Procedure (M&P) to (1) appease subscribers by rolling trucks and (2) improve their metrics by shortening call times and getting on to the next call despite the best intentions of managers
Call centers that continue to allow non-standard experiences – where CSRs have too many views, too much data, and their “favorite” screen/tool, thus reducing impact of M&Ps and training
Lack of cross-group feedback cycles – preventing proper knowledge transfer for top service impacting use-cases that would prevent issues in the future
Lack of sophisticated integration between tools/groups, preventing true interpretation of the customer experience

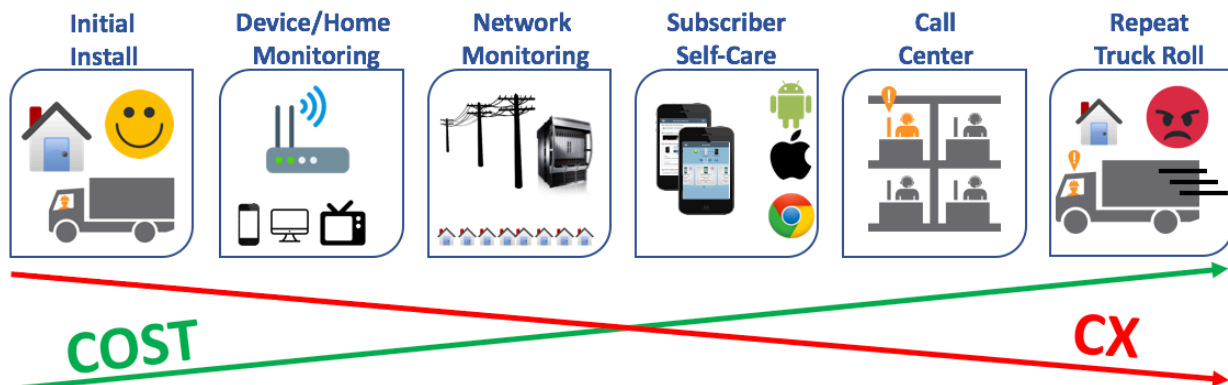


Figure 7 – Cost, CX, and the Impairment Management Lifecycle

What are some ways that this can start to be addressed? Improvement opportunities for these areas are reviewed in sections below. More important than emulating one specific program is establishing and maintaining an organizational mindset of incremental improvement through common goals, automation, technology, and measurement. Projects tagged with saving the world are doomed to fail – instead, small wins justify continued resources and maintain team motivation.

2.2. Organizational Strategy

CX efforts relate to groups across the organization and are made stronger when coordinated. To be successful, progressive companies will need to view customer satisfaction as an executive leadership role and that the role of “Chief Customer Champion” will become more commonplace. [Walker]

Here are two recent success stories from the telecommunications industry:

Rogers Communications: Through intense focus on the customer, Rogers has reduced customer complaints by 50% between 2014 and 2016 [Forbes, Davis]. Deepak Khandelwal, Rogers Chief Customer Officer, crisscrosses Canada to communicate his frontline perspective directly with call center staff and field techs.

A large tier-1 telco in North America: This organization established an executive-level role focused on improving the customer experience. Operations leads and consultants interviewed for this paper stated that an immediate effect was realized as cross-group roadblocks were removed and efforts aligned, resulting in up to 5% reduction in OPEX on key programs.

Why does this work?

- Having a common leader with unified goals helps reduce the urge for individual teams to prioritize protection of their own budget and P&L over the larger company goal
- Resources can be shared and project redundancy reduced, minimizing a dynamic where one team is “taking the hit” for another team’s benefit
- More transparency and measurement drives data-driven identification of trouble spots and stops teams from blaming each other for service issues

- A thoughtful solution architecture can be established and incrementally improved, allowing foundational technology and services to be shared. An example of this might be a common telemetry collection and storage program that feeds analytics and IT efforts on multiple programs

While all this sounds like great news, only 39% of companies have at least 1 senior-level executive leading the CX charge [Forbes]. If not an executive position, operators should consider establishing a strong leader who is empowered to unify CX efforts, establish clear goals, learn from industry lessons, make mistakes, and, more importantly, commit to a cadence of continued incremental improvement over time.

The specific programs and proposals captured in this paper should be viewed through the eyes of a Customer Experience leader. Think bigger than your realm.

Real Examples: Proactive Issue & Call Prevention

A key goal of this paper is to discuss how a more integrated and auto-interpreted Situation Analysis helps transform operations success. Those topics will be discussed and it is critically important to include strategies for proactively preventing the call in the first place, thus lowering OPEX and increasing CX.

Figure 8 illustrates several layers of operational protection that should be in place; groups driving these areas should be chartered with aggressive and insatiable call/issue prevention. This layered graphic will act as a guide through preventative strategies discussed below.




Figure 8 – Issue & Call Prevention is the Top Priority

3. Field Technician: The Front Lines

The world of the field technician is a world of competing priorities: They must complete high quality work and provide a focused personalized experience to the subscriber – while at the same time work as fast as possible and get to the next job. Operations managers reading this paper have no doubt seen technicians do whatever they need to do to get out of the subscriber home and on to the next job, even if it means someone will need to fix service issues after the fact. There are certainly technician-based

challenges, but technicians are not to blame. Most of the larger opportunities for improvement are based on repeatable and systematic operations programs created and driven by field ops management.

3.1. Case Study: Preventing Wasted Trouble Calls With IVR

Technician productivity can be increased by using automation to maximize the prevention of wasted I&R trouble calls. In one region of a sizeable cable operator, an Interactive Voice Response (IVR) system was used to automatically call ahead to subscribers with a trouble call scheduled for the following day. (Specific numbers are purposely withheld to help maintain the anonymity of the operator). Subscribers answering the call were asked if they still had an issue, and were given the option to cancel or reschedule the work order by pressing numbers on their telephone. Of those that answered, an average of 10% cancelled or rescheduled each day, leading to a daily prevention of truck rolls that would have resulted in *Customer Not Home* or *Reschedule Request* scenarios. This one simple program saved the operator over \$2 million annually.

3.2. Case Study: Field Work Certification Program

A business process analysis was performed on three cable markets at a mid-tier cable operator suffering from abnormally high repeat truck roll volume – driven by a large number of customer calls and complaints. One of the early findings was that most of the customer calls occurred within one or two days after completion of a residential trouble call.

It was then determined that technicians, often contractors, were financially incentivized to maximize the number of jobs completed per day. Without proper quality/metrics programs in place and limited management visibility, technicians would rush through I&R work and on to the next job, often leaving the subscriber with impaired service (low SNR, high CER) or resolving one service while negatively affecting another.

To address these challenges, a field work certification program was created whereby technicians were required to use a mobile platform to certify the work they had just completed and archive the results for future reference. The interface was designed to provide raw data for those that could understand it, and it also allowed for a simple Pass/Fail indicator on the most important metrics. This simplified view took the guess work out of completion results and allowed technicians to move quickly.

Records could be viewed later and aggregated into management metrics reports. Additionally, work quality and repeat rates were added as key factors for compensation packages.

Several dramatic results occurred, as outlined in the table below.

Table 2 – Key Findings of an Install/Repair Certification Program

#	Key Finding	Impact to Business
1	Techs more effective	Having better tools and recognizing the focus on increased transparency, technicians were better prepared and more thoughtful about on-site work completion / quality. This was also helped by the normalized mobile interface, giving the

#	Key Finding	Impact to Business
		technician a common look and feel no matter what vendor/model of cable modem was involved.
2	Records helped future techs	Future technicians were better prepared for trouble calls at the same location since they could (a) look back at previous work certificates and evaluate the state of service during previous work orders, (b) read technician comments, and (c) see how the environment had changed.
3	84% Fewer Calls to Call Center	<p>A dramatic reduction in repeat trouble calls to the call center was recognized. Since technicians were now able to clearly see lingering issues in the home, they were guided to resolve them before closing the work order.</p> <p>A key part of this success was that field techs were given the extra time required to properly validate & cleanse the home service during the same truck roll.</p>
4	82% fewer repeat truck rolls	<p>A dramatic reduction in repeat truck rolls was recognized. This was primarily driven by first-visit resolution on previous truck rolls.</p> <p>To accomplish this, field managers had to improve handling of work orders in jeopardy (jobs scheduled later are at risk of starting late) when a tech had to stay longer at a job.</p>
5	Better management visibility	With the archived data, managers had better insights into which technicians were closing jobs without properly resolving key home issues, allowing them to better target training and performance programs. Key data was available to allow for correlations between device models, firmware, and RF performance.
6	CSRs viewed certificates in their M&P	Although not optimized for the call center, CSRs and Level 2/3 troubleshooting/support engineers used these historical certificates for improved visibility when handling customer complaints. They could see that the service was certified and working well at the time of work order completion.
7	False negative RMAs reduced	One way CSRs would appease customers was to have the subscriber's CPE replaced with a new device. While this may have helped and end the support call quickly, it did not resolve service issues and unnecessarily increased RMA load. Since there were fewer subscribers calling with issues, there was less opportunity for a CSR to send a replacement CPE.

It's recognized that these reductions in calls and truck rolls were dramatic; this organization had a long way to go. That said, even for well-trained and well-armed field teams, this type of a program can have a strong impact with a positive ROI.

For operators that have a certification program in place – or who don't want to impact the field technician's work load, variations on this program have been implemented across other global operators – with positive results. For example:

- 1) **Auto-Baselining:** One operator has launched a program to automatically create work certificates for CPEs deployed within the past two days. This is accomplished with a small amount of scripting and leveraged real-time web services typically available with most modern enterprise-class DOCSIS surveillance platforms. Through this effort, a service baseline exists for every device – which can be referenced during future troubleshooting events.
- 2) **Pre/Post Trouble Call Certificates:** Another operator also leverages validation/certificate web services by automatically creating them before and after each trouble call. How it worked: The night before a scheduled trouble call, a work certificate was created for any DOCSIS device with a scheduled trouble call the following day. Following closure of the work order, another work certificate is created, providing visibility into the impact of the work.

Notice that in the two cases above, the technician did not need to do extra work onsite other than perform their regular field I&R duties well. One way to get started is to focus programs like these on VIP/Business subscribers or those purchasing the premier product/speed package.




Figure 91 – Case Study: Impact of a Birth Certificate Program

4. Proactive Network Management

Cable operators have historically had some sort of script to look for sudden drops in online modem counts. While basic tools like these are helpful, operations teams are unfortunately alerted reactively –

after customer impact has occurred. Today, enlightened cable operators are working to become more proactive, focusing more programs that attempt to prevent outages. Hard work by CableLabs has really expanded the discussion.

A great deal of investment is being made by leading cable technology vendors to deliver a fully proactive solution, with the goal of transforming the way operators maintain the plant and prioritize work. The new goal is to find and fix network impairments without subscribers ever knowing there was an issue.

4.1. Case Study: Impact of a Proactive Network Management Program

In a project with a large cable operator, a proactive network management program was launched. This particular deployment covered a region with over 2 million DOCSIS modems. While this operator is quite cable & tool savvy, they had the following goals:

- 1) Help them evolve their maintenance program to be more proactive.
- 2) Leverage a system that automated much of the work done by their network analysts.
- 3) Correlate the tsunami of network symptoms & events into the root cause, squelching the volume of data.
- 4) Lessen the need for “eyes on glass”, driving automation and notification into daily processes.
- 5) Assist them in determining the probable location of impairments.
- 6) Detection of granular events, allowing pin-pointed truck rolls. This minimizes service impact since techs now knock fewer people offline – while disconnecting network elements during the maintenance window. (Think: work on a tap, not a node.)

Table 3 – Key Findings of a PNM Program

#	Key Finding	Impact to Business
1	Shift from Reactive to Proactive	Organization was able to shift daily processes and workflow from Reactive to Proactive. To find targeted network areas to work, they relied less on surprise call spikes from the call center and more on automated fingerprinting of RF impairments.
2	More automation in NOC meant less time searching in the field	<p>Proactive Network Maintenance (PNM) analytics & tools cut through mountains of data to help determine impairment location, while preparing the majority of the data for the work order. Analysts had advanced tools to fill gaps where necessary.</p> <p>Maintenance work allocation:</p> <ul style="list-style-type: none"> • FROM: 20% in NOC, 80% in Field • TO: 80% in NOC, 20% in Field <p>While there will always be a need for expert line techs in the field, this shift allowed field teams to focus on execution and smarter maintenance strategy rather than roaming searching for the source of impairments.</p>

#	Key Finding	Impact to Business
3	Planning Cycles	Shifted from weekly maintenance planning schedule to a daily standup. The team was able to prioritize quickly and address service interrupts more readily.
4	Home wiring	Previously, truck rolls were sometimes wasted as CSRs and analysts struggled to separate in-home issues from neighborhood impairments. Analytics produced a “home score”, which clearly separated in-home wiring issues from neighborhood plant issues. These homes were also seen as a potential source of noise ingress. Operations processes were modified to attach tickets for individual homes to line tickets, so several issues could be resolved with one truck roll.

4.2. Global Trend: Separate Proactive Team

While there is always a need for a quick-reaction team to handle network outages, there is a growing global trend among enlightened cable operators to establish a *separate proactive plant maintenance* team.

This team is tasked with more strategic efforts: finding large-impact maintenance opportunities; identifying performance characteristics among plant hardware vendors/models that impact multiple regions; improving proactive impairment detection strategies and tools; and working on the big challenges that frequently-interrupted outage teams would not have time to focus on.

At a few operator sites, separate proactive NOCs have been built – allowing customized war-room views to be displayed on the wall, different views into RF telemetry to be evaluated, and maintenance truck rolls to be dispatched. The prevention of outage-based distractions allowed this team to look further ahead. The results are clear: Dramatic reduction in trouble calls and truck rolls to the plant.

Not all operators are ready for this step –the important point here is to recognize that focused time and resources are needed to transform cable operations from reactive to proactive. As previously mentioned, steady incremental progress is more important than attempting a Hail Mary project. As a way to get started, a few cable operators have established a tiger-team to brainstorm proactive solutions, evaluate tools, and, to be actionable, rotate line technicians in/out of the tiger team periodically to give them and the team a fresh perspective.

5. Monitoring the Home Network

A growing number of operators are choosing to aggressively deploy residential gateways or Wi-Fi-enabled modems. Working with these operators day-to-day, it’s clear the demand to meet time-to-market milestones is intense. With few exceptions, early Wi-Fi launch strategies follow an approach of *get-the-service-out-there-and-manage-it-later*. This has magnified the complexity of operator-customer expectations greatly. As reviewed in earlier sections, subscribers struggle with seemingly simple network management tasks and are flooding call centers.

In mid-2015, one large cable operator with an extremely aggressive Wi-Fi market launch experienced the operational impacts shown in Figure 20. Such support loads are not sustainable and will demand an equally aggressive monitoring strategy.



Figure 20 – Early Operations Impact of an Aggressive Wi-Fi Program for a Large Operator

Globally, it's early days for automated Wi-Fi analytics and proactive management, resulting in a desperate need for structured programs to (a) expand visibility into the home network and (b) leverage automation to help reduce manual surveillance and support efforts. Some ideas on how this can be impactful appear below, and later sections will review how intelligent insights could be used in the call center.

5.1. Remote Home Surveillance & Wi-Fi QoE

As was done for cable modems during early DOCSIS days, operators are now expanding telemetry collection, on-demand tools, and data storage strategies to include Wi-Fi and home network data. This is the starting point for IT teams, so that infrastructure and tools can be built.

The volume of surveillance data required to represent today's home network has grown – leveraging protocols like TR-069, TR-181, and XMPP for visibility to attached consumer devices. Various approaches are being discussed for how to determine problems in the home network and there are lots of places where problems can occur.

Providing algorithms and specific object identifiers is not the goal of this paper – however, the general concepts should be clear. If we start by attempting to characterize a single home, one reasonable approach is to look for issues that affect all devices in the home –and then get more specific from there. The general flow is described here with a more specific example in Figure 31.

- (1) Identify critical issues that impair home Wi-Fi service (all clients likely impacted).
- (2) Identify pocket issues that could impair normal operation for multiple devices (users will likely have a mixed quality experience).
- (3) Examine individual devices to determine impact.

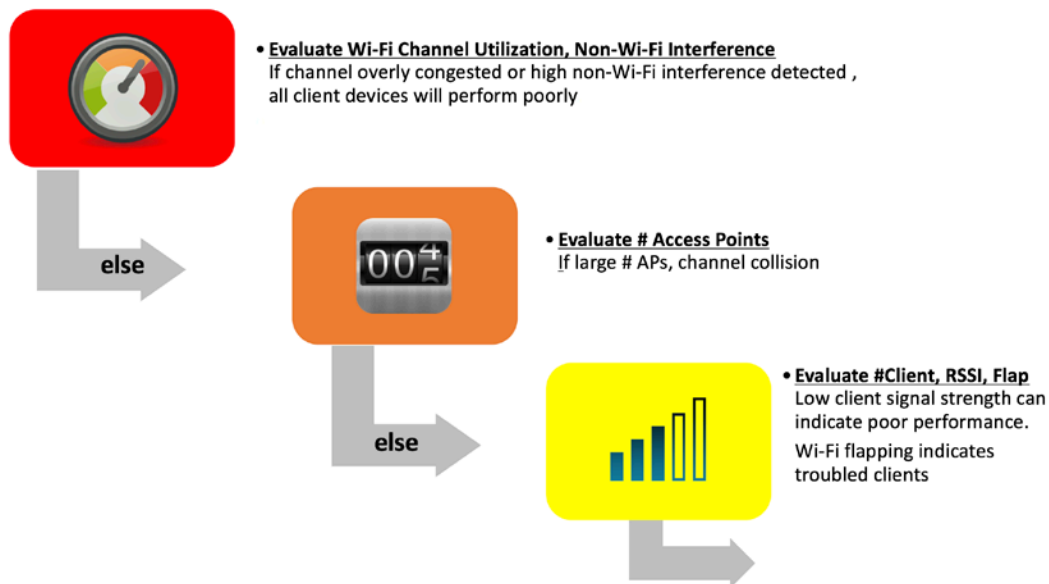


Figure 31 – Early Thoughts on Home Wi-Fi QoE

Additional examples of remote automations to be considered for healthy Wi-Fi service:

- 1) **Security**: Automated identification of gateway security issues creating service risk, and automated workflows triggered to resolve them.
- 2) **Old Devices**: Identifying older 802.11 devices which, on certain gateway models, slows service down for all attached devices.
- 3) **Low Signal**: Sustained low signal strength (RSSI) for attached devices (careful to make use of “last seen” time to filter out offline devices that some gateways may remember).
- 4) **Flapping Devices**: Flag devices continually connecting/disconnecting from the gateway.
- 5) **Too Many Devices**: Look for very large number of devices attached to the gateway.
- 6) **Interference**: Determine regional interference issues impacting Wi-Fi (correlate to geographic areas).
- 7) **Throughput**: Heavy channel load or link bandwidth may indicate over-use of limited bandwidth resources.
- 8) **Wi-Fi “Birth” Certificates**: Manual or scripted service certificates for connected devices, Signal Strength, AP configurations, and more.
- 9) **Service Package**: Correlating gateway configuration against the billing system to determine if subscriber is getting less/more features than they are paying for.
- 10) **Temporal Issues**: Track patterns over time, find repeated cycles of Good / Bad service.

These are just a few examples of opportunities for automated Wi-Fi surveillance. The output of these types of programs is critical for being able to manage home service profitably.

6. Subscriber Self-Care

6.1. Subscribers are Ready

Tools allowing subscribers to help themselves are more common today than ever before and provide a huge opportunity to prevent calls to the call center. What used to be as simple as a web page for Frequently Asked Questions (FAQ) is often now an interactive system or mobile app. In a survey published in 2013, 70% of consumers expect companies to offer some sort of self-service application [Van Belleghem].




Figure 42 – Consumers Are Ready for Self Care

The self-help trend is seen everywhere in society: people use ATMs and mobile apps instead of bank tellers, choose the self-checkout isle when buying groceries, pump their own gas, and order coffee via their phone. Nearly 3 out of 4 consumers prefer to solve customer service issues on their own [Aspect], and Gartner predicts that by 2020, customers will manage 85% of their relationship with the enterprise without interacting with a human. In fact, 65% of consumers surveyed in a 2015 study for Aspect said they felt good about themselves and the company they were doing business with when they could resolve a problem without talking to customer service.

6.2. Wi-Fi Calls Dominate the Call Center

The growing self-help trend is very good news for operators. Working with a large cable operator in their call center, it was observed that 30% of technical calls (e.g. non-billing) were Wi-Fi related. Of those, over 30% were resolved through subscriber education, with customer's often asking "*What is my Wi-Fi Password?*", and "*What is my Wi-Fi SSID?*" Another 20% were resolved by rebooting a CPE/gateway.

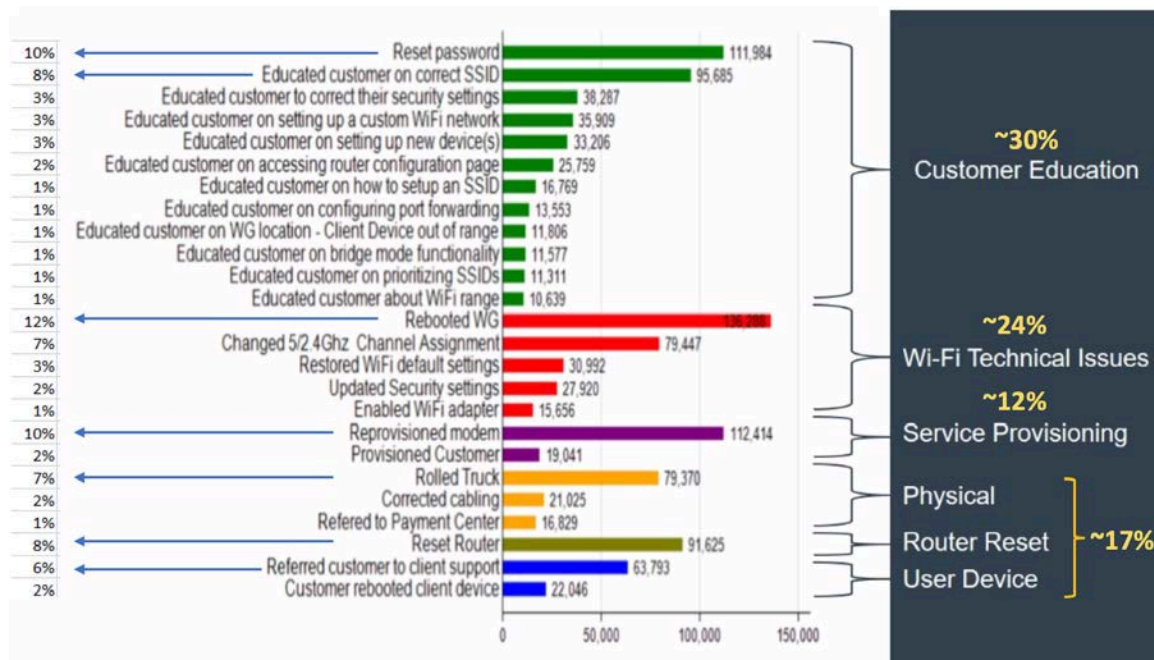


Figure 53 – Wi-Fi Resolution Codes in the Call Center

These use cases are perfectly suited for a subscriber self-help application. Part of the value comes from the fact that a self-help app user interface offers a normalized experience. It looks virtually the same no matter what gateway vendor the operator chooses to deploy.

ARRIS is deploying its Wi-Fi self-help portal app with a tier-1 telco in Latin America and expectations for call deflection are big. Rebranded to match the look and feel of the provider, subscribers can easily perform basic home network management functions from their mobile phone – like those discussed above as top call reasons including: check gateway status, view Wi-Fi password, enable guest Wi-Fi, view the home network, review top FAQ articles, and much more. The operator has a mixed deployment of multiple Wi-Fi gateway vendors – yet the subscriber self-care app looks the same.



Figure 64 – Intuitive Wi-Fi Self-Service Tools

6.3. A World Without Self-Care Apps

Let's take a moment and review what a subscriber must do to configure their Wi-Fi gateway without an integrated app:

- 1) Get on a PC physically wired to their router.
- 2) Use Google to determine the default non-routable IP address of their gateway model (and understand what an IP address is).
- 3) Open a browser on the same network and type in the IP address.
- 4) Subscriber must remember that the gateway admin UI user/password is not the same as their Wi-Fi SSID password, their email password, nor their computer/PC password.
- 5) Use Google again to determine the default admin user/password for their specific gateway vendor.
- 6) If they get this far, they need to then spend an hour with Google deciphering Wi-Fi security settings and why their child's Xbox won't connect to the Internet.

Self-care apps simplify the experience for subscribers managing their home network and will have a significant impact on cable operations and customer satisfaction. Later in the paper, read about how the mobile app provides the perfect platform for operators to deliver proactive service alerts.

A View from the Call Center

At this point in the paper, we've reviewed various layers/programs involved in minimizing issues that drive subscriber calls, reviewed some new ideas about automating issue detection, and offered a few alternative methods for handling issues when they are not avoidable. When all that is still not enough, impacted subscribers reach out to the operator. How do we handle this? Why are we still struggling to keep pace with call volumes and customer satisfaction? It's now time to talk about the call center.

7. Context: The Call Center Universe

To help readers understand call center solutions and related challenges, it will help to review a bit about the call center environment.

Subscribers: A common theme heard throughout interviews with call center leaders was, by the time a subscriber is calling the call center, they are looking to connect with a human. One top complaint consistently heard from callers is, "I hate that IVR system", after being forced to navigate the IVR menu gauntlet – or repeat to a human information already entered. In fact, 89% of customers get frustrated when they need to repeat information – often to multiple representatives as they get passed around [Accenture].

One example of why it's so critical to properly receive, route, and resolve customer calls is shown in a study at a medium-sized Latin American MSO. It was determined that 13% of their customers cancel within the first year. Of those, 32% cancelled within the first 45 days. Their data showed a clear correlation indicating that customers who cancelled early also called the call center multiple times. This provides a special opportunity to intercept and provide special care since it's cheaper to maintain an existing customer than acquire a new one.




Figure 7 – Cancellations Happen Early On, Correlate to Calls

Call Center Technology: Call centers typically use a workflow tool to guide CSRs through a call flow. This tool prompts the CSR to ask questions, enter subscriber responses, and be guided on to next steps through the tool’s branching logic. This drives a slow step-by-step verbal exchange, relying on the subscriber for device and performance information. Modern versions of these tools may be integrated into a basic set of network data to show current & historical performance details, which can be useful for basic problems. In an *Omni-Channel* environment, subscriber can connect with support resources via email, chat, social media, knowledge base, etc.

The goal is to architect these channels to leverage a common set of underlying services for a consistent, positive support experience.

CSR: CSRs are in a challenging environment. Callers are upset. CSRs are often lightly trained, have a high turnover rate, have little domain expertise, and are expected to quickly move through calls toward a resolution. CSRs may feel pressure to complete the call or appease the subscriber. Multiple call center managers reported that, “If a CSR wants to roll a truck, they know how to make it happen.” It’s not realistic or cost-effective to treat CSRs as an army of service analysts, applying their tribal knowledge on every call.

7.1. What’s the Problem?

The core problem with all this is as follows:

1. Automated surveillance and analytics programs are not delivering enough actionable insights on service health, HFC plant, home network, devices, impairments, or the subscribers themselves.
2. Correlation of information cross-realm is not happening, preventing larger insights from emerging.

3. Today's information is not presented in a way that can be quickly and clearly understood, pushing much of the work onto the CSR, NOC engineer, or data consumer. The resulting information-overload causes people to stop in their tracks and ask, "Where do I start? What's going on?"

Scenarios:

- Is it useful to have the CSR shown a page full of numbers if half of them are within limits – and *none of them is the real issue*?
- Is it better to show the CSR a page full of colorized icons instead of raw data – and let them decide what to do? ...or is that on-the-fly analysis too error prone and tough to make consistent?
- Why would we look at live subscriber SNMP data and claim that "everything's fine" when the data is available to show there is a real intermittent QoE issue every day at 10 am to 12 pm daily?
- Why would we allow trucks to be repeatedly sent to homes when data is available showing us their poor QoE is due to heavy network utilization on their stream channel, or a nearby plant impairment?
- Is that group of modems offline due to an outside plant issue or an IP Scope issue in the back office?
- If we can see that the house Wi-Fi gateway has an active iPad v1 running 802.11b and is slowing down the Wi-Fi for everyone in the house, can that be shown in plain English on the CSR screen?
- For the subscriber complaining of home network issues, should CSR tools automatically show what's happening elsewhere in the house and neighborhood? While we have them on the phone with a high speed data question, should we also take an extra minute and address the other two issues we see on their home network? These insights can be extended to views given to NOC engineers.

This is not an exhaustive list of examples – and probably not even the best ones, yet the point is clear. We're suffering from data overload, and punchline "underload" – and cable operators are paying the price. What is frustrating for operations teams is that many scenarios are not hard to model – but there hasn't been enough focus on modeling them in a systematic and scalable way and presenting them in a straightforward cohesive view.

Intermittent Issues: Conversely, there are also issues that are too challenging for quick casual human inspection: intermittent / flapping issues. There's just too much data and too many devices, which is a perfect opportunity for automation. In these scenarios, service for a subscriber may fluctuate between good/bad intermittently and may follow a clear time-based pattern. These are especially frustrating for everyone because they are hard to understand/explain and, with today's lightly integrated CSR views, subscribers often get the CSR response "things look OK right now." After 4 calls in 2 weeks, the subscriber is ready to cancel service. Now imagine that at a larger scale where area interference impacts 25 subscribers, the source of which is neither the home nor physical plant impairment. This story summarizes the challenge:

An operations VP told the story of a ham radio operator with powerful transmission capability, severely impacting certain RF frequencies. His usage pattern was fairly regular. The situation triggered numerous customer calls & truck rolls – but those impacted did not always call the call center, and not all at the same time. It took careful manual evaluation of call records and network data to narrow it down to a target cause/area. During visual inspection of the neighborhood, the large antenna was seen and the situation was addressed.

Cable Operators are Trying: Many operators are trying to build their own CSR tool to expedite the call experience. Lots of great work has been done, but in the last 18 months, at least 8 different operators have said they aren't seeing the impact they expected. Often, their efforts were focused on creating a consolidated view, aggregating lots of data into a few screens. Some operators created a peer tool to be used to complement their scripted workflow product, requiring the CSR to swivel-chair between several windows. In most cases, a lot of good data was provided – but still required the CSR to interpret the overall situation. Almost none of the operators had made any attempt to capture a thoughtful profile of the subscriber, showing if they've called for the same issue, showing other issues occurring in the home, special call handling directives, etc.

8. An Integrated Approach

A big challenge is that the insights needed by a CSR can't always be prepared just-in-time after a subscriber calls in. Such an approach lengthens calls and puts unrealistic and unnecessarily heavy performance burdens on IT systems. The content below attempts to paint a picture of a more intuitive and integrated system. The point of this is not the individual scenario/logic shown but that the visibility and flow is simplified to give contextual insights and actionable next steps to the CSR.

8.1. Imagine a Call Center Where...

- **Before a subscriber calls** in, systems in the background are constantly tracking anomalies, generating insights, tracking stats on calls per account per type, prepping key insights for use by other systems – not just exposing data
- **Caller ID** is used to detect that the subscriber is calling from home and automatically populates the CSR view with Account and Situational data
- **Interpreted Situation Summary**: Rather than see a page of data, the CSR is immediately shown a brief, intuitive situation summary, with the most important data first – including selections from the following sections:
- **Subscriber Experience Profile**: A top element of this situation report is a caller profile – meant to inform the CSR of the subscriber's engagement & experience.
 - **QoE Metrics**: An ongoing home QoE score is maintained - allowing the CSR to see a summary metric of subscriber experience for the past 4 hours, 1 day, 7 days, and last month. If necessary, the CSR can drill down to see individual QoE scores for each service: Internet; Voice; Video; Wi-Fi; and Security.
 - **Call Pattern**: A simple view shows how many times this customer has called in the last 3 weeks and whether those calls were for the same event. The CSR can acknowledge an ongoing issue and quickly get up to speed. The subscriber is happier, because they don't feel they are starting from square one on every call.
 - **Truck Rolls to Home**: Similar to above, the system shows highlights of recent trouble calls with techs visiting the home, including if one is scheduled soon.
 - In one scenario, it's clear the subscriber is a new customer and has called 3 times for Wi-Fi. The CSR is now empowered to stay on the phone as long as necessary to identify and resolve issues. Although more expensive, it's cheaper than fielding multiple calls, negative word of mouth, and potentially a subscriber cancellation.
- **Punchlines Over Data**: The situation summary also identifies likely service problem/causes, with an attempt made to order them by likelihood of impact and prioritized to match the call reason when necessary. (Wi-Fi, Video, Voice, or Internet)

- Rules of engagement: no raw data shown, colored/iconized per-KPI summary available as a drill down if user is authorized.
 - The system is designed to show only the most likely scenarios (constantly refined through disciplined & vigilant feedback and modeling by the call center leadership, working with the NOC and Field Ops).
 - For slightly better trained CSRs, permissions would allow drill down past summary info into carefully organized information. That said, even the data is curated to meet specific use cases. A design gate-keeper protects this UI real estate. For example, rather than show raw graphs of historical CER and SNR data, drill-down views might show only Green/Yellow/Red indicators for last hour/day/week. They may also be relabeled “Data Quality” and “Signal Quality” to help them be understood. (Of course, a power user’s view would show all data.)
- **Issues Outside the Home:** Often, subscriber service issues are caused by problems outside the home. Here are a few scenarios:
 - *Active neighborhood plant issues* are shown clearly at the top of the page. CSR explains that there is an issue impacting the area, work is in progress, ETA to completion is 2 hours, and offers to auto-alert the subscriber when the event is complete (SMS, email, or IVR call – no humans). Auto-alert setup occurs with one click. In such a scenario, a trouble call truck roll is not allowed.
 - *High Utilization in Neighborhood:* More subtle scenarios are also modeled: CSR is directly informed by the system that there is dramatic and sustained high bandwidth consumption on the same channels – leading to a degraded experience. Truck rolls to the home are not going to help, neither will a modem reset, nor drop-shipping a new device. In a perfect world, engineering has already been alerted and a ticket would already exist to address this capacity issue – with a targeted conservative completion time.
 - Power supply issues impacting the neighborhood would also land here.
 - CSR asks subscriber if they would like to be alerted to future large outages impacting them. SMS/Text is an option, yet preferred mechanism is alerts via the mobile app. This allows the subscriber to control the volume of alerts.
 - Actions are available allowing the CSR to “Refresh Ticket Status” which does not display technical data, just updates.
 - **Issues Not Outside the Home:** When the system doesn’t find outside issues causing problems, a prioritized list of home issues would be shown, in simple terms.
 - The system knows which are most impactful and alerts those in red
 - The subscriber’s CPE device configuration and service class are compared to billing information to make sure customer is getting what they pay for
 - System checks that data has actually been flowing upstream and downstream, data is useful when discussing outage or upsell scenarios. Verifies that all bonded channels are working and no “partial service” situation exists
 - Wi-Fi and Home Network risks are identified. The goal is to not present a view of every consumer device and status – but have anomalous findings indicated. Many Wi-Fi QoE scenarios were discussed in earlier sections but here, the CSR sees the concise lists of risks. Priority is given to service impacting items, followed by risk items that may prove problematic in the future – such as security issues, old consumer Wi-Fi devices, not using fastest 802.11 configuration, etc.

- CSR view clearly indicates whether this caller has called recently for home network issues. If so, CSR is encouraged to stay on phone and work through critical issues, education, and security/performance items
- Patterns for time-based and intermittent home issues are identified in graphic form, with colors/icons over a timeline. This may involve tracking peak use times by watching bandwidth consumption for the overall gateway. This helps CSRs know when that one home may experience congestion. Slow or “chatty” devices are flagged since they reduce overall home Wi-Fi speeds
- For basic home network and configuration questions, CSR guides the subscriber to the Self-Service system – which teaches them to find their own answer in future
- Quick action buttons drive CSR action:
 - Send the subscriber their Wi-Fi SSID name and password
 - Enable guest Wi-Fi
 - Reset Wi-Fi password
 - Reboot modem/gateway
 - Re-enable radio
 - Speed test
- **General Actions:** Actions are described above, available for specific scenarios. Others general actions may include:
 - Send service summary to subscriber (email)
 - Escalate to advanced support (full details included automatically)
 - Escalate to SRO/maintenance ticket (full details included automatically)
 - Create Whole Home Birth Certificate (checks all home devices/services, archives results)
 - Whole Home Check (checks all home devices/services again, live)
- Since intermittent problems are tracked, timing information is used to roll a truck to the home when the issue is likely to be happening
- **Upsell:** In some situations, there may be opportunities for upselling the subscriber. In our “what if” scenario, insights have been generated showing that the subscriber is on a low-bandwidth package and consistently consumes near their maximum. Maybe the subscriber calls about “slow surf” occasionally too, setting up a good opportunity for upsell. One call center manager said that their data-driven upsell approach had his team out producing front-line sales & marketing efforts for service expansion
- **Ongoing Metrics:** Metrics tracking is in place to review device & subscriber behaviors, learn more about how subscribers really act, feeding back to office managers
 - Peak use times, distribution of consumer device types as well as which ones cause problems, problematic gateway models, firmware versions

8.2. Applying Concepts to the Back Office

The focus of this paper is CX and the call center. That said, these concepts apply to network surveillance systems as well. Imagine that instead of managing a bunch of SNMP traps, looking at lists/graphs of cable modems with high CER, or staring at raw RF plots, a plant maintenance manager was presented with automated situation analysis for a node. With all the raw data available to back it up, such a summary might say something like:

95% of the modems on this node are currently within tolerance, but 3 hours ago for a duration of 45 minutes, 35% of devices (60 of 200) were impaired and appears focused only on 36.8 MHz on the upstream. Of those 60 devices, 30 were actively in use and showed an average CER of 3% during that time.

Another regional scenario that might involve full band capture might say:

Strong interference has been detected at 36.8 MHz across 3 nodes in the same area. Aggregate QoE for those nodes has dropped from 96% positive to 60%. A maintenance ticket was recently opened for this area. Recommend check for sweep transmitter left in use.

8.3. Proactively Notifying the Subscriber

With all of these powerful insights aligned with the critical need to prevent/deflect support phone calls, operators might be tempted to create an aggressive subscriber notification program, sending subscribers SMS/texts and emails when any issues are found. This is a risky idea unless approached carefully.

Subscriber Frustrations: In discussing these concepts with call center leaders, the consensus was that subscribers would be frustrated receiving numerous alerts for situations they had no control over. In addition, by proactively alerting subscribers for every service issue, they might be alerted for an impairment even when they may have never noticed a service issue.

Exceptions: In some cases, some subscribers may want to receive alerts of service impacting issues or periodic QoS score summaries. Even with these, an “opt in” approach should be used:

- Confirmed, sustained network outages in their area (alert, clear)
- Service interruptions for business-class subscribers
- Service interruptions for high-SLA subscribers (Gold / 1 GB package)
- Managed service customers (e.g. paying for Wi-Fi managed service)
- Subscriber requested a one-off notification from the call center for a specific ticket (e.g. network outage or maintenance ticket)

A Better Way: A natural platform for such alerts is via mobile app. This strategy builds on existing support and customer engagement strategies already in motion. In this scenario:

- Subscriber can configure alert frequency and severity
- Subscriber can opt to view alerts on their schedule rather than be interrupted
- Subscriber could subscribe to specific alerts
- Education/FAQ blasts could be offered (mapping to current top call trends)
- Upsell / marketing alerts could be customized and targeted. Top bandwidth users could be offered an expanded package

9. Get Started

When asked about how to get started on some of the programs discussed here, one call center leader said, “Just get started. Start small, these small improvements add up.” As stated above when talking about the need for a CX champion, it’s more important that the organization commit to the goal; provide time; resources; make incremental improvements; and measure progress along the way.

9.1. Simple Changes go a Long Way

In a comparison of 8 separate domestic and international front-line call centers being used to handle calls for a telecommunications company, one stood out. This call center had an only slightly simpler CSR interface, the key difference being their CSRs saw a green “Thumbs Up” or red “Thumbs Down” icon indicating health of a device/component/service, whereas the other call centers showed more detailed, numerical data. The results were impressive: Comparing dispatch directives against field results (no trouble found, field issue found) most call centers dispatched trucks with +/- 50% success rate – meaning, half the time trucks were sent to the field, no trouble was found. A leader associated with the program said, “This was no more effective than a coin flip.” The call center with the simpler UI was consistently above 80% effective [ARRIS Internal Study].



Figure 86 – Comparing Technician Dispatch Success Across 8 Call Centers

At one Latin American operator, call center processes allowed CSRs to create trouble call/dispatch tickets based on verbal quality checks, but without any data-driven validation. At its peak, this call center was measured as having 89% of their dispatched trucks find no issue in the home and a post-mortem analysis that there was no CER/SNR issue. This is an extreme case, however, it clearly shows that by simply integrating a check for service impacting telemetry, the number of wasted truck rolls would drop significantly [ARRIS Internal Study].

Conclusion

Today’s operational and support programs have a long way to go. Output and interfaces show too much information and don’t create the insights needed to effectively manage a business. Much of the work gets pushed onto people, the rest just gets dropped.

What do organizations need to realize? The price will be paid somehow. Either organizations commit to becoming more proactive and savvy about how to manage their customer experience, or they’ll pay by being forced to support an unsustainable load of calls to the call center, a high rate of repeat truck rolls, and high customer churn rates.

It's clear that customer support programs have a dramatic effect on customer satisfaction – and therefore branding, subscriber loyalty, and business survival. Subscriber satisfaction is driven more by customer experience than by marketing and sales efforts.

Subscribers are now expecting high-touch, personalized, immediate support, and are no longer just using the telephone to get it. It's critical that programs be designed to focus very intensely on (1) preventing issues, (2) providing methods for subscriber self-service, and (3) carefully optimizing the customer support experience for analysis and resolution of issues.

Organizational strategy should be considered strongly to prevent individual directors from protecting their P&L at the cost of the overall organization. A true customer champion role can be a unifying force in the battle for organizational alignment.

Programs like those discussed in this paper provide important service insights that are valuable to multiple internal audiences, only one of which is the customer support center. It should be clear that the value of that core investment is magnified by its impact on multiple groups within the organization. When architecting such programs, leaders should be thinking of all the ways resulting artifacts/insights can be leveraged across the organization.

The savviest operators are starting to realize that, when choosing a vendor to work with for architecting and executing their customer experience vision, it has less to do with finding the cheapest outsourced call center or ready-made workflow tool, and much more to do with which company knowing the systems, data, hardware, and cross-organization strategies the best.

Abbreviations

Term	Definition
Call Deflection	Using various mechanisms to intercept or prevent a support call to a CSR in the call center
CER	Codeword Error Ratio
CPE	Customer Premises Equipment
CSR	Customer Support Representative
CX	Customer Experience
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable System Interface Specification
EQ	Adaptive Equalization
FAQ	Frequently Asked Questions
FEC	Forward Error Correction
HFC	Hybrid Fiber Coax
I&R	Install & Repair
IoT	Internet of Things
IT	Information Technology
IVR	Interactive Voice Response
M&P	Methods and Procedures
NOC	Network Operations Center

Omni-Channel	Implies a contact center supporting various communication paths for subscribers to get support, including online chat, text, phone, email, web, mobile app, social media
OPEX	Operational Expense
P&L	Profit and Loss
PNM	Proactive Network Maintenance
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
RF	Radio Frequency
RMA	Return Merchandise Authorization
RSSI	Received Signal Strength Indicator
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
TR-069	A technical specification that defines an application layer protocol for remote management of end-user devices, published by the Broadband Forum
XMPP	Extensible Messaging and Presence Protocol

Bibliography & References

Berg Insight (Jun 26, 2014), *mHealth and Home Monitoring*, (M2M Research Firm)

American Hospital Association (AHA) (Apr 25, 2016), “*Telehealth: Helping Hospitals Deliver Cost-Effective Care*”, Washington, DC

Walker Information (2013), *Customers 2020 Study*, Indianapolis, IA

Gartner, Inc, (2011) “*CRM Strategies and Technologies to Understand, Grow and Manage Customer Experiences*”, Gartner Customer 360 Summit

Nokia (Aug 31, 2016), “*Eliminating the Potluck Aspect of Call Center Service*”, Campbell, Alasdhair

OECD (2013), “*Building Blocks for Smart Networks*”, OECD Digital Economy Papers, No. 215, OECD Publishing, Paris

Nokia (Feb 18, 2016), “*Big changes – again – For Cable Operators’ Business*”, Davidson, Steve

Accenture, Super Office (June 14, 2017), “*32 Customer Experience Statistics You need to Know for 2017*”, Kulbyte, Toma

Gartner, Inc (Oct 20, 2003), “*Enhanced Customer Service Can Potentially Harm Your Business*”

Forbes (Oct 12, 2016), “*Why Your Company Needs a Chief Customer Officer*”, Chris Davis, Alex Kazaks, Alfonso Pulido, Contributions by McKinsey & Company

Aspect (Apr 2, 2015), “*Digital Interaction to Change Customer Service Forever*”, Jason Dorsey, excerpts from The Millennial Report, (Conversion Research)

Van Belleghem, Steven, (Jun 18, 2013), “*The Real Self-Service Economy*”, page 13

Access Network Data Analytics

(Machine Learning Applied to Cable Access Data)

A Technical Paper prepared for SCTE•ISBE by

Karthik Sundaresan

Principal Architect
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
303-661-3895
k.sundaresan@cablelabs.com

Jay Zhu

Engineer
CableLabs
858 Coal Creek Circle, Louisville, CO, 80027
303-661-3312
j.zhu@cablelabs.com

Introduction

The cable access network has large amounts of monitoring data collected across various network equipment (CMTS, CMs, STBs, APs etc.). This data reflects the state of the network, status of devices and outside plant. Some examples are PNM data (RxMER, Channel coefficients), CM and CMTS MIBs (FEC stats, Service flow stats, packet drop counts), IPDR data, etc. Now, given this data set, how can the operator identify network plant issues, reliably, proactively and automatically?

Combining the analysis of data along, with network topology and device location, it is possible to create a general view of the plant condition and isolate problem sources. This paper tackles questions around, how can the operator identify network plant issues on a single modem, across a serving group, across fiber node, and correlate data across all devices in a reliable, proactive and automatic way. Machine-learning / data-analysis applications can crunch through layers of data to identify patterns and prioritize network issues automatically, allowing operators to reduce troubleshooting and problem resolution time, thereby reducing operational costs and enhance network reliability.

Problem Statement

1. Lots of data, no knowledge

There are various areas in the cable network space, each with large sets of data, to which various Machine Learning techniques can be used to solve various problems.

Proactive Network Maintenance (PNM) involves processing vast amounts of fine-grained information about the state of the network to look for patterns that indicate impending problems. Current PNM solutions have focused on leveraging the physical-layer channel measurements performed by DOCSIS cable modems and CMTSs, along with knowledge of the HFC plant topology and well-known digital signal processing algorithms in order to locate physical degradations of the plant. ML algorithms open up this space to identify a much wider variety of network impairments, particularly when those impairments aren't predictable using a priori knowledge.

CMTS and cable modem features and capabilities can be leveraged to enable measurement and reporting of network conditions so that undesired impacts like plant equipment and cable faults, interference from other systems, and ingress, can be detected and measured. With this information, cable network operations personnel can make modifications necessary to improve conditions and monitor network trends to detect when network improvements are needed.

Some of the PNM measurements are very valuable in that they can reveal problems in the network. We illustrate a couple of examples to which we can effectively apply Machine Learning techniques. This will help in quickly and automatically identifying problems, instead of a human operator looking through the data and flagging issues manually.

2. Why Machine Learning

A CM with full-band tuner capability acts like a spectrum analyzer. The spectrum analysis reveals suck-outs, excessive tilt, frequency peaking, unwanted filters, FM (or other) radio ingress, early roll-off, etc. Each of these issues (signal patterns) can be learnt by a Machine Learning application, which can then monitor live signals to flag problems in the network.

Machine Learning algorithms can look for abnormalities from the data obtained from each CM and across all CMs in the network. Grouping of CMs with common problems can be used to identify issues affecting multiple subscribers and isolate the cause of an issue in the cable plant.

The data sources which we start off with is primarily is the PNM data available from the cable modem today. We can add to that DOCSIS MAC layer and packet statistics along with CM status objects (e.g. T3/T4 timeouts, partial service indicators etc). Understanding the channel response characteristics can also help with creating optimal D3.1 profiles. All this can be combined to define a “health metric”, a human readable shorthand to identify the status of a CM. Visualization of data is a powerful first step in understanding data and this paper will also discuss viable options to visualize data for the operators. This paper will discuss methods for anomaly detection, root cause analysis, historical behavior analysis, pattern recognition, classification, prediction and condition evaluation in the access network data. See paper references [1] & [2] for further details on how machine learning applies to access network data analytics

Data sources

3. Sources (how)

The DOCSIS 3.0 and 3.1 technologies introduce CMTS and cable modem features and capabilities that can be leveraged to enable measurement and reporting of network conditions. This helps to detect and predict undesired impacts such as plant equipment and cable faults, interference from other systems and ingress. The goal is to rapidly and accurately characterize, maintain and troubleshoot the upstream and downstream cable plant, to guarantee the highest throughput and reliability of service. See [7](DOCSIS PHYv3.1).

3.1. Data types (what)

There are a few powerful indicators for problems in a cable network. These include details such as excessive transmit power at the CM or low receive CM signal levels. There were also indicators for high packet or codeword errors, both at the CM and CMTS, and low receiver MER (modulation error rate, a measure of noise and interference in the signal. While these indicators are useful, they are a view of a single endpoint in the network and don't tell an operator exactly what is the problem. It will be useful to look at these indicators not only from the point of view of a single end-device, but holistically across the entire network to understand the root causes.

Some of the data types which are useful for such analysis are

- CM,CMTS RxMER level per subcarrier
- CM, CMTS Transmit & Receive power levels

- FEC Statistics (Correctable vs Uncorrectable)
- Channel pre-equalizer Coefficients
- Downstream vs Upstream parameters
- Network Topology
- Topology map, Latitude/longitude of the devices
- Overlay map of local Wireless channels and their frequencies (e.g. LTE channels)

3.2. PNM Data (TFTP)

CMs and CMTS can upload Captured PNM Data to the configured PNM Server via TFTP. These PNM data can include the following measurements. There are MIB objects to configure/trigger the file transfer and measurements.

DOCSIS Downstream PNM Measurements and Data include: Symbol Capture, Wideband Spectrum Analysis, Noise Power Ratio (NPR) Measurement, Channel Estimate Coefficients, Constellation Display, Receive Modulation Error Ratio (RxMER) Per Subcarrier, FEC Statistics, Histogram, and Received Power.

DOCSIS Upstream PNM Measurements and Data include: Capture for Active and Quiet Probe, Triggered Spectrum Analysis, Impulse Noise Statistics, Equalizer Coefficients, FEC Statistics, Histogram, Channel Power, and Receive Modulation Error Ratio (RxMER) Per Subcarrier.

3.3. SNMP data

The CMs and CMTSs also reflect a lot of the PNM Data measurements to via MIB objects. These include the following measurements/data objects:

- CM Downstream Objects
 - CmDsOfdmSymbolCapture
 - CmDsOfdmChanEstimate
 - CmDsOfdmMerForCandidateProfile
 - CmDsOfdmRequiredQamMer
 - CmDsOfdmHistogram
 - CmDsOfdmRxMer
 - CmDsOfdmFecSummary
 - CmDsOfdmRequiredQamMer,
 - CmDsOfdmMerMargin,
 - CmSpectrumAnalysisMeas
 - CmDsConstDispMeas,
 - CmUsPreEq,
 - D3.0 full band capture
 - CmDsHist
- CM Upstream Objects
 - CmUsOfdmaEqualizerCoefficients
- CMTS Downstream Objects
 - CmtsDsOfdmSymbolCapture
 - CmtsDsOfdmNoisePowerRatio
- CMTS Upstream Objects
 - CmtsUsOfdmaActiveAndQuiet Probe
 - CmtsUsOfdmaImpulseNoise
 - CmtsUsOfdmaHistogram
 - CmtsUsOfdmaRxMerPerSubcarrier
 - CmtsUsSpectrumAnalysis

Solution Approach

4. Understand the problem (Manual)

4.1. Data Visualizations

Data visualization describes the effort to help understand the significance of data by placing it in a visual context. Patterns, trends and correlations that might go undetected in text-based data can be exposed and recognized easier by data visualization.

A picture is worth a thousand words, only when the story is best told graphically rather than verbally and the picture is well designed. One could stare at a table of numbers all day and never see what would be immediately obvious when looking at a good picture of those same numbers.

A table of CM RxMER values can do two things extremely well: it expresses the MER values precisely and it provides an efficient means to look up values for a particular subcarrier and time. But if we're looking for patterns, trends, or exceptions among these values, if we want a quick sense of the plant state contained in these numbers, or we need to compare whole sets of numbers rather than just two at a time, a table approach fails.

Instead of studying single snapshots of topical issues or events in ever-greater detail, however, we can create “datascoptes” that can be used to zoom in and out of large data sets in search of new understanding.

In this paper, we are primarily working with CM RxMER values obtained from a live field trial DOCSIS 3.1 plant. The DOCSIS 3.1 OFDM Channel here is 96 MHz wide, i.e. 1920 (50Khz) subcarriers. The number of Active sub carriers within the channel is ~1750 subcarriers. Due to errors in the data capture, or misalignment in the channel set up across different CMTs, we chose to manually best align the MER values across the different CMs. (set of Data subcarriers were shortened to less than 1750.)

The primary goal of the data visualizations here is to understand the current Channel and plant conditions and observe patterns visually for a CM over a time period and for sets of modems at a given time and for sets of modems across time. The idea is to see if there are common issues to identify (and the answer is yes.)

The idea is to take the raw data and create visualizations that help understand the plant state. We used a mix of 2d and 3d plots, to understand patterns over frequencies, over time and over topology. We are learning from the data about the kinds of issues present and how best we can automatically detect them.

We are using mainly RxMER data, as that was the data available to us, but some of these principles (visualizations and automatic anomaly detection) will apply to other data types as well.

4.2. CM RxMER Visualizations

The below Figures show the CM RxMER Data plots. The x-axis is the sub-carrier frequency, the y-axis is the time over which the captures were done, and the vertical z-axis is the actual reported RxMER db values per sub-carrier.

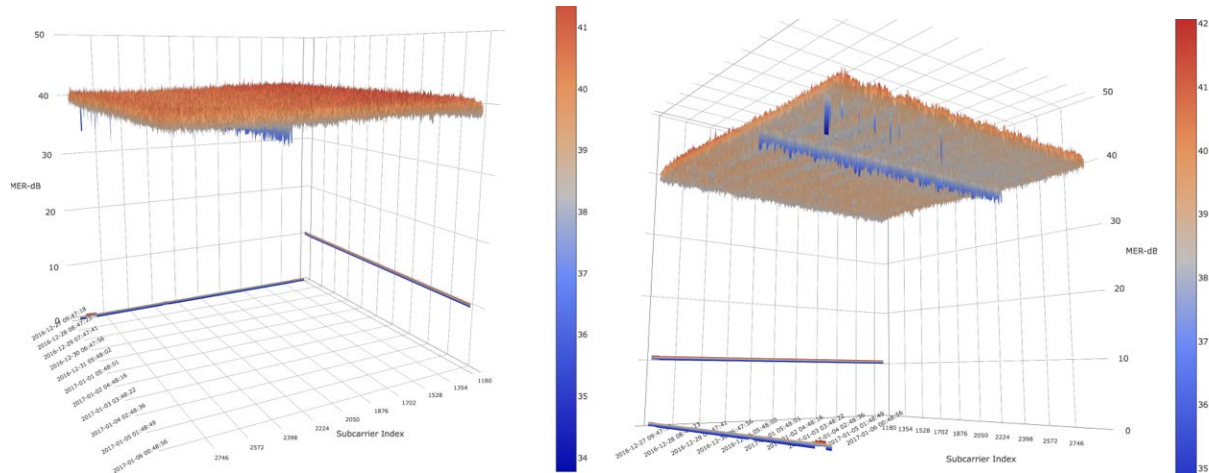


Figure 1 – CM RxMER data over frequencies and time (Same CM)

Figure 1 shows the CM's RxMER values which looked to be centered around 40dB, most of the values are near that, though one can see a small dip in the RxMER values around subcarrier index 2398-2224

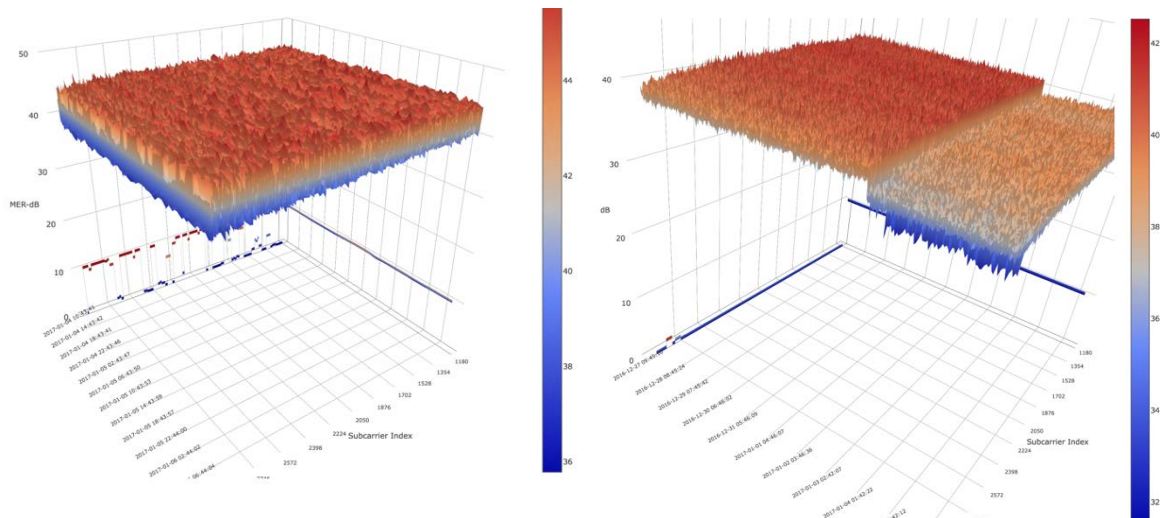


Figure 2 – CM RxMER (2 different CMs, variations and sudden drop)

Figure 2 shows the two different CM's RxMER values, the one on the left look to be centered around 41dB, but show a lot more variability in frequency and in time (max of 45db to min of 37 db. The second

modem shows RxMER values which are centered at 40db for a few weeks, and then jumps down to ~37db going forward. This essentially means some event happened in the network causing a drop in the RxMER. (Was this the effect of perhaps a splitter added ahead of the CM in the home network, or could it be something else ?)

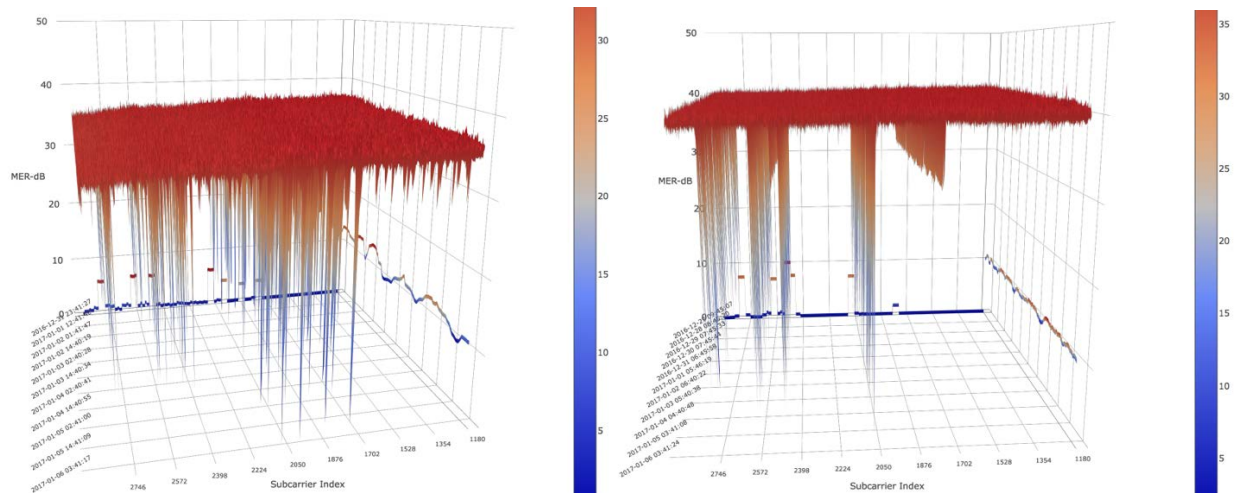


Figure 3 – CM RxMER data , different examples of issues

Figure 3 shows the two more CM's RxMER values, with each CM showing issues at specific sets of subcarriers, and in some cases over time.

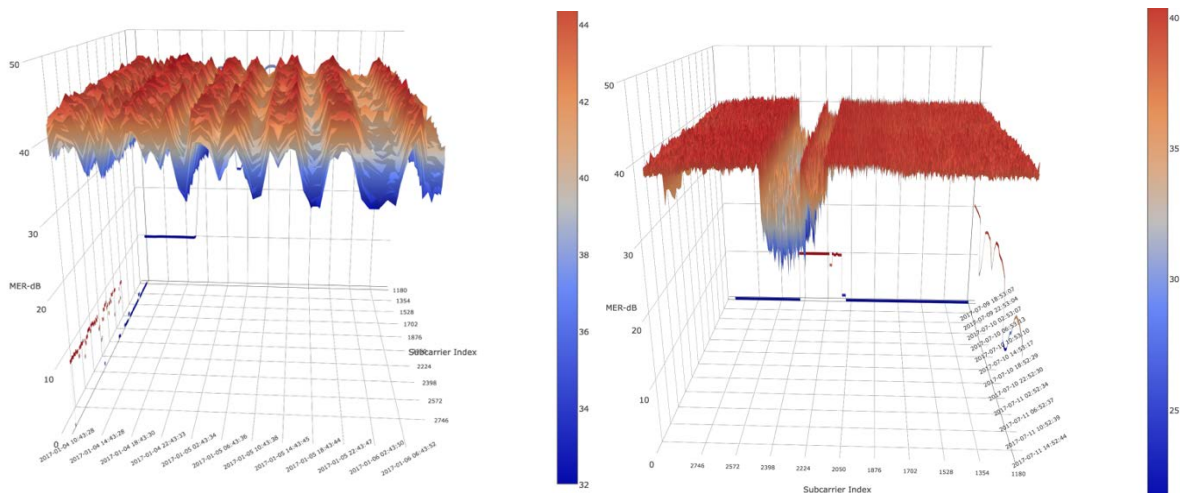


Figure 4 – CM RxMER data (oscillations over time, interference)

Figure 4 shows the two more CM's RxMER values, with the left CM showing issues in a somewhat periodic fashion, with the average RxMER value oscillating up and down over time. The second CM is showing issues over a certain set of frequencies, in this case this happens to be about 300 subcarriers or 15 MHz wide. Now as example of the possibilities of automatic issue identification: Some of the wider LTE channels are either 5,10,15 or 20 MHz wide, so if one could co-relate this interference issue seen on a set of CMs at a specific frequency to the local LTE frequency/channel map corresponding to the CM's

physical locations, one might be able to classify this issue as LTE interference and then go look for ways to mitigate it.

Hopefully the above sub-set of CM RxMER plots gives the reader a bit of understanding into the plant and how the plant response varies over frequencies and time.

4.3. Map based visualizations

Another type of visualization is geographical-location based. Here the idea is to overlay the per-CM PNM data on a map, where plant layout, fiber node locations and the modem locations are known and mapped. In the figures shown below the CMs are grouped and analyzed hierarchically under the fiber node serving that location.

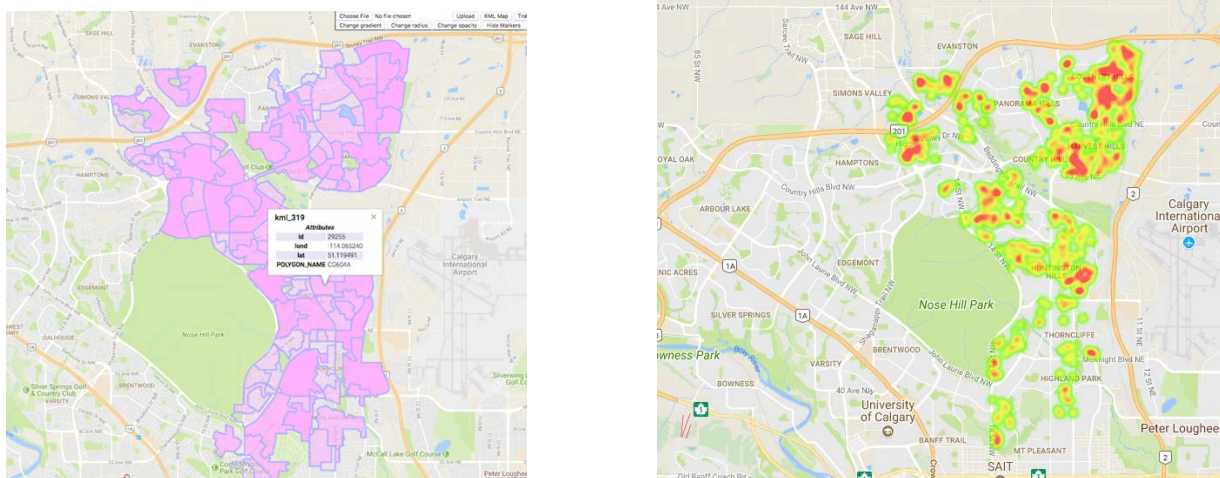


Figure 5 – CMTS Serving FiberNode locations, Heat Map of CM Rx Power level

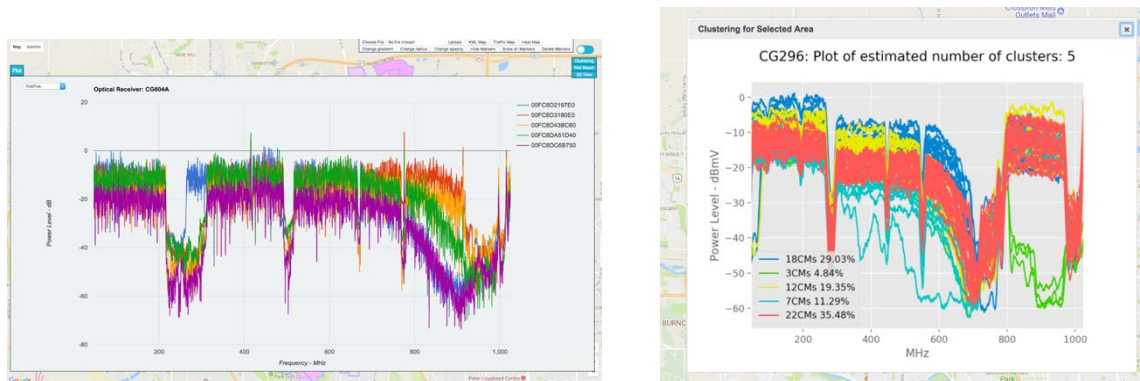


Figure 6 – CM Rx Power level over 1GHz spectrum, & K-Means clustering of all CMs based on RxPower level

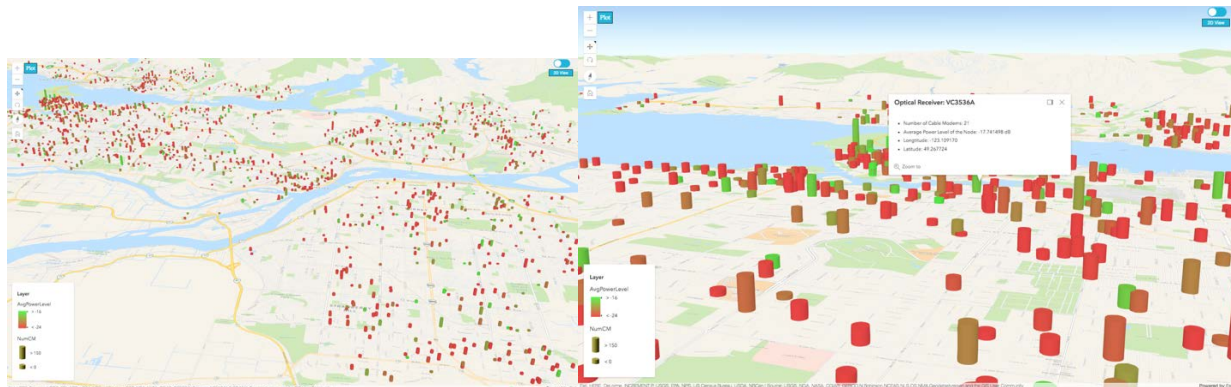


Figure 7 – Map of Fiber nodes with CM Rx Power level grouped together

5. Machine learning approach to Anomaly detection

5.1. Preparing the data

Machine learning algorithms learn from data. It is critical to feed the learning algorithms the right data for the problem one wants to solve. Even if there is good data, one needs to make sure that it is in a useful scale, format and meaningful features are included.

The process for getting data ready for a machine learning algorithm can be summarized as: Select Data, Preprocess Data, Transform Data.

The selection step is concerned with selecting the subset of all available data that you will be working on. One needs to consider what data is actually needed to address the problem.

Three common data preprocessing steps are formatting (data may be in a proprietary file format and you would like it in a relational database or a text file), cleaning (removal or fixing of missing data.) and sampling (a smaller representative sample of the selected data).

Three common data transformations are scaling, attribute decompositions and attribute aggregations. This step is also referred to as feature engineering.

Data preparation can involve a lot of iterations, exploration and analysis.

5.2. Machine Learning approach

The machine learning approach to automatic anomaly detection we have taken in this paper is as follows:

Using combined history of all CM's RxMER data, we identify the different anomalies in the CM's RxMER data, and use that to train a Machine learning classifier. We use this classifier to quickly predict and the label issues on new incoming data. We then analyze the patterns in these anomalies to detect system wide patterns and generate operational information/knowledge for the operator.

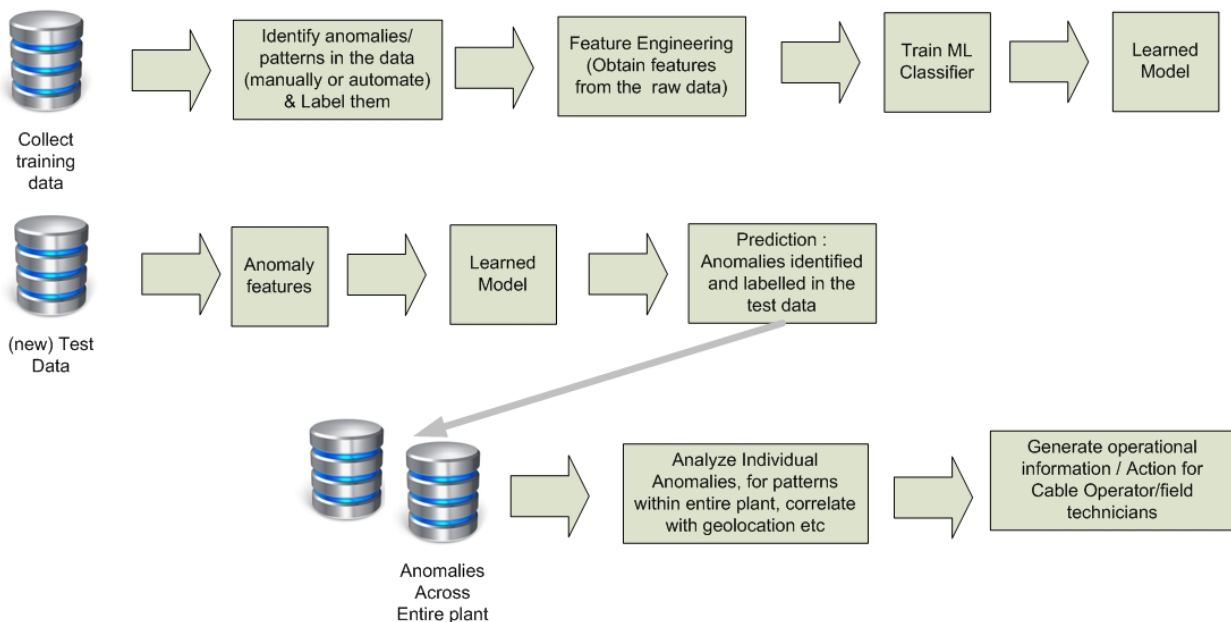


Figure 8 – Machine learning approach

We divided the data samples we had into training (70%) and test (30%) sets. To train a machine learning classifier we need events identified and labeled for the classifier to train on.

One could do this manually but this would be very kind and labor-intensive asset. One would need to walk through thousands of samples from each cable modem to identify issues and then label them. Manual labeling is too slow and doesn't allow for experimental changes as one goes through the process.

We decided to implement an anomaly detector which will automatically identify issues in the CM's RxMER data sample and label them into somewhat generic buckets of issues. In this paper we chose tilt, roll-offs, Sharp MER drops, and wide MER drops as a starting point in terms of the set of issues we wanted to identify. In the future, this set of labels or issues can be extended two specific patterns which we want to identify in the RxMER data.

Once the anomaly detector detects and labels patterns, we reduce the data by extracting the features of the raw data. In this case, we are extracting features from the anomaly such as the width, the depth, the average drop in RxMER, etc. to create a subset of features which the ML classifier will learn on instead of all the raw data.

5.3. Initial Anomaly Detection: Sliding Median & threshold comparison

A moving average is commonly used with time series data to smooth out short-term fluctuations and highlight longer-term trends or cycles. The threshold between short-term and long-term depends on the application, and the parameters of the moving average will be set accordingly. Viewed simplistically it can be regarded as smoothing the data. From a statistical point of view, the moving average, when used to estimate the underlying trend in a time series, is susceptible to rare events such as rapid shocks or other anomalies. A more robust estimate of the trend is the simple moving median over n time points.

A moving median is less sensitive to outliers, where an outlier is usually a single point in time series that is very different from all others, which may be due to some kind of error. Moving median filter simply removes outliers from the result, where moving mean/average always takes into account every point. However, moving median can be even more sensitive to short-term significant spikes that span several points, especially when they span more than half of the moving window.

The current implementation for anomaly detection on a CM's RxMER values is as follows:

For a single CM RxMER sample, maintain a sliding window with a set of different window sizes. The window slides through all subcarriers on both directions from a low frequency to high frequency and from high frequency to low frequency. It generates a list of median MER values of each window position, which is smoother than the original sample and can be used for shape comparison to extract anomaly events (dips) from the original MER sample very precisely. With different window sizes, the sliding median algorithm can look at the sample from different scales to cover as much anomaly events as possible. In the current implementations, we use 2 window sizes of 200, 800 subcarriers. The other option is to use a Savagol filter (see reference [6], for a description of Savitzky-Golay Filtering) which removes very high frequency noise from the data.

In this paper we chose the sliding median algorithm, it is used to detect the anomalies, essentially any deviations below the sliding median RxMER value when comparing values from the original sample. The algorithm flags any subcarrier values below with threshold and these are marked as anomalies.

5.4. Feature Extraction

If there are many independent features that each correlate well with the class, machine learning is easy. On the other hand, if the class is a very complex function of the features, one may not be able to learn it. Often, the raw data is not in a form that is amenable to learning, but one can construct features from it that are useful. It is advantageous to reduce the number of features considered to focus on a subset of particular interest; this is called feature selection. After anomaly detection, for each anomaly, we reduce each sample to a set of features which define that sample. This process is feature engineering or dimensionality reduction. Using the best and the least features to describe a learning problem is very important in Machine Learning, this is primarily for improving accuracy of the analysis. It also reduces measurement costs, creates faster systems with less memory, and allows simpler interpretation of the results.

Here we take a few relevant features which describe the anomaly accurately instead of the whole raw anomaly data sample. In particular, the features we use are the area occupied by the anomaly, the width (in subcarriers), lowest MER value, the average MER value at the bottom of the anomaly, the start MER level & end MER levels of the anomaly etc. In order to get rid of noise and to let the classifier be focused on certain aspects from the anomaly events, once the anomaly detector finds and extracts anomalies from the RxMER sample it immediately extracts features / regions of interest by analyzing the event width, starting MER value and ending MER value (relative), MER level of the bottom of the event, and the event score (the accumulative area / db below to the threshold) etc.

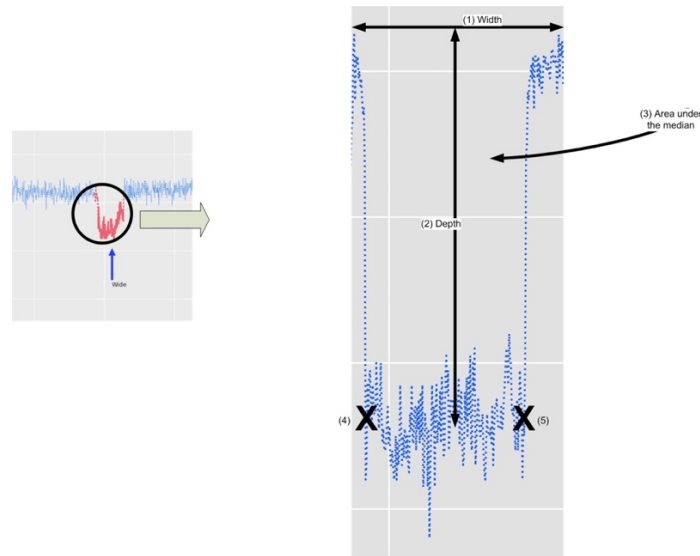


Figure 9 – Feature engineering approach

These extracted features from each event can be used as the input (after labeling) to train and test an intelligent classifier.

5.5. Data labeling

Unlabeled data consists of samples of natural or human-created artifacts that you can obtain relatively easily from the world. Labeled data typically takes a set of unlabeled data and augments each piece of that unlabeled data with some sort of meaningful "tag," "label," or "class" that is somehow informative or desirable to know. ML problems start with data — preferably, lots of data (examples or observations) for which you already know the target answer. Data for which you already know the target answer is called labeled data. In supervised ML, the algorithm teaches itself to learn from the labeled examples that we provide.

The labeling approach we take here can be based on clustering of the basis of the raw samples or clustering based on extracted features. Clustering based on the extracted features, was faster and more reliable. We use a simple K, means clustering to manually label the groups and combine them into larger ones. With the fact that similar patterns exist in the anomalies we detected, we can cluster the anomalies into a few major groups. The noise in the anomalies can be an issue when trying to use K-means to automatically generate clusters with a small value of K. Increasing the value of K can prevent considerable impact from the noise since the K-means will create much smaller groups, and therefore, makes much less clustering mistakes and reduces the number of samples, which would help generate much cleaner visualizations for manual identification and labeling.

For the smaller clusters generated by K-means, we manually check each of them and group similar ones into larger clusters. For those smaller clusters who have too much noisy samples and mis-clustering samples, we remove them to keep the dataset clean. We then give each large cluster a numeric label and they are ready to be used as the input for training classifiers.

For now, we are limiting the work to 4 simple labels, the idea in the future is to increase the label space. Temporarily, we have 4 labels for all the anomalies samples. However, the anomalies should be classified more clearly in the future. For example, from our observations, the wide anomalies could be separated into a few smaller groups because some of them have slightly different patterns and they could be caused by different interference sources.

Based on thresholds for each of the anomaly types, we differentiate each incoming sample as tilt (up vs down), roll-offs (left end or right end), sharp MER drops, and wide MER drops). Each sample from the anomaly detector is labeled as one of these anomalies. The thresholds are set such that some of the samples are discarded, so that with the labelled samples we use to train Machine learning classifiers, we use very clean data to learn on. A manual review of these labels is performed as a sanity check.

The labeling operates in stages, each stage identifying the labels for that type of anomaly.

5.5.1. Synthetic data generation

Synthetic data is an alternative when you don't have enough actual data samples for training. Once a general pattern has identified in the actual data, which is distinctive and can be described, one can create pseudo-random synthetic data samples which follow the same pattern characteristics. This synthetic data is useful to train a machine learning classifier to learn the properties of known features, for future identification.

By learning from the data, it's possible to abstract the similarity of patterns and guess what they should look like with random changes. By doing so, we developed out synthetic RxMER data generator that can randomly create anomaly events (sharp anomalies, wide anomalies, tilts, roll-offs) and put them into RxMER samples that have different levels of noise and different center of MER values. When there is not enough data or when the real samples don't have much variability, we may use the synthetic data generator to provide a larger sample space for the classifier to get trained. With experiments, we have proven that the synthetic data acts very similarly to the real data during the training process and successfully prevented overfitting of the classifier, as the feature extraction can abstract the most important information from anomaly events, and get rid of the non-real parts from the synthetic data.

5.6. Classifier (Model)Training

The fundamental goal of machine learning is to generalize beyond the examples in the training set. This is because, no matter how much data we have, it is very unlikely that we will see those exact examples again at test time.

In Decision Tree Learning, a new example is classified by submitting it to a series of tests that determine the class label of the example. These tests are organized in a hierarchical structure called a *decision tree*. The training examples are used for choosing appropriate tests in the decision tree. Typically, a tree is built from top to bottom, where tests that maximize the information gain about the classification are selected first.

Support Vector Machine (SVM) is a supervised machine learning algorithm which is mostly used in classification problems. Each data item is a point in n-dimensional space (where n is number of features) with the value of each feature being the value of a particular coordinate. Classification is performed by finding the hyper-plane that differentiate the classes well.

In this paper, we chose Convolutional Neural Networks (CNN) as the basis for our machine learning model. (See reference [5] for details on CNN).

5.6.1. Neural networks & Deep learning

Neural Networks are models that are inspired by the structure and function of biological neural networks. They are a class of pattern matching algorithms that are commonly used for regression and classification problems. Deep Learning methods are a modern update to Neural Networks that exploit abundant cheap computation. They are concerned with building much larger and more complex neural networks. Deep Learning is a type of Neural Network Algorithm that takes metadata as an input and processes the data through a number of layers of the non-linear transformation of the input data to compute the output. Deep Learning is about learning multiple levels of representation and abstraction that help to make sense of data, the algorithm automatically grasps the relevant features required for the solution of the problem. In Deep Learning Neural Network, each hidden layer is responsible for training the unique set of features based on the output of the previous layer. As the number of hidden layers increases, the complexity and abstraction of data also increase. It forms a hierarchy from low-level features to high-level features. The most popular deep learning algorithms are: Deep Boltzmann Machine (DBM), Deep Belief Networks (DBN), Convolutional Neural Network (CNN), Stacked Auto-Encoders.

Convolutional Neural Networks are a category of Neural Networks that have proven very effective in areas such as image recognition and classification. ConvNets have been successful in identifying faces, objects and traffic signs apart from powering vision in robots and self-driving cars.

5.6.2. Prediction / Signature Recognition

The feature space is very much reduced and most of classification algorithms can handle it well. Although we can use SVM or KNN instead of using Neural Networks and they will work just as well as NN does for our current problem, we still use the deep architecture (Neural Networks) that offers statistical scalability, computational scalability and human-labor scalability over shallow architectures (KNN, SVMs..) to prepare for the increasing of the data amount and complexity in the future.[3]. The classifier we use now is a 6-layer convolutional Neural Network implemented in Keras with tensorflow as the backend.

We are using two separate classifiers, one to identify tilt issues, as they are across the whole CM RxMER sample, and another classifier to identify all other anomalies. The data is split up as 70 % for training and the remaining 30% as test. The neural network is using a batch size of 200 (number of samples). Within each epoch, the model is saved and retrained. The idea is to train the model, and then plug it back in and see if it really works. We are running 36 epochs until we reach the required training accuracy, i.e. the training loss per class. In each epoch we also validate against the Test data set, for accuracy (validation loss) against the test data. During each epoch, the accuracy increases. We stop training when the testing loss starts increasing.

5.6.3. Save Model (preservation)

Once the classifier gets trained and the model gets preserved, we add an additional stage to the anomaly detection process to use the classifier to identify what the anomalies are likely to be. Therefore, we can collect the information / ratings provided by the anomaly detector with details about the impact, location

(on frequency), width, and group (signature recognition) of each anomaly happening on each CM, and this information can be used for higher level clustering and root cause analysis.

5.7. Event signature Identification

For all new test data incoming from the network we perform the following steps.

- Anomaly detection : As a preprocessing step, we use the anomaly detector to pick out outliers in the test data.
- Prediction / Identification / Labelling: The anomalies detected in the above step are sent to the Neural network model, to classify it as a particular type of issue. The anomalies are now labeled as a specific type, along with (subcarrier) location of the anomaly, the width etc.

The below sets of figures show the labels generated by the machine learning algorithm.




Figure 10 – Clean CM data, no issues identified.

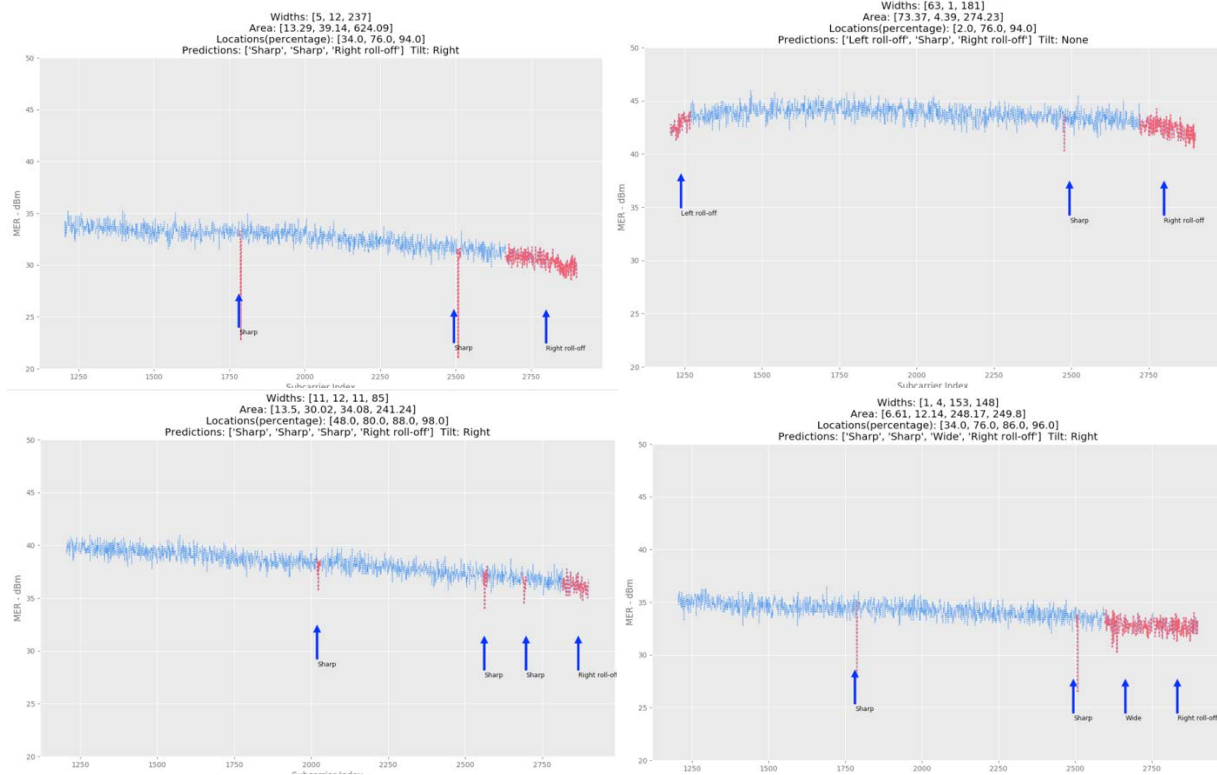


Figure 11 – Issues Identified: Tilt, Roll off

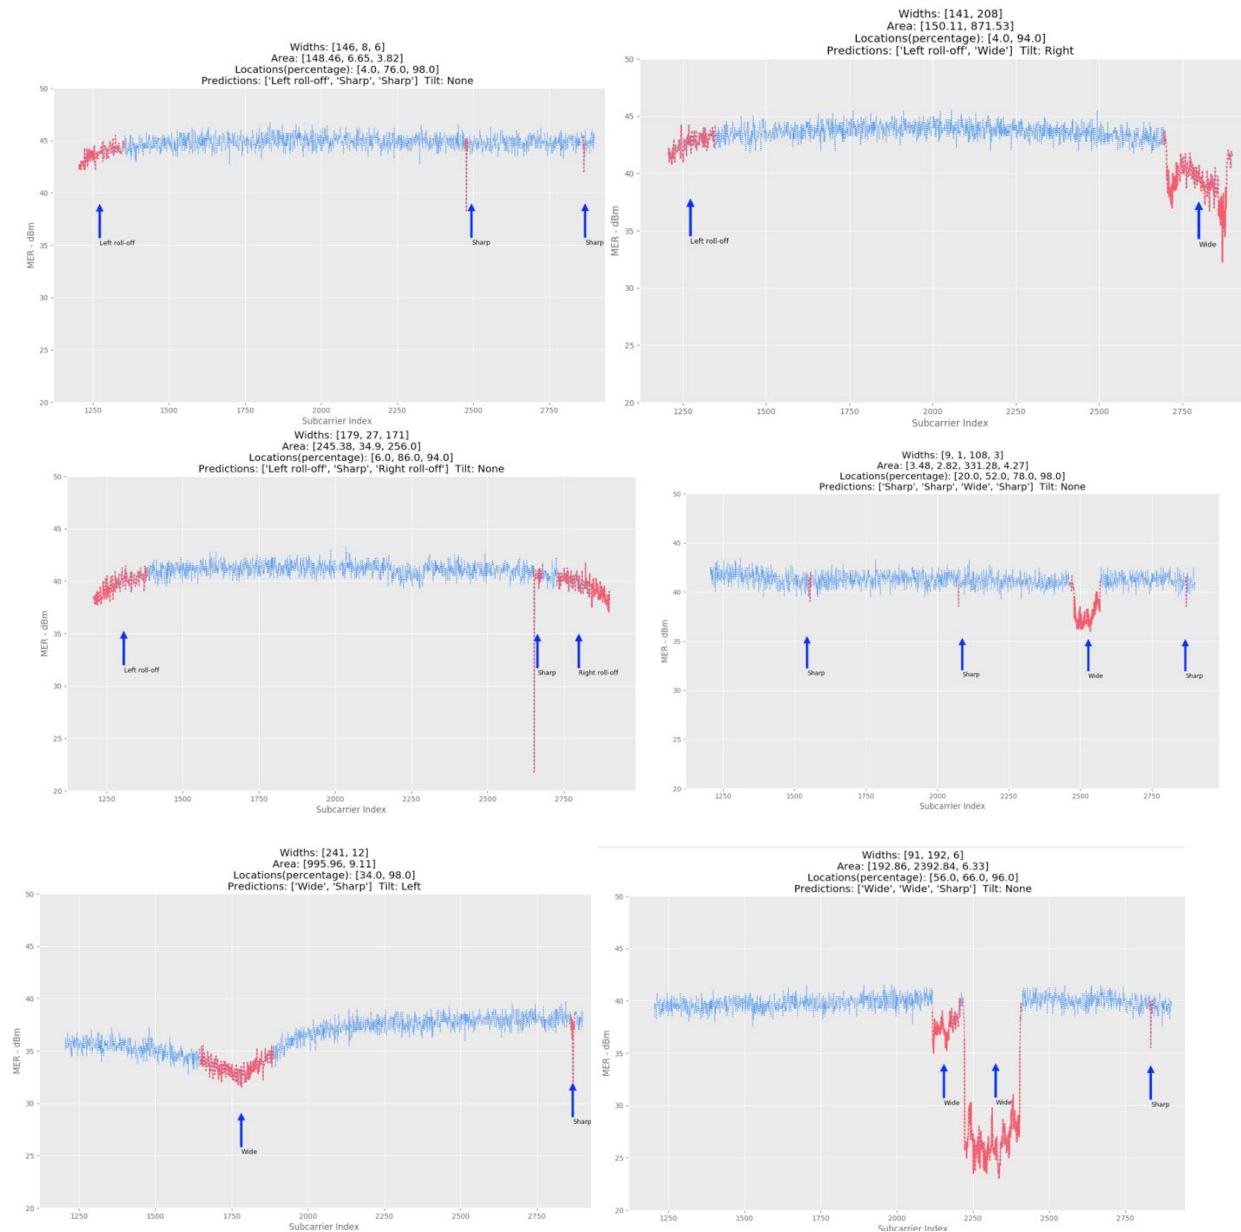


Figure 12 – Issues Identified: Wide dips, Sharp dips

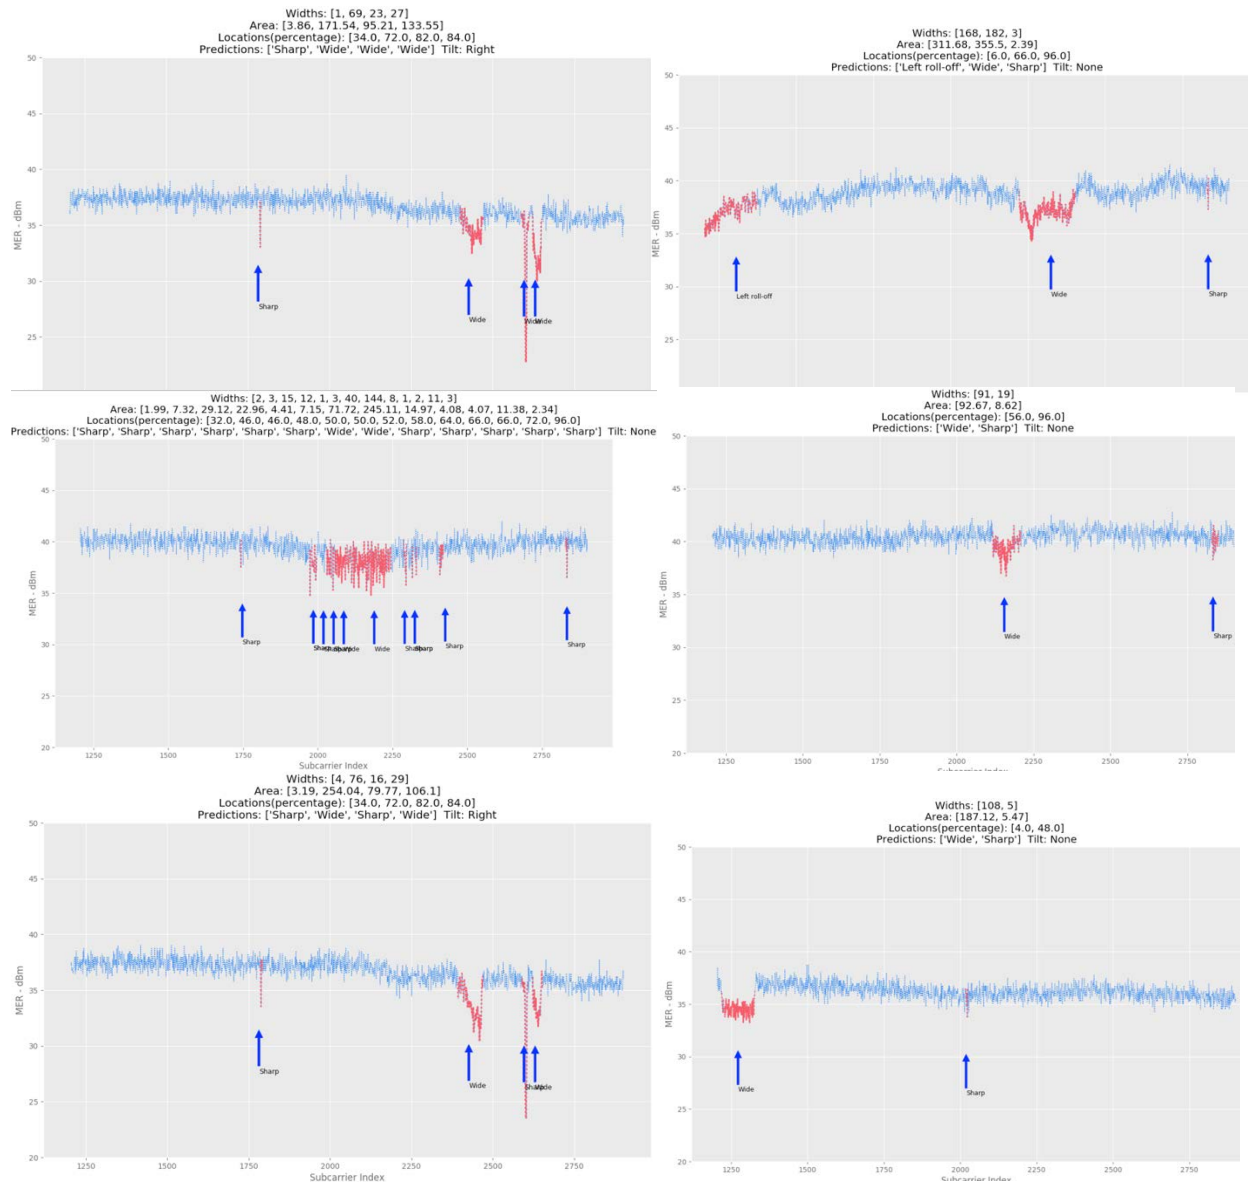


Figure 13 – Issues Identified: Various

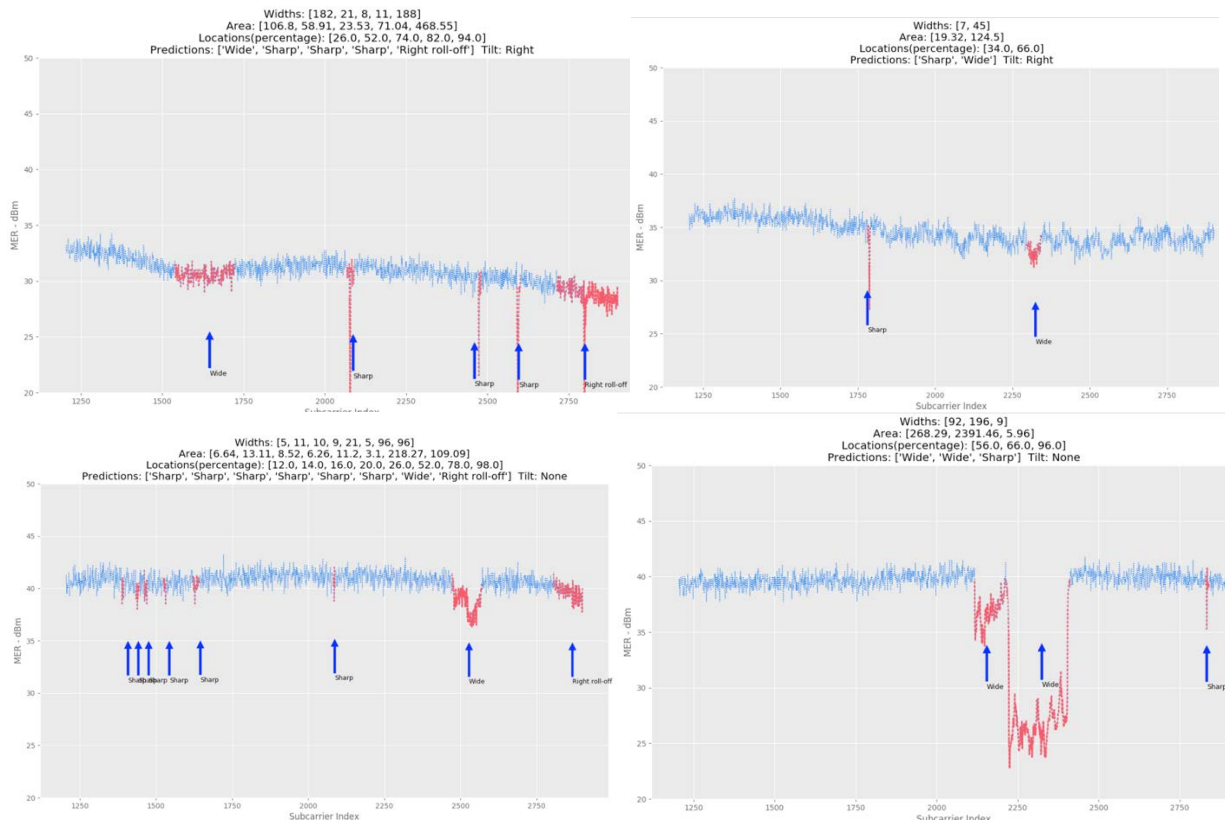


Figure 14 – Issues Identified: Various

5.8. Opensource tools used

There are many opensource tools can be used to test or create machine learning algorithms. Scikit-Learn is a convenient library that can be used to experiment with machine learning algorithms easily. We use Scikit-Learn to test the performance of different algorithms and the efficiency of our feature extraction algorithm. There are limitations of Scikit-Learn that it is not very scalable and not able to use the power from GPUs when training Neural Networks, but it's still a very good tool to do experiments with.

For loading and writing files, we use pandas. It's a very well developed opensource data analysis tool and it can be used to perform data preprocessing.

Instead of sticking with Python's matplotlib, we use Plotly to create plots and perform data visualization for the most part. Plotly is very powerful and easy to use across multiple languages such as Python and Javascript. We use Plotly to generate interactive 2D plots, 3D plots to help our data analysis process.

For constructing our neural network, we use Keras with Tensorflow as it's backend. Keras is a very useful library that abstracts the implementations of powerful machine learning libraries such as Tensorflow and Theano to help researchers and developers to easily and quickly create and test

different structures of Neural Networks. The limitation of Keras is that it's only able to use the power from a single GPU. This limits its usability when implementing large Neural Networks with large datasets. Some other libraries would support to utilize multiple GPUs to have shorter training time, e.g.: Theano. Those opensource libraries can be used in the future when the architecture of Neural Networks is mature and we need to deal with datasets that are very large.

6. Identifying global patterns

The Machine Learning classifier above gives us quick identification of issues across the DOCSIS 3.1 OFDM channel. The next step is to agglomerate these individual CM issues across all the CMs on the channel. Looking at the issues as seen by all the cable modems will give a view into the health of the network and the plant.

- Across topology
 - One view when looking at the data as a whole would be to identify issues across the plant topology. For example, the algorithms may be able to identify that many of the issues have a common root at some location in the cable plant.
- Across a specific geolocation
 - Another way to look at data as a whole would be to identify issues which correlate to specific geographic locations. For example, the algorithms may be able to correlate the issues seen in certain cable modems in certain geographic locations to say the LTE Channel map and frequencies for that location.
- Across CMTS
 - Another way of grouping issues in the data would be to group issues based on the fiber node are the CPS to which the modems are associated with.
- Across CM types
 - Another view of the issues could be to group issues seen by specific modem manufacturers together. There may be patterns which can emerge along this feature as well.

Root cause analysis: Analyzing the data across these global plant views, gives the operator a chance to understand the root causes of various issues. A simple histogram of the issues seen across a channel group by frequencies would be an indicator of interference issues, or other plant issues, on those sets of frequencies.

The figure below shows the simple proof of concept in which we are building a cable plant with modems associated with the right hierarchical elements in the network to topology. Once those issues in each of the modems are identified, we start looking at issues across groups of modems and present in various views the issues overlaid on the network map.




Figure 15 – Issue visualization, based on anomaly detection and pattern recognition

Conclusion

This paper presents the machine learning approach to automatically identify issues in the cable plant. Machine learning classifiers can quickly identify issues on new data samples based on previously seen anomalies. Starting from PNM data such as RxMER Data, we can identify anomalies which we use to train a Machine learning classifier. Once trained a neural network can reliably identify anomalies in new test data. These anomalies form a base layer of information about the plant. Consolidating data across multiple modems gives the operator a very good view into the various issues affecting the cable plant

Acknowledgements

Many thanks to Paul Schauer (Comcast) and Nader Foroughi (Shaw) for sharing data captures from their DOCSIS networks. Without the data they so graciously shared, this analysis would be based on synthetic data, which would be only be -3dB as much fun.

Abbreviations

bps	bits per second
CM	Cable Modem
CMTS	Cable Modem Termination System
FEC	forward error correction
HFC	hybrid fiber-coax
Hz	hertz
RxMER	Receive Modulation Error Ratio

Bibliography & References

Applications of Machine Learning in Cable Access Networks, INTX 2016, Karthik Sundaresan, Nicolas Metts, Greg White;(CableLabs;) Albert Cabellos-Aparicio (UPC BarcelonaTech).

A Comprehensive Case Study of Proactive Network Maintenance, SCTE 2016, Larry Wolcott, John Heslip, Bryan Thomas, Robert Gonsalves

Scaling Learning Algorithms towards AI, 2007 , Yoshua Bengio and Yann LeCun,
<http://yann.lecun.com/exdb/publis/pdf/bengio-lecun-07.pdf>

How to Prepare Data For Machine Learning, by Jason Brownlee <http://machinelearningmastery.com>

An Intuitive Explanation of Convolutional Neural Networks, 2016 by Ujjwal Karn.
<https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/>

Filtering and Smoothing Data <https://www.mathworks.com/help/curvefit/smoothing-data.html>

[DOCSIS PHYv3.1] DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I11-170510, May 10, 2017,Cable Television Laboratories, Inc.

GIS A Success Story - Facilitating a Customer Journey from Design to Connection Using GIS

An Operational Practice prepared for SCTE•ISBE by

Erin Hayes

Director of Corporate Construction

Midco

3901 N Louise Ave.

Sioux Falls, SD 57107

605-274-2932

erin.hayes@midco.com

Introduction

Our industry is undergoing a massive digital transformation that is evolving in the science of where and who our customers are. In order to begin this digital era within our outside plant data, Midco started by converting from an archaic AutoCAD/Lode environment to a GIS platform. The conversion to GIS technology allows Midco to fulfill customer serviceability inquiries more efficiently and optimizes network design for future customers. Midco was able to simultaneously convert to a GIS platform and bring coax design functionality in-house, resulting in lower design costs and faster design turnaround time.

For several years, Midco outsourced all system design, at significant cost, to perform these design and as built functions. Midco wanted to utilize an alternative method to design which would allow us to bring this functionality in-house. The structure enhanced facilitation of our core GIS fiber management system, through development of RF/AC Power Design, ICOMS Integration software and conversion.

Content

Geovisualization allows you to view and map the location of objects, structures, specific street addresses, latitude and longitude data, varying devices and can be accurately situated on a map using geocoding methods.

Patterns are often more clearly observed when viewing mapped data. GIS provides a very effective means for graphically conveying complex information. Layouts created with a GIS are extremely useful when included in reports and presentations.

Midco has been working in this environment since 2008 when we purchased a Fiber Management System. With a touch of a few keys, our software and use of Esri (Environmental Systems Research Institute) ArcGIS applications have allowed us to map, and model every fiber cable, buffer tube, fiber strand and circuit. This data allows us to view and identify any type of usage throughout our Northern Plains Network.

Esri is a software development and services company providing geographic systems software and geodatabase management applications. Midco currently utilizes ArcGIS Desktop, ArcGIS Server, ArcGIS Portal and several other extension licenses. Schneider Electric's ArcFM, Fiber Manager and Designer Hybrid Fiber Coax (DHFC) product suite extend functionality of the core Esri platform to provide specialized tools for telecommunications companies.

These tools have permitted us to quickly respond with accuracy to several large scale Request for Proposals and expansion of residential and commercial areas while applying cost analysis to determine hotspots, as well as payback.

Our GIS database links all of our organization's digital data together based on a location, such as addresses, customer and network data of all capacities. This enables all departments of our organization to

have access to, and share the same data, and ensure all departments and individuals are using the most up-to-date information. Enhanced access to better quality and time-relevant data helps our organization discover new patterns in our data and the ability to make crucial business decisions.

Midco's GIS designer tool, DHFC allows us to provide instant feedback to the designer with clear understanding of what parts of the design need attention. DHFC's knowledge of the component catalog and infrastructure allows for extremely simple workflows with minimal clicks for design and redesign. All of this is driven by a centralized specification catalog where the specs drive the calculation engine on the client. The architecture allows for better ROI of the infrastructure being designed.

We have many home grown systems at Midco, most of which we have already integrated into our GIS environment. There have been many wins for Midco such as exciting spatial analysis which allows us to make smarter financial decisions at a much faster pace. Our design turnaround is one week versus 30 days and we're working in a single enterprise platform for engineering, network, GIS analysis and viewing of outside plant.

Midco is blazing the trail in GIS. The next beneficial step was to have the ability to provide our Account Executives instantaneous Good Faith Estimate (GFE) quotes for business sales. GFEs used to take 48+ hours, giving our competitors the upper hand. The Service Qualification App affords us the ability to provide automated quotes while in front of the customer which will significantly improve our close rate. Development of this application was in line with our on net, near net strategy. The application provides an in depth analysis of an area giving us the ability to capture additional business opportunities.

Customer Journey

Midco built a solution that automates the process of creating GFEs (Good Faith Estimates) that is required for Midco's service qualification workflow. The application was architected to integrate with our existing Esri ArcGIS for Server and Portal platforms. No additional investment was required other than development of the Serviceability App.

The Fiber Serviceability App captured additional on net, near net business opportunities. The application has the ability to provide a single cost to the business requesting service and/or estimate costs of building the business area. Previously GFEs are performed manually by designing the business area in-house then providing a bill of material. The Fiber Serviceability App performs these functions, saving time and resources.

In 2017, we've experienced a 35% increase in GFEs which made Midco's model even more favorable.

Originally there was a misconception that the new Master Address Database (MAD) provided the same information as the Fiber Serviceability App. MAD currently contains fiber pricing information about business opportunities; however, MAD fell short in providing reliable GFEs for a number of reasons:

- **MAD** – Blind geocoding addresses and unable to ensure the address is getting placed in the proper position.
 - **Service Qualification App** – Allows the user to visually confirm the location.
- **MAD** – Only accounts for underground routes.
 - **Service Qualification App** – Has the intelligence to determine underground or aerial routes. Aerial is a less expensive option.

- **MAD** – Limited due to prepopulated addresses and cost calculations; therefore, if an address is not in MAD, it will not have an address associated.
 - **Service Qualification App** – When an address is entered, it instantaneously provides the fiber cost calculation or allows the user to manually place the point.
- **MAD** – Unable to determine multiple route choices due to prepopulated addresses.
 - **Service Qualification App** – Captures multiple route choices which allow us to visually see the best route that will pass additional business opportunities.
- **MAD** – Unable to automatically create a PDF map.
 - **Service Qualification App** – Allows a user to create a pdf for upload into CRM (Customer Relationship Manager).

The Fiber Serviceability App provides us instantaneous quotes and a better close rate. Not only does this benefit our future customers, it also saves time, resources and ability to visually see the bigger picture.

Connecting our Customers Using GIS

An enterprise GIS provides broad access to geospatial data and applications throughout the organization. The advantages to deploying an enterprise GIS include:

- Using a common infrastructure for building and deploying GIS solutions
- Extending geospatial capabilities to an enterprise community
- Improving capabilities of other enterprise systems by leveraging the value of geographic information
- Increasing overall operating efficiency using GIS across our organization
- Enable viewing of geospatial information via mobile devices in the field
- Integration with other enterprise systems

GIS experts maintain control of the information and applications, yet productivity skyrockets as more users have access to geospatial information. Geospatial information can also be integrated with other enterprise applications to geoenable executive analysis and decision-support systems.

Midco delivers digital maps online, replacing paper map requests and quickly realized how much easier it is to click a check box to populate a basemap with data rather than fumble with existing AutoCAD maps to get the same result. Through training and partnering with our other departments, we were able to adapt our users to a more efficient digital workflow.

Midco understood our Strategy needed to be robust to accommodate our many field technicians. Our field requirements were:

1. Must be mobile
2. Easy to use
3. Quickly find features on map
4. View all records based on type of technical review necessary (installers, technicians, PM engineers, construction teams, etc.)
5. Trace RF and fiber
6. Make notes and drawings that can be saved to the database
7. Be able to view aerial photography
8. Be able to integrate GPS tracking

Geovisualization is one of the most powerful features of GIS and can be used to improve departmental collaboration. The overall benefits of spatial analysis:

- Enhance reports benefiting all departments
- Visualization aspect of data display
- Analyze data that will be interactive
- Capturing RF levels at the tap will allow us to identify plant problems before a truck roll is required
- Interaction between the NOC, dispatch and Field Operations will ensure plant performance
- Difference and Heat mapping will display customer preference and spending patterns
- Customer data will provide us the opportunity to upsell our products
- Customer call volumes, by Region, will visually identify concerns and allow Customer Service and other teams to respond accordingly
- Trouble call reports will be thematically displayed identifying other potential patterns of concern
- Any data the Call Center currently analyzes can take advantage of the system
- Other areas that will benefit from completion of this project included NOC software, Facilities, Capital and Operating performance, by Region, amongst many others
- Further investigation and development through APIs will allow our fiber management system to interact with network information
- Centralized AV system, providing a quicker customer response

Conclusion

Through the use of GIS, Midco has the benefit of an in-house mapping scheme and is no longer dependent on outsourced labor. After our initial investment and implementation, costs for mapping have decreased substantially and provides an efficient, accurate, streamlined process for all users all while delivering services to our customer base expeditiously.

Abbreviations

GIS	geographic information systems
FM	fiber manager
DHFC	designer hybrid fiber-coax
Esri	environmental science research institute
ROI	return on investment
GFE	good faith estimate

CRM	customer relationship manager
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

The New CX Standard: Location Data-Based Models for Driving Cost Savings and Improving Customer Satisfaction in Field Service Customer Journeys

An Operational Practice prepared for SCTE•ISBE by

David Troll

Senior Vice President, Sales and Customer Operations

Glympse

1424 11th Ave. #300

Seattle, WA 98122

(216) 334-9454

david@glympse.com

Introduction

Face-to-face field service interactions are the most critical step of the cable customer's journey. Because these occur in a customer's home, they are intimate, memorable and often the most powerful reference a customer will have by which to measure a brand. However, field service interactions are historically viewed by customers and field service technicians alike as frustrating, high-friction events - primarily due to their uncertain and unpredictable nature.

Successful field service interactions can launch long-term, loyal customer relationships and increase employee satisfaction and retention. But one mistake can create a cord-cutter, and frustrated customers can likewise drive costly employee churn.

Field service success, in this context, means delivering phenomenal customer experiences. Doing so requires 1) equipping both field service technicians *and* customers with access to high caliber tools and technology *and* 2) empowering them to effectively engage with one another during the appointment journey. That's not to say that only perfectly executed field service appointments are successful. It's impossible to account for every variable and thus impossible to give all customers precise information (service windows, arrival time, completion time, etc) and guarantee results.

Rather, successful field service engagements require humanizing two-way interactions and multi-party coordination in order to respect everyone's time and emphasize the "service" component. More simply, customer and employee satisfaction increases if you make a reasonable commitment, keep everyone informed, and meet the commitment or communicate issues early. Traffic patterns and human error will render any estimated arrival time obsolete, for cable technicians or taxis and ride-hailing apps, but continuous updates and real-time visibility can make even that experience positive. Likewise, advance communication and flexible customer reschedule options can remove the technician's anxiety when approaching the door of a subscriber who has waited hours past their expected appointment time.

Cable operators that excel in delivering insight-driven, customer-centric field service don't just get happy customers and employees: these multi-system operators (MSOs) also reap revenue and operational benefits as a direct result of deeper engagement between customers and technicians. Customers that are fully informed about a technician's arrival are:

- less likely to call customer care asking "Where's my tech?"
- less likely to abandon the appointment, requiring another truck-roll or losing the subscriber completely
- more likely to facilitate the technician's quick access, resulting in faster completion times

- eligible to be billed immediately following an installation (accelerating or retaining revenue that would be rescheduled or lost)

A location-based approach to customer experience (CX) in field service delivers concrete financial impact.

Content

1. Learning Outcomes

The described customer-oriented approach to field service can be achieved by leveraging mobility, on-demand location data and popular consumer devices and engagement channels. Cable operators can simultaneously create satisfied, long-term customer relationships and drive significant savings by focusing on two best practices.

- I. Effectively coordinate *where* and *when* the people and resources required for a successful field service engagement will be.
- II. Communicate this critical *where* and *when* information to stakeholders holistically - before, during and after a field service event, *and* via each subscriber's preferred communication channel(s).

This operational paper demonstrates how to leverage location-based data to build an “Uberized” view of field service - a digital, interactive experience focused on the customer and built with their needs in mind. This live insight improves customer and employee experiences in field service operations, and simultaneously drives revenue growth and cost savings.

2. Customer Experience Market Landscape

Today's consumers have more options and less brand loyalty. Personalized experiences, transparency, responding appropriately in emotionally charged situations and just knowing when to be front-and-center in a crucial customer service moment versus when to take a back seat, are critical differentiators for service organizations.

2.1. Customer Experience Must Be a Priority

Most companies understand that customer experience should be a priority. However, many don't understand what it truly takes to build a customer-centric business model, don't have the resources to fund necessary improvements, or simply can't keep up with the rapidly changing demands of consumer and business customers.

2.2. Customer Experience is Inconstant

Consumers are continually redefining their standards for “good customer service.” In part, this is based on the new norms established by emerging consumer technologies. Other times, competition can disrupt the status quo and encourage customers to take a more critical look at the level of service they should expect from the brands with which they do business. Today, competition extends far beyond direct competition from other cable operators and even over-the-top (OTT) providers. Cable providers are service providers;

they are being compared to every company with a product or service that's delivered to a customer's home or business.

Customer service channels continue to morph and fragment at an astounding pace. Ten years ago “omni-channel” meant phone + email and maybe a website FAQ. Yet some organizations still face significant challenges in effectively executing on those three basic channels. Single channel self-service will soon be obsolete, as customers will expect robust self-serve capabilities that fluidly transition across any device or channel throughout their journey and even during “independent” service interactions.

2.3. Customer Experience is Future-Oriented

Industry experts point to several shifts in the future of customer service.

2.3.1. CX Will Become “Micro”

In its “2017 Predictions: Dynamics That Will Shape The Future In The Age Of The Customer,” Forrester Research predicts: “CX will go micro to design signature moments that win the hearts, minds, and wallets of customers.”

No matter how carefully the complete customer journey is crafted, a customer's loyalty may come down to a single interaction. It will be important to understand which moments in the customer experience journey are likely to be the most visible, most anxiety-inducing or represent the greatest opportunity for additional revenue capture - and craft a high-value and/or frictionless experience around that moment of truth.

2.3.2. Digital CX Strikes a Balance Between Continuous and Unobtrusive

“Digital experience and engagement will draw people into nonstop virtual interactions,” according to Gartner's 2017 strategic digital disruption predictions.

While the micro CX trend requires preparation around likely moments of truth, delivering continuous digital CX requires giving customers a constant, real-time stream of contextually relevant information about their interactions with you. The trick is to make data available without overwhelming customers, so they can choose the most relevant moments to engage. Customers will want to interact periodically, at their own discretion. While they may appreciate occasional nudges to check back for significant updates, interruption won't be tolerated.

2.3.3. Self-Service Morphs Into Proactive Service

Originally designed to deflect costly phone calls, self-service is now the *preferred* channel for many customers for simple inquiries. As previously mentioned, standard self-service portals have become status quo and people often expect self-service on any channel - a smart phone, social media app, smart watch or a chat bot on a Slack plugin.

For businesses already offering multiple self-service options, the next level is predicting and proactively solving probable issues. Proactive communication with customers about activities that impact them directly can be a powerful first step toward proactive service. If you can anticipate customer needs and keep them informed about how you will - or how you already have - solved a problem, it's a big win. Just remember: customers want proactive service, not intrusive service. There's a massive difference between

making a timely recommendation for a plan that better fits a customer's consumption and usage habits and bombarding someone with ads for irrelevant upgrade offers.

2.3.4. *Messaging Becomes Oxymoronic: Automated and Humanized*

Interacting via popular messaging applications is not new, but continued consumer adoption will require an enhanced customer experience via these “concise” channels. Text-based updates are good. Interactive, multimedia experiences embedded within popular channels are better. Being present where customers spend their time, and engaging them in meaningful two-way conversations is the only way to win the service race. It will be critical to build interactive digital experiences that not only inform, but also allow customers to change elements of a service interaction that don't fit their preferences, or that just are not feasible - such as rescheduling an appointment when soccer practice gets changed. The alternative is to speak loudly and definitively through a virtual megaphone, and hope customers are still listening on the other end.

2.4. Customer Experience is a Competitive Differentiator

In the “good old days” the technology platform required to deliver residential video, data and voice services was expensive, as was the creative process that produced the content that drove adoption of these services. That meant a huge barrier to entry for challenger brands, leading to relatively few over-builders. Today the world is flat: the digital revolution has democratized access to the services and content, as well as the tools of production. Anyone can make a video and reach millions of viewers via YouTube or Twitter. So while a fast pipe into homes and businesses is an advantage, it's no longer *the* advantage.

2.4.1. *Incumbent or Innovative*

Even worse, those assets require continuous care and maintenance, forcing incumbents to focus time, energy and capital on their existing business. As noted by Clayton Christensen in *The Innovator's Dilemma*, “The very decision-making and resource- allocation processes that are key to the success of established companies are the [ones] that reject disruptive technologies.” Meanwhile, new entrants don't have these same legacy advantages...or constraints. Many will fail because the high cost of entry, but those that succeed do so because they solve things differently. *They innovate!*

Taxi and car-rental companies maintain vehicles and train employees, while Uber and Lyft use social feedback to ensure a quality experience for the riders and drivers. Hotel chains refurbish rooms and invest in rewards programs while Airbnb has a near limitless supply of inventory but minimal overhead (1 percent of the employee-base of Hilton or Marriott and no hotels). *Service is the primary differentiator*, and the digital disruptors have led the way in this.

2.4.2. *Apples & Oranges is Apples & Apples*

This high bar permeates every industry. Customer service in cable is not compared with other cable providers (of course) or just telcos and satellite companies. The modern, digital consumer compares every experience against “all service providers.” The iPhone interface affects how consumers perceive the latest release of an electronic programming guide, down to the font choice. Netflix recommendations on what to watch or Amazon “people who bought X also bought Y” suggestions set the expectation that everything comes with an expert opinion and curated recommendation.

The *experience* has become a central part of the product or service. Brands are defining themselves based on the experiences they deliver. Consider home furnishing retailer West Elm, who launched a line of hotels to extend the experience customers have with its products. The new paradigm is centered on data-driven, predictive and personalized user experiences. Consumers don't confine their judgement of a good customer experience to each industry, which means you can't afford to only measure yourself against your direct competitors.

If you are not measuring yourself against legendary service leaders like Nordstrom's, Amazon, Disney, and Trader Joe's, you are not looking through your customer's eyes. The pay-TV and ISPs who achieve a 64-65 industry average ACSI rating should look to the 83 of Breweries and Internet Retailers, or the 87 of TV and Video Players for inspiration. Maybe you won't win hearts and wallets with 2-day free shipping, but there are lessons to be had.

2.5. Customer Experience is Challenging

Delivering outstanding service has grown significantly more complex.

2.5.1. Field Service Challenges

A good field service experience (for customers) is dependent on two things: whether they are informed (and thus, empowered) and whether their need was met during the first visit (first call resolution). Even with sophisticated automation and optimization, field service is inherently unpredictable. It's impossible to provide a precise arrival time, every time. And the farther out you are from the appointment window, the more things can change. Assigning the right technician with the right skills and parts won't prevent unforeseen complications. And even when the field service technician arrives within the service window and is equipped to do the job, there's still a risk as to whether the customer will be there and prepared for the appointment. Success depends on continuously engaging customers so they are invested, and so they trust and respect their service provider enough to be around for the appointment.

2.5.2. Cable Industry Challenges

Beyond the standard field service complexities, cable operators face several unique barriers. First and foremost, the persistent "cable guy" reputation can create a negative experience through confirmation bias. A customer expecting a late, unprepared, poorly equipped and unprofessional technician is sure to find one of those things. Particularly in cable, where pop culture has transformed the frontline service professional into a villain, this can be detrimental to retaining skilled technicians. Adding to this legacy of poor service is the challenge of outdated technology systems - often siloed, overly customized and not user-centric. Technicians struggling to find accurate customer service history or get accurate home diagnostics are starting at a disadvantage. Again, it's much harder to retain technicians who don't feel empowered with modern tools to help them be effective on the job (especially when they compare those tools to their Galaxy S7s and iPads).

Despite these challenges, it *is* possible to provide customers with a positive experience, and to respond quickly when things don't go smoothly. The goal should be to build a customer engagement model using data that's not only common across these disparate systems and people, but which will continue to be relevant as customer preferences and technologies morph and change. This paper explains how **location** data plays a critical role in that successful approach.

2.6. Historical Approach

Cable operators and software vendors alike have attempted to solve the core challenge – colloquially called the “waiting without knowing” or “where’s my technician?” problem - in several ways.

2.6.1. Workforce Expansion

“Throwing bodies at the problem” remains a popular approach. More people can do more work, be on-time more often, etc. But adding in-house technicians is a commitment, and securing contract labor requires planning and effort. In either case, people are expensive.

2.6.2. Routing & Scheduling Optimization

Optimization creates both field and back office efficiencies that, theoretically, create more accurate service windows and better customer-facing communications. An efficient workforce arrives on-time and prepared more often. But many things can undermine this including adoption curves, compliance gaps, and as noted above, the dynamic nature of field operations. Without a way to continuously communicate updated appointment information, there’s still a risk of a customer-technician miss.

2.6.3. Statistical and Predictive Modeling

Likewise, some technological approaches rely on statistical analysis and complex modeling to provide a much narrower service window and more accurate ETA. This can improve scheduled windows (to a point) depending on the accuracy of the historical inputs and appointment details. But this puts the field organization on the hook for a tighter schedule with less margin for error. And without trust in the service provider, and effective two-way communication channels, customers may not trust the prediction, leading to frustration, inbound calls or abandoning the appointment all together.

2.6.4. Call Center Automation

One tactic is to wait until customers reach out, then divert inbound “Where’s my tech?” calls to an automated answering service. Some service providers even offer status updates via a self-service portal. This is more cost-effective than having live-agent support handle calls, but it requires significant upfront deployment costs and internal change management to ensure new processes are adopted. It can also frustrate customers.

2.6.5. Proactive Customer Notifications

“Don’t call us, we’ll call you.” Outbound reminder calls, emails and text messages can help inform customers. However, adoption of (high-attention) SMS is still very low while other channels suffer from low contact rates. Transactional notifications do not foster rich engagement, while only reminding customers of the original service window does not add value. Even “on the way” alerts can often surprise customers with little advanced warning, or fail based on poor field compliance.

3. Opportunities

This paper proposes that MSOs leverage location-based data as the foundation for building rich digital customer journeys that improve awareness, engagement and satisfaction throughout the end-to-end field

service lifecycle. Case examples demonstrate this approach results in improvements to four key areas: customer satisfaction, customer engagement, operational efficiencies and revenue growth.

3.1. Key Performance Metrics

Not every performance metric will be relevant to every business model. Following is a list of common metrics that have been used to measure the impact of location-driven customer experience journeys for field services. Some overlap and some are complementary. Choose three to five that are realistic for you to impact and measure closely, and which will serve as the best indicators for whether you'll achieve your desired outcome (e.g. customer experience, revenue generation, cost savings).

3.1.1. Customer Satisfaction

3.1.1.1. Customer ratings

- **Net Promoter Score (NPS):** indicator of long-term or transactional customer happiness based on “willingness to recommend” question; typically measured in absolute value (-100 to 100) and progress is seen based on point increase
- **Customer satisfaction (CSAT):** indicator of short-term customer happiness based on sliding scale responses of satisfied and very satisfied; typically measured in absolute value (0 to 10 or 100) with progress seen based on point increase
- **5-star feedback rating:** measures a customer's satisfaction with a given service interaction, notably the location-driven digital journey for field service engagement

3.1.1.2. Retention/customer churn

The ultimate measure of customer loyalty, or lack thereof; service providers should measure year-over year (YoY) or month-over-month (MoM) retention of subscribers/households, RGUs, \$/£/€ etc.

3.1.1.3. Response speed (to feedback)

Although this is a driver of satisfaction and not a measure of it like the items above, minimizing your response time to issues (days out booking on repairs, follow-up on low ratings or complaints) will likely correlate with customer sentiment.

3.1.2. Customer Engagement

3.1.2.1. “My Account” activity

Measures user engagement with a service provider's web-based portal or customer service application.

3.1.2.2. Customer contact information (provided or updated)

Demonstrates a customer's willingness to share personally identifiable information (PII), useful for future outreach.

3.1.2.3. Notification opt-in

Signals a customer's trust in your brand based on their willingness to receive continuous updates regarding service or appointment status.

3.1.2.4. Engagement by appointment stage

Provides insight regarding when customers are most interested in appointment-related information, and can be analyzed to identify important customer "moments of truth" on which to focus.

3.1.2.5. Customer-initiated communication

Additional chat and phone requests represent initial cost increases, as does deploying automation and self-service tools, but ultimately these drive savings by offsetting wasted truck rolls due to customer no-shows, deflecting expensive calls to customer care, and preventing revenue loss due to new customer abandonment.

3.1.2.6. Customer no shows

Directly impact re-roll costs and revenue attrition; this can also indicate a customer's lack of respect for your company/brand.

3.1.3. Operational Efficiencies

3.1.3.1. Inbound "Where's my technician?" calls

Decreases result in savings multiplied by the cost per call to customer care; some service providers also measure improvements in average handling time (AHT) based on CSRs having ready access to field service information to share with customers.

3.1.3.2. Customers not at home (reschedule for another day)

With a misaligned customer schedule, technicians waste time and costs go up with a second truck-roll; improvements here directly impact operating margins and speed time to revenue (for installation and upgrade appointments).

3.1.3.3. Suspended appointments (same-day returns)

Frequently due to a missing part, skill or customer schedule misalignment, this increases costs due to repeat truck-rolls and decreases first-time completion rates, imposing a second visit on customers for a single issue.

3.1.3.4. First-time installations or resolutions

A key indicator of operational efficiency as well as a high correlation with customer satisfaction scores, it also delivers faster time to revenue.

3.1.3.5. Task execution time

Decreases resulting from available and prepared customers (no re-rolls, no waiting for access) creates efficiencies resulting in field productivity gains.

3.1.3.6. Daily completions

Increases result from faster task execution time and ability of technicians to troubleshoot for missing parts, equipment and expertise in the field.

3.1.3.7. Service window met

One of the best “all in” metrics indicating operational health, this often correlates strongly with customer satisfaction.

3.1.4. Revenue Growth

3.1.4.1. Cancellations (lost revenue)

Reducing new customer cancellations directly prevents lost revenue; making narrow appointment window commitments and meeting them helps prevent new subscriber abandonment.

3.1.4.2. Rescheduled appointments (delayed revenue)

Increases can drive short-term cost increases but can prevent cancellations and unrecoverable revenue.

3.1.4.3. Time to revenue

Accelerates in relation to higher first-time completion rates.

3.1.4.4. Missed service level agreement (SLA) penalties

Impact profitability but can be minimized through improvements to operational performance as previously described.

3.1.4.5. Expansion/upsell success

Many studies have shown that growing with an existing customer is much easier than acquiring a new one; increasing the spend of your established customer base is a good indicator of whether they are satisfied with your current service, and you can take advantage of customer engagement with the digital customer experience to surface timely offers.

4. Approach

Don't rely only on pre-planning, reactive re-optimization or outmoded methods of customer communication. This approach assumes that everything will *not* go as planned. This paper shows how to use location data and location sharing to first coordinate the arrival of all the necessary people and equipment at the service location, and then create an interactive digital experience to inform customers and field service teams alike with continuous updates and real-time insight.

4.1. Overview of Steps

At the highest level, there are four core steps in this approach.

- I. Collect the necessary data to understand WHERE all required people and resources are as the impending field service event approaches.
- II. Leverage this contextual location intelligence to assess WHEN the required elements will arrive at the service destination.
- III. Keep everyone informed, starting early and ongoing.
 - a. The primary focus should be on keeping customers informed
 - b. Back office staff like customer care representatives and dispatchers can also benefit from this same insight to ensure maximum operational efficiencies (when to reassign, send help, who to send, etc.).
 - c. Field supervisors can use this same information to provide onsite assistance or conduct spot-checks as needed.
- IV. Enable continuous bidirectional customer updates with an omni-channel approach, so they can choose how they engage with you.

4.2. Required Resources

Resources required to successfully leverage location data to deliver an Uber-like customer engagement include people, places & things (data, devices and applications). It's precisely the mobile and distributed nature of these required resources that make them ideally suited building blocks for developing location-driven customer experience journeys.

4.2.1. *People, Places & Things*

4.2.1.1. *Customer(s)*

Inherently mobile despite appointments assigned to their stationary home or business; connected via many web-enabled devices.

4.2.1.2. *Field service technician(s)*

In-house or contractor employees tasked with completing a set of appointments (installations, maintenance, repairs, trouble calls, etc.) each day; want to use the same tools at work that they use at home.

4.2.1.3. *Customer care representative or dispatcher*

Tasked with creating service work orders and assigning them (sometimes with help from an automation solution) to field service technicians; also typically fields inbound customer calls regarding appointment status updates.

4.2.1.4. *Service garage and/or parts depot*

Facility where field teams acquire the vehicle, equipment, parts and inventory required for the day's (or several days') work; may also include ancillary facilities that house spare parts and inventory slated for return.

4.2.1.5. *Parts and inventory*

Bulk, in truck stock or assigned to a customer or location; also, might be distributed amongst teams and depots.

4.2.1.6. *Vehicle*

Method of transportation, mobile storage and frequently a hub of power, telemetry and sometimes technician connectivity.

4.2.2. *Data*

4.2.2.1. *Real-time location of primary field service representative*

Can be collected from a mobile client application installed on a technician's device (custom-built or off-the-shelf field service or fleet management application), telematics device installed in the technician's vehicle; specialized vehicle hardware (e.g. car topper, OEM navigation system, sensor, etc.), GPS data from a web-enabled device, etc.

Must be accurate and have negligible delay.

4.2.2.2. *[Optional] Real-time location of secondary field service representative(s)*

May be necessary to coordinate a part swap or on-site assistance to ensure a first-time completion.

4.2.2.3. *Service destination*

Geocoded address or latitude/longitude position of a customer's home or business, or a network component requiring service.

4.2.2.4. *Route and traffic data (and related corporate policies)*

Can be obtained from any mapping provider. However, special considerations for routing preferences should be considered in advance and factored into planning.

For example, service appointments requiring special equipment, or operations policy might call for right-side routing, in which the vehicle's right side must arrive facing the final destination.

4.2.2.5. *[Optional] Available and sourceable inventory*

Available to the technician and back office dispatch.

4.2.2.6. Other Optional Data

Data which is not required but which can contribute to a more robust digital experience for the end-customer

- Promised service date and time window
- Work order or delivery identifier
- Selected work order or delivery details (sender, weight etc.)
- Additional customer preferences (preferred method of entry, security clearance, etc.)
- Technician information (photo, name and unique ID #)

4.2.3. Devices

In this methodology, devices serve two primary purposes. First, they are used to collect the required data. The approach can work regardless of the device used to collect location data, as long as that data can be collected continuously and in near real-time. Common devices used include smartphones, tablets, vehicle-mounted fleet management hardware and sensors.

Second, determine which devices you expect customers to use to engage in the digital appointment experience. A web-based approach can empower customers with insight on any Internet-enabled device from a PC to a smartphone. For emerging devices including wearables, automobile head units, smart home electronics and even home IoT devices such as Amazon's Echo, you'll need to identify the appropriate API calls or embed the appropriate applications to receive and communicate technician location and ETA updates.

4.2.4. Applications

None of the following applications are required. However, these can be helpful in automating certain steps of the detailed procedure covered in the next section of this paper.

4.2.4.1. Field Service Management (FSM)

Actions completed within the FSM can be leveraged to automate proactive customer notifications and updates to the customer-facing digital experience.

4.2.4.2. Customer Relationship Management (CRM)

Actions completed in the CRM (e.g. creating a work order, booking a field service appointment, confirming a new appointment, or assigning a work order to a field service resource) can be leveraged to automate proactive customer notifications and updates to the customer-facing digital experience.

Additionally, the CRM can supply optional data such as service window, appointment details, customer address, etc.

4.2.4.3. Enterprise Resource Planning (ERP)

Useful for supplying accurate insight regarding parts and inventory, equipment, SLAs, warranties, etc.

4.2.5. Communications Methodology

Consider whether you will host the customer-facing “Uberized” experience independently, or as an embedded element in a customer self-service portal.

Additionally, determine a preferred method (or multiple methods) for communicating with customers to 1) set the initial expectation that they’ll have access to a new digital experience to track their pending field service appointment and 2) provide timely notifications when their technician’s location and ETA change significantly. Notifications might come as SMS, MMS, email, a My Account push notification or a combination of these. Customer contact information will need to be collected and/or updated and verified. Decide whether to manage these notifications internally or outsource to a third party.

Finally, determine whether it’s important to make real-time technician location available and searchable to field service team members, managers, dispatchers and customer care representatives - and if so, the best way to display this insight. For example, do you want location and ETA data made available to an IVR in order to provide updates to customers that call in looking for an update?

4.3. Detailed Procedure

The following describes how to apply the core steps to deliver insight about the status, ETA and technician location to a customer for their field service appointment.

Step One: Collect the necessary data to understand WHERE all required people and resources are as the impending field service event approaches.

1A: Secure a location data stream for the field service technician (from mobile device, field service management application, fleet tracking hardware, etc.).

1B: Associate location data with real-time route and traffic data.

1C: Scrub technician personally identifiable information (PII) to ensure their personal contact information and any non-standard customer-facing information is not shared without consent.

1D: If delivery of a spare part or assistance from a nearby team-member is required, acquire location data stream(s) for additional resources.

Step Two: Leverage this contextual location intelligence to assess WHEN the required elements will arrive at the service destination.

2A: Set the current location and appointment address as a journey endpoint.

2B: Calculate real-time ETA based on technician’s dynamic location.

2C: Collect new location data and update real-time ETA following preset intervals.

2D: If additional resources are dispatched (per 1C), calculate ETA for each individually, and as a group.

2E: Confirm (through location data source) the technician’s arrival onsite.

Step Three: Keep everyone informed, starting early and ongoing. The primary focus should be keeping customers informed, but back office staff like customer care representatives and dispatchers can also benefit from insight in order to ensure maximum operational efficiencies as well as to keep customers informed.

3A: Secure customer opt-in for receiving appointment related communications.

3B: Collect, update and/or verify customer contact information and preferences.

3C: Build and launch a web-based viewer containing modular components for displaying relevant appointment data throughout the complete appointment lifecycle. Baseline components should include an ETA countdown/estimator and a map view of the technician's location (which will start as an approximate location and progress to a precise live location for the assigned technician).

[The process for acquiring data to build this experience is described in steps one and two.]

Optional components include the ability to confirm, cancel or reschedule an appointment, add a pending appointment to a personal calendar, view technician photo and other relevant information, as well as service window or SLA, and work order details.

3D: Best practices indicate that brand-specific experience design as well as messaging prompting customers to check back for updates help achieve higher levels of engagement.

3E: The link hosting the web-based experience should remain consistent throughout the entire appointment lifecycle, though what's displayed on the web viewer should be frequently updated.

3F: Immediately following the creation of a new work order and/or appointment, send a confirmation notification to the customer. The notification should contain the appointment-specific link. At this stage, the web viewer can confirm the appointment date and time and include text explaining that a live view and ETA for the technician will be available on the day of the appointment. It's easiest if this action is prompted through integration with the CRM or work order management system, but can also be managed manually by back office staff.

3G: On the morning of the appointment, update the web viewer to display the estimated service window and language prompting the customer to check back for a live view and ETA closer to the start of the service window. Send another notification. It's easiest if the update and notification are triggered by the assignment of a work order to a specific technician in a dispatch or field service management console, but can also be managed manually for organizations using simple calendaring tools for dispatch.

3H: Upon completion of the previous appointment (or upon departure from the field service depot or warehouse), leverage location data stream and ETA calculations from steps one and two to display real-time technician location on a map, and ETA countdown in minutes. Send a third customer notification. It's most effective to launch the viewer update to "live" status through integration with a field service management application that can indicate when a field service technician has completed the previous appointment or has updated his status to "en route" or "traveling."

3I: Display arrival confirmation on web viewer once technician is on site.

3J: Leverage web viewer engagement and screen-space to request customer feedback regarding the overall appointment experience. For an advanced approach, use integrations with customer service and social marketing tools to enable notifications to customer care to reach out to unhappy customers or enable satisfied customers to share good feedback with their friends.

3K: Expose the web-based journey viewer to dispatchers and customer care representatives.

3L: Ensure technician compliance (for the most accurate location sharing experience) by securing their support early. Include technician leadership in planning sessions and leverage their expertise to help you determine the most accurate data sources.

3M: Consider including personalized advertisements as part of the digital appointment tracker to promote upsell and cross sell activity.

Step Four: Enable continuous customer updates with an omni-channel approach so they can choose how they engage with you.

4A: Establish and refine primary web-based viewer experience. At a minimum, customers should be able to engage with the experience you've built via any standard web browser on any standard Internet-enabled device.

4B: Research which non-standard channels are most commonly used by your customer base - and prioritize. For example, companies with app-heavy customer care approaches might consider how engagement with a My Account app might be increased if available on a smartwatch or automobile head unit. A tech savvy customer base may prompt you to explore translating this web-based experience into insight that's accessible through voice-enabled IoT devices or chat bot-driven messaging applications.

4C: Develop partnerships and research how APIs and SDKs can be utilized to push live technician location and ETA updates to these preferred devices and channels.

4.4. Practical Applications

The location-based approach for improved customer experience is demonstrably useful for the most obvious cable field service use case, in which a customer must be present and prepared at their home or business in order to complete a successful field service. However, this model can also be applied to two other use cases in order to increase first-time resolution, which directly impacts both customer satisfaction and operational efficiencies.

Use Case #1: B2C/B2B service event in which the customer must provide the technician access to the property and ensure all obstacles are removed (e.g. barking dog is put away or security clearance is provided to restricted zones).

Use Case #2: Technician requires additional talent/skills from a colleague or manager, or requires specialized equipment (e.g. an extension ladder) not currently in his or her possession.

In this use case, add an additional technician/truck icon to the web-based customer journey viewer to alert the customer that help is required. Use the standard process, but simply display both relevant resources in a single viewer.

Use Case #3: Technician requires a part to achieve first time resolution which he does not currently have access to in the vehicle or at the job site.

A recent Aberdeen Group report found the top reason for a failed service visit is parts unavailability. If this is a common challenge in your field service organization, consider modifying the location-based CX model to also support better sourcing of spare parts in type of case.

Start by logging all inventory (assets, parts, consumables) and making it discoverable based on its location or affiliation with a vehicle. This will provide back-office and technician visibility into all parts options. If you integrate your parts management system with your fleet management or field service system and enable parts lookup and ordering directly from the technician's field service application, you can achieve maximum efficiency.

Next, determine if you'll also allow the customer to track the delivery of the spare part(s). If so, and if the spare part will be delivered by a courier service or a colleague, simply add the additional resource's location and ETA as in use case #2. If so and the technician must pick up the spare part himself from a warehouse, parts depot or from a colleague's job site or vehicle, the approach will be different. You can display the pickup location as a new point of interest in the live map web viewer, and share the technician's location as he travels there prior to the final service destination as part of the "live" location sharing customer journey phase. Or, create an additional phase in the web-based customer journey that confirms the technician is currently "picking up a part" at an approximate location and adjusts the ETA countdown accordingly.

If you don't want to show this logistical coordination to customers at all, simply enable technicians to share location with one another (using any number of consumer location sharing applications). For this approach to deliver maximum customer satisfaction, technicians will still need a method for communicating parts-related delays to customers, managers and dispatchers.

5. Tracking Results

5.1. Recording of Results and Analysis Best Practices

Prior to launching your new model for location-based customer experience, select the KPIs most important to your organization. As outlined in the Key Performance Metrics section of this paper, this model can generate results across four key categories: customer satisfaction, customer engagement, operational excellence and revenue growth. Focus on the metrics that will best support your business case, or which align most closely with your corporate goals.

Certain performance metrics can indicate success in multiple areas. For example, an increase in first time completions will often translate into increased NPS as well as operational savings calculated by multiplying the average cost for a truck roll by the reduction in required re-rolls.

Next, set benchmarks for your selected KPIs. This will be key for demonstrating results. If you don't have access to benchmark measurements for your core KPIs, it's important to define success prior to launching your new customer engagement model. Use a combination of industry averages and aspirational targets. For example, if part of the new initiative is to also start measuring NPS for the first time in the history of

your organization, work with your team to develop a realistic goal for the first week, month and year. Plan for fluctuations.

Don't expect to see significant improvement overnight. Establish a cadence for recording results and assessing change to demonstrate continued progress. Measure customer engagement with notifications, the web-based interactive experience and live map viewer by stage for each appointment - but also assess whether engagement levels are trending up or down from a week to week and even month to month basis. Likewise, record whether location was successfully shared for each appointment, as this can be a key factor impacting many common KPIs.

For metrics like the number of appointments completed overall, average job completion time or number of appointments completed on time, record daily results and monitor incremental improvements. Be sure to account for seasonality and other initiatives that may impact your results. Savings from inbound call volume deflection and revenue growth generated from prevented cancellations can be impactful on a weekly basis. Likewise, some organizations have been able to assess results and identify areas for improvement with weekly analysis of NPS. Overall cost savings (resulting from re-rolls and operational efficiencies) are most effectively measured and compared on a month-to-month basis.

Customer satisfaction levels and anecdotal feedback should be measured for every service interaction - and you'll need a benchmark for how quickly you're able to respond to negative or highly positive feedback. A method for quickly recognizing negative feedback can help you identify employees who need skills training or service coaching. Don't forget to monitor overall improvement (or irregular dips) in customer satisfaction across the entire field service organization on a weekly and monthly basis.

5.2. Success Examples

Companies who embrace the location-driven model for digital customer engagement have seen near-instant benefits.

Table 1 – Examples of Results

Industry	Location	Solution	Benefit
Cable	Europe	Partial notification process w/iterative updates including a live map technician view	+14-points NPS lift (at peak) in 8 weeks
Cable	U.S. National	Complete notification process with iterative updates including a live map technician view	\$10s of thousands monthly in wasted truck rolls; increased monthly revenue (achieved by preventing cancellations and reschedules)
Furniture Retailer	Southern California	One SMS with a link to a live tracker and ETA countdown	10% drop in customer not-at-home

Results data provided by Glympse.

5.3. Troubleshooting

Though offering customers a live map view and ETA countdown for their field service appointment (continuously updated with strategic notifications) can deliver phenomenal benefits, you'll want to keep refining your processes in order to achieve maximum results.

The most common errors when building digital CX hubs centered on location sharing and tracking are:

- Live location sharing fails during the last mile.
- Customers don't engage with the location-based digital journey.

Consider the following methods to prevent or overcome these errors.

5.3.1. Live location sharing fails during the last mile.

Technician compliance is the primary factor impacting whether location is/is not shared (and thus, whether the ETA is correct) on the day of service. Involve technicians, their leadership and union representatives early in the process to secure buy in. Overcome privacy objections by demonstrating how technician contact information will be anonymized. Explain the benefits to field workers of a better-managed customer journey and how improved customers' perceptions will set-up technicians for success at the job site. Use this project to demonstrate your commitment to empowering and retaining technicians by offering them the most updated technology. Provide training and establish habits to ensure technicians are logged into the appropriate application and have the appropriate GPS settings enabled on their smart devices.

Technician non-compliance is not always a result of unwillingness to cooperate. If you're relying on a CRM or field service management application to trigger customer-facing location sharing, check to ensure technicians, care representatives and dispatchers regularly comply with the actions in those systems that will drive your new digital customer engagement initiative.

Technology can also be the culprit preventing location-sharing, so be sure to check the end-to-end tools and systems you are employing, from devices/carriers with poor connectivity, to deficient network components that fail to pass the necessary data, to poorly designed integrations that inconsistently update the required systems. There are many potential points of failure to consider and avoid.

5.3.2. Customers don't engage with the location-based digital journey.

Even though customers are familiar with digital journey viewers for ride sharing apps and simple food and package delivery trackers, this experience may be unexpected coming from their MSO. Secure opt-in, preferences and update customer contact information as part of the appointment booking process. Remember to set expectations early by telling customer's they'll get a live ETA and map view of their technician on the day of service, and remind them as the appointment draws closer. The most common reason customers don't engage is because they don't understand the potential benefit of doing so. Education is key.

Likewise, compliance among call center representatives can make or break an effort to explain the new location-sharing digital journey. Using systems (e.g. mandatory contact information fields) and processes

(e.g. tracking and reporting on sign-up compliance or incentives for top performers) to promote their role in customer awareness is a key step in driving engagement.

Others notice high customer engagement during certain stages in the customer-facing journey, but much lower engagement in other stages. You can use this insight as an opportunity to identify which areas of the holistic field service process may need improvement from both an operations and service execution perspective. Or, lukewarm engagement during certain stages in the digital journey might simply indicate where to scale back on your proactive communication and/or digital journey.

Conclusion

Location data and the application of that data to create interactive, location-sharing experiences can provide end-customers with critical insight about the status of a field service. Armed with a continuously updated ETA and confirmed by a live map view, customers are empowered to better manage their schedules and believe that you respect their valuable time. The result: customers are there, prepared and have a positive attitude regarding field service representatives, and technicians, improving the technician experience and odds of success. Not only can you use this approach to improve overall customer satisfaction and NPS, but a robust CX hub that enriches technician location sharing activity improves customer engagement with your brand, and through more efficient channels. Engaged, informed customers reduce customer care costs and improve field team efficiency, resulting in operational savings that can easily justify the cost of this and future customer service initiatives.

Abbreviations

B2B	business-to-business
B2C	business-to-consumer
CRM	customer relationships management technology
CSAT	customer satisfaction
CX	customer experience
ERP	enterprise resource planning system
ETA	estimated time of arrival
FAQ	frequently asked questions
FSM	field service management software
MMS	multimedia messaging service
MoM	month-over-month
MSO	multi-system operator
NPS	Net Promoter Score
OTT	over-the-top
PII	personally identifiable information
RGU	
SMS	short message service
SLA	service level agreement
YoY	year-over-year

Bibliography & References

Forrester, *2017 Predictions: Dynamics That Will Shape the Future In The Age Of The Customer*, <https://go.forrester.com/wp-content/uploads/Forrester-2017-Predictions.pdf> (accessed December 16, 2016).

Gartner, *Top Strategic Predictions for 2017 and Beyond: Surviving the Storm Winds of Digital Disruption*, <https://www.gartner.com/doc/3471568/top-strategic-predictions-surviving-storm> (accessed December 16, 2016).

American Customer Satisfaction Index:
http://www.theacsi.org/index.php?option=com_content&view=article&id=148&Itemid=213

Aberdeen Group, Field Service 2016: Strengthen the Team and Bond with Your Customers,
http://v1.aberdeen.com/launch/report/research_report/12665-RR-field-service-customer.asp

Christensen, Clayton. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business Review Press, 1997. Print

The New Customer Care Experience

Moving from Scripted Dialogs to Automation, Omni-Channel and Predictive Analytics

A Technical Paper prepared for SCTE•ISBE by

Marc Bellini

Vice President, Sales Engineering

Nokia

1380 Rodick Road, Markham, ON. L3R 4G5. CANADA

+1-416-821-4268

marc.bellini@nokia.com

Introduction

In the United States alone companies spend \$112 billion on call centers each year,¹ but only half of all customer service issues get effectively resolved. In response, customer care is evolving from a static model of ‘reactive customer care’ — associated with high operations costs — to an intelligent, automated, and predictive model that reduces costs and ensures that subscribers are happy with services that just work. Cable multiple system operators (MSOs) can make life easier for subscribers, by offering an omni-channel customer care solution that includes compelling and effective self-care tools and much-improved agent-assisted care.

The time is coming fast when customer care systems, leveraging humanoid interfaces powered by machine learning, artificial intelligence (AI), and bots become so advanced that we enter the era of autonomous care. To facilitate this transformation, MSOs need to embrace next-generation customer care technologies.

Interactive bots provide an ideal interface for customers experiencing common issues that have simple solutions, such as connecting to Wi-Fi, resetting forgotten passwords, or checking on the status of a scheduled technician appointment. For many MSOs, it is these simple issues that drive a large volume of help desk calls. Interactive bots can also be used to create powerful self-care and augmented customer care solutions, which power web-based chat or instant messaging tools, as well as behind-the-scenes support for customer service representatives (CSRs).

Proactive bots — combined with advanced, predictive analytics — can determine when a customer is experiencing connectivity issues. Rather than waiting for that customer to seek technical support, remedial action can be taken proactively. The customer issue is resolved automatically, without the need for any interaction between the customer and the help desk. This is known as autonomous customer care and it is the desired objective for both MSOs and customers alike. According to industry analyst firm Analysys Mason, autonomous customer care is the fastest-growing sub-segment of customer care, with sales forecast to reach 1.486 billion USD by 2020².

For MSOs, a reduction in the number of help desk calls that require human intervention reduces support costs and frees up CSRs that can be trained to handle more complex tasks or to provide premium technical support. For customers, it results in faster response times and a better customer experience.

¹ “Customer service of the future is powered by artificial intelligence,” Brandon Buckner, December 1, 2016; <https://www.ibm.com/think/marketing/customer-service-of-the-future-is-powered-by-artificial-intelligence>

² “Customer care systems: worldwide forecast 2016–2020”, Analysys Mason, 20 September 2016; <http://www.analysismason.com/Research/Content/Reports/Customer-care-forecast-Sep2016-RMA11/>

Using Technology to Change the Fabric of Customer Care

Intelligent virtual assistants (IVAs) — such as those found in interactive audio hubs, such as Google Home and Amazon Echo — are becoming very popular with consumers. These IVAs provide hands-free, instant access for making reservations, playing music, or controlling household lighting. As a part of the Internet of Things (IoT), these devices are always on and can be awakened with a simple one- or two-word greeting, followed by a verbal instruction. For example, “Alexa, what is the weather going to be like today?” or “Okay Google, how much is 100 Euros in American dollars?” Worldwide, Ovum forecasts that 192 million interactive audio hubs will be in use by 2021.³

These devices have also introduced consumers to artificial intelligence (AI), machine learning, and natural language processing (NLP) technologies.

1. Artificial Intelligence (AI): Increasing efficiency and minimizing customer frustration

Artificial Intelligence (AI) has the potential to transform customer care, by making processes more intelligent. AI refers to software algorithms designed to simulate human intelligence by thinking, reasoning, planning, predicting, learning, and solving problems. AI is getting a lot of attention lately owing to a convergence of a few different factors:

- improvements in computer processing power — a trillion-fold increase in the last 60 years;
- declines in the cost of processing data — thanks to cloud and virtualization technologies; and
- increased volumes of data — that needs to be analyzed quickly if it is to provide value.

Initial applications of AI include language translation programs (like Google Translate), eCommerce applications (like Amazon, which makes product recommendations) and IVAs (like Google Home and Amazon Echo). With a market size of \$100B by 2025 it is clear that AI is not just another fad.⁴

By combining AI with other technologies — including NLP and machine learning, both discussed later in this paper — powerful bots can be created and applied to customer care. Using AI technology, vast repositories of data can be analyzed, creating insights that can be used to deliver personalized services, power proactive care solutions, and empower ever-smarter CSRs.

2. Machine Learning: Welcome to the machine

At its core, AI is a series of algorithms that require human programming and, as a result, AI only knows what it is taught. The brains behind AI, and its continued advancement, is machine learning. Machine

³ “2017 Trends to Watch: The Digital Consumer Landscape”, Ovum, March 2017; <https://www.ovum.com/research/2017-trends-to-watch-the-digital-consumer-landscape/>

⁴ “Understand The Spectrum Of Seven Artificial Intelligence Outcomes,” by R “Ray” Wang, Software Insider, September 18, 2016; <http://blog.softwareinsider.org/2016/09/18/mondays-musings-understand-spectrum-seven-artificial-intelligence-outcomes>

learning allows algorithms and computers to learn from data. It is the science of giving computers the ability to learn without being explicitly programmed.

AI and machine learning are related concepts, but it is important to note that not all AI techniques use machine learning and that machine learning is used for other things besides AI, such as decoding genetic sequences.

Machine learning works with structured data to detect patterns that provide useful insights. Everyday examples are personalized recommendations from services like Netflix. In the context of customer care, machine learning can classify a subscriber's issue and intelligently present the best solution. Each customer issue that is processed contributes to the knowledge system, resulting in a more robust data set over time. This process of continuous improvement allows customer care systems to better classify issues and to route them more quickly and intelligently with each subsequent transaction. Eventually, machine learning allows the knowledge system to acquire more knowledge than any one human expert could ever possess.

3. Natural Language Processing: Harnessing the power of the spoken word

Natural Language Processing (NLP) uses AI to find patterns within large datasets to recognize language. One of the applications of NLP is with bots embedded in IVAs, which has introduced consumers to a 'screenless' user interface. Google recently announced that 20 to 25 percent of queries on its Android devices are voice searches. This is predicted to reach 50 percent — across all platforms — by 2020.⁵ Further, about a third of Amazon Echo users use the devices three times (or more) every day.⁶

As the accuracy of speech recognition reaches 95 percent (and beyond), the problem of comprehension evaporates and we begin to interact with bots as if it were a person rather than a device. Some of the benefits of using voice instead of other interfaces include:

- convenience: hands-free, instant access when hands are occupied or when focus is on another task, like driving or cooking;
- speed: on average, humans speak 150 words per minute but can type only 40; typing speed decreases further when using mobile devices; •
- novelty: many consumers find this new interface cool and exciting, but it needs to be accurate and reliable or this level of excitement will not be maintained; and
- reliability: according to Google, at the end of 2016, their NLP engine could recognize nearly 10 million words, with 90 percent accuracy.⁷

This hands-free interface is one of the elements that makes bots applicable to customer care. Instead of having to download, install, and access an application on a mobile device, visit a web site, or call a help

⁵ "Has voice control finally started speaking our language?", Rhodri Marsden, The Guardian, December 4, 2016; <https://www.theguardian.com/technology/2016/dec/04/voice-control-amazon-echo-digital>

⁶ "Voice Is the Next Big Platform, and Alexa Will Own It", Jessi Hempel, Backchannel, December 19, 2016; https://backchannel.com/voice-is-the-next-big-platform-and-alexa-will-own-it-c2cf13fab911?mbid=synd_digg#.8ohqs4369

⁷ "Mary Meeker: voice-controlled tech set for exponential rise in next few years", Danny Yadron, The Guardian, June 1, 2016; <https://www.theguardian.com/technology/2016/jun/01/mary-meeker-voice-controlled-tech-boom-technology-predictions>

desk and navigate a maze of structured voice menus using an interactive voice response (IVR) system, bots provide a nimble and consumer-friendly approach.

3.1. Natural language understanding (NLU): Determining intent

The term natural language understanding (NLU) is often used interchangeably with the term NLP, but NLU is an important subtopic of NLP that deals with language comprehension. While NLP lets people and machines communicate with each other, NLU is used to determine how to best handle unstructured inputs and to convert them into a structured form that a machine can understand and act upon.

When engaged in a conversation, humans are (for the most part) able to handle mispronunciations, contractions, colloquialisms and other idiosyncrasies associated with language. Machines are less capable of dealing with unpredictable inputs and NLU is used to improve a machine's ability to understand language. In fact, while NLP is reaching 90 percent accuracy, NLU typically struggles to achieve 60 percent accuracy. Understanding intent is more difficult than speech recognition.

NLP is sometimes used as an umbrella term that refers to the systems that work together to handle all interactions between machines and humans. An effective NLP system can process what is said, analyze it, comprehend its meaning and determine the intent, establish the appropriate action, and respond in language the user will understand.

Enhancing the omni-channel customer experience

One of the biggest trends in the modern call center is 'omni-channel' customer care. More than just a proliferation of ways that customers have to seek information or assistance — such as visiting a web site, using instant messaging/chat, sending email, engaging via social media, using custom mobile applications or calling the help desk — a true 'omni-channel' customer experience requires that each of the available channels be integrated.

If one channel doesn't lead to resolution, the customer must be able to access another channel without having to duplicate actions or re-enter information. An omni-channel experience eliminates the time needed to recapture lost information when moving between channels. By integrating bots, existing omni-channel customer care solutions are enhanced, providing further improvements to the customer experience.

4. Interactive Bots: Changing expectations in customer care

Unlike a CSR, a bot can analyze real-time and historical data before initiating a response. Interactive bots can provide an ideal interface for customers experiencing common issues that have simple solutions, such as resetting forgotten passwords or checking on the status of a technician appointment. For many MSOs, it is these simple issues that drive a large volume of help desk calls. For consumers, this prevents waiting in call center queues, navigating through Q&As on websites, or finding the right app on their smartphone.

Interactive bots can also provide behind-the-scenes support to CSRs; what is known as 'augmented customer care'.

5. Improved Reactive Care; Providing CSRs with an ‘extra set of hands’

When it comes to reactive, agent-assisted care, the art is in handling the escalation as fast and as accurately as possible. Here too AI comes increasingly into play to provide better diagnostics and accurate insights.

It can be difficult for CSRs to multi-task when speaking with a customer. Augmented customer care solutions utilize interactive bots in the background, to review network and device information, access a library of use cases, pinpoint customer issues, and present CSRs with resolution options – all in real time.

6. IVR Integration: Leveraging cutting-edge technology for improved problem resolution

When calling the help desk to report a problem, or to seek assistance, most consumers are hoping to speak with a live person. To start, however, most interactions begin with an IVR system. Most callers are frustrated with existing IVRs because they lack the ability to understand intent. One of the benefits of AI is its ability to sift through large volumes of unstructured data. One of the challenges, however, is for NLP to understand the literal meaning of customer requests. It is essential that NLPs determine customer intent.

In the context of customer care, intent represents the purpose found within a subscribers’ statement, question, or request; whether provided with a voice command or via a web-based chat or instant messaging tool. For example, when a subscriber asks “What is my guest Wi-Fi network password?”, the customer care system must be able to interpret the specific intent of that request. To do this, the phrase must be parsed and given structure:

```
need:guest {intent} / need:Wi-Fi {intent} / password {intent}
```

Then, the intent must be mapped to an appropriate remediation procedure or proactive action.

7. Empowering connected consumers to help themselves

Owing to the increasing amount of human-computer interaction, some people are more comfortable accessing information and seeking assistance using self-service portals. Customers feel empowered when they can resolve problems themselves.

Recognizing the importance of providing flexible support options, some MSOs have launched self-service care applications designed to help consumers troubleshoot and resolve issues themselves. Interactive bots can provide an effective interface for these tools, in addition to web-based chat or instant messaging tools, like the one embedded in Facebook.

Self-service apps have proven to be very effective, in some cases reducing the number of calls to the help desk by as much as 75 percent and reducing the number of truck rolls — that is, sending technicians to consumers’ homes to troubleshoot and resolve issues — by more than 25 percent.

8. Dynamic Intelligent Workflows

For reactive customer care, CSRs typically use guided troubleshooting processes — sometimes called workflows. Workflows provide step-by-step instructions that ensure best practices are followed and that all agents have a consistent approach to problem resolution. These same workflows can also be used for self-service care applications.

Typically, the steps included in a workflow are fixed, with a pre-defined sequence based upon a series of educated guesses. In reality, the most appropriate actions will differ from call to call, depending on the customer context. Until recently, however, workflows were not able to adapt to changing contexts; it was seen as too much effort to have a workflow respond to various options.

Using machine learning — which collects information on each successfully (and unsuccessfully) completed workflow — and adapting the sequence for every customer's unique situation, dynamic intelligent workflows can predict the optimal sequence of tasks that should be taken to resolve specific issues.

The effectiveness of these workflows can (and should) also be monitored. Further, this data should be stored and analyzed, resulting in a series of best practices that can be leveraged for future calls.

Algorithms use all of the information available — workflow history, customer information and network status — to prescribe specific workflows to agents using a recommendation engine that selects the next-best action (NBA) that has the highest probability of resolving a customer issue in the shortest time.

Instead of the fixed sequence that characterized workflows in the past, dynamic intelligent workflows start with a common set of introductory steps. Then, based on the available data — some of which is collected in near real time — quickly diverge into customized paths.

As a result, not only do all workflows get continuously optimized, but each individual workflow has the highest probability to resolve a customer issue in the shortest time. This enables faster response times, reduced support costs and a better customer experience. It also simplifies the workflow design since special cases based on context do not need to be hard coded into the workflows.

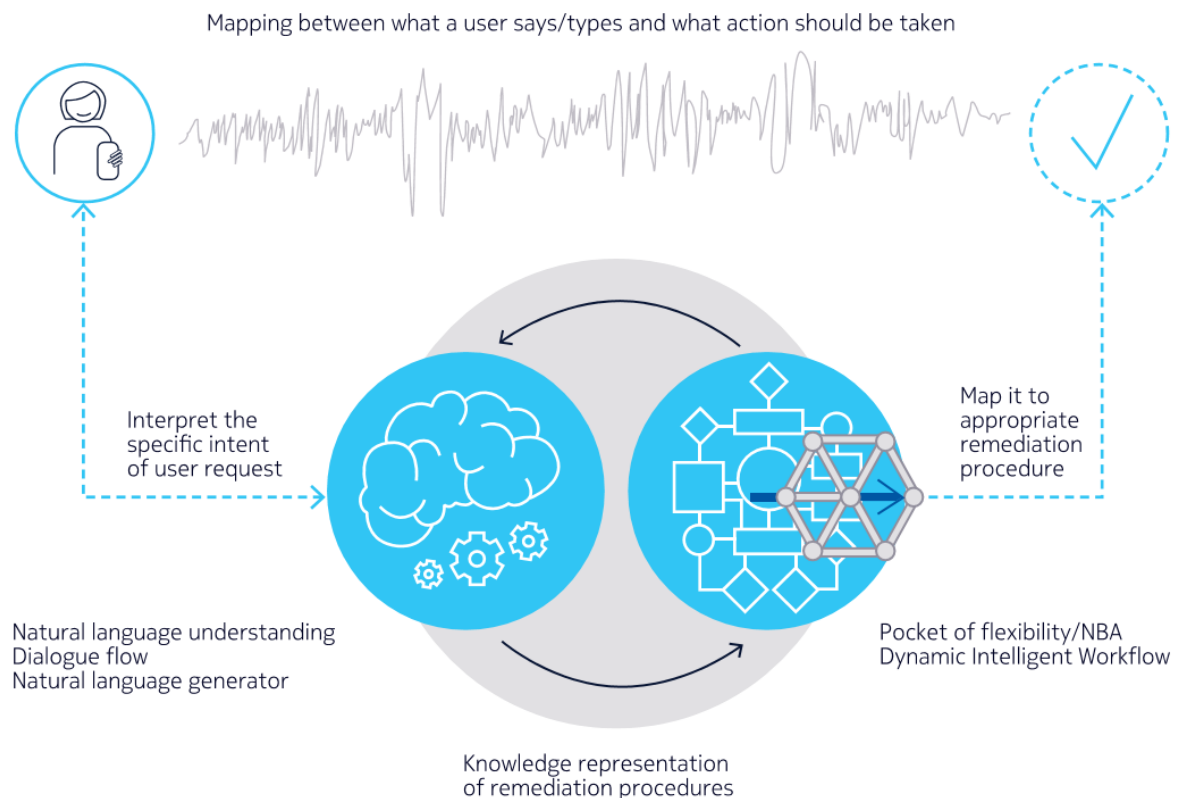


Figure 1 - Using technology to harness the power of the human word.

9. Personalization

MSO networks produce vast amounts of data every minute. Using this data, machine learning and AI technologies provide the foundation for advanced customer experience metrics. The ability to deal with varying amounts of data in real time will allow a much more detailed sensing of network and service conditions and the resulting user experience in real time.

In addition to being used for enhancing the omni-channel customer experience, analytics, AI, and machine learning can also be used to understand subscriber behavior and deliver personalized services.

By analyzing network and subscriber data, insights can be generated that allow for the compilation of subscriber profiles that can be used to classify subscriber issues, present the best solution, and provide a basis for targeted marketing activities that deliver the right offers to subscribers on the right devices at the right time.

Providing a personalized experience — combined with reduced complexity around customer care and offering self-service — are essential to increasing revenue, minimizing customer churn, and maximizing customer satisfaction.

Automating customer care

Advanced, predictive analytics today can determine that a customer is experiencing connectivity issues. Rather than waiting for that customer to seek technical support, remedial action can be taken proactively. The customer issue is resolved automatically, without the need for any interaction between the customer and the help desk.

10. Anticipating service disruptions using predictive analytics algorithms

Most MSOs rely on customer complaints (usually calls to the help desk) to learn about service disruptions. This is because existing systems, like network operations support systems (OSSs), cannot easily identify issues with the access network, customer premises equipment (CPE), or applications.

It's also important to note that 96 percent of customers don't complain after a poor customer experience, but 91 percent of them will switch providers. These customers are referred to as "silent churners" and they represent a major challenge for MSOs. If a customer doesn't call to complain, how do you know they even have a problem?

What is required is another way to identify and diagnose issues; without the customer having to contact the help desk. That would present a tremendous opportunity. Analytics provide the means to move toward proactive care, by capturing and storing data from the network, CPE, trouble tickets and more. Through analysis of this historical data, algorithms can be developed to better predict service disruptions and take proactive actions to address issues before the customer notices or calls in. Although it would be ideal for the network to always know when service degradations occur, it's often the spike in customer calls that is the first indication of a problem.

Based on this reality, algorithms can be used to track incoming calls to the help desk, correlate the rate of calls to each network element or service to the expected level of calls. The algorithm can then detect (in real time) when outage spikes are starting to happen and identify the offending network entity or service without needing to look at customer ticket records or check on past service disruptions. This is called a "call anomaly detection" algorithm.

The process starts with an examination of all calls received by CSRs. When a sudden burst of calls (a spike) is identified, the algorithm correlates the calls with the network and service topologies. Using real-time statistical signal processing algorithms, calls concerning possible service disruptions affecting multiple subscribers are categorized and separated from other calls being received. This call anomaly

detection can discover the outage within minutes and update IVR systems to play the appropriate message, thereby ensuring the call centers are not overwhelmed and can stay focused on solving the real issues.

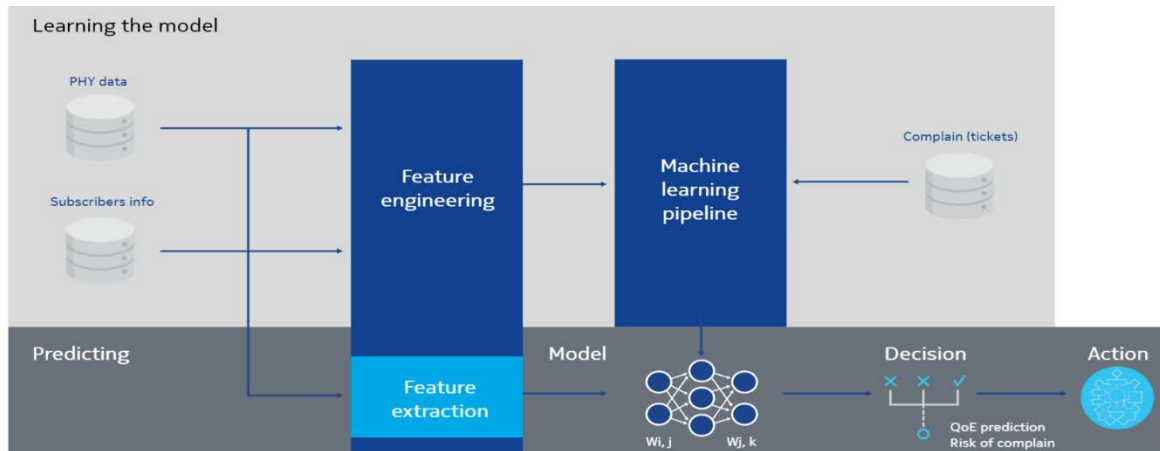


Figure 2 - Using predictive analytics to anticipate service disruptions.

The algorithms must be constantly tested, updated and refined using real-time data from the network and customer calls. In addition, the ability to implement actions automatically provides the MSO with the flexibility to adapt the actions according to the MSO's desired business processes. This is also known as autonomous customer care and it is the desired objective for both MSOs and customers alike.

11. Proactive Bots for automated problem detection and correction

Unlike a CSR, a bot can analyze real-time and historical data before initiating a response. Proactive bots, combined with analytics and machine learning, can identify service-affecting issues and fix them automatically, without any interaction between the customer and traditional support channels. The time is coming fast when customer care systems, leveraging bots, become so advanced that we enter the era of autonomous customer care.




Figure 3 - Proactive bots, combined with analytics and machine learning, can identify service-affecting issues and fix them automatically.

The real value for autonomous care will be found within the extensive library of use cases, known as the knowledge system. It will be the ability to match subscribers' intents to the appropriate remediation procedures (found in the knowledge system) that will provide the key to unlocking the evolution toward autonomous care.

As the knowledge system continues to improve with each subsequent transaction, it will eventually acquire the ability to perform complex inferences on that knowledge (reasoning), to the point where it approaches human-like intelligence. The result is a reduction in the number of help desk calls that require human intervention. This will free up CSRs that can be trained to handle more complex tasks or to provide premium technical support.

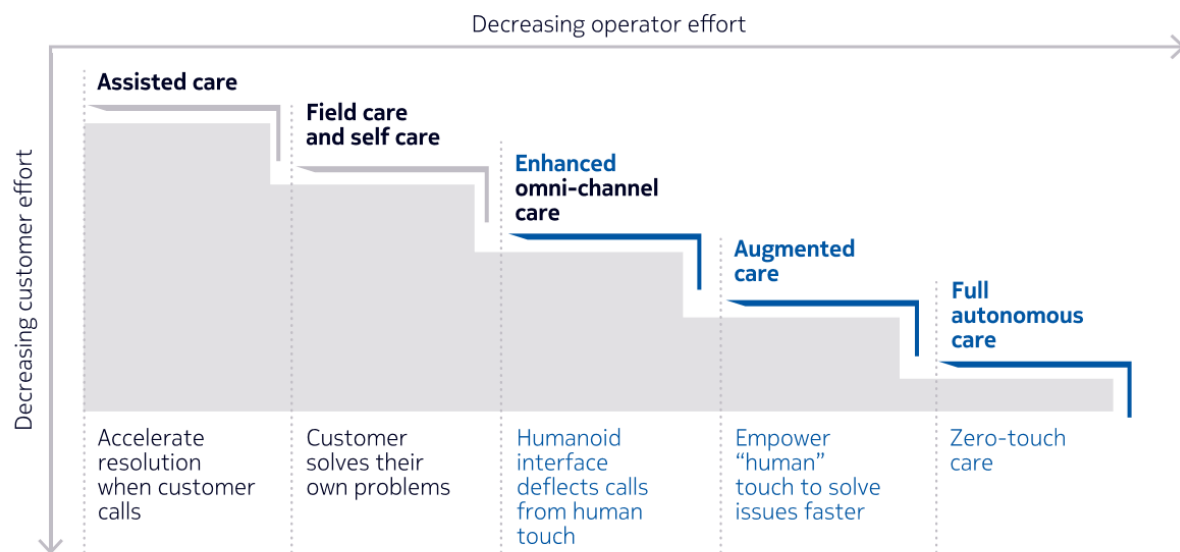


Figure 4 - The evolution of customer care, moving from assisted and self care, to autonomous care, leveraging new technologies and the power of the knowledge system.

By effectively mining and analyzing the vast quantities of network and subscriber data, MSOs can gain valuable insights about customers' experiences, preferences and predicted behaviors. These analytics provide real-time business intelligence that can be applied at every customer touch point. The first step is to determine the types of data that should be collected based on what the data is going to be used for. The data needs to be very specific to the customer care business; not just collecting data and storing it in data warehouses for its own sake. Ideally, the objective is to collect data that is connected to the network and mapped to various key performance and key quality indicators.

Once collected, this data can be used by various AI technologies to automate different tasks. For example, data can be collected from all managed CPE. An examination of that data might highlight an issue with a specific brand of CPE, for example. Then, rather than waiting for subscribers to seek technical support using traditional channels, action can be taken proactively. Customer issues will be resolved automatically, without the need for any interaction between the customer and the help desk.

This data can also be used to create the “self-healing network”, which allows MSOs to accumulate a list of known issues, map them to available solutions (found in the knowledge system), then allow for the automatic remediation of common issues that affect a significant number of customers. These common transactions fall in to what some MSOs call the 80:20 rule; 80 percent of the calls are related to 20 percent of the troubleshooting use cases.



Figure 5 - Data can also be used to create the “self-healing network”, mapping known issues to available solutions, then remediating common issues would otherwise affect a significant number of customers.

With these assets in place, a customer care system can respond automatically or assign interactions to appropriately-skilled CSRs. The system also allows customers to escalate to a live agent at any time. Gartner predicts that, by 2020, 10 percent of initial engagement requests will be taken by bots, up from less than one percent today.⁸

Use cases

Below are two use cases that illustrate the use of bots to resolve customer issues more quickly and efficiently.

12. Wi-Fi Network Password

One of the biggest call drivers for many MSOs are when subscribers forget their Wi-Fi network passwords. As they are often set once and forgotten, this information is not always readily available. This

⁸ “Predicts 2017: CRM Customer Service and Support”, Gartner, November 2016;
<https://www.gartner.com/doc/3505517/predicts--crm-customer-service>

is especially true for subscribers that set up sub-networks, typically for guests. This is a great use case to automate with your IVA.

- Subscriber: “Alexa, what is my guest Wi-Fi network password?”
- Bot: “I can send that to you. Do you want me to send it via e-mail or text message?”
- Subscriber: “By text message, please.”
- Bot: “Okay, a text message with your guest Wi-Fi network password has just been sent.”

13. Technician Locator

Another common call driver for many MSOs is when subscribers are awaiting the arrival of a technician at their home to either install new equipment or troubleshoot a technical problem. This is another great use case that can be automated with your IVA.

- Subscriber: “Okay Google, when will the GlobalComms technician arrive?”
- Bot: “Your GlobalComms technician, Lauren, is en route. She will arrive at your home no later than 1:15pm.”

Conclusion

Contact centers are playing an increasingly important role in the context of the overall customer experience. Tolerance for legacy customer care solutions is waning. There is an appetite for change and emerging technologies are generating substantial interest with consumers.

When done right, customer care can be a unique differentiator that delivers great value. Ineffective customer care, on the other hand, is very expensive and can turn loyal customers into detractors. Customer care solutions that leverage bots, AI, machine learning and NLP can enable the kind of efficient interaction that consumers are demanding.

Making customers more self-sufficient, providing CSRs with tools that result in faster, more accurate customer care and ultimately resolving issues automatically not only makes customers happier, but it also generates a number of significant business benefits for MSOs, such as fewer help desk calls, more efficient agents, lower customer churn, and fewer truck rolls; all leading to reduced support costs and improved profitability.

Abbreviations

AI	artificial intelligence
CPE	customer premises equipment
CSR	customer service representatives
IoT	Internet of Things
IVA	intelligent virtual assistant
IVR	interactive voice response
NBA	next-best action
NLP	natural language processing
NLU	Natural language understanding
OSS	operations support system
MSO	multiple system operator

Bibliography & References

“2017 Trends to Watch: The Digital Consumer Landscape”, Ovum, March 2017; <https://www.ovum.com/research/2017-trends-to-watch-the-digital-consumer-landscape>

“Customer care systems: worldwide forecast 2016–2020”, Analysys Mason, 20 September 2016; <http://www.analysysmason.com/Research/Content/Reports/Customer-care-forecast-Sep2016-RMA11>

“Understand The Spectrum Of Seven Artificial Intelligence Outcomes,” by R “Ray” Wang, Software Insider, September 18, 2016; <http://blog.softwareinsider.org/2016/09/18/mondays-musings-understand-spectrum-seven-artificial-intelligence-outcomes>

“Has voice control finally started speaking our language?”, Rhodri Marsden, The Guardian, December 4, 2016; <https://www.theguardian.com/technology/2016/dec/04/voice-control-amazon-echo-digital>

“Voice Is the Next Big Platform, and Alexa Will Own It”, Jessi Hempel, Backchannel, December 19, 2016; <https://www.wired.com/2016/12/voice-is-the-next-big-platform-and-alexa-will-own-it/>

“Mary Meeker: voice-controlled tech set for exponential rise in next few years”, Danny Yadron, The Guardian, June 1, 2016; <https://www.theguardian.com/technology/2016/jun/01/mary-meeker-voice-controlled-tech-boom-technology-predictions>

“Customer service of the future is powered by artificial intelligence,” Brandon Buckner, December 1, 2016; <https://www.ibm.com/think/marketing/customer-service-of-the-future-is-powered-by-artificial-intelligence>

“Predicts 2017: CRM Customer Service and Support”, Gartner, November 2016; <https://www.gartner.com/doc/3505517/predicts--crm-customer-service>

Proactive Network and Technical Facilities Monitoring Using Standardized Scorecards

An Operational Practice prepared for SCTE•ISBE by

Dr. Franklin Lartey
Senior Manager of Technology
Cox Communications
6305 Peachtree Dunwoody Rd, Atlanta, GA 30328
305-900-6601
Franklin.lartey@cox.com

1. Introduction

Proper measurement is important in making informed business decisions that improve customer experience. Several studies such as Lartey, McGinn, and Diponzio (2016), and Marut (2016), Snow and Weckman (2016) have repeatedly demonstrated the importance of measurement in telecommunications. Yet, measurement is still a complex issue in the broadband industry due to the plethora of measurement possibilities within this environment. Indeed, every card, port, or sub-channel on any communication device provides a number of measurement opportunities adding to the complexity of measurement in the above stated environment.

The current article presents a simple measurement system that seeks to proactively identify issues on the broadband network infrastructure including technical facilities and resolve them prior to any noticeable impact to customers. By so doing, this article will contribute to the increase in network reliability and availability, the decrease in trouble calls and truck rolls, and the increase in customer satisfaction and loyalty. To achieve its goal, the article presents a set of scorecards deployed by a multiservice operator (MSO) to proactively monitor its data, transport, DOCSIS, video, and telephony networks as well as its technical facilities. It also discusses the key performance indicators (KPI) used for each network and the rationale behind their selection, followed by a presentation of the operational model implemented to ensure proper monitoring. The operational model describes the process and structure that allowed the operator to proactively resolve identified issue. Finally, the article presents some real-life achievements of the standardized scorecard system, providing an opportunity for replication and automation within the industry.

2. Measurement Theory

Knowingly or not, human beings and animals use measurement on a daily basis for differentiation. For example, the biggest and strongest animal rules the pack; the fastest car wins the race. As explained by Tao (2011), measuring is a fundamental concept in Euclidean geometry. In that regard, a solid body can be measured in three different dimensions, respectively using the length on one dimension, the area on two dimensions, and the volume on three dimensions. The notion that every physical item has a value that can be measured without significantly influencing the item is well accepted in classical physics under the term *measurement theory*.

Measurement theory is the part of physical theory that determines the empirical and operational content of concepts used. Busch and Lahti (2008) posit that measurements are both operational procedures and physical processes. As operational procedures, they define the observables of the theory and as physical processes, they are also subject to the laws of physics. Some examples of measurement presented by Roberts (2009) includes brightness, intelligence, loudness, mass, preference, and temperature. Roberts puts temperature and mass in the well-developed science such as physics; he classifies intelligence as measurement in less-well-developed sciences such as psychology.

The theory of measurement is extensively discussed by Hand (1996) in his article “Statistics and the Theory of Measurement” published in the Journal of Royal Statistical Society. Confirming what precedes, Hand identifies three categories of measurement theory namely (1) representational measurement theory, (2) operational measurement theory, and (3) classical and other theories of measurement.

Representational measurement theory seeks to numerically represent attributes of objects, events, and even substances. In the creation of the standardized scorecard, individual measurements are taken from devices or obtained through a physical measurement activity. For example, the number of available rack spaces is an example of application of the representation measurement theory.

Operational measurement theory defines measurements in terms of specific operations that produce a number. An example of the application of this theory is the scorecard discussed in this article. Indeed, the scorecard takes a number of measurements and produces the overall score based on an internal process using weighted scores of identified categories.

Classical and other theories of measurement are concerned with subjective interpretations among others. An example of application is the determination of the cleanliness of a technical facility, which is a subjective judgement of the site's cleanliness.

All these categories of measurement theory have various schools of thought. In the broadband industry, various measurements exist such as signal level, noise level, speed, or bandwidth. The measurement retained will depend on the selected application and usage.

3. Key Performance Metrics Used

This article discusses network and technical facilities monitoring in a cable broadband organization also known as MSO or multiservice operator. In such setting, three main services are delivered to the end customers namely data, voice, and video services (Lartey, McGinn, & Diponzio, 2016). Data services mainly include Internet connections at various speeds, up to and even greater than one gigabit per second. Video services encompass audio visual programs sent to the end users using broadcast or narrowcast technologies. Voice services provide phone functionalities to customers using circuit or packet switching technologies. All these services are generally aggregated within a master telecommunications center (MTC) which can also be classified or called secondary telecommunications center (STC), regional data center (RDC), hub, head end, or simply technical facility depending on the size, functionality, and location of the facility on the network. Most networks are constructed in a 4-level hierarchy: national, regional, metro, and access as shown on figure 1 and confirmed by Noll (n.s.) and Schmitt (2012). Technical facilities house the inside plant (ISP) portion of the network and are the purpose of this article. Elements related to the outside plant (OSP) might be mentioned here but are not the focus of the article.




Figure 1 - Representation of a broadband network

1. Technical Facilities

A technical facility houses all the equipment necessary to the good functioning of a network, including power and cooling required for the equipment, hence the name critical system facility. Devices in a technical facility need power. This is generally provided in three main forms: commercial, generated, or stored. Commercial power is provided by the power company and the amount of power entering the facility should support its peak load. For that reason, the operator should have an acceptable indicator of the total power consumed in every technical facility. This metric is important in preventing overheating and other catastrophic events as explained by Rasmussen (2012).

Generated power is produced by a generator on site. Depending on the sizes of the facility, there can be as many as two or three redundant generators or even none. In situations where there is no generator, a good practice consists of prewiring the facility with a special connector for a portable generator. In the event of a power outage, the facility functions on battery until a generator is brought on site. Just as with commercial power, the generator should be sized to support the peak load of the facility.

For its proper operation, a technical facility needs to store power to use as needed. That can be achieved using an uninterruptible power source (UPS) or a bay of batteries. Like with commercial power and generators, stored power can support a given load needing to be measured.

Every device in a technical facility is powered using either alternating current (AC) or direct current (DC). In general, both types of current are needed in a technical facility. Indeed, some network devices operate on DC while others operate on AC power. Inverters and rectifiers are needed for proper AC and DC powering in a facility. Inverters convert AC to DC while rectifiers produce AC either from commercial power, generator, or batteries. The overall relationship is depicted in figure 2.

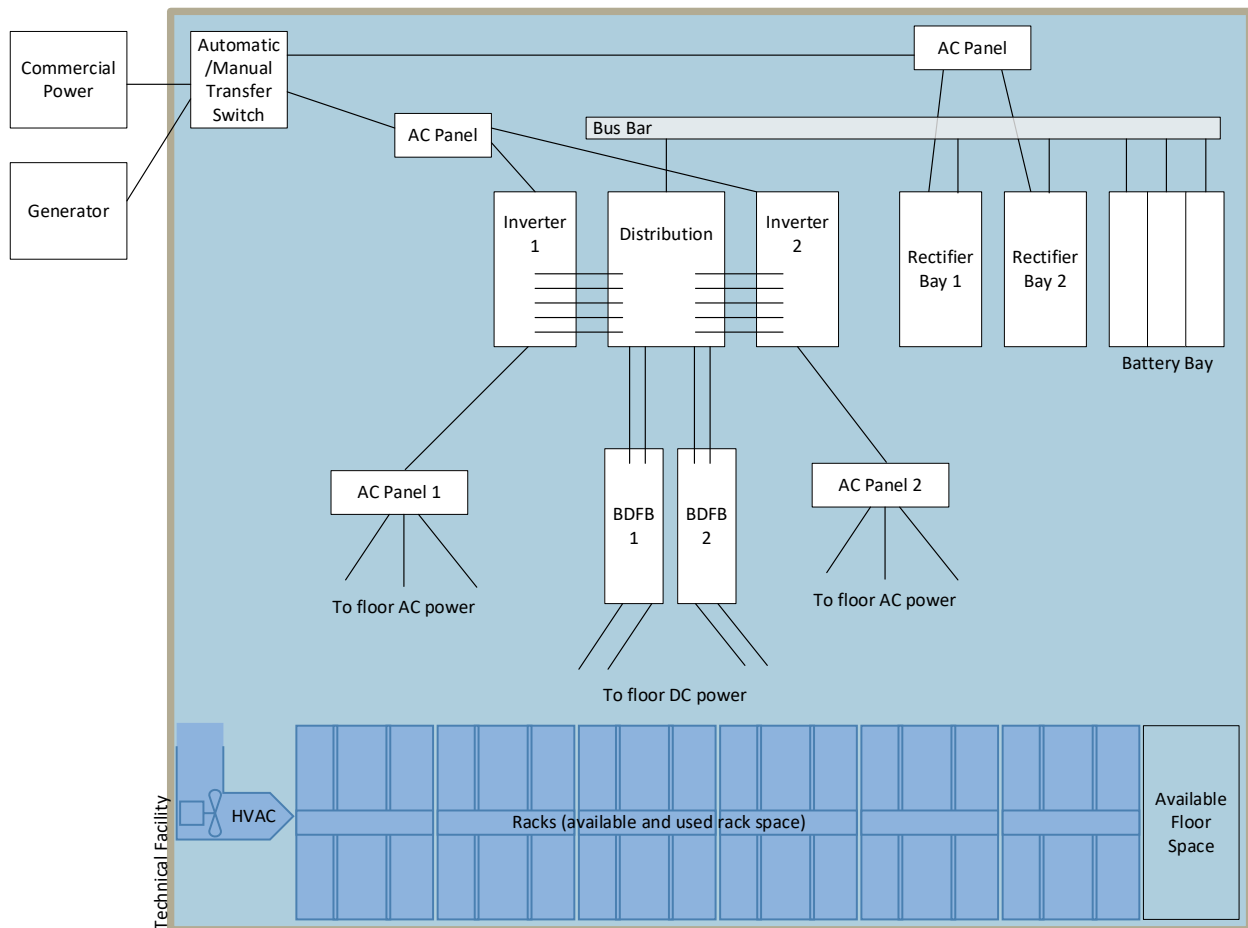


Figure 2 - Technical facility components

By using power, all devices in the technical facility produce a quantity of heat. It is thus necessary to have a heating, ventilation, and air conditioning system (HVAC) for continuous operation of the devices. To ensure these devices will run appropriately during periods of intense heat, the capacity of the HVAC needs to be appropriate and depending on the criticality of the facility, redundant HVAC systems could be required.

Regardless of its size, the availability of rack space or floor space is an important indicator in a critical facility. This metric is useful when new equipment needs to be installed or when an obsolete equipment is replaced by a newer and larger one. In addition, because cleanliness is key in the proper operation of equipment, it is a metric to consider. In the case presented here, cleanliness was added as a subjective observation. The role of technical facilities being to enable network continuity, technical facilities are inter-connected on the network. Each facility houses the necessary devices allowing the propagation of information on various networks such as data, transport, data over network system interface specification (DOCSIS), video on demand (VOD), switched digital video (SDV), telephony, and much more. The overall scorecard also includes measurements of these networks and their related devices.

2. Data Network

The data network is composed of interconnected switches and routers residing in different technical facilities. Traffic from various devices are aggregated by switches and sent to routers to be forwarded across the network routing boundary onto different locations including the Internet. As a result, if a port on the router is congested and depending on the configuration, traffic starts getting slower due to packet retransmission. A congested router port in the headend has the potential of creating bad customer experience due to slow traffic. For this reason, the data scorecard identifies routers of importance and captures information related to port capacity and utilization.

Network growth being a constant with today's bandwidth hungry applications, port availability could be a metric to consider. While the current scorecard did not take this metric into account, it is good practice to include such measurement for awareness. Nonetheless, its absence did not impact the proactive nature of this system because problems were identified long enough to take the necessary corrective actions.

As a summary, the data scorecard used the router ports' bandwidth utilization as the main metric. It identified the main system routers in every technical facility, the ports carrying the main payloads, and captured the values related to the traffic on these ports.

3. Transport Network

The transport network helps interconnect technical facilities using technologies that allow information transportation over very long distances. Such technologies include synchronous optical network (SONET), dense wavelength division multiplexing (DWDM), time division multiplexing (TDM), optical transport network (OTN), and much more. The medium of choice is the fiber optic cable for long distances and local cross-connects, even though copper cables are also used for local cross-connects. The transport network provides various functionalities including information transport, multiplexing, redundancy in case of fiber cut, switching, and management.

The main transport equipment on the network is an optical network element (ONE). Examples of network elements include as the Cisco ONS 15454 SONET/SDH multiservice transport platform (MSTP), the Cisco ONS 15600 multiservice switching platform (MSSP) (Cisco, n.s.), the Fujitsu Flashwave 9500 Packet Optical Networking Platform, and the Fujitsu Flashwave 7500 metro/regional multiservice ROADM or reconfigurable optical add drop multiplexer (Fujitsu, n.s.), just to name a few elements used in the case presented in this article.

On the network considered here, information is transmitted between optical network elements through the DWDM system. The system uses 40 wavelengths also known as lambdas of 10G, 40G, 100G, OC48, and OC192. For that reason, the number of available wavelengths was a metric used in the scorecard.

Knowing that a wavelength could transport Ethernet or an optical carrier (OC) signal, the transport scorecard would measure an OC wavelength deeper and leave the Ethernet carrier to be measured by the data scorecard at the router's port. As such, besides identifying critical circuits and measuring their available lambdas, the transport scorecard also measured available synchronous transport signal (STS) in optical carriers and available cross-connects in network elements.

4. DOCSIS Network

The Data Over Cable System Interface Specification (DOCSIS) network is also known as the Cable Modem Termination System (CMTS) network. Its main role is to provide Internet to residential as well as commercial customers. DOCSIS is a protocol developed by CableLabs to specify communication on the hybrid fiber coaxial (HFC) cable at the data link layer of the OSI model as well as the physical layer (signal modulation and transmission). The CMTS is a network device that controls all communications on the DOCSIS network. Example of CMTS devices include the Cisco uBR10000 (Cisco 10K), Arris C4, Casa C3000 series, and Motorola BSR 64000 just to name a few. New versions of CMTS are known under the acronym CCAP or Converged Cable Access Platform. The CCAP decentralizes some functions of the CMTS, offers video distribution, and supports DOCSIS version 3.1 that currently provides up to a gigabit per second of Internet bandwidth to end users (Sundaresan, 2015). Examples of CCAP include Cisco's CBR-8 and Casa's C100G.

One CMTS serves many HFC nodes and every node serves many homes and businesses. In the case presented in this article, an HFC node served an average of 550 homes, with some nodes serving over a thousand homes. The main issue on the CMTS network is congestion due to shared bandwidth which creates a negative customer experience. As such, a proactive way of improving customer experience is to identify congestion prior to observable effects to customers and the metric retained for this measurement is bandwidth utilization.

5. Video Network

The video network allows operators to acquire, process, and transmit video to subscribers over the HFC network using the same plant as that used for the CMTS network. Video services include broadcast and on-demand services. In the broadcast category, a copy of a specific program is sent to all subscribers at the same time. In this configuration, all subscribers see the same program and the same advertisements as they are sent over the network. Vasudevan, Liu, and Kollmansberger (2008) identify two types of broadcast services namely digital broadcast and switched digital video. Digital broadcast sends all video programs to the subscribers' set-top box (STB), while switched digital video (SDV) sends video programs when a subscriber or the group of subscribers request it. With this technique, SDV optimizes the use of the spectrum because it allocates frequencies on the spectrum only to channels that are being watched. Even though SDV optimizes the use of spectrum, proper allocation needs to be made or customers might experience poor services where they cannot tune to the desired program and have to keep trying until bandwidth becomes available. For that reason, the scorecard includes bandwidth utilization by service group (SG) as key metric for SDV.

On-demand services are different from broadcast services in the sense that the requested program is sent to only one subscriber. In their article on IPTV architecture for cable systems, Vasudevan, Liu, and Kollmansberger (2008) identify two types of on-demand services namely video on demand (VOD) and network-based personal video recorder (nPVR) which is also known as cloud DVR. nPVR was not considered for the scorecard because the network presented here used DVR (digital video recording) which stored recorded programs directly on the recorder and did not require network access to watch them.

Video on demand (VOD) stores movies on a server on the network and when a subscriber selects a movie to watch, the content streams to the subscriber's STB where the movie is displayed on the screen. In this configuration, bandwidth is allocated on the network for that movie alone. If there is no bandwidth

available, the subscriber would have to try over and over, creating a negative customer experience. For that reason, bandwidth utilization on the VOD network was retained as key metric on the scorecard.

6. Telephony Network

The telephony network allows phone subscribers to call other phone subscribers or to receive phone calls from anywhere in the world. The phone subscribers in the case presented here are connected to the network by the same HFC plant used for the data and video networks. Just like the other networks, the telephony network contains interconnected network elements that process and transport voice information, connecting a caller to a recipient who could be on off network. Some of these network elements include the main switches, the gateways, and the cross-connects.

In the network of interest by this article, there were many Digital Multiplex Systems (DMS) and CallServer 2000 (CS2K) switches. Along with the switches were the Tellabs Titan 5500 Digital Cross-Connect System (DCS/DACS), the Tellabs Titan 532 DACS, and the Nuera gateways 4K and 8K. To ensure good customer experience, all these elements needed to be measured. For the DMS/CS2K, the selected metrics included SPM (spectrum peripheral module) port availability, IWSPM (interworking spectrum peripheral module) peak utilization, end point capacity, TVCID AIN triggers (television caller identification advanced intelligent network trigger), and SIP SOC RTU peak utilization (session initiation protocol - software optionality control - right-to-use). These were identified as metrics properly addressing the performance and capacity of the switches.

For the gateway controllers (GWC), the metrics selected were the trunk side GWC BHCA (busy hour call attempts) utilization and the line side utilization. For each 5500 DACS, the number of STS available was the key metric and for the 532 DACS, the number of available ports was the key metric of capacity and utilization.

7. RF Spectrum

All three services provided by the MSO (Internet/data, video, and phone) are simultaneously delivered to the end customer using the same medium: the coaxial cable (Lartey, 2015). Regardless of the medium, signal is always transmitted on specific frequencies; a frequency being the speed of the vibration or number of wave cycles in one second. For example, sounds audible to human beings are generally accepted to be in the frequency range between 20 and 20,000 hertz (Hass, 2003). The hertz (Hz) is the measurement unit of frequency and 20,000 Hz is the same as 20 KHz (Kilo Hertz). Figure 2 represents a wave cycle and figure 3 shows an example of a signal being propagated at the speed of 3Hz or 3 wave cycles per second. In the millions, the frequency is expressed in megahertz (MHz) and in the thousands of millions, it is expressed in gigahertz (GHz).




Figure 3 - Representation of a wave cycle for information transmission




Figure 4 - Representation of a signal with a frequency of 3 Hz

Just like any communication medium, the coaxial cable allows the transmission of analog signal on different frequencies. A range of frequencies is known as a spectrum and the HFC plant currently operates on frequencies between 5 MHz and 1 GHz upgradable to 1.2 GHz and even 1.8 GHz in the near future (Urban, 2010). Part of this spectrum is used for upstream signals from the subscriber to the headend and the other part is used for downstream communication from the headend to the subscriber. The spectrum from 0 to 5 MHz is generally too noisy and currently unusable. In conventional US DOCSIS networks, the spectrum from 5 to 42 MHz is used for upstream communication while the downstream channels use the spectrum from 54 MHz to 1 GHz. The spectrum from 42 to 54 MHz is used as a guard-band where a diplexer or diplex filter installed in the HFC node separates the upstream return from the downstream service as explained by Watts (2011) and Urban (2010). However, European and contemporary DOCSIS standards offer more variations for frequency spectrum allocation.

Because there is a limited number of usable frequencies on the coaxial plant, managing the spectrum is an important task contributing in improving customer experience. For example, a new television channel or additional DOCSIS bandwidth can be added only if there is available spectrum. For that reason, the HFC spectrum was included in the standardized scorecard and the main metric is the number of available

channels. This number is expressed in MHz for the upstream and in 6 MHz chunks for the downstream, assuming SC-QAM. This choice was based on the fact that a 6 MHz spectrum can carry either one analog video program, 12 standard definition video programs, 3 high definition video programs, two 3D television programs, or one DOCSIS channel at 38.8 Mbps. This channel paradigm includes OFDM blocks in DOCSIS 3.1, however that is not currently within the scope of this paper.

4. Detailed Procedure and Results

The scorecard is a measuring system presenting the health of all networks based on the common principle stipulating that “if you cannot measure it, you cannot improve it”. Indeed, there needs to be a sense of measurement to determine improvement. Because the broadband infrastructure includes technical facilities and many different networks, there are also different scorecards: one for technical facilities, one for each type of network, one for RF spectrum, and the final system scorecard is a standardization or normalization of all the others. In other words, the following scorecards will be created on a regular basis: (1) technical facilities; (2) data network; (3) transport network; (4) DOCSIS network; (5) video network; (6) telephone network; (7) RF spectrum network; and (8) standardized system scorecard.

Each scorecard uses specific key metrics presented in the previous section. Three scoring states are defined for reporting each metric: green, yellow, or red. The green state means the system is performing as expected for that metric. Yellow means the system is performing below expectations and suggests caution. Red implies that something is wrong and needs to be fixed. Each scoring system is presented in what follows.

1. Technical facilities scoring system

Technical facilities are classified into four tiers based on five main criteria: (1) impact to the network in case of facility failure; (2) monthly recurring revenue (MRR) generated in the area served by the facility; (3) number of customers served by the facility; (4) value of the building constituting the technical facility; and (5) type of equipment and number of DOCSIS serving groups in the technical facility. Table 1 shows an example of classification, with tier-1 sites being the most critical to network operations and customer impact.

Based on the technical facilities classification, a scoring grid was created with points allocated to each key metric by facility tier as shown on table 2. For each facility, every key metric is scored based on the specifications of the telecommunication building critical system tool (TBCST), an example of which is represented on table 3.

Table 1 - Classification of technical facilities in the region

Criteria	Tier-1	Tier-2	Tier-3	Tier-4
Impact of the technical facility in case of catastrophic failure	Significant	Significant	Moderate	Moderate
Revenue or monthly recurring charges (MRC) generated in the area served by the facility (all services).	≥ \$2,000,000	≥ \$1,000,000 and < \$2,000,000	≥\$250k and < \$1,000,000	N/A

Criteria	Tier-1	Tier-2	Tier-3	Tier-4
Number of customers served by the facility	> 20,000	Between 7,500 and 20,000	> 1,000 and <7,500	N/A
Value of the building (not including equipment housed)	> \$1 million	Between \$500,000 and \$1 million	> \$150,000 and < \$500,000	N/A
Type of equipment in the facility	Voice Gateway, CS2K, RDC eqt.	≥ 17 CMTS service groups	≥ 1 CMTS service group	N/A

Table 2 - Technical facility scoring grid

Tier - Health	HV AC	Pwr	Gen.	DC Plant	Batt.	UPS / Inv.	Space	Monit.	Clean	Fire Sup.	Score
1 - G	8	6	8	7	7	6	4	2	1	1	50
1 - Y	4	3	4	3	3	3	2	1	0	0	23
1 - R	2	1	2	1	1	1	1	0	0	0	9
2 - G	5	3	5	4	4	3	2	2	1	1	30
2 - Y	2	1	2	2	2	1	1	1	0	0	12
2 - R	1	0	1	1	1	0	0	0	0	0	4
3 - G	2	1	2	2	2	1	2	1	1	1	15
3 - Y	1	0	1	1	1	0	1	0	0	0	5
3 - R	0	0	0	0	0	0	0	0	0	0	0
4 - G	-	1	-	1	1	1	-	-	1	-	5
4 - Y	-	0	-	0	0	0	-	-	0	-	0
4 - R	-	0	-	0	0	0	-	-	0	-	0

Table 3 - Scoring of key metrics

Key Metric	Measurement Unit	Green	Yellow	Red
AC Service (Power)		<60%	60- 80%	>80%
AC Service Main Breaker Size	Amps @ 277/480V, 3 Phase			
	Amps @ 120/208V, 3 Phase			
	Amps @ 120/240, 1 Phase			
Annual Peak Load	Kilowatts			
Generators		<60%	60- 80%	>80%
Generator Capacity	Kilowatts			
DC Plant and Battery		≥ 4h	2h - 4h	< 2h
DC Power Plant Size	Amperage			
DC Power Plant Load	Amperage			
DC Power Plant Battery Amp Hours	Amp Hours			
UPS Size	Kilowatts			
UPS Output Load	Kilowatts			

Key Metric	Measurement Unit	Green	Yellow	Red
<i>Inverters</i>		<60%	60- 80%	>80%
Inverter Size	240V Single Phase Kilowatts			
	120V Single Phase Kilowatts			
	120/208V 3Three Phase Kilowatts			
Inverter Load	Amperage			
<i>HVAC</i>		<60%	60- 80%	>80%
Total Tons Building HVAC	Tons BTU			
Total Tons Headend HVAC	Tons BTU			
<i>Floor Space & Monitoring</i>				
Available Floor Space	Rack locations			
Video monitoring system installed	Judgement			
Site Cleanliness	Judgement			

An example of the technical facilities scorecard is shown on figure 5. As presented, both sites 5 and 13 have space issues. Site 6 has saturated commercial power based on its peak load, and site 15 needs a bigger generator because equipment deployed have outgrown the existing generator capacity. The overall technical facility system score is 89 and the health of the system is coded yellow, suggesting caution.

TECHNICAL FACILITIES SCORE CARD

Oct-yy

Facility	Tier	HVAC	Power	Gen	DC Plant	Battery	UPS/Inv	Space	Monitoring	Cleanliness	Fire Supp	Score
Site 1	1											46
Site 2	1											43
Site 3	1											46
Site 4	1											48
Site 5	1											39
Site 6	1											43
Site 7	1											50
Site 8	1											42
Site 9	1											46
Site 10	1											46
Site 11	1											50
Site 12	1											46
Site 13	1											35
Site 14	1											42
Site 15	1											44
Tier 1 Totals												44.4
Site 16	2											30
Site 17	2											30
Site 18	2											30
Site 19	2											27
Site 20	2											28
Site 21	2											28
Site 22	2											30
Site 23	2											26
Site 24	2											27
Site 25	2											28
Site 26	2											27
... List Truncated ...												
											Total Score	89.95

Figure 5 - The Technical Facilities Scorecard

2. Data Network Scoring System

To score the data network, the most critical links were first identified. These are aggregation links carrying customer and service traffic. Examples of such links include the connections between the RDC and the MTC or leased circuits for redundant Internet connectivity. These links, also known as interconnects, were grouped by type (backbone, leased, spur, ring) and their utilizations were individually reported based on the utilization report from Tivoli Netcool Performance Manager (TNPM).

An interconnect is considered green if its utilization ratio is less than 60 percent of its capacity at the 95th percentile. All interconnects being protected on an east-west path, this allows for the support of most traffic on one side in the event of a failure of the other path. Interconnects are marked yellow if their

utilization ratio is between 60% and 75% at the 95th percentile while those with utilization above 75% are marked red.

On the final data scorecard, green interconnects add 100% of their weighted value based on the number of interconnects in the category or type. Yellow interconnects add 50% of their weighted score and red interconnects add nothing. The sum of all these scores results in a total score showing the health of the data network. A total score above 95% is green; between 80% and 95% is yellow; and less than 80% is red. The total score is then standardized and reported on the system scorecard for the data network.

Figure 6 shows an example of a data scorecard resulting from sub-scorecards of selected links. In this representation, none of the data links is red which means the data network is in a healthy state, scoring 97.96 over 100 possible points.

DATA SCORECARD				Oct-yy
	Total Interconnects	Green	Yellow	Red
RDC Backbone	10	10		
Leased Circuits	4	2	2	
Spurs	9	8	1	
East K Ring	43	43		
West K Ring	27	27		
Metro Ring A	17	17		
Metro Ring O	37	34	3	
Total	147	141	6	0
Percentage	100%	95.92%	4.08%	0.00%

Monthly Score to Report	97.96
--------------------------------	--------------

<i>Legend</i>	
Green	interconnects add 100%
Yellow	interconnects add 50%
Red	interconnects add 0%

Figure 6 -The Data Network Scorecard

3. Transport Network Scoring System

The principle of the transport network scorecard is similar to that of the data network scorecard but instead of Ethernet links, the transport scorecard identifies lambdas, STS', and cross-connects of importance as explained in the previous section related to the transport network. For the first category, all lambda routes are listed and their utilization scored. A lambda or wavelength is considered utilized once it

is allocated for a specific purpose. For example, a 10G lambda can be allocated to carry data traffic between the MTC and the RDC.

A lambda route is red if its utilization is above 75% of capacity. The capacity of a lambda route is the total number of equipped lambdas on that route. The lambda route is yellow if its utilization is between 60% and 75% and green if its utilization is below 60%. All the lambdas in the transport network are currently configured in a ring topology to allow redundancy.

Just like lambdas, STS' are organized in rings. Each STS ring is listed under its master ring name along with the STS' name and characteristics (e.g. OC-48, 2xOC48, OC-192). A ring listed as 2xOC-48 suggests that a second OCx48 was activated on that same ring with the same parameters as the first. Each STS ring is scored green, yellow, or red using the same utilization criteria like the lambda (<60%, 60-75%, >75%) and the total number of green, yellow and red STS rings is reported on the scorecard.

The third category on the transport scorecard is the cross-connects report. Cross-connects were a major source of failed or delayed projects due to chassis constraints and lead time to acquire and deploy new cross-connect cards such as the Cisco XC, XCVT, XC10G, or XC-VXC cards. For that reason, including them on the scorecard was deemed a proactive measure for improving customer experience by increasing on-time delivery of new services. As such, cross-connect report was added to the transport scorecard to identify shortcomings and take proactive action. This sub-scorecard identifies every transport network element (NE) or chassis with cross-connect capability and lists the percentage of VT1.5 and STS-1 consumed in the chassis. Using the same scale like the previous categories, each NE's VT1.5 and STS-1 are scored green, yellow, or red. Any red score makes the chassis red. If both VTs and STS' are green, the chassis is green else, the chassis is yellow. The result of this sub-scorecard is reported in the transport scorecard as a sum of chassis and a sum of each color. For the final score, any green or yellow percentage score increases the total score while a red percentage score reduces the final score.

The overall transport network score is obtained by adding the scores for lambda, STS, and cross-connects as shown on figure 7. If the final score is 90 or more, the transport network is green. A score below 90 but greater or equal to 75 makes the system yellow. A score below 75 makes the system red.

TRANSPORT SCORECARD

Oct-yy

	Region-K	Region-A	Region-O	Total	Score
<i>Lambda</i>					
Total Lambda routes	8.00	3.00	4.00	15.00	
Green	7.00	2.00	3.00	12.00	
Yellow	1.00	0.00	0.00	1.00	
Red	0.00	1.00	1.00	2.00	0.87
<i>STS</i>					
Number of STS Rings	37.00	19.00	13.00	69.00	
Green	29.00	11.00	8.00	48.00	
Yellow	2.00	2.00	4.00	8.00	
Red	6.00	6.00	1.00	13.00	0.81
<i>Cross-Connects</i>					
Total shelves	67.00	8.00	25.00	100.00	
Green shelves	60.00	8.00	24.00	92.00	
Yellow shelves	4.00	0.00	1.00	5.00	
Red shelves	3.00	0.00	0.00	3.00	0.97

Monthly Score to Report 88.28

Legend

Green lambda, STS, or crossconnects maintains score
Yellow lambda, STS, or crossconnects maintains score
Red lambda, STS, or crossconnects reduces score

Figure 7- The Transport Network Scorecard

4. DOCSIS Network Scoring System

The basis of the DOCSIS or CMTS scorecard is the weekly average utilization report of all radio frequency (RF) signals or carriers allocated to this network in general and to Internet in particular. Indeed, signals related to the DOCSIS set-top gateway (DSG) for set-top out-of-band communication, telemetry, and management are not taken into account in this scorecard. The list of DOCSIS carriers (RF signals allocated to DOCSIS) is grouped by serving group (SG) also called service group, representing the group of nodes served by a group of RF signals on the upstream or on the downstream. The upstream and downstream scores are calculated separately using a common principle. Each carrier is considered over-utilized if its utilization at the 95th percentile is 70% or above. Based on this report, a SG is made red if any one of its carriers is over utilized. The list of red SGs will later be used during the analysis to address any pending issue on the network.

Because a SG resides on one CMTS as two CMTS' do not serve the same node, the CMTS is cored based on the ratio of red SG compared to the total number of SG it contains. A CMTS chassis is green if less than 5% of its carriers are over-utilized; it is yellow if the ratio of over-utilized carriers is between 5 and 10%, and red if it is greater than 10%. This view helps identify chassis creating the most pain on the

network. At the system level, a market or region will be green if less than 5% of its carriers are over-utilized; yellow if over-utilized carriers are between 5 and 10%, and red if overutilization is above 10%. The overall system score is the average of the upstream and downstream system scores as presented on figure 8.

DOCSIS CAPACITY SCORECARD

Location	Week-1	Week-2	Week-3	Week-4
Region A	92.26%	91.74%	94.79%	94.76%
Region K	97.01%	96.78%	97.22%	97.70%
Region O	96.95%	97.48%	98.19%	98.04%
Region S	98.32%	92.26%	100.00%	98.92%
Total	96.52%	96.38%	97.19%	97.32%
MONTHLY		96.85%	STD SCORE	
			93.70	

Downstream Carriers	
Location	Count
Region A	1265
Region K	4664
Region O	3259
Region S	104
Central	9292

Upstream Carriers	
Location	Count
Region A	741
Region K	2102
Region O	2761
Region S	84
Central	5688

KEY:

< 5% of Total Carriers are Over Utilized

< 10% of Total Carriers are Over Utilized

>= 10% of Total Carriers are Over Utilized

Figure 8 - The DOCSIS Network Scorecard

5. Video Network Scoring System

The video capacity scorecard includes scores related to the VOD and SDV service group utilization. The VOD successful sessions and SDV tune-in success rate are part of the video reliability scorecard and will not be addressed in this article. To create the VOD and SDV capacity scorecards, the list of respective SGs is generated using the respective element management systems (EMS), along with the average peak utilization for the selected period based on a 15 minute-interval. This calculation is based on the total bandwidth available for the SG, knowing that a VOD or SDV AQM is 38.8 Mbps. As such, a SG with 4 QAMs is 155 Mbps, one with 5 QAMs is 194 Mbps, and one with 8 QAMs is 310 Mbps.

A SG is marked over-utilized if its average peak utilization is greater than 75% for SDV or 85% for VOD. The difference in the maximum utilization ratio derives from the manufacturer recommendation. Indeed,

75% is the threshold recommended by the SDV solution provider. It is the limit at which additional QAMs or new SGs are required to maintain a good customer experience. The final SDV and VOD scores are calculated for each region using the formula:

$$1 - \frac{\text{Total Over Utilized SG}}{\text{Total SG}}$$

The system score is calculated the same way using overall system values. An example of the video scorecard is presented on figure 9. Any market or system score above 95% is marked green; a score between 90 and 95% is marked yellow, while any score below 90% is considered red. The final system score is standardized and reported on the overall system capacity scorecard for the given period.

VIDEO (VOD & SDV) CAPACITY SCORECARD

SDV Service Groups with Utilization above 75%

Location	September-yy		October-yy	
	Total SG	SG >75%	Total SG	SG >75%
Region A	296	4	320	9

SDV Service Group Utilization Performance

Location	September	October	Variance
Region A	98.65%	97.19%	-1.46%

VOD Service Groups with Peak Utilization Greater than 85%

Location	September-yy		October-yy	
	Total SG	SG >85%	Total SG	SG >85%
Region A	320	9	320	0
Region K	353	69	388	82
Region O	208	10	298	9
Region S	16	0	16	0
Central Region	897	88	1022	91

VOD Service Group Utilization Performance

(Percentage of SG under the 85% threshold for the reporting period.)

Location	August	September	October	Variance
Region A	94.93%	97.19%	100.00%	2.81%
Region K	80.00%	80.45%	78.87%	-1.59%
Region O	85.10%	95.19%	96.98%	1.79%
Region S	100.00%	100.00%	100.00%	0.00%
All Regions	86.71%	90.19%	91.10%	0.91%

Total Utilization Performance for the Month	91.10%
---	---------------

Normalized Score to Report for the month	82.19
--	--------------

Green	95% to 100% VOD Service Group Utilization Performance
Yellow	90% to 95% VOD Service Group Utilization Performance
Red	<90% VOD Service Group Utilization Performance

Figure 9 - The Video Network Scorecard

6. Telephony Network Scoring System

The telephone scorecard is based on metrics identified as determinants of the telephone network capacity. These metrics include SPM port availability, IWSPM peak utilization, CS2K end point capacity, TVCID AIN triggers, SIP SOC RTU peak utilization, trunk side GWC BHCA utilization, line side GWC BHCA

utilization, DS1 capacity on gateways, and STS available on DACS. Information on availability is based on utilization. For example, if 25% of ports are utilized, then 75% of ports are available.

The telephone scorecard is organized in four main categories namely (1) DMS/CS2K; (2) Gateway controller utilization; (3) Nuera 4K/8K Gateways; and (4) DACS. All categories are weighted the same, each contributing at 25% of the total telephony score. Each category contains a set of weighted key metrics as shown of table 4. The weight of a metric in a category is determined by its potential impact to customers in case of saturation. An example of the telephony scorecard is presented on figure 10.

Table 4 - Telephony Scoring Grid

Category / Metric	Metric Weight	Category Weight
<i>DMS/CS2K</i>		25%
SPM Port Availability	20%	
IWSPM Utilization Peak	10%	
CS2K End Point Only Capacity	30%	
TVCID AIN Triggers	10%	
SIP RTU SOC Peak Utilization	30%	
Category Totals	100%	
<i>Gateway Controller Utilization</i>		25%
Trunk Side GWC BHCA Utilization Only	60%	
Line Side GWC BHCA Utilization Only	40%	
Category Totals	100%	
<i>Nuera 4K/8K Gateway</i>		25%
DS1 Capacity Market-W	40%	
DS1 Capacity Market-J	30%	
DS1 Capacity Market-F	30%	
Category Totals	100%	
<i>DACS</i>		25%
5500 STS Available	70%	
532 Ports Available	30%	
Category Totals	100%	
<i>OVERALL SCORE</i>		100%

TELEPHONY CAPACITY SCORECARD

Key Telephony Metrics (Region-K)	Aug-yy	Sept-yy	Oct-yy
DMS/CS2K			
SPM Port Availability	4	4	4
IWSPM Utilization Peak	25.72%	25.66%	25.78%
CS2K End Point Capacity	82.28%	81.59%	80.41%
TVCID AIN Triggers	73.27%	73.09%	72.40%
SIP SOC RTU Peak Utilization	50.95%	52.51%	83.35%
Category Totals	60.00%	60.00%	37.50%
Gateway Controller Utilization			
Trunk Side GWC BHCA Utilization Only	9.39%	9.44%	9.33%
Line Side GWC BHCA Utilization Only	27.25%	27.14%	27.44%
Category Totals	100.00%	100.00%	100.00%
Nuera 4K Gateway			
DS1 Capacity Market W	63.69%	65.20%	65.20%
DS1 Capacity Market J	85.11%	85.11%	85.11%
DS1 Capacity Market F	48.21%	48.21%	48.21%
Category Totals	77.50%	77.50%	77.50%
DACS			
5500 STS Available	32	31	32
532 Ports Available	36	36	39
Category Totals	82.50%	82.50%	82.50%
Total Score Region-K	80%	80%	74%

Key Telephony Metrics (Region-O)	Aug-yy	Sept-yy	Oct-yy
... Scorecard Truncated ...			
Total Score Region-O	83%	83%	83%
TELEPHONY SYSTEM SCORE	81%	81%	78%

Overall Scoring System
Green: 95% to 100%
Yellow: 80% to 95%
Red: Below 80%

Figure 10 - The Telephone Network Scorecard

7. RF Spectrum Scoring System

The objective of RF management is to ensure the availability of this limited resource for the deployment of new programs or services. For example, if an operator wants to add additional downstream DOCSIS

bandwidth, that will need to come from the current RF spectrum. To this effect, the RF spectrum scorecard takes into account the existing channels on the spectrum, looks at the planned gains and needs for a period, then computes the number of channels available by the end of the period. For this scorecard to be effective, regular collaboration with boundary partners in charge of equipment or programming that need spectrum is required. Such boundary partners include video engineering, DOCSIS engineering, telephony engineering, product, marketing, as well as government and public affairs departments.

The spectrum scorecard considers each market or spectrum boundary, looks at the total number of subscribers served in the market, then creates a weight that is equivalent to the proportion of subscribers. Based on the spectrum needed by market, the total spectrum capacity remaining by the end of the period is calculated. If the number of RF channels remaining at the end of the period is positive, the market is green and scores its weight. If this number is negative, the market is red, scoring a quarter of its weight. If it is zero, the market is yellow, scoring $\frac{3}{4}$ of its weight. The total system score is the sum of all the market scores. It is green if above 95%, red if below 90%, and yellow if between 90% and 95%. This score is later standardized using the system's normalization algorithm and reported in the system scorecard. Figure 11 shows an example of RF spectrum scorecard.

Downstream Spectrum Allocation (expressed in 6 MHz carriers)

	Market 1	Market 2	Market 3	Market 4		Market 11	Market 12	Market 13	Market 14	Market 15
QAM CATEGORY										
Analog Carriers	59	62	62	58		59	58	59	57	60
SD Carriers	12.8	11.7	11.1	11.3		24	23	23.2	23.8	24.2
HD Carriers	32.5	33.5	34	33.5		38.5	40.5	39.5	39	42
3D Carriers	0.5	0.5	0.5	0.5		0.5	0.5	0.5	0.5	0.5
VOD Carriers	5	5	5	5		4	4	4	4	8
SDV Carriers	8	8	8	8		0	0	0	0	0
DOCSIS	8	8	8	8		5	8	8	8	8
Circuit Switched Telephone	0	0	0	0		0	0	0	1	1
Control/Signaling	1	1	1	1		1	1	1	1	3
Unusable	0	0	0	0		0	0	0	0	0
Current QAMs Used/Unusable	126.8	129.7	129.6	125.3		132	135	135.2	134.3	146.7
Max QAMs	134	134	134	134		152	152	152	152	153
Current Unused QAM capacity	7.2	4.3	4.4	8.7		20	17	16.8	17.7	6.3
Planned Spectrum Usage	-1.5	-1.5	-1.5	-1.5		-9.5	-6.5	-7.6	-6.5	-4
Planned Spectrum Gains	1	1	1	1		3	3	3	3	1
Surplus/Deficit QAM Capacity	6.7	3.8	3.9	8.2		13.5	13.5	12.2	14.2	3.3

Subs served				
Weight	3%	1%	6%	1%
Downstream Score	3%	1%	6%	1%

... Truncated ...

2%	3%	7%	22%	34%	100%
2%	3%	7%	22%	34%	100%

Upstream Spectrum Allocation (expressed in 1 MHz carriers)

	Market 1	Market 2	Market 3	Market 4
Status Monitoring	0.3	0.3	0.3	0.3
Sweep Return	0.8	0.8	0.8	0.8
Digital Box Return	0.384	0.384	0.384	0.384
Circuit Switched Telephone	0	0	0	0
CHSI Return	9.6	9.6	9.6	9.6
Upstream Spectrum Usage	11.084	11.084	11.084	11.084
Max Upstream Bandwidth (MHz)	37	37	37	37
Available Upstream Bandwidth (MHz)	25.916	25.916	25.916	25.916
Upstream Spectrum EOY Planned Need	0	0	0	0
Upstream Spectrum EOY Planned Gain	0	0	0	0
Surplus/Deficit Upstream Capacity	25.916	25.916	25.916	25.916

Market 11	Market 12	Market 13	Market 14	Market 15
0.3	0.3	0.3	0.3	0.496
0.8	0.8	0.8	0.6	0.9
0.192	0.192	0.192	0.192	0.384
0	0	0	5.632	14.636
9.6	19.2	19.2	9.6	16
10.892	20.492	20.492	16.324	32.416
37	37	37	37	35
26.108	16.508	16.508	20.676	2.584
0	0	0	0	0
0	0	0	0	0
26.108	16.508	16.508	20.676	2.584

Upstream Score	3%	1%	6%	1%
-----------------------	-----------	-----------	-----------	-----------

2%	3%	7%	22%	34%	100%
----	----	----	-----	-----	-------------

Figure 11 - The RF Spectrum Scorecard

8. The System Capacity Scorecard

The system capacity scorecard standardizes individual scorecards previously presented to an overall system view. According to the online version of Oxford Dictionary, a standard is something used as norm or model for comparative evaluations and standardization consists of making something conform to an established standard (Oxford University Press, 2017). For the overall system scorecard, a standard was created and helped transcribe scores from individual scorecards that use different grading systems to the established standard. In the standard scorecard, a green score or healthy sub-system is graded between 90 and 100. A yellow sub-system is graded between 80 and 90, and a red sub-system is graded below 80. A

representation of the score converter is shown in figure 12. The scale of the subsystem scorecard is stored in the column labelled “From”, and that of the system scorecard is stored in the column labelled “To”. The subsystem’s score is entered in the “Current Score” and automatically computes the “Score to Report”. Figure 13 shows a view of the system scorecard corresponding to a standardized summary of each of the scorecards presented previously along with a trend line representing the evolution of the scores for the previous three months.

Video

Current Score

94.50

→

Score to Report

89.00

	From	To	Score
Green Max	100.000	100	
Green Min	95.000	90	
Green Increment	0.00500	0.01000	
Yellow Max	95.000	90	
Yellow Min	90.000	80	89.00
Yellow Increment	0.00500	0.01	
Red Max	90.000	80	
Red Min	0.000	0	
Red Increment	0.09000	0.08	

Figure 12 - Score Standardization

SYSTEM CAPACITY SCORECARD

CAPACITY SCORECARDS	Weight	Owner	Jul-yy	Aug-yy	Sept-yy	Oct-yy	Trend
DATA	15%	Data Engineer	97.96	98.98	98.98	97.96	
TRANSPORT	15%	Transport Engineer	87.61	88.28	88.28	88.28	
CMTS	15%	CMTS Engineer	86.38	83.13	91.59	93.70	
FACILITIES	15%	Technical Facilities Eng.	89.52	89.47	89.95	89.95	
RF SPECTRUM	15%	Video Engineer	100.00	100.00	100.00	100.00	
VIDEO (VOD, SDV)	15%	Video Engineer	77.24	78.71	80.38	82.19	
TELEPHONY	10%	Telephone Engineer	80.21	80.83	80.83	78.44	
SYSTEM CAPACITY TOTAL			88.83	88.87	90.46	90.66	

Green	Between 90 and 100
Yellow	Between 80 and 90
Red	Less than 80

Figure 13 - The System Scorecard

5. Analysis and Troubleshooting

The standardized system scorecard serves as a levelling tool in representing all scorecards using a common metric. By itself, it helps understand the evolution of the network capacity over the months and identify areas of concern. For example, a category that consistently stays red or one that suddenly becomes yellow or red such as telephony in figure 13 calls for explanations. For that reason, there is a process to share the findings of the scorecard in a meeting setting.

On a monthly basis, the Network Operations leaders and boundary partners meet for a presentation of the network's health. To make the meeting interesting to everyone, leaders of the different networks are also owners of their respective scorecard findings and action items, even though the scorecards are created by engineers in the planning team. The presentation of the scorecards is done in the form of analysis structured around three main points:

- Month to month performance comparison and improvements implemented
- Areas of concern and lessons learned
- Plan to stay or become green (projects to kick off, processes to change, etc.)

The structure presented above helps organize discussions and provides guidance to properly analyze the findings of the scorecards. Regardless of the system's color, each red item in the subtending scorecard needs to be identified, understood, and resolved in order to make the system become or stay green. For that same reason, items marked yellow should be discussed and monitored.

When an issue of interest is identified, the planning engineer creates trouble tickets and assigns them to a subject matter expert (SME) or a person selected by the technology owner. The technology owner is responsible for following the tickets to resolution with help from the planning engineer, and reports from the trouble tickets are shared with the teams during various project and team meetings.

8. Areas of Further Investigation

While the deployment of a standardized scorecard allows the operator to be proactive and resolve issues prior to customer complaints, its efficiency depends on the identification of the correct metrics and follow-up of issues to their complete resolution. In other words, time is spent solving a problem that is not yet real to the end users and the company needs to identify and assign appropriate resources to undertake such tasks.

Even though the scorecard presented here obtains all its data from automated sources, the individual scorecards are compiled using Microsoft Excel. This is time consuming and provides an opportunity for improvement. The main issue faced with such opportunity is the vast number of platforms to query and assemble data. Indeed, all the subtending scorecards collect data from different sources requiring different types of expertise.

In addition to automation, the scorecard needs to be adapted for new technologies. For example, the deployment of DOCSIS 3.1, fiber deep, and remote PHY would create various levels of complexity due to the number, size, and structure of service groups. Further research needs to be conducted to confirm this assertion.

To make the standardize scorecard system more efficient, there needs to be a correlation between the improvement seen on the scorecards and customer experience measured in terms of number and types of trouble tickets, Net Promoter Score (NPS), or customer satisfaction as discussed by Lartey, Hargiss, and Howard (2015).

Conclusion

This paper presented a method for monitoring an entire broadband or telecommunications network using standardized scorecards. The main objective in doing so is to proactively resolve network issues prior to negative customer experience. In doing so, the paper first presented a literature review on measurement theory, followed by an overview of the broadband network. The overview resulted in the identification of key performance metrics in each of the broadband networks and critical infrastructures. These networks and infrastructures included technical facilities, data network, transport network, DOCSIS network, video network, telephony network, and RF spectrum. After presenting the rationale for the selection of the key metrics on each network, the paper presented detailed procedures on how to create the respective scorecards, along with real example of each of the scorecards. Finally, the paper presented an operational analysis and troubleshooting process before ending with a discussion on areas of further investigation.

Regardless of the level of automation, human intervention is always required to analyze issues and implement corrective measures or else, existing issues can become chronic problems defined by Sasisekharan, Seshadri, and Weiss (1994) as “problems that are likely to continue in the immediate future

without diagnosis and repair” (p. 453). Indeed, leaving little issues unresolved will certainly result in bigger issues in the future. To keep networks at peak performance, humans need to understand their current state and take action for their future evolution.

Abbreviations

AC	alternating current
AIN	advanced intelligent network
AP	access point
BHCA	busy hour call attempts
bps	bits per second
CCAP	converged cable access platform
CMTS	cable modem termination system
CS2K	CallServer 2000
DACS	digital cross-connect system
DC	direct current
DCS	digital cross-connect system
DMS	digital multiplex system
DOCSIS	data over network system interface specification
DSG	DOCSIS set-top gateway
DVR	digital video recording
DWDM	dense wavelength division multiplexing
EMS	element management system
FEC	forward error correction
GHz	gigahertz
GWC	gateway controller
HD	high definition
HFC	hybrid fiber-coax
HVAC	heating, ventilation, and air conditioning
Hz	hertz
IPTV	Internet protocol television
ISBE	International Society of Broadband Experts
ISP	inside plant
IWSPM	interworking spectrum peripheral module
KHz	kilohertz
KPI	key performance indicator
MHz	mega hertz
MSO	multiservice operator
MSSP	multiservice switching platform
MTC	master telecommunications center
NE	network element
NPS	net promoter score
nPVR	network-based personal video recorder
OC	optical carrier

ONE	optical network element
OSP	outside plant
RDC	regional data center
RF	radio frequency
ROADM	reconfigurable optical add drop multiplexer
RTU	right-to-use
SCTE	Society of Cable Telecommunications Engineers
SDH	synchronous digital hierarchy
SDV	switched digital video
SG	service group or serving group
SIP	session initiation protocol
SME	subject matter expert
SOC	software optionality control
SONET	synchronous optical network
SPM	spectrum peripheral module
STB	set-top box
STC	secondary telecommunications center
STS	synchronous transport signal
TBSCT	telecommunication building critical system tool
TDM	time division multiplexing
TNPM	Tivoli Netcool Performance Manager
TVCID	television caller identification
UPS	uninterruptible power source
VOD	video on demand

Bibliography and References

- Busch, P., & Lahti, P. (2008, July). Measurement theory. *Philosophy of Science*. Retrieved from <http://philsci-archive.pitt.edu/4108/1/measurement-theoryBuschLahti%28v2a%29.pdf>
- Cisco. (n.s.). *Cisco ONS 15600 Series*. Retrieved from Cisco.com: <http://www.cisco.com/c/en/us/products/optical-networking/ons-15600-series/index.html>
- Fujitsu. (n.s.). *Fujitsu Network FLASHWAVE 9500*. Retrieved from Fujitsu.com: <http://www.fujitsu.com/global/products/network/products/flashwave-9500/index.html>
- Hand, D. J. (1996). Statistics and the theory of measurement. *Journal of the Royal Statistical Society*, 159(3), 445-492. Retrieved from <http://www.lps.uci.edu/~johnsonk/CLASSES/MeasurementTheory/Hand1996.StatisticsAndTheTheoryOfMeasurement.pdf>
- Hass, J. (2003). *What is frequency?* Retrieved from An Acoustics Primer, Chapter 5: <http://www.indiana.edu/~emusic/acoustics/frequency.htm>

- Lartey, F. M. (2015). *Increasing promoters in the residential broadband service industry: Relationship between customer satisfaction and loyalty using ordinal logistic regression*. ProQuest LLC. doi:UMI 3682580
- Lartey, F. M., Hargiss, K., & Howard, C. (2015). Antecedents of customer satisfaction affecting broadband loyalty: An implementation of SERVQUAL and NPS. *International Journal of Strategic Information Technology and Applications*, 6(1), 27-42. doi:10.4018/IJSITA.2015010103
- Lartey, F. M., McGinn, R., & Diponzio, N. (2016). Reducing the cost of fiber-to-the-home brownfield deployment through the implementation of core extraction on an HFC network. *Journal of Network Operations*, 1(2), 78-90. Retrieved from www.scte.org
- Marut, D. (2016). How to apply statistical monitoring, online learning, and measuring & control mechanisms to use less energy in a non-homogeneous network. *Cable-Tec Expo. '16*. Philadelphia, PA: SCTE•ISBE.
- Noll, K. A. (n.s.). *Hybrid fiber-coaxial networks: Technology and challenges in deploying multi-gigabit access services*. Nanog.org. Retrieved from <https://www.nanog.org/sites/default/files/08-Noll.pdf>
- Oxford University Press. (2017). *Standard*. Retrieved from Oxford Living Dictionaries: <https://en.oxforddictionaries.com/definition/standard>
- Rasmussen, N. (2012). *Power and cooling capacity management for data centers*. The Schneider Electric White Paper Library. Retrieved from http://www.apc.com/salestools/NRAN-6C25XM/NRAN-6C25XM_R3_EN.pdf
- Roberts, F. S. (2009). Measurement theory with applications to decisionmaking, utility, and social sciences. In G.-C. Rota, *Encyclopedia of Mathematics and its Applications*, vol.7 (pp. 1-420). New Brunswick, NJ: Cambridge University Press.
- Sasisekharan, R., Seshadri, V., & Weiss, S. M. (1994). Proactive network maintenance using machine learning. *AAAI Technical Report*, 453-462.
- Schmitt, M. (2012, March 12). *Cable network overview*. Retrieved from [ieee802.org: http://www.ieee802.org/3/epoc/public/mar12/schmitt_01_0312.pdf](http://www.ieee802.org/3/epoc/public/mar12/schmitt_01_0312.pdf)
- Snow, A., & Weckman, G. (2016). Statistical Methods for System Dependability: Reliability, Availability, Maintainability and Resiliency. *NexComm*. Lisbon, Portugal: International Association of Research and Industrial Academia (IARIA).
- Sundaresan, K. (2015). Evolution of CMTS/CCAP Architectures. *Spring Technical Forum*, pp. 1-8. Retrieved from <http://www.nctatechnicalpapers.com/Paper/2015/2015-evolution-of-cmts-ccap-architectures>
- Tao, T. (2011). *An introduction to measure theory*. Providence, Rhode Island: Library of Congress.
- Urban, D. (2010). Comparison of techniques for HFC upstream capacity increase. *Spring Technical Forum Proceedings*, 153-178. Retrieved from www.nctatechnicalpapers.com

Vasudevan, S. V., Liu, X., & Kollmansberger, K. (2008). IPTV architectures for cable systems: An evolutionary approach. *IEEE Communications Magazine*, 102-109.
doi:10.1109/MCOM.2008.4511657

Watts, K. (2011). *Best practices for proactively maintaining your return paths*. Indianapolis, IN: JDSU.
Retrieved from SCTE:
<http://www.gcsccte.org/presentations/2011/SCTE%20Best%20Practices%20for%20Proactively%20Maintaining%20Your%20Return%20Paths%20v10.pdf>

Simplifying Field Operations Using Machine Learning

Applications of Machine Learning to Multiple System Operators (MSOs)

A Technical Paper prepared for SCTE•ISBE by

Sanjay Dorairaj

Senior Director, Comcast Innovation Labs
Comcast Corporation
1050 Enterprise Way, Sunnyvale, CA 94089
215-900-2349
Sanjay_Dorairaj@cable.comcast.com

Chris Bastian

CTO, SCTE
140 Philips Road, Exton, PA 19341-1318
cbastian@scte.org

Bernard Burg

Sr Manager, Data Science, Comcast Applied Ai
Comcast Corporation
1050 Enterprise Way, Sunnyvale, CA 94089
Phone: 408.900.8575
Email: Bernard_Burg@cable.comcast.co

Nicholas Pinckernell

Data Scientist
Comcast Corporation
1899 Wynkoop #550, Denver, CO 80202

Introduction

Every so often in the history of our evolution, humans discover something so important that it propels us into a new plane of technological and intellectual superiority. Over two million years ago, the **Stone Age** helped us build tools that established us as the dominant species on this planet. Much later, the **Bronze Age** (circa 3500 BC) and the **Iron Age** (circa 1200 BC) catapulted us to new levels of technological sophistication through the introduction of coin-based currencies, faster means of transport, durable manufacturing and construction and numerous other developments. This laid the foundation for the **Industrial Age** (circa 1700 AD), which ushered in the age of mechanized agriculture, mass transportation and electronic communication. The invention of the computer and the internet in the later parts of the 20th century heralded the dawn of the **Internet Age**. Individuals anywhere on the globe could now communicate and exchange information with one another. And much like Ray Kurzweil's Law of Accelerating Returns [1], the Internet Age is hardly over. Now, we find ourselves at the cusp of two back to back, tightly coupled events that are also bound to be of equally great historical significance - the **Age of Big Data** and the **Age of Machine Learning**.

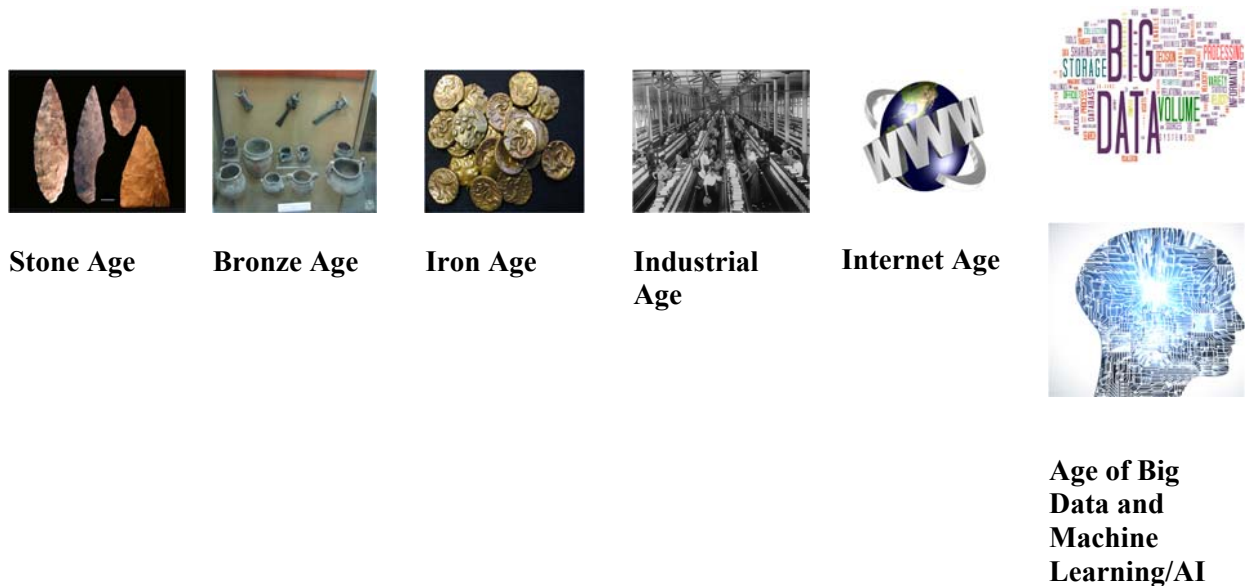


Figure 1 - From the Stone Age to the Age of Big Data and Machine Learning

The explosion in data aka “Big Data”, is a direct result of the exponential improvements in computing power and storage, with similar decreases in their cost [2]. This fueled an abundance of both personal and organizational data. The chart, below, provides a dramatic portrayal of the rapid growth of data over just one decade. Despite all of this data, the insights that we were able to generate has been limited by decades-old statistical and mathematical techniques and there wasn't much innovation in this field. The advent of Machine Learning has propelled us forward, by offering techniques that transform the big data into a veritable gold mine of valuable insights.




Figure 2 - Cambrian Explosion of Data (Source: Patrick Cheesman)

This paper is about machine learning - its definition and its applications. It especially examines the relevance of machine learning from the perspective of the cable's multiple system operators (MSOs). While there have been some attempts in technical and trade literature to pinpoint the benefits of machine learning to cable service operators, there has not yet been a holistic treatment of the subject, to our knowledge. This paper is an attempt to fill that gap.

1. Machine Learning Overview

Definitions of machine learning tend to compare it with traditional statistical methods. Leo Breiman, one of the pioneers and early evangelizers of machine learning, talked about the two cultures of statistical modeling - the data modeling culture and the algorithmic culture [3]. In the **data modeling approach**, which could be compared to traditional statistical approaches, the model assumes an underlying stochastic process. Inferences are made using techniques such as linear and logistic regression. Sample sizes are determined based on concepts founded in probability and inferential statistics and generally tend to be a tiny portion of the population size. **Machine learning** or the **algorithmic approach**, on the other hand, does not assume the existence of a well-defined process to the underlying data. Instead, it treats the model as a black box. Machine Learning algorithms such as neural networks and decision trees try to decipher the underlying patterns in the data using methods similar to that of maximum likelihood estimation. These algorithms typically require large amounts of data to yield good predictions.

Table 1 - Data Modeling vs Machine Learning approach

	Data Modeling Approach	Machine Learning Approach
Sample data requirements	Low	High
Constraints	Several	Few
Validation	Goodness of fit, residual examination	Performance on an independent test data set
Multiple variable prediction accuracy	Low	High
Data Interpretation characteristics	Linear or curvilinear patterns that can be approximated as functions	Complex non-linear patterns

Traditional inferential statistics has found its niche in several areas such as predicting an election outcome or predicting the effects of a new medication on a population. They perform well when the number of predictor variables are low. As the number of predictor variables increase, these models tend to break down. This is because of the large number of constraints these models are required to satisfy to yield valid predictions [4]. As the number and diversity of predictor variables increase, it becomes more and more difficult for these constraints to be met. On the other hand, machine learning algorithms are capable of dealing with complex processes and millions of predictor variables. The key requirement for machine learning to be successful is a data-rich environment, and the explosion of data in organizations today has proven to be instrumental in the increasing popularity and success of machine learning.

What role does machine learning play in Artificial Intelligence (AI)? AI is an overarching term that encapsulates all attempts to instrumentalize technology with the ability to think and act independently, much like humans do. It refers not only to the software and algorithms that renders this capability but the hardware and control systems as well. Machine Learning can be viewed as the subset of AI technologies that deals with pattern recognition.

A crucial advantage that humans have over existing computing platforms is our ability to make inferences from a complex set of input events. For example, our eyes are sophisticated enough to visually process information in three-dimensional space and recognize objects and emotions with little difficulty. Another example is our ability to look at a multi-variable time-series chart and immediately identify the anomalies present. The intelligence that enables us to excel at these tasks can be traced down to our uncanny ability to leverage our historical knowledge to perform real-time pattern matching. Machine learning and its derivative technology – Deep Learning, render computing platforms with pattern matching skills. In some cases, they are far superior to humans because they can process numerous parameters and complex underlying processes in an almost unbounded manner, limited only by computing and storage costs. In addition to the strong reliance on mathematics and statistics, machine learning is also strongly tied to software development, since the amount of data that it needs to be successful requires the use of state of the art software development methodologies.

The below diagram succinctly captures the overlapping areas of knowledge that data science comprises of - computer science, mathematics, statistics and domain expertise. The *unicorns* in the middle refer to those data scientists who possess the rare combination of all these skills.

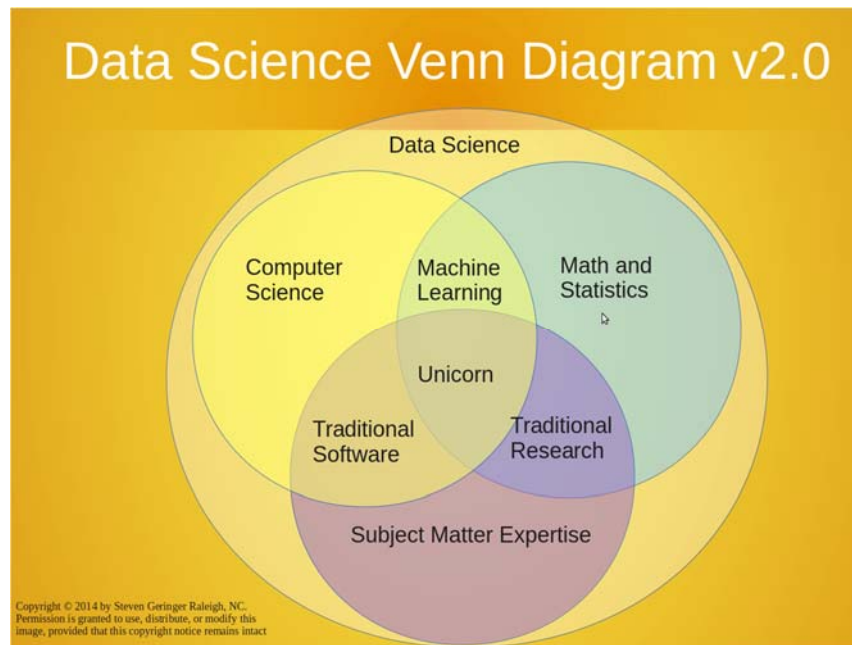


Figure 3 - Data Science Venn Diagram

2. Applications of Machine Learning for MSOs

In this section, we define general classes of machine learning algorithms and discuss how these classes of algorithms can add value to service providers.

The general classes of machine learning algorithms

1. Classifiers
2. Clustering Algorithms
3. Recommender Systems
4. Anomaly Detection Algorithms
5. Linear Regression

Classifiers are used to discern similarities among sets of data and assign them to categories based on their similarity. Examples of classification could be identifying objects in a video frame, identifying the underlying sentiment in a customer service message – happy, upset or neutral, or, associating a log message from a set-top to a specific error class. The technologies powering classifiers range from the simple - decision trees and random forest, to the very complex - deep neural networks. The choice of technologies used are typically functions of the level of complexity and the number of features in the underlying data. Image classification has been shown to benefit greatly by technologies derived from neural networks such as convolutional neural networks (CNNs).

Clustering algorithms group similar data into clusters. They are typically used to group data that share similar characteristics or to look for significant deviations in data. For example, clustering algorithms could be used to look at smart home data and create user profiles based on shared behavioral characteristics – for example, early risers, late risers and so on. Clustering algorithms range from the simple such as K-Means

clustering to more advanced algorithms such as agglomerative hierarchical clustering that may require additional tuning for optimal performance.

Recommender systems are a class of algorithms that make user or product recommendations based on historical usage or behavioral data. They can be used to suggest movies to users based on what those users have watched in the past or based on what users with similar viewing habits may have watched. For example, if two users like *Star Wars* and one of the users has watched *Dark Matter*, another sci-fi series, then the recommender would suggest *Dark Matter* to the other user. In a similar way, they can also be used to recommend products that users would like to purchase. In the case of customer service, they can recommend actions that the customer service representative can take in each situation based on past actions. A popular method of building these recommendations is using an algorithm called Collaborative Filters.

Anomaly Detection algorithms are similar to classification algorithms except that they typically only deal with cases where there just two classes of data exist and where one class occurs with an extremely low frequency. If the anomalies are relatively large, then clustering algorithms can be used; however, if anomalies are very few, joint probabilistic methods to model the rare events are more appropriate. Anomaly detection can be used to look for events such as billing fraud and device errors in cases where device failure is rare.

Linear Regression algorithms are used to make predictions about continuous variables. An example could be predicting customer churn rate or predicting bandwidth utilization. Linear Regression and Classifier algorithms share similar characteristics with respect to the technologies that are used. Where they differ is while classifier algorithms are designed to maximize the separation between dissimilar data points to allow for classes to be determined, linear regression algorithms interpret results in a continuous manner. One other point to note is that ML-based linear regression models are typically interpolative, traditional statistical linear regressions models are both interpolative and extrapolative. This only points to usage and does not imply that the traditional model is superior to the ML-model in cases where extrapolation is required.

Table 2 summarizes the above discussion.

Table 2 - Machine Learning algorithms

Class of Algorithms	Description	Technology Examples	Applications
Classifiers	Assigns data to categories based on similarity to other data.	Random Forest and Neural Networks	Sentiment Analysis, Image Classification
Clustering Algorithms	Groups similar data into clusters	K-Means, Hierarchical Clustering	User profiles and anomaly detection
Recommender Systems	Make recommendations based on historical data	Collaborative Filtering	Product recommendations
Anomaly Detection	Identify rare events	Joint Probabilistic modeling	Billing fraud detection
Linear Regression	Predict values for continuous variables	Linear Regression	Churn rate prediction

Discussed below are a few machine learning concepts and ideas that are also important to the successful application of machine learning.

2.1. Supervised versus Unsupervised Learning

As mentioned earlier, machine learning algorithms seek to find underlying patterns in data and mathematical ways of representing those patterns. The mathematical representation is referred to as a **model**. This search for patterns leads to two broad classes of machine learning – **supervised** and **unsupervised** learning. Given any set of data, if a machine learning algorithm is asked to determine underlying patterns in an autonomous manner, then that form of machine learning is known as unsupervised learning. Examples of unsupervised learning are (1) building consumer behavior profiles from customer call data and (2) classifying defect data into groups based on similarity between defects. Supervised learning is guided learning. In this case, the data also includes a parameter known as a **label** that captures the system response to a given set of input parameters. In determining customer churn for example, the available data will include the number of issues seen by a customer on a set-top on any given day. These are the predictor variables. In addition to the predictor variables, supervised algorithms require a label that could indicate whether the customer tried to cancel service that day. Supervised algorithms can be used to build a model that can predict the probability of a customer cancellation from the predictor and response variables. These algorithms have been shown to be effective in improving customer diagnostics, optimizing call centers, increasing the efficiency of truck rolls, and pro-active network healing.

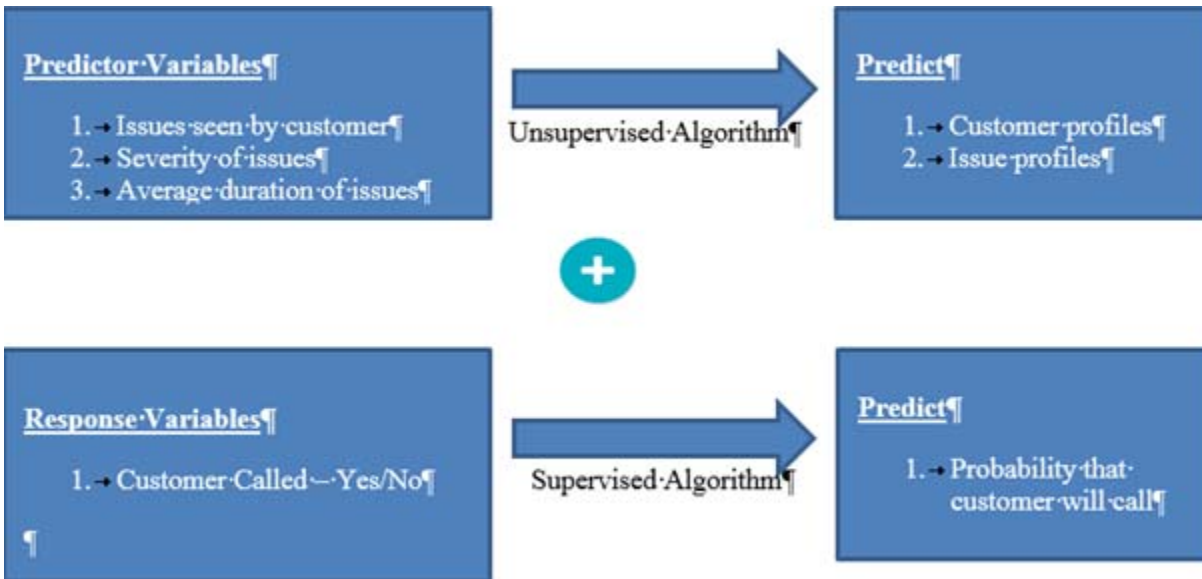


Figure 4 - Data Requirements for Supervised and Unsupervised Algorithms

2.2. Training set, Test set and Hyperparameters

Data used to build machine learning models is typically broken down into subsets - the **training data** and the **test data**. Training data is used to train the algorithm and allow it to build a model for the underlying data. Typically, the algorithm contains a number of tunable parameters, called **hyperparameters**, that are used to optimize the performance of the model. For example, when trying to use a clustering algorithm to build customer profiles, one of the hyperparameters is the number of clusters. In our examination of the resulting clusters from such an algorithm, we may notice that a certain cluster count yields a more optimal set of clusters than another cluster count. In a similar way, other hyperparameters can also be tuned till an optimal model is obtained. A key success factor for a machine learning model is to ensure that the training set and the test set are kept completely separate. This ensures the absence of any kind of bias during model generation. For this reason, hyperparameter tuning is not done using the test set, but rather, the training data is subdivided into a training set and a cross-validation set, and the cross-validation set is used to validate hyperparameters.

2.3. Feature engineering

There are two types of parameters that come into play with machine learning. The first type is referred to as a predictor variable and the second type is referred to as a response variable. Predictor variables are variables that are used to make predictions and response variables are the prediction. In image recognition, for example, pixels in an image are the predictor variables and the predicted class (cat, dog, flower etcetera) is the response variable. Similarly, when predicting the likelihood of a customer call, predictor variables could include the state of the set-top box modem and state of the infrastructure. In this case, whether the customer called, given the set of predictor variables, would be the response variable.

The selection of predictor variables is a crucial part of machine learning since the quality of the predictor variables ultimately determines the quality of the prediction. Predictor variables are also referred to as **features**. Feature selection is in itself a complex process and it has spawned a whole separate branch of

machine learning called **feature engineering**. Feature engineering usually involves two types of activities (1) Reducing the set of all possible features into a set of features suitable that are better predictors of the output class and (2) Transforming or extending the set of available features with new features that are more suitable for the particular machine learning task. A popular method to transform one set of parameters into a smaller set of better predictor variables is called Principal Components Analysis (PCA).

2.4. Ensemble approaches

Often, when doing machine learning, the algorithms taken separately do not yield the best results. However, when combined with other machine learning algorithms or even multiple instances of the same algorithm, the quality of the results tends to improve, this is referred to as the **ensemble approach** to machine learning and this method is quickly gaining popularity in the machine learning community. The random forest algorithm is such as example. Sometimes, results from different algorithms such as random forest and support vector model (SVM), may be combined to yield a better classifier. Software tools include features that offer the programmatic selection of the best ensemble models through trial and error. A successful demonstration of the ensemble approach is the Netflix Prize which went to a team of machine learning engineers that developed the best algorithm using a similar ensemble approach [5].

2.5. Online versus offline algorithms

In certain cases, machine learning models may need to be built in real-time or **online** mode. For example, recommender systems need to process incoming events in real-time and provide recommendations based on the current state of the system. In this case, the model will need to be updated in real-time to ensure that the recommendations are up to date. In cases such as anomaly detection however, it may not be necessary to build a real-time model and an **offline** model is sufficient. In this case, models are built when data is available and refreshed with lesser frequency, perhaps on the order of weeks or months.

Depending on the type of application, an online or an offline model may be required. Not all machine learning algorithms work in an online model, so therefore, if choosing the online learning route, it is important that an algorithm that supports online learning is selected.

3. Operational Efficiency Improvements Using Machine Learning

As discussed above, there are several applications to machine learning. Some of the applications such as recommender systems, campaign management systems, market analysis and so on are revenue generating. Other applications have to do with cost optimization. These include customer call prediction, churn prediction, fault prediction, capacity planning and so on. In this section, we focus on the potential for machine learning to improve the operational efficiency of an organization.

Listed below are a set of reasons establishing how machine learning can help with operational efficiency goals.

- Cable system operators have a lot of data sources (understatement!) with valuable information about the state of the system
- These data sources are currently used only for basic- to medium-level analytics tasks, such as relative frequency comparison, difference computations and advanced visualizations.
- Predictive analytics using machine learning can help flag customer service issues in advance, presenting operators an opportunity to fix them before they disaffect service

- Machine learning tools can also be used to perform root cause analysis to identify underlying issues and recommend remediation actions
- When ML insight is deployed in development and field tools, it helps drive down call volume and truck rolls, thereby decreasing operational costs related to these activities

3.1. Machine Learning – A New Operations Paradigm

Machine Learning is a new paradigm of operations. This is especially true for field technicians who stand to benefit the most from this tool. Field technicians are used to certainty. When a DOCSIS monitoring device is plugged into an outlet, the expectation is that the spectral signature that they see is exactly what is present. The same goes for other measures such as signal loss, signal-to-noise ratio, signal levels and so on. This is a deterministic paradigm where what is reported is exactly as it is.

Machine Learning solutions are different. They do not provide answers that are a 100 percent guaranteed to be true. What you get is an answer and a probability associated with that answer being true. For example, in the case of a spectral impairment, the machine learning solution may say that there is a 95 percent chance of the signal containing a wave impairment. How should the field technician or the network operations center react to a probabilistic result? There are known methods of handling uncertainty and these are all based on an application's aversion to false positives.

Evaluating performance of machine learning models involves balancing cost reduction, customer satisfaction and model complexity. A large volume of repair calls implies that **small improvements can yield sizeable cost saving**. Consider the below example

- 1 million repair calls a month at a hypothetical \$10 per call implies a monthly cost of \$10m per month.
- A **1 percent reduction** results in a 100-thousand-dollar monthly saving and an annual saving of approximately **1.2 million dollars**

Machine learning also provides a means for tuning the model to yield a desired false positive rate. Reducing the number of false positives would however drive down the number of true positives, so there is a tradeoff that must be made. The examples below show two use cases

- A destructive self-healing action such as a reboot would require higher precision; we should therefore minimize false positives to reduce disruption to the customer
- A non-intrusive self-healing action such as a billing change would allow for lower precision, as false positives influence the overall result in the same manner as false negatives.

Similarly, in the case of the spectral impairment detection case study considered in this paper, the machine learning algorithm would assign similar probabilities to each of impairment that it detects and the field operations and the network operations center should have a strategy to deal with this information in a meaningful manner.

4. Typical Development Methodology

The development and deployment of machine learning within an organization typically takes place as two parallel, though connected, workstreams. The **first workstream** 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 5 shows the two workstreams and Figure 6 shows a high-level view of the machine learning model lifecycle.

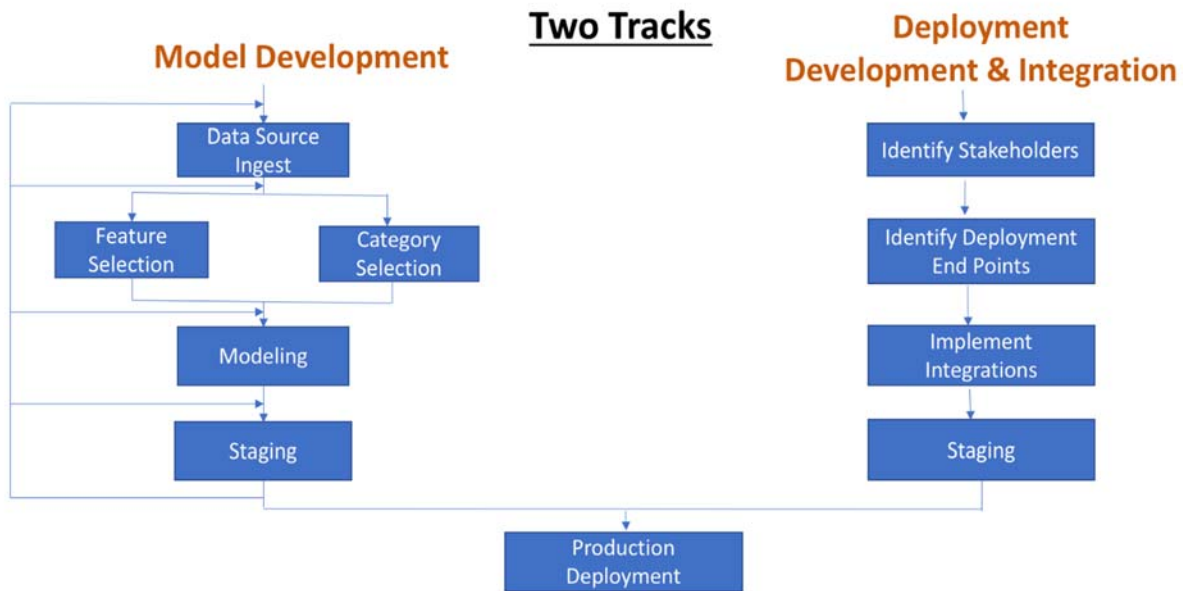


Figure 5 - Machine Learning Workstreams





Figure 6 - Machine Learning Process (Source: Crisp Industry Standard Process for Data Mining -CRISP DM)

5. Case Study: Spectral Impairment Detection

Cable operators monitor the use of the spectrum for every device (e.g. cable modem). Such measurements give a state of the communication between the network infrastructure and the device.

The goal of this method is to automatically characterize these spectra by labeling all their impairments. This is instrumental to: 1) Assess the performance of the RF spectrum, 2) Consider variation over time and temperature; 3) Standardize automation & detection of anomalies, and 4) Remove subjectivity and manual interpretation by technicians.

Experts have identified 15 impairments for which automatic detection would bring a competitive advantage. Each of these impairments exhibits an identified cause, and is linked to a repair action that improves the performance of the RF spectrum. For example:



is a *resonance peaking* caused by an amplifier problem



is an *adjacency/alignment* caused by a head-end problem



is a *suck-out* caused by an in-home problem

Figure 7 - A few RF spectrum impairment samples

The 15 plant-related impairments ripe for detection and subsequent correction include: Suck-outs, Notches, Tilt (and direction), Ripples / Waves, Off-Air Ingress, Foreign carriers, Wideband / Edison, Roll-off, Resonance / Peaking, Filters, Leveling, Adjacency / Alignment, Power Summary, Distortion / Intermod, and Pilot-to-Channel ratio.

6. Design Approaches

The accuracy of the spectral impairment detector currently in production is low, with only 5 impairments being detected. The new impairment detection described is significantly more accurate, targeting the detection of 10 of the 15 known impairments.

To enhance the accuracy of spectral impairments interpretation two methods are being pursued. Each of them will result in a much higher impairment classification accuracy.

Mathematical modeling: Spectral data is modeled through traditional signal processing methods, extracting features characterizing each of the 15 impairments in a direct, static mapping.

ML models: An ML algorithm learns dynamic mappings between the features extracted by the mathematical model and the impairments. As such, ML uncovers optimal solutions, fine-tuning each feature to its best use within a context. This comes at the cost of labeling huge quantities of data to perform the supervised learning. Addressing this issue, the team built a labeling engine to crowd-source labeling within Comcast.

6.1. Mathematical Modeling

Each of the impairments described in the figures below are detected by a corresponding set of features.

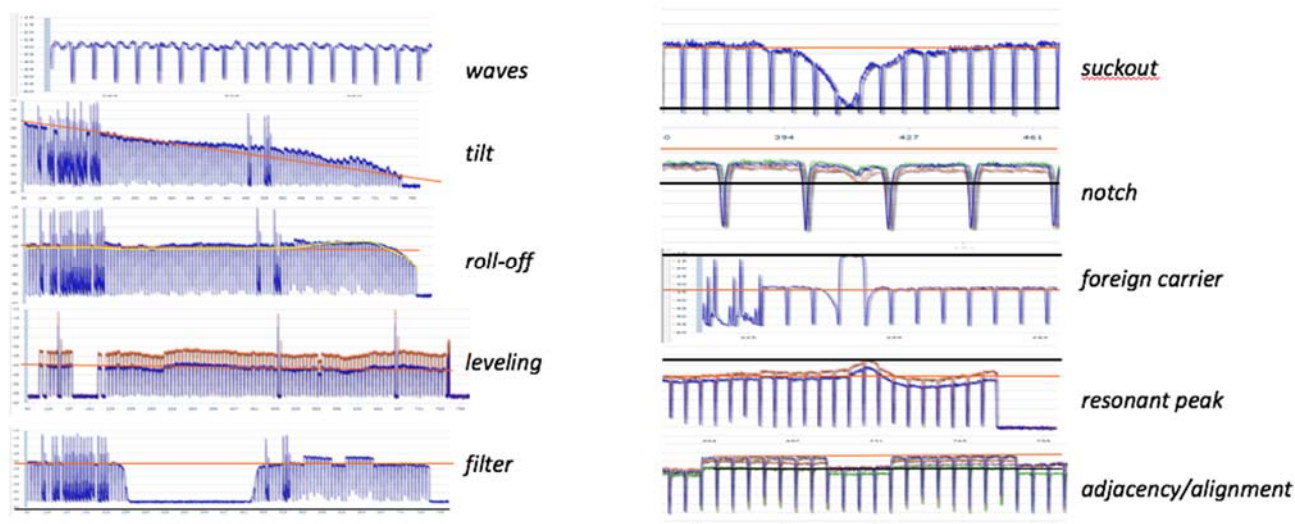



Figure 8 - Spectral Impairments and Their Shapes


Some of the impairments, like roll-off, filters or suck-outs are very impactful to end customers, even preventing them from accessing some channels. Other impairments, like waves, off-air-ingress or tilt slow transmissions down. All of these impairments are linked to known causes. Their diagnostic is key to the performance of Comcast's operations, and is of particular use to field technicians, because it allows them to pinpoint the cause of poor performance, or installation malfunctions.

The proposed approach is based on noise-resistant feature detection. Two data representations are used in parallel. The spectrum representation uses the complete spectrum, with a sampling at 117 kHz. The channel representation characterizes each TV channel which corresponds to a 6 MHz sampling. Channel representation is used and well understood by technicians.



Spectrum

Represents complete spectrum by samples of 117kHz



Channels

Represents each channel of 6MHz

Figure 9 - Spectrum, Channels and Features

The features are independent from each other and are oblivious to the frequency at which they appear, and their combination allows detection of impairments. Impairment detection methods are also independent from each other, allowing their results to be combined. Thus, this overall detection method allows fine tuning both features and impairments, independently. This flexibility permits the introduction of new features, as well as new methods for impairment detection, without affecting existing detection methods.

6.1.1. Feature Detection

The program extracts similar features for both channel and spectral representations.




Figure 10 - Feature Detection

A plateau, with frequencies between 120MHz and 750MHz, is considered for the analysis. A linear approximation of this plateau offers stable features $y = ax + b$ to assess the flatness of the plateau and its height. A similar approach is undertaken with higher degree approximations of the same plateau. From this plateau, positive and negative peaks are detected. The shape around these peaks is an important feature, as some of the peaks are formed by single channels, whereas others have parabolic shapes --illustrating that many channels are affected simultaneously.

6.1.2. Example of impairment detection: Tilt and roll-off

A tilt is well approximated by a linear signal. The roll-off, in contrast, shows a fast decrease of the signal amplitude at high frequencies, and is better modeled as a parabola.

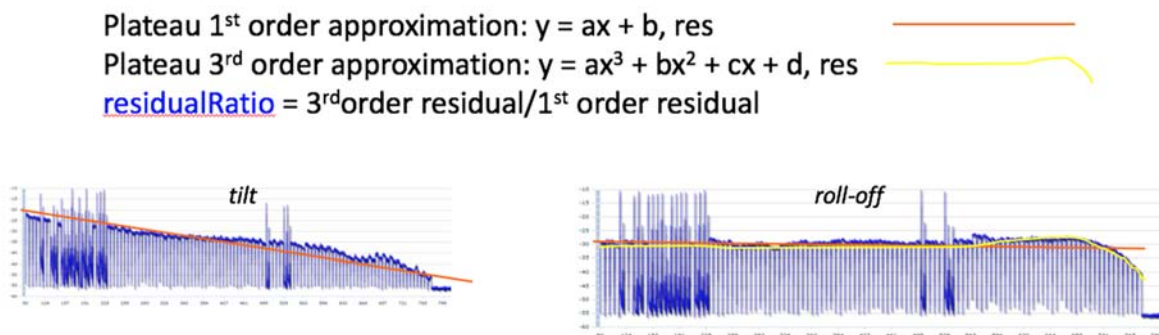


Figure 11 - Tilt and Roll-off detection

One solution differentiating a tilt from a roll-off is to make the ratio between the residuals of the 3rd order approximation and the 1st order approximation. If the ratio is near to 1, the impairment is a tilt since no fast decrease of the signal amplitude at high frequencies was detected.

Example of impairment detection: Suckout, notch, foreign carrier, resonant peak

From a feature extraction point of view, suckout and notch impairments are seen as signal dips, whereas foreign carrier and resonant peaks represent crests in the signal amplitude.




Figure 12 - Suckout, Notch, Foreign Carrier and Resonant Peaking Classification

The suckout is a large dip spanning several channels, whereas the notch is a tiny dip that cannot be seen in the channel view.

The foreign carrier is a sharp, single channel peak in the signal, whereas the resonant peak is a shallow peak spanning across several channels.

6.1.3. Results

The presented method returns a complete impairment diagnostic including all impairment instances detected on a spectral signal.

In Figure 13, the presented method discovers a combination of wave, tilt, suckout at 429MHz, and suckout at 445MHz. In the figure, suckouts are annotated with a red cross-circle:




Figure 13 - Prediction Example – Wave, Tilt and Suckout

One of the advantages of this mathematical modeling is its simplicity: Each of the impairments is linked to a few features extracted via signal processing. The fine tuning of these features warrants some experimentation and skills, e.g. for defining the threshold making the difference between the tilt and the roll-off. These parameters are largely independent and can be fine-tuned independently. But static tuning might not be the optimal solution.

6.2. ML Models: Towards an Optimal Solution

ML models can be used to bring mathematical models into a new dimension. Instead of having a finely tuned mathematical model working with static parameters -- like the threshold making the difference between the tilt and the roll-off -- imagine having an ML algorithm that *dynamically* fine tunes these parameters, according to the expected output. The great advantage of ML is that it uses algorithms such as linear regression and classifications to determine the best parameter settings. ML is capable of optimizing thousands of parameters that are far beyond the capabilities of what humans can fine-tune.

However, ML comes at a huge cost in this setting. ML works best in supervised learning, so, data needs be labeled. The features that were treated independently in the mathematical model are now mixed together, leading to a combinatorial explosion. Labeling 15 features into 6 buckets (e.g. none, tiny, small, medium, large, huge) leads to 15^6 possibilities = 11 million to hit each bucket at least once. This domain is most probably sparsely populated; however, this simple calculation shows that labeling data is a daunting task, way beyond human capabilities. At a first glance, hundreds of thousands of labeled data could and should be generated.

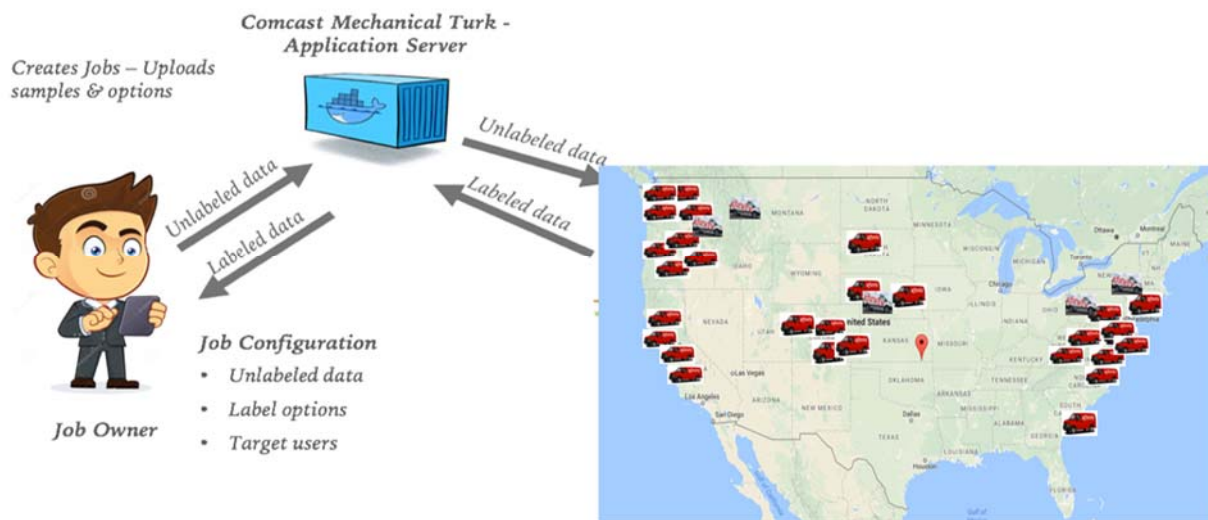


Figure 14 - Comcast Labeling Machine Solution Overview

The good news is, Comcast and other, like-minded MSOs have thousands of experts in the field capable of labeling this data. These are our industry's technicians. The idea is to farm the labeling task out through a "Turk mechanism" [6]. In this case, the unlabeled data is provided to the Comcast Mechanical Turk server that crowd-sources the data to be labeled to the technicians. At any time, an underutilized technician can access data and label it, through an API accessed by an app on their device.

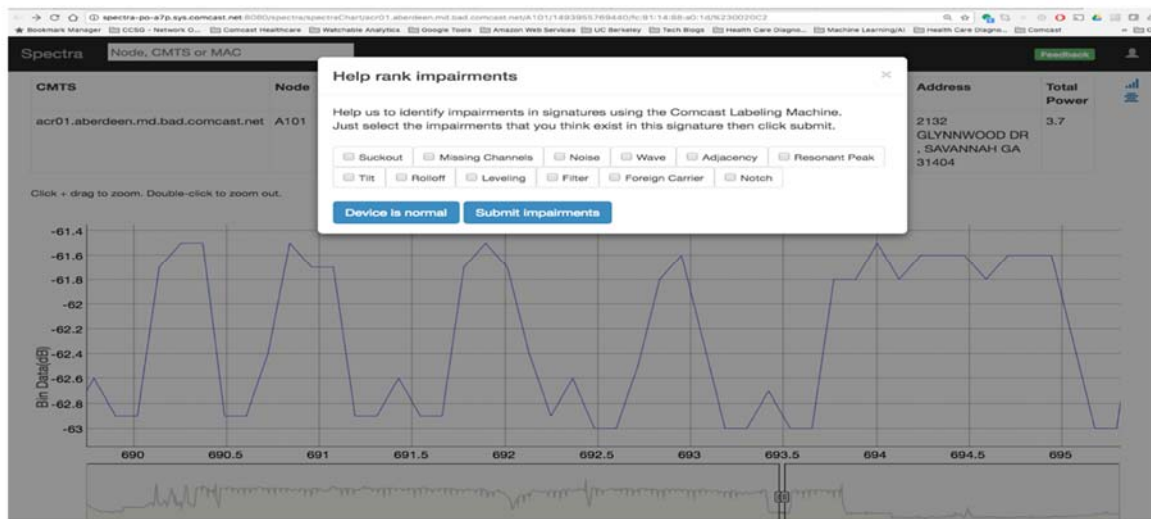


Figure 15 - Comcast Labeling Machine – Labeling Impairments

The graphic interface shown above allows a quick annotation, prior to sending the data back to the labeling Turk. The same data can be labeled independently by several technicians to improve the label quality through a vote or an averaging process. The collected data can be reviewed at the labeling engine prior to sending it to the ML algorithm.

Conclusion

Comcast and similar service provider companies have a lot to gain through the applications of machine learning, especially in the area of improving operational efficiencies.

A key takeaway from this paper are the concepts of machine learning as they relate to multiple service operators. Especially important is the new paradigm under which machine learning operates – that of probabilistic expectations and the move away from determinism.

“The best is the enemy of the good”.

The search for answers that are a hundred percent guaranteed to be true can stifle our ability to be successful because it makes us resistant to innovative approaches to problem solving, such as machine learning. Machine learning provides us not just an answer but also the probability associated with the outcome. We, as cable operators, should begin to appreciate the value of such results and have processes that can educate folks on how to use this information. Only then, can we reap the true benefits of machine learning.

This paper also looked at how spectral impairments in the RF spectrum can be predicted using two approaches, one based on straight-forward mathematical modeling and another based on machine learning. Mathematical modeling is similar to a rule-based approach where patterns in the RF spectrum are predicted based on how well they fit certain mathematical functions. The mathematical functions are built using one or more observations. The main drawback of the mathematical model is its inability to scale to accommodate a larger set of representative spectral impairment patterns. Machine learning trains numerous sets of labeled spectral impairment observations and uses this method to build a model for spectral impairment detection. It can also leverage the feature selection work done using the mathematical model.

Given the vast amount of training data that the machine learning model has seen, it is able to better discern subtle differences in spectral waveforms and consequently leads to better predictions. In addition, it is much more maintainable than the mathematical modeling approach since learning to identify a new spectral impairment is simply a matter of adding the new spectral impairment data to the training set and rebuilding the model. The presence of labeled data or an easy method to label the data would further simplify the machine learning approach. The mathematical modeling approach on the other hand is harder to maintain because it requires an expert to generate new functions to recognize a new impairment.

Both the mathematical model and the machine learning model can nicely coexist, with the predictions that they each make serving to contribute to reinforce the overall prediction accuracy.

Abbreviations

AI	Artificial Intelligence
CNN	Convolutional Neural Network – A type of neural net that has been very successful in image classification
DOCSIS	Data Over Cable System Interface Specification – The protocol used to carry data traffic over cable infrastructure
ML	machine learning
MSO	Multiple Service Operators – typically refers to cable providers
PCA	Principal Component Analysis – A type of machine learning algorithm used for feature transformation
RF	radio frequency
SVM	Support Vector Model – A type of machine learning algorithm used for classification

Bibliography & References

The Singularity is Near, By Ray Kurzweil, ISBN - 9780715635612

Declining storage costs, mkomo.com, URL: <http://www.mkomo.com/cost-per-gigabyte>

Statistical Modeling: The Two Cultures, Leo Breiman , URL:

<https://projecteuclid.org/euclid.ss/1009213726>

Gauss-Markov assumptions for linear regression models, URL:

https://en.wikipedia.org/wiki/Gauss%E2%80%A3Markov_theorem

The Netflix Prize, URL: <http://www.netflixprize.com/>

Amazon Mechanical Turk, URL: <https://www.mturk.com/mturk/welcome>

The Imperative of Customer-Centric Operations

An Operational Practice prepared for SCTE•ISBE by

Anis Cheikhrouhou

Principal Consultant
Nokia, Bell Labs Consulting
1 Route de Villejust
Centre de Villarceaux
91620
Nozay
France
+33 1 6040 4381

Anis.Cheikhrouhou@bell-labs-consulting.com

Jim Davenport

Principal Consultant
Nokia, Bell Labs Consulting
3519 Redding Road
Upper Arlington, OH 43221
+1 614-357-1022

Jim.Davenport@bell-labs-consulting.com

Anish Kelkar

Principal Consultant
Nokia, Bell Labs Consulting
600 Mountain Avenue,
Murray Hill, NJ 07974, USA
+1 732-208-9692

Anish.kelkar@bell-labs-consulting.com

Executive Summary

Customer satisfaction drives better performance in the market and increased company value. Unfortunately, as is evident from lagging NPS scores, traditional service providers including MSOs have not performed as well as other industries despite allocating a proportionally much larger share of human resources in customer care. In contrast, web scale providers allocate very limited human resources for customer care, but have much stronger customer intimacy.

To make matters more challenging, customer preferences are evolving rapidly. Access to content is being diversified away from a single static subscription service delivery toward more granular consumption. Consumers and enterprises are demanding a more transactional model for personalized, context aware services that adapt rapidly to changes in need or demand.

To meet these current and evolving challenges, MSOs must transform to a customer centric operations. The paradigm shifts from managing a network to delivering a service that delights the customer. Service customization, user control of their services and automatic adaptation of service delivery are the new requirements for operations. In this model of operations with capability to take autonomous actions, it is critical to define the optimum level of human responsibilities and the insights that they need to perform their functions.

This paper presents four core principles of future operations: customer-centric automation, hyper-scale analytics, an open and programmable architecture, and cognitive awareness.

While MSOs have moved aggressively towards automation of the manual operational functions, **customer centric automation** will enable MSOs to assure and fulfill dynamic services that adapt to changes in network state and customer demand. Deploying **hyper-scale analytics** at the edge and across all business functions helps to predict, minimize and potentially prevent customer impact due to network outages as well as to rapidly isolate root causes.

MSOs have started partnering with each other to deploy services across their footprints. Adopting an **open and programmable architecture** with API exposure to 3rd parties will accelerate deployment and reduce costs for such future services. **Cognitive awareness** and self-learning capabilities in the network enable delivering a personalized service to customers by learning from their past usage of service and their current context (in transit, location, home/business/school, etc.). Autonomous, dynamic actions based on the state of demand and resources will prevent customers from being impacted from network and service issues and address demand surges faster than humanly possible.

This new operating model should either be neutral or reduce operational costs for the MSOs. We expect mean time to restore service and service cycle time will be less than 5 minutes compared to multiple hours or days. Today the goal for network availability is typically 99.999%. With the adoption of this architecture, the goal is to deliver services that are available all the time. From a customer perspective, we should aim to have at least 60% of services be fully customizable, and at least 80% of interactions should be done through self-care. No Faults Found truck rolls will be minimized to less than 5% because analytics will identify root cause before the truck rolls are launched.

These benefits are achievable, but MSOs need to transform and make foundational changes. On the people side, MSO's need higher skilled roles with data sciences and software automation skills. Processes need to be updated so that they drive standardization and enable automation. Data will need to

become available freely across all business silos. Tickets and workflows management will need improved discipline so that correct data is entered consistently and these artifacts can be used across the organizational silos for training cognitive platforms.

1. Introduction

Across all industries, customer satisfaction is a key parameter for achieving success in the market place. **Figure 1** shows the performance of an actively managed portfolio of companies with the highest ACSI (American Customer Satisfaction Index, n.d.) scores for long positions and lowest ACSI scores for short positions. The financial impact of high customer satisfaction is dramatic.




Figure 1 - Impact of Customer Satisfaction on Market Performance

This is a clear indication that customer satisfaction defines success in the market place. Typically, the measure of customer satisfaction has been Net Promoter Scores (NPS). Bell Labs Consulting performed an analysis of several industries comparing NPS with human resources investments in customer care. (Weldon, 2015)




Figure 2 - NPS comparison

Traditional service providers (including MSOs) invest proportionally more in customer care, however their NPS scores lag significantly behind other industries (Figure 2). Compare this with web scale companies who allocate minimal human resources to customer care, but have much stronger relationships with their end customer. Now, we can argue that MSOs have taken a tremendous amount of effort over the past few years to improve customer satisfaction. In addition, with the adoption of new cloud-based services and adoption of DevOps, MSOs have become agile and flexible in deploying new services. However, a linear path of evolution will not deliver the quantum leap in customer satisfaction that we desire. MSOs must transform their operations to meet the following key requirements:

- Offer customizable, dynamic services which can be consumed when and where desired
- Enable partner MSO services rapidly across their footprints
- Give customers increased control to meet their service needs
- Predict changes in network state, impact to service quality and prevent negative customer impact
- Adapt services to customer demand
- Personalize service delivery to customer context

In this paper, we will cover the principles of the future customer centric operations that will meet these requirements and enable achieving the goal of delivering services that delight the customers.

2. Principles of the future customer centric operations

The four principles of the future customer centric operations are:

1. **Customer-centric automation** for fast and frequent adaptation of services to customer needs
2. **Hyper-scale analytics** to predict, isolate and remedy customer-impacting degradations
3. **Open and programmable architecture** for personalization and interoperability of services
4. **Cognitive awareness** for optimization of services within multiple dimensions of context

2.1. Customer-Centric Automation

As networks and services become increasingly complex the ability for humans to efficiently operate them decreases dramatically. Many operations processes can be automated but the most important are the ones that directly impact the customer experience. This “customer-centric” automation of assurance and fulfillment activities ensures customers can get what they want and the quality they expect.

In the future, services from traditional MSOs and telecommunication companies will be expected to adapt dynamically to changing demand and underlying conditions like network state. Web scale companies do this adaptation today to their services. Degradations to service quality or threats to quality will drive preventive actions to optimize and prioritize network and service resources. The adaptation will occur automatically, thus reducing the need for the customer to call in to the provider.

Customers will request additions or changes to their services through omnichannel options, e.g., apps, web sites, text and Twitter messages. Automated fulfillment and service orchestration with zero touch provisioning will reduce that average service cycle time to meet their requests from days to minutes.

Table 1 - Metrics Impact of Customer-Centric Automation

Metric	Today	Future
Service Cycle Time	6-10 days	< 3 minutes
Mean Time to Restore (Sev 1, 2 w/o construction)	1-4 hours	< 1 minute
Average Handle Time	~5 minutes	Replaced by Self Care
Service Quality Indicators	Not consistent	> Business Goals
Self Service Effectiveness	Not measured	> 90%

2.2. Hyper-Scale Analytics

MSOs are starting to deploy use cases for analytics, with focus on getting additional insights from their network or the customer experience. With hyper scale analytics, MSOs will deploy analytic capabilities across the entire edge and core network and spread through the entire operations. The data lake has multiple sources like element management systems (EMS), customer premise equipment logs, network operations tickets, customer care tickets and workflows, billing systems, marketing intelligence and others.

Hyper-scale analytics will analyze event streams from multiple sources to detect changes from normal behavior caused either due to state change in the network or fluctuations in customer demand. It identifies potential reasons for the abnormality and suggests possible resolutions. In case of impairments that cannot be recovered automatically, like fiber cuts or incidents at single point of failures, analytics isolates the

root cause and provides an estimate of the customer impact. This will enable MSO operations to prioritize recovery actions.

Implementation of analytics can have a significant impact on MSO operational performance especially on customer impact time. However, to implement network analytics, MSOs need to improve the accuracy of their network state. Table 2 identifies the key metrics that measure the effectiveness of analytics adoption.

Table 2 - Metrics improved by hyperscale analytics

Metric	Today	Future
Customer Impact Time	Not measured	< 3 minutes
Service adaptation latency	Not existent	< 1 minute
Network state accuracy	50-60%	> 99.9 %

2.3. Open and Programmable Architecture

Our networks need to become open and programmable so that our customers and MSO partners can easily provision, deploy and adapt their services without any human intervention. We envision that in the future, MSOs will collaborate heavily to deploy common services across each other's foot prints. This is feasible only by providing a rich and updated framework of network APIs for easy integration and provisioning of services. For consumers and enterprise customers, self-service portals provide similar self-serve functions albeit with an easy user interface of web portals or mobile apps. The open and programmable architecture will also enable partners and customers to update and adapt their services in an automated fashion.

These actions will realize operational cost savings and will be a tremendous asset to improve customer satisfaction. In addition, easy, do-it-yourself capabilities for service chaining will allow customers to develop their personalized bundle of services, often resulting in an upsell of services.

As MSOs move towards programmable networks, the key metrics that should be measured and improved are in the table below.

Table 3 - Metrics for open and programmable architecture

Metric	Today	Future
% Customizable services	Very few	> 60%
Mean time to update 3 rd party service	N/A	<2 hours

2.4. Cognitive Awareness

Self-learning capabilities will become increasingly important as network operations become more autonomous. As an example, when faced with a network impairment, the new operational model will first try to autonomously correct the issue or mitigate the impact. If that fails, then the service operations center is presented with the impairment along with the customer impact and prioritized mitigations or resolutions. Once the human decides and acts to successfully resolve the impairment, the self-learning

algorithms update to include the resolution in their repertoire. This concept, called augmented intelligence will allow humans to focus on higher value tasks of continuous improvements and policy management, and machines take care of mundane tasks. The other use examples for cognitive awareness include personalization of service by self-learning of customer context, service usage and dynamic service adaptation based on current and historical data.

The metrics that will be impacted by cognitive awareness are shown in Table 4.

Table 4 - Metrics for Cognitive awareness

Metric	Today	Future
Network and Service availability	99.999%	100%
% incidents self-healed	Not measured	> 99%
Service adaptation rate	N/A	< 30 sec
Contextual Offer acceptance rate	N/A	>30%

3. Future vision of customer centric operations

The future vision of the new operating model aligned with these principles is shown in **Error! Reference source not found..**

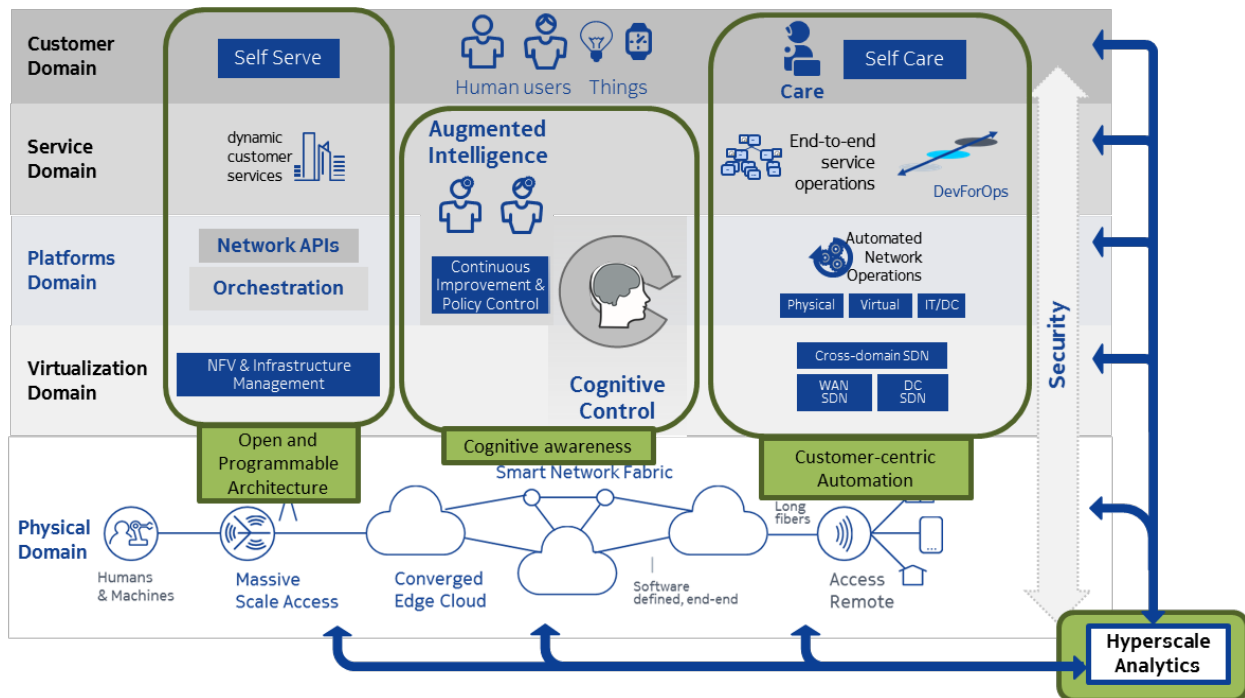


Figure 3 - Future customer centric operating mode

The physical domain incorporates the traditional “network”. This is the domain that has traditionally been managed by network operations. The physical domain includes hardware assets for both the physical and virtual network components.

The domain of virtualization includes capabilities that enable cloud based services. This architecture does not differentiate between OpenStack versus integrated solutions and applies equally to both solutions.

The platform domain provides the “programmable” capabilities that are crucial to service adaptation. It includes rapid and dynamic orchestration of resources as well as access to the state of the network. This domain provides API exposure capabilities that can be exploited for deploying customizable services, both by consumers, enterprises as well as partner MSOs.

We have traditionally operated networks by managing individual network elements. However, as we move towards cloud-based dynamic services, it is paramount that we manage end-to-end services to ensure quality of experience to the user. The service domain orchestrates service delivery and assurance using dynamic service models that deterministically measure quality of experience. DevForOps (Weldon, 2015) is the agile methodology used for rapidly deploying customer services in partnership with MSOs vendors.

The customer domain provides increasing levels of self service and self-care to give customers control over their services in an omnichannel environment.

While a lot of the normal events should be handled autonomously without minimal human intervention, the system provides humans augmented intelligence that enables them to manage policies and continuous improvements.

This operating model impacts all facets of operations, some of which are covered in section **Error! Reference source not found..**

4. Key steps to transform to the future customer centric operations

The adoption of customer-centric operations requires an evolution of the operating model in all 4P's: People, Process, Platforms, and Performance.

Transformation first starts with people. The change is best achieved by staff who embrace innovative disruption and with data science, automation and software skills. The implementation of analytics requires a huge culture shift in operations towards a data driven culture that can trust data available to it and use it as a tool for improvement.

Services will be designed to be inherently flexible in nature. A potential option is to design base services with options that can be programmed either through APIs or customized through service. Designs will need to change to deploy services that look and feel the same to the end user, but are available to them ubiquitously- on and off MSOs foot print, and with local customizations. The service design needs to

adapt dynamically to feedback from the network that indicates any changes in how the services are being delivered or consumed.

Processes must focus on continuity of service management and agile DevForOps delivery. Before the cognitive capabilities are deployed, it is critical that MSOs test operational readiness to prevent unexpected outcomes. Current operational processes like configuration management, inventory management, capacity planning need to be redefined with abstraction of resource, service and product layers and standardized across the MSO.

The top priorities for platforms includes enriched self-service portals, omnichannel workflows and self-learning automated capabilities in every business function. Analytics need to be included as a key requirement in service design, rather than bolted in during subsequent lifecycle stages. Data needs to be universally accessible within the MSOs environment across organizational boundaries and with adherence to high security standards.

In today's operations, we often run into examples where data is either not accurate or not freely available across organizational barriers. Analytics works best when the data provides a full perspective of operations. This means that data from marketing, network, customer care and operations needs to be available freely to the analytic use cases. Any organizational silos that prevent this from happening need to be understood and resolved with optimum solutions that adhere to regulatory considerations.

Present ticketing systems and workflow tools work inconsistently across organizational boundaries. There are examples where short cuts are taken such as tickets may not have the right data describing the impairment or may be closed with incorrect resolution codes. In many cases, care metrics have huge, unexplainable discrepancies in Average Handle Time and First Call Resolution. Agents in care, NOC and dispatch organizations need to be trained to be consistent and disciplined in entering information in tickets and using workflows. The metrics and data entered in these systems will be used for training self-learning, cognitive platforms. Hence emphasis should be to have accurate information in these artifacts to prevent "garbage in, garbage out".

We expect that MSOs will need additional investments in the transformation program to a customer centric operating model. However, these investments will quickly realize business value. *The use of automation, self-learning and analytics will reduce time spent by humans in these activities and improve customer satisfaction which is a direct barometer of market value.*

5. Use case - Future SD-WAN service

The Software-defined wide area network (SD-WAN) service is a popular early MSO offering to enterprise customers. From an operational cost, flexibility and customer experience perspective, these services are a huge upgrade over the traditional WAN services. A customer-centric architecture enables more advanced SD-WAN services. The future operations enable these services to become dynamic and rapidly adaptable. An example of this behavior is shown in **Figure 4**.

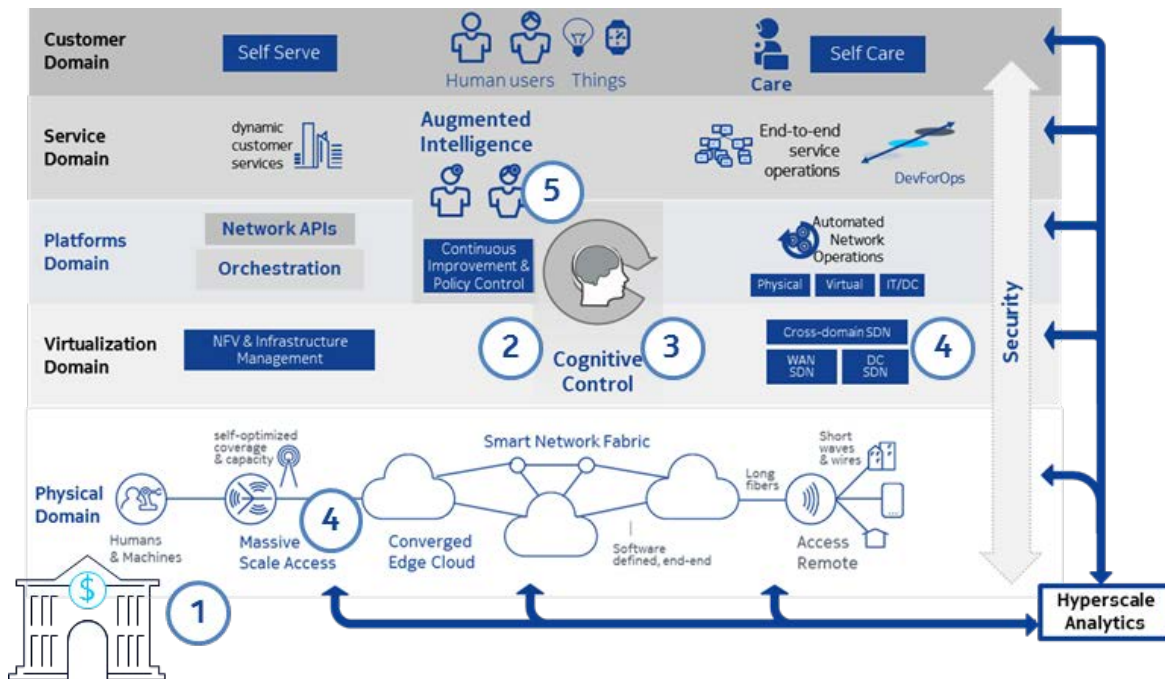


Figure 4 - Dynamic fulfillment of SD WAN

With the principles of customer-centric operations, a peak requirements event can be addressed proactively:

1. **Detect** recurring patterns of higher bandwidth and QoS demands at an enterprise site through analytics
2. **Correlate** increased bandwidth and QoS needs with the underlying enterprise application
3. **Determine** the optimum resource allocation within current SD-WAN service policy to deliver the required bandwidth and QoS
4. **Orchestrate** just-in-time required bandwidth and QoS at target site
5. **Self-learn** from customer behavior to suggest new rule-based policy for SD-WAN service considering customer applications and target sites

The four key principles of the new operating model collaborate to detect changes in customer behavior, automatically deploy needed resources and prevent customer impact.

All these dynamic changes are done autonomously and governed by business policies that are currently in place with the enterprise customer. However, the customer centric operation signals humans that a change of business policies might be needed over time, especially as it learns more about the customer behavior.

6. The future starts now

This paper integrates many capabilities into a cohesive architecture aligned around the four principles. But these capabilities are already being used in many instances either within MSOs or with fixed or wireless service providers. Web-scale companies are clearly far ahead on the road to implementing these capabilities.

It isn't hard to find leading service providers adopting these principles to meet their business challenges and opportunities. Telefonica is combining customer-centric automation of network management with AI capabilities with the goal of averting problems before they impact customers. (Telefonica rolls out AI based network management, 2017). AT&T's Network 3.0 Indigo strategy clearly embraces the evolution to principles of hyper-scale analytics and cognitive awareness. They expect their AI systems to predict when something is about to fail and use self-learning to improve their abilities over time. (AT&T wants to use AI as a crystal-ball, 2016). AT&T envisions building closed communities for data sharing and collaboration using their Network 3.0 Indigo platform. (Data Communities on AT&T Network 3.0 Indigo, 2017) Dialog (Axiata) is exposing core assets via APIs for external consumption to create "friction-free interoperability and partnership (Dialog API case study, 2016).

7. Conclusion

Consumer expectations are changing in this rapidly evolving business environment. Web-scale companies are setting new expectations for customer self-service, rapidly adapting customer experiences, and rate of innovation. MSOs are increasingly competing with these non-traditional providers of high-speed data, video content, and voice services. Operations is a strong influencer of customer satisfaction with a service. To meet customer expectations, MSOs must invest to transform to a new customer-centric paradigm in their fulfillment and assurance operations. *The option of in-action presents higher risk, as current manual operations will not be able to sustain the burden of massive scale and customer expectations, leading to much larger business challenges.*

8. Metrics definitions

Metric	Definition
Service Cycle Time	Average time to complete a work order related to a service
Mean Time to Restore (Sev 1, 2 w/o construction)	Mean time to restore an incident (severity 1 and 2 outages only) without accounting for any construction activities
Average Handle Time	
Service Quality Indicators	Indicators related to Service quality such as voice mean opinions score (MOS) or data throughput
Customer Impact Time	Number of customers impacted multiplied by elapsed event time from detection to restore
Service adaptation latency	Elapsed time for implementing service changes and updates
Network state accuracy	Percentage incidents when errors reported by various controllers (SDN, MANO, inventory systems, etc.) during task execution
% Customizable services	Number of customizable services divided by the number of total offered services
Mean time to update 3 rd party service	Mean time to update a 3 rd party service
Network and Service availability	Percentage of network/service available time

% incidents self-healed	Percentage of incidents that undergo a fully automated process to restore service/functionality
Service adaptation latency	Elapsed time for implementing service changes and updates
Contextual Offer acceptance rate	Percentage of contextual offers accepted by subscribers per month

9. Abbreviations

AI	Artificial intelligence
API	Application programming interface
ASCI	American Customer Satisfaction Index
DC	Data center
DevForOps	Development for Operations
EMS	Element management system
IT	Information technology
MSO	Multiple system operator
NFV	Network function virtualization
NPS	Net promoter score
QoS	Quality of service
SD-WAN	Software-defined wide area network service
SDN	Software defined networks
WAN	Wide area network

10. Bibliography & References

American Customer Satisfaction Index. (n.d.). Retrieved from <https://www.theacsi.org/national-economic-indicator/financial-indicator>

AT&T wants to use AI as a crystal-ball. (2016). Retrieved from <http://www.computerworld.com/article/3096986/artificial-intelligence/att-wants-to-use-ai-as-a-crystal-ball.html>

Data Communities on AT&T Network 3.0 Indigo. (2017). Retrieved from The Policy Forum at AT&T: <http://policyforum.att.com/wp-content/uploads/2017/03/Data-Communities-on-ATT-Network-3-Indigo.pdf>

Dialog API case study. (2016). Retrieved from https://www.telco2research.com/articles/EB_Dialog-API-case-study

Telefonica rolls out AI based network management. (2017). Retrieved from <http://www.computerweekly.com/news/450417040/Telefonica-rolls-out-AI-based-network-management>

Weldon, M. K. (2015). *The Future X Network: A Bell Labs Perspective*. CRC Press.

When Does the DOCSIS 3.1 TaFD Feature Increase the Capacity of My Network?

A Technical Paper prepared for SCTE•ISBE by

Ayham Al-Banna, Ph.D.

Engineering Fellow

CTO Office – Network Solutions, ARRIS

2400 Ogden Ave., Suite 180

Lisle, IL 60532, USA

630-281-3009

Ayham.Al-Banna@arris.com

Tom Cloonan, Ph.D.

CTO

CTO Office – Network Solutions, ARRIS

2400 Ogden Ave., Suite 180

Lisle, IL 60532, USA

630-281-3050

Tom.Cloonan@arris.com

Abstract

The traffic engineering benefits of the TaFD feature of OFDMA channels are evaluated for different network scenarios based on real-world assumptions & measurements from a network operator. MSOs are looking at this feature as a potential savior for symmetrical service offering. It will be shown that deploying the TaFD feature may or may not benefit the overall capacities of HFC networks depending on the network assumptions and conditions. The specific conditions will need to be studied and analyzed to decide whether TaFD adds value or not. Migration alternatives are proposed for cases where a symmetrical service offering is desired but TaFD may not help..

1. Introduction

As MSOs start planning the deployment of DOCSIS 3.1 in the US direction to increase US capacities, they wonder whether the DOCSIS 3.1 Time and Frequency Division (TaFD) Multiplexing feature will yield increased network capacities. Part of the challenge is that MSOs are trying to increase their peak rate service offering without increasing the size of the US spectrum due to the prohibitive cost and operational complexity of changing the split across the whole network. Therefore, they look to the TaFD feature as a promising solution to help them address the customer demand and competitors that offer faster services. While previous work [1] showed the advantages of the TaFD when compared to other US scheduling schemes, this article focuses more on analyzing the value of the TaFD feature from a capacity gain viewpoint given real world network scenarios.

The capacity gain that is expected from the DOCSIS 3.1 TaFD feature is a function of many variables. These include, but are not limited to, the following: the service group size, subscribers' traffic engineering statistics, desired peak rate service offering, US spectrum size, OFDMA channel parameters, size and number of SC-QAM channels, TaFD switching overhead, etc. This paper takes the above parameters into consideration as it analyzes any potential capacity gain and improvement in the peak rate service offering as a result of deploying the TaFD feature.

The analysis in this paper is based on multiple scenarios and network conditions provided by a partner network operator (referred to as MSO X for the rest of the paper) as part of a joint collaboration, where a study was performed to evaluate the potential benefits of deploying the TaFD feature. The goal was to provide guidance as to whether deploying the TaFD feature made sense given the network conditions, planned service offering, and subscriber BW usage. Another goal of the study was to provide migration alternatives if/when the TaFD feature does not provide the desired capacity gains.

This paper is organized as follows. Section 2 defines the problem statement and provides the TaFD capacity analysis for multiple scenarios. Alternative migration strategies are provided in Section 3 and Section 4 concludes the paper.

2. TaFD Capacity Analyses

2.1 Problem Statement

The US spectrum configuration in the MSO X network is composed of two QAM64 6.4 MHz channels and two QAM64 3.2 MHz channels as shown in Fig. 1 with a total estimated plant capacity of ~75 Mbps for the whole US Service Group (USSG). The current US peak rate offering for existing DOCSIS 3.0

subscribers is 12 Mbps. The bottom line question is: what is the maximum peak rate (Tmax) that can be offered for the new DOCSIS 3.1 subscribers if the TaFD feature is deployed as shown in Fig. 2? The answer to this question is critical due to the pressure of speed war battles with competition!




Figure 1 - Existing configuration of the US spectrum with 5-42 MHz plant




Figure 2 - Deploying the TaFD feature with a 20 MHz OFDMA channel that overlaps with the existing SC-QAM channels in 5-42 MHz plant

As described in the introduction section, the analysis of capacity gain & potential service offering as a result of deploying the TaFD feature is function of many parameters and variables. In the following two subsections, we evaluate two different scenarios (Scenarios A & B) with slightly different assumptions.

2.2 Scenario A Analysis

Scenario A was based on a measured histogram for the USSG BW utilization levels that is shown in Fig. 3. It can be seen that >90% of USSGs have 60% BW utilization or less. Moreover, the histogram of number of subscribers per USSG for the USSGs shown in Fig. 3 is provided in Fig. 4. It can be seen that >90% of USSGs have 180 Subscribers or less. Based on these statistics, the average throughput per subscriber (Tavg) was estimated to be $60\% \times 75 \text{ Mbps} / 180 = 250 \text{ kbps}$.




Figure 3 - Histogram of USSG BW utilization levels




Figure 4 - Histogram of number of subscribers per USSG

In order to estimate the capacity of the OFDMA channel capacity that will provide the TaFD feature, it was also important to obtain some plant measurements. The distribution of US SNR values measured on MSO X plant was as shown in Fig. 5. Although the provided SNR values were the peak SNR values per USSG and number of samples that formed the distribution was not large enough to form a Gaussian-style probability density function (pdf), the histogram (with mean SNR of ~35 dB) still was considered to be usable for the purpose of rough capacity analyses. The estimated OFDMA channel spectral efficiency values for different channel widths is given in Table 1 and assumed the following parameters

- 2K FFT with 50 kHz subcarrier spacing
- Cyclic prefix = 1.875 μ sec
- Frame size = 10 OFDMA symbols
- Guard band of 0.5 MHz on each side
- Pilot pattern 2
- Large Codeword size
- MSO SNR margin of 2 dB
- QAM256 modulation order (two modulations order better than QAM64 used with the currently deployed SC-QAM channels that have a spectral efficiency value of 4.15 bps/Hz [2])

Per Table 1, observe that the total capacity of the 20 MHz OFDMA channel, shown in Fig. 2, can be estimated as 20 MHz*5.79 bps/Hz ~ 116 Mbps. Obviously, this value assumes that the OFDMA channel has access to the spectrum 100% of the time and therefore would not take any TaFD switching overhead into consideration.




Figure 5 - Histogram of peak US SNR values per USSG

Table 1 - Spectral efficiency of the US OFDMA channel

OFDMA Channel (MHz)	10	20	40	60	80	96
Spectral Efficiency (bps/Hz) at PHY Level	5.40	5.79	5.98	6.04	6.08	6.09

In order to proceed with the TaFD analysis, the following assumptions are made:

- 5-42 MHz plant with total usable spectrum of 20 MHz
- TaFD spectrum is 19.2 MHz (overlapping spectrum between SC-QAM and OFDMA). OFDMA can use 0.8 MHz of spectrum that is not overlapped with SC-QAM
- D3.0 spectral efficiency (PHY) = 4.15 bps/Hz [2]
- D3.1 spectral efficiency (PHY) per Table 1
- T_{avg} is equal for SC-QAM and OFDMA subs = 250 kbps (calculated previously)
- Currently offered T_{max} for SC-QAM = 12 Mbps. The SC-QAM subs will continue to be offered a T_{max} of 12 Mbps
- T_{avg} for the OFDMA subscriber is equal to T_{avg} for a SC-QAM subscriber
- Total number of subscribers (N_{sub}) = 180
- TaFD switching and scheduler inefficiency overhead = 20%. This includes overhead due to guard time, guard bands, and scheduling inefficiencies
- MAC overhead = 5%
- ARRIS QoE Formula is used to estimate the needed capacity as shown in Fig. 6:
 - Needed Capacity = $N_{sub} * T_{avg} + 1.2 * T_{max}$




Figure 6 - ARRIS QoE formula used to estimate the needed capacity

Given the above assumptions, the total usable plant capacity using the TaFD feature in a 5-42 MHz plant as a function of the percentage of deployed DOCSIS 3.1 CMs is shown in Fig. 7. This total capacity takes the TaFD switching overhead into consideration and therefore it can be observed that the total capacity of the plant drops, when a small percentage of DOCSIS 3.1 CMs get deployed, due to the 20% TaFD switching overhead assumption. The corresponding peak rates that can be offered to DOCSIS 3.1 subscribers as a result of deploying DOCSIS 3.1 and the TaFD feature are shown in Fig. 8. The results in Fig. 8 includes various cases where bonding between SC-QAM and OFDMA channels is supported or not. There are a few key points to note from these results:

- 25 Mbps can be offered using SC-QAM-only channels.
- Tmax service beyond 25 Mbps for DOCSIS 3.1 subscribers cannot be offered using 5-42 MHz spectrum without utilizing the channel bonding feature in addition to the TaFD feature.
- With bonding enabled, at least 55% of subscribers need to be D3.1 subscribers in order to offer higher than 25 Mbps service.

These conclusions do not present the TaFD feature as an attractive transition technology because a large percentage of DOCSIS 3.1 subscribers need to be deployed in order for the feature to start adding value above what SC-QAM can offer alone.




Figure 7 - Total usable plant capacity with TaFD in 5-42 MHz plant for Scenario A




Figure 8 - Peak rates offered to DOCSIS 3.1 subscribers with TaFD in 5-42 MHz plant - Scenario A

The above results make the reader wonder why the TaFD had marginal benefits in 5-42 MHz plants with scenario A. This can actually be attributed to the following:

- OFDMA is less efficient with small channel widths
 - 20 MHz is just about one fifth of the maximum channel size of 96 MHz defined in the DOCSIS 3.1 PHY specifications. [3]
 - Most existing subscribers are DOCSIS 3.0 subs. This leaves little room for spectrum access for OFDMA subscribers.
 - Not only is the OFDMA channel width small, but it also time-shares the spectrum with SC-QAMs resulting additional reduction in throughput.
 - As the number of DOCSIS 3.1 subs increases (and number of SC-QAM subs decreases), the OFDMA performance slightly improves.
- TaFD Switching overhead is significant compared to the gained capacity
 - Switching between OFDMA and SC-QAM channels requires guard time and/or guard bands. [1]
- There is a slight capacity gain from using OFDMA but....
 - The benefit of this gain is marginalized by the switching overhead.
 - Reducing the switching overhead with optimized implementation can increase the value of the TaFD feature.

One of the transition plans of MSO X was to move to an US mid-split of 5-85 MHz and increase the number of SC-QAM channels as shown in Fig. 9. Therefore, a similar analysis was performed for the 5-85 MHz split option and the results are shown in Fig. 10. Note that the analysis considered multiple options, where one option assumed that the number of SC-QAM channels is as described in MSO X's transition plan in Fig. 9 deployed together with an OFDMA channel that overlaps with those SC-QAM channels via the TaFD feature. The other option is a reduced number of SC-QAM channels (to the

minimum that is needed to accommodate DOCSIS 3.0 subscribers). In the latter option, when the number of SC-QAM channels is reduced, the reclaimed spectrum is given back to the OFDMA channel, which does not overlap with the SC-QAM channels, and therefore does not use the TaFD feature. This is essentially a Frequency Division (FD) operation, where different channels occupy different parts of the spectrum.

The analysis results show that the system performance when reducing the number of SC-QAM channels to the absolute minimum that is needed to accommodate DOCSIS 3.1 subscribers and using OFDMA channels in other parts of the spectrum outperforms the option where the number of SC-QAM channels is increased and the TaFD feature is utilized. In particular, the results show that in 5-85 MHz plant, the offered T_{max} with the FD option (FD operation between OFDMA and reduced number of SC-QAM channels) can be as high as 213 Mbps with 5% of the subscribers being DOCSIS 3.1 subs. This T_{max} value is more than the T_{max} of 184 Mbps that can be offered with the TaFD option (using TaFD & bonding between an OFDMA channel that overlaps with large number of SC-QAM channels) with 5% of the subscribers being DOCSIS 3.1 subs. Similarly, with 95% of the subs being DOCSIS 3.1 subs, T_{max} value for the FD option is 227 Mbps compared to 200 Mbps with the TaFD option.




Figure 9 - MSO X original plan to move to 5-85 MHz plan and increase the number of SC-QAM channels– Scenario A




Figure 10 - Peak rates offered to DOCSIS 3.1 subscribers with TaFD in 5-85 MHz plant - Scenario A

Given the above analysis for scenario A, it can be concluded that it is more efficient to stick to SC-QAM technology in 5-42 MHz plants when the percentage of DOCSIS3.1 subs is small (less than ~50%) given the fact that a Tmax value of up to 25 Mbps can be offered using SC-QAM only. It can also be concluded that offering higher Tmax values (for symmetrical services) in scenario A will likely require a migration to 5-85 MHz spectrum and deployment of DOCSIS 3.1. In this case, it is better to use frequency division operation between an OFDMA channel and a small number of SC-QAM channels, which will yield Tmax values as high as 213 Mbps with initial deployments.

2.3 Scenario B Analysis

Scenario B analysis will show that the conclusions can differ depending on the analysis assumptions. The spectrum configuration of scenario B is shown in Fig. 11. Only a few parameters/assumptions were changed from scenario A to create scenario B as shown in Table 2. All other parameters are kept the same including the total number of subscribers per USSG of 180 subs. Note that scenario B assumes that the ODMA channel starts at 13 MHz (not 5 MHz) due to the noisy conditions and potential existence of multiple Out Of Band (OOB) signals in that part of the spectrum (5-13 MHz).




Figure 11 - Spectrum configuration for scenario B with 5-42 MHz plant

Table 2 - Parameters assumed differently for scenarios A & B

Scenario A	Scenario B
Tavg = 250 kbps per sub	Tavg = 84 kbps per sub
US spectrum from 20 to 40 MHz	US spectrum from 13 to 40 MHz
4 US QAM64 SC-QAM channels (2X3.2 + 2X6.4 MHz) between 20 and 40 MHz. OFDMA starts at 20 MHz	4 US QAM64 SC-QAM channels (2X3.2 + 2X6.4 MHz) between 20 and 40 MHz. OFDMA starts at 13 MHz with variable bit loading (QAM64 in 13-20 MHz, QAM256 otherwise)
No DOCSIS channels below 20 MHz	QAM64 for OFDMA channel between 13 and 20 MHz

The results of the analysis for scenario B in 5-42 MHz plant are shown in Fig. 12, which shows the total net capacity of the plant after TaFD switching overhead is taken into consideration. Note that the OFDMA-only capacity refers to the capacity gained from non-overlapping spectrum between 13 and 20 MHz. It can be observed that the US capacity for the 5-42 MHz plant increased to 107 Mbps when 5% of the subscribers are DOCSIS 3.1 subs. This increase in capacity represents a gain of about 43% when compared to plant original capacity of 75 Mbps. Similarly, the US capacity for the plant is increased by 85% when all subscribers are DOCSIS 3.1 subscribers.

The corresponding peak rates for Fig. 12 are shown in Fig. 13. It can be observed that 50 Mbps peak rate service can be offered using SC-QAM-only channels. Also, for initial deployments where 5% of the subscribers are DOCSIS 3.1 subs, TaFD with no bonding between SC-QAM and OFDMA channels provides about 28% gain in Tmax (i.e., 64 Mbps). On the other hand, TaFD with bonding between SC-QAM and OFDMA channels provides about 52% gain in Tmax (i.e., 76 Mbps).




Figure 12 - Total usable plant capacity with TaFD in 5-42 MHz plant for Scenario B




Figure 13 - Peak rates offered to DOCSIS 3.1 subscribers with TaFD in 5-42 MHz plant - Scenario B

TaFD analysis for scenario B with 5-85 MHz mid-split plant was also performed in a similar fashion to what was done for scenario A in the previous section. The 5-85 MHz plant configuration for scenario B is shown in Fig. 14. Similar to what was done earlier, the TaFD analysis considered two options: the FD option and the TaFD option (refer to the previous section for the description of those options).

It can be observed from the analysis results shown in Fig. 15 that the system performance when reducing the number of SC-QAM channels to the absolute minimum that is needed to accommodate DOCSIS 3.1 subscribers and using OFDMA channels in other parts of the spectrum (i.e., FD option) exceeds the

system performance where the number of SC-QAM channels is increased and the TaFD feature & bonding is utilized (TaFD option). In particular, the results show that in a 5-85 MHz plant, the offered Tmax with the FD option can be as high as 273 Mbps with 5% of the subscribers being DOCSIS 3.1 subs. This Tmax value is more than the Tmax of 245 Mbps that can be offered with the TaFD option with 5% of the subscribers being DOCSIS 3.1 subs. Similarly, with 95% of the subscribers being DOCSIS 3.1 subs, Tmax value for the first option is 277 Mbps compared to 250 Mbps with the second option.




Figure 14 - MSO X original plan to move to 5-85 MHz plan and increase the number of SC-QAM channels – Scenario B




Figure 15 - Peak rates offered to DOCSIS 3.1 customers with TaFD in 5-85 MHz plant - Scenario B

Given the above analysis for scenario B, it can be concluded that moving to DOCSIS 3.1 and utilizing the TaFD feature in 5-42 MHz plant can yield increased plant capacities and peak rate offering. This increase ranged from 28% (64 Mbps via TaFD without bonding) to 52% (76 Mbps via TaFD with bonding) for initial deployment when the percentage of DOCSIS 3.1 subs is 5%. Observe that this percentage gain can be deceiving because while the percentage gain could be high (up to 52%), the offered Tmax of 76 Mbps may not be enough for symmetrical services. Therefore, the MSO will have to perform business case analyses to decide whether these achieved Tmax values justify the spending and operational complexity of this feature/scenario. Note that moving to 5-85 MHz split, where the number of SC-QAM channels is reduced to the minimum needed to accommodate DOCSIS 3.0 subs, deployed along with an OFDMA

channel via frequency division operation provides the optimal solution to offer symmetrical services. This less operationally complex solution can provide T_{max} values up to 273 Mbps, which is 446% gain compared to the 50 Mbps that can be offered via 5-42 MHz plant and SC-QAM only channels and about 11% gain compared to the 245 Mbps that can be offered via 5-85 MHz plant and TaFD/bonding features.

3. Migration Alternatives to Offer Symmetrical Services

The results in the two subsections above suggested that moving to 5-85 MHz split is more beneficial than utilizing the TaFD with bonding feature when symmetrical service offering is the goal on mind. The ability to provide symmetrical service offering can only be improved by moving to high-split (5-204 MHz). However, some MSOs cannot sacrifice the DS spectrum for the sake of US spectrum in the near future. Therefore, the 5-85 MHz mid-split option might be a good compromise for those MSOs in the next few years.

Moving to 5-85 MHz option does not have to be done across the whole network at once. The MSO may choose to selectively change the plant to support 5-85 MHz spectrum. This means performing 5-85 MHz surgical split upgrades only in areas where symmetrical service offering is required due to demand or competition. The MSO can then choose to leave the rest of the network unchanged especially if the subscribers are happy and competitors are out of that market.

In an effort to offer symmetrical services in a specific area, once the MSO migrates that part of the plant to 5-85 MHz split, the next step is to allocate just enough spectrum for SC-QAM channels to satisfy legacy customers (the ARRIS QoE formula can be used to do those estimates). Once that is performed, the remaining spectrum may be used for OFDMA and DOCSIS 3.1 may be deployed. Channel bonding between SC-QAM channels and OFDMA can also be used to offer even higher peak rates.

Note that, in addition to the above methodology, the MSO can choose to perform selective subscriber migration to DOCSIS 3.1. In this scheme, the MSO can choose to move heavy (high usage) subscribers and super users (high peak rate subscribers) from legacy channels to DOCSIS 3.1 OFDMA channels. That, in turn, reduces the pressure on SC-QAM channels and therefore the number of those SC-QAM channels can be further reduced leaving more spectrum for the more efficient OFDMA channels. The decision-tree flow diagram of the above technique is shown in Fig. 16.

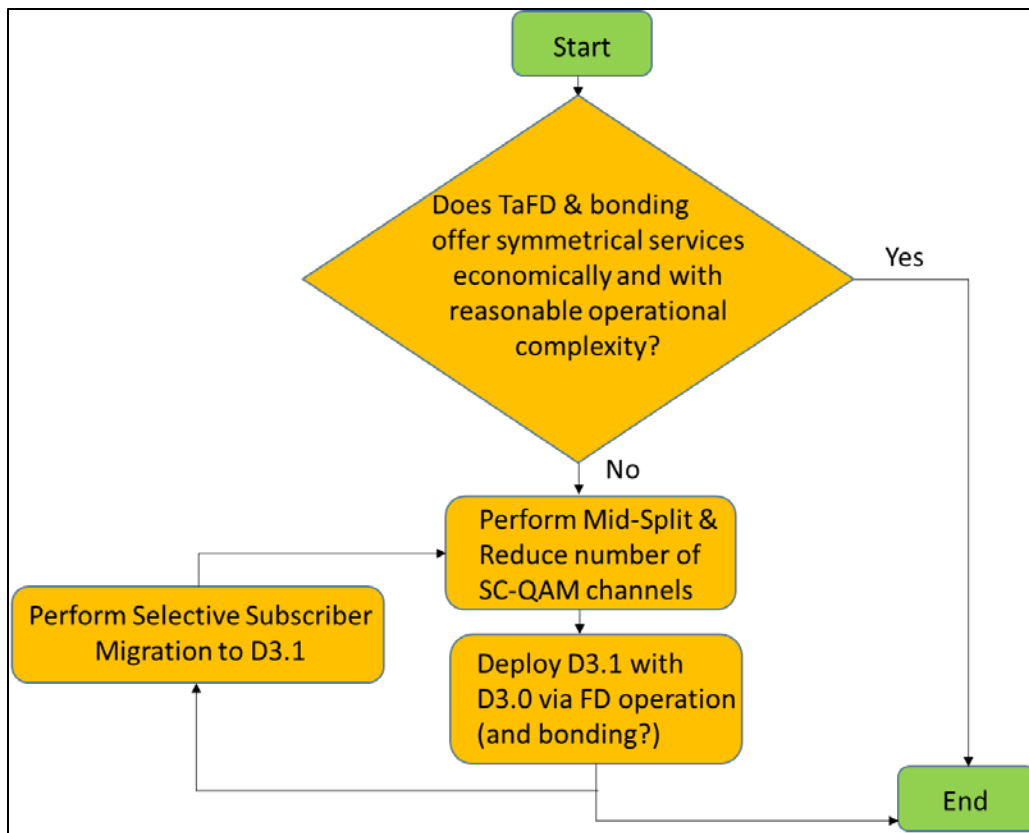


Figure 16 - Decision-tree flow diagram to offer symmetrical services

4. Conclusions

The paper provided traffic engineering analyses for the TaFD feature of OFDMA channels. It was shown that the value added by the TaFD feature depends on the network assumptions. The paper showed that the TaFD feature did not offer any value in some network scenarios but provided benefit in others, especially if used with the channel bonding feature. The MSO will have to perform business case analyses in order to decide whether the benefits & peak rates provided by the TaFD feature are worth the investment or whether investing in other alternatives like mid-split migration in areas where symmetrical service offering is needed. The paper recommended a selective migration to the mid-split option especially when high US peak rate offering for symmetrical services is desired. A decision-tree flow diagram was proposed for the migration process, where selective subscriber migration to DOCSIS 3.1 was proposed to provide an added benefit.

Acknowledgement

The authors would like to thank Sunil Mansharamani of ARRIS for his contributions to this paper.

Abbreviations

BW	Bandwidth
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
DOCSIS	Data Over Cable Service Interface Specifications
FD	Frequency Division
FFT	Fast Fourier Transform
HFC	Hybrid Fiber Coax
kbps	kilobits per second
Mbps	Megabits per second
MSO	Multiple System Operator
Nsub	Number of subscribers
OFDMA	Orthogonal Frequency Division Multiple Access
pdf	probability density function
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
SC-QAM	Single Carrier QAM
SNR	Signal to Noise Ratio
TaFD	Time and Frequency Division
Tavg	Average throughput per subscriber during busy hour
Tmax	Peak rate
US	Upstream
USSG	US Service Group

Bibliography & References

- [1] Ayham Al-Banna et. al., “*US Scheduling In The DOCSIS 3.1 Era: Potential & Challenges*,” SCTE Cable-Tec Expo 2015.
- [2] Ayham Al-Banna et. al., “*The Spectral Efficiency of DOCSIS® 3.1 Systems*,” SCTE Cable-Tec Expo 2014.
- [3] Cabelabs, “*DOCSIS 3.1 Physical Layer Specification – CM-SP-PHYv3.1-I10-170111*,” January 2017.

1024 to 4096 Reasons for using D3.1 over RFoG:

Unleashing Fiber Capacity by Jointly Optimizing D3.1 and RFoG Parameters

A Technical Paper prepared for SCTE•ISBE by

Venk Mutalik
Engineering Fellow
ARRIS Inc.
15 Sterling Drive
Wallingford, CT 06492
203-494-6556
venk.mutalik@arris.com

Brent Arnold, ARRIS

Benny Lewandowski, ARRIS

Phil Miguelez, Comcast

Mike Cooper, Cox

1. Introduction

DOCSIS® 3.1 (D3.1) offers exceptional capabilities for broadband service providers to enhance capacity and throughput both in the downstream (DS) and upstream (US) directions. While there has already been considerable activity in deploying D3.1 in the downstream, operators are in the early phases of rolling out D3.1 in the upstream.

While sharing many similarities with DOCSIS® 3.0 (D3.0), D3.1 standards differ in important aspects – OFDMA operation, higher orders of modulation, increased flexibility in channel width, and most notably in burst upstream operation. There are many new things to consider in D3.1 including basic parameters such as Cyclic Prefix (CP), Roll-off Period (RP), OFDMA FFT size, minimum performance requirements, and encompassed spectrum. While D3.1 allows for a large encompassed spectrum affording significantly increased capacity, it also allows for very small mini-slots thus leading to a rather large dynamic range of RF inputs to existing RFoG ONUs. It is therefore an appropriate time to understand the capability of currently deployed RFoG plant and endeavor to jointly optimize D3.1 and RFoG parameters to ensure robust throughput in the downstream *and* the upstream.

In this paper, we describe relevant D3.1 parameters and link them to physical characteristics of an RFoG network. We examine performance of individual ONUs using stand-alone D3.1 test equipment as well as analyze D3.1 initial range and register protocols in a realistic D3.1 CMTS environment. We next consider a complete multi-ONU RFoG environment with multiple simultaneously transmitting cable modems and ONUs and introduce a new way for analyzing ranges of error free operation windows of D3.1 upstream systems. We conclude with a discussion of real world application of analysis presented in this paper for legacy deployed RFoG systems and offer practical suggestions for green-field D3.1 RFoG deployments.

2. A Word about OBI and D3.1

Optical Beat Interference (OBI) is a profound issue in RFoG reverse path (Mutalik, Schemmann, Maricevic, and Ulm – 2015 INTX NCTA Spring Technical Forum) and affects the system in debilitating ways. The subject of OBI and of its effects on cable systems as well as methods to mitigate and eliminate its deleterious effects has been extensively reported in the context of D3.0 over RFoG (ZorluOzer, Mutalik, Vieira, and Chrostowski – SCTE Cable-Tec Expo 2015). While mitigating OBI could have been argued to be adequate for D3.0 systems, an effective elimination of OBI is a requirement for D3.1 systems, as will be made clear in subsequent sections of this paper. Briefly, typical bonded D3.0 upstream (US) systems have at most 4 cable modems (CMs) that can transmit simultaneously, whereas in D3.1 systems, up to 40 CMs could transmit simultaneously. In this environment, the probability of OBI is vastly increased leading to significant impairments in system throughput (packet loss). Packet losses affect TCP/IP throughput quite profoundly, thus affecting efficiency of the network and with a drastic limitation on capacity.

Fortunately over the years, the cable industry has developed two very effective solutions for the elimination of OBI.

2.1. Wavelength Selective ONUs (WSO Option)

OBI occurs when two or more ONUs, at near identical wavelengths (WLs), transmit simultaneously and reach the same upstream receiver. Since US transmission is in burst mode, ONU lasers typically exhibit wavelength drift at laser start up, which can significantly increase the OBI occurrence. Therefore,

selecting WLs of ONUs that do not intersect even if all ONUs are in burst mode would effectively eliminate OBI. Furthermore, this would enable the use of passive optical splitters as originally envisioned in RFoG deployments. Furthermore, this concept allows for distributed splits in the field thus providing substantial flexibility in the plant. However, in this approach, the passive loss and the potential for additional noise at the headend receiver from a larger number of simultaneously transmitting ONUs could limit available SNR and thus limit capacity. These effects are however a part and parcel of the D3.1 environment, and are described in detail in subsequent sections. Innovative ways of building the ONU and the headend receivers and configuring of the RFoG network would help alleviate some of these concerns. In this paper, this option is referred to as the WSO option.

2.2. Multi Diode Receivers (MDR Option)

OBI can also be eliminated by restricting the light of each ONU to an individual photo-diode (PD). Thus multiple ONUs can transmit simultaneously with no opportunity for OBI to occur. With innovative optical and electronic design techniques, these MDRs can also be designed to fully support PON wavelengths. Since the MDRs are placed at the location of a traditional RFoG splitter, the light levels entering the individual PDs from the ONUs are quite high, and thus when retransmitted to the headend receiver provide a substantial SNR advantage, potentially allowing one to take advantage of higher order QAM modulation.

However, MDRs by their nature are active devices and require powering at their locations and cabinet or strand accommodations. Oftentimes, since powering is provided, these MDRs also include optical amplification to extend the link and/or provide higher power levels in the DS to the ONUs. This is referred to as the MDR option.

In this paper, we assume the elimination of OBI and now proceed to describe D3.1 parameters, compare these to D3.0, and describe their interplay with established RFoG standards in the cable industry such as the SCTE IPS 174.

3. D3.1 Numerology Upstream and Downstream

D3.1 uses Orthogonal Frequency Division Multiplexing (OFDM) in the DS and the US, whereas D3.0 used Single Carrier QAM (SC-QAM). Furthermore, while D3.0 was capped to SC-QAM256 in the DS and SC-QAM64 in the US, D3.1 allows for much more complexity of modulation, all the way up to OFDM4096. Thus, unlike in the D3.0 environment where the capacity is capped for a given RF spectrum, a higher MER in the D3.1 environment can enable higher modulation formats and therefore meaningfully increase capacity of the network. This is a paradigm shift in that higher MER for D3.0 served to provide performance margin, higher MER for D3.1 provides increased capacity (provided a higher modulation format can be achieved).

In the DS, D3.1 operates in 4K or 8K FFT modes, affording channel bandwidths that can span 24 MHz to 192 MHz. The subcarrier spacing for the 4K mode is 50 kHz, and for the 8K mode is 25 kHz, while the symbol duration is 20us for the former and 40us for the latter. These parameters and the performance requirements for various OFDM constellations are summarized in the table below.

Table 1 – D3.1 DS and US Numerology

D3.1 Downstream Parameters Summarized (from CM-SP-PHYv3.1)		
Mode	4k	8k
Channel BWs	24MHz to 192MHz	
Subcarrier spacing	50kHz	25kHz
Symbol duration	20us	40us
Cyclic Prefix	0.9375, 1.25, 2.5, 3.75, 5us	
Rolloff Period	0, 0.3125, 6.25, 0.9375, 1.25us	
OFDM Constellation	CNR AWGN 1GHz	CNR AWGN 1-1.2GHz
4096	41.0	41.5
2048	37.0	37.5
1024	34.0	34.0
512	30.5	30.5
256	27.0	27.0
128	24.0	24.0
64	21.0	21.0
16	15.0	15.0

D3.1 Upstream Parameters Summarized (from CM-SP-PHYv3.1)		
Mode	2k	4k
Channel BWs	10MHz-96MHz	6.4MHz - 96MHz
Subcarrier spacing	50kHz	25kHz
Symbol duration	20us	40us
Cyclic Prefix	0.9375, 1.25, 1.5625, 1.875, 2.1875, 2.5, 2.8125, 3.125, 3.75, 5.0, 6.25 us	
Rolloff Period	0, 0.3125, 0.625, 0.9375, 1.25, 1.5625, 1.875, 2.1875 us	
OFDMA Constellation	CNR AWGN	
4096	43.0	
2048	39.0	
1024	35.5	
512	32.5	
256	29.0	
128	26.0	
64	23.0	
32	20.0	
16	17.0	
8	14.0	
QPSK	11.0	

For the US, 2K or 4K FFT mode can be selected, which respectively have carriers spaced 50kHz or 25kHz apart and with symbol durations of 20us and 40us.

3.1. Time and Frequency Scheduling: Role of the mini-slot

In D3.1, multiple subcarriers are joined together to create a mini-slot. For the US, 8 of the 50kHz subcarriers in the 2k mode, or 16 of the subcarriers in the 4k mode, which in either case amounts to 400 kHz wide frequency spectrum, is considered as a mini-slot. All D3.1 operations (save a few which do not affect our discussion) are done in multiples of the 400 kHz mini-slots. Thus 400 KHz represents the minimum amount of RF spectrum that may be present at the RF input of the ONU in RFoG operation. Depending upon the traffic utilization, the D3.1 CMTS scheduler allots RF spectrum on a case-by case basis in the time and frequency domain to the various CMs for transmission. This spectrum is always allotted in 400 kHz bandwidth increments.

In the time domain, the D3.1 standard allows for multiple symbols to be joined together in a frame; just how many symbols can form a frame depends upon the encompassed spectrum and other modes of operation. The CMTS allots slots for transmission to the CMs in duration of frame sizes. Since the CMTS schedules in real time, the CMs generally turn off per frame and turn on again at a frequency spot that the CMTS allots them. While the frame sizes can vary from 6 to 32 symbols per frame, the frame size is chosen so as to optimize throughput and latency.

Thus, the 400 kHz mini-slot in the frequency domain, along with the consecutive set of symbols that comprise a Frame, in the time domain, form a unit of well-defined entity that carries data in the D3.1 protocol.

3.2. Cyclic Prefix and Rolloff Period

The HFC plant is a significant asset of the MSOs, but is comprised of many cascaded active and passive devices. Each of these connections in the plant is a potential source for reflections and when many such reflection points exist, they can degrade upstream transmission quality. To minimize the effects of these reflections, D3.1 standard uses the Cyclic Prefix (CP). A part of the end of each symbol is prepended to the start of the same symbol (and a part of the start of the symbol is appended to the same symbol) thus ensconcing the symbol within its own parts. At the CMTS, the main signal from the CM and all of its echoes are received and are auto-correlated with the redundant parts discarded and the main symbol information utilized for demodulation.




Figure 1 – Illustrating the D3.1 US Symbol and CP, RP

The design and implementation of the CMTS receiver is a source of innovation and has a direct effect on the CP needed. It should be noted here that a large CP would enable resilience to large amounts of reflection; the time dwelled on CP is essentially ‘dead-time-on-the-wire’ and directly reduces efficiency. This is all the more important since the CP is appended to *each* symbol. Thus the efficiency in real terms is

$$\text{Efficiency} = \text{Symbol Time} / (\text{Symbol Time} + \text{Cyclic Prefix})$$

The graph below shows the drop in efficiency with Cyclic Prefix for the 2k and 4k modes for the US D3.1 transmission.




Figure 2 – US D3.1 Efficiency as a function of Cyclic Prefix

While typical HFC networks have reflections at such a magnitude that they require a CP, of say 2.5us, the fiber plant of RFoG is relatively clean, with minimal reflections implying that it may require a shorter CP duration than the HFC plant, thus affording higher efficiencies. In reality this may not be the case as we will explain in subsequent sections. Briefly, this has to do with the laser turn on time as provided by current standards for the RFoG ONU. In any case, reducing the CP has a direct effect on the efficiency, almost as much as an increase in the order of modulation. There is a case to be made for viewing the CP values as mediating the MER values that are normally used to establish modulation order and capacity.




Figure 3 – Illustrating a D3.1 Frame: Note the Overlaid RPs

In addition to the Cyclic Prefix explained above, there is a Rolloff Period (window size), RP, within the CP which gently ramps the data up at the symbol start and ramps it down at the symbol end thus shaping the symbol in the time domain and maximizing channel capacity by sharpening the edges of spectrum of the OFDMA signal. The RP is always less than the CP and has to follow the values listed in Table 1. Generally, the next symbol has its RP coexistent with the current RP, thus introducing only the net CP value as the appendage factor that determines efficiency. Just as the RP does not affect efficiency, neither does the frame size, because the CP affects each symbol regardless of the frame size. The total Frame Length is thus

$$\text{Frame Length} = (\text{Symbol Time} + \text{CP}) * (\text{Symbols/Frame})$$

In the current paper, we have used the 2k mode (symbol length of 20us) and various CP and RP values. However, by way of example, a CP of 1.875us with an RP of 0.9375us with a frame length of 8 symbols/frame has been a common configuration. This configuration would have yielded an efficiency of 91% and a frame length of 175us.

3.3. The Dizzying Choices of D3.1

With the myriad of customizable options listed above, beginning with the mode of operation (2K/4K for the US and 4K/8K for the DS), the choice of the encompassed spectrum (7.4MHz to 96 MHz for the US and 24MHz to 192MHz for the DS), the modulation formats (up to 4096QAM for US and DS) and the CP choices (0.9375-5.0us for DS and 0.9375-6.25us for the US) and the RP choices (0-1.25us for DS and 0-2.1875us for US), it is clear the D3.1 offers a dizzying array of choices for the operators. For this paper, however, we have been selective in our approach to analyze performance of D3.1 over RFoG. We begin with understanding the current RFoG spec and how the role of 400 kHz mini-slot affects the laser turn on operation. Then we analyze the choice of CP and RP and how they affect turn on timing of the US laser. Then with this understanding, we choose CP and RP, and the RF levels that would enable transmission of D3.1 over RFoG and evaluate the dynamic ranges of operation for the WSO and MDR options. Finally, we offer practical suggestions on enhancing the capacity of D3.1 over RFoG systems making full use of the fiber asset deployed.

4. The Basics: Putting D3.1 and RFoG Together

As indicated earlier in the paper, the industry standard for RFoG is defined in the SCTE IPS 174. This standard, first conceived around 2005 defines various aspects of the ONU, such as US and DS WL ranges, ONU laser output wavelength and power, ONU Laser turn-on and turn-off RF levels and various timing specifications. It is critical to understand the current RFoG standard and its interplay with D3.1 parameters to enable robust D3.1 over RFoG. We begin with a quick summary of the interplay of D3.0 and the current RFoG standard.

4.1. D2.0, D3.0 and the SCTE IPS 174

The current SCTE standard worked well in the D2.0 environment where there was no significant system impairment due to the effects of OBI, but in the D3.0 mode, OBI severely impacted the performance resulting in the proliferation of WSO or the MDR options described earlier. The current standard is summarized below.

Table 2 – RFoG Standard Table Summarized

Simplified Summary of RFoG SCTE IPS 174		
ONU DS WL Range	1540-1565 nm	
ONU US WL Ranges	1610+/-10 nm	
ONU Laser P on	3+/-1.5 dBm	
ONU Laser P off	<-30 dBm	
ONU Laser Turn-On time	<1.3 us	
ONU Laser Turn-Off time	<1.6 us	
ONU Laser Turn-On RF	+7 to +16 dBmV	
ONU Laser Turn-Off RF	+1 to -8 dBmV	
OMI at Total Power	35% +/-3dB @39dBmV	
Nominal RF Input Level/6.4MHz	33 dBmV (17.5% OMI)	
Nominal Number of RF Channels	4	

Since D3.0 begins its preamble transmission 1.5us after the initiation of ramp up, the 1.3us timing introduced by the standard continued to be adequate after the elimination of OBI. Furthermore, the RF levels were chosen so that the ONU would be triggered by legitimate signals, but stop ingress and spurious noise from erroneously turning on the ONU. Since the RFoG standard requires a rather high nominal RF input at 33dBmV/6.4MHz, the turn-on trigger remains quite high (although the standard specifies +7 to +16dBmV as turn-on threshold, most ONUs commonly have a 13dBmV +/-3dB as their turn-on threshold). Additionally, D3.0 operates in burst mode with up to 4 bonded SC-QAM64 RF channels, each typically set to 6.4MHz of encompassed spectrum. Thus the total RF load is normally less than 26MHz and requires a rather modest MER to support 64QAM upstream operation. Finally, only a maximum of 4 ONUs are likely to transmit simultaneously, and with OBI elimination, the turn-on timing, the turn-on level, and the modest SNR requirement enables fair transmission of D3.0 over the current standard.

4.2. Understanding D3.1, D3.0 and the SCTE IPS 174

A move to D3.1 requires a finer understanding of the RFoG standard and how it impacts the system. At its heart, we need to understand and resolve the three dynamic ranges simultaneously. These are described next.

4.2.1. Optical Receiver Dynamic Range

Typical RFoG optical receivers are designed with low noise, high sensitivity and high gain since the optical input can be quite low. In typical D3.0 applications, the optical input to the headend receiver is around -19 dBm/ONU and even if all 4 ONUs turn on simultaneously, the optical input does not exceed -13dBm in total. Typical RFoG receivers can handle up to -10dBm of optical input but higher optical level could overload the receiver and cause unwanted effects.

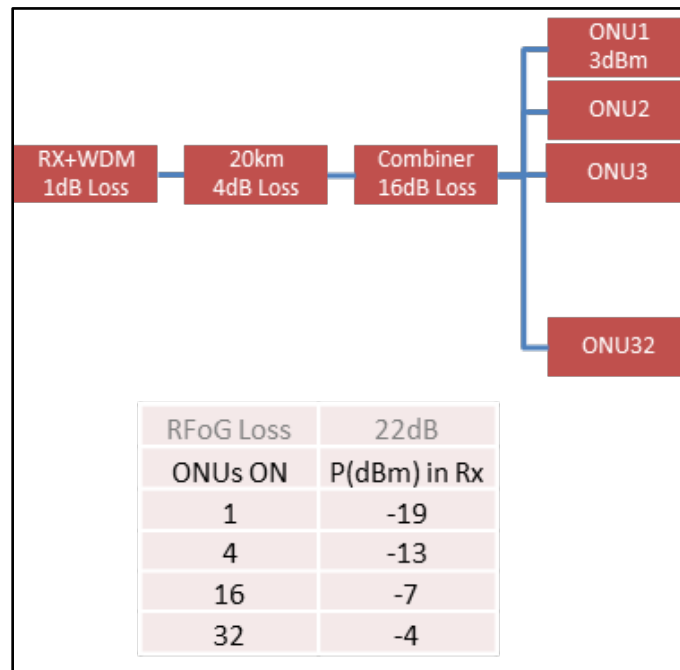


Figure 4 – Illustrating Optical Rx Dynamic Range

In a D3.1 environment the same -19dBm/ONU of power could result in up to -7dBm of power for 16 ONUs on simultaneously and up to -4dBm of power if all 32 ONUs turn on simultaneously. These power levels introduce the risk of overloading typical RFoG receivers. Of course, the optical level overall could be decreased, but that would lead to a lower operational MER resulting in reduced system performance.

There are two ways to overcome this effect. One way would use standard (non-RFoG specific) return receivers. Even though these have slightly higher EIN and lower gain, their lower sensitivity enables a higher optical power and ensures an operational dynamic range. Alternatively, we would use an MDR in the continuous mode. In this case, the optical level to the headend receiver is independent of the optical level at each of the MDR distribution ports and thus ensures a wide operational dynamic range.

4.2.2. ONU RF Input Dynamic Range

The operational range of an ONU is a significant determining factor of the overall dynamic range in the system. At the low RF input levels, it is limited by the SNR or MER availability and at high RF input levels it is limited by the ONU laser clipping. In a D3.0 environment where there are typically only 4 RF SC-QAM channels, the total RF level is only 6dB higher than the RF level of each individual SC-QAM.




Figure 5 – Illustrating ONU RF Input Dynamic Range

However, for D3.1, the RF transmission occurs in mini-slot sizes of 400 kHz, therefore the smallest RF input can be at 400 kHz wide and the largest RF signal could occupy the entire 96 MHz band. The implications of this are that the total RF level at the ONU RF input may be 24dB higher than the RF level of a 400 kHz mini-slot. This is illustrated in the figure below:




Figure 6 – Illustrating Typical RF Levels: D3.0 left and D3.1 right

Note however that practical US systems today are limited to 42MHz for the low splits, 85MHz for mid-splits and 204MHz for high split applications. Therefore the RF levels depicted here are for illustration purposes and the exact RF dynamic range is determined by specific encompassed spectrum in each system.

Recall that the typical implementation of the RF turn-on level is 13 ± 3 dBmV. As a result, we would need at least 16 dBmV in 400 kHz to turn the ONU Laser on. However, if all 96 MHz (note that in a 204 MHz upstream, there may be up to two 96 MHz OFDMA channels) of available spectrum were utilized, the total power would be 40 dBmV, which is close to the onset of clipping. Notice that in 400 kHz case, the system performance is not limited by the available SNR or MER, which likely would be more than sufficient at RF levels below 16 dBmV/400 kHz, but that it is limited primarily by the laser turn on itself. Overall, this is a potential limitation that could substantially limit system operational range.

There are two possible ways to handle the wide RF range requirement. In the first way, we would limit the total RF encompassed spectrum. Limiting the spectrum to 25.6 MHz for example as is currently the case, the RF level range is 18 dB and therefore affords 6 dB additional operational dynamic range. The other way would be to let the ONUs run in Continuous mode (CTM). Note here that only the ONU laser is always on, but the CM is still operating in burst mode. This will radically open up the dynamic range since the laser is always on; there is no requirement to limit operation to 16 dBmV minimum RF. As mentioned earlier, systems would provide sufficient SNR or MER below 16 dBmV/400 kHz. This is illustrated in the NPR curve below.




Figure 7 – Illustrating a Typical NPR curve and RF Turn-On Levels

The figure above illustrates a typical NPR curve of an ONU with -19dBm input at the headend receiver with RF levels represented per 6.4MHz bandwidth. The shaded red box illustrates the 13+/- 3dBmV/400kHz laser turn-on level prorated to a 6.4MHz bandwidth. It is clear that the RF input could be lower than the lowest RF laser turn on level and the ONU would still have produced sufficient SNR to resolve 1024QAM signals. If 256QAM were required, the system could have gone even lower. However, following the SCTE standard results in a dynamic range that is rather limiting.

4.2.3. The MER-BER Dynamic Range

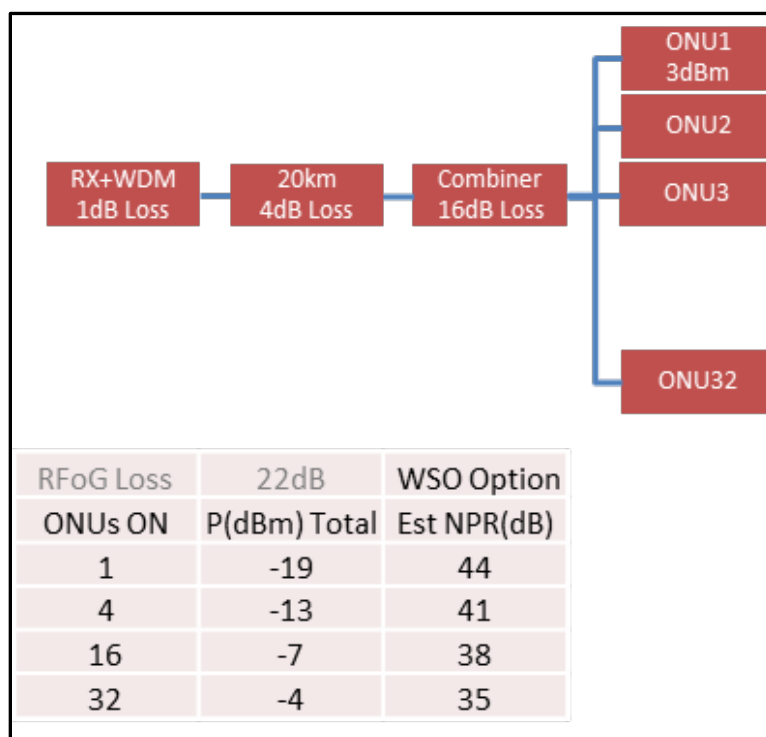


Figure 8 – Illustrating the MER-BER Dynamic Range

While the above section makes a compelling case for a continuous mode operation, if all ONUs operated in the continuous mode (CTM), then it could result in a substantial increase in the overall noise level at the headend receiver. As we have shown in earlier sections, with 16 ONUs simultaneously on, the optical level increases by 12 dB, bringing with it an increase in the noise floor due to the additional shot and RIN noise of individual ONUs, thus affecting the left side of the NPR curve shown in Figure 7. In practice the RF dynamic range improvement obtained by operating in continuous mode could be tempered by this additional noise. However, one thing to note here is that D3.1 uses the powerful low density parity check (LDPC) codes that work exceptionally well in AWGN environments (not so much in clipping), so that the system may function even below the specific performance levels specified in the D3.1 standard subject to the CMTS implementations. In other words, faced with a choice of dynamic range limitations due to laser clipping or the noise floor, one would always want to choose to operate closer to the noise floor and take potential advantage of LDPC.

5. The Single ONU Simulated D3.1 Tests

Testing a new protocol such as D3.1 with its dizzying array of customizations in a multi-ONU CMTS setting is fraught with choices. It was necessary to find a way to accomplish a large set of tests to help establish the operational parameters that could then be used in a multi-ONU CMTS test bed. Therefore, we began our testing with a single ONU in a test bed with Rohde & Schwarz CLGD-FSW test gear. Briefly, the Cable Load Generator (CLGD) can generate frames of burst US OFDMA test signals with various CP and RP values. When the CLGD is connected to the FSW signal analyzer thru the ONU and a

Headend Receiver, the FSW can read the incoming burst mode signal and provide MER values (but no BER values at this time).

5.1. CLGD-FSW Test Configuration

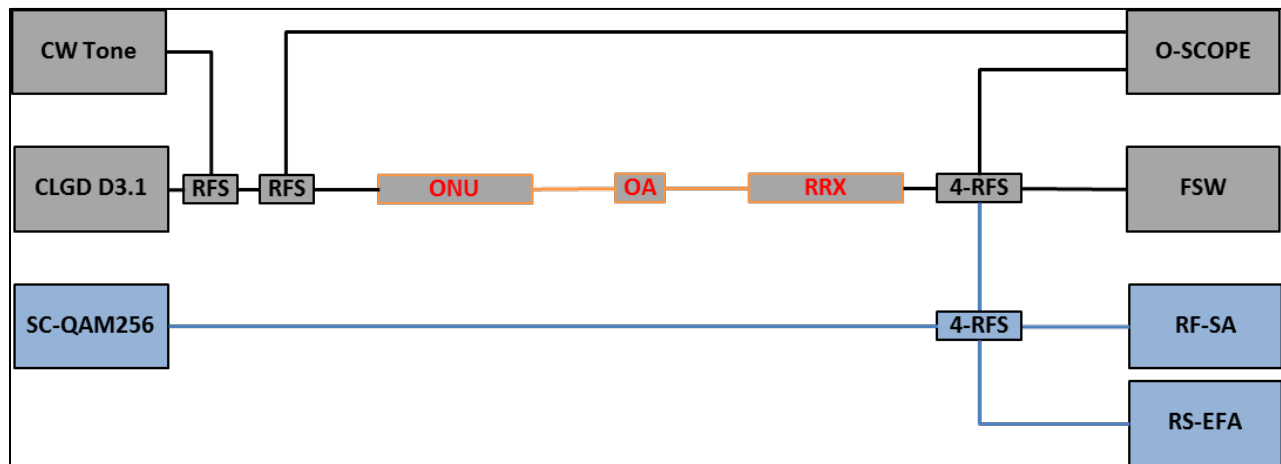


Figure 9 – CLGD-FSW Single ONU Test Setp

The test bed comprised a CW tone generator that was coupled in with the CLGD so that we could enable CTM operation of the ONU when needed. A part of the input signal was available on an O-Scope for time domain analysis. The RF level out of the headend receiver was split and was made available to the O-Scope and to the FSW. We also had a known SC-QAM256 ‘probe’ channel for reference, one that combined with the output of the headend receiver so that we could compare the D3.1 burst signals to a known continuous signal as the experiment progressed.

The CLGD would create 5000 frames of US data and send them thru the ONU while the FSW would record the mean MER as well as the maximum and minimum MER values. One could then vary the encompassed spectrum, the RF level, the OFDMA profiles, the FFT modes, the CP, and the RP via CLGD commands. In total, over 50 files were created that could simulate the set of options that D3.1 provides.

Furthermore, the CW tone could enable the ONU to be in BTM or CTM mode of operation. Finally, the ONU was customizable and the turn on timing could be adjusted as needed. Each of these settings would then be tested using the options described above. All in all, we tested hundreds of combinations to understand the operational set for D3.1 testing over multiple ONUs and over a real CMTS.

The first set of tests used a standard SCTE IPS 174 compliant ONU, with a turn-on time around 1.3us. The CP was increased from 0.9375 all the way to 6.25us as specified by the D3.1 spec. In each case, the RP value was set to 0.3125us. The ONU was then set to the CTM with the CW tone and the CLGD passed the frames across the ONU to the receiver and were analyzed by the FSW. As expected, the mean, maximum and minimum MER values in each test were very close to each other indicating that the performance was unaffected by the choice of a CP.

However, when the ONU operated in a burst mode, the minimum MER value fell quite dramatically each time the test was run with a CP value less than the laser turn on time. Even with the CP value well above

the laser turn on time, the FSW continued to show a small reduction in the minimum MER relative to the mean MER. The reduction in minimum MER even with large CP values indicates that the laser turn-on process itself impacts performance and is briefly examined in a subsequent section. Changes in RP did not seem to affect the performance significantly.

5.2. Single ONU Measured Results




Figure 10 – Measured MER for Single ONU using the CLGD and FSW Test Equipment

The results above indicate that the CP values for standard SCTE IPS 174 ONUs that operate in the BTM must be chosen above 1.5us, preferably at least 1.875us to accommodate the laser turn-on and turn-off requirements. No such restrictions need be observed in the case of ONUs that operate in CTM. As explained earlier, larger CP values reduce throughput efficiency, therefore values of CP between 1.875 to 2.5us may be a tradeoff for robust and efficient systems.

6. Additional Considerations

6.1. Initial Range and Register

Initial Ranging and Registering in D3.1 differs slightly from the same process in D3.0. In D3.0, the range and register signals are of a lower modulation format and require a lower MER than the data signals themselves. In D3.1, the range and register signals use the same modulation profile as the data signals and therefore, adequate MER must be established before the ranging & registering process can be completed. It is therefore critical that the system robustly support the required MER before the CMs are set to range and register.



Figure 11 – Time Domain Captures of Range and Register Signal and Signals at 50Mbps

While ranging and registering is a multi-step process, different than the one used to transmit data signals, the CM still uses the ranging signals in the 400 kHz mini-slot. Presented above is a time domain capture of the range and register signal on the left. As can be seen, the signal amplitude is rather limited when compared to a burst of 50Mbps of traffic which has been captured in the time domain graph shown on the right. Therefore, the RF levels into the ONU must feature an RF level which is high enough within 400kHz to trigger the ONU laser on for successful range and registering. Range and register signals are sent periodically in addition to station maintenance, therefore the MER of the system must be maintained to prevent time out conditions.

6.2. Laser Turn on Effects: CP, RP and OMI

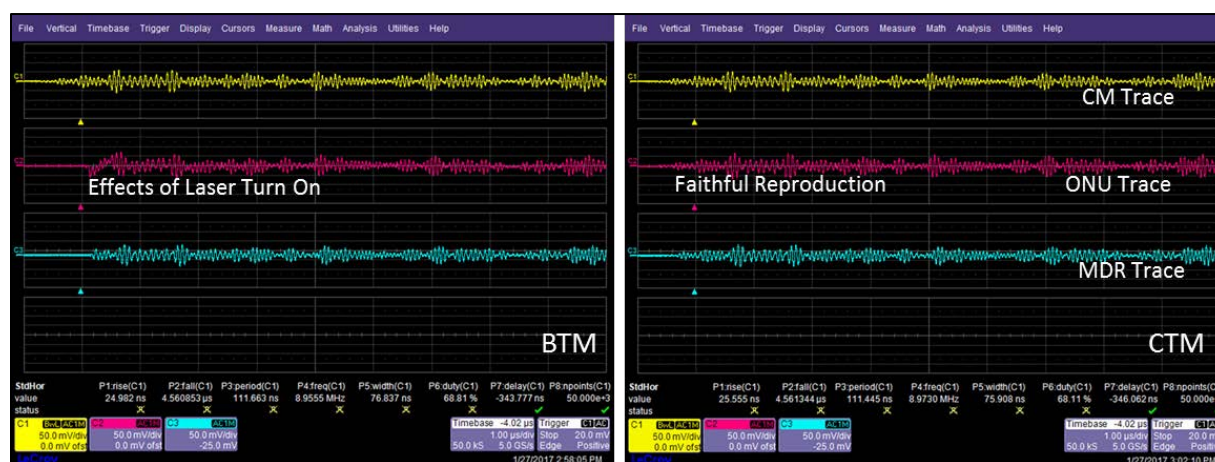


Figure 12 – Time domain capture of the effects of the Laser turn on

Presented above are time domain captures that zoom in on the initial few micro-seconds of the ranging signal. We configured the O-scope to provide outputs of the time domain signal from the CM, the ONU directly connected to the receiver and the ONU thru a coupler to the MDR (with the output of the MDR connected to a receiver). On the left is the time domain capture with the ONU in the BTM and on the right is the same time domain capture but with the ONU in the CTM. The effects of laser startup are

visible in the figure on the left. It is seen here that a part of the signal is cut off and a part of the signal is experiencing RF level variation consistent with laser turn on.




Figure 13 – OMI of an ONU at Turn-on and potential for Collateral Clipping

There is a time delay (100-500ns) before the laser begins the process of turning on. From initial laser turn on to maturity (90% of stated output power), the laser experiences a very high OMI. While this high level of OMI is experienced only for a short period of time (500 – 1000ns), the laser experiences significant clipping during this turn-on time (note that the RP may mitigate this effect to a small extent). If the CP exceeds the laser turn on time, then the CMTS is still able to lock to the symbol and get information within the frame. However, if the laser turn on time exceeds the CP, then substantial packet loss is observed.

Since all ONUs share the same headend receiver, if one laser is clipping, a different laser that fully turned on (and otherwise not clipping itself), could experience packet losses due to the collateral effects of the other ONU(s) with lasers in the initial turn-on process. This effect also occurs in D3.0 RFoG systems, but could happen more frequently in D3.1 RFoG systems due to the possibility of a higher number of ONUs turning on and off. Effects of collateral clipping are difficult to detect, in systems that are already noise and OBI prone, but must be accounted for in multi-CM environments.

6.2.1. A word about CTM: Ingress and CP

Ingress can be a significant concern. Whereas the burst mode operation of the ONU and the RF input level threshold limit ingress, we have seen that these also limit dynamic ranges. One way to potentially resolve this would be to reduce the laser turn-on threshold as a means to increase operational dynamic range. However even with said decrease in the laser turn-on levels, there is a fair amount of low frequency noise laser clipping that could inherently limit the operational range. Readers are urged to contemplate the benefits and constraints described here in the context of their own networks in deciding upon CTM or BTM operation.

We have seen earlier that appropriate selection of Cyclic Prefix values are required for the smooth operation of D3.1. However, we have also seen that large values of CP are ‘dead-time-on-the-wire’ and limit efficiency.




Figure 14 – Illustrating the Role of CP in enabling throughput

As can be seen in the chart above, a lower value of CP would enable a higher throughput at lower SNR or MER. Alternatively, at the same SNR, one could achieve higher throughput. In the above example, the ability to migrate from 2.5us CP to 0.9375 CP could result in a half order of improvement in modulation profile. We anticipate that HFC systems will generally need a higher CP value set to accommodate the multiple reflections present in the HFC plant. However, in an RFoG environment where these reflections are minimized, operating in CTM may allow for the use of lower CP values thus providing additional throughput if needed.

Table 3 – Trade Off Analysis of BTM and CTM Operation

	Effects	ONUs in BTM	ONUs in CTM
Ingress, Noise Funneling Observed	No	No	Yes
Noise Floor Rise Observed	No	No	Yes
Laser Turn On Clipping	Yes	Yes	No
Collateral Clipping	Yes	Yes	No
CP Values Restricted: Laser Turn On Time	Yes	Yes	No
Dynamic Range Restricted: Laser Turn On Level	Yes	Yes	No

Notice that since an ONU in CTM will always transmit a mature optical signal, the performance is independent of the CP length, the laser does not suffer turn on clipping and does not exhibit the effects of collateral clipping. As mentioned once before, and indicated in the table above, these trade-offs must be contemplated in the context of ingress mitigation before a decision for the CTM or BTM mode transmission is made.

7. System Tests

7.1. Test Configuration

After understanding single CM/ONU performance, we now move to test in the multi-CM/ONU environment.

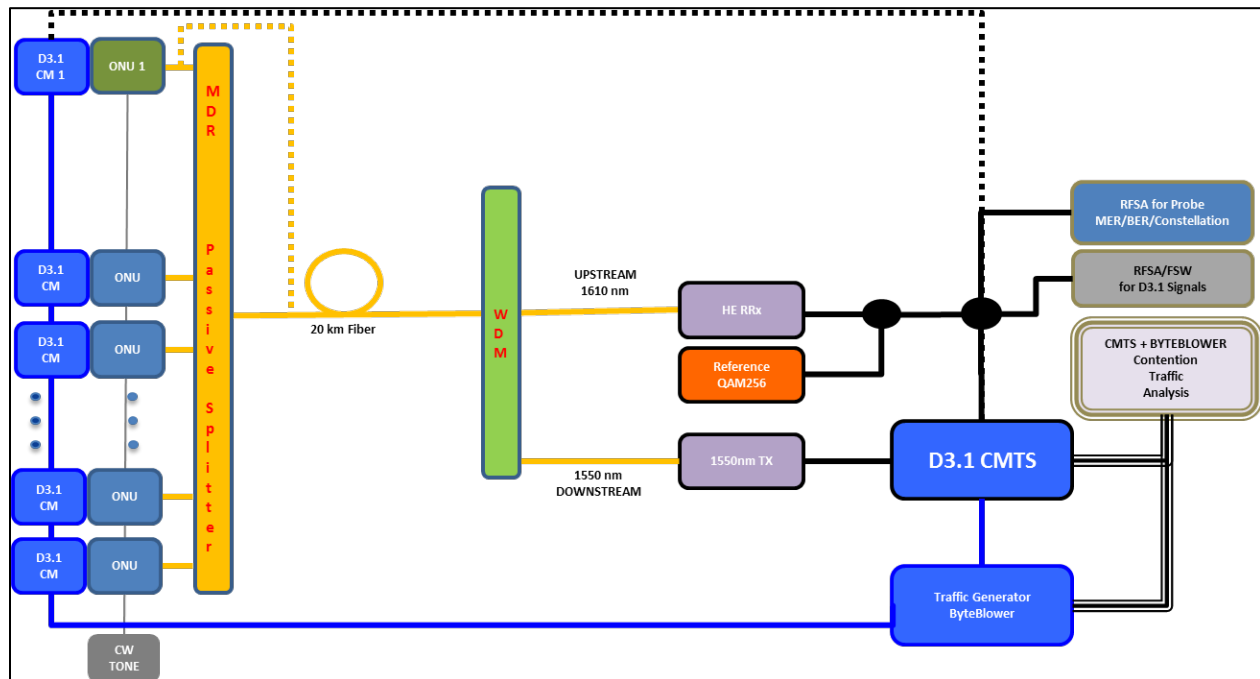


Figure 15 – 16 CM/ONU System Configuration

Presented above is the test configuration used for testing. We used the SB8200 D3.1 compatible CMs for our tests, with wavelength selective CP801TU for the ONUs. The ONUs selected all have a laser turn-on level of around 13dBmV. Each ONU optical output was split by a 3dB splitter that enabled an MDR and a regular optical splitter to be accommodated. The output of the MDR or the splitter traversed 20 km and passed thru a Wavelength Division Multiplexer (WDM) to reach the regular HFC upstream headend optical receiver. The output of the receiver was split with one output going to a set of displays and another entering the UCAM card of the E6 D3.1 CMTS. We also introduced a steady probe channel of SC-QAM256 at the output of the receiver so that we may evaluate RF levels and also identify egregious clipping events as they occurred during test. The DCAM card of the E6 CMTS was connected to a Directly Modulated 1550nm transmitter and connected to the WDM thus completing the two way link.

A CW tone generator was filtered and split so that each ONU would receive a steady RF signal to force it in CTM as needed. This steady tone was then combined with the D3.1 output as presented as input to the ONU. A ByteBlower traffic generator supplied traffic to the set of CMs thru a switch and also analyzed the traffic as it came back to the CMTS for frame loss and other information. The entire information was fed to an Excel based analysis engine that would then determine the maximum, minimum frame losses of each ONU to determine if the said test was successful, the criteria of success being a less than 1% maximum frame loss over any ONU under test. The CMTS itself can provide a set of diagnostic

information, most importantly the RF levels out of the CMs for the initial setup and continued operation of the ONUs.

7.2. System Simulations




Figure 16 – Illustrating NPR estimates for the MDR and WSO cases

Presented above are NPR simulations for the MDR and WSO options. These simulations are presented for cases of 1, 4 and 16 ONUs transmitting simultaneously. Recall that the WSO option is completely passive while the MDR option requires powering at the MDR location, and that both options eliminate OBI. It is seen here that the NPR decreases with more ONUs simultaneously operating. While more ONUs reduce NPR, the statistics might suggest that the time duration of simultaneous operation may be quite limited and it might be more likely that fewer than the maximum number of ONUs are on generally in the BTM operation. Note however that we have not taken into account in this NPR simulation any of the other effects we have discussed in previous sections such as self or collateral clipping.

In all of the tests below, we have used a maximum of 16 CM/ONUs, this is because current state of the art allows for 16 WL options for the WSO option. However, we have allowed for a 32 way splitter (both for the MDR and WSO) test bed in anticipation of higher quantities of tests as technology evolves. The optical input to each MDR port is 0dBm, while the optical input to the headend receiver is -6dBm, leading to a total link estimate of 29dB. In the WSO case, the output of the splitter traverses 20 km and reaches the headend receiver at -19dBm, thus leading to a 22dB link.

There are two dashed horizontal lines, one is at 35.5dB and identified as for OFDMA1024 operation, the other dashed line 6dB below is for the OFDMA256 operation. Note that on the noise side, if the ONUs were operated in the BTM, the performance may be limited by laser turn-on levels rather than by the availability of signal to noise margin. Thus the familiar improvement in the NPR dynamic range that would accrue by reducing the NPR requirement from OFDMA1024 to OFDMA256 would not result in an improvement in operational margin for a real system because the ONU would not have turned on below the RF level. However, in the CTM case, even with a reduction in NPR as more ONUs are added, since the ONUs are always On, even with the increase in noise floor, the RF levels that were previously not available due to turn-on level limitations would now operate and provide sufficient NPR to enable wider dynamic ranges. Remember that this is a lab environment without the presence of ingress and therefore

the observed dynamic range may be optimistic. To that end, additional tests would be required to obtain more rounded evaluation of the CTM and BTM systems.

For US tests in this section, we used 22MHz of encompassed spectrum from 20MHz to 42MHz. The mode used was the 2k FFT, with a CP value of 1.875us and an RP value of 0.9375us. The frame size was with 8 symbols and 175us in duration. The traffic used was UDP/IP with uniformly distributed 64-1500 Byte packets. This was chosen to simulate the type of traffic encountered with 64 Byte packets simulating TCP 'acks' for OTT content like Netflix and 1500 Byte packets simulating YouTube uploads.

7.3. Measured Results

7.3.1. RF Levels in the 16CM/ONU System

Table 4 – RF Levels into the ONUs for the MDR and WSO Systems

RF Levels dBmV/1.6MHz	MDR System ONU RF in	WSO System ONU RF in
<i>CM-ONU 1</i>	27.5	27.0
<i>CM-ONU 2</i>	27.0	26.5
<i>CM-ONU 3</i>	25.3	24.8
<i>CM-ONU 4</i>	28.3	28.3
<i>CM-ONU 5</i>	28.3	27.8
<i>CM-ONU 6</i>	27.3	27.3
<i>CM-ONU 7</i>	27.5	27.0
<i>CM-ONU 8</i>	28.0	27.8
<i>CM-ONU 9</i>	27.3	26.8
<i>CM-ONU 10</i>	26.3	25.8
<i>CM-ONU 11</i>	25.3	25.0
<i>CM-ONU 12</i>	26.8	26.3
<i>CM-ONU 13</i>	26.5	26.5
<i>CM-ONU 14</i>	29.0	28.8
<i>CM-ONU 15</i>	23.8	24.0
<i>CM-ONU 16</i>	28.3	28.3
Mean RF	27.0	26.7
Max RF	29.0	28.8
Min RF	23.8	24.0
SD	1.4	1.3
Max-Min	5.3	4.8
<i>Mean RF 4</i>	27.0	26.6
<i>Max-Min 4</i>	3.0	3.5

When the CMTS is connected and the CMs are allowed to range and register after establishing DS and US MER, the CMTS diagnostic RF levels into each ONU are tabulated as shown above. It is common to see around 6 dB of variation in RF levels across the ONUs. This is due to ONU unit to unit variations, link variations for individual ONUs to the CMTS, and the resolution of the CMTS long loop ALC itself.

This variation is further enhanced when temperature and aging effects are taken into account. These variations cause the NPR curves of individual ONUs to be offset horizontally and vertically, and sometimes to shift continually in response to environmental factors. Therefore, the effective range of operation of the system could be considerably reduced from what would appear to be the dynamic range of an individual ONU.

7.3.2. Dynamic Sliver and Dynamic Range

In this paper, we introduce the concept of a ‘Dynamic Sliver’ to take into account the above mentioned variations. The Dynamic Sliver is defined as the range of attenuation that may be applied at the CMTS input that enables all individual CMs and ONUs to operate with essentially no frame loss. With this understanding, we tested a 16 CM RFoG configuration over a 20 km link and into a D3.1 CMTS. A 16 CM option was chosen because it allows for both the WSO and MDR options to be tested.

7.3.3. Detailed Test Scenarios

As part of the analysis, we tested the following 16 conditions

1. 1024OFDMA with 10Mbps/CM leading to total of 160Mbps traffic for 16 ONUs
 - a. MDR with ONUs in CTM operation
 - b. MDR with ONUs in BTM operation
 - c. WSO with ONUs in CTM operation
 - d. WSO with ONUs in BTM operation
2. 256OFDMA with 8Mbps/CM leading to a total of 128Mbps traffic for 16 ONUs
 - a. MDR with ONUs in CTM operation
 - b. MDR with ONUs in BTM operation
 - c. WSO with ONUs in CTM operation
 - d. WSO with ONUs in BTM operation
3. 1024OFDMA with 1Mbps/CM to simulate a lightly loaded system for 16 ONUs
 - a. MDR with ONUs in CTM operation
 - b. MDR with ONUs in BTM operation
 - c. WSO with ONUs in CTM operation
 - d. WSO with ONUs in BTM operation
4. 1024OFDMA with 10Mbps/CM with 4 ONUs simulating an RFoG system with minimal penetration
 - a. MDR with ONUs in CTM operation
 - b. MDR with ONUs in BTM operation
 - c. WSO with ONUs in CTM operation
 - d. WSO with ONUs in BTM operation




Figure 17 – Measured System Dynamic Slivers and Dynamic Ranges

7.3.4. Discussion of Test Results

Presented above is a visual summary of the collated test results. As indicated, the point of reference is the attenuator setting just prior to the CMTS. An increase in attenuation results in the RF levels to the ONUs to increase and vice-versa. With this understanding, it is noted here that the ranges represented in hues of Green refer to the ‘Dynamic Sliver’ that is available for the 16 ONU system, whereas the additional Yellow hues represent what would have been the dynamic range had a single ONU been used with data traffic but with all other ONUs transmitting light.

By way of example, in test condition 1.a, all 16 ONUs are operational in CTM, but only one ONU is passing traffic, and therefore that ONU would have perceived the dynamic range as wide as from the attenuator setting of -12 to +3, a full 16dB Dynamic Range. However, in a system, due to the RF level variation described earlier, all ONUs together have an error free operational range of the CMTS attenuator setting of -9 to 0dB, thus indicating only a 10dB of usable dynamic range. As fewer ONUs are operating, the available Dynamic Range increases.

Test Condition 1 indicates that the on-set of clipping for 1024OFDMA with 10Mbps loading has occurred at the attenuator setting of 3, which corresponds to an RF level of 26dBmV/1.6MHz. This translates to 32dBmV/6.4MHz, and 38dBmV/25.6 MHz, fairly close to the onset of clipping as seen by previously measuring pre FEC BER with SC-QAM256 testing for this ONU. Furthermore, from observing results from 1.b, which indicates BTM operation, it is seen that the ONUs did not turn on until the RF input level reached 19dBmV/1.6MHz of power, this is equal to 13dBmV/400kHz, which is the specified RF input level required to turn on these ONUs. It is seen here that the CMTS Dynamic Sliver is drastically reduced for BTM operation primarily due to the SCTE IPS 174 spec that specifies the laser turn-on parameters. In the 1.c test case, the wavelength selective ONUs in CTM operation have a naturally lower dynamic range and consequently a smaller Dynamic Sliver. However, the dynamic sliver in the 1.c case still exceeds the 1.d case, which is limited by the laser turn-on.

When 256OFDMA is used, it is observed that the maximum throughput is reduced slightly from 10Mbps to 8Mbps, with a slight improvement in the clipping performance of the system. While the CTM operated ONUs are able to take advantage of the lower operational NPR needed for OFDMA256 operation relative to the OFDMA1024 operation, the ONUs in the BTM operation are unable to take advantage of this due to the RF turn-on level restrictions and as a result have a rather restricted Dynamic Range and Dynamic Sliver. The effect of CTM operation can be more important as the modulation complexity decreases.

For test condition 3, it is noted that as the utilization decreases and the traffic becomes lighter, the onset of clipping is delayed by a few dB, since the ONUs are generally transmitting less RF spectrum (please see Figure 7). However, there is no improvement observed on the noise side consistent with the NPR curves and our earlier discussion. Thus the improvement in Dynamic Range and consequently of the Dynamic Sliver is all due to the delayed onset of clipping as a result of lower utilization and the more robust nature of 256OFDMA relative to 1024OFDMA.

Finally for test condition 4, we see that as we reduce the number of ONUs to simulate a very lightly penetrated system, the Dynamic Range of the CTM system becomes better due to an improvement in signal to noise with lesser number of ONUs (please see Figure 16). The Dynamic Sliver also improves, because of reduced link variations using these 4 ONUs selected for test (in this example). Note however, as expected, there is no improvement in the onset of clipping and of the laser turn-on.

It is anticipated that this set of tests may be augmented in the future with a larger set of ONUs and with expanded test conditions.

8. Conclusions

This paper represents the first step towards gaining a complete and full understanding of the implications of supporting D3.1 US capability in existing and to be deployed RFoG Systems. The discussion in this paper and the set of test results presented are very important beginnings of establishing D3.1 performance standards over RFoG.

The increased performance requirements characteristic of D3.1 along with the ability of a significantly higher number of CMs to transmit simultaneously make the matter of OBI elimination a much more urgent matter. While D3.0 and existing RFoG standards have been a good match, D3.1 and its many different aspects such as OFDMA operation, higher orders of modulation, increased flexibility in channel width, Cyclic Prefix, Roll-off Period, OFDMA FFT size, minimum performance requirements, and encompassed spectrum, call for the operator to carefully consider complex tradeoffs to optimize operating conditions for their RFoG networks.

In this paper, we described relevant D3.1 parameters and linked them to the physical characteristics of an RFoG network. We examined performance of individual ONUs under various D3.1 parameters to understand optimum operating conditions of ONUs in burst mode. We then utilized two well-known methods of OBI elimination in a complete multi-ONU RFoG environment which was subjected to the vast array of D3.1 parameters, and we observed performance in the BTM and CTM ONU modes of operation.

D3.1 enables the industry to enhance its capacity and flexibility, while RFoG presents a way for Cable operators to continually push fiber deeper into their networks. In support of this goal, it is anticipated that this set of tests may be augmented in the future with a larger set of ONUs and with expanded test conditions.

9. Acknowledgments

The authors would like to thank their ARRIS colleagues Tom Cloonan, Ayham Al-Banna, Michael Dehm, Larry Spaete, Ron Miller, Adam DeBelle, Ketan Gadkari, Dave Hartrum, and Vip Rathod for their valuable contributions to the paper. It is also a pleasure to acknowledge Doug Jones, James Lin and Aaron Quinto of the CableLabs for helpful discussions.

10. Abbreviations

4-RFS	4 Way RF Splitter
AGC	Automatic Gain Control
ALC	Automatic Level Control
BER	Bit Error Rate
BTM	Burst Mode Operation
CLGD	DOCSIS® Cable Load Generator
CMTS	Cable Modem Termination System
CMTS	Cable Modem
CP	Cyclic Prefix
CTM	Continuous Mode Operation
CW	Carrier Wave
D2.0	DOCSIS® 2.0
D3.0	DOCSIS® 3.0
D3.1	DOCSIS® 3.1
dBm	Decibels relative to one milli-watt
dBmV	Decibels relative to one milli-volt
DS	Downstream
EIN	Equivalent Input Noise
FFT	Fast Fourier Transform
GHz	Gigahertz
HFC	Hybrid Fiber Coax
IPS	Interface Practices Subcommittee
kHz	Kilohertz
LDPC	Low Density Parity Check
Mbps	Megabits per Second
MDR	Multi-Diode Receiver

MER	Modulation Error Ratio
MOD	Modulation
MHz	Megahertz
NPR	Noise Power Ratio
ns	Nanoseconds
OA	Optical Amplifier
OBI	Optical Beat Interference
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OMI	Optical Modulation Index
ONU	Optical Network Unit
O-Scope	Oscilloscope
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RF-SA	RF Spectrum Analyzer
RFoG	RF over Glass
RP	Roll-off Period
RRx	Return Receiver
RS-EFA	Rohde & Schwarz EFA Test Receiver
Rx	Receiver
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SNR	Signal to Noise Ratio
UDP	User Datagram Protocol
US	Upstream
us	Microseconds
WDM	Wavelength Division Multiplexer
WL	Wavelength
WSO	Wavelength Selectable ONU

11. Bibliography & References

CableLabs, “DOCSIS 3.1 Physical Layer Specification – CM-SP-PHYv3.1-I11-170510”, May 10, 2017

CableLabs, “DOCSIS 3.0 Physical Layer Specification – CM-SP-PHYv3.0-I13-170111”, May 10, 2017

SCTE Interface Practices Subcommittee, “ANSI/SCTE 174 2010 – Radio Frequency over Glass Fiber-to-the-Home Specification”

Sebnem ZorluOzer, Ph.D., Venk Mutalik, A. Vieira, Ph.D., and J. Chrostowski. “*From OBI and SNR to Ookla and Netflix: How Network Impairments affect Customer Perceptions: The role of Leading and Lagging Indicators as We Evolve HFC to FTTP*”; SCTE Cable-Tec Expo 2015

Doug Jones, and Curtis Knittle. “*Operational and Efficiency Impacts of the DOCSIS 3.1 Upstream over RFoG*”; CableLabs Technical Brief, March 2017

Venk Mutalik, Marcel Schemmann, Zoran Maricevic, and John Ulm. *“The Yin and the Yang of a Move to All Fiber: Transforming HFC to an All Fiber Network while Leveraging the Deployed HFC Assets”*;
2015 INTX NCTA Spring Technical Forum

The Universality of Modulation

A Technical Paper Prepared for SCTE•ISBE by

Tom Williams
Principal Architect
Cable Television Laboratories, Inc.
858 Coal Creek Circle
Louisville CO 80027-9750
303-661-3486
t.williams@cablelabs.com

Introduction

This paper examines the common properties of modulation methods that use orthogonal basis functions, including single carrier (e.g., 64-QAM or 64-state quadrature amplitude modulation), multi-carrier (e.g., OFDM (orthogonal frequency division multiplex)), and spread spectrum. It also explains a time-frequency swapping technique that reveals OFDM and single carrier modulation are duals of one another, if the time axis and frequency axis are relabeled. This time-frequency swapping technique allows new modulation candidates to be created, such as duobinary (partial response signaling) OFDM, which has valuable properties for the cable plant.

Universality of Modulation

1. Modulation Techniques

Modulation techniques are used at carrier frequencies to send digital data over a distance, either by wires, wireless, or optically. Three popular modulation methods are single carrier, multi-carrier, and code division multiple access. All three have been used on cable networks at one time or another.

Orthogonality between signals is a property that allows one signal, comprising symbols, to be clearly received without interference from other signals' symbols.

$$\sum x(n) \cdot y(n) = 0 \quad [\text{eq. 1}]$$

Two variables, x and y, are orthogonal over a range if:

Sum $x \cdot y = 0$ over a range in $y \neq x$. [footnote 1]

Different modulation techniques use different techniques to achieve orthogonality.

In cable, single carrier modulation has been used extensively on downstream signal paths; examples include 64-QAM and 256-QAM. ATDMA (advanced time division multiple access), essentially a burst-mode single carrier transmission technique, has been used on upstream signal paths. Single carrier modulation comprises a time series of voltage impulses (symbols) that have been filtered to limit interference with other frequency bands. Figure 1 is a voltage vs. time diagram showing five $\sin(x)/x$ impulses with uniform time shifts. The illustrated different symbols are the same value, although shifted in time. The symbols can have any positive or negative values and have real-only values or be complex. The sampling instants are illustrated with five vertical lines. The waveforms do not interfere with each other because at each sampling instant, one symbol is at its peak while the others are passing through zero. Thus, orthogonality is maintained. For this system to work optimally, linear distortions, such as echoes, need to be removed prior to sampling, typically using an adaptive equalizer. Otherwise, the responses from the other symbols will not be at zero at the sampling instant and they will contribute distorting energy.

Figure Note: A low speed data source is exclusive-or'd with a high-speed PN (pseudo-noise) sequence to produce an output that appears noise-like.

Figure 3 is a block diagram for DSSS showing the transmission and reception of a spread-spectrum signal. At the receiver, a PN generator must be using both the same code as the transmitter, and be synchronized with it. With S-CDMA DOCSIS functionality, each circular time shift (excluding the initial chip which is not shifted) produces another basis function that is orthogonal to all the other shifts.




Figure 3 - Diagram for Direct Sequence Spread Spectrum

OFDM is used in DOCSIS® 3.1 technology as well as in numerous wireless standards. With OFDM many different subcarriers that are all harmonics of a fundamental are used to obtain orthogonality. Figure 4 shows an OFDM waveform with only four such subcarriers. Each of the harmonically-related subcarriers has a different magnitude and phase value. When all four subcarriers are combined (summed) for transmission, the result is a single composite signal. However, orthogonality allows the original subcarriers to be separated at the receiver, usually with a fast Fourier transform (FFT). A guard interval (GI), which is also known as a cyclic prefix (CP), is made by copying samples from the end of the signal and pasting it onto the beginning. This is done so that equalization can be performed with a circular convolution (or equivalent) and no interference is suffered from the previous OFDM block if there is an echo on the channel. This assumes that the echo is shorter than the GI.




Figure 4 - An OFDM Signal in the Time Domain Containing Four Subcarriers

Note that in Figure 4, an OFDM signal in the time domain contains only four subcarriers, first, second, third, and fourth harmonics.

Figure 5 show the same OFDM signal of Figure 4, but in the frequency domain.




Figure 5 - The OFDM signal of Figure 4 Viewed In The Frequency Domain

Figure 6 shows a spectral plot of an OFDM signal that is affected by a deep frequency-selective fade. This fading occurs frequently in wireless channels where a sum of echo components cancels the signal, or at least at some subcarrier frequencies. This is an environment ideal for OFDM where the faded subcarriers lost in the noise can be recovered using forward error correction (FEC).




Figure 6 - A received OFDM Signal in the Frequency Domain

Figure 7 is a diagram of time vs. frequency showing common impairments. This diagram helps with the understanding of the effects of different impairments on different modulation types.




Figure 7 - Time and Frequency Relationships Between Common Impairments

Random thermal, or Gaussian, noise is present at all frequencies and all times, so there is no modulation technique that has a relative advantage in AWGN (additive white Gaussian noise). The Shannon Hartley Theorem states the maximum data capacity in a channel with AWGN [6]. If the noise is non-white (i.e., its spectrum is non-flat), the maximum capacity can be determined from the “water-pour” method of transmit power distribution. In cable systems, having a non-flat signal-to-noise ratio occurs due to cable loss varying with frequency, and nonlinear distortion products, which are random noise-like if the distortion was created by digital carriers.

Burst noise occurs locally in time, but often has a wide spectrum. Single carrier, with FEC, can be effective against burst noise for correcting corrupted TD (time domain) symbols. However, an OFDM receiver will perform a FFT on the sequence and spread the burst noise contamination to all FD (frequency domain) symbols.

A continuous wave (CW) interferer can be continuous in time but localized in frequency. OFDM with FEC can repair the localized damage to a limited number of subcarriers. However, with single carrier modulation the CW interferer will affect all symbols, turning a constellation point into a donut shape.

Likewise, a deep frequency-selective fade can be overcome by OFDM with FEC, as mentioned previously.

There are other important considerations for selecting a modulation technique for a particular radio frequency (RF) signal path, such as tolerance to frequency offsets and tolerance to phase noise (both of which increase the cost of local oscillators), and peak-to-average power ratio, which makes transmitters consume more power, decreasing battery life. Furthermore, receiver designers have a number of design tricks, or “secret sauces” to mitigate the effects of impairments, such as noise cancelers.

Let's make a signal to transmit from symbols:

$$E = [e_1, e_2, e_3, e_4 \dots e_j] \quad [\text{eq. 2}]$$

Where E is the signal to be transmitted, and e_n are the individual component symbols.

We can formulate a modulation matrix, C, with rows and columns, where the rows are (hopefully) orthogonal to each other. *Each of the three modulation techniques described previously is nothing more than a different set of row functions, also known as orthogonal basis functions.* For single carrier, the modulation matrix is simply an identity matrix, with a single diagonal row of 1s and 0s elsewhere. For DSSS the rows may be from a Walsh matrix, with a single circular shift between rows. For OFDM, the rows may be complex exponentials (sine and cosine waves), where the first row is the first harmonic, and the second row is the second harmonic, etc. This is illustrated in Figure 8.

$$C = \begin{bmatrix} c(1,1) & c(1,2) & \bullet & \bullet & \bullet & c(1,k) \\ c(2,1) & c(2,2) & \bullet & \bullet & \bullet & c(2,k) \\ c(3,1) & \bullet & \bullet & \bullet & \bullet & c(3,k) \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ c(j,1) & c(j,2) & \bullet & \bullet & \bullet & c(j,k) \end{bmatrix} \quad [\text{eq. 3}]$$

The principle of orthogonality between rows is restated as:

$$\sum_{n=1}^{n=k} c(x, n) \cdot c(y, n) = 0 \quad [\text{eq. 4}]$$

For all rows where $x \neq y$.

So, a signal for transmission, F, is made by simply multiplying the input sequence by the multiplication matrix, C:

$$F = E \cdot C = [f_1, f_2, f_3, f_4 \dots f_j] \quad [\text{eq. 5}]$$




Figure 8 – Eight-Row Matrix as Sines And Cosines, Making OFDM Modulation

In Figure 8, cosine waves are solid and sine waves are dashed. X[1] forms an upper sideband and X[7] forms a matching lower sideband.

2. Rotation of Time and Frequency Axes

Delving deeper into the time and frequency relationships between symbols, single carrier modulation and multicarrier OFDM may be viewed as fundamentally a same modulation technique, apart from a 90-degree rotation in the time-frequency plot.

There is a duality between time and frequency that can be observed in discrete Fourier transform (DFT) and inverse discrete Fourier transform (IDFT) equations:

$$f[k] = \frac{1}{N} \sum_{n=0}^{N-1} F[n] e^{+j \frac{2\pi}{N} nk}$$

[eq. 6]

$$F[n] = \sum_{k=0}^{N-1} f[k] e^{-j \frac{2\pi}{N} nk}$$

[eq. 7]

The differences between equations are only scale factor and a negative sign in front of the complex exponential. The equations are almost the same. If you were shown a set of transform pairs, you could not identify which plot was time and which was frequency.

Figure 9 is a time-frequency plot. With a single carrier signal, such as a pulse amplitude modulation (PAM) signal, each symbol is very short in time, wide in bandwidth, and a next symbol occurs sequentially in time. An OFDM subcarrier is narrow in frequency and long in duration. Many OFDM subcarriers operate simultaneously in time. In Figure 9, there are 32 time symbols or 32 frequency symbols. By rotating a PAM transmission 90 degrees, you have an OFDM transmission, and vice-versa.




Figure 9 - 32x32 Block Time-Frequency Plot

Likewise, by rotating many TDMA single carrier sequential transmissions from several transmitters you have several OFDMA simultaneous transmissions from several transmitters. [2]

While OFDM (without TD tapering) has out-of-band energy splatter, this characteristic is analogous to PAM signals having a $\text{sine}(x)/x$ response in time if the channel rolloff factor (α) is small or zero. So OFDM, using TD tapering, does the analogous operation as PAM using a rolloff factor, α . [3]

So a symbol could be observed to be 32x1 or 1x32. Note that dispersion, or other linear distortion, occurs along the time axis, but not the frequency axis. Dispersion along the frequency axis would be indicative of non-linear distortion.

Performing a 90 degree rotation on a sequence is nothing more than performing a FFT on a sequence, and a -90 degree rotation can be done with an inverse fast Fourier transform (IFFT) on the sequence.

This rotational view of modulation techniques reveals some interesting cases if studied. For example, conventional duobinary (partial response) transmission is well-known. [4] If it is rotated 90 degrees on a time-frequency axis, you can obtain a new transmission method: frequency domain (FD) duobinary, or duobinary OFDM. This is simulated in Figure 10, which is a screen shot from a digital oscilloscope with

an internal FFT. The TD plot (on top) has an envelope shaped like a half-cosine, and the FD plot (created by the oscilloscope) is flat, just a time-frequency dual to conventional duobinary.




Figure 10 - FD Duobinary Block Transmission

If all symbols have a same magnitude, it has a desirable property of abrupt drop of energy out-of-band, naturally reducing interference with neighboring channels. Also, the burst's energy rises and falls gradually in time, causing less interblock interference if an uncorrected short echo is on the channel. On the negative side, the noise performance is poorer, due to being fundamentally duobinary, but that deficiency is ameliorated somewhat by getting more symbols per second within a given bandwidth. It also has a non-flat power vs. time which makes it less desirable for use with power-limited transmitters, such as cell phones. See Appendix A – Detail on Duobinary Modulation.

The desirable characteristic of a gentle rise and fall in transmit power level is a result of duobinary modulation summing each subcarrier with the one next to it, provided it has a same magnitude. This desirable characteristic would be reduced if subcarriers with random or non-equal magnitudes were summed.

One possible application for duobinary OFDM is for very narrow bandwidth OFDM transmissions with a small number of subcarriers, such as ham radio, where the spectral splatter could cause adjacent channel interference. Another use is signaling with a small number of bits in a narrow bandwidth, such as “acks” or acknowledgements.

Conclusion

For three common modulation methods, using orthogonal basis functions can be considered to be the same process as matrix multiplication, where the rows are orthogonal to each other. Likewise, by relabeling the time and frequency axes, single carrier modulation and OFDM multicarrier modulation are comparable.

Abbreviations

ATDMA	advanced time division multiple access
AWGN	additive white Gaussian noise
BPSK	binary phase shift keying
CDMA	code division multiple access
CP	cyclic prefix (see also GI)
CW	continuous wave
dB	decibel
DFT	discrete Fourier transform
DOCSIS	Data-Over-Cable Service Interface Specifications
DSSS	direct sequence spread spectrum
FD	frequency domain
FEC	forward error correction
FFT	fast Fourier transform
GI	guard interval (see also CP)
I	in-phase
IDFT	inverse discrete Fourier transform
IFFT	inverse fast Fourier transform
PAM	pulse amplitude modulation
PN	pseudo-noise
PRS	partial response signaling
Q	quadrature
QAM	quadrature amplitude modulation
OFDM	orthogonal frequency division multiplex
OFDMA	orthogonal frequency division multiple access
PRS	partial response signaling
QPSK	quadrature phase shift keying
RF	radio frequency
SC	single carrier
S-CDMA	synchronous code division multiple access
SCTE	Society of Cable Telecommunications Engineers
TD	time domain

Bibliography & References

- [1] DOCSIS 2.0 Radio Frequency Interface Specification, CM-SP-RFIV2.0-C02-090422, page 71
- [2] US patent 5,815,488
- [3] Digital Telephony, 3rd Edition, John Wiley and Sons, by John Bellamy, Appendix C.
- [4] Digital Telephony, 3rd Edition, John Wiley and Sons, by John Bellamy, Appendix C, pp. 185-188, and pp. 310-311.

- [5] Digital Telephony, 3rd Edition, John Wiley and Sons, by John Bellamy, Appendix C, pp. 587-592
- [6] Communications in the Presence of Noise, by Claude Shannon, reprinted Proceedings of the IEEE, VOL 26, NO. 2, 1998

Appendix A – Detail on Duobinary Modulation

A basic modulation technique, such as BPSK (binary phase shift keying) can be created by connecting a periodic series of positive or negative impulses to a lowpass filter having a $\text{sinc}(x)/x$ impulse response, as illustrated in the top left side of Figure 11 (and in Figure 1). This produces a raised cosine frequency response on the modulated signal on the top right side of Figure 11. The abruptness of the FD roll-off is a factor commonly called “alpha,” and depends on damping applied to the $\text{sinc}(x)/x$ waveform. Duobinary modulation is simply a different impulse response. The impulse response lasts over two symbol periods, not one, as illustrated on the lower left side of Figure 11. In the frequency domain the response is cosine, not raised cosine, and is illustrated in the lower right side of Figure 11.

Passing a two-level complex (I and Q) signal through a duobinary filter produces a 9-PRS (partial response signaling) signal as illustrated in Figure 12. To compare the power of the two signals, the quadrature phase shift keying (QPSK) signal has four equally probable states. The 9-PRS signal has a single state in the middle with a probability of 0.25, four high power corner states with a combined probability of 0.25, and four intermediate power levels for a combined probability of 0.5.

If the voltage difference between A and B is assumed to be 1.0, the power of the 9-PRS constellation is $.25*0 + .5*1.0 + .25*1.414^2 = 1$ watt. If the voltage difference on the QPSK constellation is set to be 0.707 between points C and D, the QPSK power is also 1 watt. So, a noise vector required to make a slicing error on the QPSK signal is .707 volt, and 0.5 volt on the 9-PRS signal, a difference of 3dB.

Looking at the spectrum in Figure 10, observe that three power levels are visible, the peak subcarrier level, an intermediate subcarrier level, and a zero-power subcarrier level, where the energy drops to the origin.

For conventional TD duobinary, the occupied bandwidth of a QPSK signal, relative to 9-PRS, is greater by the channel roll-off factor, commonly called alpha. [6] For DOCSIS single carrier modulations, the value is around 5% more for the downstream and 25% more for the upstream. For FD duobinary, if the number of subcarriers is the same, with equal subcarrier spacing, the occupied bandwidth will be the same. The bandwidth advantage for FD duobinary is a more abrupt drop of energy out of band, allowing closer carrier spacing.

Other candidate modulations for duobinary OFDM are discussed in ref. [4].




Figure 11 - Impulse and Spectral Responses Comparison of Single Carrier (QPSK) and Duobinary (9-PRS)




Figure 12 - Relative Power Calculations for Error Thresholds

Figure Note: Both signals have the same RF power. On duobinary, not all states are equally probable.

Footnotes:

1. For complex numbers, this would be $\text{Sum } x \cdot \text{conj}(y)$
2. It is also possible to use non-orthogonal signals for communications. For example, a non-orthogonal spread-spectrum signal can work in the presence of other signals by taking advantage of spreading gain. In this case, interference is non-zero, but hopefully tolerable.
3. It is also useful to view equalization as another matrix multiply.
4. If the sine-shaped TD duobinary signal illustrated as the top trace in Figure 9 had a cyclic prefix inserted, the waveform would have a fish outline. That is, the shape would gain a tail.

Best Practices for DOCSIS 3.1 Phase Noise Design in the Remote PHY Node

A Technical Paper Prepared for SCTE•ISBE by

Zhuo Zhao

HW Design Engineer
Cisco System
16F, Building C, Yishan Road, Shanghai, China
86 21 24222117
zhuzhao@cisco.com

Jiayou Meng

Technical Leader
Cisco System
16F, Building C, Yishan Road, Shanghai, China
86 21 24057572
jiameng@cisco.com

Haibin Tang

HW Design Engineer
Cisco System
16F, Building C, Yishan Road, Shanghai, China
86 21 24222110
haitang@cisco.com

Introduction

This paper presents a best practice for DOCSIS 3.1 phase noise design related to lower cost and higher phase noise performance in remote PHY node/shelf products. The following figure shows a remote PHY device (RPD) block diagram, which is applicable to a remote PHY node or remote PHY shelf.

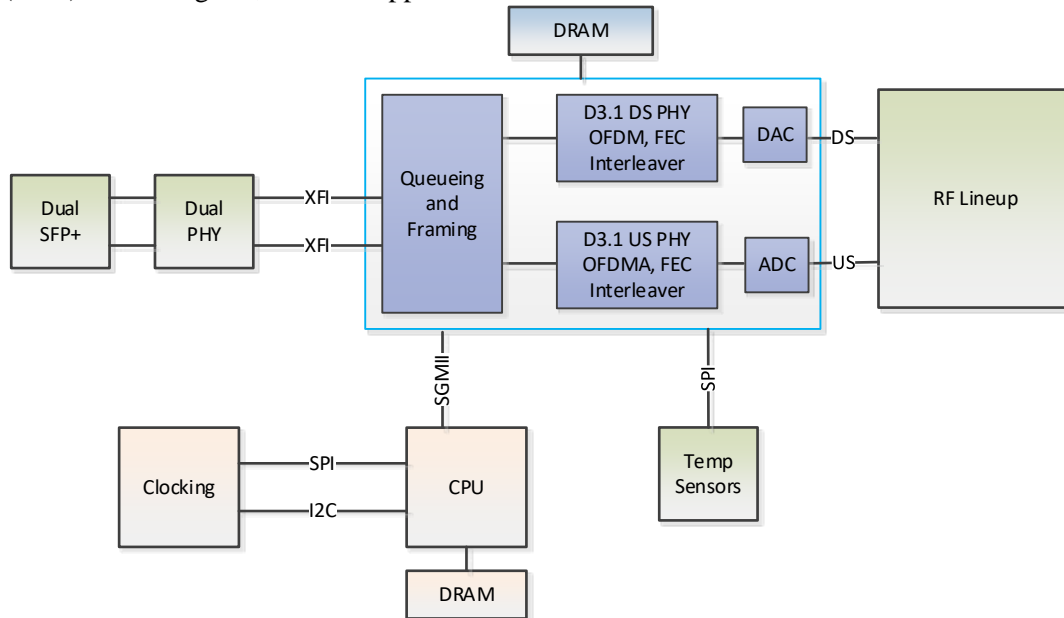


Figure 1 - Remote PHY Device Block Diagram

Clocking

There are two main clock domains in the RPD:

- **DOCSIS system clock domain:** This is the main data path clock. When the RPD is working in synced clock mode, this clock needs to be synchronized with the CCAP core data clock by R-DTI/IEEE 1588. When the RPD is working in timing re-stamping mode, this clock domain would base on the RPD local TCXO clock source. The DOCSIS timing is based on this clock domain. All of the 10 gigabit Ethernet (GE) interface reference clocks, PHY3.0/3.1 DOCSIS clock (204.8 MHz) are in this clock domain and clock margin is not supported.
- **Local system clock domain:** local CPU/CPLD/CPU_DDR/FPGA_DDR/PCIe/GE/SGMII clocks need to sync with the DOCSIS clock; those clocks would come from one local 25 MHz TCXO.

Figure 2 shows the RPD clock block diagram.

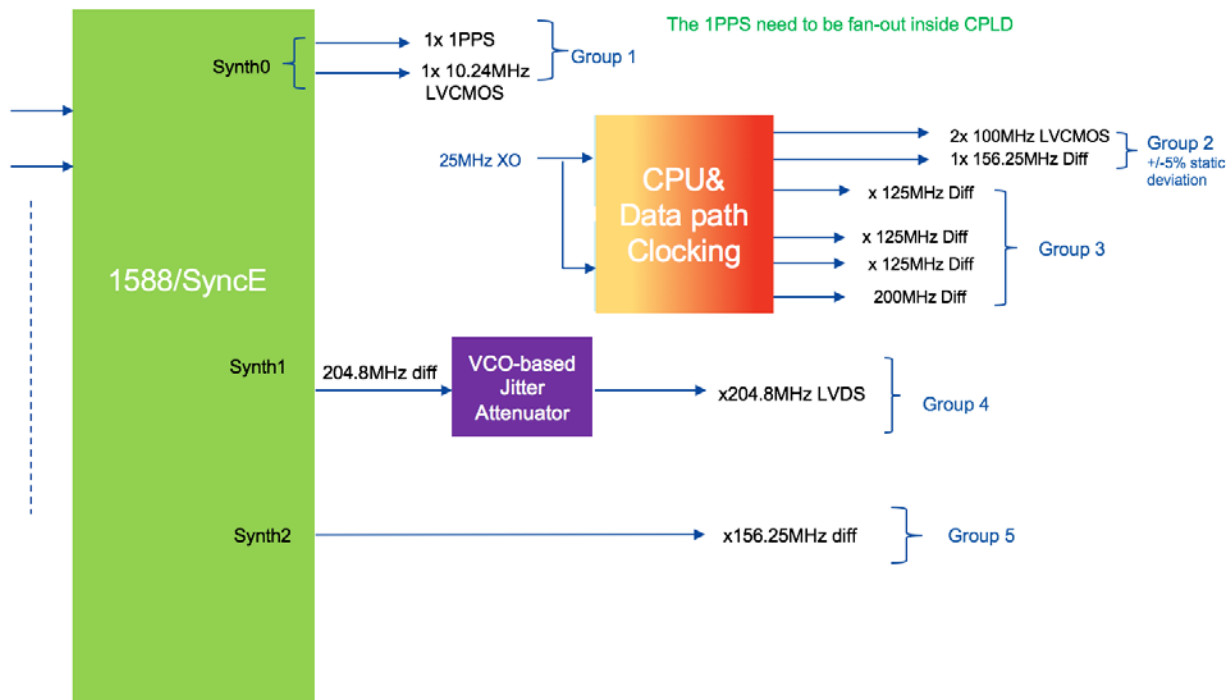


Figure 2 - Clocking block diagram

- PHY3.1 Phase Noise Requirement**

A clean clock is a necessity when working with high performance D/A or A/D-converters in the remote PHY node/shelf product. For DOCSIS applications, the clock noise requirement is driven by DRFI phase noise requirements as well as the noise and spurious requirements as stated in Table 1.

Table 1 - PHY3.1 phase noise requirement

Phase noise, double sided maximum, Full power CW signal 1002 MHz or lower	1 kHz - 10 kHz: -48 dBc 10 kHz - 100 kHz: -56 dBc 100 kHz - 1 MHz: -60 dBc 1 MHz - 10 MHz: -54 dBc 10 MHz - 100 MHz: -60 dBc
Full power 192 MHz OFDM channel block with 6 MHz in center as Internal Exclusion subband + 0 dBc CW in center, with block not extending beyond 1002 MHz [CW not processed via FFT]	1 kHz - 10 kHz: -48 dBc 10 kHz - 100 kHz: -56 dBc
Full power 192 MHz OFDM channel block with 24 MHz in center as Internal Exclusion subband + 0 dBc CW in center, with block not extending beyond 1002 MHz [CW not processed via FFT]	100 kHz - 1 MHz: -60 dBc
Full power 192 MHz OFDM channel block with 30 MHz in center as Internal Exclusion subband + 7 dBc CW in center, with block not extending beyond 1002 MHz [CW not processed via FFT]	1 MHz - 10 MHz: -53 dBc
Adjacent channel spurious signals and noise	Fc+/-3.75-9 MHz -62 dBc
All other channels spurious signals and noise	47 MHz to 1218 MHz -73 dBc

The "all other channels" specification in Table 1 sets the limit on the wideband noise floor. Hence for wideband noise, the requirement at the DAC output is -73 dBc within a 6 MHz bandwidth, which translates to -141 dBc/Hz. When referenced to the total signal power, this number remains constant as the number of channels is increased.

Assume that the DAC is operated at an update rate of 4 gigasamples per second (GSps); a SNR over Nyquist of 48 dB is required.

• Wideband Jitter Requirement

It is well known that the signal to noise ratio of a DAC converting a single tone is

$$SNR[dB] = 20 \cdot \log\left(\frac{1}{2\pi f_c \sigma_f}\right)$$

where f_c is the signal frequency and σ_f is the RMS jitter.

For a finite bandwidth, random, stationary signal with bandwidth f_b , it can be shown that the SNR out of the DAC, due to clock noise will be equal to

$$SNR[dB] = 20 \cdot \log \left(\frac{\sqrt{12}}{2\pi\sigma_\tau \cdot \sqrt{f_b^2 + 12f_c^2}} \right)$$

Using the previous, and assuming that the jitter is caused by white noise and is translated to white noise in the Nyquist bandwidth, noise density can be expressed as

$$N[dBc/Hz] = -20 \cdot \log \left(\frac{\sqrt{12}}{2\pi\sigma_\tau \cdot \sqrt{f_b^2 + 12f_c^2}} \right) - 10 \cdot \log(f_{DAC}/2)$$

Given that the SNR applies to the Nyquist bandwidth, it follows that noise density will decrease by 3 dB as the clock rate is doubled, assuming that the amount of jitter remains constant.

The maximum allowable jitter to achieve a given noise density is calculated as:

$$\sigma_t = \sqrt{10^{N/10} \cdot f_{DAC}/2 \cdot \frac{3}{\pi^2 \cdot (f_b^2 + 12f_c^2)}}$$

Assume a signal frequency of $f_c = 1$ GHz, an update rate $f_{DAC} = 4$ GSps. The maximum noise density for DRFI is -141 dBc. Assume 10 dB margin to this specification to account for sinc attenuation (approx. 1 dB for MAX5882), additional noise sources and margin to the specification, you find the maximum jitter $\sigma_t = 0.2$ ps.

• DOCSIS PHY3.1 Chip Phase Noise Requirement

Table 2 shows an example from silicon PHY3.1 chip vendor for phase noise requirement.

Table 2 - DOCSIS3.1 PHY chip phase noise requirement

Reference clock Phase noise (SSB)	@ 1 kHz	-116 dBc/Hz
	@ 10 kHz	-134 dBc/Hz
	@ 100 kHz	-143 dBc/Hz
	@ 1 MHz	-144 dBc/Hz
	>10 MHz	-150 dBc/Hz
Ref. Clock jitter	Integrated, 1 kHz to 20 MHz	225 fs, RMS
	Integrated, 100 Hz to 20 MHz	275 fs, RMS

100 kHz spot phase noise specification derived from the downstream worst-case frequency band [100 kHz to 1 MHz] phase noise specification = -60 dBc.

There are multiple clock solution designs for RPDs.

Table 3 - RF Clock Design Options

Key chip	Key Design	Solution
LMK04828	Jitter Cleaner Single PLL mode with external VCXO 204.8 MHz clock output	Solution 1: VCXO-Based Jitter Cleaner Dual PLL Design
LMK04616	20.48 MHz input, Dual PLL mode (1st stage 122.88 MHz VCXO + 2nd stage internal VCO) 204.8 MHz clock output	
HMC7044	20.48 MHz input, Dual PLL mode (1st stage 122.88 MHz VCXO + 2nd stage internal VCO) 204.8 MHz clock output	
ADF4002	Frequency synthesizer with external VCXO 204.8 MHz clock output	Solution 2: Frequency synthesizer + VCXO

Solution 1: VCXO-Based Jitter Cleaner Design

PHY3.1 silicon vendor has strict requirement on the 204.8 MHz clock for ADC/DAC and 156.25 MHz clock for 10 GE interface.

The IEEE 1588 clock recovery from DPLL has high phase noise and high spurs, so a jitter cleaner device is needed for “cleaning” the 1588 clock.

• Jitter Cleaner

Traditionally, the RPD will use a jitter cleaned chip with an external 204.8 MHz VCXO.

There would be two PLL loops, PLL1 uses an external VCXO and loop bandwidth to provide low jitter clean clock, and PLL2 would be used for frequency generation. See Figure 3.

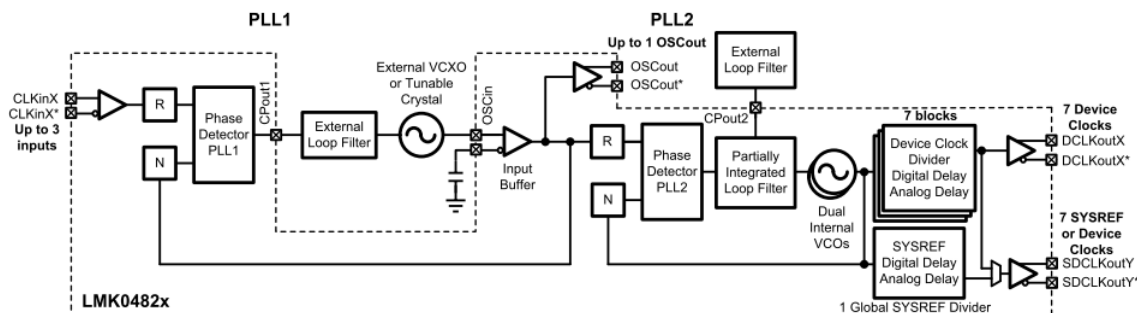


Figure 3 - VCXO-Based Jitter Cleaner Block Diagram

The dual loop PLL architecture of the LMK04828 provides the lowest jitter performance over a wide range of output frequencies and phase noise integration bandwidths. The first stage PLL (PLL1) is driven

by an external reference clock and uses an external VCXO or tunable crystal to provide a frequency accurate, low phase noise reference clock for the second stage frequency multiplication PLL (PLL2).

PLL1 uses a narrow loop bandwidth (typically 10 Hz to 200 Hz) to retain the frequency accuracy of the reference clock input signal while at the same time suppressing the higher offset frequency phase noise that the reference clock may have accumulated along its path or from other circuits. This “cleaned” reference clock provides the reference input to PLL2.

Ultra-low jitter is achieved by allowing the external VCXO or crystal’s phase noise to dominate the final output phase noise at low offset frequencies and the internal VCO’s phase noise to dominate the final output phase noise at high offset frequencies.

On the RPD, PLL2 is not used and is bypassed. Only PLL1 is used to generate the 204.8 MHz clock for the PHY3.1 silicon chip

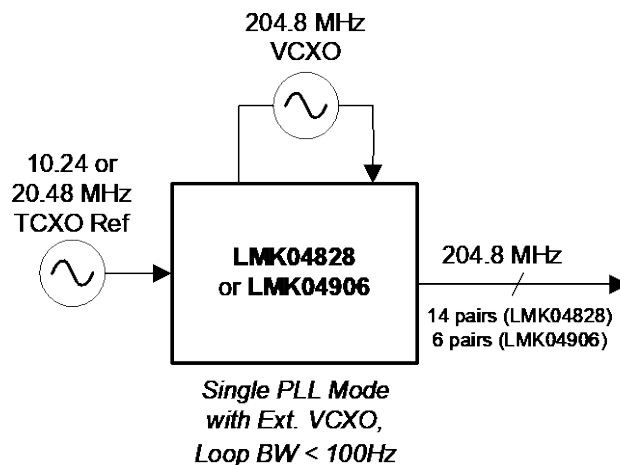


Figure 4 - Set-Up in Lab

- **PLL Simulation**

A single PLL configuration with TCXO & VCXO noise models is shown in Figure 5.

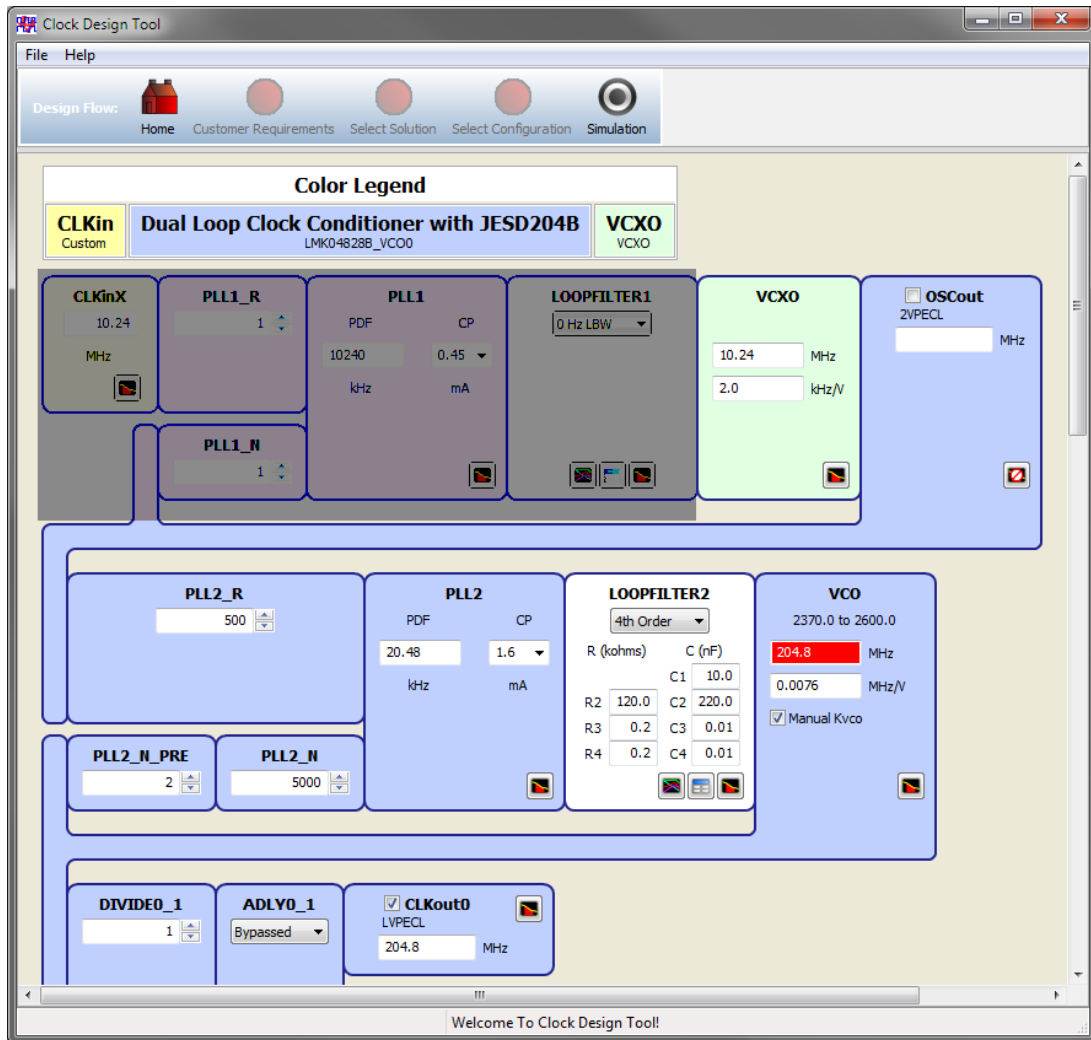


Figure 5 - Single PLL Configuration with TCXO & VCXO Noise Model

Temperature compensated crystal oscillator specification is shown in Figure 6.
We use Vectron's VT-820 as a simulation model (Figure 7).

Parameter	Symbol	Min.	Typ	Max	Units
Output Frequency	f_o	8		45	MHz
Supply Voltage, ¹ (Ordering Option)	V_{DD}	+1.8, 2.8, +3.0 or +3.3			V
Supply Current, 8 to 19.999MHz 20.000 to 31.999MHz 32.000 to 45.000MHz	I_{DD}			1.5 2.0 2.5	mA
Operating Temperature, (Ordering Option)	T_{OP}	0/55, -10/60, -20/70, -30/80, -30/85, -40/85			°C
Stability Over T_{OP} , (Ordering Option)		$\pm 0.5, \pm 1.0, \pm 1.5, \pm 2.0, \pm 2.5, \pm 3.0, \pm 3.5 \pm 4.0, \pm 5.0$			ppm
Initial Accuracy ² , "No Adjust" Option				± 1.0	ppm
Power Supply Stability, $\pm 5\%$ change				± 0.2	ppm
Load Stability				± 0.2	ppm
Aging				± 1.0	ppm/yr
Pull Range, (Ordering Option)	TPR	$\pm 5, \pm 8, \pm 10, \pm 12, \pm 15$			ppm
Control Voltage to reach Pull Range		0.5		2.5	V
1.8V option		0.3		1.5	V
Control Voltage Impedance		500			Kohm
Output Level ³	V_o p/p	0.8			V
Output Load				10K 10pF	
Phase Noise, 10.000MHz					dBc/Hz
10Hz			-91		
100Hz			-116		
1kHz			-137		
10kHz			-149		
100kHz			-150		
Start Up Time				2	ms

Figure 6 - Vectron TCXO Datasheet

Load Custom Phase Noise

Enter phase noise data for block: VCXO

	Offset (kHz)	Phase Noise (dBc/Hz)
1	0.01	-91.0
2	0.1	-116.0
3	1.0	-137.0
4	10.0	-149.0
5	100.0	-150.0
6	1000.0	-150.0

Frequency of block: 10.24 MHz

Buttons: Clear/Reset Noise, <-- Set Noise, Load Noise from File, Close

Figure 7 - TCXO Noise Model

Low noise and low jitter VCXO data are shown in Figure 8 and Figure 9. We use the RAKON: RVX7050M 204.800 MHz as simulation.

7.0 SSB Phase Noise

Parameter	Typ.	Max.	Unit	Test Condition / Description
a. 10Hz offset	-67		dBc/Hz	25°C
b. 100Hz offset	-97		dBc/Hz	25°C
c. 1kHz offset	-125		dBc/Hz	25°C
d. 10kHz offset	-145		dBc/Hz	25°C
e. 100kHz offset	-155		dBc/Hz	25°C
f. 1MHz offset	-155		dBc/Hz	25°C
g. 10MHz offset	-156		dBc/Hz	25°C

Figure 8 - RAKON VCXO Noise Model

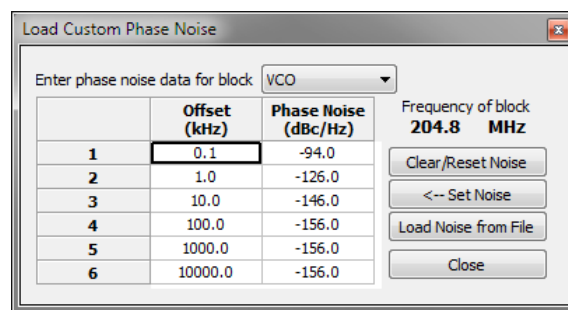


Figure 9 - VCXO Noise Model

PLL loop filter characteristics are shown in Figure 10.




Figure 9 - PLL Loop Filter Characteristics

LVDS output phase noise / jitter are shown in Figure 11.

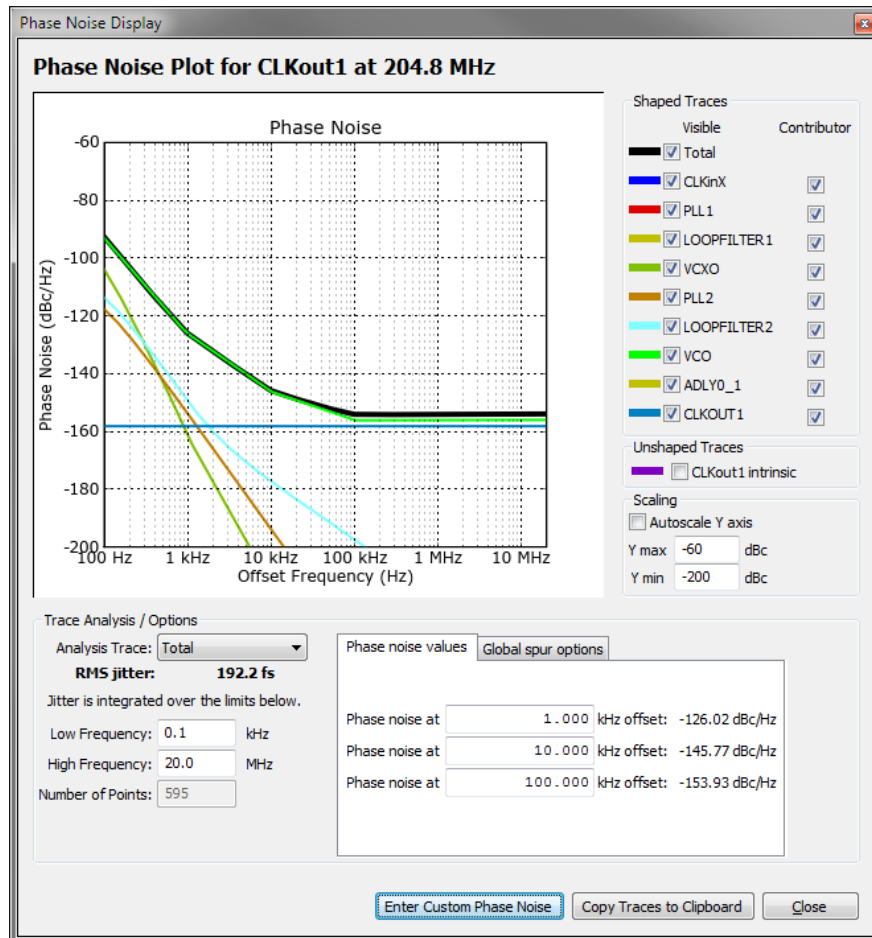



Figure 10 – PLL Phase Noise Simulation


- **Test Result**

Please see 204.8 MHz Output Phase Noise test results in Figure 12 and Figure 13, showing good margin over phase noise and integration jitter specification.



	Reference clock Phase noise (SSB)	@1kHz	-116dBc/Hz
		@10kHz	-134dBc/Hz
		@100kHz	-143dBc/Hz
		@1MHz	-144dBc/Hz
		>10MHz	-150dBc/Hz
	Ref. Clock Jitter	Integrated, 1kHz-20MHz	225fsec, rms
Measurement		Integrated, 100Hz-20MHz	275fsec, rms
	R-PHY specification		

Figure 11 - PLL Phase Noise/Jitter (1 kHz to 20 MHz) test result



	Reference clock Phase noise (SSB)	@1kHz	-116dBc/Hz
		@10kHz	-134dBc/Hz
		@100kHz	-143dBc/Hz
		@1MHz	-144dBc/Hz
		>10MHz	-150dBc/Hz
	Ref. Clock jitter	Integrated, 1kHz-20MHz	225fsec, rms
		Integrated, 100Hz-20MHz	275fsec, rms
Measurement		R-PHY-specification	

Figure 12 - PLL Phase Noise/Jitter (100 Hz to 20 MHz) test result

Solution 2: Frequency Synthesizer + VCXO

The ADF4002 frequency synthesizer is used to implement local oscillators in the up-conversion and down-conversion sections of wireless receivers and transmitters. It consists of a low-noise digital phase

frequency detector (PFD), a precision charge pump, a programmable reference divider and programmable N divider. The 14-bit reference counter (R counter), allows selectable REFIN frequencies at the PFD input.

• Block Diagram

A complete PLL can be implemented if the synthesizer is used with an external loop filter and voltage controlled oscillator (VCO). In addition, by programming R and N to 1, the device can be used as a standalone PFD and charge pump. See Figures 14 through 16 for block diagram, test configuration, and test results.

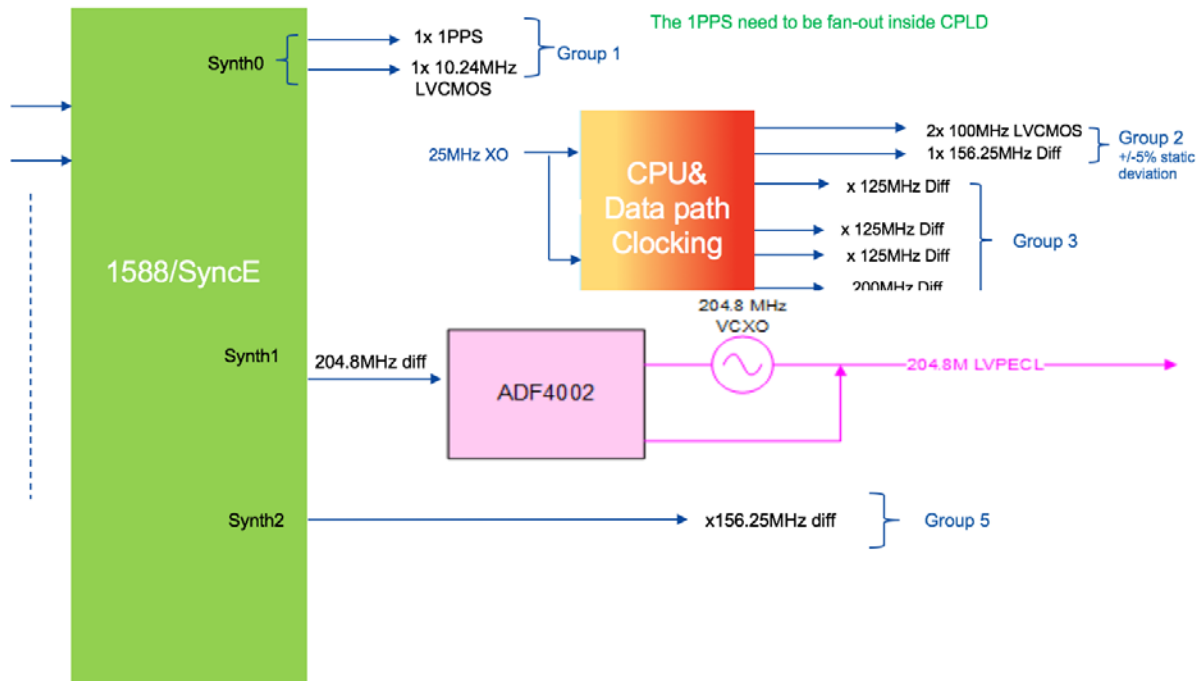


Figure 13 – ADF4002 with External VCXO

- **PLL Configuration**

Select Device and Connection

Main Controls

Registers

Sweep and Hop

Other Functions

Features

RF Settings

Reference Frequency: 204.8 MHz
☐ Automatic ☒ Manual

RF VCO Output Freq.: 204.8 MHz
PFD Frequency: 2048 kHz
Channel spacing: 2048 kHz

R: 100
N: 100

B

P

A

12

8

4

x

+

x

PFD (kHz)

2048

=

RFout (MHz)

204.8

N = 100

Settings

Charge Pump Setting 1: 5.0 mA
Charge Pump Setting 2: 5.0 mA
Charge Pump Gain: 0
Charge Pump Tri-State: Disabled
FastLock: Disabled
Timeout: 3 PFD Cycles
Phase Detector Polarity: Positive

Counter Reset: Disabled
Lock Detect Precision: 3 cycles
Power Down: Normal Operation
ABPW: 2.9 ns
Muxout: Analog Lock Detect

Device in use:

ADF4002

Software version:

7.7.4

Latches/Registers

0x 190

Write R Counter Latch

0x 6401

Write N Counter Latch

0x 1F80D2

Write Function Latch

Write All Latches (Function > R > N)

Write Initialization Latch

Figure 14 - ADF4002 Register Configuration

- **PHY3.1 Phase Noise@RPD Port Test Result**




Figure 15 – PHY3.1 phase noise test @RPD RF port

Conclusion

This paper has highlighted best practices for DOCSIS 3.1 phase noise design in the remote PHY device with two clock solutions. It also proved that we can get good margin for not only phase noise but also for jitter performance considering the tough requirements from both PHY3.1 silicon chip and CableLabs PHY3.1 specification.

Abbreviations

A/D	analog-to-digital
ADC	analog-to-digital converter
CCAP	converged cable access platform
CW	continuous wave
D/A	digital-to-analog
DAC	digital-to-analog converter
dB	decibel
dBc	decibel carrier
DOCSIS	Data-Over-Cable Service Interface Specifications
DPLL	digital phase-locked loop
DRFI	[DOCSIS] Downstream Radio Frequency Interface [Specification]
FFT	fast Fourier transform
fs	femtosecond
GE	gigabit Ethernet
GHz	gigahertz
GSps	gigasamples per second
Hz	hertz
HW	hardware
IEEE	Institute of Electrical and Electronics Engineers
ISBE	International Society of Broadband Experts
kHz	kilohertz
LVDS	low voltage differential signaling
MHz	megahertz
OFDM	orthogonal frequency division multiplex
PFD	phase frequency detector
PHY	physical layer
PLL	phase-locked loop
ps	picosecond
RF	radio frequency
RPD	remote PHY device
RMS	root mean square
SCTE	Society of Cable Telecommunications Engineers
SNR	signal-to-noise ratio
SSB	single sideband
TCXO	temperature controlled crystal oscillator
VCO	voltage controlled oscillator
VCXO	voltage controlled crystal oscillator

Bibliography & References

DOCSIS 3.1 Physical Layer Specification CM-SP-PHYv3.1-I11-170510 TI 204.8 MHz test result
Cisco RPD HW design specification

Unified Architectures for Remote PHY Backhaul and 5G Wireless Fronthaul

A Technical Paper prepared for SCTE•ISBE by

Yuxin (Eugene) Dai, PhD
Principle transport Architect
Cox Communications
6305 Peachtree Dunwoody, Atlanta, GA 30328
404-269-8014
Eugene.dai@cox.com

1. Introduction

Remote PHY (RPHY) deployment started in 2017 in some MSO networks. RPHY works well in deep fiber architecture; it puts RF modulation devices deep in the field, attached to the N+0 coax outside plant, while keeping MAC and higher layer devices in the center of the network, e.g., in the Headend or data center. One of the main advantages of RPHY architecture is the efficiency. By moving RF modulators closer to the customers, higher order modulations such as 2K and 4K QAM can be used, therefore increasing the network efficiency. However, there are challenges for RPHY architecture. Besides the technical challenge of synchronizing RPHY and MAC, it also poses challenges on the backhaul network for RPHY. The communication protocol between RPHY and MAC is Ethernet at a rate of multiple Gigabit/s. To backhaul large amounts of high-speed Ethernet traffic in MSO's access networks is a challenge.

Another deployment in the communication industry is coming 5G wireless network services. Although most of the MSOs do not directly provide wireless services (not including Wi-Fi) to the end-customers today, many of them are contracted by wireless carriers to provide wireless backhaul and/or fronthaul services in their networks. The amount of data to backhaul 5G wireless is expected to increase greatly and pose a challenge to MSO's metro and access networks.

RPHY and 5G backhauls are new to MSO; it poses big challenges to MSO's access network. To design a unified and converged access network to backhaul RPHY and 5G traffic and at the same time align with the MSO long term migration direction to passive all fiber access networks is an even bigger challenge.

In this paper, we first analyze MSO's current access network architectures for DOCSIS and PON services, then discuss the options for RPHY and 5G wireless backhauls. We then propose two unified fiber access network architectures. One is based on DWDM and the other is based on PON, to backhaul RPHY and 5G wireless traffic.

2. MSO's Access Networks Today

Generally speaking, MSO's access networks today consist of HFC (hybrid fiber coax), active optical Ethernet Network (AON) and PON (passive optical networks) as shown in Figure 1.

Legacy HFC, although under constant upgrade, is mainly providing DOCSIS protocol (D3.0, D3.1 etc.) based cable data, voice and video services. Point-to-Point (P2P) active optical Ethernet is occasionally used for business customers. PON, either GPON (2.5G), XGS-PON (10G) or 10G EPON are used in recent years to provide high-speed services, for example Gigabit to both residential and business customers. Among the above-mentioned access networks, HFC is mostly deployed. The deployment of PON, which is considered as the migration direction, is increasing. AON, active optical Ethernet, has limited deployment in MSO networks.

From OSP (outside plant) point of view, HFC has P2P fiber from a cable headend to an optical node in the field; and P2MP (point-to-multiple-point) coax cables from a node to the end-users as shown in Figure 1.




Figure 1 - MSO access architectures today

AON has P2P fiber from a headend to an end-user. A PON ODN (optical distribution network) has a P2MP fiber topology from a headend to the end-users.

3. Remote PHY and RPHY backhaul

3.1. Remote PHY in a nutshell

RPHY technology brings RF modulation deep in the field closer to the coax OSP. As a result, higher order modulations are possible for RPHY to increase efficiency, for example 4K and 8K QAM (quadrature amplitude modulation) are possible. DOCSIS MAC is physically located at the Headend or in data center servers. Gigabit to 10G Ethernet interfaces are assumed between RPHY and DOCSIS MAC. The RPHY and MAC are synchronized via SyncE, IEEE1588, etc.

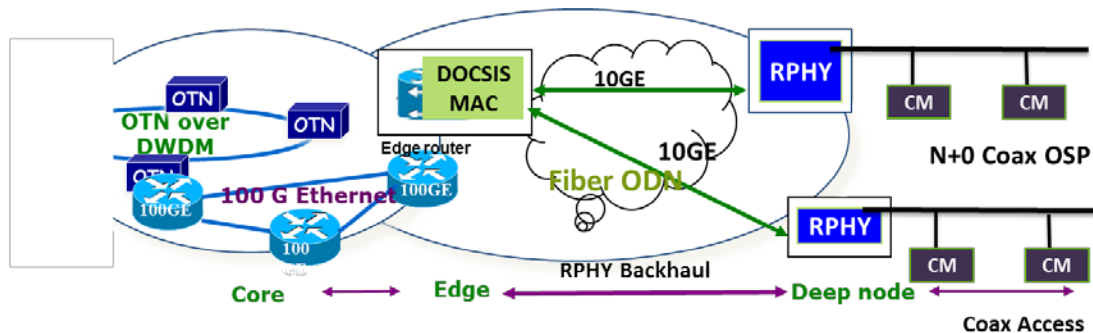


Figure 2 - Remote PHY

3.2. The network segment of RPHY

A RPHY node could be deployed deep in the access network segment, as shown in Figure 3. A HFC optical node normally feeds an “N+x” P2MP (point-to-multi-point) coaxial cable plant, where the “x” in donates the number of RF amplifiers in the coaxial cable plant and “N” represents the HFC optical node.

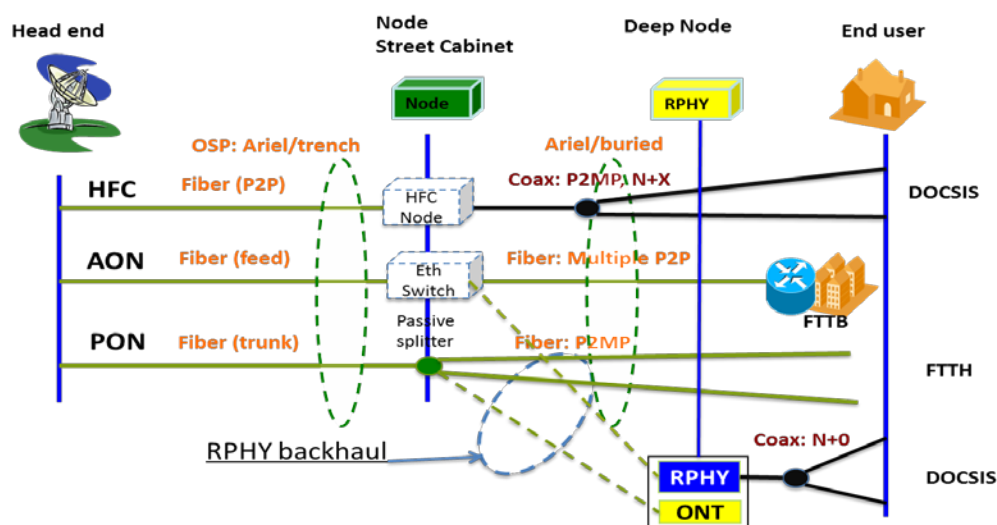


Figure 3 RPHY in MSO network

A RPHY often feeds a passive P2MP coaxial cable plant, which also is referred to as “N+0”. Therefore, a RPHY node is located between a HFC optical node and the cable end-users. In this scenario, a RPHY node is a type of deep optical node.

3.3. RPHY backhaul

Under MSO’s access network today, in reference to Figure 3, AON and PON can be used to backhaul RPHY traffic. In AON backhaul for RPHY, an Ethernet switch in the deep node location is used to aggregate RPHY traffic and uplink to the headend or data center via a P2P optical fiber. The advantage of AON backhaul is that it supports aggregation. The drawback is that there are active devices in the field. Since there are few AON in MSO’s access network today, AON backhaul is not considered in the paper.

PON, such as XGS-PON, 10G EPON or NG-PON2, for RPHY backhaul is feasible. Since PON has been deployed in MSO access networks for many years and it is considered as the migration direction, PON RPHY backhaul provides a converged solution. We’ll discuss PON for RPHY backhaul in section 6.

Another solution for RPHY backhaul is using passive DWDM which requires to the construction of a new DWDM OSP in the access network. It will also be discussed in section 5.

4. C-RAN 5G wireless fronthaul

4.1. 5G wireless fronthaul and backhaul in a nutshell

Several MSOs have been contracted to provide fiber connectivity for 3G and 4G fronthaul for wireless carriers for years. CPRI (Common Public Radio Interface), also called RoF (Radio over Fiber) has been used for 2G, 3G and 4G wireless fronthauls for wireless carriers for many years. The network for wireless backhaul and fronthaul architecture is called C-RAN (Centralized Radio Access Network), as shown in Figure 4.




Figure 4 - C-RAN for wireless backhaul and fronthaul

Fronthaul fiber segments start from wireless BBU (Baseband Unit) to RRH (Remote Radio Head) running CPRI protocol. The fiber distance of a fronthaul segment is normally < 20km. The backhaul fiber segment is from BBU to the packet core network. The backhaul network is the packet data network, the fiber distance varies. The BBU is normally located at the wireless data center.

With the roll out of 5G services in the near future, there will be a huge increase in demand of fiber or DWDM wavelengths for fronthaul CPRI traffic.

4.2. Similarities of RPHY backhaul and 5G wireless fronthaul

There are astonishing similarities in network topologies for RPHY backhaul and 5G wireless C-RAN fronthaul. Figure 5 shows the network topologies of RPHY backhaul and 5G wireless fronthaul fiber networks.

- The fronthaul distance from wireless BBU to RRH is < 20 km. Statistics data shows that about 80% of RPHY are within 20 km distance from a cable headend.
- DOCSIS RPHY nodes and 5G small cell RRH are more closed to the end-users than their corresponding previous generations.
- Both RPHY backhaul and 5G wireless fronthaul networks can be built with P2P fiber, passive DWDM, active DWDM and high-speed PON.
- Both RPHY backhaul and wireless fronthaul need high-speed data transport networks, for example 10Gbps or 25Gbps rates.

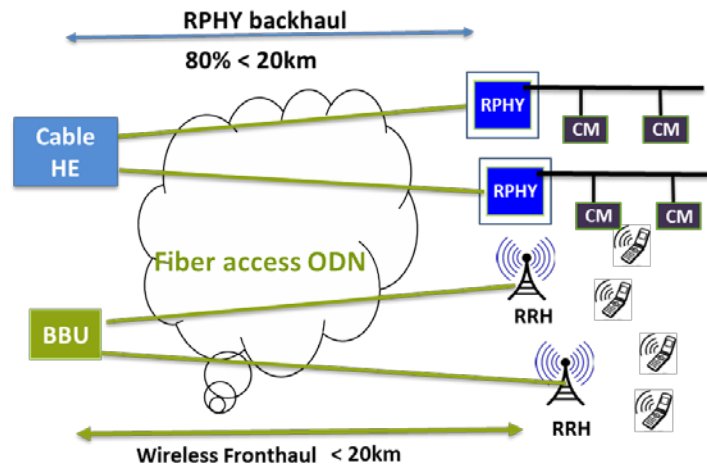


Figure 5 - Fiber networks for RPHY backhaul and 5G wireless fronthaul

Therefore, it is possible to design a unified fiber access network, whether built on DWDM or PON, to support both RPHY backhaul and 5G wireless fronthaul. For MSO, a common passive DWDM architecture for both RPHY and 5G wireless fronthaul benefits from the large volumes in wireless backhaul market.

5. A novel passive DWDM architecture for RPHY backhaul and 5G wireless fronthaul

5.1. High level requirement for common DWDM architecture

Firstly, the unified DWDM system should have low cost as well as flexibility in channel plan. This requires a passive DWDM architecture with no optical amplifications. The use of an optical amplifier, such as SOA (Semiconductor Optical Amplifier) or EDFA (Erbium-doped Fiver Amplifier), will not only increase the cost dramatically but will also impose an inflexible channel plan. “Red-blue” filters are needed for bidirectional EDFA which excludes some channel plans such as interleave. It should support at least 20 km reach without optical amplification to meet the 5G wireless fronthaul reach and 80% of the RPHY reach.

Secondly, it should provide 1:1 fiber protection on DWDM trunk fiber. For a 32 channel system, the DWDM trunk supports 32 RPHY with 4096 end-users assuming one RPHY has 128 end-users. Therefore, 1:1 fiber protection is important for RPHY backhaul and it is required for 5G wireless fronthaul. Loss in the DWDM trunk will result in loss of many adjacent 5G cells which is not acceptable in wireless communications.

Thirdly, it should align with the migration direction to PON. The common ODN (optical distribution network) for DWDM supports overlay with all PONs. This means the common ODN should enable coexistence with GPON, or XGS-PON, or NG-PON2, or 10G EPON or 25Gbps/100Gbps EPON without any modifications.

Finally, the solution should be standard based.

5.2. Architectural challenges

There are a number of architecture challenges to design a common DWDM system optimized for RPHY backhaul and 5G mobile front-haul. For example, one challenge is the coexistence with all PONs, e.g. GPON, XGS-PON, NG-PON2, 10G EPON, 100G EPON on a common ODN. This requirement stems from the uncertainty of the future technologies, and the fact that no one really has a crystal ball. Another challenge is to keep DWDM system passive with at least 20km fiber reach and compensate the losses from DWDM filters, optical protection switch, coexist filters, etc.

Therefore, innovative architectures and new component technologies are needed.

5.3. A novel DWDM with PON overlay architecture

Figure 6 shows a system diagram the unified DWDM architecture for RPHY backhaul and 5G wireless fronthaul. This system has following characteristics:

1. The DWDM system coexists with all PONs that are standardized today, including GPON, XGS-PON, 10G EPON, and NG-PON2. This is achieved by two new coexistence schemes.
 - a). The coexistence of all PONs with the DWDM system is achieved by using a C band bandpass filter instead of conventional PON coexistence filter at the OLT side, see the left side of Figure 6. The C band bandpass filter passes the C band wavelengths from the common port to one output port that connects to a DWDM MUX and reflects all other wavelengths to another output port that connects to an OLT. All the PONs that expect NG-PON2 use wavelengths out of C band, for example GPON uses 1310nm/1490nm, XGS-PON uses 1270nm/1578nm and 10G EPON uses 1270nm/1577nm for upstream and downstream. Therefore, all PON wavelengths will be separated from the C band wavelengths. As the result, if any one of the above mentioned PON connects to the OLT port in Figure 6, it will coexist with the DWDM system.

NG-PON2 uses part of C band wavelengths for upstream. The 4 TWDM channels in NG-PON2 are located at the lower C band from 1524nm to 1534nm with 200GHz channel spacing. If the DWDM wavelength plan avoids this portion of C band, then the DWDM system can coexist with NG-PON2 as well.

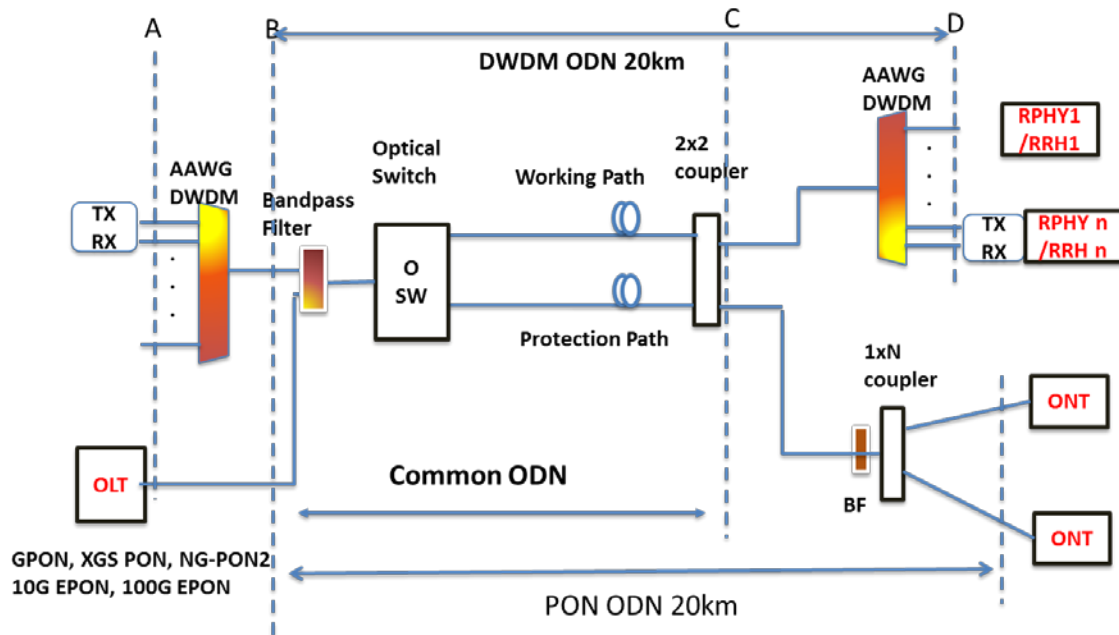


Figure 6 - A novel DWDM with PON overlay architecture

b). In the field, the coexistence of the DWDM system with all PON is realized by using 2x2 passive optical couplers with a C band blocking filter before the PON splitter, see the right side of Figure 6. The blocking filter could be the C bandwidth filter with the C band passing output not used.

2. The system supports 1:1 trunk fiber protection for DWDM and PON. The 2x2 passive optical coupler has two functions. One is to provide trunk fiber protection, and the other is to enable coexistence of DWDM with PON.
3. Gaussian AWG (Arrayed Wavelength Grating) is used for DWDM MUX and DMUX to lower the insertion loss caused by DWDM filters. Take a 40 channel DWDM system, compared with a thin film filter, Gaussian AWG can save at least 6 dB in link budget for a pair of MUX/DMUX.
4. The DWDM system can support 50 channels in C band with 100GHz channel spacing. Since there is no EDFA, no guard band is needed for “red-blue” filter. As the result full C band can be used.
5. Using 10G SFP+ EML transceivers, the passive DWDM system support at least 20km reach with the possibility to extend to 40 km reach with reasonable margins.

As being mentioned in 5.2, one of the challenges to keep the DWDM system passive (no EDFA) with at least 20 km fiber reach is lower passive loss. By using an innovative coexistence architecture, the link budget saving is about 3 dB. Combine the insertion saving from using Gaussian AWG, the total link

budget saving is about 9 dB. This is critical to keep the DWDM system passive, otherwise EDFA would be required.

5.4. New component technologies

To keep the system passive while supporting at least 20 km reach, the low loss Gaussian AWG is assumed in the system architecture shown in figure 7. There are two issues with AWG, one is temperature stability; and the other is that the Gaussian AWG has a narrower passing band. However, these two problems can be solved with a recent development in optical component technologies.

AWG has many advantages over thin film filters, such as

- Constant insertion loss independent of channel counts, therefore has lower loss for large channel counts.
- More uniform loss between channels.
- Cyclic prosperity.

However, there is one main drawback for AWG – temperature instabilities. The wavelength variation for AWG is ~ 11 ppm/degree C. In comparison, the wavelength variation for thin film filters is ~ 1 to 2 ppm/degree C. As the result, an AWG normally requires, or as an option, temperature control by thermoelectric (TE) devices, or by environment control. This limits the applications of AWG in outdoor and/or in extreme environments.

During the first decade of year 2000, research on temperature insensitive Athermal AWG (AAWG) made progress in labs. Like AWG, AAWG is also based on silica technology. It uses a material that has a different thermos expansion coefficient than that of silica to compensate for the reflection index change of silica caused temperature variations. Figure 7 shows an AAWG design [1].




Figure 7 - Athermal AWG

In this design, a copper plate was used to compensate the thermal expansion of silica, the operating temperature range tested was -30 to +70/degrees C. Athermal AWG products are available today.

The second issue of Gaussian AWG/AAWG is the narrower pass band. Figure 8 shows a comparison of Gaussian AAWG with a flat top AAWG and thin film filter. The 1 dB pass band of Gaussian AAWG is ≥ 0.2 dB. In comparison, the 1 dB pass band of a flat top AAWG is ≥ 0.38

dB and the 0.5 dB pass band for the thin film filter is ≥ 0.22 dB. In order to use low loss Gaussian type AAWG, the laser transmitter needs to narrow wavelength variations. General speaking, an EML (electro-absorption modulator laser) has to be used. The EML is one type of external modulated laser, it has more wavelength stability than a direct modulated laser (DML). Take 10G EML SFP+ for example, the wavelength accuracy is

- EOL: ± 0.1 nm;
- BOL: ± 0.04 nm.

BOL and EOL represent the “beginning of life” and “end of life” respectively. The EOL of wavelength drift of the 10G SFP+ is within the 1dB pass band of the AAWG.

Moreover, to mitigate the dispersion penalty, EML is also preferred for ≥ 20 Km reach 10G systems.

Table 1 - Comparison of AAWG and thin film filter

Parameter	Unit	AAWG Type		Thin Film
		Gaussian	Flat top	
Channel Spacing	GHz	100 GHz		100 GHz
Channel Count		40	40	32
Wavelength Accuracy	nm	± 0.05	± 0.05	± 0.05
1 dB pass band	nm	≥ 0.2	≥ 0.38	≥ 0.22 (0.5 dB)
3 dB pass band	nm	≥ 0.4	≥ 0.58	
Insertion loss	dB	≤ 3.5	≤ 6.0	≤ 5.4 (~6 for 40ch)
Uniformity	dB	≤ 1.5	≤ 1.5	≤ 2

Together, the Gaussian AAWG and EML transmitter enable the low insertion loss DWDM system dispatched in Figure 6. The passive DWDM system supports 20 km to 40 km of reach with up to 50 channels in C band. Comparing with optically amplified solution (using EDFA), the passive DWDM RPHY backhaul and wireless fronthaul; system has lower deployment and operational cost.

6. PON backhaul/fronthaul architecture

The advantage of the DWDM system discussed previously is that it is built on mature technologies and most of the optical component needed can be found off-shelf. On the other hand, if RPHY is considered as a transitional solution – deep fiber nodes with final migration path to PON, then using PON for network backhaul and fronthaul may provide a converge solution.

In the past few years, there has been increasing deployment of PON in MSO access networks, from GPON to 10G EPON. Since PON ODN is already deployed or start to deploy in MSO OSP, it would be beneficial to use PON ODN for RPHY backhaul from evolution to all fiber access networks point of view. However, migrate all coax OSP to PON ODN may take time, and need innovated coexistence solutions [2][3].

6.1. NG-PON2 for RPHY backhaul

TDM PON, such as GPON (2.5G) has been used by telecom carriers for 2G, 3G and 4G wireless backhaul and fronthaul for years. In recent years 10G PON (10G EPON, XGS-PON) and NG-PON2 (40G) are on the market. They are the candidates for 5G wireless backhaul and fronthaul. IEEE is

developing 25G/100G EPON. These high-speed PONs are suitable for RPHY backhaul applications. Since the NG-PON2 is the higher rate PON on the market today, it is used to illustrate PON RPHY backhaul.

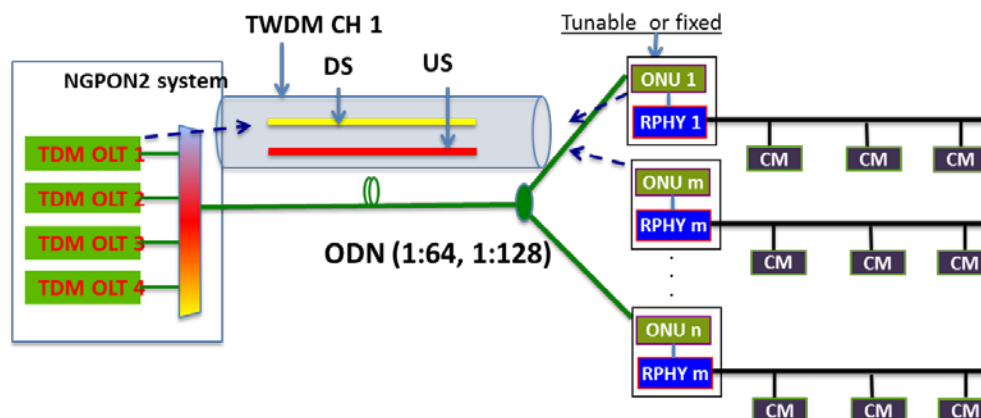


Figure 8 - NG-PON2 PON for RPHY Backhaul

NG-PON2 for RPHY backhaul architecture is in Figure 9. NG-PON2 is a hybrid TDM PON (10Gb/s/ch) and WDM (TWDM). The WDM PON architecture is based on tunable optics. It supports 4 10Gbps TDM PON channels (expandable to 8 TDM channels) and a number of P2P WDM PON channels. The downstream channels of NG-PON2 are in the L band from 1596 nm -1603 nm, and the upstream channels are in the C band from 1524 nm to 1544 nm. In current implementations, the TDM PON in NG-PON2 has 4 channels with 200 GHz channel spacing.

The cost of the NG-PON2 is a concern today due to the high cost of tunable lasers and tunable filters. However, it is feasible to implement NG-PON2 with fixed optical if bond 4 channels are used to form a 40Gbps PON with 4 pairs of 200 GHz or 100 GHz DWDM channels. This low cost solution is practical and could be standardized.

100G EPON (802.3ca), which has currently been developing at IEEE, is also a 4 channel WDM-TDM PON. The discussion of NG-PON2 for RPHY backhaul also applies to 100G EPON to some degree.

Figure 9 shows a NG-PON2 system for RPHY backhaul. In the field location, a RPHY is co-located with an NG-PON2 ONU. The RPHY and NG-PON2 ONU could be integrated on a circuit board, or connected vis external Ethernet interface. The Ethernet traffic from a RPHY is encapsulated into NG-PON2 frames to transport to the OLT. In this architecture, a NG-PON2 system also serves as an aggregation device. Compared with P2P DWDM solution, PON for RPHY backhaul is more efficient, it takes advantage of the statistical gain of bandwidth without using active Ethernet switches.

Another benefit of multi-channel PON such as NG-PON2 for RPHY back haul is traffic balance between the OLTs. In the tunable NG-PON2 case, an ONU can send traffic to multiple OLTs with fast tuning laser transmitter in principle. In the case of fixed optics with channel bonding, it is naturally a statistically traffic balanced solution.

6.2. NG-PON2 with 1:1 protection for RPHY backhaul

As discussed in in section 5.1, a DWDM RPHY backhaul and wireless fronthaul system prefers 1:1 fiber protection. This requirement also applies to PON for RPHY backhaul and wireless fronthaul. Figure 10 shows a NG-PON2 RPHY backhaul system with 1:1 protection on trunk fiber and 4:1 protection on OLT.

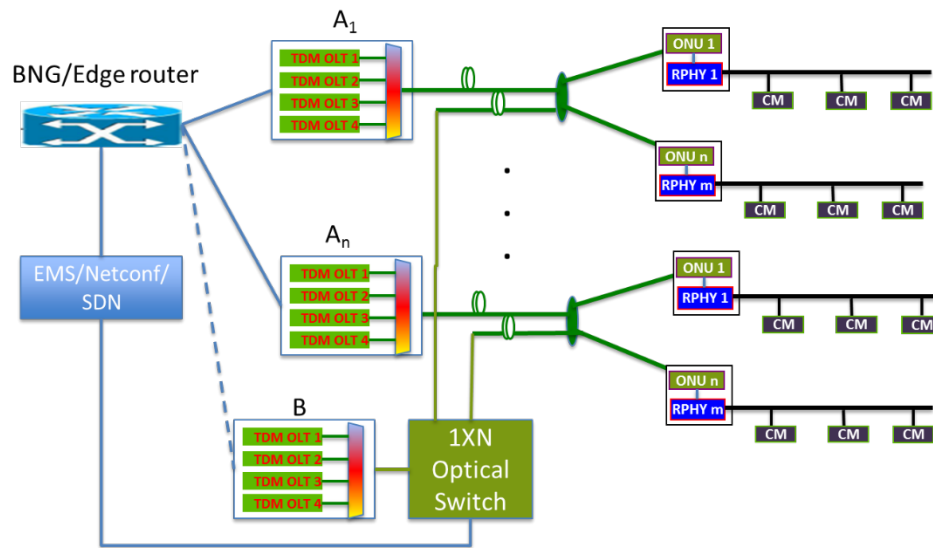


Figure 9 - PON for RPHY Backhaul with Protections

7. Conclusions

RPHY backhaul and 5G wireless fronthaul networks have similar network topologies, it is feasible to build unified architectures for both network applications. Two unified architectures are proposed. One is a unified passive DWDM architecture that coexists with all types of PON. The other architecture is based on high-speed PON. The DWDM is a mature technology for mobile fronthaul and RPHY backhaul today and the evolving high-speed PON may provide a lower cost solution for tomorrow.

8. Abbreviations

8.1. Abbreviations

AAWG	Athermal AWG
AON	Active Optical Ethernet Network
AWG	Arrayed Wavelength Grating
BBU	Baseband Unit
CPRI	Common Public Radio Interface
C-RAN	Centralized Radio Access Network
DOCSIS	Data over Cable System Interface Specifications

DS	downstream
EDFA	Erbium-doped Fiver Amplifier
EML	Electro-absorption modulator laser
GPON	Gigabit Passive Optical Networks
HFC	Hybrid Fiber Coax
ODN	Optical Distribution Network
OLT	Optical line terminal
ONU	Optical Network Unit
OSP	Outside Plant
PON	Passive Optical Networks
QAM	quadrature amplitude modulation
RPHY	Remote Physical Layer Device
RoF	Radio over Fiber
RRH	Remote Radio Head
SOA	Semiconductor Optical Amplifier
US	upstream

9. Bibliography and References

[1] J Hasegawa and K Nara, “ Development of Wide Operating Temperature Range (-30 to 70/C) Athermal AWG Module with high Reliability”, Furukawa Review, No 30 2006.

[2] Yuxin (Eugene) Dai, “Unified Residential and Business Services Access Network Architecture for MSOs”, SCTE Cable-Tech EXPO 2010 Technical Forum

[3] Yuxin (Eugene) Dai “Beyond RFoG – DOCSIS PON Backhaul for Smooth Migration to PON”, SCTE•ISBE

Fiber Deep Networks and The Lessons Learned From The Field

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

Introduction

Fiber Deep architectures are becoming a natural next step in the evolution of the HFC network. As the need for bandwidth continues to grow, fiber will continue to be pushed deeper into the network. The constant expansion of fiber has already been witnessed in both telephony and cellular networks. The movement from DSL to ADSL and ultimately to VDSL2⁺ had the same goal - drive fiber closer to the customer and improve bandwidth at the customer premise. The same can be also said for the evolution of wireless networks. As the networks have upgraded from 1G to 4G, it has been about getting fiber closer to the customer and improving the amount of data available to watch the latest sports game or facetime with a friend. The world is fixated on data, and fiber is a medium which can help meet the demands of today and into the foreseeable future. The key is balancing the cost of deploying fiber against the demand for bandwidth.

As we look at how to balance the cost of deployment, a key factor to be considered is the amount of bandwidth available. The differentiating factor of the HFC network over competitive technologies is the wonderful coaxial cable. The natural form of the coaxial cable provides strong shielding of interfering signals from the center conductor which transmits both video and data. This shielding is a key differentiator which allows the coax to transmit significantly more data over the same distance versus a traditional twisted pair network. As can be seen in the current DOCSIS 3.1 technology, the HFC architecture using a coaxial last mile has the ability to provide up to 10 Gbps using today's spectrum defined for DOCSIS 3.1. Today, due to spectrum allocation of video and data, capacity available is 1-2 Gbps. . In comparison, the latest deployable twisted pair technology G.FAST is capable of delivering only 800 Mbit/s over 300 feet. As can be seen, the coax cable has a significant amount of bandwidth available which makes it an excellent medium for the last mile.

With this in mind, the next focus is on minimizing deployment costs of new technologies. To do this effectively, a fiber deep architecture effectively utilizes the current last mile while providing higher bandwidths and a significant amount of additional network benefits.

Fiber Deep Networks

A Fiber Deep Network typically refers to an HFC network which has had all line amplifiers removed and fiber nodes are therefore close to a customer's premises. The general term for this type of network is N+0 (Node plus Zero Amplifiers). Typically, a fiber deep node will serve between 50 and 128 homes.

In a typical network migration, fiber nodes are split as the capacity of the node reaches its limit. The capacity can be referring to either upstream or downstream data bandwidth. A general rule of thumb is to split a node when either the upstream or downstream capacity reaches 80% of the total available capacity. This provides the necessary time to attain the proper permits and get the work scheduled before the node reaches full capacity before customers notice congestion.



So why go all the way to N+0 instead of adding fiber nodes and removing amplifiers as the bandwidth requires? The primary reason is to design a network which is easily upgradeable to next generation technologies. As a network is slowly upgraded over time (node splitting), the new network which emerges becomes less optimally segmented and potentially is not easily upgraded to next generation technology due to a mix of evolving technologies and legacy networks solutions coexisting in adjacent sections of the network. When a network is properly designed to N+0, the fiber deep nodes can be strategically placed to be within specified distances of all homes. Designs are generally optimized for maximum reach from the fewest nodes for economic reasons, and a good general rule is that nodes will be 1000 feet of the furthest home that is being served. This allows an upgrade path to Remote PHY, DOCSIS 3.1 full duplex, and potentially FTTH for customers looking interested in fiber connectivity and speeds.

Another benefit of making the leap to a fiber deep network is simplification of operational procedures and practices. For large networks, an entire city can be upgraded within a year. Once the new network is in place, all maintenance procedures and operational practices will be consistent within that city and typically within a serving area of a regional office. Technicians will enjoy not needing to learn how to maintain five different versions of the same fiber deep node which will also reduce errors in the field, will not need to maintain amplifiers, and have less distribution cable.

Probably the most important benefit will be a strong increase in customer satisfaction. Based on current bandwidth calculations, a properly deployed fiber deep network will provide enough bandwidth for the next 7-10 years for 90% of households passed. For the 10% that still require more bandwidth, upgrade paths to Remote PHY and FTTH are available.

Fiber Deep Serving Groups

The challenging nature of today's HFC network lies in the wide variety of architectures which exist. Even today small isolated sections which are Node + 10 Amplifiers (N+10) can still be found. These areas are rare but they do still exist and compound the challenge of upgrading a network. Over the years, most operators have been working aggressively to bring their networks into an N+3 or N+4 state. In the most bandwidth-congested areas, an N+2 network may exist to provide higher bandwidth services. As mentioned previously, a strength to a focused fiber deep strategy is converging all of these different network types into a single uniform architecture.

So where does one start? Most operators tend to start with homes passed by a node as the base criteria. Today's networks tend to pass 500 – 2500 homes per node. The goal for a fiber deep network should be between 50 and 128 homes passed per node. The 128 is typically used as a maximum due to its alignment with both RFoG equipment and FTTH OLT equipment. A typical OLT port support 32 homes thus 4 OLT ports will support up to 128 homes. Similar in a RFOG environment, 128 homes is the maximum amount of homes which a single RFOG transmitter can support. Once a range for the number of homes passed has been decided, new serving groups can be overlaid onto the existing architecture and a single node placed within each of those serving groups. The goal will be to center the node as much as possible with the serving group while still utilizing as much of the existing coax as possible. Installation of new coax should be considered as a secondary option due to the additional cost associated with the installation. Existing amplifiers will be removed and replaced with a shunt to allow power to pass through but not RF. Once the new serving groups are overlaid, the power plant needs to be redesigned to support both the installation of the higher powered nodes plus the removal of the amplifiers.

An operator has basically two options to upgrading their power network. The first option is to redesign the power grid to optimize energy losses and reduce power loss within the coax. The second option is to maintain existing power supply locations and drive power from these locations to the new fiber deep nodes. Let's take a look at both of these options individually.

1. Distributed Powering

There have been a few trials performed over the past two years focused on many factors of a fiber deep network; however, one of the main focuses of the trials has been the final location of the power supplies. The initial theory was to redistribute the power supplies to minimize loss of power in the network and to redesign the network to maximize power efficiency. Due to the excess of power added to the network back when circuit switch telephony was a hot topic, a typical power supply runs today between 50% and 60% of its maximum capacity. Even with a 20% safety factor designed in, there is a significant amount of power available to be used.

The second variable in the equation is the reduction of power as the network transitions from an amplifier rich network to a node rich network. To demonstrate this, a N+3 network today will typically have seven amplifiers cascaded from a single leg of a node. Most nodes today have four output ports which equates to 28 amplifiers downstream of a node. Using 35 watts as typical power draw for an amplifier, the power draw on this amplifier segment is roughly 980 watts. Adding the 60 watts for the node give a total of 1040 watts for the node plus the amplifiers. Assuming that a N+3 network will require six fiber deep nodes at 140 watts per node, the new power load of the fiber deep network will be 840 watts. Theoretically, there is 200 watts of left over power.

A quick initial look at powering shows that theoretically it makes sense to redistribute the power and optimize the network at the same time.

2. Centralized Powering

So, why are we discussing maintain the power supply locations and not moving them to better optimize the network? As is normally the case, theoretical and practical do not always match. The key to the challenge is the practicality of moving power supply locations. Unfortunately, the same old problems continue to haunt today's network expansions primarily with the main problem being the permitting of new sites. It is not even the cost which becomes the major issue. Based on the trials performed, the biggest factor was the timing of getting the new permits and scheduling the work. With one of the trials, delays in permitting cost typically three to six months of delay. The delays became so significant that work was cancelled until after permits could be acquired which then put another delay into the overall schedule as work to install equipment and do the fiber deep upgrade was not scheduled until the permits were obtained.

This became such an issue that one major operator put moving the power supply as the absolute last option when designing a fiber deep network. Deploying 23 Ohm power feeder cable is considered a better option over moving power supply locations.

When considering powering of the new fiber deep nodes, it is generally a good idea to use a 140 Watts as the new load of the fiber deep node. If there is not enough power initially available, there are a few options to provide the additional power to the node. The first thing to consider is pulling power from an adjacent section of the network if there is additional power available. A secondary option would be to

increase the diameter of the cable to .875" to reduce resistance in the hardline cable. A third option would be to increase the capacity of the power supply available. For example, 15 Amp power supplies can be upgraded to 18 Amp power supplies. 18 Amp power supplies can be upgraded to 24 Amp power supplies with 240VAC feeds.

In a nutshell, the key lesson learned over the past two years of fiber deep trials and network expansion is that it is more important to maintain the existing power supply locations to meeting network upgrade timelines and budget.

Remote PHY

The trials are just starting today for Remote PHY upgrades. As is normally the case, the next architecture upgrade is right on the heel of the current technology. The nice thing about Remote PHY is that it fits nicely as the next step for fiber deep. Fiber deep reduces the home passed per node. Remote PHY takes the next step and moves the PHY chip from the CMTS to the node. This has multiple advantages which include both the headend and the performance of the plant. For the headend, there is a significant reduction in space over current and prior technologies freeing up valuable space for new services. For the node, it has a significant impact on available bandwidth, transmission speeds, and architecture flexibility. With new services like small cells, there is still debate today around current architectures with DOCSIS 3.0 technologies having the speed and latency requirements to support 5G wireless networks. The benefit which Remote PHY provides is 10GbE connectivity which enables the possibility of edge computing/processing happening in the node which could help deliver on low latency requirements. A combination of Remote PHY and DOCSIS 3.1 technology provides a network which has all of the benefits of real-estate, backhaul, and backup powering to be a strategic advantage in new and emerging business to business technologies.

So what are the challenges to Remote PHY in the outside plant? Today, the key challenge is with heat dissipation especially in ground level deployments. The current generation of node enclosures are designed to support 60 Watt loads. The new generation of Remote PHY nodes can be loading the network with up to 180 Watts. Initial testing has shown that a standard node enclosure does not perform well with a 180 Watt load installed inside of it. Testing results have shown that internal temperatures of the node can reach higher than 170° F on a sunny day with 90° F ambient temperatures. Newer generation of node enclosures are designed to support the higher loads and provide more airflow to keep the node cool. It is critical to perform a testing regime which mimics the extreme conditions for the intended network to prevent unexpected infant mortality conditions.

Over the years, many operators have created rules around never putting anything on the network which could potentially pull more than 150 Watts. This was done to protect the integrity of the network and reduce unexpected issues in the field. Due to the number of instances in the past of large loads taking down a portion of the network that these rules were put in place. Now, the latest generation of Remote PHY nodes is forcing us to rethink these rules.

So, how are large loads better managed? There are a few key items to be considered when understanding the impact of a large load on the network. Remember one of the first equations taught when taking any type of electronics course: " $P=VI$ " Power is equal to the voltage times the current. This equation is extremely important to understand the impact of the larger load on the network. Assuming that the voltage right at the power supply is 90V will allow us to quickly calculate that the current draw of a 180 Watt load right at the power supply is 2A. Unfortunately, the loads are not typically placed right at a

power supply. Now, we have moved further out into the network to a spot where it would be ideal to place a node. If the voltage at this location is 60 Volts, the new current draw will be 3 Amps. But wait, is it really 60 volts anymore? Unfortunately, placing the load on the network has increased the current draw through the coax cable. Now it is time to apply Ohms law: “ $V=IR$ ” For this example, the resistant of the coax cable is assumed to be 2 Ohms. The original current in the cable is 6 Amps. Therefore, the original voltage drop across the cable is 12 Volts. Now with the new load, the current increases to 7 Amps and the voltage drop becomes 14 Volts. If the voltage at the node location was original 60V, it will now be 58V. The impact of adding a large node can ripple through out a network. The further away from a power supply the load is applied to the network, the bigger the affect that the load will have. If the impact is too far away, a network can be destabilized and crash.

To maintain a healthy network when building a fiber deep network, there are a few key lessons to take into consideration. The first lesson is to plan for the larger 180 Watt Remote PHY nodes but design the network for the 140 Watt fiber deep nodes. When possible, attempt to future proof the network to simplify the deployment of Remote PHY nodes. Always maintain a minimum 50 VAC at the fiber deep node to reduce the impact of the load onto the network. Fiber deep network upgrades do not typically have a major impact on the amount of power used by the network. Individual sections can vary widely depending on the condition of the current network. There have been upgrades where power supplies were decommissioned and upgrades which require more power. To date, there did not appear to be a strong correlation as to why some network upgrades required additional power supplies and some did not. As deployments continue, a correlation may emerge.

To Fiber or Not to Fiber?

For the past fifteen years, there has always been one looming question out there. Does it make sense to skip right to the end state and put fiber all the way to the residence? Even after all of the advancements in fiber deployments, it is still a tough question to answer today. The cost of deploying fiber has come down significantly since Verizon jumped in two feet first but has it come down enough to make it the first choice for all applications? To help answer this question, the first step will be to look at a cost comparison between upgrading a current HFC network to fiber deep versus jumping straight to FTTH. Then follow on with a discussion between the two deployment types.

1. FTTH – FDH + Hardened Drop Terminal

There are many ways to build a fiber-to-the-home (FTTH) network and to keep it simple, the most common type of network deployed in the US has been chosen. In this model, fiber is run from the headend via a fiber ring into a neighborhood. The fiber ring uses a branch splice terminal to tap off up to 48 fibers which feed a 288 fiber FDH. 18 of the 48 fibers are used to feed splitters, the remainder are used for point-to-point connections. The FDH cabinet can feed up to 36 hardened drop terminals (8 fiber ports per terminal). Smaller hardened drop terminals (HDTs) can be used, however, 8 fiber terminals were used for this model. The last connection to the home is a single pre-connectorized fiber drop. These are used between the home and the HDT. There are a few key reasons why this network type has been chosen.

1. This has been the predominant network type chosen by most large telco and cable operators
2. This network type typically has the lowest total cost of ownership. This is primarily due to the fact that general contractors can install the drop cable minimizing the use of fiber splicing technicians
3. The multiple connection points within the network provides easy access for trouble shooting.

4. Future upgrades are easier to implement due to the fiber optic splitters and customer connection points co-located within the FDH cabinet

With all of the positives, there are a couple of drawbacks to consider.

1. The FDH cabinets tend to be large due to the number of splitters and end user connection points.
2. The up-front cost of the deployment is typically higher

The model is generic in nature and should only be used as a reference point. It is built around a perfectly laid out housing addition which has four homes per block and all laid out in a perfect grid style pattern. This is done to make lengths of fibers and drops easier to calculate. The goal of the model is not to have a perfect cost but to have a comparison which can be used to assess the two network styles. As the geography of the residential neighborhood gets more complex, the assumption is that the cost will increase for both types of networks in a similar fashion.




Figure 1 - Residential Layout

For the FTTH model, two factors need to be taken into consideration: the cost to “pass” a house and the cost to “connect” a house. When investigating the total cost to deploy a network, it is very important to make an assumption on the percent of homes which will be connecting to the network. Based on this percent, different network types should be considered. For example, if there is a very low connection rate expected, then a heavier spliced network should be considered. The additional cost to connect a home will be small compared to the cost of deploying a network. For a more typical expected connection percentage (>35%), the cost to connect a house becomes the more significant factor and will drive an overall lower TCO. Especially if there is a medium to high churn rate, then simplifying the connection to the home is critical to keeping ongoing cost lower.

When purchasing a FTTH network, it important to fully understand the way the network was built. To keep initial costs lower, companies will build a network which uses splicing as a way to connect

customers. This becomes a burden for the operator who manages the network and is required to turn customers on and off.

There are four areas of cost which the model is focused on. The first is the material cost to pass 5000 homes. The second is the labor cost to install the materials for the 5000 homes. The third is the material cost to connect a single home and the final is the labor cost to connect the home.

The following figure shows how the FDH and HDTs are laid out for the 288 homes. The yellow rectangles represent houses and the black lines represent roads.




Figure 2 - FTTH Residential Layout

The total cost to pass 5000 homes is as follows:

Material Cost:	\$310,000
Labor Cost:	\$650,000
Total Cost:	\$960,000

~\$192 per home passed

The material cost to connect a home includes the hardened drop cable, the home fiber patch cord, the ONU, and the demarcation box at the side of the house. The labor includes the cost to install a drop aerially to the house, the installation of the demarcation box, the splice to connect the two cables plus an hour of support for the installation of the ONU.

The total cost to connect a home is as follows:

Material Cost: \$162
 Labor Cost: \$185
 Total Cost: \$347

The following table highlights the total cost to connect 5000 homes based on the projected percent of homes connected.

Table 1 - FTTH Total Cost

%	Total Cost	%	Total Cost
5	\$1,042,062	55	\$1,910,062
10	\$1,128,862	60	\$1,996,862
15	\$1,215,662	65	\$2,083,662
20	\$1,302,462	70	\$2,170,462
25	\$1,389,262	75	\$2,257,262
30	\$1,476,062	80	\$2,344,062
35	\$1,562,862	85	\$2,430,862
40	\$1,649,662	90	\$2,517,662
45	\$1,736,462	95	\$2,604,462
50	\$1,823,262	100	\$2,691,262

Based on a FDH with 8 fiber HDT, it would cost \$1.82 million to pass 5000 homes and connect 50% of those homes (~\$365 per home passed).

2. Fiber Deep: N+3 → N+0

Now for comparison, a “typical” HFC N+3 network is converted to N+0 by pushing fiber further into the network by eliminating all amplifiers and utilizing DWDM splitters at the original node location to support the addition of new nodes downstream of the original node location.

The following figure represents a simplified N+3 network. A single node is supporting a network of amplifiers which provide service to 288 homes. Each active supports 16 homes via four drop terminals.

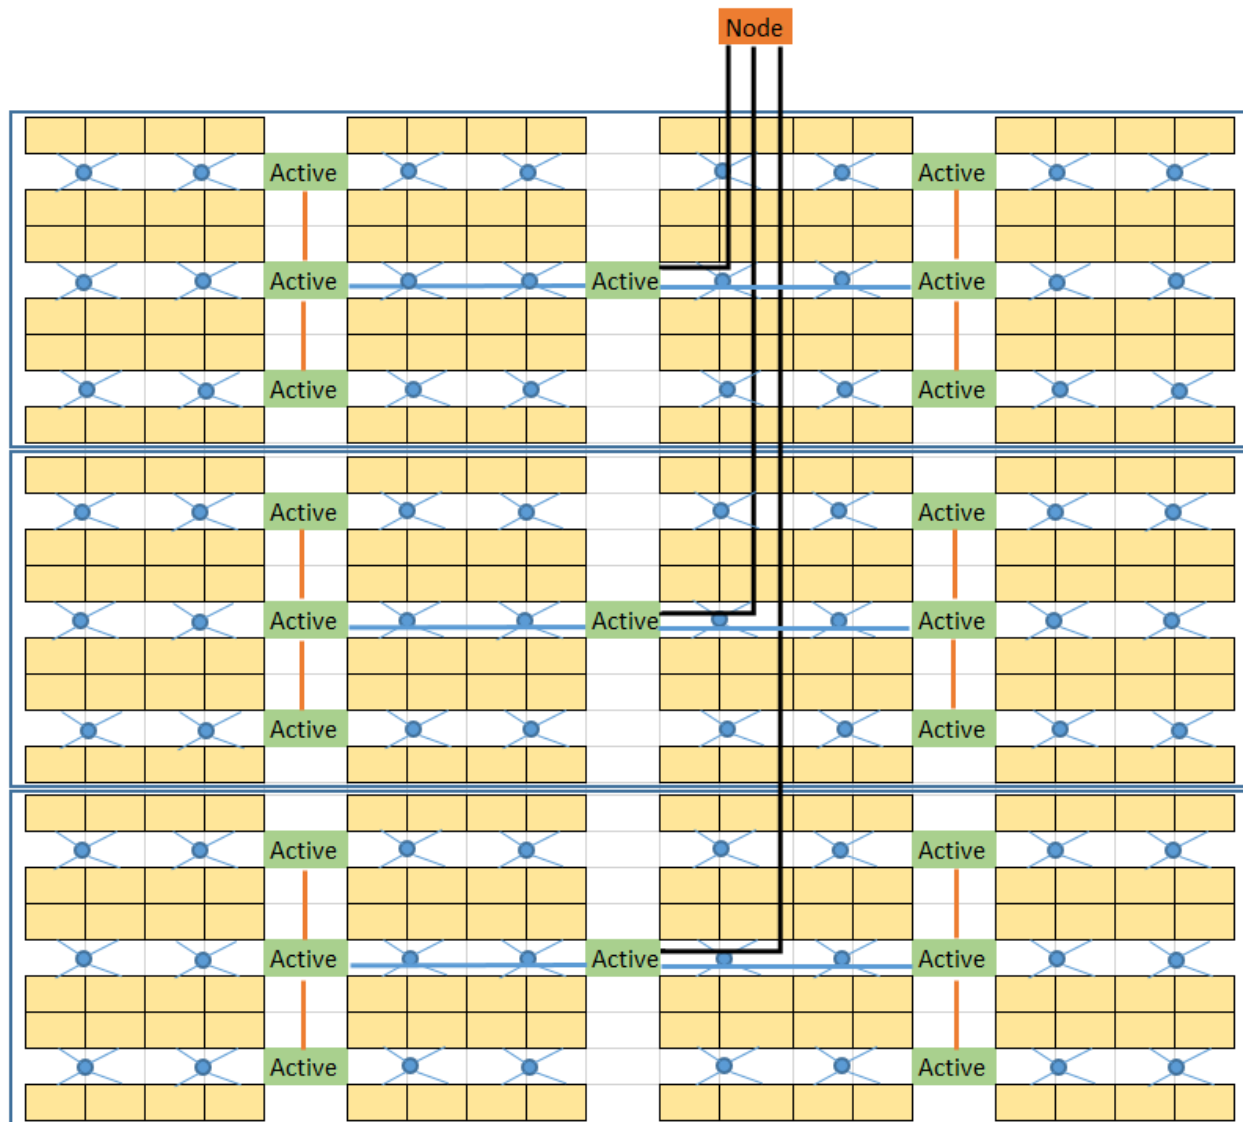


Figure 3 - N+3 Network

To move to a fiber deep network, the initial node is replaced with a DWDM splitter and typically 12 fiber or 24 fiber cable is laid between the old node location and the new fiber deep locations. To keep the model simple, individual fibers from the original node site were placed to the new fiber deep locations. In practice, a 24 fiber cable may be placed with four to six fibers dropping at each new node location. If properly done, this would reduce the cost of deploying fiber to the new locations.

Figure 4 below shows the new layout of the N+0 network. The 20 amplifiers from the N+3 network have been removed and replaced with 6 fiber deep nodes. The model assumes that power supply locations are not moved and five power bridging circuits are utilized to move power from one line to another when needed within the 288 home serving group. The new N+0 network has 48 homes passed per fiber deep node.




Figure 4 - N+0 Network

Similar to the FTTH model, both material cost and labor cost were calculated to create a comparison. For material, the cost of fiber deep nodes, 12 fiber cable for feeding the new nodes, fiber splice closures, DWDM Splitters, and the power inserters/ bridge circuits were used. For labor calculations, node installation, splicing of the 12 fiber cable along with splicing of the DWDM splitters, fiber splice closure installation, and power inserter installation was utilized.

The total cost to convert the N+3 network to N+0 for 5000 homes is as follows:

Material Cost:	\$442,500
Labor Cost:	\$206,000
Total Cost:	\$648,500
~\$130 per home passed	

The model used an average of rates from a couple of sources to calculate the installation costs and was based fully on an aerial plant.

The question is now whether or not it is worth doing a one-time cost of \$1.5 million to get to a full fiber network or invest \$650 thousand to upgrade the current network. There are a few more factors which should be looked at when deciding if it is worth going fiber deep or FTTH.

A new growing focus on business-to-business services is driving a strong interest with the power, backhaul, and real-estate provided by the HFC network. Traditionally, the power requirements of the nodes and amplifiers was the primary reason to invest in a reliable powering network. 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 a single unit 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.

When strictly discussing a fiber deep network, the total power used will typically be less than the power required today. When doing a fiber deep upgrade, it is important to consider SWISS type of initiatives, these initiatives will drive a higher demand of the network with some of the current small cell opportunities being in the range of 300 – 400 watts at a single location. It is important to not reduce powering locations within the network when doing upgrades and redesigns. Power is a critical component which will differentiate the HFC network from the traditional telco network.

Conclusion

Financially, it makes sense to continue the upgrade of the HFC network to fiber deep. The bandwidth and flexibility which a fiber deep network provides is second to none in the world. With DOCSIS 3.1 upgrades, the practical viability of the HFC network is solid for the foreseeable future. Planning ahead while deploying a Fiber Deep network will enable an operator to have enough spare fiber for fiber-to-the-radio, fiber-to-the-home and additional new services. The coax cable is such a simple and wonderful design which will continue to separate the HFC network from the traditional telco network. This network will be the future for both residential broadband as well as new services such as small cells, WiFi, IoT and Security & Surveillance. These type of services will fuel the industry for many years to come.

Abbreviations

N+0	Node plus zero amplifiers
N+3	Node plus three amplifiers
HFC	hybrid fiber-coax
DWDM	Dense Wavelength Division Multiplexing
B2B	Business to Business
SWISS	Small Cell, WiFi, IoT, Security & Surveillance
SCTE	Society of Cable Telecommunications Engineers
DOCSIS	Data Over Cable Service Interface Specification
HDT	Hardened Drop Terminal
FDH	Fiber Distribution Hub
FTTH	Fiber to the Home
P	Power
V	Voltage
R	Resistance
I	Current
10 GbE	10 Gigabit Ethernet

Bibliography & References

Giving HFC a Green Thumb; John Ulm / Zoran Maricevic, Arris Corporation

DOCSIS 3.1 Downstream Early Lessons Learned

A Technical Paper prepared for SCTE•ISBE by

John J. Downey
CMTS Technical Leader
Cisco Systems Inc.
RTP, NC
919-392-9150
jdowney@cisco.com

Introduction

DOCSIS 3.1 (D3.1) systems have been deployed in operating cable networks and some lessons have been learned. Allocating dedicated spectrum for new D3.1 cable modems (CMs) is a scary proposition and cable operators know what's at stake. To compete with the marketing hype of fiber-to-the-home (FTTH), the cable industry has many tricks up its sleeves such as D3.1.

The D3.1 downstream (DS) has been the focus for the last year with the upstream (US) just around the corner. D3.1 offers speeds required by our customers and the tools and feature-set to make it work properly for the highest availability possible. A few features that help achieve "self-healing" are: DS resiliency, graceful profile management, mixed-modulation profiles, and more. For on-going operational monitoring, the DOCSIS standard has built-in proactive network maintenance (PNM) features to identify and locate hybrid fiber-coax (HFC) impairments quickly before customers even know they exist.

This paper and presentation explore the intricate details of real-world scenarios of DOCSIS 3.1 deployments using best practices to optimize the cable modem termination system (CMTS) for maximum throughput even in an impaired environment: the type of environments every cable operator eventually experiences.

Content

1. General Basics

The D3.1 DS uses orthogonal frequency-division multiplexing (OFDM) with spectrum options of 24 to 192 MHz in block/channel widths. The actual spectrum is 204.8 MHz, but that is not what is seen from an RF standpoint when viewing on a spectrum analyzer. A byproduct of this is that one may notice CM time offsets 20 times larger because DOCSIS 1.x, 2.0 and 3.0 are based on 10.24 MHz clocking. This D3.1 204.8 MHz clocking can create a misconception of much higher time offsets.

D3.1 DS modulation schemes range from 16-QAM to 4096-QAM with even higher options possible in the future. Five data profiles can be supported for modems to choose from with automatic profile selection. Some of the profiles could consist of mixed modulation with five to seven different modulation schemes in the entire DS block/channel. This may be advantageous for known plant roll-off issues above 750 MHz or 860 MHz, etc.

Exclusion bands are also possible to avoid ingress or create a null space to inject a carrier for amplifier automatic gain control (AGC), leakage testing, or alignment tones for rough balancing.

D3.1 CMs can cross-bond traffic between legacy DOCSIS carriers known as single carrier QAM (SC-QAM) signals with one to two OFDM channels. This will be the typical deployment scenario for many years to come to support legacy CMs and provide very high peak speeds for D3.1 CMs with 1 Gbps as a goal. The aim for these D3.1 CMs is to prefer the OFDM ch before utilizing the SC-QAM signals so as not to "starve" out legacy CMs. "Primary capability" can be either a SC-QAM or the OFDM block with its coinciding physical layer link channel (PLC).

At this time, few systems have any purely D3.1 service groups (SGs). SC-QAM signals will probably be present for many years for D2.0 and D3.0 CMs. With that said, it doesn't mean a SC-QAM signal is necessary for D3.1 CMs to operate, but it is advantageous to cross-bond for the following reasons:

1. Bigger channels are always best. Higher peak rates can be achieved and less need for load balance to occur from legacy CMs and between D3.0 bonding groups.
 - a. D3.1 prefer OFDM for traffic before SC-QAM signals are used anyway, so no concern for "starving" legacy CMs.
2. Having a SC-QAM primary allows DS resiliency to work a bit better. More on that later.
3. SC-QAM signals for voice over IP (VoIP) traffic and other low-latency queuing (LLQ) traffic may be needed.
4. Not all SC-QAM signals should be primary-capable when we start having 24 and more channels. Almost 2-3 Mbps of overhead is created when a SC-QAM signal is primary.
5. In case of an emergency situation like overheating, one could shut down the OFDM first as it could be a major contributor.

Current D3.1 CMs on the market support 32 SC-QAM plus 2 OFDM blocks. They also support up to 4096-QAM (also called 4K-QAM) even though 8K-QAM and 16K-QAM are options in the specification. CM spectrum support is 1 GHz with most at 1.218 GHz. 1.794 GHz is an option in the spec, but may or may not be pursued depending on full duplex DOCSIS (FDX), which is under development at Cablelabs.

In regards to US support, D3.1 CMs support eight SC-QAM signals along with two 96 MHz (max) orthogonal frequency-division multiple access (OFDMA) blocks. Even if the CM US chipset supports a return path upper frequency limit of 204 MHz, the internal duplex filter may be hardware or software limited to 42 MHz or 85 MHz.

Another potential issue to be aware of is DS power levels supported from a CMTS connector. This is specified by the DOCSIS Downstream Radio Frequency Interface (DRFI) spec. The level range vs channel loading is listed in Table 1. From this table, one can see that more spectrum (chs) located on a DS connector, whether video, legacy SC-QAM or D3.1 OFDM, etc., will lead to lower max output available and precautions need to be taken. Headend wiring may need to be shortened or changed from mini coax back to Series 6, passives swapped/replaced, etc.

Table 1 - Max Carrier/OFDM Loading to Channel Level Range

Max Carrier	No OFDM	24 MHz OFDM	48 MHz OFDM	96 MHz OFDM	192 MHz OFDM	384 MHz OFDM
8	41 – 50	39 – 48	37 – 46	35 – 44	32 – 41	29 – 38
16	37 – 46	36 – 45	35 – 44	34 – 43	31 – 40	29 – 38
24	35 – 44	34 – 43	34 – 43	32 – 41	31 – 40	28 – 37
32	34 – 43	33 – 42	32 – 41	31 – 40	30 – 39	28 – 37
48	31 – 40	31 – 40	31 – 40	30 – 39	29 – 38	27 – 36
64	30 - 39	30 – 39	29 – 38	29 – 38	28 – 37	26 – 35
96	28 – 37	28 – 37	27 – 36	27 – 36	26 – 35	25 – 34
128	26 – 35	26 – 35	26 – 35	26 – 35	25 – 34	24 – 33
158	25 – 34	25 – 34	25 – 34	25 – 34	24 – 33	Not Possible

The ranges listed in Table 1 are in dBmV and the maximum value is 1 dB above the DRFI spec. A CMTS with a tilt function could help engineers achieve flat inputs to the DS optical lasers with options of linear or non-linear tilt along with power level offsets as well.

2. Spectrum Allocation and Thoughts

Current deployments are looking to extend the DS out to 1 GHz or 1.218 GHz for actives/passives. D3.1 allows a higher option of 1.794 GHz, but current CMs don't support it.

The D3.1 192 MHz block(s) starts at 108 MHz optional with a 258 MHz starting frequency mandatory. The D3.1 CM's filter and/or scanning table may negate any D3.0 operation < 261 MHz center frequency for SC-QAM signals. Potential ingress sources reside throughout the spectrum from long term evolution (LTE) 4G mobile to Multimedia over Coax Alliance (MoCA), garage door openers, off-air broadcasters, etc. Proper planning for channel allocation/location in the spectrum should be done for the intended market.

US decisions also affect the low-end of the DS. One also needs to consider DS amplifier AGC/ALC and potentially legacy settop box out-of-band control channels. Figure 1 displays potential spectrum allocation and ingress sources along with typical duplex filter splits.




Figure 1 - DS Spectrum Allocation

3. Interference Resiliency Testing

To prove the resiliency of D3.1 with its added robustness from low density parity check (LDPC) forward error correction along with time and frequency interleaving, some testing was done with SC-QAM signals simulating wideband ingress. Some baseline configs and steps are listed below:

- An OFDM channel with 192 MHz width using 25 kHz subcarrier spacing.
- A DS bonding group with one SC-QAM plus OFDM block with both as primary.
- Data traffic set to ~ 1.4 Gbps through a few CMs.
- A test CM forced to use 4096 QAM.
- An interference source comprising eight SC-QAM signals from a DS port and combined with the working port.
- Variable padding added to the interfering signal starting with 30 dB.
- A speed test was observed to make sure no packets were lost.
- Receive MER (RxMER) readings were documented along with the “break point”.

Table 2 is taken from the DOCSIS 3.1 PHY spec’s Table 7-41. This table lists the suggested operational RxMER for the corresponding modulation schemes and bit loading (bits/symbol) for that modulation. These values are used by the CMTS to make the decision for graceful profile management.

Table 2 - RxMER to Bit Loading Mapping

RxMER (in ¼ dB)	RxMER (in dB)	QAM	Bit Loading
60	15	16	4
84	21	64	6
96	24	128	7
108	27	256	8
122	30.5	512	9
136	34	1024	10
148	37	2048	11
164	41	4096	12
184	46	8192	13
208	51	16384	14

Note: It's in this author's opinion that these thresholds are ~6 dB too conservative and it may be worth having a correction factor of 2-4 dB configured in the CMTS.

Figures 2 through 6 display the test results from various ingress levels and show RxMER vs subcarrier frequency. Figure 2 is the baseline 192 MHz wide OFDM block with no ingress. The average RxMER for all the active subcarriers is 47.4 dB.




Figure 2 – 47.40 dB Avg RxMER with No Interference

Note: It's this author's opinion that this modem is reporting the pilot RxMER without the proper 6 dB correction factor and hence reading 6 dB higher than it should.

Figure 3 shows an average RxMER of 41.88 dB with the ingress activated with 20 dB padding vs the 30 dB start value. As can be seen in the figure, the eight SC-QAM signals of ingress are causing lower RxMER values of ~ 28 dB in that region. With such low RxMER and affecting $48/192 = 25\%$ of the block, the internal threshold table would suggest this modem drop from 4K-QAM to 256-QAM!




Figure 3 - 41.88 dB Avg RxMER with 20 dB Padding

Interesting enough, this modem did not drop any packets and continued at 4K-QAM at full throughput. All active subcarriers were being used at 4K-QAM even in the region where RxMER values were ~ 28 dB! This is believed to be possible because of the fast Fourier transform (FFT) functionality along with LDPC and time and frequency interleaving.

Note: The secondary observation was a very slight RxMER decrease for all subcarriers even where there was no ingress. The belief here is that the ingress is being spread across all subcarriers even though it is actually localized in specific spectrum.

Figure 4 shows the average RxMER dropping to 39.13 dB and all the subcarriers getting affected in addition to just the ingress localized subcarriers.




Figure 4 - 39.13 dB Avg RxMER with 5 dB More Ingress

Note: All RxMER values dropped after the RxMER in the 48 MHz interference spectrum dips below ~ 25 dB. This was observed even when the ingress was just one SC-QAM signal.

Figure 5 continues this trend and we still did not observe any packet drops!




Figure 5 – 37.14 dB Avg RxMER with 8 dB More Ingress

Figure 6 increase the ingress so that the localized subcarriers drop below 20 dB RxMER and still operating without dropped packets. Now all the other subcarriers are definitely getting affected and almost exceeding where 4K-QAM threshold is listed. The avg RxMER is definitely past the threshold and very near the actual “break-point”.




Figure 6 - 35.59 dB Avg RxMER with 10 dB More Ingress

The following observations were made:

- Steady-state interference did not affect throughput as anticipated even at very low RxMER readings. Possibly an added effect of LDPC, freq interleaving and FFT functionality.
- Overall subcarrier RxMER values dropped when the impaired subcarrier RxMER got < ~25 dB.
 - Actual impaired spectrum may be a deciding factor, but not observed.
 - Not an issue when exclusion bands used.
 - Possibly an added effect of FFT functionality and/or overdrive.
- Some CMs appear to incorrectly report RxMER ~ 6 dB higher than other CMs and may be based incorrectly on pilot readings.
- Mixed-modulation profiles probably not necessary for ingress, but for roll-off, maybe.
- The PLC uses 16-QAM and is very robust, but avoiding known ingress spectrum is still a best practice worth following.

4. Graceful Profile Management

The CMTS incorporates a feature defined in D3.1 as profile management. This can be done with a software defined network (SDN) application or within the CMTS itself. The CMTS transmits to the CM

on a control profile (profile 0) until the CM RxMER information for all active subcarriers is received by the CMTS to determine the optimal configured profile. The CMTS periodically polls each CM to gather RxMER of every active subcarrier and recommends the best profile configured in the CMTS.

Even if thresholds are aggressive, too lenient, or no updated RxMER readings from a CM, a “catch-all” is still used and that is the processing of a cm-status message of “unfit profile”.

From testing results and “real-world” deployments, the following deviations from Cisco CMTS defaults are suggested: (refer to your CMTS vendor for specific command syntax)

1. Changing modulation for an entire block for a small amount of subcarriers is not in our best interest. The default is to ignore 2%, but increasing that to 10% is suggested.
 - o `cBR8(config)#cable downstream ofdm-prof-mgmt exempt-sc-pct 10`
2. The internal threshold table is a bit conservative. The default RxMER offset is 0 dB with a value in quarter dB steps. Increasing that to 12 is suggested. $12/4 = 3$ dB correction factor.
 - o `cBR8(config)#cable downstream ofdm-prof-mgmt mer-margin-qdb 12`

One can also statically map a CM to particular data profile, if so desired.

- `cBR8(config)#cable down ofdm-flow-to-profile profile-data <1-5> mac-addr <>`

D3.1 CMs can lock to four profiles plus the control profile and store this in NVRAM. The CMTS and CM can support more, but the CM will need a dynamic bonding change (DBC) to support the fifth and could cause dropped packets. With this knowledge, it may be best to:

3. Use 256-QAM for the control profile and 4K-, 2K-, and 1K-QAM profiles along with one mixed profile (1K-QAM with 256-QAM) for potential roll-off issues.
 - o **Note:** CMs will only use these profiles if needed and will automatically go in and out of upgrade/downgrade.

Figure 7 is a sample output from a D3.1 CM showing current subcarrier RxMER readings. It shows the percent of subcarriers that would be able to run a given modulation. **Note:** This is not based on the D3.1 threshold table.

```

CM> /cm_hal/ofdm_analyzer 32 0
-----|-----|-----
RxMER   | Max   | RxMER Histogram for 2749 Subcarriers
dB       | Bitload |
-----|-----|-----10-----20-----30-----40-----50-----60-
30 dB | 1K-QAM |
31 dB |         | *
32 dB | 2K-QAM | *
33 dB |         | *
34 dB |         | *
35 dB | 4K-QAM | *
36 dB |         | *
37 dB |         | ***
38 dB | 8K-QAM | *****
39 dB |         | *****
40 dB |         | *****
41 dB | 16K-QAM | *****
42 dB |         | *****
43 dB |         | *

```

Figure 7 - Bit Loading Information from D3.1 CM

5. DS Resiliency / Partial Mode

The D3.1 OFDM PLC uses 16-QAM and is very robust, but it's still a best practice to find "clean" spectrum to avoid potential issues. If the OFDM channel is used as the primary channel, losing the PLC would be catastrophic. Cross-bonding with SC-QAM signals may be in our best interest to provide higher peak speeds and have a means for the CM to enter partial mode.

When a D3.1 CM reports via a cm-status message a non-primary (secondary) RF channel impairment for SC-QAM or OFDM, the following things happen on a Cisco CMTS:

- The CM is marked `p-online` for easy identification.
- If the RF channel impairment is below the configured DS resiliency thresholds, the D3.1 CM's service flows are moved to a resilient bonding group (RBG) or a narrowband (NB) interface (primary DS).
- If the RF channel impairment exceeds the configured DS resiliency thresholds, the impaired RF channel is temporarily de-activated from all the bonding groups (BGs).

The CM can move in and out of partial mode automatically. This feature is independent of the graceful profile management feature.

6. Capacity

Another tool in the toolbox is the use of Powerboost™ and the peak-rate TLV. Comcast trademarked the name, but it is a simple manipulation of the DOCSIS standard for CM rate limiting. This can be exploited to satisfy speed tests without over-provisioning by the typical 10%.

Over-provisioning is typically done to satisfy speed tests and the inherent difference in reporting at Layer 2 of the Open System Interconnection (OSI) model vs Layer 3. The CM and CMTS are reporting at Layer 2, but speed tests are reporting at Layer 3. The Layer 2 Ethernet overhead of 18 bytes per frame is not being counted, leading to a misconception on the end-users part of lower speed.

Figure 8 displays the amount of peak speed achieved while only over-provisioning this CM by 2%. The customer was sold a 500 Mbps service. The CM file was configured with 510 Mbps max rate, 600 Mbps peak rate, and 70 MB DS max burst.




Figure 8 – Powerboost Example

This achieved approximately 6 seconds of 600 Mbps peak speed. The DS max burst and peak rate in the CM file could be manipulated even further to achieve longer times and/or peak rates.

7. OFDM Settings to Maximize Speeds

To maximize D3.1 speeds, the following settings are suggested (example values are for a Cisco CMTS):

- `cyclic-prefix 192` - Make this value as low as the HFC plant will support. 192 is lowest while 1024 is the default.
- `pilot-scaling 48` - Keep at the lowest setting, which is the default of 48.
- `roll-off 128` - Make as large as possible but must be less than the cyclic prefix value.
- `subcarrier-spacing 25KHZ` - Less overhead for 25 kHz vs 50 kHz.
- `profile-control modulation-default 256-QAM` - Configure 4k, 2k, & 1k-QAM data profiles with maybe one mixed profile.
- `profile-ncp modulation-default 64-QAM` - Make the next codeword pointer (NCP) as high as the plant supports.
- `guardband-override 1000000` - OFDM guardband override in 50 kHz increments.

The cyclic prefix and roll-off settings are likely the parameters that will need to be adjusted based on individual HFC plant characteristic as well as the OFDM channel width. As mentioned previously, the roll-off is required to be less than the cyclic prefix per the D3.1 PHY spec. A large roll-off value is desirable as this decreases the guardbands (taper regions) providing more spectrum for active subcarriers. However, to have a larger roll-off, a larger cyclic prefix setting is needed which leads to additional overhead on every subcarrier. It is not possible to optimize both of these settings.

In general, for smaller OFDM channel widths, more speed will be achieved by maximizing the number of subcarriers with a higher roll-off setting even though it creates more overhead per subcarrier with a higher cyclic prefix.

The Cisco CMTS also has a D3.1 DS guardband override feature. Remember, the actual DS is 204.8 MHz and the CMTS will null out a bunch (according to spec) to make 192 MHz, then there is a taper region for maybe 188-190 MHz of actual usable bandwidth. It was decided that the channel width (actual spectrum) would not exceed 192 MHz and allow the guardband to be overridden from the default settings which are based on cyclic prefix and guardband settings. Some settings can make the guardband 1.5 to 2 MHz on each end. If you have no adjacent carriers or even if another OFDM block adjacent, you may get away with less guardband without interfering with each other. The guardband must be symmetrical on the left and right.

You can override the default behavior where the guardbands are based on the roll-off setting. The larger the roll-off setting, the smaller the guardbands needed and therefore more of the OFDM channel can be used for transmitting data. The roll-off must be less than the cyclic prefix per the PHY specifications. A lower cyclic prefix is desirable as it reduces overhead and allows the OFDM channel to run faster. Your HFC plant needs to have minimal micro-reflections to support this lower cyclic prefix. Using a low cyclic prefix means you have to use a lower roll-off which will create larger guardbands. You end up with two variables that can't both be optimized. The override feature allows you to use lower cyclic prefix and roll-off but then set the guardbands manually.

The whole reason Cisco defaulted the guardband size based on roll-off is because that is recommended in the PHY specifications as to prevent adjacent channel interference. The values were picked based on testing. If you start overriding these values, you can cause adjacent channel interference. It's not

recommended to change the guardband setting from the default unless the customer understands the consequences.

The maximum an active OFDM channel can be is 190 MHz so you will be limited to 1 MHz guardbands if configuring a full 192 MHz channel. If you configure a smaller channel size, you can use no guardband – but again look out for adjacent channel interference.

These settings will achieve the highest speeds, but some settings may need to be adjusted for real plant settings. Table 3 shows the speeds that can be achieved with various settings.

Table 3 - OFDM DS Speed Estimates (25 kHz)

Channels	Spectrum	DOCSIS 3.0	DOCSIS 3.1 (25 kHz subcarrier)		
		256 QAM	1024 QAM	2048 QAM	4096 QAM
4 channel	24 MHz	151 Mbps	172 Mbps*	189 Mbps*	206 Mbps*
8 channel	48 MHz	302 Mbps	373 Mbps*	410 Mbps*	448 Mbps*
16 channel	96 MHz	603 Mbps	776 Mbps*	853 Mbps*	931 Mbps*
24 channel	144 MHz	905 Mbps	1178 Mbps*	1296 Mbps*	1414 Mbps*
32 channel	192 MHz	1206 Mbps	1584 Mbps**	1742 Mbps**	1910 Mbps**
	2x192 MHz		3168 Mbps**	3484 Mbps**	3802 Mbps**

* **25 kHz subcarriers**, running same modulation, 1.175 MHz guardbands, roll-off 192, cyclic prefix 256, 2 x NCP (64-QAM)

** **25 kHz subcarriers**, running same modulation, 1.725 MHz guardbands, roll-off 128, cyclic prefix 192, 2 x NCP (64-QAM). The red line in table 3 indicates the channel width inflexion point where more speed is achieved via manipulation of roll-off and cyclic prefix. The guardband could be reduced to 1 MHz as well for even more speed.

8. General Thoughts for D3.1 Upstream (US)

I would be remiss to not mention D3.1 US at least in brief. D3.1 US uses OFDMA. Even though this feature is not deployed in production networks to date, lab testing is on-going and some observations have been made.

Most cable plants are limited to 5 MHz to 42 MHz with some upgrades happening to expand to 85 MHz. Because of the limited spectrum and the need to still provide services to legacy CMs, D3.1 US is not being heavily pursued. Until higher US speeds > 50 Mbps are required, US expansion to 85 MHz or even 204 MHz may be stalled. Part of the D3.1 spec allows time and spectrum sharing between SC-QAM US and OFDMA, which could allow higher peak speeds for D3.1 modems at the expense of more overhead.

For the interim, it's this author's belief that this time sharing also known as time and frequency-division multiple access (TaFDMA) may not be utilized at first and cross-bonding of four advanced time-division multiple access (ATDMA) with an OFDMA channel in an 85 MHz plant would suffice. Even with only one SC-QAM signal bonded with the OFDMA signal provides benefits as listed here:

- The CM would still have T4 multiplier of 2 allowing RF tech maintenance of up to $30 \times 2 = 1$ minute without losing DS lock from a T4 timeout.
- The CM would be more resilient in that it would have at least two channels doing station maintenance (SM).
- D3.1 US power is based on a formula for spectrum used and not just how many chs are in the US bonding group. So, max Tx doesn't change much when adding a small SC-QAM.
- Having a SC-QAM signal with its own scheduling may be simpler for VoIP and other scheduled flows that are latency and jitter sensitive.

One major observation made when deploying D3.1 CMs was how the US Tx level is being reported. The US max Tx level is 65 dBmV and translates to 53 dBmV for same bandwidth as four 6.4 MHz ATDMA chs. This value is 51 dBmV/ch for D3.0 without the Cablelabs' engineering change notice (ECN) for extended power. The D3.1 CM reports its US Tx level based on 1.6 MHz of bandwidth, which leads to a 6 dB correction factor compared to a 6.4 ATDMA channel.

Note: When replacing a D3.0 CM with a D3.1 CM (even in D3.0 mode) in the same location, you may notice the US Tx level reports 6 dB lower! This is normal and expected.

More data and testing needs to be gathered for US resiliency and how modems are assigned different interval usage code (IUCs). This means modems could use different modulation depending on the CMTS upstream RxMER for the respective subcarriers. Depending on the number of IUCs supported, this may negate the usage of lower spectrum with lower modulation so as to save those IUCs for modems with poor performance and nothing to do with lower spectrum.

Conclusion / Summary

DOCSIS 3.1 is being deployed today with great success; even better than expected. With features such as graceful profile management and resiliency, deploying higher modulation schemes has become more realistic. With the introduction of remote PHY architectures, expected performance increase in RxMER, and potentially less amplifier cascades, even higher modulation schemes will be possible for the DS and

US. When even higher speeds are required for the upstream, operators have a choice to change the duplex filter split to 204MHz/254 MHz and stay with D3.1 CMs or upgrade to FDX with compatible CMs. The FDX spec is under development at Cablelabs as of this writing.

In regards to on-going monitoring of plant performance, one can utilize the CM full bandwidth capture (FBC) for DS “sweepless sweep” and ingress testing and verification. One could activate CMTS RF and use that for test signals. Other aspects of PNM can also be utilized such as US spectrum viewing at the CM and CMTS along with US pre-EQ information. In the case of remote PHY, that information will be gathered at the remote PHY device (RPD). One could also use test equipment with a built-in CM to balance amplifiers by looking at CM Tx levels as they work their way downstream from the node/RPD. This may be necessary until US sweep is supported with a distributed access architecture (DAA) system.

Everyone should be aware that correctable forward error correction (FEC) errors with D3.1 will show very high, if not 100%! This is expected when using OFDM with LDPC and such large block sizes. Even though there are potentially thousands of subcarriers, they are all processed with the FFT.

Note: There is no special CM file needed for D3.1, but the CM must be in multiple transit channel (MTC) mode also referred to as US bonding. Sometimes US level issues could force the CM to non US bonding and the CM will not register properly and may not even downgrade to D3.0 or 2.0 mode. Always keep in mind that performance and features supported can vary dramatically with CM firmware.

Fiber deep architectures and remote PHY will allow much higher speeds via the higher order modulation schemes supported with D3.1. Node plus 0 could also facilitate an easier transition to FDX. Coax attenuation is our biggest hurdle at higher frequencies, so limiting the coax to eventually drop cable only (think fiber-to-the-tap) is a proposition that still works in our favor.

By utilizing CMTS features for robustness and “self-healing”, we can successfully operationalize these more complex architectures, multiplexing technologies and modulation schemes. Some of these features include: utilization load balance (2.0 & 3.0), US & DS resiliency, dynamic modulation, graceful profile management, etc.

Future SDN of OFDM profile management may not be as critical as first thought, but it will give us even more granularity when needed or justified. Be sure to utilize PNM to be more proactive in the monitoring of your customers’ quality of experience (QoE).

Abbreviations

4G	4 th generation mobile technology (LTE)
AGC	automatic gain control
ALC	automatic level control
ATDMA	advanced time-division multiple access
avg	average
B	bytes
b	bits
BG	bonding group
bps	bits per second
ch	channel

CM	cable modem
CMTS	cable modem termination system
coax	coaxial cable
corr	correctable
D1.x	DOCSIS 1.0 & 1.1
D2.0	DOCSIS 2.0
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
DAA	distributed access architecture
dB	decibel
DBC	dynamic bonding change
dBmV	decibel millivolt
DOCSIS	Data-Over-Cable Service Interface Specifications
DRFI	Downstream Radio Frequency Interface
DS	downstream
ECN	engineering change notice
FBC	full bandwidth capture
FDX	full duplex DOCSIS
FEC	forward error correction
FFT	fast Fourier transform
freq	frequency
FTTH	fiber-to-the-home
Gbps	gigabits per second
HFC	hybrid fiber-coax
Hz	hertz
IP	internet protocol
IUC	interval usage code
LDPC	low density parity check
LLQ	low latency queue
LTE	long-term evolution (4G)
Mbps	megabits per second
MER	modulation error ratio
MHz	megahertz
MoCA	Multimedia over Coax Alliance
MTC	multiple transit channel = US bonding
NB	narrowband
NCP	next codeword pointer
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OSI	Open Systems Interconnection
PLC	physical layer link channel
PNM	proactive network maintenance
QAM	quadrature amplitude modulation
QoE	quality of experience
RBG	resilient bonding group
RF	radio frequency

RPD	remote PHY device
RTP	Research Triangle Park, real-time transport protocol
Rx	receive, receiver
SC-QAM	single-carrier QAM
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined network
SG	service group
T4	timer 4 = 30 second CM timer for DS lock
TaFDMA	time and frequency division multiple access
TLV	type length value
Tx	transmit, transmitter
US	upstream
VoIP	voice over IP

Bibliography & References

Google Hangout - #1 DOCSIS Podcast (The only one)

<http://volpefirm.com/>

<https://plus.google.com/u/0/+Volpefirm/videos>

DOCSIS 3.1 Configurations for HFC and RFoG

A Technical Paper prepared for SCTE•ISBE by

Doug Jones
Principal Architect
CableLabs®
858 Coal Creek Circle
Louisville, CO 80027 USA
303-641-6563
d.jones@cablelabs.com

Introduction

The use of DOCSIS® 3.1 technology on a Radio Frequency over Glass (RFoG) network (SCTE 174) [1], requires special attention because the DOCSIS 3.1 upstream technology can increase the likelihood of optical beat interference (OBI) which can interrupt communication on the return path.

An RFoG network uses an RFoG optical network unit (R-ONU) with an upstream laser to transmit DOCSIS signals over the fiber network. OBI is caused on an RFoG network when two conditions are met: 1) two or more R-ONUs operate at substantially the same wavelength, and 2) those R-ONUs simultaneously transmit. Preliminary testing and analysis completed at CableLabs shows that the highly efficient DOCSIS 3.1 upstream can increase the occurrence of OBI because of how many cable modems (CMs) (e.g., R-ONUs) can simultaneously transmit.

The DOCSIS 3.1 upstream introduces a new technology called orthogonal frequency division multiple access (OFDMA) which supports multiple CMs transmitting at once in order to make more efficient use of upstream spectrum. CableLabs has observed in laboratory testing that both usability and efficiency of the DOCSIS 3.1 upstream is directly impacted by OBI. Further, observations have confirmed that the DOCSIS 3.1 downstream configuration requires no modification because of how RFoG technology operates; there is a single laser transmitter in the downstream therefore OBI will not be generated.

Content

1. A Brief History of RFoG

RFoG allows a cable system to use existing headend infrastructure and consumer-premises equipment while putting an all-fiber network solution in place for future services. Around 2005 there was a competitive push for fiber-to-the-home (FTTH) networks and RFoG was devised to allow an FTTH network option for cable operators.

The SCTE started the RFoG technology standardization process in 2007, with the introduction of the project authorization request (PAR) for a specification which eventually became SCTE 174. The work was undertaken in the SCTE interface practices subcommittee (IPS) working group 5 and resulted in a standard in 2010.

As things turned out, cable continued to innovate and to this day the cable network remains competitive. However, the cable industry found another reason to use RFoG technology, that being lower construction cost for rural environments as compared to a hybrid fiber/coax (HFC) network. In areas where the drop cable exceeded 500 feet or so, an amplifier would be needed at the start of the driveway just to get the signal to a subscriber's home making for relatively more expensive HFC network construction costs as compared to an RFoG network.

In today's market, cable is still looking toward extending fiber deeper as well as constructing greenfield FTTH networks. And there is continued discussion about the latest generation of DOCSIS 3.1 technology running over an RFoG network.

An example RFoG network is shown in Figure 1 where a traditional coax-based cable headend is fed into an outside plant that is all fiber optic cable, and then inside the home the signal is returned to coax cable. What allows this is the R-ONU in the home that converts between fiber optic cable and coax cable.

Figure 1 also shows multiple R-ONUs connecting to the fiber optic network. Traditionally the number of R-ONUs is up to 32 per network which is based on both the optical power levels and loss associated with a 32-way optical splitter in the FTTH network. An interesting consequence of inserting this FTTH network into the middle of what is otherwise a traditional cable network is the cable components on either side do not need to be aware of the RFoG technology. That is, the cable operator does not need to make adjustments in either the headend or in the home in order to use an RFoG network.

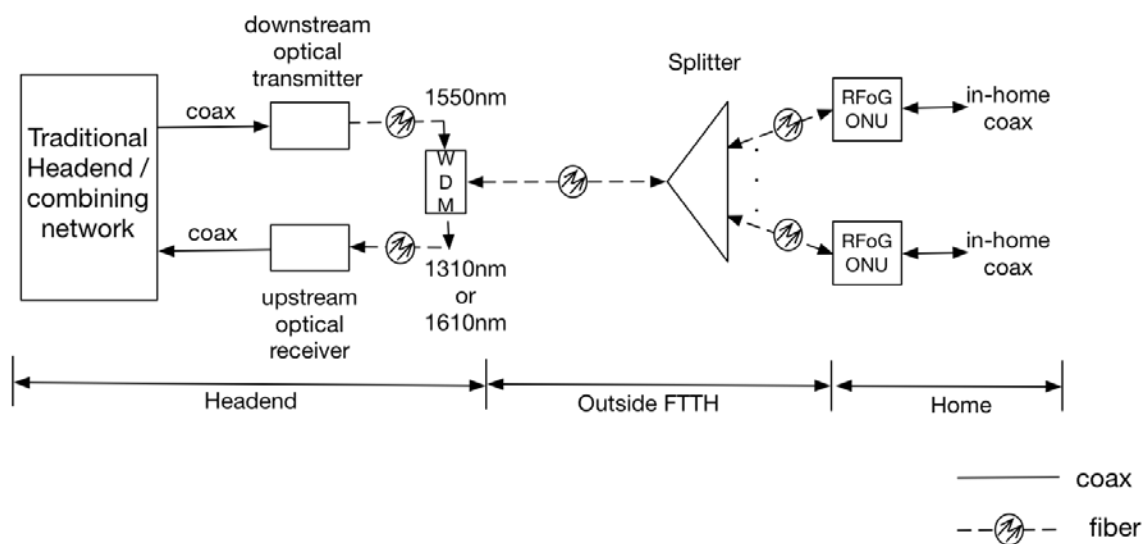


Figure 1 – Typical RFoG Network

RFoG technology is specified in SCTE 174 which places requirements on the R-ONU customer equipment. SCTE 174 includes an introduction to optical beat interference (OBI).

2. Optical Beat Interference

OBI is not specific to RFoG technology, and rather is a general optical phenomenon that occurs at a photodetector when two optical signals that overlap in wavelength are received. In this case, OBI is a beat note, i.e., a signal at the difference of those two optical wavelengths, and in an RFoG network has the effect of raising the upstream noise floor and impacting upstream communications.

In an RFoG network, OBI is an upstream phenomenon; OBI does not occur on the downstream. As shown in Figure 2, OBI is generated in the upstream optical receiver and appears on the coax that feeds the upstream combining network and can impact all upstream signals from any R-ONU transmitting at that time when the OBI is generated. With OBI, all that matters are the upstream wavelengths of the R-ONU lasers. The RF signals being generated in the home do not matter, just that at least two R-ONU return lasers are on at the same time and two of those optical wavelengths are close enough to generate

“optical beating” in the photodetector, the physical and mathematical basis of which will be addressed in the next section.




Figure 2 – OBI generation

If the optical signals simultaneously hitting the optical receiver are far enough apart in wavelength, there won't be beating. But if the optical signals are close enough in wavelength, beating occurs. Regular manufacturing tolerances of the upstream lasers can be enough to either have OBI or not have OBI on any particular RFoG network. Since the laser wavelength changes when the temperature changes, OBI and come and go throughout the day as temperatures change.

Upstream transmissions on cable plant can come from a couple of different sources, and any of these can cause the R-ONU optical laser to turn on. Common sources of return signals on the cable plant are:

- DOCSIS cable modems
- Legacy set-top boxes (e.g., SCTE 55-1 [2] and SCTE 55-2 [3])
- Legacy Cable Phone (e.g., Arris Cornerstone® Voice Port™ technology)

Any of these return RF signals on the coax can cause the R-ONU laser to activate to send an optical signal to the optical receiver.

3. Heterodyne and Beat Signals

Heterodynes and beats are related phenomenon and are technologies that have been beneficially used in communications networks, including cable TV, for decades. Beats and heterodynes are created by a frequency mixer, which is a nonlinear electrical circuit that creates new signals from the two signals applied to it. In its most common application, two signals are applied to a mixer and it produces new signals at the sum and difference of the original frequencies. Other frequency components may also be produced in a practical frequency mixer.

Heterodyning is a signal processing technique that creates new signals by combining two known signals in a mixer. Heterodyning can be a very useful technology and has been used for years in cable systems. Included in the references section is a 1967 paper [4] given at the NCTA conference about the use of heterodyning in CATV channel processing. Heterodyne production is generally very controlled in a precise circuit by applying both the input signal and a precisely controlled local oscillator signal to a mixer, and using filtering and amplification to produce the desired new signal.

Beats are also caused by applying two input signals to a mixer, however the process is generally not controlled. It is a naturally occurring phenomena related to the physical capabilities of the mixer. A photodiodes used in an optical receiver can act as and is what causes OBI in an RFoG network.

A schematic for a typical mixer is shown in Figure 3. The input to the mixer are two signals at different frequencies where sometimes frequency₂ is shown as a local oscillator (LO). The output signals are called heterodynes and consist of one or more signals at different frequencies than the input signals.




Figure 3 – Typical Mixer

Heterodyning within the mixer is based on the trigonometric identity:

$$\sin(A) \sin(B) = \frac{1}{2} \cos(A - B) - \frac{1}{2} \cos(A + B) \quad (1)$$

On the left-hand side of the equation are the two input signals at different frequencies, A and B, which are “mixed” to create the two new signals at frequencies (A-B) and (A+B), on the right-hand side of the equation. The new signals at frequencies (A-B) and (A+B) are called heterodynes.

Mixers vary and are classified by their functionality:

- An unbalanced mixer, in addition to producing the heterodyne signals, allows both input signals to pass through and appear at the output.
- A single-balanced mixer is designed such that either one or the other of the input signals is available at the output (but not both) as well as the heterodynes.
- A double-balanced mixer is designed such that neither of the input signals and only the heterodyne signals appear at the output.

Products using mixer and heterodyne technology are used for channel frequency conversions are generally carefully constructed, including detailed filters and amplifiers to protect against unwanted signals and only the desired signal is at the output.

Cable uses mixer / heterodyne technology, for example, in channel upconverters. A coaxial cable used by a cable television system can carry many television or QAM channels all at the same time because each channel is given a different frequency, i.e., EIA channel, so they don't interfere with one another. At the cable headend, upconverters move an incoming television channel (or QAM signal) to a different carrier frequency to fit within the channel plan on the coax. Channel upconverters do this by mixing the television signal frequency, f_{CH} with a local oscillator at a different frequency f_{LO} , creating a heterodyne at the sum $f_{CH} + f_{LO}$, which is combined onto the cable.

4. OBI as a Beat Noise

In RFoG networks the optical receiver, typically a photodiode, acts as an unbalanced mixer which can cause the creation of optical beat products when it gets hit simultaneously by two R-ONU lasers whose wavelengths are close enough to each other. The beats that cause OBI are a result of unintentional and uncontrolled heterodyning in the optical receiver that raises the noise floor of the return path which can wipe out the return path for as long as the two R-ONUs are transmitting together.

A schematic for a typical unbalanced mixer is shown in Figure 4.




Figure 4 – Unbalanced Mixer

In this case, the wavelengths of the optical signal are sourced from the R-ONUs and when converted to frequency are in the hundreds of Terahertz (THz), or a million times the MHz used on cable systems.

The equations to convert between frequency and wavelength are:

$$f = \frac{c}{\lambda} \quad (2)$$

$$\lambda = \frac{c}{f} \quad (3)$$

where:

f = the frequency of the signal in Hz.

λ = the wavelength of the signal in meters.

c = the speed of light in a vacuum (299,792,458 meters/second).

A signal with wavelength of 1610 nm (1610×10^{-9} meters) has an approximate wavelength of 186 THz.

The OBI is the difference product, or ($f_1 - f_2$). The sum product, ($f_1 + f_2$) will be in the THz range which will not cause interference in the MHz range. However, subtracting a THz from a THz, when those values are very close, can result in a number in the MHz, and that is the OBI signal that causes degradation to the RF signals. Since the OBI signal is at the output of the optical mixer, that noise is carried through the rest of the optical receiver circuitry and is present on the coax that goes to the upstream combining network.

5. The Physics of Photodiode OBI Generation

This section goes into the detail of deriving equation 1, and shows how a photodiode can create beats.

A diode can be used to create a simple unbalanced mixer that produces both the original frequencies and their sum and difference heterodynes. The derivation begins with the Shockley diode equation which shows the current through an ideal diode as a function of the voltage, V_D , across it. The Shockley diode equation is:

$$I = I_s \left(e^{\frac{V_D}{nV_T}} - 1 \right) \quad (4)$$

where:

I is the diode current.

I_s is the diode reverse bias saturation current (essentially a constant, depends on temperature).

V_D is the voltage across the diode.

V_T is the thermal voltage (essentially constant, depends on temperature).

n is the ideality factor of the diode (a constant).

e is a mathematical constant, sometimes known as Euler's number.

It is important to note that the current and voltage are not linearly related, rather, that the voltage is raised as an exponential to e meaning the diode is a nonlinear device. Voltage and current through a diode are not linear, rather, they are related exponentially.

When deriving how the beats are created, it's not important to know these values, rather the equation with be manipulated to show how a pair of input sine waves (signals) will result in new sine waves at the sum and different frequencies.

The exponential function of the Shockley diode equation can be expanded into a Taylor series as shown below:

$$e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!} \quad (5)$$

Which can be approximated by the first few terms of the series:

$$e^x = \left(\frac{x^0}{0!} \right) + \left(\frac{x^1}{1!} \right) + \left(\frac{x^2}{2!} \right) + \left(\frac{x^3}{3!} \right) + \dots \quad (6)$$

On the right-hand side of the equation, in the first term both numerator and denominator, x^0 and $0!$ are equal to one therefore the first term on the right-hand side of the equation is equal to one. Using this equation 6 can be simplified as:

$$e^x = 1 + x + \frac{x^2}{2} + \frac{x^3}{6} + \dots \quad (7)$$

Which can be rearranged as:

$$(e^x - 1) = x + \frac{x^2}{2} + \frac{x^3}{6} + \dots \quad (8)$$

Now the parenthetical on the left side of the equation is in the form of the parenthetical on right-hand side of the original Shockley diode equation (equation 4). From equation 4 ignoring the three constants I_S , V_T , and n , the original Shockley equation can be approximated as:

$$I = \left(x + \frac{x^2}{2} + \frac{x^3}{6} + \dots \right) \quad (9)$$

Now assume that diode is in a circuit with a resistor in series with it, and the current, I , will generate an output voltage, the output voltage will have the form:

$$v_0 = \left(x + \frac{x^2}{2} + \frac{x^3}{6} + \dots \right) \quad (10)$$

Now the derivation is in final form, and assume two input voltages, $v_1 + v_2$ (from the two R-ONUs), are applied to the diode. The output voltage can be approximated as:

$$v_0 = (v_1 + v_2) + \frac{1}{2}(v_1 + v_2)^2 + \frac{1}{6}(v_1 + v_2)^3 + \dots \quad (11)$$

If those two voltages are sinusoids of different frequencies:

$$v_1 = \sin(A) \quad (12)$$

$$v_2 = \sin(B) \quad (13)$$

Substituting these signals into the equation yields:

$$v_0 = (\sin(A) + \sin(B)) + \frac{1}{2}(\sin(A) + \sin(B))^2 + \frac{1}{6}(\sin(A) + \sin(B))^3 + \dots \quad (14)$$

To show the creation of new frequencies, the analysis will just look at the term on the right-hand of the equation that is raised to the second power.

$$(\sin(A) + \sin(B))^2 = \sin^2(A) + 2 \sin(A) \sin(B) + \sin^2(B) \quad (15)$$

To finish the analysis, notice that on the right hand of the equation, the middle term is essentially the same as equation 1 which is repeated here.

$$\sin(A) \sin(B) = \frac{1}{2} \cos(A - B) - \frac{1}{2} \cos(A + B) \quad (16)$$

Manipulating the Shockley diode equation shows how new frequencies can be created when two sine waves, the outputs of the R-ONUs, are mixed at the photodiode in the headend optical receiver. And based on equation 14 it can be seen that the creation of new frequencies is not “clean”, rather, multiple new frequencies can be created based on the squared, cubed, and the following terms raised to higher exponents that will generate multiple new frequencies which is the noise that is called OBI. Additionally, with RFoG technology it is possible to have more than two sine waves hitting the optical receiver. With upstream channel bonding in DOCSIS 3.0, a typical return path has up to four carriers, meaning four R-

ONUs can transmit at once. In the next section will be a discussion of the OFDMA upstream available in DOCSIS 3.1 and how even more R-ONUs can be transmitting at once. This result is not comparable to what happens in a channel upconverter, where complex circuitry makes for the precise generation of heterodynes. Rather, with RFoG, the heterodyning process within the photodiode is uncontrolled and just creates noise.

6. Lab Setup and Analysis

6.1. Lab Setup

The lab setup is diagramed in Figure 5 and includes an FTTH network with thirty-two R-ONUs, each with a single DOCSIS 3.1 CM attached. Various R-ONUs were used in the analysis; however, never more than thirty-two at once. For each R-ONU, the wavelength of the return laser was found using an optical spectrum analyzer and within the group of thirty-two R-ONUs there were several within close range of each other, capable of generating OBI.

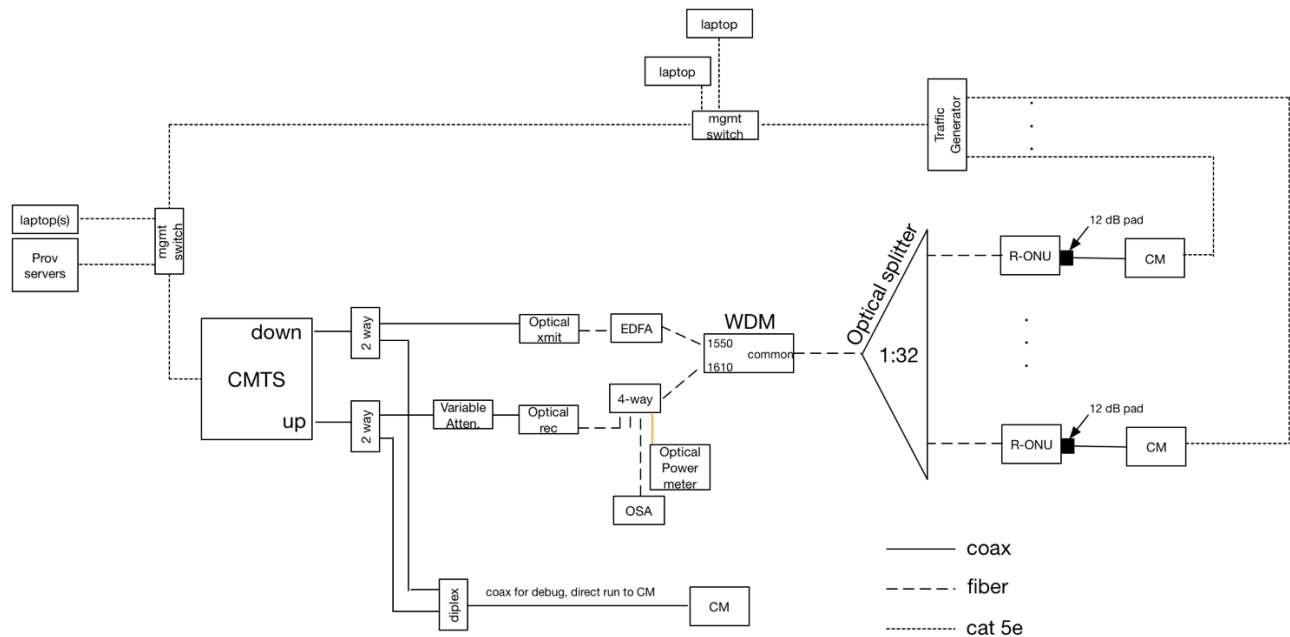


Figure 5 – RFoG Lab Setup

DOCSIS 3.1 CMs are capable of both OFDMA and ATDMA return. When in ATDMA mode, four carriers were used; each 6.4 MHz wide; modulated at 64 QAM; and centered at 17.3 MHz, 23.7 MHz, 30.1 MHz, and 36.5 MHz.

When in OFDMA mode, two channel widths were used; 10 MHz and 24 MHz. Modulations were varied between 64 QAM and 1024 QAM; however, most work was done with 64 QAM as the focus was on reliable communication so errors could be attributable to OBI.

Table 1 lists the nominal upstream laser wavelengths of the thirty-two R-ONUs used in the testing, from low to high wavelength

Table 1 – R-ONU Wavelengths

R-ONU	Laser Wavelength
1	1609.99
2	1610.71
3	1610.83
4	1610.95
5	1611.08
6	1611.34
7	1611.53
8	1611.55
9	1612.01
10	1612.40
11	1612.51
12	1612.56
13	1612.80
14	1612.83
15	1612.87
16	1612.89
17	1612.89
18	1613.10
19	1613.20
20	1613.30
21	1613.52
22	1613.60
23	1613.64
24	1613.72
25	1613.74
26	1613.85
27	1613.93
28	1614.11
29	1614.15
30	1614.98
31	1615.08
32	1615.58

Figure 6 graphically shows the distribution of the R-ONU upstream wavelengths.




Figure 6 – R-ONU Wavelengths

There is a distribution, due to normal manufacturing variation, of approximately 5.5 nm in the upstream lasers (1610 nm through 1615.58nm). The R-ONU wavelengths that are close together are candidates to cause OBI. Specifically, there is a group of five wavelengths that overlap at 1612.80 nm and 1612.89 nm, as well as a couple other groups of two wavelengths closely spaced together. Note that the wavelengths are not static; they can change with time and temperature. The creation of OBI is hard to predict and can be even harder to troubleshoot.

6.2. OBI and DOCSIS Analysis

When testing with DOCSIS equipment, the presence of OBI was tied to the occurrence of uncorrectable Forward Error Correction (FEC) codewords as reported by the CMTS. An uncorrectable codeword is one which cannot be corrected by the FEC; the assumption that OBI has caused so much noise on the return path that the codewords have been corrupted beyond the ability of the FEC to recover them. The traffic generator also reported packet loss; however, the lab analysis is based on uncorrectable codewords.

With DOCSIS 3.0 trying to induce OBI, specifically with 4 ATDMA carriers over a range of packet sizes and data rates, uncorrectable codewords ranged from 4 – 6% of the traffic under a wide range of scenarios.

With DOCSIS 3.1 trying to induce OBI, under specific conditions up to 80% uncorrectable codewords were observed. The worst case of OBI was observed when both the DOCSIS return channel and the traffic generator were purposely configured to maximize the possibility of multiple simultaneous upstream transmissions. The configuration consisted of OFDMA frames supporting 60 simultaneous minislots, the minislots large enough to carry a 64 Byte packet, and all modems sending upstream 64 Byte packets aligned on the same time interval. This is not a common traffic pattern, specifically aligning packet transmission on the same time interval.

Generally upper layer protocols can recover. Specifically, TCP transport enables re-transmissions of data. Hence, in the presence of a small percentage of data loss (uncorrectable codewords), there is a possibility that customers may not notice OBI. However, UDP will be affected, and Voice over IP (VoIP) phone calls use UDP transport. In the presence of OBI customers may hear drop-outs or stuttering on their VoIP calls. Also, some gaming applications can use UDP transport and OBI would impact performance of gaming.

6.3. Inducing OBI

Figure 7 shows a lab setup that was used to induce OBI. Note there are no CMs or CMTS needed to induce OBI which is just a result of closely spaced optical wavelengths hitting a photodiode on the upstream portion of the optical path.

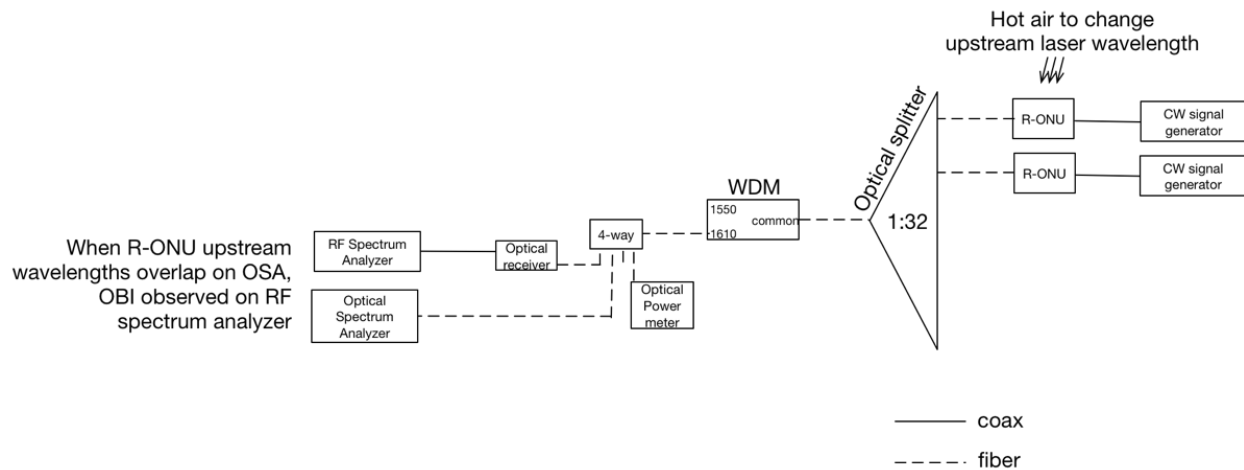


Figure 7 – OBI Lab Setup

In the OBI lab setup, just two R-ONUs were used and were chosen to have upstream laser wavelengths that were close together. A hair dryer was used to change the temperature of one of the R-ONUs to make its laser change wavelength and overlap with the other R-ONU. When the wavelengths overlap, then OBI is created in the optical receiver and can be observed on the RF spectrum analyzer. The results were captured on both an RF spectrum analyzer and an optical spectrum analyzer.

Figure 8 shows both RF and optical spectrum analyzer screenshots before the R-ONU wavelengths were induced to overlap. On the RF spectrum analyzer, one R-ONU has a CW carrier at 10 MHz and the other R-ONU has a CW carrier at 40 MHz. Note that OBI has nothing to do with the RF frequency on the coax, just the overlapping wavelengths at the optical receiver. In this figure, no OBI is observed in on the RF spectrum analyzer. On the optical spectrum analyzer, note how close the two wavelengths are; however, that they do not overlap.




Figure 8 – Wavelengths close, no OBI

Figure 9 shows that same setup a few seconds later as the wavelengths of the R-ONUs overlap and OBI is created at the optical receiver. On the RF spectrum analyzer, the peaks at 10 MHz and 40MHz are still visible, however, the noise floor has risen approximately 25 dB.

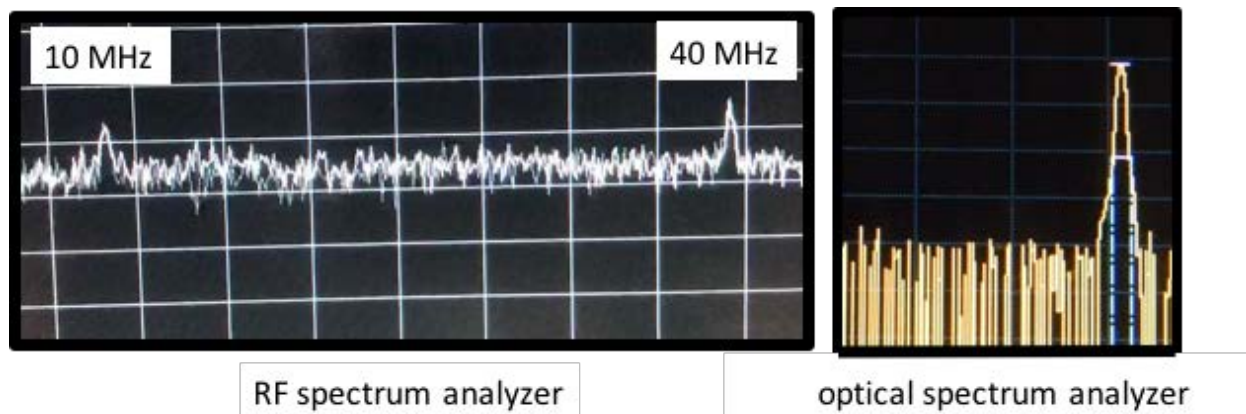


Figure 9 – Wavelengths overlap, OBI present

Note the broad spectrum of the noise caused by the OBI which impacts the entire return path (in this case, 5 – 42 MHz). It only takes two simultaneously transmitting R-ONUs to generate OBI at the optical receiver; however, data would be lost from any other CMs transmitting at the same time even though those R-ONU wavelengths are far enough apart to otherwise not cause OBI.

6.4. Aggregate Optical Receive Power

An issue that may become important with DOCSIS 3.1 upstream over an RFoG networks is the aggregate transmit power if multiple R-ONUs are simultaneously operating. Using log math, and assuming each R-

ONU outputs +3 dBm (several output levels are allowed, check SCTE 174), then the aggregate power of multiple transmitting R-ONUs is shown in Table 2.

Table 2 – R-ONU Aggregate Transmit Power

Number of R-ONU's transmitting	Aggregate Power
1	3 dBm
2	6 dBm
4	9 dBm
8	12 dBm
16	15 dBm
32	18 dBm

On a typical North American return path (5 – 42 MHz) with four DOCSIS 3.0 carriers and the possibility of 4 R-ONUs simultaneously transmitting, the RFoG network had to be designed assuming an aggregate transmit power of +9 dBm. With a DOCSIS 3.1 upstream where it is possible that more than four R-ONUs are simultaneously transmitting, the aggregate power can be higher. The network design should take this into account to ensure the aggregate optical transmit power from multiple simultaneously transmitting R-ONUs is not overloading the optical receiver.

7. DOCSIS and Simultaneous CM Transmissions

As discussed in previous sections, OBI is caused by two (or more) R-ONUs simultaneously transmitting at essentially the same wavelength. Two R-ONUs can transmit at the same time when multiple two-way services are operating. DOCSIS technology has supported multiple upstream carriers since the DOCSIS 1.0 specification, and in addition, there can be both legacy set-top box return channels and legacy cable phone systems. These various upstream services can lead to two R-ONUs transmitting at the same time.

Cable data service has been evolving, and it has been the continued increase of cable data speeds over the last few years that has really driven the deployment of multiple upstream DOCSIS channels. DOCSIS 3.0 technology introduced upstream channel bonding and drove wide adoption of multiple upstream DOCSIS channels. The first DOCSIS 3.0 modem deployments in North America began in early 2008, about the same time as when the SCTE effort to standardize RFoG technology was getting underway.

While DOCSIS technology has allowed multiple upstream carriers for years, there is still a limit of one CM transmitting at a time on a DOCSIS carrier. For example, according to the DOCSIS 3.0 MULPI specification [5], each DOCSIS 3.0 upstream carrier has an associated Upstream Bandwidth Allocation Map (MAP) that describes all upstream transmission opportunities on that carrier. The MAP message allows only one CM to transmit at a time on a carrier.

The first DOCSIS 3.0 modem deployments in North America began in early 2008, about the same time as when the SCTE effort to specify RFoG was gaining steam. The first DOCSIS 3.0 CM certifications happened at CableLabs Wave 58. Note that a DOCSIS 3.0 CMTS can schedule a single CM to transmit on more than one DOCSIS carrier at a time.

The DOCSIS 3.1 specification [6] introduces orthogonal frequency division multiple access (OFDMA) technology on the upstream, and the number of modems that can simultaneously transmit can increase

dramatically. In lab testing with 32 CMs on an optical splitter, it was possible to get over 20 modems (and their associated R-ONUs) to simultaneously transmit though the exact number depends on the CMTS scheduling algorithm. OFDMA as used in DOCSIS 3.1 is fundamentally different as compared to upstream QAM used DOCSIS 3.0 and earlier, as OFDMA technology is designed for even more efficient usage of the upstream spectrum by having multiple CMs transmit at the same time.

In DOCSIS 3.1, the OFDMA upstream is comprised of a continuous progression of OFDMA frames and each OFDMA frame contains a number of simultaneous minislots that is based on the amount of spectrum allocated to the OFDMA carrier.

As an example, in DOCSIS 3.1 when using a 2048 fast Fourier transform (FFT) size and the minimum of 10 MHz of spectrum, there will be 25 simultaneous minislots. This can be derived by understanding that the 2048 FFT size uses a subcarrier spacing of 50 kHz, meaning there are $10 \text{ MHz} \div 50 \text{ kHz}$ per subcarrier = 200 subcarriers. Per the DOCSIS 3.1 MULPI specification [7], there are 8 subcarriers per minislot, yielding 25 simultaneous minislots. In DOCSIS, a minislot is the smallest unit of upstream bandwidth allocation. The CMTS assigns minislots for CMs to transmit in.

The capacity of a minislot, how many Bytes of information it can carry, is based on both how many OFDMA symbols are in an OFDMA frame and the modulation of that symbol.

A typical OFDMA configuration used in the lab analysis is shown in Figure 10 and uses 24 MHz of OFDMA upstream, which at a 2048 FFT size (50 kHz subcarrier spacing), results in 480 subcarriers available. At 8 subcarriers per minislot, this configuration supports 60 simultaneous minislots.




Figure 10 – DOCSIS 3.1 Upstream OFDMA Frame Structure

Based on an OFDMA frame width of 16 OFDMA symbols per frame and each subcarrier operating at 64 QAM, each minislot was capable of carrying approximately 77 Bytes of data. When testing with upstream 64 Byte packets and including DOCSIS overhead, then one TCP ACK would easily fit within

two minislots. Thus, depending on CMTS scheduling and OFDMA overhead, up to 30 CMs (and their associated R-ONUs) could be transmitting at the same time.

In summary DOCSIS 3.0 upstream channel bonding allowed several CMs to simultaneously transmit on the upstream, up to one modem per carrier. On a typical North American cable network with four upstream carriers, there can be up to four CMs simultaneously transmitting. Depending on configuration, DOCSIS 3.1 allows even more CMs to simultaneously transmit on the upstream, depending on traffic loading. The bottom line is that when a DOCSIS 3.1 upstream is operated on an RFoG network there is a higher probability of OBI.

8. Managing OBI

If OBI is present, it will cause data transmission errors on the return path for as long as the two (or more) offending R-ONUs are transmitting simultaneously. All it takes is two R-ONUs transmitting at the same time to cause OBI that will wipe out the signals from any other R-ONUs which happen to be transmitting at that same time.

OBI only happens when two conditions are met; those being at least two R-ONUs at substantially the same wavelength and transmitting at the same time. The ONUs used for this testing were all specified to transmit at 1610 nm, however, there was still enough variation between most of those R-ONU upstream laser wavelengths that they would not cause OBI. It was only when certain pairs of R-ONUs were transmitting that OBI was observed.

An observation then is that OBI may not occur on an RFoG network if the natural variation of R-ONU upstream laser wavelengths are spaced appropriately. However, the next group of R-ONUs may have one or more pairs of R-ONUs that do transmit at a substantially similar wavelength and OBI can occur. This can make OBI a difficult issue to track down. OBI may not happen on all RFoG networks and when it does occur OBI is transient in nature, occurring only when the right two R-ONUs are transmitting at the same time.

Assuming the RFoG network is designed for thirty-two R-ONUs (due to optical loss on the fiber network), having fewer R-ONUs on that network can also lower the chance that there will be two R-ONUs using the same frequency. There are just fewer R-ONUs on that network so less of a chance that two are at the same wavelength. However, running the network below capacity could increase the overall cost of that installation.

Likewise, there are cable architectures that include multiple CMs in the home. In this case there is a higher chance that the R-ONU would be transmitting because of more return path transmitters (CMs) in the home. And in the case of DOCSIS 3.0 or DOCSIS 3.1 technology, there can be multiple simultaneous transmission opportunities which if two R-ONUs are on the same frequency could lead to a higher likelihood of OBI on that network.

If OBI is present, mitigation techniques typically fall into two categories; solutions in the optical domain and solutions in the DOCSIS domain.

Solutions in the optical domain rely on techniques to manipulate the wavelengths of the R-ONUs that hit the optical receiver. If overlapping wavelengths can be minimized or eliminated, then the possibility of OBI can be reduced.

Solutions in the DOCSIS domain typically focus on reducing either the number of CMs that can transmit simultaneously, or on recognizing groups of CMs (i.e., R-ONUs) that cause upstream errors when they transmit at the same time, and not scheduling these modems to simultaneously transmit.

These two different types of solutions have associated pluses and minuses. This paper does not attempt to an in-depth analysis of the various solutions on the market. An operator planning to operate a DOCSIS 3.1 upstream over an RFoG network should consider looking into solutions in the case that OBI is observed on that network.

Conclusion

DOCSIS 3.1 upstream technology can increase the probability that OBI will occur on an RFoG network. OBI will not occur on all networks. However, the current direction of the industry of considering RFoG networks for new build opportunities, deploying more upstream transmitters (e.g., DOCSIS CMs) in a home, and migrating to the DOCSIS 3.1 upstream can all increase the probability of OBI. Understanding how OBI occurs is a first step in considering alternatives to mitigate the effects of OBI, and researching appropriate mitigation techniques.

Abbreviations

ATDMA	advanced time division multiple access
CM	cable modem
CMTS	cable modem termination system
CW	continuous wave
dB	decibel
dBm	decibel milliwatt
dBmV	decibel millivolt
DOCSIS	data over cable service interface specifications
EIA	electronic industries alliance
FEC	forward error correction
FFT	fast Fourier transform
FTTH	fiber to the home
HFC	hybrid fiber/coax
Hz	hertz
IPS	interface practices subcommittee
ISBE	International Society of Broadband Experts
ITU	International Telecommunications Union
LO	local oscillator
MAP	DOCSIS upstream bandwidth allocation map
MHz	megahertz
nm	nanometer
OBI	optical beat interference
OFDMA	orthogonal frequency division multiple access
PAR	project authorization request
QAM	quadrature amplitude modulation

RF	radio frequency
RFoG	radio frequency over glass
R-ONU	RFoG optical networking unit
SCDMA	synchronous code division multiple access
SCTE	Society of Cable Telecommunications Engineers
THz	terahertz
TV	television

Bibliography & References

- [1] Radio Frequency over Glass Fiber-to-the-Home Specification, ANSI/SCTE 174 2010
- [2] Digital Broadband Delivery System: Out of Band Transport Part 1: Mode A, ANSI/SCTE 55-1 2009
- [3] Digital Broadband Delivery System: Out of Band Transport Part 2: Mode B, ANSI/SCTE 55-2 2008
- [4] *Comparison Of Demodulator-Modulator Versus Heterodyne Signal Processing For CATV Head Ends*, G. Rogeness, Proceedings of the NCTA, 1967
- [5] Data-Over-Cable Service Interface Specifications; DOCSIS 3.0 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0-I30-170111
- [6] Data-Over-Cable Service Interface Specifications; DOCSIS 3.1 Physical Layer Specification, CM-SP-PHYv3.1-I10-170111
- [7] Data-Over-Cable Service Interface Specifications; DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I10-170111

Guidelines for Cable Facility Climate Technology Optimization

Cooling Optimization for Edge Facilities

An Operational Practice Prepared for SCTE•ISBE by

John Dolan

Senior Guideline Specialist
Rogers Communications Inc.
8200 Dixie Road, Brampton ON CA L6T 0C1
519-852-5666
john.dolan@rci.rogers.com

Daniel Howard

Director, Consulting Services
Hitachi Consulting
2512 Parkdale Place NE
404-625-1593
daniel.howard@hitachiconsulting.com

Arnold Murphy

President
SCTi
3476 Galetta Road,
Arnprior, ON CA K7S 3G7
613-558-4415
a.murphy@sct-inc.com

Ken Nickel

Executive Vice President
Quest Controls, Inc.
870 Emerald Bay Rd., Suite 307
South Lake Tahoe, CA 96150
530-600-4570
knickel@questcontrols.com

Dave Smargon

COO
AIRSYS North America
915 De La Vina Street,
Santa Barbara, CA 93101
805-312-7549
dave.smargon@air-sys.com

Introduction

The SCTE ISBE 2020 Program identifies Edge Facilities (Class D as defined by SCTE 226) as by far the largest contributor to an MSO's energy consumption. This is because they significantly outnumber Data Centers and Regional headends. See Figure 1.

The focus of this paper will be on the *edge facilities* and it will provide insight into creating an overall guideline for facility climate technology optimization.




Figure 1 - Power at the Edge

The main elements to be discussed for a guideline are:

- How to manage air flow
- How to create an energy management plan and process
- What data to gather
- How to develop a baseline.
- How to consider the various cooling solutions
- How to perform measurement and validation
- How to pursue rebates
- How to evaluate the return on investment
- How to consider new or improved cooling technologies.

Following a guideline for *edge facilities* can reduce alarms and/or downtime, lower energy consumption and cost, improve energy efficiency, and improve the power margin. This is intended for cable operator local engineering and operations personnel, as well as corporate sustainability teams.

The paper was a team effort not only from the 5 principle authors but I would like to recognize the significant contributions from the following people:

Curtis Stiles, Time Warner Cable; **George Gosko**, Hitachi Consulting; **Supriya Dharkar**, Hitachi Consulting; **Jake Yu**, AIRSYS North America; and **Ed Kaye**, AIRSYS North America.

Guidelines

1. Developing a Cooling Energy Management Plan

A cooling energy management plan is a process with the final goal to reduce the cost to cool a watt of information technology (IT) equipment load. Under the assumption that many current edge facilities are overcooled and have undesirable mixing of hot air with cold air, a cooling energy management plan to reduce mixing, focus the cooling on the IT equipment, and enhance heating, ventilation and air conditioning (HVAC) equipment and controls can lead to significant energy and operational expenditure (Opex) savings as well as improvements in facility robustness.

Ultimately this plan should be applied to an entire portfolio of edge facilities to have the greatest impact on reducing the cost of cooling, and should involve common measures that can cost-effectively be applied at scale. Changes that improve the cooling effectiveness and efficiency at facilities in the portfolio should be identified in the plan, implemented, and measured to demonstrate energy consumption and cooling effectiveness before and after the implementation should be made. This will help validate the improvements and reductions of cooling costs, especially when energy incentives are sought.

At a minimum, the cooling energy management plan must:

- Categorize and characterize sites (*edge facilities*)
- Follow a standard methodology
- Clearly define metrics to evaluate solutions and validate them as successful
- Identify areas of concern
- Incorporate all rebates/incentives where available

The plan is essential in proving the changes improved energy efficiency in a manner that is compliant with both standards as well as the requirements of the energy incentives.

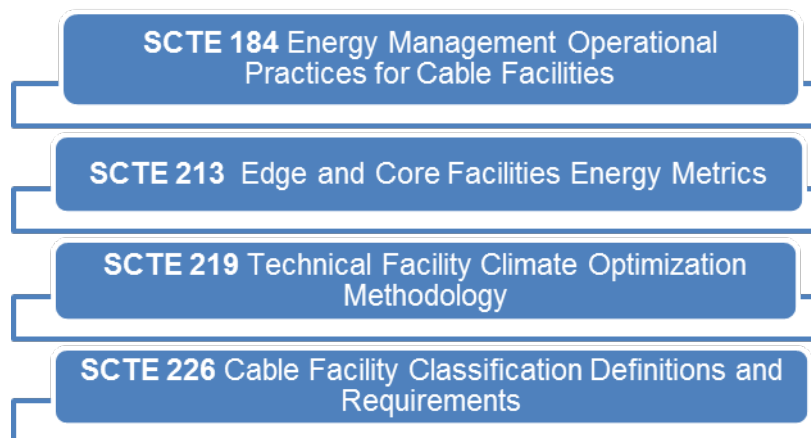


Figure 2 - Relevant SC TE Energy Standards for Facility Climate Technology Optimization

2. Process, Standard Methodology

An Energy Management Plan is a process. Figure 3 provides an example of a standard methodology to follow to implement a plan.

First, there is the overall **plan** for site or sites. The **baseline** data must be collected for each site so that there is clear data to compare any changes to. Then, the **measurement** of changes begins but a **remediation** point may be reached where there are issue occurring that are affecting the data, such as an IT refresh, that will skew results. This allows time for changes to be made to the plan. Once the site is stable the result on climate of the changes can be accurately measured and the **report** is made. Finally the analysis will determine if the changes were effective in meeting the climate goals.

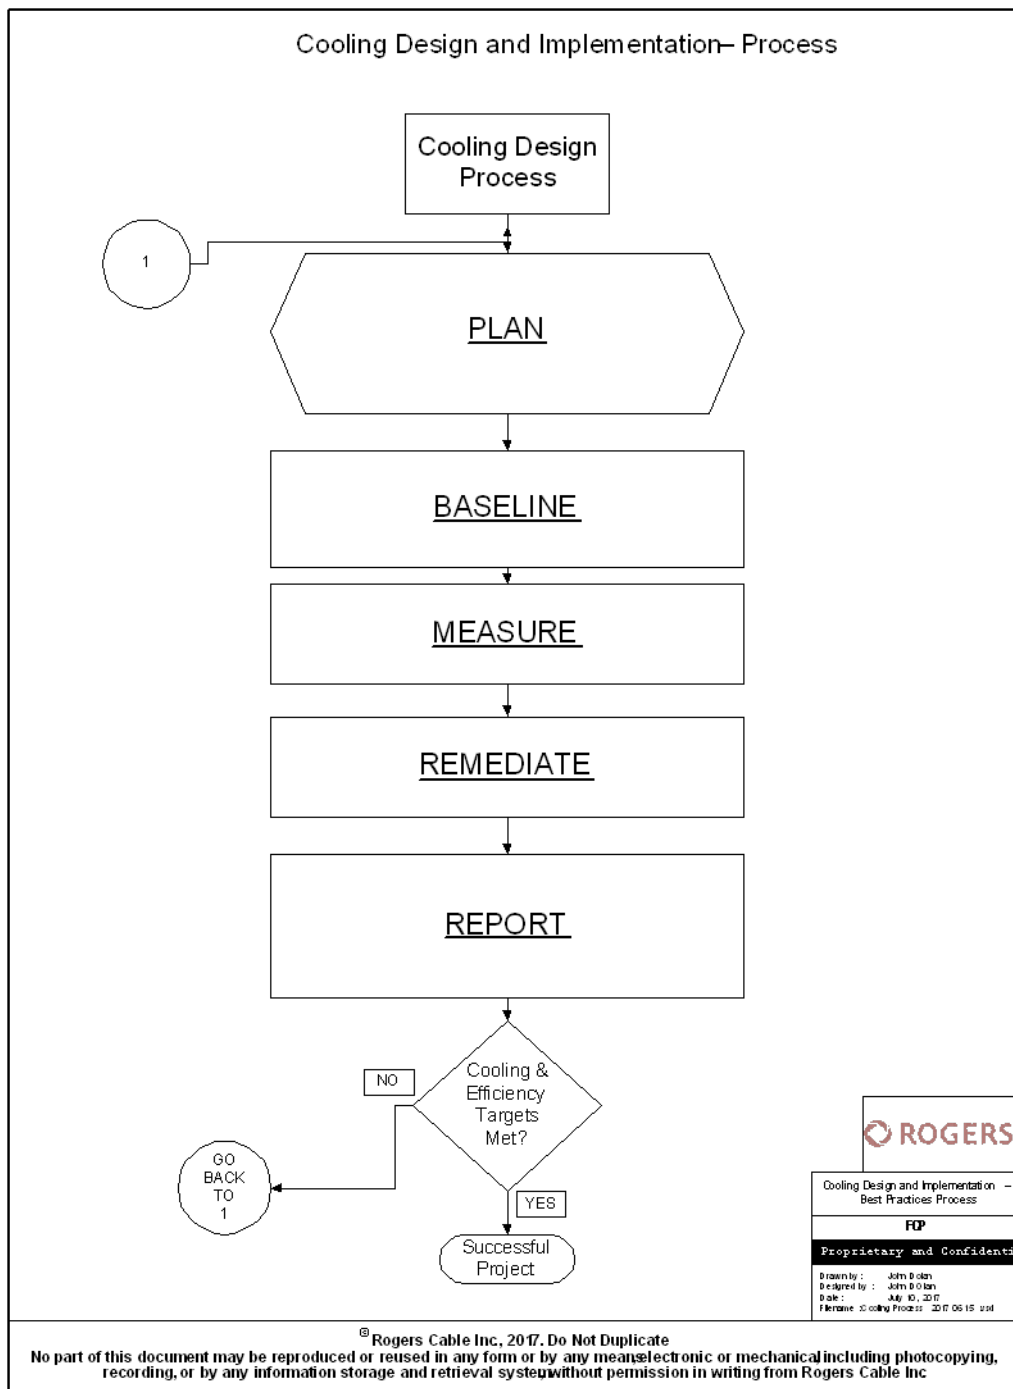


Figure 3 - Energy Plan Methodology (Used with Permission from Rogers)

3. Facility Climate Data Collection

It is extremely important to collect data both before and after implementing energy conservation measures to verify the achievement of the cooling energy management plan goals, confirm estimates of energy savings from vendors and modeling activities, and most importantly to ensure that the rack inlet

temperatures are maintained at recommended levels while lowering the cost of cooling. Many improvements such as airflow optimization, refrigerant replacement, advanced controls, and newer, more efficient HVAC units will yield immediate improvements that can be seen within a month after implementation. Others such as HVAC economizers may require data collection over cooler months where the technology is most applicable to visibly yield savings.

Data collected on the facility climate is useful not only in measuring climate technology efficiency but they are also useful in:

- Reducing alarms
- Reducing hot spots and cold spots
- Reducing temperature variations across the facility
- Improving computer room air conditioner (CRAC) redundancy

Specific data recommended for capture as part of a cooling energy management plan and implementation are provided in this section.

3.1. Rack Inlet Temperatures

Rack inlet temperatures should be characterized, either via direct thermometer measurements, remotely monitored temperature sensors, or via infrared thermal imaging. The goal of airflow optimization is to create as uniform rack inlet temperature profile across the facility as possible and maximize the temperature difference between the HVAC supply and return. Measurement of the rack inlet temperatures, ideally at the top, middle and bottom of each rack, and subsequent characterization of the distribution of rack inlet temperatures across the facility, is preferable prior to, and after implementing airflow optimization measures.

After optimizing airflow, energy savings are realized by gradually raising the HVAC set points. Rack inlet temperatures must be tracked during this process to ensure that temperatures do not exceed recommended values as the set-point is raised. Exact values to be maintained are cable operator specific, but in general the American Society of Heating, Refrigeration and Air Conditioning Engineers (ASHRAE) TC 9.9 guidelines are used.

The following chart shows the ASHRAE temperature and humidity limits (recommended and allowable) on a psychrometric chart. The red section area shows the ASHRAE recommended range for Class A1- A4 facilities.




Figure 4 - Psychrometric Chart showing ASHRAE Temperature and Humidity Levels

The following table shows the equipment environmental specifications for air cooling according to the ASHRAE 2015 Thermal Guidelines. ASHRAE recommends a dry bulb temperature for Class A1 to A4 facilities to be in the range of 64.4 °F to 80.6°F.

Table 1 - ASHRAE 2015 Thermal Guidelines

Table 1. Air-Cooled Data Center Classes (Product Operation)			
Class	Dry Bulb, °F	Max. Dew Point, °F	Humidity
Recommended			
A1 to A4	64 to 81		42°F dp to 60% rh and 59°F dp
Allowable			
A1	59 to 90	62.6	20 to 80% rh
A2	50 to 95	69.8	20 to 80% rh
A3	41 to 104	75.2	10°F dp and 8 to 85% rh
A4	41 to 113	75.2	10°F dp and 8 to 90% rh

Table 2. Liquid-Cooled Datacom Facility Classes (Product Operation)			
Typical Infrastructure Design			
Class	Main Cooling Equipment	Supplemental Cooling Equipment	Facility Supply Water Temperature, °F
W1	Chiller/cooling tower		36 to 63
W2		Water-side economizer	36 to 81
W3	Cooling tower	Chiller	36 to 90
W4	Water-side economizer (with dry-cooler or cooling tower)	N/A	36 to 113
W5	Building heating system	Cooling tower	>113

Source: ASHRAE (2012a).

While the ASHRAE guidelines give a range of rack inlet temperatures that is acceptable, it should be noted that it is also desirable to minimize the variation in rack inlet temperatures, even if all are within the ASHRAE range. If there are many racks at the upper limit and many racks at the lower limit, the set point

cannot be changed in either direction without causing some racks to get outside the recommended range. If instead, all racks are brought to the midrange temperature via AFO, it is then possible to raise the setpoint and achieve energy savings without any racks exceeding the recommended inlet temperature range. A process of continuous improvement in cooling quality would thus involve identifying racks that are at the upper and lower extremes of range, taking AFO measures to bring those racks closer to the same temperature as the other racks and thus providing more uniform cooling to all racks, and then raising the setpoint of the HVAC system to achieve further energy savings.

There are several methods for collecting rack inlet temperature data: some modern IT devices provide both inlet and exhaust temperature measurements. If such measurements are available, they should be captured. If, as more typical, such data from the IT equipment is not ubiquitously available, then rack inlet temperatures should be measured either directly with a handheld fast response digital thermometer or via strip thermometers that are placed at the top, middle and bottom of each rack, or using a thermal imaging camera. The latter is often the most convenient since both rack inlet and exhaust temperatures can be quickly measured and used to calibrate computational fluid dynamics (CFD) models subsequently. Infrared cameras can now be had for under \$1000. The other use of such data is to identify any hotspots at rack inlets that should be addressed via airflow optimization, and for quick visual verification of CFD model output graphics with actual thermal images

If the set points are to be raised following airflow optimization, the inlet temperatures should be characterized daily as the set-points are incrementally increased by one to two °F at a time to insure the facility reaches a steady state temperature distribution and is still within ASHRAE recommendations. For facilities that are poorly insulated and thus more sensitive to outside air temperatures (as can be seen from daily HVAC energy consumption vs. outside air temperature), it is also recommended to characterize the rack inlet temperatures again during the hottest months following implementation of airflow optimization and raising of HVAC set-points to ensure that sufficient cooling continues to be provided to the IT equipment.

Figure 5 depicts an example of rack inlet air temperature characterization as a baseline prior to implementing airflow optimization measures. In this example, airflow direction is not shown, but an important discovery was that the cooling problem turned out to be not as was originally thought. Many of the blue cold racks have inlet temperatures well below the ASHRAE limit of around 80 °F. In fact, only one rack appeared to have a hot spot, and the problem turned out to be an issue of how CRACs were initially installed. As is typical in cable *edge facilities*, even if there are alarms from hotspots at individual pieces of IT equipment, this facility was over cooled, raising the cost of cooling per watt as well as creating issues with forecasted capacity. Solving the hotspot issue subsequently eliminated the hotspot and permitted the addition of IT load without requiring additional cooling.

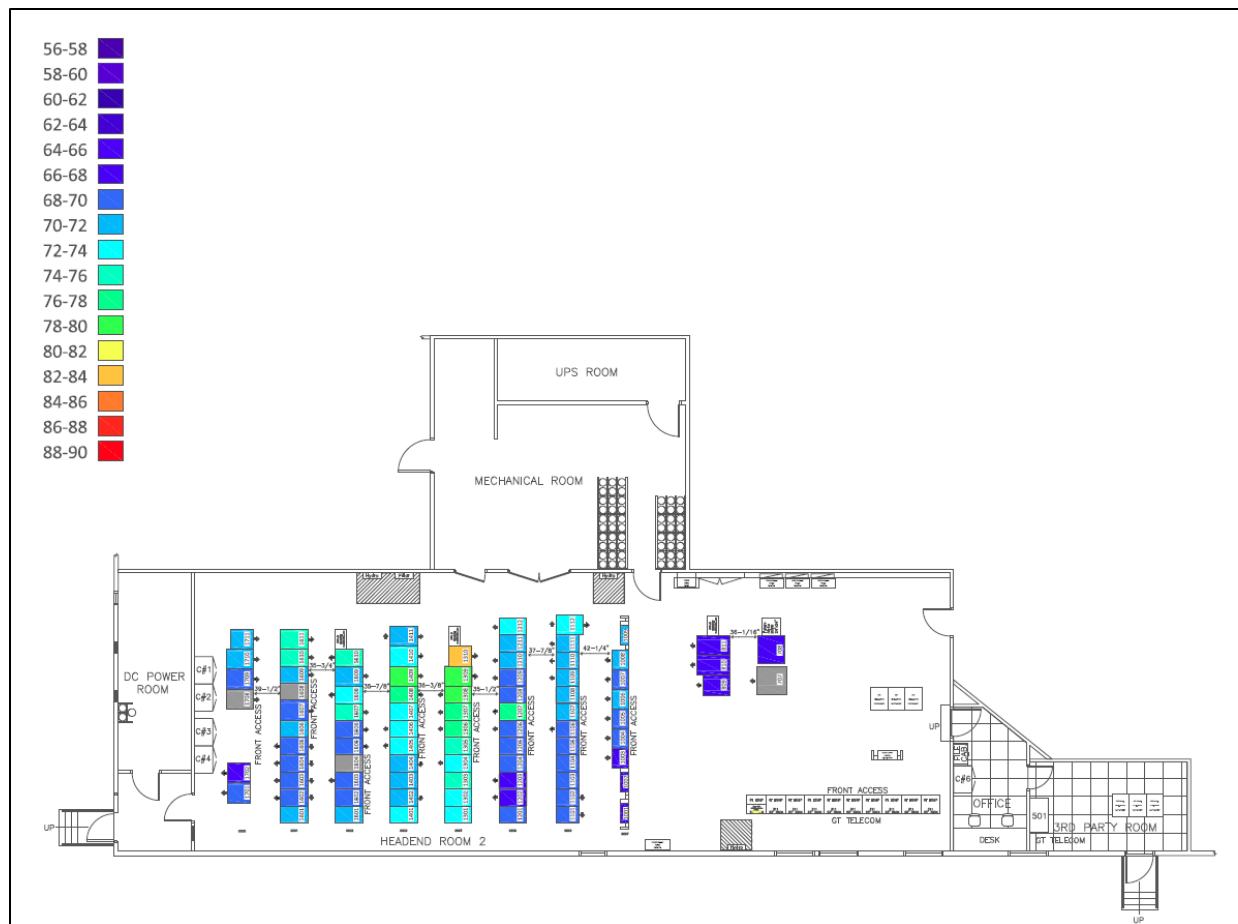


Figure 5 - Example of Rack Inlet Air Temperature Baseline Characterization

3.2. Energy Consumption of HVAC and IT Equipment

Sub-metering to separately characterize both HVAC and IT energy consumption, along with the total facility energy consumption allows calculation of the power usage effectiveness (PUE) per SCTE 213 [S213] as described in the section on PUE below, and ultimately confirms that the cooling cost per watt of IT equipment has indeed been reduced, even if the total wattage of IT equipment to be cooled has changed.

If the IT equipment load as well as other non-HVAC loads are constant and will continue to be constant for several months after implementation of the cooling energy management plan, then it is possible to use the facility utility bill to see the impact of energy conservation measures on reducing the HVAC energy consumption. Unfortunately, due to the consistent growth of services this is seldom the case in cable *edge facilities*; new equipment is constantly being added, or at a minimum, additional cards are added to existing equipment chassis, and ideally older equipment is being decommissioned and removed. The result is that the IT equipment heat load in cable *edge facilities* changes rapidly enough that sub-metering of HVAC equipment is usually required, in addition to sub-metering of the total IT equipment load.

Sub-metering of HVAC equipment can be accomplished with a variety of commercially available devices, with the addition of certain types of HVAC controllers that also monitor energy consumption, or

with monitoring options that can be included when new HVAC systems are installed. Sub-metering may be as simple as sub-metering the direct current (DC) plant (if all IT equipment is fed from the DC plant), or by installing individual sub-metering devices if both DC and alternating current (AC) plants are used to power IT equipment. If the lighting and plug loads in the facility are minimal, or easily characterized/estimated, then it is also possible to use the sub-metering of all HVAC systems to subtract that contribution from the total facility energy consumption along with estimates of the lighting and plug loads to estimate the IT equipment load and track changes that will affect the cost per watt to cool the IT equipment.

SCTE 213 recommends measuring energy consumption of two of the three values of (total building, IT load, and HVAC load) and calculating PUE at least every 15 minutes. If seeking reductions in peak demand costs, which can be a significant component of the total energy bill, it is recommended to measure these values more frequently, on the order of every minute, in order to observe the value and duration of peaks in energy usage that may be addressed by modern control algorithms and technology.

3.3. HVAC Performance

To characterize the performance of the HVAC/CRAC systems before and after implementing energy conservation measures and for performing CFD modeling, the following data should be captured:

- Complete nameplate specifications of all HVAC/CRAC units, including
 - System model and serial number
 - Date installed and date of any major upgrades or retrofits
 - Cooling tonnage provided
 - Refrigerant type and amount
- Physical condition of the HVAC units
- CRAC supply and return air temperatures
- Dimensions and type of supply diffusers and return vents
- Air flow direction (especially for directional diffusers/vents)
- Outdoor air temperature
- Cubic feet per minute (CFM) air flow
- Air pressure differential
- Humidity
- HVAC as-programmed controls and sequences of operations

A small temperature difference between supply and return temperatures often indicates significant mixing of hot and cold air is occurring in the facility, for example, or can indicate issues with the HVAC units themselves.

3.4. Facility Dimensions and Characteristics

If CFD modeling is anticipated as part of a cooling energy management plan then detailed locations and dimensions of racks, openings, HVAC unit location, supply and return details, ducting, large bundles of cabling, gaps in the floor and racks themselves, and so on are required for input into the model.

Additional data will be required such as: what the heat load of each piece of equipment (or the entire rack) and HVAC unit specifications. The airflow direction of the IT equipment can also be important to characterize individually since occasionally some equipment vents differently from other equipment in the rack. It is also important to characterize the existing use or absence of containment such as blanking

panels, plastic curtains/panels, and so on. While general layouts of facilities are available prior to implementing energy conservation measures, a site visit is often best for accurate and complete data capture. The lack of accuracy in facility layouts is due to either incomplete detail in the available layouts, or recent changes in IT equipment and installation or removal of entire racks, both of which can have a significant impact on airflow and cooling effectiveness, and thus the baseline characterization of the site.

4. Placement of Equipment

The placement of IT equipment in a facility has an enormous impact on the cooling cost per watt of IT equipment. Keeping cold supply air separate from hot return air is the key goal, and this can only be achieved via a rigid hot/cold aisle discipline of rack row orientations and IT equipment placement within the racks. Unfortunately, many cable *edge facilities* lack this kind of rigid hot/cold aisle discipline. While some cable operators are moving IT equipment to achieve a rigorous hot/cold aisle discipline in *edge facilities*, others are taking a more gradual approach and adding new equipment in hot/cold aisle manner while removing decommissioned equipment in mixed aisles and blanking openings to reduce mixing in the process. When dealing with legacy facilities even when hot/cold aisle discipline is implemented, it is often only partially implemented to avoid the service disruption that might result from moving IT equipment to achieve such discipline.

In cable *edge facilities*, an entire rack row is often exhausting into the inlets of an adjacent row, and in other cases, one or more pieces of IT equipment are exhausting into what would otherwise be a pure cold aisle with inlets of other IT equipment. Another common departure from a rigid hot/cold aisle discipline occurs when IT equipment that vents side to side is located in racks with front to back type airflow. In some cases, entire racks of IT equipment that vents side to side are exhausting hot air directly into the inlets of an adjacent rack. Luckily many IT manufacturers of these side to side flow devices also offer shrouds that can be installed to convert the airflow to front to back type flow, however these types of configurations should be done during initial equipment installation as modifications can be quite difficult to achieve in a crowded rack row.

Often IT equipment placement is done for the convenience of minimizing cabling efforts rather than optimizing cooling effectiveness. This can result in a new, higher powered (and often denser) piece of equipment heating an aisle that is already challenged for airflow. It can be preferable to place newer, higher powered devices in an aisle that is sparsely populated to distribute the heating load rather than one that is densely populated with other high-powered devices.

Given the increasing heat density of IT equipment and the impact equipment placement has on cooling efficiency, the following should be kept in mind as new equipment is installed:

- Instill hot/cold aisle discipline as new equipment is installed and old equipment is removed
- Consider and/or model the cooling impact of new high-powered equipment *before* it is installed
- Maintain proper setbacks for equipment racks and HVAC units
- Do not block either return or supply air flow with racks and cabling
- Use 100% of existing rack space in lieu of installing a new rack unless it is part of a move toward hot/cold aisle discipline
- Use blanking panels throughout the facility as well as other containment methods such as brushes and end strips/paneling.

Regarding setbacks, an example mandate would be a 2 foot minimum setback from equipment racks/rows to keep them free of cabling to improve airflow and so that ducting can be added/modified for better airflow if the aisle is later contained. Another example setback would be from HVAC supply and return vents so they are not blocked and airflow is maintained.

Ideally equipment deployment should be a joint effort between the facilities and IT workforces. The facilities workforce should be proactive in recommending new equipment locations based on the current state of cooling in the facility. As new equipment is added and old equipment removed, the facility should constantly be striving to achieve a hot/cold aisle discipline/standard to minimize the cooling cost per watt of IT equipment.

5. Airflow Management

Airflow management or specifically airflow optimization (AFO) involves taking measures to reduce mixing of hot and cold air, more effectively transport cooling to IT equipment inlets and heat from IT equipment exhausts to HVAC returns, eliminate hotspots at rack inlets, and make the distribution of rack inlet temperatures more uniform so that HVAC set-points can be raised without exceeding the upper limit at any particular rack or IT device in the facility.

Airflow optimization refers to finding ways to reduce the mixing of hot and cold air by better channeling the cold air to the rack inlets and ensure there is a balance between the air required to cool equipment and total air supplied. Excess supply air results in air bypass meaning the conditioned cold air does not achieve any cooling of IT equipment. Given the chaotic nature of airflow in many cable *edge facilities* and the fact that air patterns are not visible, it is recommended that CFD (air flow) modeling be used to visualize the air flow patterns and to identify problem areas.

5.1. Computational Fluid Dynamics Modeling Tools

A good commercial CFD modeling tool features a graphical user interface, an advanced solver, and powerful visualization and reporting capabilities. It solves the three-dimensional form of the fluid flow and energy equations to predict:

- The velocity, pressure, and temperature fields in the air space under the raised floor (under-floor plenum)
- The airflow rates through the perforated tiles
- The velocity, pressure, and temperature fields above the raised floor or slab, i.e., in the critical space of the facility with the IT equipment

The CFD tool outputs should have been validated against measurements from actual facilities. CFD modeling software can be very compute-intensive, so a high-performance laptop or desktop computer or alternately a cloud-based software solution should be used for the calculations. The package used should include good technical support, and it is recommended to explore several options before selecting a package. Software costs should be balanced against labor time required to set up and run simulations.

Most commercial CFD modeling packages will be able to handle both slab-type and raised floors, as well as ceiling plenum or non-plenum type airflow. Even if the raised floor space is not used for airflow, but instead is used for cabling (which occurs often in cable edge facilities), major openings in the raised floor can be modeled for their impact on airflow.

In the following subsections, an example will be presented of how CFD modeling can highlight airflow management issues and be used to test alternative solutions before making actual changes in airflow in a cable edge facility.

5.2. CFD Baseline Characterization

The baseline CFD model of an example cable edge facility was developed using facility layouts and dimensions along with locations and airflow directions of the following: racks, IT equipment, HVAC equipment, supplies and returns, ducting, and any major obstructions to airflow such as cabinets, shelves or cable bundles. Rack types and dimensions were also used, noting the absence of blanking panels, brushes or other means of covering gaps that can contribute to mixing. For racks with uniform distributions of similar equipment, the CFD modeler may choose to model the rack as a single entity with a given heat load, or may model the actual equipment in the rack with individual power consumption, airflow direction, and presence of lack of blanking panels on the rack. Since hot air can often flow over the tops of rack rows into cold aisles, it is important to capture the rack heights as well as other rack dimensions in the model.

Many cable operators have lists of equipment found within the facility along with the nameplate power consumption of the equipment that can be used to build rack models in CFD modeling packages. Caution in using this data must be exercised as nameplate power consumption levels of IT equipment are generally in excess, sometimes by over twice the actual power consumption of IT equipment and equipment lists can be out of date. Thus, the CFD modeler uses a combination of thermal images captured on site with previously determined de-rating factors known to be common for cable edge facility equipment. Finally, the CFD model parameters can be calibrated from the baseline CFD model of the facility to match measurements from the site itself and reproduces known issues at the site. Figure 6 depicts the baseline CFD model of this cable edge facility.




Figure 6 - Baseline CFD Model of a Cable Edge Facility

In Figure 7 hotspots are shown in red meaning the inlet temperatures are above ASHRAE maximum recommended level of 80.6°F. Dark blue areas represent spaces where air temperature is below the ASHRAE recommended minimum of 64.4°F. There are three return grids along the wall behind the header, one of which was inactive (see the three grey top ducts; the one with a red X was inactive).

The yellow and red regions in the figure indicate locations that are at the top of the ASHRAE standards for data center acceptable temperatures. There were no alarms from this high heat due to the high temp alarms being located far away on building walls and the lack of thermal sensors in the equipment in these racks. The red arrows in the figure indicate heat movement as it flows through the racks, which is happening due to the lack of blanking panels. Hot spots of up to 95 °F were seen, again with no alarms.

In this example, there was more than adequate cooling capacity: The 3 RTUs (roof top units) provide up to 264 kW of cooling, although only two were actually being used for cooling. The calibrated IT equipment load was determined to be 68 kW, which was far less than the nameplate values for this facility. Nonetheless, lack of good airflow management/distribution meant hot spots were present even given the over-cooling. Since the average total building load was 131 kW, the PUE was estimated at 1.93 for this facility for the baseline CFD model.




Figure 7 - CFD Model Results After Airflow Optimization

5.3. CFD Model After Airflow Optimization

As can be seen, there are several hotspots in rack inlets that should be mitigated, and further there is a lot of mixing of hot air from rack exhausts with cold air coming from the overhead ducting before it gets to the rack inlets. There is a very non-uniform distribution of rack inlet temperatures, which also results

from the mixing of hot and cold air in the facility. Further, the thermostats were located on the walls of facility, in some cases far away from the inlets of key racks in the facility. These were the issues to address in airflow optimization for the facility using CFD modeling.

A variety of airflow optimization tactics were explored for this facility which resulted in a proposed design for airflow optimization involving the addition of new ducting, end aisle containment and blanking panels. Figure 8 shows the improvements in airflow management that resulted from the design. Hotspots have been eliminated, mixing of hot and cold air significantly reduced, and the range of rack inlet temperatures has been reduced so that the inlet temperatures in the facility are much more uniform. In essence, with the same cooling capacity, the facility has become much cooler, and the set-point can be raised to achieve energy savings and reduce the cost of cooling per watt of IT equipment.

5.4. CFD model after airflow optimization and after raising the set-points

In addition to the airflow management changes just described, the thermostats in the CFD model were moved to the cold aisles and additional thermostats were added so that the air temperature could be controlled where it matters most - at the rack inlets. Next the set-points of the HVAC systems were raised in the CFD model until just before the point where ASHRAE limits were exceeded on any particular rack. The results are shown in Figure 8.




Figure 8 - CFD Model Results After Airflow Optimization and Raising the Set-Point

After raising the set-point by 6 °F, the inlet temperatures are slightly higher but still below ASHRAE upper limits, and some hotter areas at rack exhausts can be seen. Very hot exhaust aisles are normal and acceptable in modern facilities; it is the mixing of hot with cold air that is undesirable. With a 6 °F increase in set-point PUE was reduced from 1.9 to 1.7, the HVAC energy consumption was estimated to be approximately 17% lower, thereby significantly reducing the cooling cost per watt of IT equipment in

the facility. Also, and importantly, elimination of equipment inlet hotspots avoiding any customer impact hours.

6. Beyond Airflow Management

Even if other energy conservation measures such as HVAC optimization, upgrades or replacements are planned, it is important to do airflow optimization first. This is because airflow optimization can reduce the tonnage of cooling capacity and thus the cost of HVAC upgrades or replacement, and can also enable true HVAC redundancy in the facility via ducting additions so that in the event one HVAC unit fails, cold air is still provided to the IT equipment. HVAC upgrades and replacements with more efficient technologies are covered in a later section. In this section, other optimizations of existing HVAC systems beyond airflow optimization will be covered.

Once the airflow has been optimized and HVAC set-points have been raised to achieve energy savings, the following additional optimizations of existing HVAC systems are possible and have been shown to reduce the current draw of existing HVAC systems and thereby lower the energy consumption of the HVAC systems, in addition to potentially extending the lifespan of existing HVAC units:

- Installation of add-on or stand-alone economizers to existing HVAC systems
- Deep cleaning of coils and other system maintenance items to ‘true-up’ the HVAC system back to its nominal operating condition
- Disabling or putting into lag mode any existing HVAC units that are deemed unnecessary after airflow optimization. (Lead-lag means that one HVAC unit activates first (the “lead” unit), then another activates after a "lag" only if necessary and may only activate if the first unit fails entirely. Alternately, the HVAC systems can be sequenced to distribute the runtime across the units while still reducing energy consumed since fewer units are simultaneously running.)
- Replacement of older refrigerants such as R22 and R407A with modern, more efficient refrigerants
- Installation of advanced HVAC controllers that can reduce the compressor on-time required to achieve the same level of cooling

Figure 9 below shows an example flow chart for exploring additional HVAC optimization measures that reduce the HVAC energy consumption following airflow optimization.

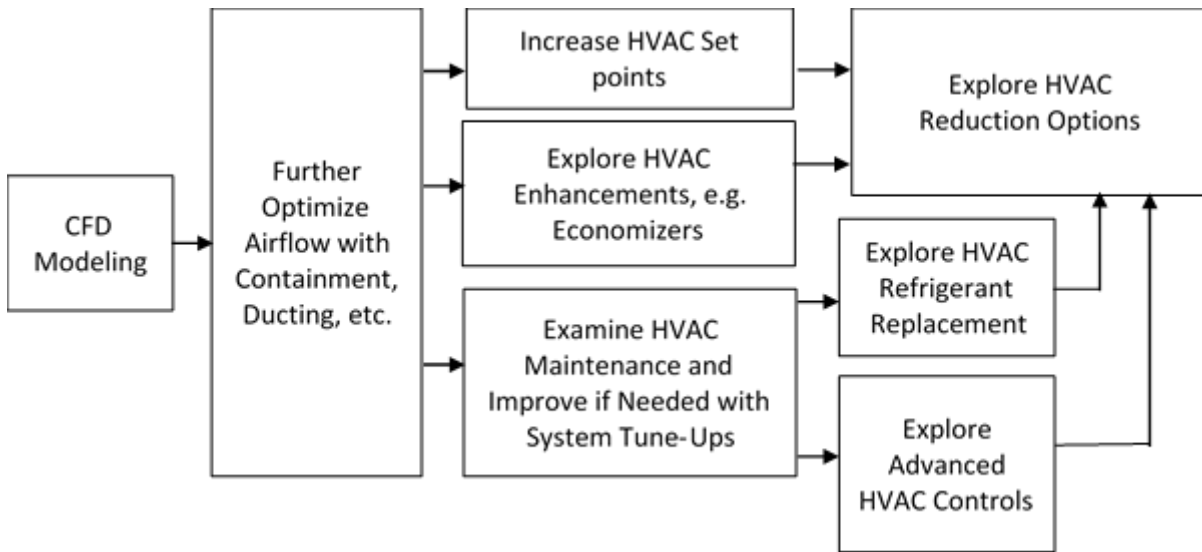


Figure 9 - Energy Conservation Measures Beyond Airflow Optimization

7. Power Usage Effectiveness

ANSI/SCTE 213 2015 [S213] describes how the power usage effectiveness (PUE) is to be measured in cable *edge facilities*. The PUE is defined as:

$$\text{Power Usage Effectiveness} = \frac{\text{Total Critical Facility Energy}}{\text{IT Equipment Energy}}$$

The above definition of PUE presumes the availability of detailed sub-metering data on the components of the facility energy consumption (IT, HVAC, lighting, plug load), or at least the availability of accurate IT and HVAC load sub-metering data. One would typically measure HVAC consumption at the CRAC units using sub-metering devices, and measure the IT load at panels of DC plants. Unfortunately, even these minimal measurements are often not available in current cable *edge facilities*, and the only data available is the total building energy consumption from the utility bill and the IT load from DC plant monitoring. Further, many cable *edge facilities* are either part of a larger facility, or have significant portions of the facility dedicated to office space that is separate from the critical equipment spaces, or have multiple zones that are independently cooled within the facility. These variations, along with the rapid changes in IT load often seen in cable *edge facilities* can significantly complicate the calculation of PUE. Nevertheless, those skilled in the art can accurately estimate the other components of energy consumption and compensate for non-critical spaces in the facility based on an energy audit of a site, detailed equipment lists, and so on. An example of this approach is shown in Figure 10 below for a cable headend. Whichever approach is used, it must be identically applied to all facilities, and preferably calibrated by actual detailed measurements from at least one facility.




Figure 10 - Example of Energy Consumption by Percentage for a Cable Headend Facility

The distribution of energy consumption shown in Figure 11 is quite typical for cable edge facilities, but can vary greatly if the facility has the corner conditions just described or has already received the benefits of energy conservation measures to reduce the HVAC energy consumption as a percentage of overall facility consumption.

ANSI/SCTE 213 2015 [S213] describes how to convert from kWh to kW, provides examples of PUE calculations for cable *edge facilities* and describes how frequently to measure PUE. Especially for cable facilities that are not well insulated, the PUE can be quite sensitive to outdoor air temperature, as is seen in Figure 12. Higher outdoor air temperatures create an additional heat load on the facility that requires additional cooling for the same IT equipment load, thereby raising the PUE of the facility.




Figure 11 - PUE vs Outdoor Air Temperature in Degrees Celsius for a Facility Before and After Aisle Containment (see Bibliography & References)

In this case, the behavior of PUE vs. outdoor air temperature was plotted before and after implementing airflow optimization using aisle containment, and the effect on PUE is dramatic. Changes in outdoor air temperature and changes in the IT load and the HVAC load from drastic IT load increases can all skew the PUE measurement. Hence careful analysis of the actual data for a site is required to accurately determine the impact of energy conservation measures.

8. Incentives

Energy rebate incentives are available to help offset the cost of many energy efficiency projects for network facilities through state and provincial utility companies. A broad range of projects are considered eligible and include for example replacing lighting with LED's, upgrading cooling units with variable speed fans, replacing energy inefficient cooling with more efficient units, installation of control systems, or replacing electrical equipment with natural gas powered. Each state has different programs and incentives so discussion with the local utility is necessary to determine project eligibility. In general incentives are available up to a maximum of 50% of the total project cost and are based on level of energy savings.

The key to successfully maneuvering the rebate process is to contact the local utility in advance of any work or commitment to purchase materials is made. Ensure there is a paper trail within your organization noting the intent to apply for incentive rebates and a paper trail with the utility companies. It is advisable to do this weeks in advance to have the opportunity to discuss the project, understand what the expectations are from the utility and agree on how the requirements such as energy metering, if required, will be done to satisfy the measurement and verification (M&V) plan requirements.

An application is submitted once agreement is reached on the project objectives, approach and M&V plan. Approval of the application can take a few weeks and may involve more discussion with the utility representatives. With most utilities once the application has been submitted work can begin on the project however if baseline measurements are required (and they often are) they must be taken before anything changes or work begins. Check with the local utility as some may require waiting until the application is approved before work can begin. Others prefer the baseline measures completed and submitted as part of the initial application.

Documentation required with the application filing generally includes:

- Description of base case and energy efficient case
- Technical specs on existing and new equipment
- Quotes, estimates or pricing from the vendor for the new equipment to show project cost
- Calculations showing the projected energy reductions (these are estimates as metering may not have been done at this point)
- A variety of forms to be signed by company representative

Costs internal to a company, such as manpower to install or remove equipment are not considered eligible project costs.

Upon completion of the project a second energy measure is required. The pre and post-measures are used to determine the value of the energy rebate. Documentation generally required for the final application submission includes:

- Invoices showing product and installation costs
- Metering results from pre and post measurements and the level of energy savings
- Disposal forms for old equipment
- Usually a variety of forms unique to the utility company to be signed by company representative

The utility company may request a site visit pre- and post-work to verify the equipment and work completed. This may be a utility representative or a third party. Requests to take pictures of old and new equipment may be made – if company policy does not allow pictures this should be pointed out the representatives during the initial application discussions.

After the final application material is submitted and the utility approves the energy rebate on average, 8-12 weeks is required for processing of payment. The company representative will be advised if any questions have arisen or if the final application has been approved. Utilities generally require the company to create an invoice to the utility for payment of the energy rebate.

9. Summary of Plan, Measurement and Verification

Energy consumption before and after climate technology optimization must be accurately characterized for internal budgeting and for external purposes such as energy rebates and incentives. Documented results will give confidence that the outcome could be replicated at other *edge facilities*.

9.1. Pilot

Pilot new technologies or approaches in representative and realistic conditions. If this is a technology that can be applied to multiple *edge facilities* then make sure you select a site that is a good representative and try to eliminate or at least account for anomalies that might impact energy savings such as any new equipment being added, unique equipment layout, etc.

Conduct your testing during a period when you expect to see the results of the technology improvement. Even though *edge facilities* have year-round cooling demands, the effect of outside air temperature will greatly impact your results.

9.2. Baseline

Establish baseline models to evaluate new technologies. Baseline measurement can be expensive for *edge facilities* but it is necessary to properly evaluate improvements. The cost of monitoring and measurement is coming down with new Internet of Things (IoT) technology which are lower cost and easier to install.

Gather baseline environmental data such as supply air temperature (SAT), return air temperature (RAT), supply air humidity (SAH), return air humidity (RAH) and rack inlet temperatures. Knowing how the facility is performing is important to understand what impact the new technology has on the operation. Network reliability cannot and must not be compromised for the sake of energy savings.

Install sub-metering on HVAC & IT equipment to measure consumption and calculate PUE. Depending on cost and how the facilities' electrical panels are configured, it might be more cost effective to meter the incoming power to the facility and sub meter the HVAC. This can be an acceptable alternative if the rest of the electrical load is inconsequential or somewhat static. The goal is to be able to measure the improvement by accurately assigning the reduction in energy to the technology being implemented.

9.3. Implement

Implement the energy conservation measures at the facility and monitor the performance. If the measures are multifaceted then consider implementing them in stages to analyze the performance of each one independently. Methodical implementation will help determine which measures have the “biggest bang for the buck” so priorities and budgets can be set accordingly.

9.4. Analyze

Gather post-implementation data and calculate energy and cost savings. Consider the following follow-on questions:

1. Did the new optimization technology perform as expected?
2. Is the return on investment (ROI) payback within company guidelines? Use the environmental data collected to evaluate the impact on the performance of the facility. I.e. was proper temperature maintained in the facility?
3. Are there any inlet hotspots that need to be addressed?

9.5. Refine

Is there confidence the results can be repeated in other *edge facilities*? What worked well, what didn't? What should be done next? Use the data gathered to refine goals and design larger programs for the entire footprint.

9.6. Ongoing

Continue to capture data for 'health checks', changing trends, and to explore new opportunities. Ongoing monitoring will help to ensure the new measures are operating properly and that changes in performance can quickly and accurately be evaluated to prevent degradation in savings over the long term. *Edge facilities* are surprisingly dynamic and require constant monitoring to keep them at optimum performance.

10. Cooling Technologies

In this section, we provide top level insight into choosing the primary method for cooling a facility containing electronic equipment as well as choosing an economizer technology that can be coupled with a primary cooling method. We provide the reader a definition of each primary method and each economization option as well as a decision guideline that steers the reader towards the best efficiency.

It is not always possible or practical to simply choose the most efficient cooling approach as there are a series of other trade-offs that must be considered. These tradeoffs are outlined in this section in a qualitative manner which looks at the following elements of each choice:

- **Efficiency** → which methods typically deliver cooling capacity while consuming least amount of electricity.
- **Modularity** → Do you have options for growing capacity that creates an advantage like space utilization or control & coordination over simply adding more HVAC systems
- **Installed Cost** → This includes both relative equipment and labor costs and compares each option's upfront implementations costs.
- **Limitations** → This is an attempt to focus some attention on less tangible items. Limitations refers to the requirements for implementing a technology and whether the equipment room can accommodate each technology. For example, some sites do not have access to running water. There are several cooling and economizer solutions that require running water to work so these choices
- **Maintenance** → This characteristic compares the effort/cost to maintain each system. The simpler and lower cost to maintain the more preferable.

10.1. Primary Cooling Options

One or a combination of the methods below is typically used as primary cooling method to offset heat load year-round

10.1.1. Direct Expansion

Direct expansion (DX) uses a compressor to drive a refrigeration cycle to cool a site. It is also known as mechanical cooling. It introduces no outside air and can cool the indoor temperature cooler than the outdoor temperature. It has the least limitations but is also the least efficient.

10.1.2. Chilled Water

Chilled water (CW) system provides cold water to cool the room. After absorbing heat from the room, the chilled water returns to the chiller where the chiller removes the heat from the water using the refrigeration process. It has high efficiency but is often quite complex to implement and therefore most often cost effective for very large projects (> 100kW).

10.1.3. Adiabatic Cooling

Adiabatic (AD) cooling works by blowing the supply air through an evaporation pad. Water evaporates as the air passes through the pad, the air cools down much like the process of perspiration cools down your skin. It is very efficient and works best in dryer environments since the cooling capacity is limited when the environment has high relative humidity.

10.2. Economization Options

Note: not all economization options are compatible with each primary cooling option.

10.2.1. Direct Airside Economizer

Direct airside (DA) economizer uses fresh air to cool the room when the environmental conditions are favorable. DA is the most energy efficient way of cooling a site because it directly uses the outside air with no loss from heat exchange. In areas with lower air quality, filtration or other mitigating technology must be properly deployed.

10.2.2. Indirect Airside Economizer

Indirect airside (IDA) economizer uses some form of heat exchanger to transfer heat between inside and outside air. It has lower efficiency than direct airside economizer but does not introduce outside air into the room.

10.2.3. Pumped Refrigerant

Pumped refrigerant (PR) economizer uses an economizer pump in place of the compressor when the outside air is significantly colder than the room temperature. Similar to indirect airside economizer, it does not introduce outside air. Additionally, it does not require any large wall penetration but usually has higher upfront cost (equipment and installation).

10.2.4. Indirect Waterside Economizer

Indirect waterside (IDA) economizer typically uses one or more cooling towers and heat exchanger to offload all or parts of the chiller's load. Freeze protection must be considered for colder regions.

10.3. HVAC System Selection Guide

The following flow chart provides a process to help select among the various cooling technologies available for *edge facilities*. It provides a basic guide but does not go into the detail of such items as compressor types, air flow management or filters.

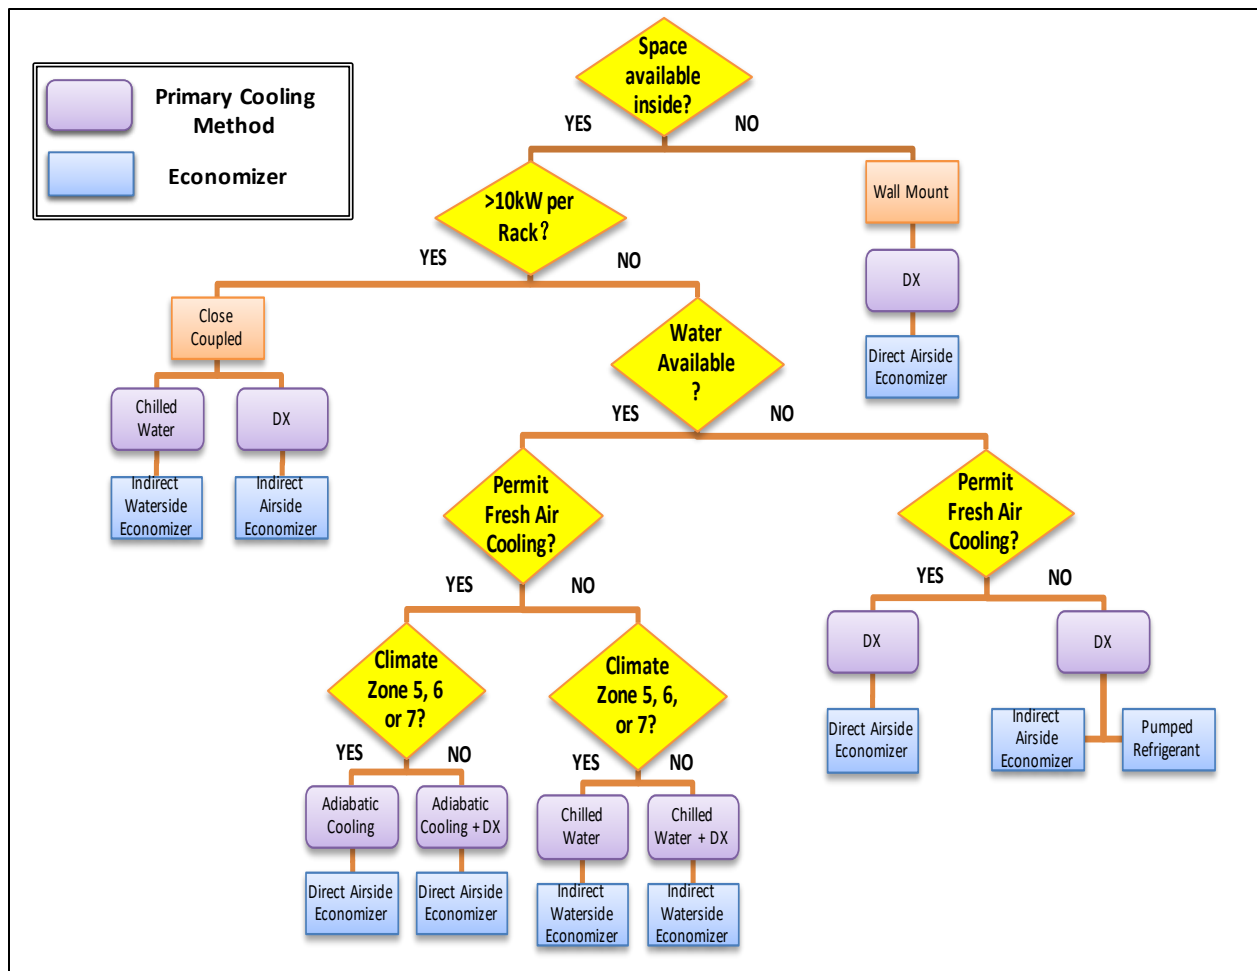


Figure 12 – HVAC System Selection Flow Chart

10.4. HVAC System Selection Guide

Figure 13 provides an overview of a qualitative analysis of the cooling technologies applicable to *edge facilities*. The scale provided is 1 to 100 with 1 being the lowest in a category and 100 being the highest in a category. For example, DX+PR cooling technology has a fairly high installed costs, but high efficiency, but also has limitations in size and is complex, modular to deploy (costs scale as you grow) but has a more parts to maintain.




Figure 13 - Comparison of various HVAC system by category

11. Financial Analysis

Two viewpoints will be considered under financial analysis, actual financial analysis and non-financial analysis focusing on other factors such as customer outage hours.

11.1. Financial Payback

The focus of financial analysis is simple payback and compares 2 options. For example, continuing to use the existing cooling technology compared to adding in and economizer (free air cooling). Payback is generally acceptable under 3 years.

Analysis can be simple payback or more complex such as NPV (net present value), IRR (internal rate of return) and ROI (return on investment).

Simple payback is often used but ROI is more useful as it includes all factors for the life of the cooling technology including but not limited to: Initial costs, rebates, installation costs, cost of power, energy costs, maintenance costs, component replacement costs and end of life costs.

Simple Payback can be calculated as follows:

$$\text{Payback Period} = A + B/C$$

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. Normally less than 3 years is acceptable.

There are more detailed well know financial analysis that can be considered for large projects: NPV, IRR and ROI.

Net present value is another way of representing a ROI. It is the present value of the cash inflows and outflows at the required rate of return of your project compared to your initial investment. This is superior to Payback as it considers the time value of money over the life of the asset. The more positive the NPV the better this is for a project. Negative values indicate the initiative should be abandoned.

Analysis of the net present value (NPV) of two options is as follows.

$$NPV = -C_0 + \sum_{i=1}^T \frac{C_i}{(1+r)^i}$$

C_0 = Initial Investment

C = Cash Flow

r = Discount Rate (different from business to business)

T = Time

The Internal Rate of Return provides the rate at which a project breaks even. This does not provide actual amounts and should be used in conjunction with NPV. A high positive % is a good indication of a good investment. Of course, a solution with a negative % ROI should be abandoned.

Analysis of IRR of two options is as generally follows but it is recommended to use one of the tools available such as Excel.

$$0 = P_0 + \frac{P_1}{(1+IRR)} + \frac{P_2}{(1+IRR)^2} + \frac{P_3}{(1+IRR)^3} + \dots + \frac{P_n}{(1+IRR)^n}$$

P_0, P_1, \dots, P_n Cash flows in periods 1, 2, ...n

IRR = Internal rate of return

Return on Investment, ROI, measures return on investment over time. A higher % indicates a good investment. Again, a very low or negative % ROI solution should be abandoned.


Analysis of ROI of two options is as follows:

$$ROI = \frac{\text{Gains from investments} - \text{Cost of investment}}{\text{Cost of investment}}$$

11.2. Example, Simple Payback and Operational Savings

An example of simple payback and operational savings is a study done by SCTi at a Rogers Facility to determine possible savings by raising the CRAC set point to reduce its energy consumption. Expectations in the industry appear to be modest at 1 – 2% savings per °F, although some have reported as high as 6% savings. This is highly dependent on initial conditions and other initiatives that may have been done a site to improve other aspects of cooling such as air flow. The results are shown in Table 2.

Table 2 - Savings by Increasing Set Point

	Energy Reduction					
	Set Point		kW Average	% kW	% kW per	
	°C	°F			°C	°F
Data Total for 4 CRAC						
2016-10-30 Baseline	21	69.8	93			
10-Dec	25	77	86			
Δ	4	7	7	7.5	1.875	1.042
Per CRAC Annual	kWh		15,330			
	\$ at \$0.13 per kWh		\$2,000			

The annualized savings per CRAC unit is \$2,000 and simple Payback is 1.25 years for this initiative. Assuming 50 similar facilities the savings would be \$1M in energy (operation) costs. Be aware that these types of studies can be complicated to plan and execute because of the number of variables affecting the energy efficiency results.

11.3. Other Non-Financial Considerations

For non-financial analysis, the Payback period is less important than improvement in other areas. It is more about increasing reliability, reducing churn, adding more customers, customer satisfaction or network uptime.

One general measure of this is customer impact hours of a cooling solution or technology. If the aim is to reduce customer impact hours as a result of cooling issues it may be prudent to, for example, invest in greater IT thermal resilience by installing more than N+1 in cooling, or add more sophisticated remote monitoring capabilities for remote sites to reduce truck rolls, maintenance costs, and provide a more predictive maintenance response.

12. Conclusion

This paper has provided the outline for an energy management plan for a process that will improve energy efficiency at edge facilities. Airflow is key and keeping supply and return air separate has the largest impact on the efficiency of the cooling system. A plan will ensure baselines are established before a project starts and that efficiencies and savings of the cooling technology solution can be clearly demonstrated. It ensures that appropriate data is gathered throughout the project to calculate the savings or payback. Finally, a plan can begin simply with airflow optimization but can evolve as required to use CFD and calculation of PUE to further improve efficiencies.

Abbreviations

AC	alternating current
AD	adiabatic cooling
ANSI	American National Standards Institute
ASHRAE	American Society of Heating, Refrigeration and Air conditioning Engineers
CRAC	computer room air conditioner
CFD	computational fluid dynamics
CFM	cubic feed per minute
CW	chilled water
DA	direct airside economizer
DC	direct current
DX	direct expansion
HVAC	heating ventilation and air conditioning
IDA	Indirect Airside Economizer
IDW	Indirect Waterside Economizer
ISBE	International Society of Broadband Experts
IRR	internal rate of return
IT	information technology
IoT	Internet of things
PR	pumped refrigerant
PUE	power usage effectiveness
MSO	multiple system operator
NCTA	National Cable and Telecommunications Association
NPV	net present value
RAH	return air humidity
RAT	return air temperature
ROI	return on investment
SAH	supply air humidity
SAT	supply air temperature
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

ASHRAE TC 9.9 Thermal Guidelines for Data Processing Environments – Expanded Data Center Classes and Usage Guidance

SCTE 184 Energy Management Operational Practices for Cable Facilities

SCTE 213 Edge and Core Facilities Energy Metrics

SCTE 219 Operational Practice, Technical Facility Climate Optimization Methodology

SCTE 226 Cable Facility Classification Definitions and Requirements

WP49-PUE The Green Grid: PUE: A Comprehensive Examination of the Metric

DKRZ News Archive:

<https://www.dkrz.de/about-en/contact/press/news-archive/increase-in-energy-efficiency-at-the-dkrz>

Energy Conservation Measure Recommendations for Cable Edge Facilities

Energy Audits and Analysis of Ten Cable Headends

A Technical Paper prepared for SCTE•ISBE by

Daniel Marut

Senior Manager of Sustainability – Energy & Technology
Comcast
1701 JFK Blvd, Philadelphia, PA, 19103
215-286-7319
Daniel_Marut@comcast.com

Daniel Howard

Director of Consulting Services, Hitachi Energy and Environmental Efficiency Group
Hitachi
2512 Parkdale Place NE
404-625-1593
daniel.howard@hitachiconsulting.com

George Gosko, Hitachi

Supriya Dharkar, Hitachi

Riebeeck van Niekerk, Hitachi

Tanner McManus, Hitachi

Michael Baselice, Comcast

Gregory Baron, US Air Force (formerly with Hitachi)

1. Introduction and Executive Summary

Comcast contracted Hitachi Consulting to explore energy conservation measures (ECMs) at five headend sites in the West Division and five in the Central Division. The task involved an on-site energy assessment, development of computational fluid dynamic (CFD) models of existing airflow conditions, and recommendations of ECMs for each headend. The effort resulted in identification of three key measures that apply broadly to Comcast headends and hubs: airflow optimization, advanced HVAC controls, and replacement of older, less efficient and ozone-depleting refrigerants. Implementing these three measures would provide a 5-year energy savings opportunity for the ten sites of just over \$1.5 million, with the annual savings being just over \$300,000. Hitachi Consulting also assessed LED lighting opportunities at the ten headends. Implementation of LED lighting and controls would provide a 5-year energy savings opportunity for the ten sites of just over \$300,000, with the annual savings being just over \$60,000.

The motivation for the effort is the fact that cable headends and hubs often do not employ the most modern cooling practices such as contained equipment aisles with hot/cold aisle discipline, which is now common in most data centers. These headends and hubs consequently have far more cooling capacity that would otherwise be needed. The challenge is to explore what could cost-effectively be done in these facilities to achieve significant energy savings in a reasonable payback period.

Detailed cost proposals from a multitude of subcontractors across all sites was not feasible for the present effort. However initial estimates indicate that payback periods on the order of 3 years or under are feasible for most sites and with sites in states with higher utility rates paying back even sooner. The estimated range of implementation costs varies from approximately \$40k to \$160k, depending mainly on the size of the site. The true cost of implementation and payback period can only be determined from piloting the ECM implementations and measuring the actual energy savings obtained in the pilots.

In addition to potential energy consumption and cost savings benefits, there are also significant performance and customer satisfaction improvements that come from having more efficient, robust, and redundant cooling in headends and hubs. The benefits of the airflow optimization, advanced HVAC controls and refrigerant replacement also improves:

- power margin
- site resiliency towards R-22 phase-out by 2020
- normalizing inconsistent temperatures across the inlet side of the equipment
- reduces overheating equipment situations with no alarms
- adds HVAC redundancy
- and extends the useful life of the HVAC technology.

All of this leads to significant operating expense (OpEx) cost reductions and improved customer satisfaction via reduced IT equipment downtime.

More optimized cooling technology can also reduce the cost of future capital investments by lowering the tonnage of cooling required in replacement projects. The reduction of total energy consumption at headends and hubs can also enable more sites to be viable for alternate energy projects that seek to reduce Comcast's dependence on the electrical grid and reduce the carbon footprint overall.

The specific energy conservation measures recommended in this effort include:

Air Flow Optimization (AFO):

- Increase utilization of blanking panels in all racks to limit the hot and cold air to a specific space and limit infiltration within the racks
- Increase utilization of top of aisle containment to limit hot air recirculation and infiltration over the top of the racks
- Increase utilization of end aisle containment (strip curtains or end panels/doors) to contain cold aisles and prevent infiltration of cold air
- Redirect and/or add additional supply ducting to deliver cold air directly into contained cold aisles
- Reposition and/or add additional return ducting to facilitate hot air return to CRAC units

HVAC Controls:

- Add advanced HVAC controls to optimize key components of the HVAC system to reduce HVAC energy consumption by 15-25%

Refrigerant Replacement:

- Install nextgen replacement refrigerants that extend the life of existing HVAC systems and can also increase efficiency and capacity over R-22 and R-407C by as much as 20%
- As part of refrigerant replacement and/or installation of advanced controls, “true-up” the HVAC equipment to address any performance issues and bring it back to nominal operation.

In this report, the results of detailed site visits, modeling and recommendations for each of the ten headend sites will be presented, followed by an analysis of the portfolio overall as well as conclusions and recommendations from the effort.

The ten Comcast headends covered by this report are:

West Division Sites

- Roseville, MN
- Hayward, CA
- Santa Clara, CA
- Beaverton, OR
- Burien, WA

Central Division Sites

- Stone Mountain, GA
- Atlanta, GA
- Jonesboro, GA
- Woodstock, GA
- Augusta, GA

The potential energy savings associated with implementation of these three ECMs at the 10 headend facilities is summarized in Table 1 below. The average utility rate for these 10 sites was \$0.081, and as stated in the introduction, when all sites are considered, the total energy cost savings over 5 years was estimated to be \$1.5M.

Table 1 - Energy Savings Analysis for ECM Implementation at All Ten Headend Sites

Facility	Size / Max IT Load (provided by Comcast)	Energy Savings (kWh)	% HVAC Energy Reduction
Roseville, MN	24,175 ft2 / 470 kW	311,133	22%
Hayward, CA	33,000 ft2 / 330 kW	213,206	26%
Santa Clara, CA	28,800 ft2 / 460 kW	284,735	17%
Beaverton, OR	13,737 ft2 / 1,100 kW	479,090	23%
Burien, WA	6,561 ft2 / 530 kW	306,214	19%
Stone Mountain, GA	25,596 ft2 / 960 kW	704,865	19%
Atlanta, GA	24,626 ft2 / 1,000 kW	751,019	19%
Jonesboro, GA	6,467 ft2 / 220 kW	170,482	36%
Woodstock, GA	6,720 ft2 / 430 kW	384,536	38%
Augusta, GA	10,245 ft2 / 210 kW	142,933	24%
All Facilities	179,927 ft2 / 5,710 kW	3,748,213	22%

Table 2 below shows statistically how the rack inlet temperature changed before and after AFO implementation for all ten headend sites. Note that 383 racks with max inlet temperatures of 80°F or more have been fixed, and if the desired maximum inlet temperature is 75°F or less, then implementing the airflow optimization ECMs brings 466 racks into the standard, even after the set point is raised in the facilities with airflow optimization.

Table 2 - AFO statistical impact on rack inlet temperature distribution after set point increase for all 10 headend sites.

Range of Max Inlet Temperature	Number of Racks *		
	Baseline	After AFO ECMs	After AFO ECMs and Set Point Increase
Above 80°F	384	1	1
Between 75°F and 80°F	222	47	139
Between 70°F and 75°F	329	246	714
Below 70°F	781	1,422	862
Total	1,716	1,716	1,716

* Number of racks for all three scenarios does not include Beaverton, OR – Phase 2 headend, Stone Mountain, GA – VPC2 and Atlanta, GA – IT Room as these rooms were not recommended for AFO due to discontinuities in hot/cold aisle layouts and therefore raising the set point temperatures would achieve

significant energy and cost savings. The table also does not include 304 empty racks from some of the sites that did not have any installed IT equipment at the time of the audit.

2. Site Audits and Analysis

2.1. Procedure

For each site, a detailed site assessment was performed, involving tasks prior to the site visit, the actual site visit, CFD modeling of the airflow in the critical spaces within the site, analysis of the results, recommendations for ECMs and finally estimation of the energy savings which would result from each ECM individually as well as the combined impact of all recommended ECMs. The following tasks were performed for each site:

- HVAC equipment: Assessments of the following units, if applicable.
 - Packaged rooftop units (RTUs)
 - Computer room air conditioning (CRAC) units
 - Wall-packaged units
 - HVAC control system, or building automation system (BAS), where applicable
- Datacom Equipment
 - Equipment racks and rows of racks that may be candidates for aisle containment
- Assessment tasks
 - HVAC
 - Review and document as-programmed controls sequences of operations for the HVAC systems and equipment;
 - Document nameplate data and physical condition of the installed HVAC equipment;
 - Review and document building automation system (BAS) user interface, graphics access, and overall system capabilities (if applicable);
 - Document location of supply and return air diffusers;
 - Document identified solutions for analysis and consideration.
 - IT Equipment
 - Statistically sample and verify rack electronics to help quantify CFD model confidence.
 - Verify floor plan and rack layout on-site against existing equipment lists and site diagrams;
 - Document blanking panels utilized on-site (via pictures/on site estimates);
 - Capture thermal images to document/verify hot zones with larger critical equipment;
 - Document CFM and temperature differential for these equipment types for CFD modeling; and
 - Measure the CFM, temperature, velocity of each mass inlet and outlet locations such as perforated tiles, return and other supply vents.
 - Measurement of the vent dimensions and the calculation of vent free-area.
 - Locations of the thermostats and other external temperature sensors.

For the CFD modeling, analysis and airflow optimization recommendations, the pre-audit data such as floor plans and equipment lists from Comcast databases were verified and corrected if needed during the site audits. Additional information on the HVAC units and supply and return temperatures/CFM were

collected while on site. All this information was input into a commercial CFD modeling tool. Next, thermal images captured at the site were used to calibrate the CFD model. The thermal images were used to adjust the power consumption (heat load) of the IT equipment as recorded in the Comcast databases to better match the conditions at the site. In addition, annual kWh existing in the form of utility bills were cross referenced as a sanity check on the overall calibrated IT load for each site. Next, several models with various AFO ECMs were created. The most viable model for optimally producing energy savings while maintaining performance was selected for each of the sites. Finally, a scenario with increased set point temperature was created such that all rack inlet temperatures are still within the ASHRAE recommended range (64°F-80°F) for realizing the energy savings resulting from airflow optimization. The following flow chart summarizes the procedure:

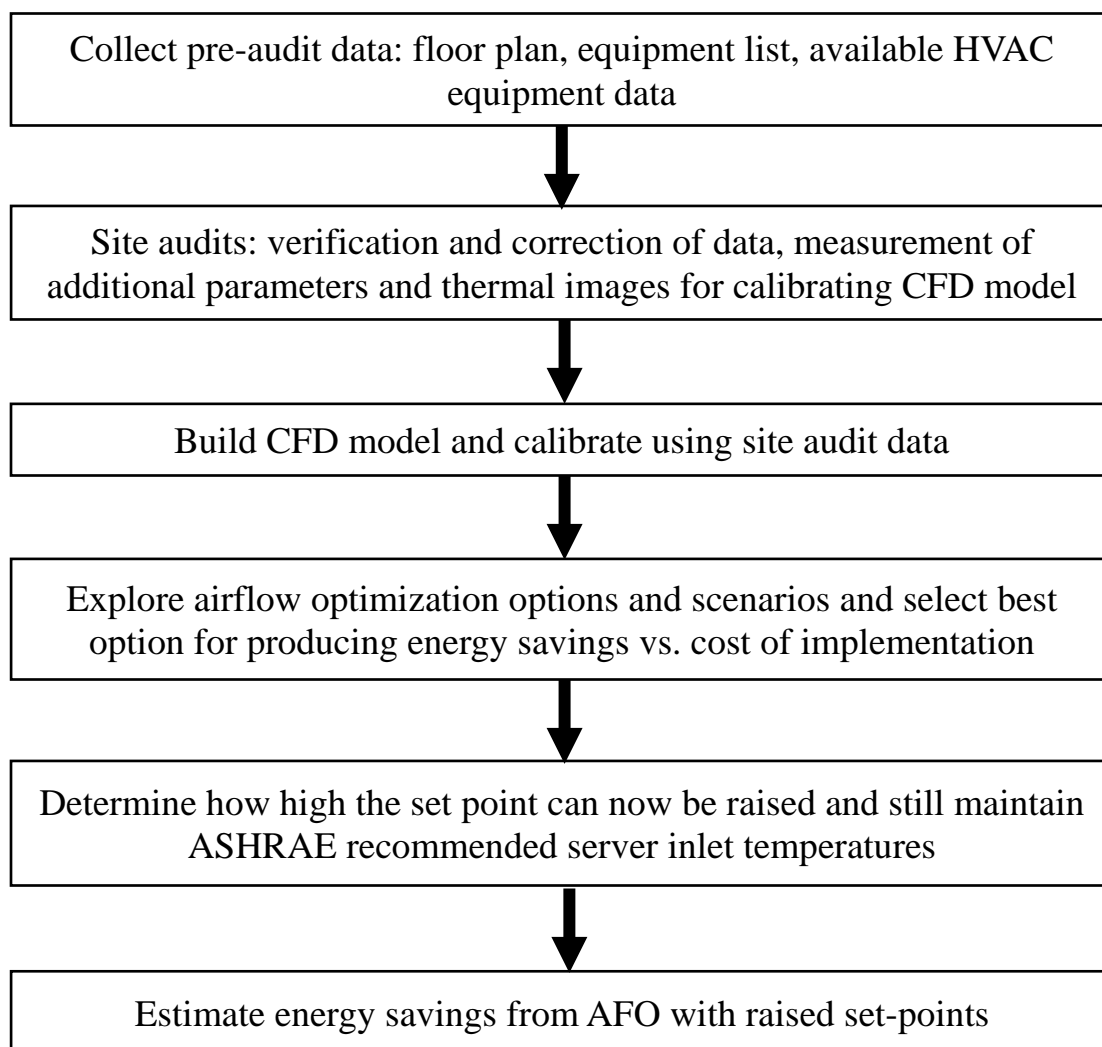


Figure 1 - Process flow chart for airflow optimization and energy savings estimation.

2.2. Example Site Assessment for Beaverton, OR

In the interest of space, only one example of a detailed site assessments and ECM recommendations will be provided here. The Beaverton facility is a one-story building located just west of Portland, OR with two headend rooms (Phase 1 and Phase 2) and an administration/office space. The building is typically occupied by Comcast staff and contractors at least 12 hours a day during weekdays with marginal occupancy on weekends.

The headend space has two (2) zones with critical equipment: “Phase 1” and “Phase 2”.

1. Phase 1 headend is comprised of 36 racks of local market equipment and is cooled by two (2) 16.5 ton Liebert CRAC units. The space has both hot-cold and mixed aisle configurations. Currently neither end-of-aisle containment nor rack containment is in place. The Phase 1 headend is cooled by two (2) – 16.5 ton Liebert CRAC up-flow units. There is no ducting; the air is simply directed from the Liebert units towards the equipment aisles.
2. Phase 2 headend is comprised of 328 racks of local market and regional data center equipment and is cooled by eight (8) 31.5 ton up-flow Liebert CRAC units. The space has limited hot/cold aisles with most of the aisles being mixed aisles where exhaust from one rack row can flow into the front of the inlet side of the adjacent rack row. There is no end of aisle containment and limited use of blanking panels. All the CRAC units have supply ducting, however there is still a lot of heat buildup, primarily due to the long distance from the sources of the heat to the returns of the Liebert CRACs. Heat also builds up since all the returns to the CRACs are at floor level and not necessarily lined up with any specific hot aisle.

The administration/office space is cooled with four roof top units (RTUs).

Figure 2 and Figure 3 below shows the current layout of the Phase 1 and 2 spaces utilized for CFD modeling, respectively, with labels depicting hot, cold and mixed aisles.

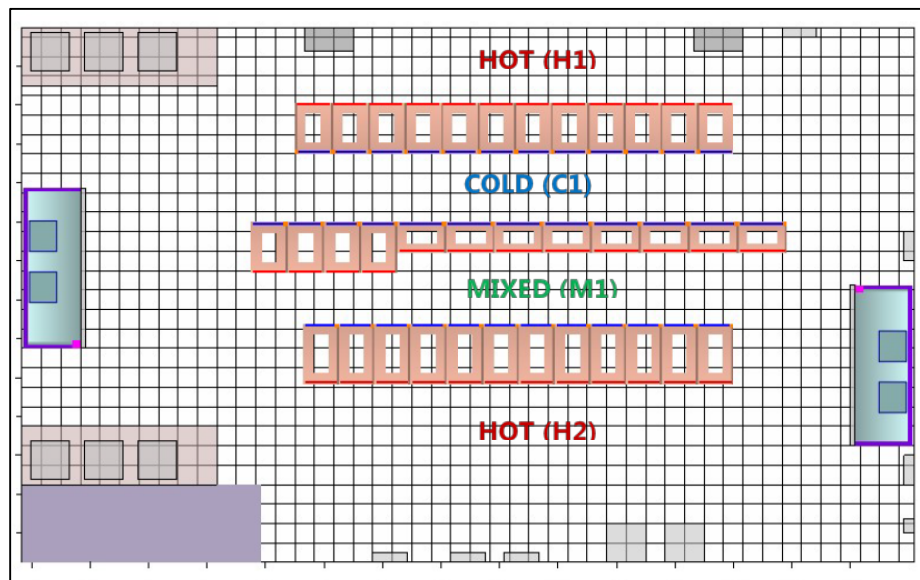


Figure 2 - Layout of the Beaverton headend – Phase 1 headend room.




Figure 3 - Layout of the Beaverton headend – Phase 2 headend room.

It should be noted that it is extremely common to find mixed aisles in telecommunications edge facilities. The fact that hot/cold aisle discipline is beginning to be established in this example facility puts it ahead of the curve and enables the kind of energy savings that are sought in these edge facilities.

Table 3 below summarizes the overall facility specifications for the Beaverton headend facility. Note that the stated IT load in Table 3 is based on summing the equipment power consumption values from the Comcast database prior to the onsite audit. Since many of the chassis of the larger IT devices are only partially populated and actual power consumption is typically less than the nameplate value, the database power value is generally higher, and can often be significantly higher than the actual IT load. For this facility the calibrated IT load determined after the site audit was calculated to be 446 kW instead of 1,100 kW. It is not uncommon for actual site IT loads to be on the order of half the nameplate/stated IT loads, however exceptions do occur, thus a methodical procedure should be used for each site modeled.

Table 3 - General facility summary for Beaverton, OR.

Gross Building Size	13,737 ft ²
Spaces/Zones	<ul style="list-style-type: none">• Headend: Phase 1 (local market) and Phase 2 (local market and regional datacenter)• AC and DC power rooms• HVAC utility room• Administrative (office, storage, restrooms, technical workspace)
Critical Net Floor Area	Phase 1: 1,083 ft ² Phase 2: 7,040 ft ² Total: 8,123 ft²
Stated IT Load	1,100 kW (Phase 1 and 2)
Calibrated IT Load	Phase 1: 33 kW Phase 2: 413 kW Total: 446 kW

2.2.1. Site Findings

2.2.1.1. HVAC Systems

The Phase I headend space uses hot aisle/cold aisle configuration in many areas, however the Phase II headend space has a significant number of mixed and non-contiguous aisles that creates challenges for AFO-based energy savings that have acceptable payback periods. This is because of the large number of blanking panels needed, plus the end aisle containment required. Further, even in well-defined hot and cold aisles, the CFD modeling revealed that hot exhaust can still pass through large rack openings to get into the cold aisle and mix with the supply air, thereby raising the intake temperatures. Blanking panels would definitely help but would reduce mixing much more effectively if coupled with a more consistent hot/cold aisle configuration, which may require either moving equipment and/or racks in the near term, or alternating waiting to add blanking panels until the existing process of decommissioning IT equipment and adding new equipment in proper hot/cold aisle manner plays out sufficiently to ensure the blanking panels accomplish the AFO goals.

All the CRAC units use either of type R-22 or R-407C refrigerants in their direct expansion (“DX”) cooling circuits, which allows the opportunity to improve both energy efficiency as well as eliminate ozone-depleting older refrigerants via replacement of these refrigerants with next generation types.

The site HVAC system types and conditions are summarized in Table 4 below.

Table 4 - HVAC system type and conditions for Beaverton, OR.

Primary Cooling Systems	<ul style="list-style-type: none"> Headend – Phase 1 room is cooled with two (2) Computer Room Air Conditioner (CRAC) units. Headend – Phase 2 room is cooled with eight (8) Computer Room Air Conditioner (CRAC) units. Admin space is cooled with four roof top units (RTUs). 				
Primary Heating Systems	2 - Trane heat pumps and 2 - Trane gas-fired RTUs				
Air Distribution	Phase 1 room: no ducting; Phase 2 room: Single integrated ducting system tied into all CRAC units				
HVAC Redundancy	Phase 1 room: numerical only; Phase 2 room: yes				
Controls	No advanced HVAC controls in use				
HVAC Equipment	Unit	Date	Age (years)	Tons	Refrigerant
	CRU-1	Aug-99	18	16.50	R-22
	CRU-2	Sep-99	18	16.50	R-22
	CRU-3	Sep-00	17	31.50	R-22
	CRU-4	Sep-00	17	31.50	R-22
	CRU-5	Sep-00	17	31.50	R-22
	CRU-6	Sep-00	17	31.50	R-22
	CRU-7	Sep-00	17	31.50	R-22
	CRU-8	Sep-00	17	31.50	R-22
	CRU-9	Sep-00	17	31.50	R22
	CRU-10	Apr-10	7	30.00	R-407C
	HPU-SR1	May-01	16	7.50	R-22
	HPU-SR2	May-01	16	7.50	R-22
	RTU-PWR2	Mar-01	16	10.00	R-22
	RTU-PWR1	Mar-01	16	10.00	R-22
Total Air Flow Demand*	Phase 1: ~10,679 CFM (CRAC Units) Phase 2: ~59,708 CFM (CRAC Units)				
Total Air Flow Supply	Phase 1: ~16,800 CFM (Based on CRAC unit nominal capacity) Phase 2: ~113,450 CFM (Based on CRAC unit nominal capacity)				

*Approximate calculation based on the IT load and an average increase of 20°F across the IT equipment

2.2.1.2. HVAC Findings

An energy audit of the site to support CFD modeling and develop ECM recommendations was performed on January 23-25, 2017. In addition to verifying the HVAC and IT equipment in use, as well as presence and location of ducting and thermostats, the condition and performance of the HVAC systems were measured and observed. Infrared thermal images of hotspots and other strong sources of heat, and supply air temperatures and flow measurements at the diffusers were also gathered to help calibrate the CFD models.

The following summarizes observations from the site audit on current site conditions and opportunities for improvement:

1. The headend rooms (Phase 1 and Phase 2) both have considerably more tons of cooling than the IT equipment load would mandate yet are challenged in maintaining a uniform cool temperature across the room. The Phase 1 room has over three times the cooling tonnage needed for the IT heat load, while the Phase 2 room has over twice the cooling tonnage required for the IT heat load and with its ducting configuration does have redundancy. The overcooling is common in edge facilities due to the significant amount of mixing of hot and cold air in these facilities.
2. There appears to be a mismatch between the rack airflow demand and supply in certain areas, which is resulting in hotter inlet temperatures in those areas (as captured by thermal images).
3. A reoccurring inefficiency is external hot to cold aisle rack recirculation, i.e., the removal or transport of heat produced back to the return intakes of the CRAC units. In many cases, the heat appears to spill over the top of the racks, or goes through and around them into the next aisle, compounding the problem.
4. The inefficiency in removing heat from the room is also demonstrated by the temperature of the return air to the CRAC units which has a range of 66⁰ F to 75⁰ F with an average of 71.5⁰ F.
5. The headend space contains some hot/cold aisle configuration, but still has many mixed aisles with racks aligned in the same direction and often right next to a hot/cold aisle.
6. The Phase 1 room lacks true HVAC redundancy; if one of the CRAC units fail or lose power, equipment in certain zones will not get the sufficient cooling they need to avoid overheating. This is also common in telecommunications edge facilities.
7. The Phase 2 room has true HVAC redundancy, but still lacks ducting to every aisle, thereby reducing the effectiveness of the HVAC redundancy.
8. Both headend rooms lack sufficient blanking panels and other measures to contain all racks.

2.2.1.3. Energy Profile

Table 5 shows a summary of the Beaverton headend energy profile collected from utility data, equipment data from the Comcast database and later calibrated, and information collected on-site.

Table 5 - Summary of energy supply/demand at Beaverton, OR.

Energy Sources	Electricity (Grid)
Calibrated IT Load	Phase 1: 33 kW Phase 2: 413 kW Total: 446 kW
Sub-Meter Data Availability	None available
Energy Utility Provider(s)	Portland General Electric
Baseline Annual Electricity Consumption (kWh)	6,404,800
Utility Rate (\$/kWh)	\$0.074
Annual Electricity Cost	\$474,319
Baseline Annual Natural Gas Consumption (therms)	N/A
Utility Rate (\$/therm)	N/A
Annual Natural Gas Cost	N/A
Estimated PUE*	1.73
*PUE listed is an estimate based on utility bills and calibrated IT load. To calculate actual PUE, sub-metering of IT and headend space is required.	

Hitachi Consulting and Comcast compiled monthly energy utility consumption and cost through past bill requests from utility providers. The chart in Figure 4 below illustrates the trend of electric use and demand for 24 recent months. There is an overall upward trend in energy consumption over the two years shown in the chart: the average daily kWh per month in January-March 2017 is about 18% higher than the values in January-March 2015.




Figure 4 - Historical monthly utility consumption for Beaverton, OR.

Table 6 below summarizes the measured and estimated values for electrical load types found in the facility and Figure 5 depicts the relative consumption amounts. The critical IT load accounts for most of and the electricity demand and consumption at the facility. This load was calculated utilizing equipment data for the site, calibrated by thermal images through the CFD model.

Table 6 - Energy load and consumption by building system for Beaverton, OR.

Baseline Electrical Load	Load (kW)	kWh
Critical IT	446	3,907,000
HVAC (compressors, fans, etc.)	240	2,104,000
Lighting	8	74,000
Other Load (UPS losses, Plug Load)	37	320,000
Total	731	6,405,000




Figure 5 - Energy load percentage by building system for Beaverton, OR.

Note that since the IT and/or HVAC energy consumption is not currently monitored at the facility, it is only possible to estimate the current PUE of the facility using the calibrated IT heat load of 446 kW and a total average load of the headend space of 731 kW. These numbers result in an estimated annualized PUE of 1.73, which means the facility presents a solid opportunity for energy efficiency improvements. The recommended installation of ECMs could help reduce the PUE value. Making use of its metering capabilities, the installation of advanced HVAC controls on all HVAC systems at the facility would also permit an accurate PUE to be determined not just as an annual average, but throughout the year. The sensitivity of the facility energy efficiency to many factors such as outdoor temperature as well as IT load changes could be accurately monitored, as could also the health of the HVAC systems and the impact of the energy efficiency improvements.

2.2.2. Computational Fluid Dynamics (CFD) Modeling Results

2.2.2.1. Baseline CFD Model: Phase 1 Headend

Figure 6 below shows the CFD model results for existing airflow conditions in the Phase 1 headend room at Beaverton, OR. The baseline model was calibrated based on the thermal images captured at the site. A wide range of temperatures across rack inlets were observed due to the lack of containment and lack of hot/cold aisle discipline. The mixed aisle M1 between rows 16 and 17 contains several CMTS devices, the intakes of which are exposed to the hot exhaust of the equipment in aisle 17. Fortunately, since the CMTS devices are located at the bottom of the racks in row 16, they are still getting relatively cool air, as can be seen in the thermal image below in Figure 7. However, if additional devices are added to these racks towards the top of the rack, they will be exposed to the hot exhaust more directly. Alternately, if

additional equipment is added to row 17 such that additional hot air exhaust is fed into aisle M1 and no airflow optimization is deployed, it is possible that the CMTS devices would have significantly increased intake temperatures.

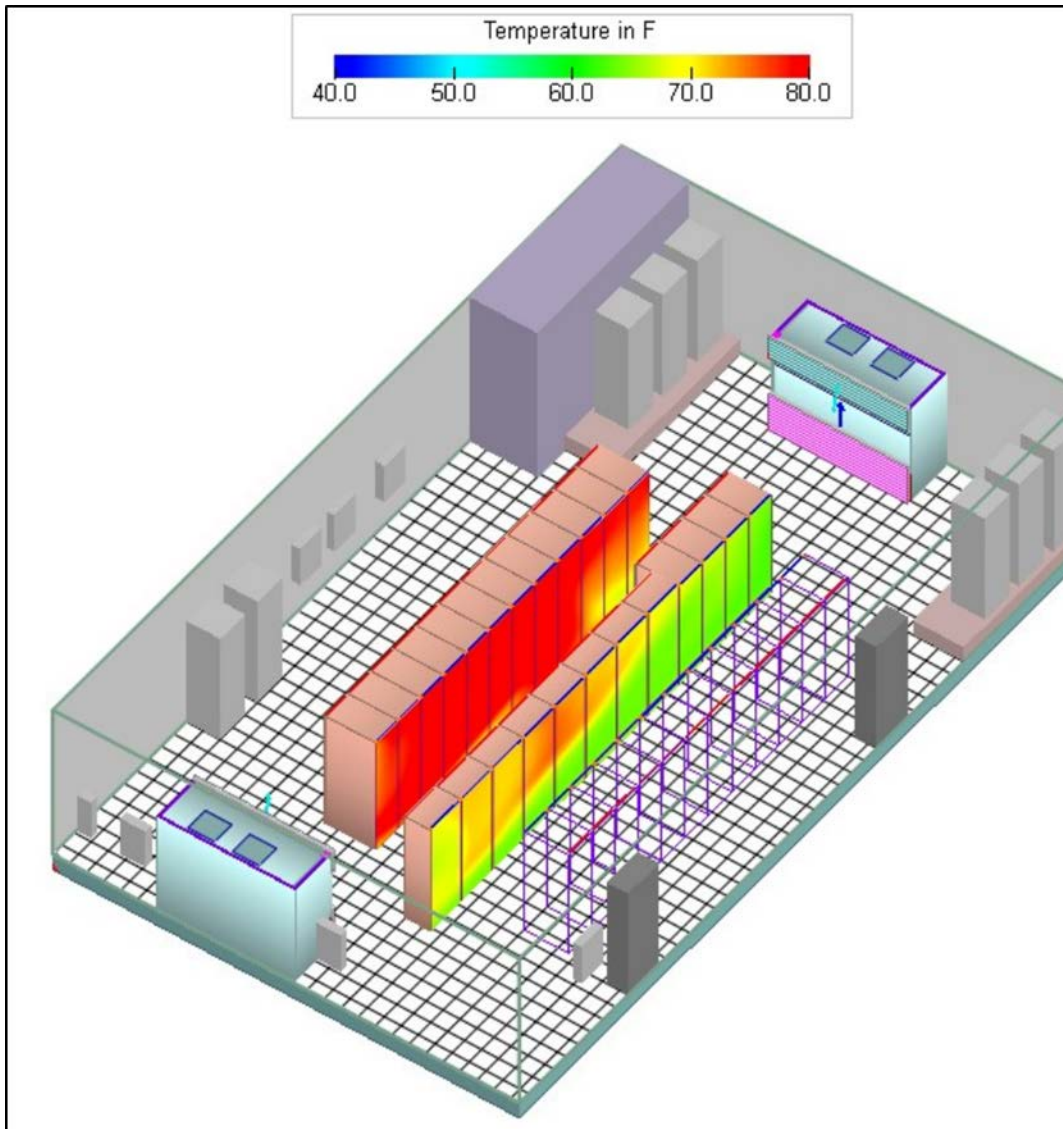


Figure 6 - 3-D model of rack inlet temperature distribution for Beaverton, OR – Phase 1 headend room.




Figure 7 - Thermal (left) and standard (right) images of rack rows 16 and 17 in the Phase 1 headend room.

2.2.2.2. Air Flow Optimization ECM Recommendations: Phase 1 Headend

Hitachi Consulting utilized CFD modeling to design and optimize the airflow by iteratively building in various efficiency measures which allowed for the channeling of hot air more efficiently toward CRAC unit return vents and away from equipment inlets. Better cooling for critical equipment and elimination of hot spots was the overarching goal and once the areas exhibiting these issues were corrected and considerably improved, set points could be raised. For the Beaverton Phase 1 headend, Hitachi Consulting recommended the following AFO measures:

- Top and side containment of both the cold aisle C1 and the mixed aisle M1 via:
 - Addition of blanking panels;
 - Addition of top rack containment panels.

Figure 8 below shows a CFD model screen capture of the improvement in the rack inlet temperature across the Phase 1 headend facility after the AFO recommendations are implemented. The main aspect to note is the decrease in recirculation of hot-air over the tops of racks back into the cold (C1) and mixed (M1) aisles. These AFO recommendations result in uniform rack inlet temperatures across all racks and better containment of cold and mixed aisles to eliminate hot spots.




Figure 8 - 3-D model of rack inlet temperature distribution with AFO recommendations for Beaverton, OR – Phase 1 headend room.

2.2.2.3. Achieving AFO Energy Savings: Phase 1 Headend

With hot spot elimination and uniform distribution of rack inlet temperatures based on Hitachi Consulting's AFO ECM recommendations, the set point temperature can now be raised 6°F to achieve energy savings. Figure 9 below shows the CFD modeling results of raising the set point temperatures after the recommended AFO ECMs are implemented. Note that the recommendations include movement of the thermostats to the aisles. As shown in Figure 9 below, the temperatures of the hot aisles have increased. However, the cold and mixed aisles are still maintained, the rack inlet temperature distribution is uniform and all rack inlet temperatures are still within the ASHRAE recommended range for class A1 to A4 data

center spaces (64°F-80°F). Hitachi recommends that set point be raised gradually (1-2°F per day) to avoid any alarms to the equipment.




Figure 9 - 3-D model of rack inlet temperature distribution with AFO recommendations and after raising set point 6 °F for Beaverton, OR – Phase 1 headend room.

2.2.2.4. Summary of AFO Impact: Phase 1 Headend

Figure 10 below shows a summary view of Beaverton Phase 1 airflow: (1) at current airflow baseline, (2) optimized after AFO implementation, and (3) optimized after AFO implementation with an increased set point by 6°F (while maintaining inlet temperatures under 75°F throughout the racks). Table 9 shows statistically how the rack inlet temperature changed before and after AFO implementation.

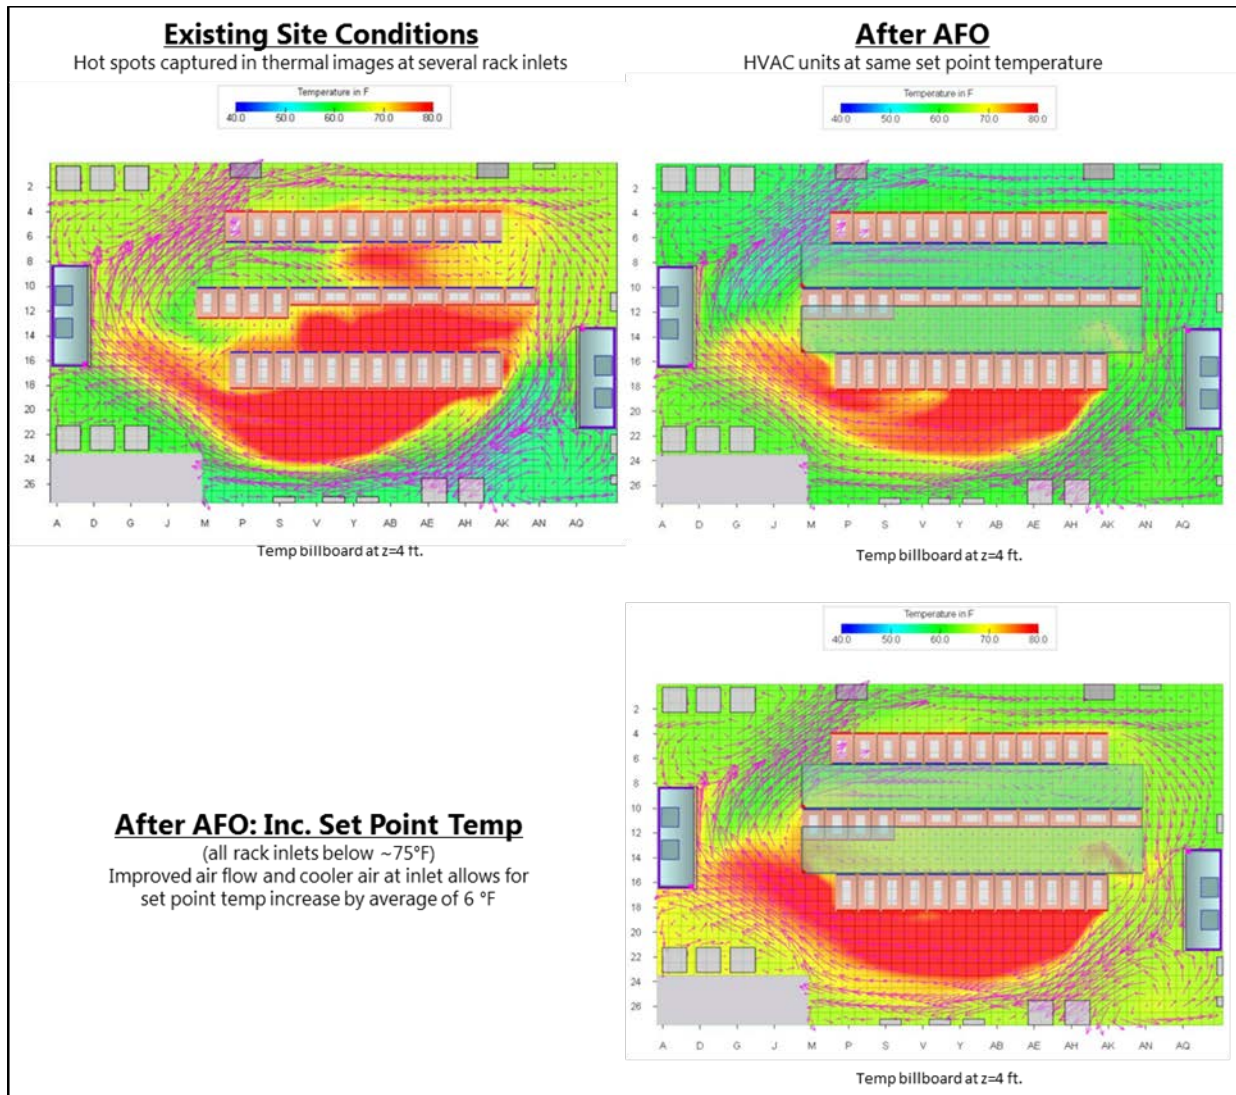


Figure 10 - Temperature distribution at 4 feet for Beaverton Phase 1 headend room, including (1) baseline, (2) + AFO recommendations, and (3) + AFO recommendations and raised temperature set point.

Table 7 - AFO statistical impact on rack inlet temperature distribution after 6°F set point increase for Beaverton Phase 1 headend room.

Range of Max Inlet Temperature	Number of Racks*		
	Baseline	After AFO ECMs	After AFO ECMs and 6°F Set Point Increase
Above 80°F	9	0	0
Between 75°F and 80°F	6	0	0
Between 70°F and 75°F	9	0	7
Below 70°F	6	30	23
Total	30	30	30

*Does not include 6 empty racks that do not have any equipment (total 36 racks).

2.2.2.5. Baseline CFD Model: Phase 2 Headend

Figure 11 below shows the CFD model results for existing airflow conditions at Beaverton, OR – Phase 2 headend. The baseline model was calibrated based on the thermal images and CFM measurements captured at the site. A few hot spots can be seen in the figure, as well as a few racks with significantly lower inlet temperatures than most of the racks in the room.




Figure 11 - 3-D model of rack inlet temperature distribution for Beaverton, OR – Phase 2 headend room.

2.2.2.6. Air Flow Optimization ECM Recommendations: Phase 2 Headend

Hitachi Consulting utilized CFD modeling to design and optimizes the airflow by iteratively building in various efficiency measures which should allow for the channeling of hot air more efficiently toward CRAC unit return vents and away from equipment inlets. Better cooling for critical equipment and elimination of hot spots was the overarching design goal. This site exhibited some challenging hot spots that would require some row and rack reconfiguration, combined with AFO measures to be fully effective. For the Beaverton Phase 2 headend, Hitachi Consulting recommends the following AFO measures:

- Full containment of all racks via blanking panels to prevent recirculation of heat within racks
- Top containment over areas of high heat production to reduce its effect on adjacent aisles

Figure 12 below shows CFD model capture of the improvement in the rack inlet temperature across Phase 2 headend room after the AFO recommendations are implemented. AFO recommendations result in more uniform rack inlet temperatures across all racks and better containment of cold aisles to eliminate hot spots.




Figure 12 - 3-D model of rack inlet temperature distribution with AFO recommendations for Beaverton, OR – Phase 2 headend room.

2.2.2.7. Achieving AFO Energy Savings: Phase 2 Headend

Achieving energy savings and an acceptable payback of AFO ECM implementation was challenging for the Beaverton Phase 2 headend room due to the significant investment required in blanking panels and containment and the fact that there are only one out of eleven contiguous rows with a hot and cold aisle configuration. While extreme hotspots were indeed eliminated by the AFO ECMs, analysis of the modeling results showed that there were still a significant number of racks with inlet temperatures above 80 °F. Consequently, it was not recommended to raise the set points in this headend, and that means there will be no energy savings from the airflow optimization, rather only performance improvements.

The suspected causes for the AFO ECMs being insufficient to permit raising the set points in this situation are:

1. The depth or length of the aisles is roughly 88 feet and while the supply is ducted to bring cold air to the areas, there is no medium or return ducting to remove the heat or bring it back to the returns of the CRACs.
 - a. This is evidenced by the average return temperature of 70°F, with a range of 66°F to 75°F
2. Insufficient hot/cold aisle configuration with mixed aisles interspersed throughout the headend.
3. The Liebert CRAC units are designed to have the return air inlets at floor level which makes it difficult to remove the heat. Often, due to the nature of thermal stratification, the low return grilles pull in the cool air meant for the IT equipment.
4. The heat produced from the equipment in the center of rows either stagnates or recirculates over the tops of racks, creating hot air that cannot easily move back to the CRAC return vents.

The performance improvements from adding blanking panels throughout the facility nonetheless significant since they create channels and effectively block exhaust air from recirculating within the racks which helps the overall cooling. Should Comcast decide to install the blanking panels and top containment, it was recommended that a more detailed equipment and HVAC audit and review of this site be performed to confirm that the conditions of this room are such that achieving significant energy savings and further airflow optimization may require redesign and renovation of the room. A second examination would be required if a large increase in the IT heat load is planned. Such redesign and renovation was outside the scope of the present effort.

2.2.2.8. Summary of AFO Impact: Phase 2 Headend

Implementing the recommended AFO ECMs eliminates severe hot spots and makes the temperature distributions and rack inlet temperatures more uniform. Figure 13 below shows a summary view of the Phase 2 room air flow: (1) at current airflow baseline and (2) optimized after AFO implementation. Table 8 shows statistically how the rack inlet temperature changed before and after AFO implementation.




Figure 13 - Temperature distribution at 4 feet for Beaverton Phase 2 headend room, including (1) baseline and (2) + AFO recommendations.

Table 8 - AFO statistical impact on rack inlet temperature distribution after AFO ECM implementation for Beaverton Phase 2 headend.

Range of Max Inlet Temperature	Number of Racks*	
	Baseline	After AFO ECMs
Above 80°F	62	58
Between 75°F and 80°F	89	84
Between 70°F and 75°F	65	72
Below 70°F	16	18
Total	232	232

*Does not include 96 empty racks that do not have any equipment (total 328 racks)

2.2.3. HVAC Energy Conservation Measure Recommendations

In addition to the AFO ECMs just presented, the following additional energy conservation measures are recommended as a result of the site visit and energy audit.

2.2.3.1. Addition of Advanced HVAC Controls

Advanced HVAC controls are recommended for all ten (10) CRAC and four (4) RTU units in the facility to provide complete HVAC monitoring/health checks, and to increase energy efficiency of the affected CRACs. HVAC controls are installed by licensed mechanical contractors, usually recommended by Comcast and local to the site. The installation is simple and the local mechanical contractors are trained for about an hour and then overseen by the subject matter expert (SME) on the first one or two installs. The install takes about an hour for the first one, but thereafter, the process should only take about 30 minutes. The controller node has power connections and a clamp for the energy monitoring. Temperature probes are put in the supply air, return air, and by the furnace exhaust.

The advanced HVAC controls work by optimizing key components of the HVAC system to reduce HVAC energy consumption by 15-25%. Since the controls measure the supply air, return air, and energy usage, the unit can also tell when the equipment begins to fail to provide the cooling necessary or is not operating at 100% of its nominal capability. Thus, advanced HVAC controls also provide a health-check functionality.

2.2.3.2. Refrigerant Replacement

A nextgen, more efficient refrigerant is recommended to replace all R-22 and R-407C refrigerants currently in use at the facility. The new refrigerant in each of the units will improve their lifespan, effectively increase the capacity of each unit and will increase energy efficiency and thereby provide energy savings. In addition, this ECM also eliminates ozone-depleting refrigerants at the facility making the facility compliant to upcoming regulations.

The R-22 and R407C refrigerant is reclaimed and replaced by licensed mechanical contractors, usually one recommended by the cable operator and local to the site. The refrigerant replacement process entails reclaiming the R-22 or R-407C refrigerant by vacuuming the coolant lines until the R-22 or R-407C reaches approximately 500 microns. The system is then charged with the nextgen refrigerant per the nextgen pressure temperature chart.

2.2.4. Summary of ECM Recommendations

Table 9 provides a summary of the recommended ECMs for the Beaverton headend facility.

Table 9 - Summary of ECM recommendations for Beaverton, OR.

Space	Air Flow Optimization	Advanced HVAC Controls	Refrigerant Replacement
Phase 1 Headend	<ul style="list-style-type: none"> Containment: add ceiling panels and blanking panels Raise set point 6 degrees. 	<ul style="list-style-type: none"> Add HVAC controls to both CRAC units 	<ul style="list-style-type: none"> Replace refrigerant in both CRAC units
Phase 2 Headend	<ul style="list-style-type: none"> Containment: add ceiling panels and blanking panels Raise of set point temperature not recommended, resulting in no predicted savings from AFO for the site. 	<ul style="list-style-type: none"> Add HVAC controls to all 8 CRAC units 	<ul style="list-style-type: none"> Replace refrigerant in all 8 CRAC units
Power Room	N/A	<ul style="list-style-type: none"> Add HVAC controls to (2) Trane RTU Units: RTU-PWR1-2 	<ul style="list-style-type: none"> Replace refrigerant in (2) Trane RTU Units: RTU-PWR1-2
Administrative Areas	N/A	<ul style="list-style-type: none"> Add HVAC controls to (2) Trane RTU Units: HPU-SR1-2 	<ul style="list-style-type: none"> Replace refrigerant in (2) Trane RTU Units: HPU-SR1-2

2.2.5. Summary for Beaverton, OR Headend

Deployment of the three recommended ECMs would provide a total annual HVAC energy reduction of over 479,000 kWh, thereby reducing the HVAC energy consumption by 23% and improving the power margin for the facility. Other benefits of the recommended ECMs are summarized in the next section of this report.

3. Portfolio Analysis and Recommendations

The Hitachi Consulting team completed an on-site energy and equipment audit of each of the ten headend sites similarly to the one described above for Beaverton, OR. In this section, an analysis of the entire portfolio will be presented, with trends and conclusions based on examining all ten sites.

3.1. Airflow Optimization Issues Across the Portfolio

The Hitachi Consulting team developed CFD models for each site depicting baseline air flow conditions at the facility, and identified technically and financially feasible ECMs for Comcast's consideration. The CFD modeling along with the site visits showed inconsistency in the aisle temperatures, limited hot/cold aisle discipline and generally that the spaces were overcooled to try and compensate for mixing and areas of heat. Even with this general overcooling, the CFD modeling results across the portfolio identified inlet side areas of the racks with temperatures of up to 95°F with no alarms being triggered. The following are common issues that were identified at the headends across the portfolio of ten sites covered in this effort.

3.1.1. *Insufficient Use of Blanking Panels*

Substantial portion of the empty rack units (RUs) in most of the head ends were un-blanked. Un-blanked RU spaces in the racks lead to recirculation of hot air inside of the racks. The red circle in the Figure 14 below shows this recirculation of hot air in a rack at a head end that was audited during the site visits.




Figure 14 - Recirculation of Hot Air within the Racks – in Absence of Blanking Panels (x-z plane)

3.1.2. *No Hot- Cold Aisle Configuration*

Very few facilities that were audited had complete hot/cold aisle configurations which results in mixing of hot exhaust air with the cold air being supplied to the racks. Figure 15 below shows the typical configuration noticed at the audited head end sites. The exhaust of one rack row faces the inlets of an adjacent rack row, leading to hot spots on the inlet side of those racks.




Figure 15 - Hot Exhaust Air Blowing onto the Inlet of the Racks (x-z plane)

3.1.3. Absence of Aisle Containment

Most cold aisles at the audited head end sites were not contained. Absence of side containment lets hot air infiltrate around the sides of the racks into the cold aisle. Absence of top aisle containment leads to hot air looping into the cold aisle from the top of the racks. Adding top and end aisle containment can significantly reduce the mixing of hot exhaust air with the cold air being supplied to the servers. Figure 16 (a) is a plan view capture of rows showing infiltration of hot air looping around the side of the aisles. The red circle highlights this hot air looping in. The red circle in the Figure 16 (b) below shows the infiltration of hot exhaust air into the cold aisle from the top of the racks. This infiltration is commonly seen in telecommunications edge facilities.




Figure 16 - (a) Plan view of racks showing hot air infiltration from the sides (x-y plane) (b) Infiltration of hot air from the top of racks (x-z plane)

3.1.4. Summary of AFO Recommendations

Considering all ten sites covered by this effort, the following airflow optimization recommendations apply in some manner to all ten sites:

- Increased utilization of blanking panels in all racks to limit the hot and cold air to a specific space and limit infiltration within and without the racks
- Addition of top containment where needed to limit infiltration of hot air from over the top of the aisles
- Addition of end aisle containment in form of strip curtains or end panels/doors to contain cold aisles and prevent infiltration of hot air
- Addition and/or redirection of supply ducting to focus cold supply air directly into the aisles
- Addition and/or repositioning of return grilles to facilitate hot air return to CRAC units.

It should also be stated that the airflow optimization recommendations presented in this report were developed such that they are resilient to equipment changes at each site, whether to increase or decrease the overall heat load. Further, since the recommendations include placing thermostats in the aisles, this allows controlling any decommissioned space independently and/or closing of diffusers, thereby minimizing the cooling provided to the decommissioned space as a site evolves.

3.2. Presence of R-22 Refrigerant

All sites contained at least one of either R-22, R-407C, R-134A refrigerants, and for many of the sites most HVAC units have older, outdated refrigerants that should be replaced with nextgen refrigerants. Nextgen refrigerants should be a direct drop-in replacement for refrigerants of type R-22, R-407C, and R-134A and improve performance while meeting the DOE standards as acceptable refrigerants. These nextgen refrigerants can increase the HVAC energy efficiency over R-22 and R-407C by as much as 20%, and further give the HVAC system an increased capacity that can also extend the life of the HVAC

units because the compressor does not run as much. Note, most importantly, as of 2020 it will no longer be possible to buy R-22 or install it into HVAC systems.

The number of HVAC units for each site that are recommended to have their refrigerant replaced are listed in Table 12 below.

Table 10 - Number of HVAC units with outdated refrigerants by site.

Facility	Number of HVAC Units with Outdated Refrigerants
Roseville, MN	9
Hayward, CA	16
Santa Clara, CA	10
Beaverton, OR	14
Burien, WA	13
Stone Mountain, GA	7
Atlanta, GA	6
Jonesboro, GA	4
Woodstock, GA	4
Augusta, GA	5
All Facilities	88

While the larger sites generally have more HVAC systems that would benefit from refrigerant replacement, there is no rule of thumb for converting site size or site IT load into a predictable number of HVAC units that require replacement refrigerants. A detailed audit by subject matter experts should be performed to determine precisely which units require, or could benefit from nextgen refrigerants.

As part of replacing the refrigerant, it is good practice to “true-up” the equipment, meaning the mechanical technician will go through the equipment, replacing simple pieces where needed, like contacts, belts (tension or replace), and also check for refrigerant leaks. If there is anything that cannot be fixed or replaced in 15 minutes with parts normally found on the mechanical technician’s truck, it is brought to the attention of the site property manager and scheduled to be fixed immediately such that the deployment of the new technology is not impaired.

3.3. Lack of Efficient Controls

All sites visited had HVAC systems that would benefit from advanced HVAC controls to improve energy efficiency, extend the life of the system, and provide additional sub-metering of HVAC consumption data. Advanced HVAC control units use an algorithm to optimize key HVAC components and consequently the HVAC system uses about 80% of the original HVAC energy consumed. Advanced HVAC controls can also extend the life of the HVAC system.

An added benefit to advanced HVAC controls is when the equipment is replaced, the advanced control units can be reinstalled on the new equipment to provide continued savings and longevity.

Table 11 lists the number of HVAC units for each site that are recommended to have advanced controllers installed.

Table 11 - Number of HVAC units recommended for advanced controllers by site.

Facility	Number of HVAC Units with Outdated Refrigerants
Roseville, MN	12
Hayward, CA	23
Santa Clara, CA	17
Beaverton, OR	14
Burien, WA	13
Stone Mountain, GA	51
Atlanta, GA	12
Jonesboro, GA	6
Woodstock, GA	11
Augusta, GA	10
All Facilities	169

Similarly, to nextgen refrigerants, as part of installing the advanced HVAC controls, it is good practice to have the equipment “true-up”. Since both refrigerant replacement and the installation of advanced controllers should have a system “true-up,” one of the benefits of doing both ECMs at the same time is to reduce the total number of “true-ups” required.

Also, as with refrigerant replacement, while the larger sites generally have more HVAC systems that would benefit from advanced controls, there is no rule of thumb for converting site size or site IT load into a predictable number of HVAC units that should have advanced controllers installed. Again, a site audit by SMEs is the best way to accurately determine how many units would benefit from HVAC advanced controls.

3.4. General Age Issues for HVAC Units

The last two recommendations for installation of nextgen refrigerants and advanced HVAC controls should be tempered by the following considerations for individual sites: history of critical HVAC mechanical issues, history of not attaining set points, general reliability of manufacturer (for example Liebert units often exhibit up to 25 year lifespans), preventive maintenance practices at the site, incentives by state and federal governments for HVAC replacements, and whether a particular site is slated for complete overhaul, expansion, or decommission. For example, if a site is slated for complete decommissioning in the next two years, and the cost of the ECMs recommended results in a payback period that significantly exceeds two years, facility managers may prefer to hold off on the ECMs for that particular site.

However, it should also be noted that the system “true-up” procedure that should be done as part of either refrigerant replacement or advanced controller installation has the added benefit of detecting HVAC system issues in a process-controlled manner that is not service-impacting. Thus the “true-up” can prevent a subsequent HVAC system failure that might otherwise impact service delivery. Nonetheless, the

three main ECMs recommended in this effort met the criteria of significant energy savings with a reasonable payback period.

3.5. Overstated IT heat loads

Since many of the sites lacked sub-metering of IT and HVAC power consumption, equipment lists from the Comcast database were used to estimate the IT equipment heat load for CFD modeling. Unfortunately, the lists for all the sites generally contained nameplate heat load value, and many of the larger energy-consuming devices were only partially populated and thus consume far less energy than their nameplate values. Therefore, the total IT heat load of the facility had to be adjusted or de-rated to match the actual heat load at the site. To accomplish this, thermal images of each row and rack were collected during site audit using a thermal camera, and the CFD model was then calibrated to the thermal images. The following Table 12 shows the resulting derating of IT equipment kW for each of the sites visited.

Table 12 - De-rating of the equipment/server

Site	Stated IT Load - Name Plate Heat Load (kW)	Calibrated IT Load (kW)	De-rating factor
Roseville, MN	470	253	53.8%
Hayward, CA	236	86	36.4%
Santa Clara, CA	460	129	28.0%
Beaverton, OR	1,100	446	40.5%
Burien, WA	531	195	36.7%
Stone Mountain, GA	1,031*	627	60.8%
Atlanta, GA	1,042*	448	43.0%
Jonesboro, GA	131*	68	51.9%
Woodstock, GA	352*	218	61.9%
Augusta, GA	182*	102	56.0%
Total (all sites)	5,535	2,572	46.5%

*Site IT load was corrected based on equipment seen on-site.

Thus, based on the ten sites covered in this effort, when examining the nameplate IT load of a facility that lacks sub-metering for true load measurement, a maximum of 62% of the nameplate IT load should be used for site analysis. It should be noted that the actual IT heat load could be as low as 28% of the nameplate value. This derating of stated IT load is important not only to HVAC optimization, but also to facility powering requirements and planning, and may prevent costly facility powering upgrades that could have been unnecessary.

3.6. Summary of Savings Resulting from the Recommended HVAC ECMs

The potential energy savings associated with implementation of the three ECMs at the ten headend facilities was summarized in Table 1 in Section 1 of this paper.

The benefits of the AFO, controls and refrigerant replacement go beyond energy reduction and cost savings; they also improve power margin, solve the problem of R22 phase-out by 2020 for the site, inconsistent temperatures across the inlet side of the equipment, overheating equipment with no alarms,

redundancy and compressor/HVAC life extension. Table 13 below provides a more complete list of benefits from implementing the proposed ECMs at edge facilities.

Table 13 - Summary of proposed ECM benefits for Comcast critical facilities.

	Meeting New Standards for Facilities	Improving / Maintaining Customer Satisfaction	Lower OpEx Costs		Lower CapEx Costs
			Energy Reduction	Maintenance Reduction	
Airflow Optimization	Move to hot/cold aisle discipline in all facilities	<ul style="list-style-type: none"> • True HVAC redundancy to prevent IT equipment overloads Reduce alarms and outages	<ul style="list-style-type: none"> • Reduced cooling tonnage / number of HVAC units • Eliminates hotspots Permits increasing set point	Fewer HVAC units to maintain	Higher power margin
Advanced Controls	HVAC health check monitoring	Increased visibility of HVAC performance – better able to predict failures and replace accordingly (reduce alarms and outages)	<ul style="list-style-type: none"> • Optimized HVAC runtime Peak demand reduction	<ul style="list-style-type: none"> • Extends HVAC life • Reusable on replacement equipment Fewer truck rolls	Higher power margin
Nextgen Refrigerant Replacement	Regulatory compliance for elimination of ozone-depleting refrigerants	PR benefit for Comcast customers who care about the environment	Increased HVAC capacity of existing systems (reducing compressor run time)	Extends HVAC life and reduces load on existing HVAC units	Higher power margin

Table 2 in Section 1 of this paper showed statistically how the rack inlet temperature changed before and after AFO implementation for all ten headend sites.

3.7. Energy consumption trends

As a final analysis of the entire portfolio, consider the energy consumption trends depicted in Figure 17 below for all ten sites. Note that the three largest sites, Atlanta, Stone Mountain, and Beaverton, all had significant energy consumption growth over the past two years and further the slope of the growth trend line for each is similar. Hayward and Santa Clara also had growth trends, and upon examination of the actual numbers, these two sites had the steepest growth curves from a percentage perspective, even though the absolute growth was overshadowed by the three largest sites. The five remaining sites either had no growth, or in the case of Woodstock, actually decreased energy consumption very slightly.




Figure 17 - Energy consumption trends across the portfolio of ten sites.

3.8. Lighting Opportunities

LED lighting opportunities were also assessed at the ten headends. Implementation of LED lighting and controls would provide a 5-year energy savings opportunity for the ten sites of just over \$300,000, with the annual savings being just over \$60,000. This is without incentives and with occupancy sensors. Incentives are generally available across all utilities, but were not investigated as part of this effort.

Table 14 below depicts the energy savings opportunity with LED lighting retrofit and controls at each of the ten headend sites.

Table 14 - Potential energy savings for LED lighting retrofit at all headend sites.

Facility	Annual Savings	
	Energy Savings (kWh)	% Lighting Energy Reduction
Roseville MN	68,145	76%
Hayward CA	107,265	63%
Santa Clara CA	73,381	71%
Beaverton OR	56,053	76%
Burien WA	29,295	75%
Stone Mountain, GA	137,335	75%
Atlanta, GA	119,977	74%
Jonesboro, GA	36,140	76%
Woodstock, GA	12,079	74%
Augusta, GA	12,947	80%
All Facilities	652,617	72%

4. Conclusion

Energy savings are possible for cable edge facilities, even given their diversity, historical development, and changing functionality. The headends analyzed in this study prove that with a methodical approach, these savings can be achieved across an entire portfolio. As demonstrated in the example case presented, this methodical approach means seeking not to impose modern standards at any cost, but rather applying solutions with a keen eye towards payback period, longer term site plans and more traditional benefits of energy conservation measures.

5. Abbreviations

Abbreviation	Definition
AFO	Air flow optimization
AHU	Air handling unit
ANSI	American National Standards Institute
APOP	Alternate point of presence
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
A1-A4	ASHRAE data center classes
ASR	Aggregation services routers
BAS	Building automation system
CapEx	Capital expenditures
CCAP	Converged cable access platform
CFD	Computational fluid dynamics
CFM	Cubic feet per minute
CMTS	Cable modem termination system
CO	Central office
CRAC	Computer room air conditioner
Cx	Cold aisle number
DC	Direct current
DOE	Department of Energy
DX	Direct expansion
ECM	Energy conservation measure
HVAC	Heating, ventilation, and air conditioning
Hx	Hot aisle number
IR	Infrared
IT	Information technology
kW	Kilowatt (unit of power)
kWh	Kilowatt-hour (unit of energy consumption)
LED	Light emitting diode
Mx	Mixed aisle number
NA	Not applicable
Nextgen	Next generation
OpEx	Operating expenses
PDU	Power distribution unit
PR	Public relations
PUE	Power usage effectiveness
RTU	Rooftop unit
RU	Rack unit
R 22, R410A, R 407 C, R134 A	Types of refrigerants
SCTE	Society for Cable Telecommunication Engineers
SME	Subject matter expert
SOW	Statement of work
Tons	HVAC tonnage (unit of HVAC capacity)- 1 ton = 12,000 Btu/hr.

UPS	Uninterruptable power supply
VHE	Video head end
VPC 1/2	Video processing center
XD C/O/R	XD-extreme heat density system C- chiller and pumping unit O- overhead cooling module R – rear cooling module

Network Migration Strategies for the Era of DAA, DOCSIS 3.1, and New Kid on the Block... Full Duplex DOCSIS!

A Technical Paper prepared for SCTE•ISBE by

Ayham Al-Banna

Engineering Fellow

ARRIS

2400 Ogden Ave., Suite 180

Lisle, IL 60532

630-281-3009

ayham.al-banna@arris.com

Tom Cloonan

CTO, Network Solutions

ARRIS

2400 Ogden Ave., Suite 180

Lisle, IL 60532

630-281-3050

thomas.cloonan@arris.com

Jeff Howe

VP, Systems Engineering

ARRIS

2400 Ogden Ave., Suite 180

Lisle, IL 60532

630-281-3124

jeff.howe@arris.com

Introduction

The cable industry has achieved tremendous progress in offering high speed data since the first DOCSIS specification was released in 1997. MSOs had the dual goal of meeting customers' demand for higher speeds and defending itself against competitive threats of speed wars with alternate technologies. As MSOs continue their network evolution, they are currently faced with no clear path since many options are available to augment their existing HFC networks.

For example, Figure 1 shows multiple potential evolutionary paths that the MSOs can select. The network architecture (e.g., I-CCAP/DAA/PON) is plotted against the topology which is presented here as the depth of the fiber in the network (e.g., HFC, FTTLA/FTTC, FTTT, and FTTH).

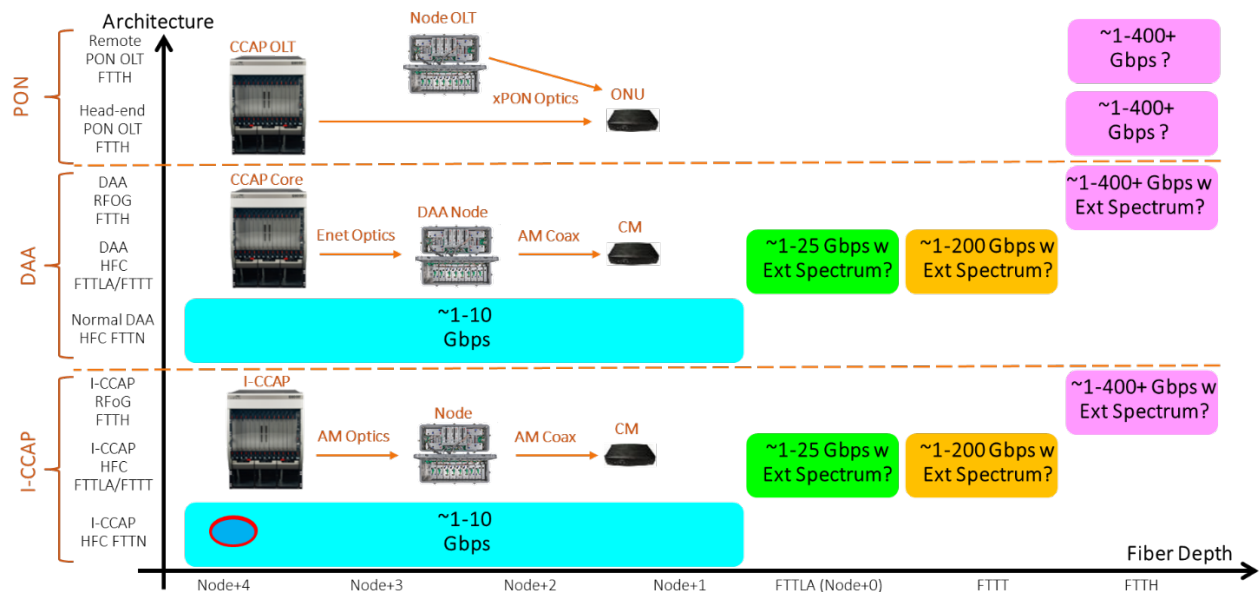


Figure 1 - Evolution of Cable Networks in the Next 2-3 Decades

The transitions between different phases of the same architecture or moving from one architecture to another will depend on the priorities and conditions within an MSO. Different MSOs may select to transition to a particular alternative at different times and different locations. For example, current HFC networks using the normal practice of node splits going to Node+0 (N+0) may be able to continue that practice until year 2025. At the same time, some MSOs may choose to move to an N+0 architecture in the immediate future. Similarly, if an Extended Spectrum DOCSIS technology develops, MSOs may choose to move to FTTT architecture as soon as 2025 in order to access even higher speeds. Finally, it is assumed that most MSOs may eventually choose to migrate their networks to FTTH over the next decade or two. Note that the capacities of all architectures (I-CCAP/DAA/PON) in an FTTH environment in the 2030 time frame are assumed to be similar (~400 Gbps+) because those architectures will likely leverage similar technologies at that time.

Given the large combinations of the various network architectures (I-CCAP/DAA/PON) shown in Figure 1 and different fiber depth topologies, selecting the appropriate architecture/topology transition path is not a trivial task. The challenge at hand is to understand the available technology enablers to assist in

selecting the appropriate transition path. These technology enablers include node splitting, DAA, DOCSIS 3.1, spectrum management and reclamation, FTTx, Selective Subscriber Migration (SSM), extended spectrum DOCSIS, Full Duplex DOCSIS (FDX), and others. This paper will examine the forces that are driving MSOs to provide symmetric multi-Gigabit per second service, the technologies that will assist them in getting to those services, and the factors that will help guide them down the alternative migration paths that are available.

Drivers Behind Gigabit per Second Services

For many years, studies have indicated that Downstream Internet traffic has been experiencing a ~50% compound annual growth rate (CAGR). For almost 35 years, this growth rate has shown itself in the Maximum Downstream Sustained Traffic rates (aka the “Billboard Bandwidths”) that service providers have offered to their subscribers. The 50% CAGR of Maximum Downstream Sustained Traffic rates is often reported as Nielsen’s Law, and is depicted in Figure 2. The same trend has also shown itself (with slightly more variation) in the Average Downstream Bandwidth Consumption rates that subscribers have consumed. Upstream Billboard Bandwidths and Average Upstream Bandwidth Consumption Rates display much more variation, and typically have shown recent CAGRs at different MSOs with growth rates less than the 50% found in the downstream. However, the upstream long term trend has been added to Figure 2. Projecting these curves out over the next 15 years indicates that a significant amount of bandwidth per subscriber is going to be needed, and MSOs need to map out their network migration strategies to meet these needs.

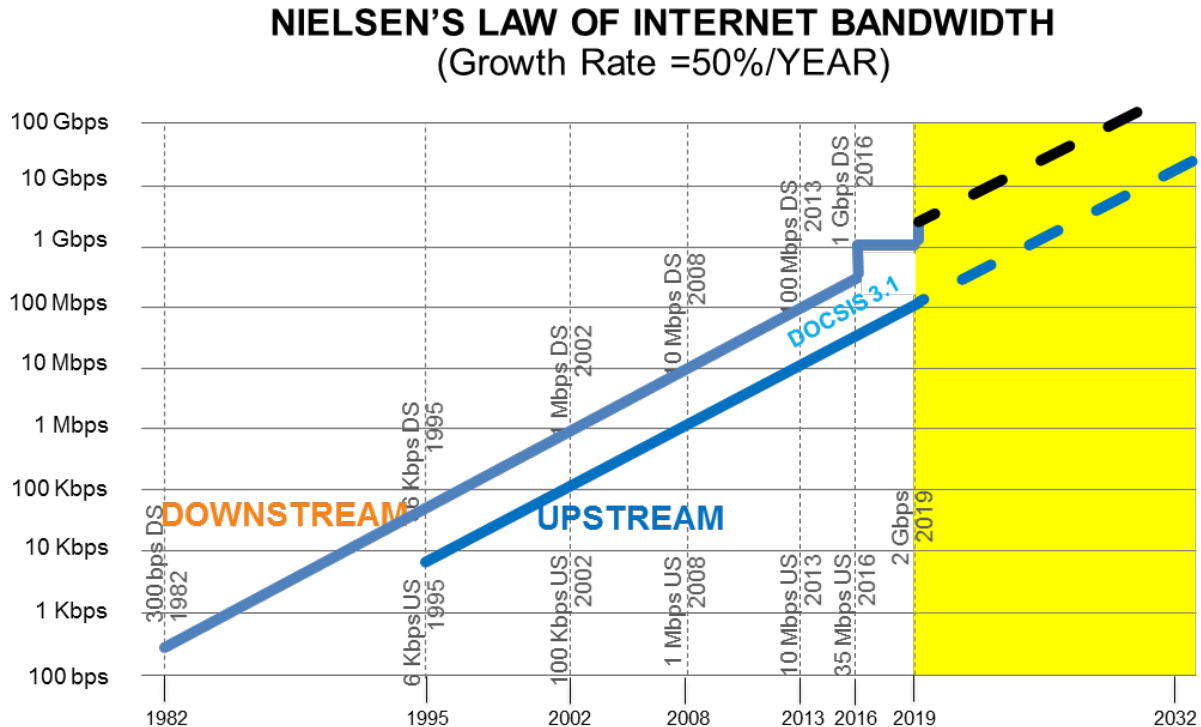


Figure 2 - Nielsen's Law of Internet Bandwidth

Although Nielsen's Law held fairly true to form over the past 35 years, there are indications that factors that may cause it to deviate from its historical trend. One area examine is the growth of upstream bandwidth. Historically, the upstream bandwidth has lagged behind the downstream usage. A lot of this was driven by how subscribers used the Internet. Users primarily accessed content on the Internet and downloaded to their PCs, either doing large file transfers or viewing web pages. The upstream traffic was typically limited to protocol acknowledgements. As a result, access protocols, such as DSL and DOCSIS, evolved along an asymmetric path.

Recent factors are causing a reexamination of this trend. PON technology has been introduced that supports symmetric bandwidth for the upstream and downstream. Although subscribers did not initially have a need for symmetry, MSOs were subjected to competitive pressure from PON providers because symmetric service was something that MSOs could not easily provide. Further, the usage of the Internet itself is seeing a shift. While historically usage was primarily in the downstream direction, new cloud based services such as YouTube that allows users to upload video, and cloud based file storage and backup services is dramatically increasing the demand for upstream bandwidth. This is resulting in a projected discontinuity in the upstream bandwidth curve where the upstream bandwidth will take a step function upward to become close to the downstream curve.

Tempering this projected dramatic jump in upstream bandwidth demands is an apparent slowing in bandwidth usage. Historically, access technology was the limiting factor in usage, in that demand for bandwidth exceeded the ability of the MSOs to provide it. Any time that the service tier was increased, the average usage went up by a corresponding amount. However, data that has been collected from several MSOs appears to indicate that although the "Billboard Bandwidth" continued to grow at the 50% CAGR, the average bandwidth during the current decade did not grow nearly that fast.

During the decade of the 2010s, the average MSO downstream bandwidth usage depicted in Figure 3 grew at a 36% CAGR and the average MSO upstream bandwidth usage depicted in Figure 4 grew at a 17% CAGR. A possible interpretation of this is that technology has finally allowed bandwidth supply to finally catch and exceed the bandwidth demand. Another possible interpretation is that traffic is becoming more bursty, due to an increasing spread between peak and average utilization.




Figure 3 - Average DS Bandwidth in the 2010s




Figure 4 - Average US Bandwidth in the 2010s

If MSO bandwidth supply has indeed caught up and surpassed the demand, MSOs may be able to slow down the growth in their advertised “Billboard Bandwidth” rates. Technologies such as DOCSIS 3.1, which increased the available upstream bandwidth, and the upcoming FDX technology which will increase the upstream bandwidth even more, will address the increase demand for upstream bandwidth. Combining upstream demands with a slowdown in the downstream growth rate to 40% results in a modified Nielsen’s Law curve, shown in Figure 5. While network migrations will present challenges for

MSOs, the modified bandwidth growth curve indicates that HFC technology will remain viable for at least another 15 years.




Figure 5 - Modified Nielsen's Law Curve

Technology Enablers Supporting Bandwidth Expansion

1. Service Group Splits

Service Group (SG) splits, sometimes referred to as node splits, have long been a trusted tool used by MSOs to reduce the bandwidth demands within a Service Group. The basic idea behind the SG split is that it divides the subscribers connected to a single SG into two smaller groups. Ideally, roughly half the subscribers are left connected to the connected in the original SG, while the other half of the subscribers are re-connected to a different new SG.

There are several different ways that an MSO can perform SG splits. In older times, when there were multiple nodes per CMTS SG, the splitting would occur solely in the Headend or hub with the addition of new CMTS ports and reconfiguration of the RF combining network. Once there is a one to one mapping between an SG and a node, the next step is to segment the node into multiple SGs (e.g. 1x1 to 2x2 to 4x4). This is often accomplished using Wavelength Division Multiplexing. Today, many nodes have already been segmented. This means the next step is to actually "split" the fiber node. This implies

pulling fiber deeper into the network and installing new nodes in the system. These new nodes may have one or more new SGs associated with them.

Thus, two separate fiber nodes (and the associated feeds for two separate fiber nodes) are required to support the bandwidth for the pool of subscribers in place of where there used to only be one, and therefore there is a cost associated with the node split.

Node splits offer no change in the Service Group bandwidth requirements for broadcast services. If the MSO needed 50 Quadrature Amplitude Modulation (QAM) channels to support broadcast video prior to the node split, then the MSO will still require 50 QAMs to support broadcast video after the node split. The signal is merely further replicated by RF splitting in the headend and sent to the new fiber node.

However, the principle benefit of the SG split is associated with Narrowcast services (Switched Digital Video (SDV), VoD, and DOCSIS). The key benefit of the split is to effectively double the capacity per subscriber for all of these Narrowcast services. SG splitting oftentimes permits MSOs to “free up” some amount of Narrowcast video QAM spectrum whenever they performed the split. However, since oftentimes the driver for the split was to increase the DOCSIS bandwidth per subscriber, the number of DOCSIS channels typically remains the same.

As node splits are performed and fiber is run deeper and closer to the subscribers, the network eventually reaches the point where the fiber node is the last active device in the outside plant. This plant topology is referred to as Fiber to the Last Active (FTTLA) and as Node+0 (N+0), as the Fiber Node has no amplifier or other active component following it. A benefit of nodes splits is that as each split occurs and the number of amplifiers is reduced, the noise contribution from amplifiers is reduced. The noise funneling effect from the multiple subscribers is reduced as the number of subscribers in the Service Group is reduced. However, once the topology reaches N+0, there are diminishing returns for doing further node splits. Reaching an N+0 topology is an important milestone for an MSO, because it is also a prerequisite for migrating to FDX technology. More in-depth discussion on node splits can be found in [CLO1].

2. Distributed Access Architectures

Some MSOs will likely be able to support their video and HSD services using Traditional Headend-based Integrated CCAPs (I-CCAP). However, other MSOs may be planning to perform node splits more rapidly than other MSOs. These MSOs may see a need to support more Service Groups than would be easily supported by an I-CCAP chassis. Adding additional I-CCAPs may cause issues related to the required power and/or rack-space. However, there is an alternate access architecture that helps to solve the problems of MSOs who have issues with the required power and rack-space within their Headends. This technique employs Distributed Access Architectures (DAAs). In addition to addressing the headend power and rack-space issue, DAA architectures are also a necessary component of implementing Full Duplex DOCSIS (discussed later). There are several types of DAAs being proposed for use in the future, and each proposal has its own sets of pros and cons [EMM1]. This paper provides a brief description for three of the more relevant ones.

2.1. Remote PHY (R-PHY)

This approach separates the PHY (Upstream and Downstream) from the headend and places the full PHY layer (including the Forward Error Correction (FEC), symbol generation, modulation, and Digital to Analog Converter (DAC)/Analog to Digital Converter (ADC) processing) into the fiber node. This requires that these functions be removed from the headend CCAPs, CMTSSs, and EQAMs. The DOCSIS

MAC processing remains in the MAC Core within the headend. This approach is slightly disruptive, as it requires many pieces of headend equipment (ex: CCAPs, CMTSs, and EQAMs) to be modified.

The R-PHY approach is an evolution of the Modular Headend Architecture (MHA) approach. But there are also many significant enhancements, such as the need to support Upstream MAC/PHY separation, the need to support new timing interfaces that work over Ethernet, and the need to add DOCSIS 3.1 support within Downstream External PHY Interface (DEPI) and Upstream External PHY Interface (UEPI). However, this approach offers benefits as well. Remote PHY helps with the nonlinear optical noise problem by using digital optics instead of analog, and it also helps with the headend power and rack-space problem. Another benefit of the Remote PHY approach is that it permits MSOs to continue to re-use their headend-based CCAPs as part of the solution. That represents a form of investment protection.




Figure 6 - Remote PHY

2.2. Remote PHY with Virtual Core (vCore)

An extension of the R-PHY approach is to virtualize the MAC Core. Rather than using dedicated hardware in the headend to provide the CCAP Core functionality, the MAC Core functionality is virtualized and run on Commercial Off-The-Shelf (COTS) compute platforms. This architecture will typically take more headend space than dedicated hardware specifically designed for the Core functionality, but this alternative provides other benefits. By decoupling the hardware from the software functionality, each can be updated independently. The virtual platform can be shared with other applications, and can be easily scaled up and down to meet demand. Since the hardware can be scaled back when not needed to meet demand, power savings can result. The architecture is based on Software Defined Networks (SDN) and Network Function Virtualization (NFV) techniques, which provides an infrastructure for rapid feature development.




Figure 7 - Remote PHY with Virtual Core

2.3. Remote MAC/PHY (R-MACPHY)

This approach places the entire upper and lower MAC (Upstream and Downstream) and the entire PHY layer functionality (Upstream and Downstream) into the fiber node. In effect, this places all of the CMTS, Edge QAM, and CCAP functions into the Fiber Node and only requires a switch or router to remain in the Headend. As a result, this approach is not as disruptive. Remote MAC/PHY also helps with the nonlinear optical noise problem, and also provides to the maximum amount of power and rack-space savings within the headend (even more than the Remote PHY approach). By placing both the MAC and the PHY in the same location, it eliminates the DEPI and UEPI protocol overhead. It is also possible that existing headend CCAPs (if appropriately modified) could be used to serve as dense Aggregation Routers (or repurposed PON OLTs) feeding the Remote CCAPs as well.




Figure 8 - Remote MAC/PHY

3. DOCSIS 3.1

DOCSIS 3.1 is a backwards-compatible augmentation to the DOCSIS 3.0 specification that provides better spectral efficiencies (more bps/Hz) and wider channels for both the Downstream and Upstream paths. The specification provides improved spectral efficiencies via many techniques, including the use of:

- Orthogonal Frequency Division Multiplexing (OFDM) modulation,
- Higher modulation orders (4096QAM and higher)
- More efficient Low-Density Parity Check (LDPC) Forward Error Correction
- Bit-loading to custom-fit the modulation orders to the varying SNRs across the spectrum of the HFC plant, and
- Multiple modulation profiles to provide different modulation rate to different CMs depending on their specific noise characteristics

Backwards compatibility is guaranteed by the fact that DOCSIS 3.0 and DOCSIS 3.1 channels can co-exist on the HFC spectrum. In addition, pre-DOCSIS 3.1 CMs will work with DOCSIS 3.1 CMTSs, and pre-DOCSIS 3.1 CMTSs will work with DOCSIS 3.1 CMs.

As a result of its power and flexibility and backwards-compatibility, many MSOs are looking to DOCSIS 3.1 to give them a boost that will extend the life of their HFC plant by (at a minimum) several years. The actual HFC plant life extension that will result from the use of DOCSIS 3.1 depends on many different factors, including the annual subscriber bandwidth growth rates, the number of node splits that are performed, the amount of investment that the MSO is willing to put into their plant to extend its spectral width, and the quality of the HFC plant (i.e. SNRs).

4. FTTx

Migrating to an N+0 architecture means pushing fiber deeper into the outside plant and closer to the subscriber. This is just one flavor of what is referred to as FTTx, where x depends on how deep into the plant the fiber goes. In the case of N+0, this is also called FTTLA (Fiber to the Last Active) or FTTC (Fiber to the Cabinet or Fiber to the Curb). There are other types of FTTx architectures that can benefit subscriber bandwidth growth.

4.1. Fiber to the Tap (FTTT)

Fiber can be taken beyond the traditional node location and could be run all the way to the subscriber tap. From this location, the coax cable run is much shorter, resulting in less attenuation that would enable an Extended Spectrum DOCSIS solution. Extended-spectrum DOCSIS refers to extending the spectrum used in cable networks above and beyond of what DOCSIS 3.1 can support [CLO2]. This can be effective in network topologies where no amplifiers or diplexers are present. The coaxial cables can support very high frequencies such as 25 GHz for RG-6 drop cables. Although attenuation will cause a reduction in the modulation orders that can be used, the extremely wide spectrum will allow much higher total bandwidths.

4.2. Fiber to the Home (FTTH), Fiber to the Premise (FTTP)

Running fiber all the way into the premises is the next logical (and final) step in running fiber deep into the network. A Passive Optical Networks (PON) is a technology that provides a direct optical link

between the headend and the subscriber home. The device in the headend is called an OLT, and the device in the home is called an ONU or an ONT. Many ONUs (or ONTs) can share a single FTTH optical feed from the OLT in the headend, so the bandwidth capacity provided by a PON is always shared by all of the ONUs (or ONTs) connected to the PON feed.

PON technologies today include bandwidth capacities such as 1 Gbps, 2.5 Gbps, and 10 Gbps. Ultimately, 40+ Gbps bandwidths will also likely be provided. For MSOs, this is an overlay technology to the DOCSIS HFC delivery system, since it does not offer any form of backwards-compatibility to DOCSIS. PON will likely be used in Business Services and MDU environments first, but it will also find great utility in servicing elite Residential subscribers as well (once Residential subscriber bandwidth demands exceed those that can easily be provided by traditional DOCSIS systems that haven't been upgraded).

PON may find a few competitors in the FTTH space. One FTTH competitor to PON is RF over Glass (RFOG). RFOG technology permits MSOs to transmit their standards RF signals (e.g. DOCSIS, MPEG-TS Video, and Analog) all the way to the subscriber homes over fiber. It requires a special ONU to be placed within each home, and the ONU is responsible for performing an optical-to-electronic conversion function (which is quite similar to the function performed by a typical fiber node). RFOG offers several benefits to MSOs. It permits MSOs to begin transitioning their HFC plant into a Fiber-to-the-Home (FTTH) plant (which is likely to be the plant of the future) while maintaining backwards compatibility with their huge existing CPE investment. RFOG eliminates the coaxial portion of the HFC plant, which can lead to improved SNRs and higher modulation orders. RFOG can extend their DOCSIS 3.1 transmission system to spectral widths that exceed the 1.2-1.7 GHz spectral limits of typical coaxial distribution systems within the HFC plant. Initial RFOG systems suffered from a type of noise called Optical Beat Interference (OBI) that is sometimes generated when multiple ONUs transmitted at the same time. However, there are now forms of OBI-free RFOG systems that eliminate this type of interference.

5. Full Duplex DOCSIS

Full Duplex DOCSIS is an enhancement to the DOCSIS 3.1 specification to enable greatly increased upstream bandwidths. The target is to be able to provide 10 Gbps downstream bandwidth and 5 Gbps upstream bandwidth within a Service Group. In order to expand the upstream bandwidth while having minimal impact on the downstream bandwidth, FDX allows certain portions of the spectrum to be used for upstream and downstream transmissions simultaneously. The spectrum from 108 MHz to 684 MHz has been designated for these bi-directional transmissions. The updated spectrum usage is depicted in Figure 9.




Figure 9 - FDX Spectrum Usage

It should be noted that the simultaneous transmission and reception of packets is from the fiber node point of view. Each individual CM will still be operating in a frequency division multiplexing (FDD) mode. CMs will be grouped together into Transmission Groups (TG). Each TG will use some channels in the FDX band as upstream channels and the other channels as downstream channels. However, one TG may be using one part of the spectrum as an upstream channel while another TG may use that same part of the spectrum as a downstream channel. In addition, usage of the spectrum for upstream and downstream within a TG can be changed over time. From a CM point of view, the FDX band operates as a Dynamic FDD system, as illustrated in Figure 10.




Figure 10 - Dynamic FDD Operation from cable modem perspective

Alternative Network Migration Paths

As described in the previous section, there are many tools that are available for the MSOs to choose from to support their network migration plans. These tools include node splitting & segmentation, DAA vs centralized architectures, DOCSIS 3.1, HFC vs. FTTH, RFoG vs. PON for FTTH, selective subscriber migration, etc. The optimal choice depends on the network parameters, offered demand and statistical distribution of subscribers among services, and MSO's restrictions (e.g., logistics/operational/resources constraints, current infrastructure, budget, etc.). Some additional factors to consider include current Service Group size and target final Service Group size, when will the transition from QAM video to IP video occur, and when will symmetric services be required. Therefore, a solution that perfectly works for one MSO may not be optimal for another MSO.

As previously discussed, the rate of downstream bandwidth growth appears to be slowing to 40% per year. In the absence of any other network changes, this implies that nodes need to be split or segmented approximately every 2.1 years in order to keep up with bandwidth demands. As the analysis in [CLO1] shows, the effectiveness of node splits is reduced each time a node is split into a smaller Service Group. This is due to the peak bandwidth of a single subscriber starts to become a more dominant effect on the Quality of Experience over the average bandwidth of all the subscribers in the Service Group. A point of diminishing returns is probably reached when the Service Group reaches around 50 subscribers. Depending on the current average Service Group size, nodes splits can provide an effective migration strategy for many years to come. Table 1 summarizes how many year it takes to reach an average Service Group size of 50.

Table 1 - Estimate HFC Plant Life Using Node Splits

Current Average Service Group Size	Years
100	2.1
200	4.1
300	5.3
400	6.2
500	6.8
600	7.4
700	7.8
800	8.2
900	8.6
1000	8.9

The above table assumes no other changes are made to the network. An additional migration strategy involves migrating how spectrum is allocated. Although High Speed Data (HSD) is the fastest growing service within an MSO's HFC spectrum, MSO-managed video services still consume the largest percentage of the spectrum today. To accommodate the growing HSD bandwidth, MSOs may look to various technology paths that offer to squeeze the bandwidth of MSO-managed video into a smaller portion of the HFC spectrum. The future will likely see different MSOs using different mixes of SD broadcast video, HD broadcast digital video, SDV, VoD, IP video, and analog video.

Over time, the analog video spectrum will be heavily reclaimed (many MSOs have already entirely reclaimed it). DTAs offer a good, low-cost technique for accomplishing that goal. Future Media Gateways with low-cost IP-STBs may also provide similar low-cost alternatives. SDV is another technique that can help to reclaim spectrum from the broadcast digital video tier, whereby video streams are only transmitted over a Service Group if a subscriber is viewing that stream. As SG sizes become smaller, SDV becomes more effective and can reclaim more legacy video spectrum.

In addition to a transition away from analog video towards digital video, and in addition to a transition away from broadcast video towards SDV, many MSOs are also looking to a transition away from MPEG-TS based QAM digital video delivery to IP based video delivery over DOCSIS. There are several reasons for this trend. Several of these reasons can be grouped together saying that DOCSIS provides better spectral efficiency over QAM digital video delivery [CLO1]. In addition, Over-the-Top (OTT) video delivery is becoming popular with subscribers and is delivered over IP. Over time, MSOs may migrate away from their managed QAM digital video delivery to their own OTT video delivery. Figure 11 depicts how downstream spectrum may migrate over time, increasing the amount of spectrum allocate to DOCSIS, which increase the amount of available bandwidth.

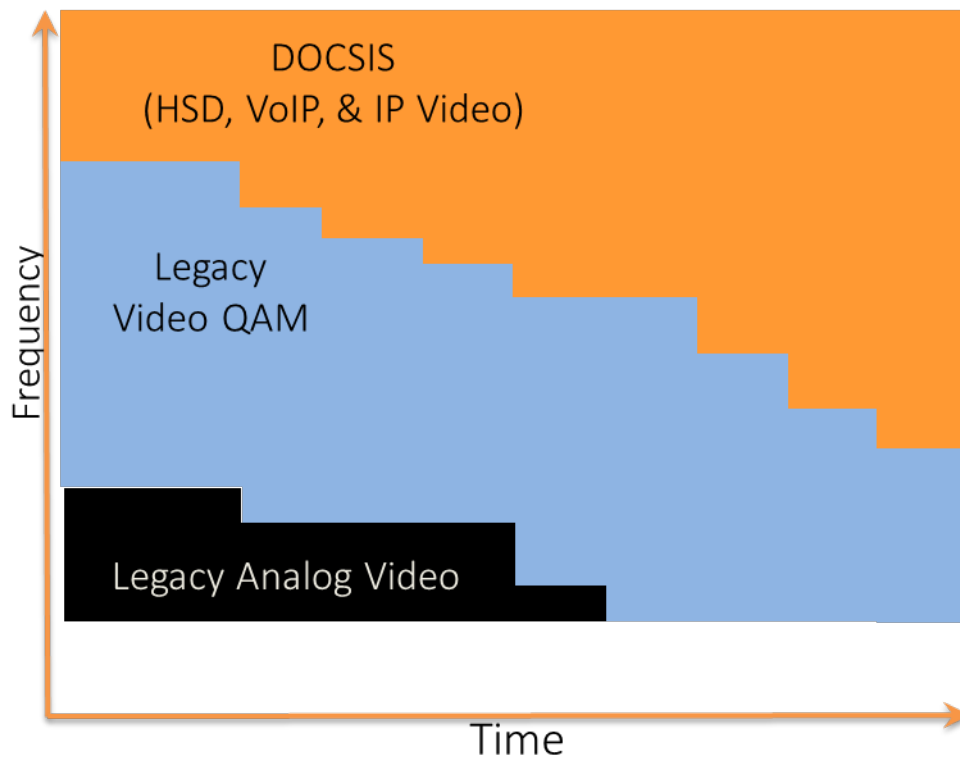


Figure 11- Downstream Spectrum Migration

MSOs also have options on how to migrate their upstream to meet subscriber demands. The migration is dependent on transitioning CPEs to DOCSIS 3.1 CMs (and eventually FDX CMs). The migration also depends on how quickly MSOs are required to dramatically increase upstream to provide symmetric services, typically to meet competitive pressures (e.g. from Google Fiber). Some MSOs may be able to meet the short term upstream bandwidth demands by migrating to an 85 MHz mid-split. Other MSOs may need to migrate quickly to providing large upstream bandwidths, and migrate to a 204 MHz high-split instead.

The starting point of a split may cause different paths to the end goal of using the FDX band for upstream. When FDX becomes available, how the usage is shared between legacy D3.1 CMs and the new FDX CMs may depend on the diplexer in the legacy CMs. If the legacy D3.1 CMs are on an 85 MHz plant, they would not be able to participate in the whole upstream of the FDX band of 108 to 684 MHz (although with a software upgrade they could share the downstream FDX spectrum with the FDX CMs). The FDX CMs would be able to use the FDX band for upstream or downstream transmissions. On the other hand, if the legacy D3.1 CMs are currently configured for a 204 MHz split, it will be able to share the spectrum from 108 to 204 MHz in the upstream direction with the FDX CMs, while the FDX CMs will be able to additionally use the spectrum from 204 to 684 for upstream bandwidth.

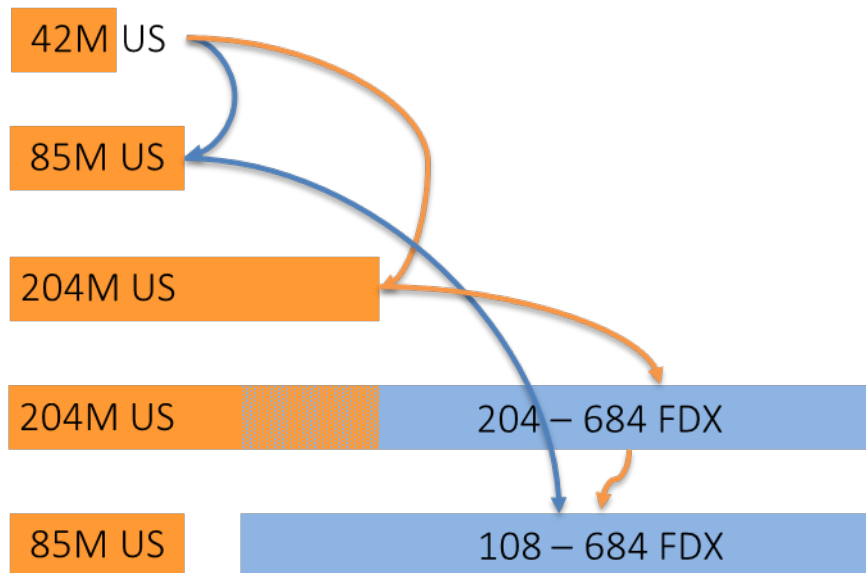


Figure 12 - Upstream Spectrum Migration

Although the FDX band is 576 MHz wide, MSOs may not desire to use all this spectrum for FDX initially, or may be limited by how much spectrum can be freed from other services. The specification for FDX allows FDX CMs to use only a portion of the FDX band for FDX channels. However, the portion of the FDX band that is not being used for FDX channels can only be filled with legacy video QAM channels. This is because FDX CMs can only transmit and receive FDX channels in the portion of the spectrum reserved for the FDX band. The possible FDX band migration steps are shown in Figure 13. Depending on other factors such as spectrum availability, upstream bandwidth demand, and FDX CMs penetrations, MSOs may choose how slow or fast to progress through these migration steps.

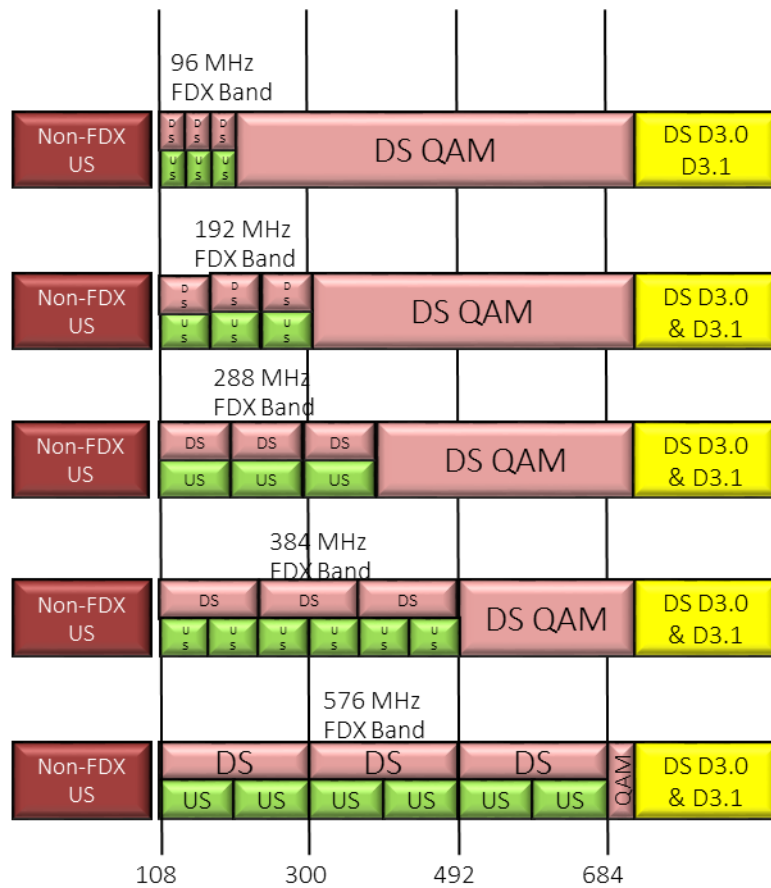


Figure 13 - FDX Spectrum Migration

Conclusion

With the new demands driving bandwidth growth such as the competitive pressure to provide symmetrical upstream and downstream bandwidth services, some see FDX DOCSIS as the answer to solving all problems. However, not all MSOs are the same, and one technology is not going to solve every problem. MSOs will require a whole toolkit of technologies and procedures to address their network migration needs. Those tools include node splitting & segmentation, DAA vs centralized architectures, DOCSIS 3.1, HFC vs. FTTH, RFoG vs. PON for FTTH, selective subscriber migration, etc. Utilizing those tools creates network migrations such as the one shown in Figure 14.

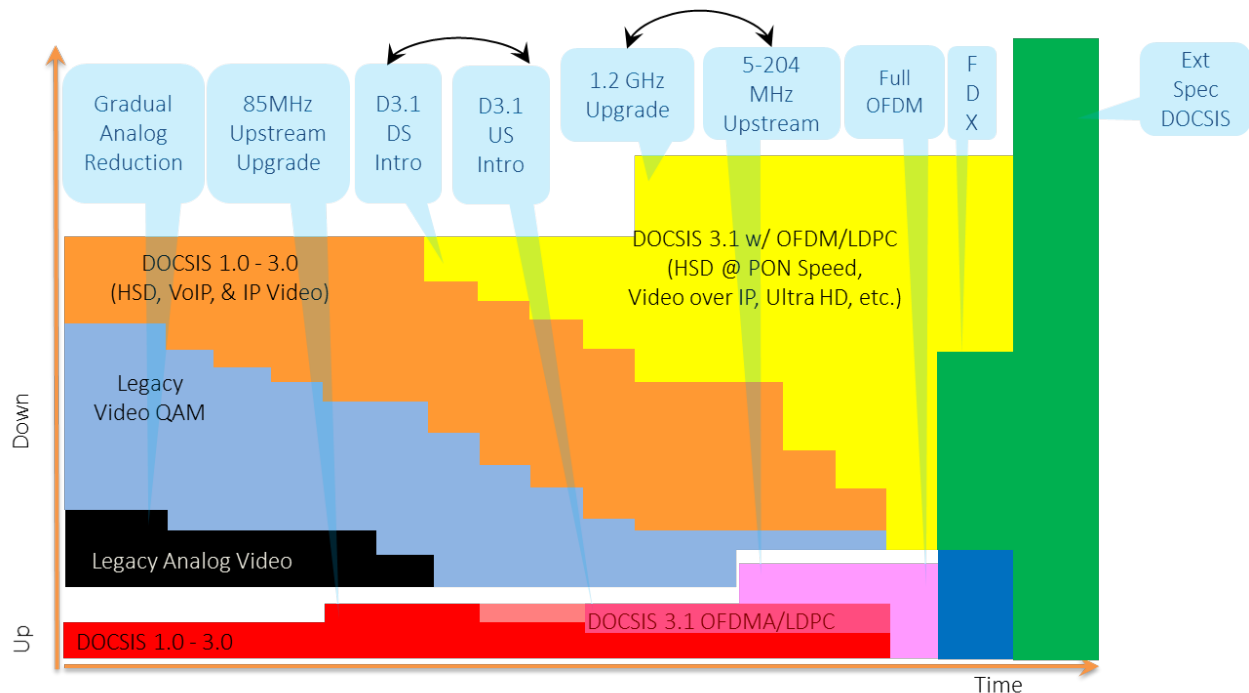


Figure 14 - Network Migration Options to Meet Bandwidth Demands

Each MSO has a unique set of circumstances, and they must apply the set of tools in a unique combination to meet their specific goals and objectives. And the MSO may also have to apply different combinations of these tools at different times for different sites.

Abbreviations

ADC	Analog to Digital Converter
Bps	bits per second
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
COTS	Commercial Off-The-Shelf
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification version 3.1
DAA	Distributed Access Architecture
DAC	Digital to Analog Converter
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DSL	Digital Subscriber Line
DTA	Digital Television Adapter
EQAM	Edge Quadrature Amplitude Modulator
FDD	Frequency Division Multiplexing
FDX	Full Duplex DOCSIS
FDX CM	Full Duplex Cable Modem
FEC	Forward Error Correction
FTTC	Fiber to the Cabinet or Curb
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTT	Fiber to the Tap
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
HD	High Definition
HFC	Hybrid Fiber Coax
HSD	High Speed Data
Hz	hertz
I-CCAP	Integrated Converged Cable Access Platform
ISBE	International Society of Broadband Experts
LDPC	Low-Density Parity Check
MAC	Media Access Control interface
MACPHY	DCA instantiation that places both MAC & PHY in the Node
MDU	Multiple Dwelling Unit
MHA	Modular Headend Architecture
MHz	Megahertz
MSO	Multiple System Operator
N+0	Node+0 actives

NFV	Network Function Virtualization
OBI	Optical Beat Interference
OFDM	Orthogonal Frequency Division Multiplexing
OLT	Optical Line Termination
ONU	Optical Network Unit
OTT	Over-The-Top
PHY	Physical interface
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
RF	Radio frequency
R-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RFoG	RF over Glass
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networks
SDV	Switched Digital Video
SSM	Selective Subscriber Migration
TG	Transmission Group
UEPI	Upstream External PHY Interface
US	Upstream
vCore	Virtual Core
VoD	Video on Demand

Bibliography & References

[CLO1] T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, and S. Chari, “Predictions on the Evolution of Access Networks to the Year 2030 & Beyond,” The NCTA Cable Show Spring Technical Forum 2014

[CLO2] Tom Cloonan et. al., “Lessons from Telco & Wireless Providers: Extending the Life of the HFC Plant with New Technologies,” Spring Technical Forum conference INTX 2015

[EMM1] M. Emmendorfer, T. Cloonan, and J. Ulm, “A Side-by-side Comparison of Centralized vs. Distributed Access Architectures,” The NCTA Cable Show Spring Technical Forum 2014

IG Discovery for 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
tonliu@cisco.com

Introduction

In legacy DOCSIS, data can only be transmitted in one direction across any part of the spectrum. Compared to the passive optical networks (PONs), a cable access network is severely limited in the maximum symmetrical data speed due to the upstream RF spectrum scarcity. Since bringing fiber to the home is extremely expensive, cable operators have searched for an alternative to deliver the multi-gigabit services promised. This need together with recent trends in the cable industry (i.e. the deployment with DOCSIS 3.1 Orthogonal Frequency Division Multiplexing (OFDM); the deep fiber migration; and the remote PHY network architecture) has resulted in the rapid development and standardization of the full duplex (FDX) DOCSIS technology. With FDX DOCSIS, the RF spectrum can be used simultaneously in both the upstream (US) and downstream (DS) directions, allowing up to 5 Gbps US service and 10 Gbps DS service over the cable access network.

In FDX communications, a system supports simultaneous bi-directional transmissions across the same spectrum. Interferences between the bi-directional transmissions therefore must be mitigated for the intended signals to be properly received. DOCSIS is a point to multi-point system, where multiple cable modems (CMs) are connected to the same Cable Modem Termination System (CMTS) port via a coax distribution line. 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. Since the source of the interference is unknown to the receiving CM, PHY layer echo cancellation cannot be used. FDX DOCSIS address this issue by grouping CMs that interfere with each other into an Interference Group (IG). CMs in the same IG must transmit or receive along the same direction at any given frequency and time. CMs from different IGs have enough RF isolations to allow simultaneous US and DS transmissions at the same frequency.

In this paper, we will discuss IG discovery, a new process introduced in FDX DOCSIS to determine the IGs based on the CM to CM interference measurement obtained via sounding. We will start by introducing the basic IG concept and the operational principles to conduct sounding. We will examine the system overhead in terms of the spectrum cost and the time to converge for sounding among a given number of CMs at the desired frequency granularity. We will then propose a set of optimization techniques to improve sounding efficiency. We further extend the solution space by incorporating an iterative IG Discovery model to allow the system to automatically adapt to the changing network environment for optimized system performance.

IG Discovery Overview

1. Interference Groups

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 in a full duplex mode. This occurs when the co-channel interference (CCI) levels at the receiving CMs are above a design threshold when a CM is transmitting simultaneously over the same FDX spectrum.

FDX DOCSIS uses a sounding procedure to measure the CM to CM CCI. During Sounding, the CMTS selects one or more FDX capable CMs as test CMs to transmit test signals on designated subcarriers, while directing other FDX capable 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.

The measured CCI, in the form of the RxMERs collected from the measurer CMs, can then be used to sort CMs into IGs. Quantitatively, given a set of CMs, cm_1, cm_2, \dots, cm_N in a service group, cm_i 's IG group, $IG(cm_i)$, can be determined, such that,

$$\text{for any transmitting } cm_j \in IG(cm_i), RxMER_{ji} < \overline{MER} \quad (1)$$

or,

$$\text{for any transmitting } cm_j \notin IG(cm_i) \quad RxMER_{ji} > \overline{MER}; \quad (2)$$

Where, $RxMER_{ji}$ is the RxMER obtained at cm_i when cm_j is transmitting test signals. \overline{MER} is the threshold designed for $IG(cm_i)$, for its member CMs to properly demodulate a target modulation scheme.

Since the path loss of the interfering signal is reciprocal in a passive coax plant, symmetrical CCI is expected between a pair of CMs, therefore,

$$\text{if } cm_j \in IG(cm_i), \text{ then } cm_i \in IG(cm_j); \quad (3)$$

However, as the RxMERs are also impacted by the noise sourced internal to a CM, the RxMER level may not be the same. Sounding is thus required at both cm_i and cm_j to accurately detect the interference.

Figure 1 shows an IG Discovery example using the RxMER measurement data listed in Table 1. The shaded cells mark out the three IGs after applying a 35dB MER (or 10 bits/subcarrier) threshold, namely IG1 for CMs under Tap1, IG2 for CMs under Tap2, and IG3 for CMs under Tap3, Tap4 and Tap5.

From the example, we can observe the following:

1. Low MERs for CMs under the same tap; for example, the MER is 6dB for CMs under Tap1, as the RF path loss between the CMs under the same tap is much less compared to the inter-tap case.
2. Low MERs for CMs under the taps close to the end of distribution line; for example, CMs under Tap 3 through Tap5 all have MER below 35dB, due to the poor coupling loss of the lower-value taps.
3. Symmetrical CM-to-CM interference indicating reciprocal path loss of the passive plant.

R E C E I V E	MER (dB)	TRANSMIT				
		Tap1	Tap2	Tap3	Tap4	Tap5
	Tap1	6	39.9	39.9	39.9	39.9
	Tap2	39.9	9.8	37.5	37.5	37.5
	Tap3	39.9	37.5	13.2	34.5	34.5
	Tap4	39.9	37.5	34.5	15.8	31
	Tap5	39.9	37.5	34.5	31	18.2

R E C E I V E	Mod Order	TRANSMIT				
		Tap1	Tap2	Tap3	Tap4	Tap5
	Tap1	-	11	11	11	11
	Tap2	11	-	10	10	10
	Tap3	11	10	-	9	9
	Tap4	11	10	9	-	8
	Tap5	11	10	9	8	-

Table 1 - CM-to-CM Interference and IG Formation

DS output power: 39 dBmV/6 MHz
US Input power: 8 dBmV/6.4 MHz

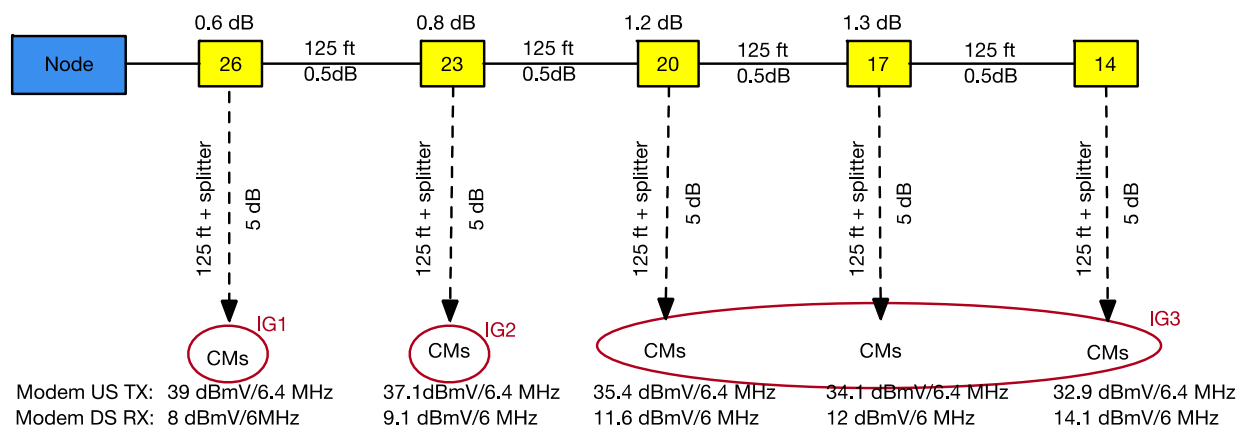


Figure 1 - CM Interference Groups over a passive coax distribution line

2. Sounding Techniques

There are two sounding methods proposed in FDX DOCSIS [3][4].

1. Sounding with OFDMA Upstream Data Profile (OUDP) test bursts
2. Sounding with continuous wave (CW) test signals

The OUDP method is intended for the deployment scenario where the legacy high-split DOCSIS 3.1 CMs, after necessary software upgrade, can share the US spectrum between 108 to 204 MHz with the FDX CMs. Since the DOCSIS 3.1 CMs cannot generate a multiplicity of CW tones as required in the CW sounding method, the DOCSIS 3.1 OUDP test bursts must be used instead as the test signals. When the OUDP test bursts are being transmitted by a test CM, other CMs that are capable to receive in this frequency band measure the RxMERs in the time and frequency encompassed by the continuous OUDP bursts. The OUDP test burst is intended to cover all DS subcarrier frequency locations by taking advantage of a faster RxMER measurement scheme to be implemented on the new FDX CMs.

The CW method is intended for the deployment scenario where the DOCSIS 3.1 CMs, after necessary software upgrade, can share the DS spectrum with FDX CMs. For example, a low-split or mid-split DOCSIS 3.1 CM can share the DS spectrum between 108 to 684 MHz, and a high-split DOCSIS 3.1 CM can share the DS spectrum between 258 to 684MHz. During CW sounding, one or multiple FDX test CMs send CW test signals at selected DS subcarrier frequency locations, while the rest of CMs, including both legacy D3.1 CMs and FDX CMs measure the MER using the DOCSIS 3.1 RxMER measurement method.

3. Spectrum Overhead

A sounding test opportunity requires spectrum resource in time and frequency for both the US and DS directions. As shown in Figure 2, in the US direction, a test signal transmission opportunity is required for

a test CM to send the test signals. In the DS direction, a test signal interference region is required to carry zero-bit-loaded symbols, to avoid any packet caused by the interference from the test signals.

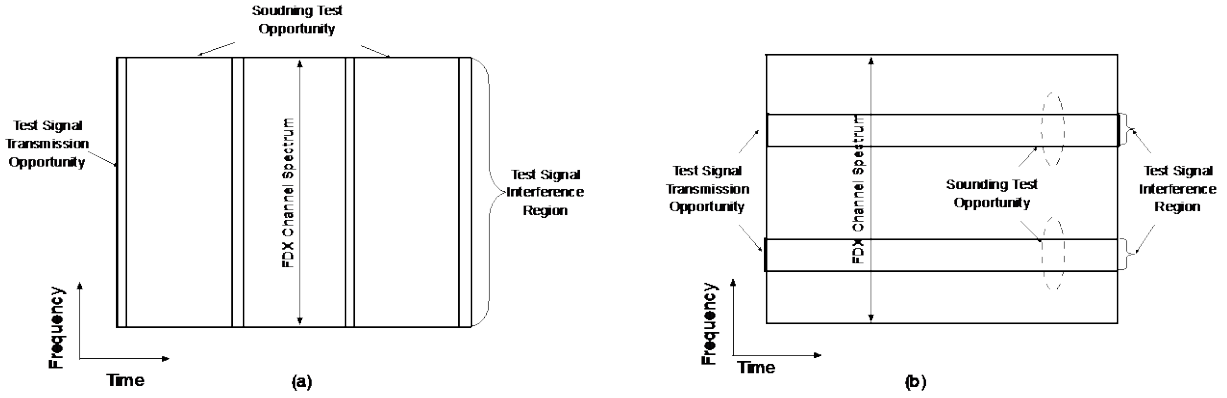


Figure 2 - Sounding Test Opportunities (a) OUDP Test Opportunities (b) CW Test Opportunities

For the OUDP sounding, a sounding test opportunity covers the entire FDX channel width in frequency and lasts about 20 to 60 milliseconds in time [4]. Thus, no spectrum can be used for traffic when the OUDP sounding burst is present on the FDX channel under test.

For the CW sounding, a sounding test opportunity includes a single CW subcarrier and a few guard subcarriers on both sides, to prevent inter-symbol interference at adjacent data subcarriers. Comparing to the OUDP sounding, a CW test opportunity occupies much narrower spectrum however lasts longer in time. It typically takes around 200 to 300 milliseconds for DOCSIS 3.1 RxMER measurement scheme to converge.

With the CW sounding, the CMTS has the option to limit the number of sounding test opportunities, so traffic can be sent using the data subcarriers outside the CW interference regions, particularly, the DS traffic to the measurer CMs, and the US traffic from a test CM if the test CM's IGs have been identified through previous sounding.

The spectrum overhead S_{avg} spent on sounding can thus be expressed as the percentage of the sounding dwell time multiplied by the percentage of the number of subcarriers budgeted for sounding,

$$S_{avg} = (Sb_{sounding}/Sb_{total}) * (T_{sounding_cycle}/T_{sounding_interval}) \quad (4)$$

where,

$Sb_{sounding}$: total number of subcarriers in all concurrent sounding test opportunities

Sb_{total} : total number of subcarriers on a given FDX channel under test,

for OUDP sounding, $Sb_{sounding} = Sb_{total}$;

for CW sounding, $Sb_{sounding} < Sb_{total}$;

$T_{sounding_cycle}$: duration of a sounding cycle to sound all intended Test CMs on a given FDX channel

$T_{sounding_interval}$: the average time interval between subsequent sounding cycles.

4. Sounding Cycle

As mentioned in the previous section, a sounding cycle includes all the necessary operational steps to identify the interference relationships among all CMs that may transmit and/or receive on a given FDX channel. As shown in Figure 3, a sounding cycle includes preparation, interference test and recovery three phases:

- Preparation Phase

To prepare for sounding, the CMTS has to ensure the FDX channel operates in the DS direction from the measurer CMs' point of view. If the FDX channel has been operating in the US direction in regarding to the measurer CMs, CMTS must switch it to the DS direction and wait for the measurer CMs to acquire the DS channel prior to sounding starts.

- Interference Test Phase

The interference test phase consists of one or more test windows. Each test window marks the time span of one or more parallel test opportunities as shown in Figure 3. In case of OUDP sounding, a single test opportunity covers the entire FDX channel width, hence the number of test windows required is equivalent to the number of test CMs. In case of CW sounding, a test window may contain multiple concurrent test opportunities arranged at difference frequency locations. These test opportunities can be assigned to one test CM or a group of test CMs to sound in parallel. The number of test windows required therefore equals to the number of parallel test groups that can be arranged among the test CMs. Parallel sounding is an optimization technique to shorten the sounding cycle.

- Recovery Phase

After the interference test is done, a recovery phase is required for the CMTS and the CM to resume regular operations. The recovery phase may include channel direction change to recover the traffic throughput prior to sounding.

The sounding cycle duration can be simply expressed as,

$$T_{sounding_cycle} = T_{prepare} + N * T_{test_window} + T_{resume} \quad (5)$$

Where,

$T_{prepare}$: sounding preparation time

N : the number of sounding test windows

T_{test_window} : duration of a sounding test window

T_{resume} : recovery time to resume FDX traffic operation post sounding.

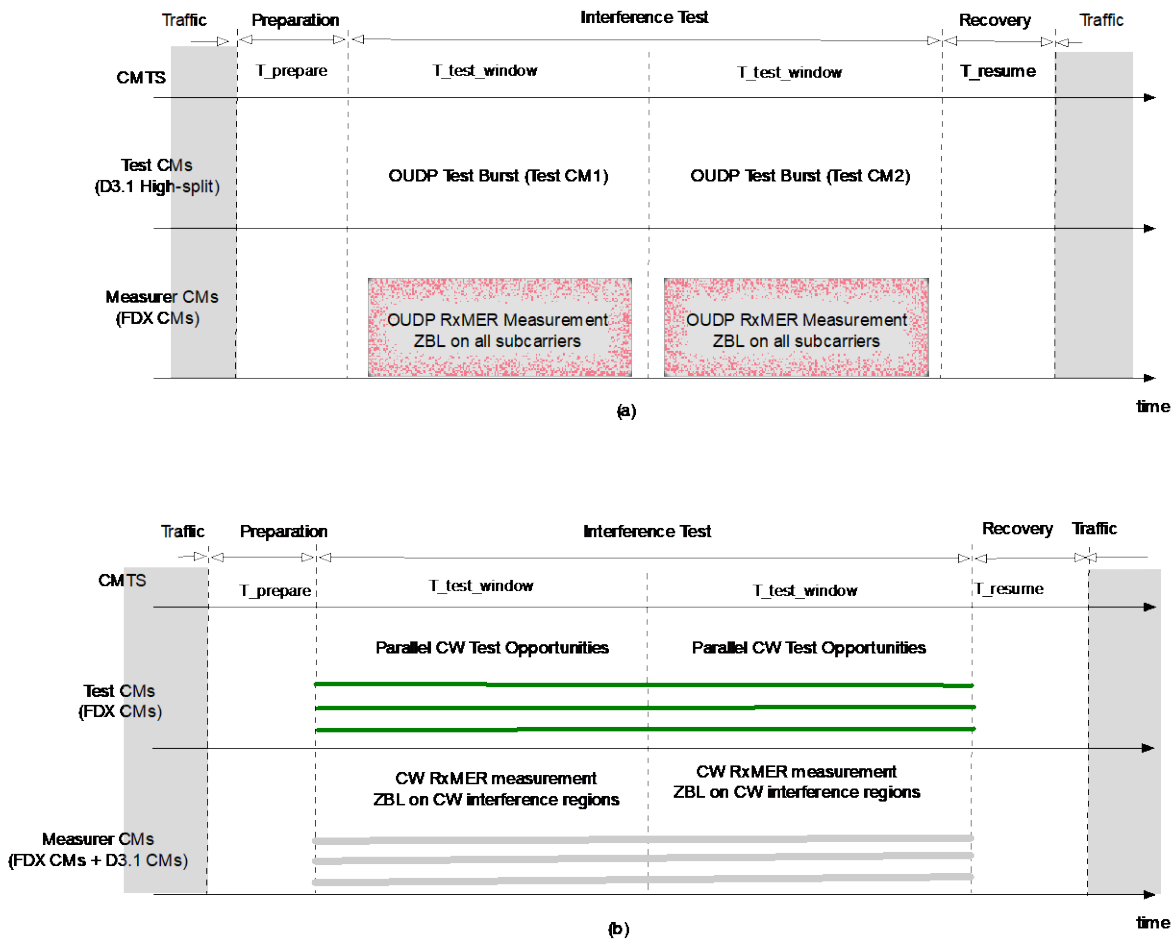


Figure 3 - Sounding Cycle (a) OUDP Sounding (b) CWT Sounding

The sounding cycle duration is a performance benchmark from the FDX operation point of view. It quantifies the FDX bandwidth access time when a new FDX CM is coming online and the traffic interruption time when there are active FDX CMs already operating on the given FDX channel prior to start of sounding.

For CW sounding, the sounding cycle duration is inversely proportional to the number sounding subcarriers at a given spectrum overhead level, as shown in equation (4). It is also impacted by the number of concurrent CW test signals that a CM can send. Figure 4 shows the sounding cycle duration in relation with the sounding subcarrier percentage and the number of CW test signals per CM.

From the chart, we can observe that at given sounding frequency granularity:

- The CW sounding cycle duration decreases as the number of sounding subcarriers increases.
- The sounding cycle duration remains the same if the number of sounding subcarriers allocated results in the same number of test windows.
- The number of CWs a CM needs to generate is bounded by the available number of sounding subcarriers. For example, there is no time advantage for a test CM to generate more than 255 CW tones if only 5% of the subcarriers can be used for sounding at any given time.

- At given frequency granularity, spectrum budget and number of CMs to sound, an optimum number of concurrent CW tests per CM exists that can result in the shortest sounding cycle duration.




Figure 4 - CW sounding duration at different sounding subcarrier percentage and number of CWs per CM

Number of test CMs	60
Number of subcarrier a CM need to sound	3976
Number of subcarriers in a CW interference region	7
Time to prepare for CW test	200 ms
CW Test window duration	500ms
Time to resume FDX operation	100ms

Table 2 - Assumed parameters for the CW sounding duration calculation

IG Discovery Optimizations

This section looks at a set of optimization techniques for IG Discovery based on the following realizations:

- Since a CM cannot be both a transmitting CM and receiving CM on a given FDX channel at the same time, sounding can be decomposed into two directional tests, namely, a transmitting test and

a receiving test that can be conducted independently. This leads to the Partial Sounding technique.

2. The frequency granularity required for sounding is bound by the MER margin acceptable to a modulation order and the corresponding correlation bandwidth in plant's frequency response. This leads to the MER sub-sampling technique.
3. IG discovery accuracy is relative to the DS spectrum efficiency. Errors in interference measurement and estimations can be compensated with lower modulation orders. IG Discovery may never complete as the interference environment keeps changing. This realization leads to the iterative IG Discovery technique.

The following subsections describe each technique in detail.

5. Full Mesh Sounding vs. Partial Sounding

Full mesh sounding is intended to proactively test all pairing permutations between the transmitting CMs and the receiving CMs. To perform full mesh sounding, the FDX channel under test must be changed to the DS direction for all potential measurer CMs. Consequently, full mesh sounding lasts longer in time and causes longer traffic interruptions. Full mesh sounding may not be desirable if the traffic condition does not permit the necessary time and spectrum required.

Partial sounding attempts to minimize the traffic impact by opportunistically pairing the test CM and Measurer CMs based on the channel direction in use. Partial sounding can be either a transmitting test or a receiving test as shown in Figure 5. The transmitting test allows the CMTS to evaluate if a new CM can transmit upstream on a FDX channel when a specific set of CMs are receiving over the same spectrum. The receiving test allows the CMTS to evaluate if a new CM can receive on a FDX channel when a specific set of CMs are transmitting upstream over the same spectrum. Based on the partial sounding, the CMTS can conditionally enable a CM's FDX service if the operation conditions match the tested scenarios




Figure 5 - Partial Sounding, (a) Transmitting Test; (b) Receiving Test

Full mesh sounding and partial sounding can be combined to provide an optimum system solution, for example applying full mesh sounding upon boot up to acquire the interference relationship base line, and applying partial sounding repetitively when a new interference condition is present.

6. Sequential Sounding vs. Parallel Sounding

Parallel sounding is used to reduce the sounding cycle duration. Parallel sounding is possible when the number of sounding test opportunities is greater than the number of test signals a CM needs to generate at a time.

The following is an example to exam the timing advantages of the parallel sounding. Figure 6 shows a service group with N (64 in this example) FDX CMs that are capable to transmit and receive on a FDX channel. The time to conduct full mesh sounding requires N CW sounding test cycles, if sounding is performed sequentially with only one CM transmitting in each test window.

Figure 7 shows a parallel sounding algorithm that sounds 8 CMs at a time. First horizontally by arranging each column of 8 CMs transmitting on different subcarrier locations while the rest of CMs in the service group measuring MER on all DS subcarriers. After this step, the only unknown interference is between different rows, so the second step is to sound vertically by arranging each row of CMs to send test signals in parallel while the rest of the CMs measure. The total number of CW test cycles with this approach is 16. Assuming each CW test cycle takes 800ms, parallel sounding in this example only takes 12.8 seconds, while the sequential sounding method would take 51.2 seconds.

Compared to sequential sounding, parallel sounding takes less time but a cost of frequency granularity. Parallel sounding is suitable to identify interferences at restricted frequency locations or form coarsely grained IGs to speed up FDX service access.




Figure 6 - Sequential sounding example




Figure 7 - Parallel sounding example

7. Complete Sampling vs. Sub-sampling

Complete sampling refers to the type of sounding in which sounding is attempted on all subcarriers of a given FDX channel. In the case of CW sounding, complete sampling can only be achieved with incremental subsampling, which may take an extended period of time to complete.

The complete sampling is generally not necessary for FDX operation. The frequency granularity required for sounding is bound by the MER margin acceptable to a given modulation order and the corresponding correlation bandwidth at a given frequency. Results from the subsampling can be directly used for IG discovery. The CCI level on the unsounded subcarriers can be interpolated with a maximum likelihood estimation with certain error margins.

Figure 8 shows a subsampling example with the measured MERs scattered across a few subcarriers. Figure 9 shows the MER interpolations in between the sparsely spaced measurement samples. For each estimated MER value, a variation range is incorporated to bound the worst-case estimations. As time progresses and more subcarriers are sounded, the cumulative subsampling approaches the full sampling with less estimation errors as shown in Figure 10.

Subsampling allows the CMTS to quickly enable the FDX operations with coarsely grained initial IGs, and incrementally refine the IG formations with continuous subsampling.




Figure 8 - Sub-sampling at selected subcarrier locations




Figure 9 - Subsampling with interpolated MER (InMER) estimations




Figure 10 - Cumulative subsampling over time

Iterative IG Discovery

The iterative aspect of IG Discovery is important. As the interference environment changes, either triggered by a new CM coming online, channel allocation change or temperature fluctuations, the system must be able to adapt, using previous computations together with any new sounding data to produce reliable IG decisions.

The iterative IG Discovery process can be modeled as a multi-stage feedback loop that constantly refines the IG decisions based on the new measurement data and the feedback for positive and negative outcomes. As shown in Figure 11, the iterative IG Discovery process includes the following four steps:

- Sounding

This is for measuring the interference between the specific transmitting and receiving CM pairs at given frequency locations. The measurement data obtained will be used for IG formation.

- IG Formation

The new measurement provided by sounding, together with previous computation results, is used to form IGs to enable FDX operation with acceptable error margins.

- FDX Operation

The FDX operation is constantly monitored. Events and statistics, such as CM population, traffic condition and signal quality are collected for IG evaluation.

- IG Evaluation

IG decisions are re-evaluated based on the operation events and statistics. The evaluation results in a new set of transmitting and receiving CM pairs and specific frequencies targeted for the next round of sounding.

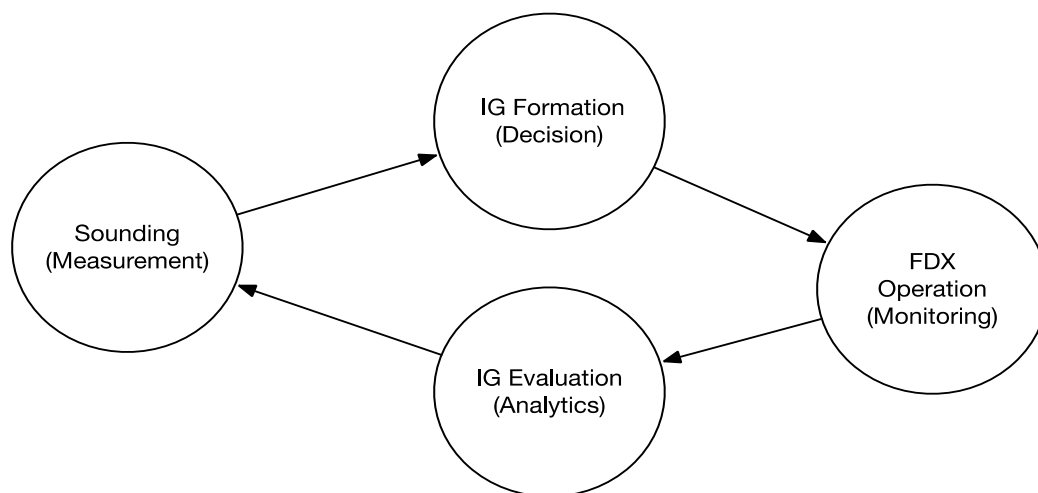


Figure 11 - Iterative IG discovery process

Conclusion

The operational requirements for IG Discovery results in conflicting design considerations, in terms of spectrum budgeting, time to convergence and the interference detection accuracies. In search for a balanced, optimization solution, a system approach is used to identify the key performance impacting elements and their tradeoff relations. Based on this, a set of optimization techniques are described including:

- partial sounding
- parallel sounding
- interference subsampling with interpolations

The solution space is further extended by incorporating an iterative process that follows a measurement – decision – monitoring – analysis feedback loop, to allow the IG Discovery to be constantly refined and adaptive to the changing interference environment.

Abbreviations

CM	Cable Modem
CMTS	Cable Modem Termination System
CW	Continuous waveform
DS	Downstream
FDX	Full Duplex
IG	Interference Group
HFC	hybrid fiber-coax
MER	Modulation Error Ratio
Hz	hertz
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing with Multiple Access
OU DP	OFDMA Upstream Data Profile
PON	Passive Optical Network
US	Upstream

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] FDX MAC EC: MULPIv3.1-x-17.1764-1
- [4] CM-SP-PHYv3.1 Annex F

Can a Fixed Wireless Last 100m Connection Really Compete with a Wired Connection and Will 5G Really Enable this Opportunity?

New Wireless Spectrum Opportunity: How will this Factor into the MSO Access Architecture with Fixed Wireless Access as a Delivery Option?

A Technical Paper prepared for SCTE•ISBE by

J.R. Flesch

Director, Advanced Technology, CPE
ARRIS
3871 Lakefield Drive
Suwanee, GA 30024
Jr.flesch@arris.com
678-473-8340

Belal Hamzeh, Ph.D.

VP, Wireless Technologies
CableLabs
b.hamzeh@CableLabs.com

Bryan Pavlich

Staff Software Engineer
ARRIS
bryan.pavlich@arris.com

Dorin Viorel

Principal Architect
CableLabs
d.viorel@CableLabs.com

David Virag

Distinguished System Engineer
ARRIS
david.virag@arris.com

Charles Cheevers

CTO/CPE
ARRIS
charles.cheevers@arris.com

Introduction

For wireless communications, this is an unprecedented time. More licensed and unlicensed radio spectrum between UHF whitespace and millimeter wave mega-block partitions is being made available for commercial interests to invest in and grow business services than at any single prior point in history. The FCC is balancing competitive access for both licensed and unlicensed exploits with innovative dynamic spectrum arbitration promoting shared access in the 3.5 GHz CBRS band. The spectral largesse across all bands has predictably drawn enthusiastic attention from all the major MSO and MNO players with service expansion (or protection) interests at stake. The 5G area of wireless connectivity at scale, 10 Gbps speeds and millisecond or less latency has set in motion a burgeoning and perhaps somewhat pre-emptive set of wireless test trials aimed at establishing both technical merit and posturing some degree of “best stewardship” of the public airwaves.

The use of millimeter wave spectrum has sparked many debates about its architecture and economies — given the physics restrictions of primarily requiring “Line of Sight” to deliver the promise of multi Gigabits of wireless delivery. It is this non-determinism of signal propagation that has generated lots of research, innovation, and testing of solutions to create and define a deployable architecture that will support both Fixed Wireless Access and mobility uses.

This paper focuses on the hot industry topic: can a Fixed Wireless Access solution be developed to compete with or augment the wired broadband solutions today? It will examine the available spectrum options for delivery of a reliable, high-bitrate wireless connection over the last few hundred meters of Front-haul as an alternative to fiber-to-the-home (FTTH). These are the cases where a newcomer wants to overlay incumbent, existing greenfield opportunities, or CAPEX considerations render the latter alternative unsound. Leverage of the best attributes of near-line-of-sight (nLOS), non-line-of-sight (NLOS), and line-of-sight (LOS) signaling will be examined. The opportunity to extend a hybrid fiber/coax (HFC) plant by means of a wireless end network overlay will be analyzed for viability for the last 200m access to a home.

The economics of Fixed Wireless access lie somewhere in the following parameters:

- Cost of the spectrum used
 - As we discussed there are licensed and unlicensed bands to consider. There is also potentially new granular band usage in the 5G space that will follow some of the directions of the CBRS solution based on software based spectrum control
- The size of the cell for bandwidth distribution
 - As we move towards high bandwidth low latency wireless broadband services — the size of the traditional cell size is likely to reduce substantially. This is a function of propagation and distance of technologies like millimeter wave as well as a requirement to provision speeds of Gbps burst levels, which will be required for new Fixed Wireless Access networks
 - While there are some FWA plays that are trying to target 80 foot and higher Towers for 5G deployments — these are not likely to be the architecture for Gbps broadband services as they may scale for coverage at lower bandwidths and MDUs in dense areas — but are expensive in tower real estate lease and also won’t deliver the higher headline Gbps speeds. This may not be required for some overlay applications — but if the price of the 200 Mbps level SLA is reduced substantially — the cost of the CPE device dominates the economics of the solution

- The backhaul distribution and connectivity
 - As the cell size gets smaller it still should deliver multiple Gbps to enough customers to make it economically viable. This makes the backhaul to the cell important for speed and scaling to meet the front haul costs to consumers. The ideal solution is that a Fiber connect to every Gbps capacity Small Cell for 5G — likely at least 10 Gbps. There are some intentions and architectures to use Wireless backhaul — also using millimeter wave technology. Today there are backhaul solutions that utilize the unlicensed spectrum in the 60 GHz range. Depending on the Line of Sight of the wireless backhaul — different technologies and frequencies can be used
- The cost of the CPE equipment
 - This is one of the main barriers and inertia contributors to using Fixed Wireless Access. Even in sub 6 GHz access technologies — there is typically an outdoor transceiver to mitigate losses of the exterior walls. For microwave frequencies — like LMDS at 10 GHz — there is also the need for an external antenna and transceiver to mount on the MDU or outside Single Family Unit. This additional cost for CPE is a much debated and analyzed problem for 5G FWA solutions. One of the primary questions for a technology solution is to try and get the CPE transceiver to be self-installable by the consumer and added typically to the upper floor window of the home. The alternative is an outdoor mounted transceiver — which increases install operational costs and has serious ergonomic issues for homeowners and home owners' associations (HOAs). There are technology plays to try and mitigate some of the outside/inside connection problems — typically targeting wireless transmission through windows/walls to try and not have to drill holes for mounting/wires

These topics will be reviewed below to review the physics, architectures, and ergonomics of Fixed Wireless Access solutions. The paper will highlight the CBRS 3.5 GHz sub 6 GHz solution as well. While not a solution typically targeted at 5G (given the 3GPP requirements for 5G (speed and latency in particular) — there is some basis in the likelihood of 5G being a dual PHY or dual standard technology. Millimeter wave is not deterministic in its performance due to the environmental and NLOS issues — and the economics of deploying to the worst condition don't work, therefore there may be likely solutions that:

- Provide Small Cell with both Millimeter wave and sub 6 GHz and CPE that support both PHY
 - This makes the solution more expensive and larger in size and higher in power consumption
 - This makes the solution more expensive for LTE sub 6 GHz transmission for example in the CBRS band — as the small cell will be typically deployed in smaller cell range to meet the speeds and millimeter wave LOS requirements
 - This makes the solution more expensive for Wi-Fi sub 6 GHz transmission as the cell size would have to scale to Wi-Fi EIRP (Equivalent Isotropically Radiated Power) coverage
 - With dual frequency devices, the solution is more reliable with fallback on sub 6 GHz LTE when the millimeter wave transmission speeds drop due to changing environmental conditions
 - Additionally, there are multiple frequency overlays that may be utilized:
 - <1 GHz for potential NB-IOT applications. Lots of low bandwidth devices at maximum range. Standards running in these frequencies include 802.11ah, 802.11af, LoRA, SigFOX, LTE-M, LTE and others

- 1 GHz-6 GHz for licensed and unlicensed usage: LTE, LWA, LAA and Wi-Fi for mobility and broadband applications
- > 6 GHz for broadband and mobility applications and high bandwidth 5G applications

The 3GPP group is defining the specification around the addition of the 5G NR spec — allowing support for all three general frequency areas above. Most of the current trialing for Fixed Wireless Access solutions has been using the LTE MAC over different frequencies. Some of the trials have been using Wi-Fi MAC over 160 GHz bands and some using 802.11ad MAC at 60 GHz frequencies. The 3GPP standards group has a goal to complete and deploy the 5G NR spec and solutions in 2018 with the 5G MAC being specified by 2019. It's generally believed that Fixed Wireless Access solutions will deploy on the 5G NR specification or earlier — and may move to support the 5G MAC as well — when this is finalized for mobility in 2019/2020. Mobility solutions for 5G are expected to be deployed by 2021.

Content

1. Making the Technical Case

1.1. The Increasing Amount of Available Spectrum in the Sub 6 GHz and Higher Frequency Millimeter Wave Bands

In 2016, the FCC made available nearly 11 GHz worth of licensed and unlicensed spectrum for both mobile and Fixed Wireless exploitation, ostensibly to seed the rapid development of 5G infrastructure. 850 MHz of the new spectrum was located around 28 GHz, 3 GHz between 37 and 40 GHz, and a full 7 GHz between 64 and 71 GHz (this latter segment comprising the only unlicensed block). The year also saw the FCC finalize spectrum sharing rules on the 150 MHz of spectrum allocated to CBRS (3.55 – 3.70 GHz, specifically) in 2015.

From a technology perspective, there is often reference to the 'sub 6 GHz bands' (comprising LTE and Wi-Fi and other services — TVWS, NB-IOT etc.) and the millimeter wave bands (from 28 GHz and higher). For the purposes of this paper we will discuss the importance of both (sub 6 GHz having range and NLOS properties and millimeter wave nLOS or LOS) and the potential for the combination of both to be enablers for reliable 5G based services for FWA and mobile.

Before we discuss the Millimeter Wave bands — it is worth looking at the new CBRS frequencies and their capabilities and relevance to the MSO — as well as the potential for them to be combined with 5G services and also the platform of Software Managed spectrum access used by CBRS (Spectrum Access Service) to potentially also be leveraged by 5G and multiple spectrum solutions.

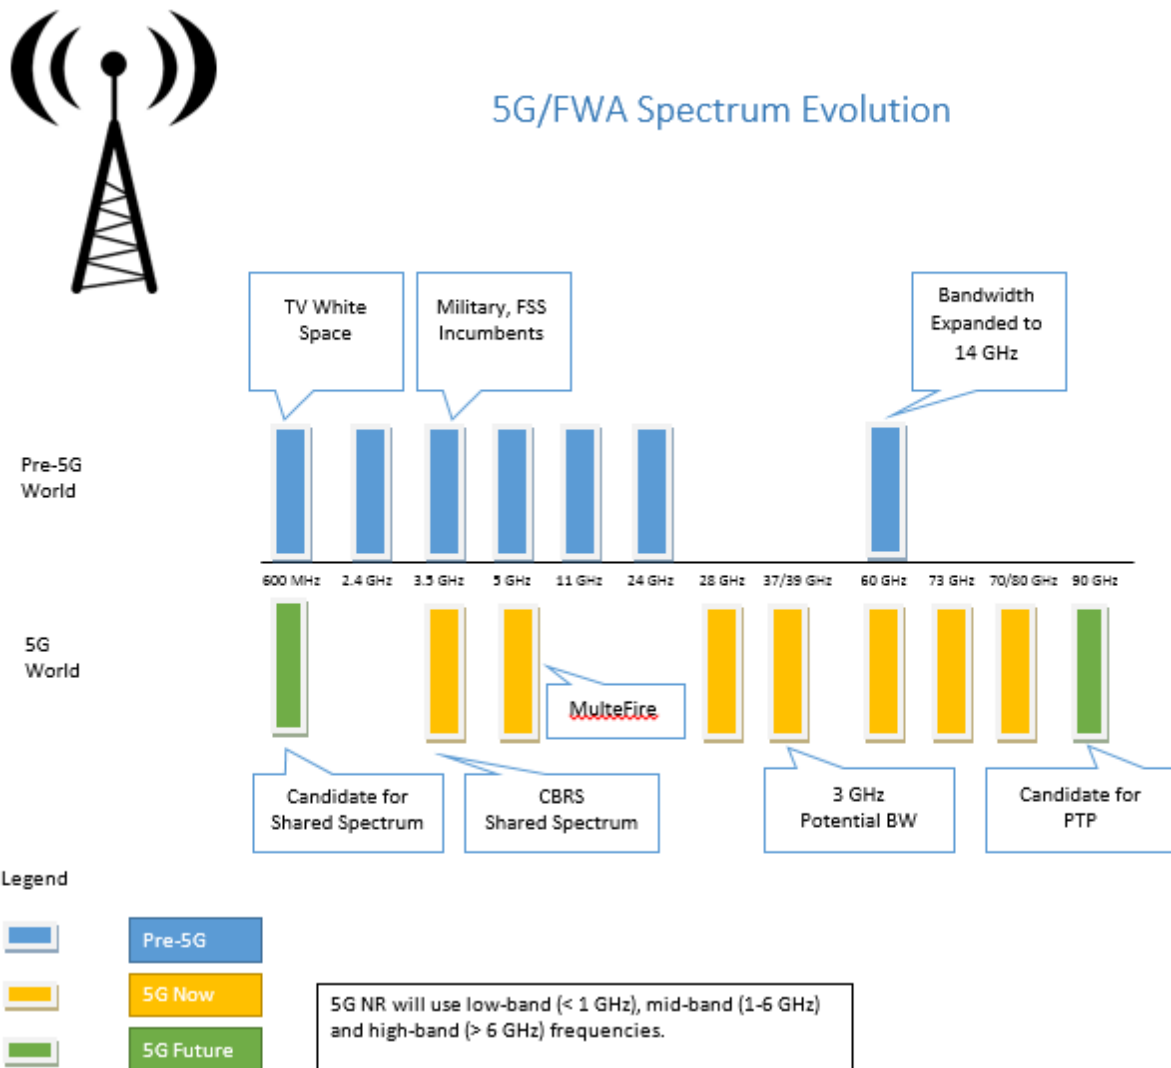


Figure 1 – 2016 FCC Spectrum Allocations

1.2. The 3.5 GHz CBRS Band

CBRS was envisaged by the FCC as an 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):



Figure 2 – Shipborne Radar Exclusion Zone (original:yellow, revised:blue)

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 operates on a highly granular geographic basis (census tract cell sized) which permits it to re-use spectrum on a tract by tract basis, 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.

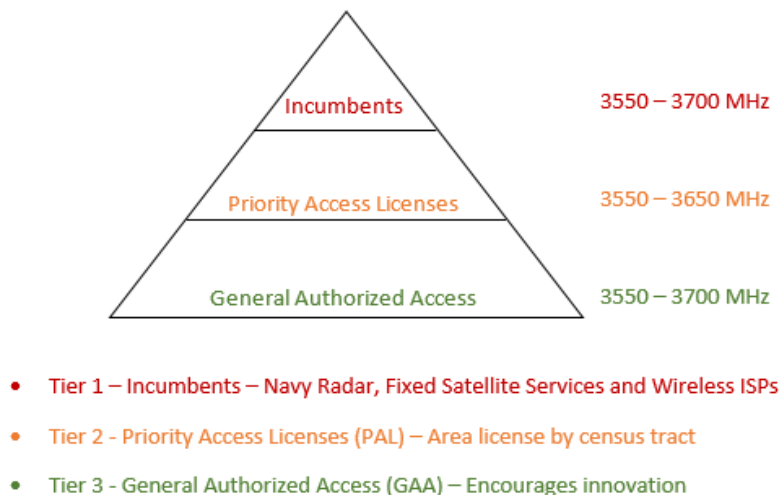


Figure 3 – CBRs Priority Tier Membership Distributionⁿ

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 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 (CBSD) 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 4 – CBSD Category A and B Power Signature Limits

1.2.1.1. Distance and Path Loss

3.5 GHz has similar nLOS behavior and propagation characteristics to other sub 6 GHz mobile carrier mid-bands.

$L = 10 n \log_{10}(d) + C$, where $n=2$ is the Path Loss (PL) Exponent in Free Space

$PL(d) = 20 \log_{10}(4\pi d \times f/c)$ Note: $20 \log_{10}(4\pi / 300,000,000 \text{ m/s}) = -147.56$

$$\begin{aligned}
 3.5 \text{ GHz, PL (200m)} &= -147.56 + 20 \log_{10}(d) + 20 \log_{10}(f) \\
 &= -147.56 + 46.02 (@200m) + 190.88 (@3.5 \text{ GHz}) \\
 &= 89.34 \text{ dB (3.5 GHz, 200 meters)}
 \end{aligned}$$

$$\begin{aligned}
3.5 \text{ GHz, PL (800m)} &= -147.56 + 20 \log_{10}(d) + 20 \log_{10}(f) \\
&= -147.56 + 58.06 (@800m) + 190.88 (@3.5 \text{ GHz}) \\
&= 101.38 \text{ dB (3.5 GHz, 800 meters)}
\end{aligned}$$

$$\begin{aligned}
2.5 \text{ GHz, PL (200m)} &= -147.56 + 20 \log_{10}(d) + 20 \log_{10}(f) \\
&= -147.56 + 46.02 (@200m) + 187.96 (@2.5 \text{ GHz}) \\
&= 86.42 \text{ dB (2.5 GHz, 200 meters)}
\end{aligned}$$

$$\begin{aligned}
2.5 \text{ GHz, PL (800m)} &= -147.56 + 20 \log_{10}(d) + 20 \log_{10}(f) \\
&= -147.56 + 58.06 (@800m) + 187.96 (@2.5 \text{ GHz}) \\
&= 98.46 \text{ dB (2.5 GHz, 800 meters)}
\end{aligned}$$

The overall path loss is similar (~3dB difference) in nature between the 3.5 GHz and the 2.5 GHz band which is to be expected. Sub 6 GHz delivery at 200m to 800m remains a viable NLOS solution that provides 100s of Mbps of broadband capability. The potential to go to Gbps peak levels with more bands can be realized with CAT16 and CAT18 solutions including Carrier Aggregation and use of Licensed Assisted Access (LAA) within the 5 GHz Wi-Fi bands. Wi-Fi solutions in the 5 GHz band with 8 spatial streams also go to multi Gbps of Wi-Fi capacity.

1.3. Requirements

1.3.1.1. LOS Delivery

Analysis of compound annual growth rate (CAGR) metrics for both average and peak data connectivity demands in the home network indicate that by 2020, service group (SG) downstream (DS) needs — for SGs with a roster of at least 100 clients — will be touching 3 Gbps.




Figure 5 – Historical and Projected Bitrate Consumption Growth at the Service Group (Cloonan's Curve)¹

The profile of connectivity will dramatically migrate from a “lean” ratio of roughly 9:1 DS:US rates to something approaching a 6:1 ratio, due to client upstream (US) IoT cloud needs (particularly in regards to security camera feeds). For Fixed Wireless Access to be able to compete with Wired Solution, it has to be able to provide Gbps peak rates on both the Downlink and the Uplink. To achieve these peak data rates in a Wireless Access network the only wireless bandwidth available with sufficient contiguous spectrum to meet 3+ Gbps SG downstream service — *if* the spectral efficiencies remain under 10 bps/Hz -- lies well into LOS-delivered (and near millimeter wave) frequencies (28 GHz, 37 GHz, 39 GHz, 60 GHz and 64-71 GHz). These frequencies offer huge amounts of bandwidth (the unlicensed bands alone in the 60 GHz range can deliver 128 Gbps), but these frequencies present some clear technical, operational, and aesthetic challenges.

For example, millimeter wave LOS does not tolerate path loss variation well (as would be the case for smoke, fog, precipitation, wind-driven foliage, or moving bodies which cross the delivery vector to the client). Adaptive modulation schemes see to it that poor SINR signature at the receiver triggers a fallback to less complex constellations (reduced information density — hence, reduced delivered bitrate) to aid the receiver in resolving and decoding the (hopefully temporarily) compromised signal. But physically moving LOS apertures between Point of Presence (POP) — typically the small cell on a pole or street furniture at < 20 feet) — and client to more open pathways to reduce the risk of signal interdiction or attenuation inevitably drives the POP above both terrain and flora masks — and probably requires a similar elevated location for the client's antennae. The increase in vertical height comes with a price tag

of extra labor hours to access devices, resolve look angles, align antennae, and verify link performance. There is also the general view that 5G must work at lower street furniture elevations to get the desired connectivity and smaller homes passed cell size. Equally for applications like 5G connected autonomous cars, the addition of 5G small cells along freeways with street lights may be the end game solution to facilitate latency and speed to car requirements. The U.S. for example has 26 million street lights all with power connections. 60% of these street lights are owned and operated by the private sector with 40% owned by city and state municipalities. Public dollars pay for all the street lights to operate with energy costs of \$2Bn alone for lighting. Figure 6 below illustrates the simplistic problem of LOS technologies where environmental conditions, terrain, and foliage can affect the performance of any transmitted signal. Higher elevations do not always solve the problem and the erection of new 80 foot and higher towers is in many ways not feasible. There will be of course a macro tower based Cell with fill ins of smaller cells on other emerging POPS — particularly street lighting and other elevated powered locations.




Figure 6 – Terrain and Foliage Masking Effects on LOS Delivery

To illustrate the problem further the following pictures (Figures 7-9) illustrate modeling of LOS millimeter wave propagation in a typical random Georgia residential subdivision at different heights of 2m/6 foot, 10m/30 foot+ and 33m/100 foot pole. The effect that trees and Foliage have is reviewed below in later sections. As you can see — even with the investment a 100 foot tower — the number of houses in the LOS ‘Look’ is still relatively low.

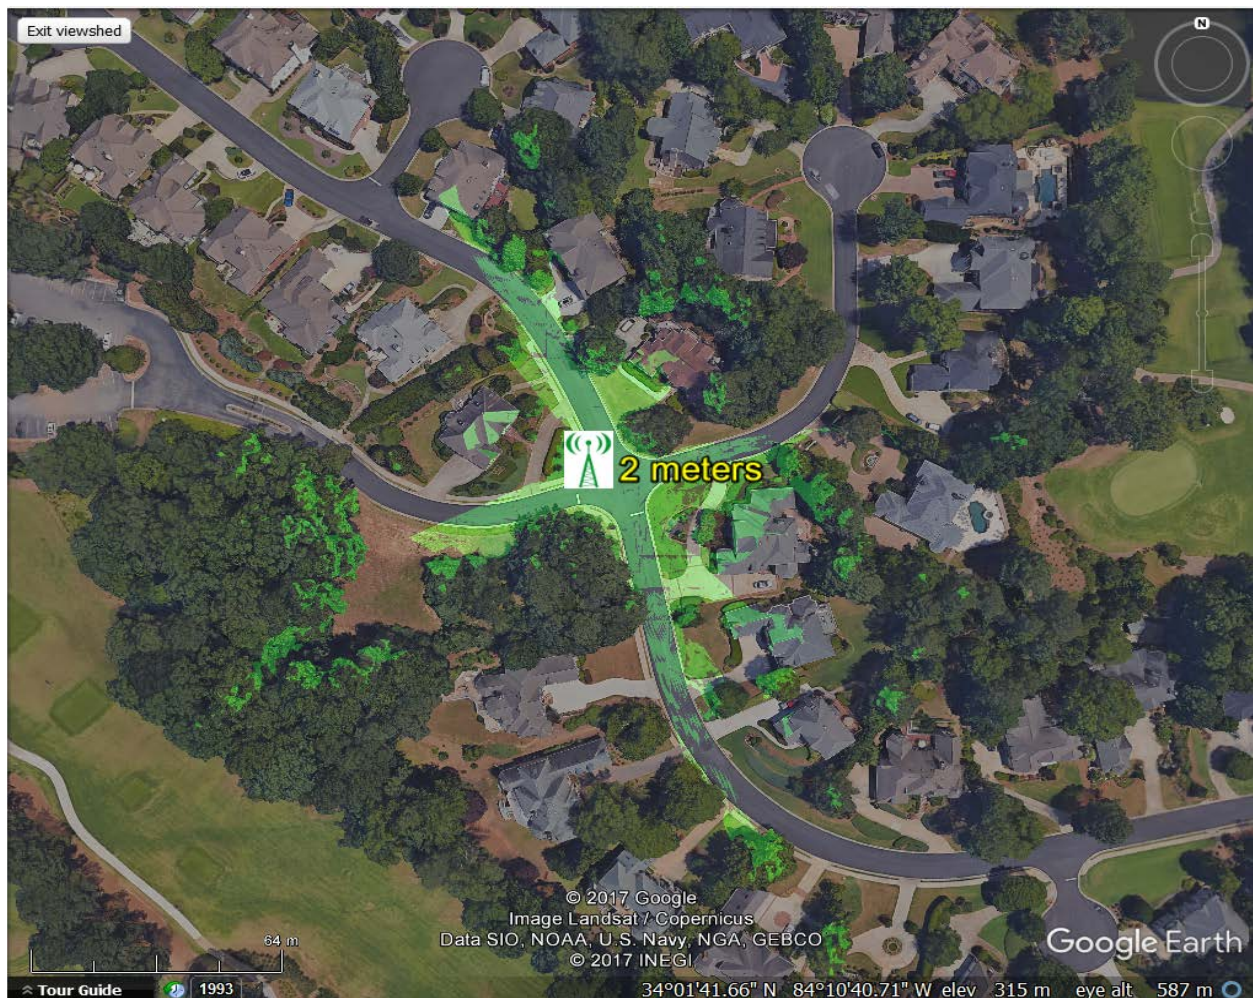


Figure 7 – LOS “Look” Mask from 2-Meter Pedestal (lime green shadow)



Figure 8 – LOS “Look” Mask from 10 Meter Mast (lime green shadow)



Figure 9 – LOS “Look” Mask from 100 Foot Monopole (lime green shadow)

This is the theoretical LOS “look” mask. In practice, there are other difficult challenges — including the following:

- Planning permission for a new tower to be erected
- Use of street furniture especially street lights as an alternative — foliage can wrap lights and private companies operate many of the lights so deals need to be made with different companies and municipalities
 - Putting more than one Cell on the Light fixture for more than one provider may not be practical and may drive a more neutral host requirement for any investment in street lights for 5G usage
- Adding backhaul capacity to the street furniture or erected tower will also cause disruption
- Adding backhaul capacity using Wireless backhaul increases the Small Cell size and power requirements and requires positioning for both best backhaul and customer LAN side connection

- The question as to whether the client device can be placed inside the consumer's home is one of the biggest open issues for 5G deployment
 - There is high desire to make 5G FWA client self-installable — to improve the economics of 5G wireless as a replacement to incumbent solutions
 - However, the reflection coefficients and penetration losses for building materials in the millimeter wave bands make this a difficult problem to solve and to get any amount of bandwidth into the home even through windows. tinted windows, double and triple glazed windows with energy efficient glass — also pose even larger problems than wooden shingles. Even using 28 GHz (which has lower atmospheric absorption compared to 60 GHz) which is comparable to 1-2 GHz for free space path loss — there is 25dB to 50dB loss when the window is metal coated. Today, external glass is tinted and coated with metal to provide a block to ultra-violet rays and improve insulation. For comparison, clear glass is as low as 2dB and Plasterboard with metal studs is about 9dB loss. Wall construct ranges from low loss 6-7dB for 1 foot of brick to 25dB for a 3 foot brick wall. This makes the challenge of a home self-install difficult at least at the multiple hundred-meter range. It suggests that 5G small cell must be closer to homes for both self-install and for Gbps capabilities. The tradeoff is to have an external technician-installed antenna and transceiver

There are other issues like HOA covenants regarding external mounted devices and their location — often this forces satellite dishes (or the equivalent outdoor 5G transceiver) to be put on the rear of house. This hinders the economics of deployment on Street Lights, potentially forces new poles to be placed around the subdivision vs. using the interior infrastructure. It does not help that millimeter LOS propagation under even pristine conditions does not survive a service radius much more than 300 meters or so (typically more like 200 meters) in order to deliver more efficient modulation schemes. Such constrained throw feeds directly into a monopole (antenna tower) density calculation where it becomes all-too-clear that these skyline-altering structures would necessarily be much more visible than cell towers -- typical spacing for these latter elements perhaps reaching a 10x more sparse seeding — variously estimated in the literature to around the 1-2 km range in suburban areas). See below Figure 10.



Figure 10 – Example of LOS (left) vs NLOS (right) Monopole Seeding in Golf Community Development

Operationally, the implied Operational Cost overhead on site planning comprises both a shared analytical element (locating the POP to best service the maximum number of clients while maintaining non-overlapped antenna patterns, for example) along with a per-client “tuning” aspect which aligns the antenna elements and determines the best placement solution at the client site from a technical and aesthetic perspective.

However, the downside can be marginalized. The LOS POP which serves the 200-300-meter client radius footprint can effectively be structured to future-proof service bitrate demands such that each POP antenna element / client antenna pairing reuses the entirety of the available millimeter bandwidth (assuming the necessary unicast service switching density in the base station and sufficient wireline backhaul bandwidth to optimize the POP). This suggests that the wireless portion of the delivery architecture can be re-plumbed to upgrade QoE without modification to the endpoint RF frontends (including antennae). The cost and size of the 5G Small Cell POP is then another aspect. Increasing the number of element arrays improves the ability to pair and serve specific clients. That is another area of research to optimize deployment economies and future proofing.

1.3.1.2. nLOS Delivery

Sub-6 GHz carrier radios can be bound with Mu-MIMO smart antenna frontends to effect steerable, beam-formed radiation patterns which can be delivered with per-user-location agility and offer adaptable tolerance for propagation path impediments — while still delivering dense constellation modulation formats over greater distances than those which put paid to mmWave LOS signaling. A canvas of the data provided in the Field Test sections indicates that the nLOS service radius at 3.5 GHz ought to approach 800 meters — and this includes penetration of dwelling walls.

The immediate takeaway is that sub-6 GHz nLOS wireless connectivity requires POP seeding density on the order of 6% or so — $(200/800)^2$ -- of that necessary to insure LOS connectivity. More notably,

perhaps, is that this elevated mounting requirement can opportunistically leverage 2nd-story “street furniture” instead of 40-80-foot tree topping monopoles required for LOS installations. Where such leverage is unavailable, shorter wooden masts, strand mounts or pedestals may be used and their location optimized only by considerations of reducing client azimuth contention (and not by clean LOS apertures) — Figure 11 below.

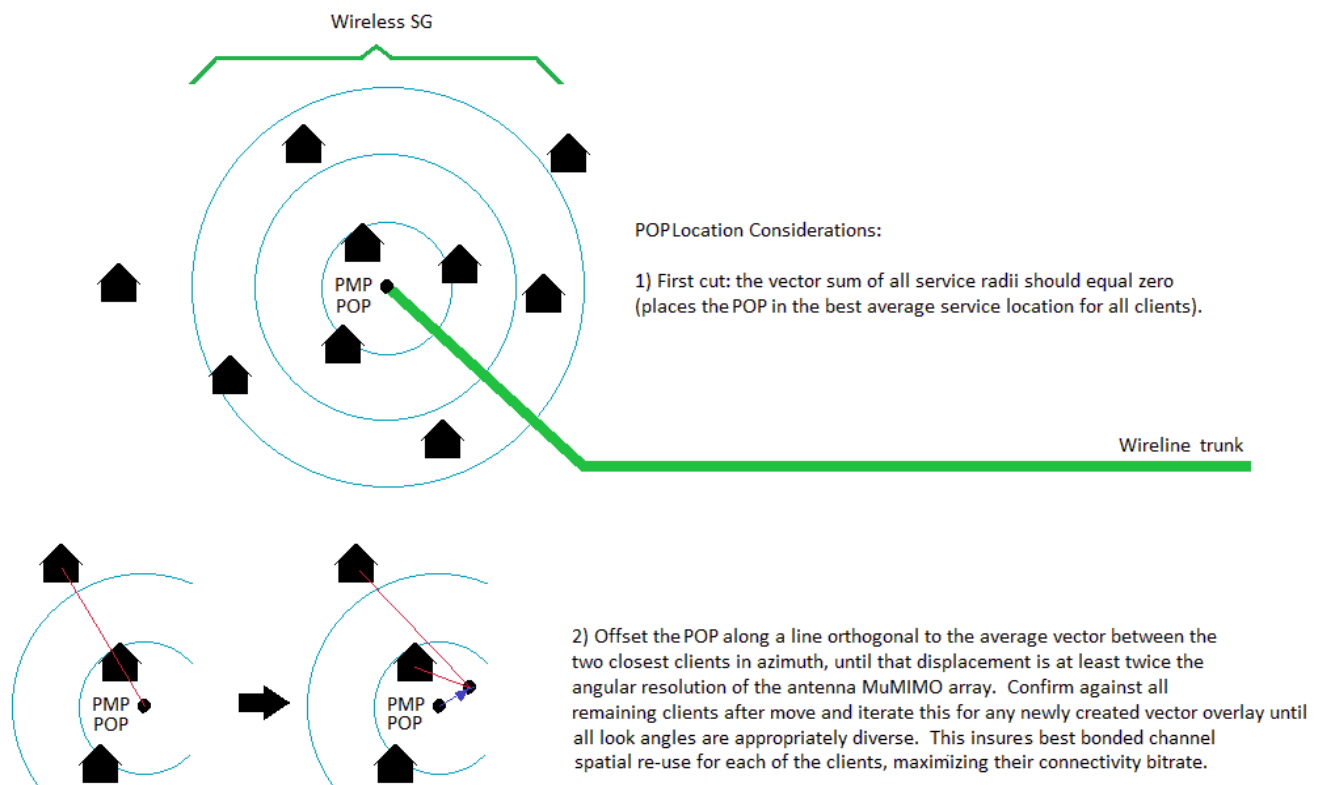


Figure 11 – nLOS Site Survey Fine Tuning of Monopole Axis to Obtain Azimuth Diversity

The remaining issue however, is that 3.5 GHz CBRS, a new useable swath of sub-6 GHz delivery due to its 150 MHz of contiguous (if access-tiered) spectrum, still cannot promise better than 1.5 Gbps if one accepts 10 bps/Hz as an efficiency asymptote. (And this is made slightly worse by the observation that only 70 MHz can be guaranteed by PAL license). Critically, however, this spectral efficiency caveat looks to be rendered “overcome by events” (OBE) by ongoing research into massive MIMO antenna arrays. These arrays offer new potential especially for the client side — allowing the client to host multiple antennas in limited footprints.

On the POP/Small Cell with multiple Antennas, it serves a set of single antenna users and the multiplexing gain can be shared by all users. The good news is that the array geometries, though of daunting size for indoor CPE operating at 3.5 GHz, are eminently suitable for base stations associated with POP antenna masts — and a seminal, demonstrable aspect of massive MIMO signaling in the Fixed Wireless domain (which eliminates sounding offsets due to Doppler effects) is that the channel’s MIMO signature can be presumed reciprocal. This implies that, within the bounds of sufficient link SINR, massive MIMO need not be symmetric on link ends to extract benefit. (Refer to the section below on Spectrum Efficiency.) See below some examples (Figures 12,13, and 14) of the different gains of an 8x16

element array and a 4x4 element array at 60 GHz. 8x16 Array yields a 25.3dBi gain — while 4x4 only a 16.4dBi gain.




Figure 12 – 8x16 Element Array and Dimensions

8 × 16 Array

Max realized gain at 60 GHz is 25.3dBi.




Figure 13 – 8x16 Array - Gain Calculations in the XY, XZ and YZ Planes – Max Gain 25.3dBi

4 × 4 Array

Max realized gain at 60 GHz is 16.4dBi.




Figure 14 – 4x4 Srray – Gain Calculations in the XY, XZ and YZ Planes – Max Gain 16.4dBi

1.3.1.3. Connections at the Home

External signal propagation characteristics largely define the architecture of the home network connection with LOS millimeter downlinks exhibiting notoriously poor dwelling penetration performance, which is the opposite of sub-6 GHz nLOS. In the case of the LOS solutions, in all but the most (accidentally) favorable of client locations (nearly adjacent to the POP with a clear field of view and predominantly fair weather), an external downlink transceiver (or, at minimum, an external antenna and downconverter) must be employed to capture the LOS external signal and convert it for presentation to the gateway (WAN) agent of an internal home network. This is still to be debated — and as mentioned above in particular if there is investment in the POP/Small Cell with closer proximity to the home perhaps 50M, more element arrays and potentially relaxing the served bandwidths to under Gbps peaks. The potential outside to inside architectural aspects are listed below:

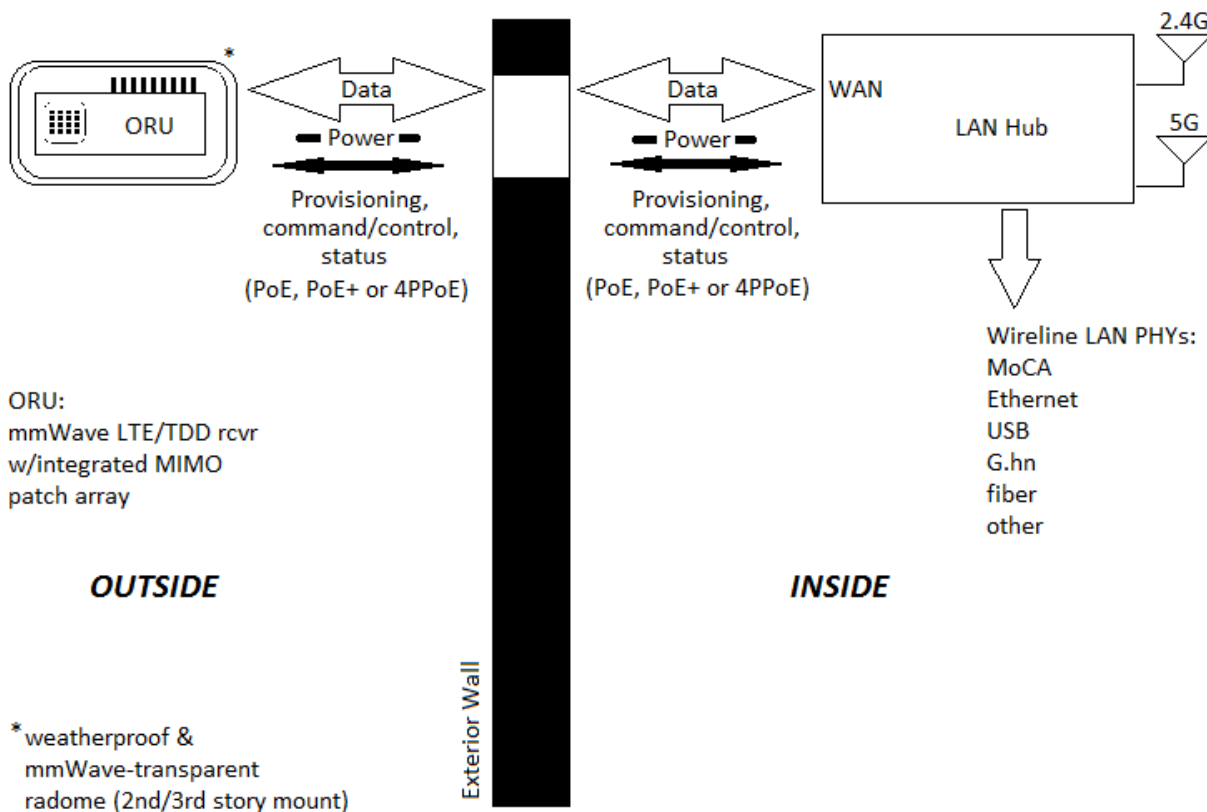


Figure 15 – Implementation Option on LOS-to-Home-Network Interfaces

There is another aspect to this outside mount and inside connection model — that is also problematic and needs to be reviewed. The most optimum location for this outside device is as high up as possible typically. If we assume that roof install is too costly — and the goal is to leverage a Window mount solution — then there are a few considerations that are being worked through.

- Locating the Baseband or Home NTU (Network Termination Unit – LAN HUB above) in a second-floor room typically means an occupied bedroom. Generally consumers don't like Gateways or Routers in bedrooms: the real estate located to them is constrained, Homeowners don't like putting devices on their bedside lockers or in baby's rooms and overall people don't like blinking LED's or bright single LED's in rooms where darkness is desired for sleeping
- The external unit itself if mounted on the window — has the challenge of powering and getting the analog to baseband signal into the home/bedroom
 - There are several companies trying to solve this problem with
 - Window sash mount solutions for conveying the power and data from inside.
 - Window mount on pane using sticky compounds or magnets
 - Inductive Powering through Window — trying to drive to 8W-10W
 - RF technologies to send data through window from outside pad to inside pad
 - Even with these solutions the residential customer must give up the aesthetics on an upstairs window and potentially the view. There are also other issue considerations such as insect screens on windows in warmer states in the USA

The requirement for external signal conversion is not mirrored in nLOS, sub-6 GHz networks. Field reports indicate that 3.5 GHz CBRS propagation, for example, tends to proxy that of 2.4 GHz — making it the home ingress equivalent of legacy 802.11 Wireless Local Loop. The upshot of this good news is that self-contained 3.5 GHz CPE with integrated Mu-MIMO antenna elements can be placed at will in the home and successfully recover signal (though a clever onboarding application would facilitate self-install by leveraging antenna MIMO data to suggest placement and optimize reception). As it proxies the WAN connect afforded by a DOCSIS Cable Gateway/Wireless AP (the sub-6 GHz radio, small MIMO and LTE/TDD receiver substituting for the DOCSIS subsystem), the cost and mechanical footprints of the NLOS gateway ought to mimic that wireline box. And the allowable transmit power of 3.5 GHz CBRS (+30 dBm or 1W) provides enough punch to be able to craft a neighbor mesh to sidecar some amount of opportunistic node+1 network attachment from a functioning neighbor's network feed in the event of failure of the to-home main LOS downlink.

1.4. Data and Analysis from various Field Tests

1.4.1.1. *Qualcomm supplied Research data^m*

Qualcomm has investigated propagation and loss behaviors for multiple bands (sub-6 GHz through millimeter) in an ongoing study aimed at determining 5G wireless technology adaptations. Bands and areas of interest are listed below.

Channel Measurements/Simulations

Scenario		Description	Frequency Band
Measurements	Materials	Various construction materials, humans, etc.	22-43 GHz
	Foliage	Various tree species	29 GHz
	Indoor	Residential (interior/exterior walls)	22-67 GHz
	Indoor	Emulated Stadium	2.9 GHz, 29 GHz
	Indoor	Bridgewater shopping mall	2.9 GHz, 29 GHz, 61 GHz
	Urban micro	Street Canyon	2.9 GHz, 29 GHz, 61 GHz
	Indoor	Office	2.9 GHz, 29 GHz, 60 GHz
	Urban micro	New Brunswick	2.9 GHz, 29 GHz, 60 GHz
Scenario		Description	Frequency Band
Simulation	Indoor	Dense office	29 GHz, 38 GHz, 60 GHz
	Urban micro	High density (Manhattan)	29 GHz, 38 GHz, 60 GHz
	Urban micro	Low density	29 GHz, 38 GHz, 60 GHz
	Indoor	Stadium	29 GHz, 38 GHz, 60 GHz

Confidential and Proprietary - Qualcomm Technologies, Inc.

Figure 16 – Propagation Study Subjects by Band

A key study area was the ability of the selected carrier frequencies to penetrate dwelling apertures or materials. As can be seen below, statistical variation over common materials, sidings, and window type were tremendous (and as such, likely defy an “average” characterization of transition losses over client distributions which could include most permutations listed in the tables).

Summary of Out-to-In Propagation Loss


•Measured penetration losses for various materials

Residential 1	
Material	Loss
Vinyl siding	~6-7 dB
Stone siding	~35 dB
Window glass	~10 dB
Plastic blinds	~2 dB

Residential 2	
Material	Loss
Plywood	~8-10 dB
Hollow sheetrock	~1-2 dB
Wood exterior wall/panel	~10 dB
Brick exterior	~30 dB
Metal doors/window frames	high


Residential & Commercial	
Material	Loss
Commercial Tinted Window	10-20 dB
Clear glass	2.5 dB
Residential Home Exterior	~9 dB

•Results consistent with other industry sources



Sources: Samsung and Nokia

Confidential and Proprietary - Qualcomm Technologies, Inc.




Sources: Qualcomm, DCM, E// and Huawei (basis of 3GPP & 5GTF O2I model)

Figure 17 – Penetration Loss by Material Construction and Frequency

In regards to the wireless propagation environment, however, it is not only client endpoint buildings but interdicting topographical features and weather which can degrade the channel. Channel fade due to weather has been well characterized over the lower parts of the bands in question due to decades of satellite link use and the resort of high ground and towers can be plainly shown to address topology contours. However — in North America in general — large swathes of geography are dominated by trees and other foliage which, depending on seasonal growth and longitude, can interrupt a good many LOS apertures between BS and client and present performance challenges. Data below is captured in two bands — the 2.9 GHz being representative of nLOS/NLOS sub-6 GHz carriers and the 29 GHz proxy's behavior at millimeter wave frequencies. The impact of deciduous and conifer trees (under gusty wind conditions) suggest that the leaf density from the conifer more frequently produces heavy link losses and these, more so at higher carrier frequencies.

Foliage (Trees) Attenuation at 29 and 2.9GHz




(*) Variations include both spatial and temporal sampling

(**) At 2.9GHz antenna aperture (224mm x 169mm) may be too large for accurate measurement in this case

Confidential and Proprietary - Qualcomm Technologies, Inc.

Figure 18 – Scattering Losses at sub-6 GHz and mmWave for Foliage

Signals incident upon exterior construction have a difficult time penetrating exterior walls — more so with brick than siding (refer to data above). However, in even the lowest loss scenarios (lap siding and the like) there is considerable gradient to the losses with frequency (~ 5 dB sub-6 GHz up to 17 dB in the unlicensed millimeter wave band).



Note: Values indicate the low 50th percentile penetration loss for the bands

Figure 19 – Exterior Lap Siding Penetration Losses over Frequency

Internal walls seem more homogeneous whether the construction is residential or office. In both cases, the loss over frequency or construction appears to be in the range of 3.5 dB.




Figure 20 – Interior Wall Penetration Losses, Home (L) and Office (R)

An overall view of interior CPE performance at distances from 50-600 meters from an outside POP/Small Cell operating at 28 GHz was performed and the probability of outage measured for both 8-element (solid lines) and 64-element (dotted lines) MIMO arrays (outage being defined as producing < 100 Mbps). In the graph below, the Y axis is the empirical cumulative distribution function (CDF — jumps of 1/n at each of the n data points) and X axis is the Throughput in Mbps.




Figure 21 – Probability of Service Outage on Inside CPE Vs Bitrate Demand and Distance to BS

As can be seen, for 50 and 100 meter distances, essentially all CPE could produce at least 100 Mbps (as opposed to 34% outage for the simple antenna array at 200m, falling to 85% outage for even a 64-element array at 600m). Even at 50 meters with the best antenna, however, a full third of the CPE population could not muster the expected 400 Mbps. Moved to 200 meters' distance, another third of the population was throttling back (70% or so experiencing difficulties).

At a fixed service distance to the POP/Small Cell of 200 meters, moving the CPE or at least the Antenna/Transceiver portions outside the dwelling bought back significant performance.




Figure 22 – Comparison of Indoor vs. Outdoor Throughput for 2 MIMO Arrays @ 200m

Getting the *outage* CPE outside the home and equipped with a 64-element array meant 95% of that originally defunct population recovered from outage to consistently produce at least 100 Mbps.

1.4.3 CableLabs Analysis and Testing

CableLabs has been engaged in multiple 5G (millimeter wave) field trials to evaluate the opportunities and limitations of using millimeter wave 5G Fixed Wireless links to provide Multi-Gbps services to the end user. Based on the developments in the 5G ecosystem, the field trials focused on the 28 GHz, 39 GHz, and 70 GHz bands. An extensive list of KPIs was collected during the field trials, but for the scope of this paper, we will limit the discussion to spectral efficiency and link length.

Channel bandwidths in the mm-wave bands are significantly larger than typically used wireless channels. For example, the 28 GHz band is divided into 425 MHz wide channels, the 39 GHz band is divided into 200 MHz wide channels, and the 70 GHz band is divided into 1.25 GHz wide channels. Thus, even under modest spectral efficiencies, such bandwidths can potentially enable multi-Gbps services to the end user.

It is worth noting that the field trials were performed using systems that are still under development, thus further performance improvement can be expected as the systems mature.

1.4.3.1 28 GHz vs. 70 GHz Performance Comparison

CableLabs conducted a millimeter wave link field trial focused on evaluating the impact of millimeter wave spectrum on link performance. 28 GHz and 70 GHz bands were selected as they represent the lower and higher bands of the millimeter wave spectrum.

The test setup mimicked a wireless drop scenario to a residential household; the transmitter and receiver were placed approximately 240 foot. apart, and a single spatial stream was used in the evaluation. The evaluation focused on the impact of simple channel impairments (single impairment at a time, no compound impairments) on the link performance relative to a line of site link.

Figure 23 below summarizes the impact of single channel impairments on the link's spectral efficiency. The link under evaluation was a SISO link, thus for MIMO links (for example 2x2 MIMO), the spectral efficiency can be scaled accordingly.




Figure 23 – 28 GHz vs. 70 GHz Spectral Efficiency

As can be seen in Figure 23, in LOS conditions and slightly obscured channels, a MIMO system can achieve ~ 7 bps / Hz, which translates to a minimum requirement of 150 MHz of channel bandwidth to be able to support 1 Gbps link capacity. As mentioned previously, current regulations provide 425 MHz wide channels in the 28 GHz band, and 1.25 GHz wide channels in the 70 GHz band.

The performance of the 28 GHz and 70 GHz links in short distances are comparable in favorable channel conditions, but the 70 GHz band is more susceptible to channel impairments, especially moisture bearing channel impairments. As seen in Figure 23, the 70 GHz link's spectral efficiency drops by 15% due to tree foliage, and the link is completely lost when going through an evergreen tree.

The difference between the 28 GHz band and 70 GHz band becomes more observable upon evaluating the cell edge of the link. Figure 24 below represents simulated results for the maximum distance from the transmitter to achieve a target modulation order. As shown in Figure 24, the link length (defined by the lowest achievable MCS) is reduced by 75% in the 70 GHz band in comparison to the 28 GHz band, which in turn means a ~16x increase in cell density to achieve the same coverage.




Figure 24 – 28 GHz vs. 70 GHz Link Distance

Based on the results from Figure 24, one would assume that deploying a wireless network based on 70 GHz band would incur a much higher deployment density in comparison to 28 GHz. This would be true if advances in antenna technology are not considered. For the same size antenna panel, a 70 GHz antenna panel can have 6 times the density of antennal elements compared to a 28 GHz antenna panel, which translates to ~7 dB gain for the same panel size. A 7-8 dB gain allows for 70 GHz links to have similar link lengths as 28 GHz links (assuming favorable channel conditions); thus 70 GHz remains a viable option for delivering multi-Gbps connectivity to end users.

1.4.3.2 37 GHz Field Trial

In addition to the 28 GHz and 70 GHz field tests, CableLabs conducted a 37 GHz field trial using millimeter wave system under development. The objective of the field trial was to evaluate the coverage and capacity of 200 MHz wide links operating in the 37 GHz band under various channel conditions.

The system under test had advanced MIMO and beamforming capabilities which are some of the fundamental features in 5G. The benefits of MIMO and beamforming can be leveraged to deliver high capacities and/or extended coverage even in the presence of channel impairments.

As shown in Figure 25, link capacities of approximately 750 Mbps were achievable in LOS conditions, which were degraded to just under 490 Mbps in adverse weather conditions. Also of particular interest was the maximum link length that can be achieved to deliver service where a LOS link extending approximately 2600 feet while delivering nearly 190 Mbps was demonstrated.

As shown in the 28 GHz and 70 GHz field trials, the 37 GHz links are also highly susceptible to channel impairments, where a 70% link capacity reduction was observed due to foliage and a reduction of approximately 90% in link capacity due to dense foliage.




Figure 25 – 37 GHz System Performance & Spectral Efficiency (values above columns)

1.4.3.3 Field Trials Summary

Based on the results of the field trials, Fixed Wireless networks leveraging millimeter wave links hold the potential of delivering high speed service to end users. Large channel bandwidths and advanced signal processing techniques such as MIMO and Beamforming are key enablers.

Nonetheless, the susceptibility of millimeter wave links to channel impairments such as tree foliage adds an amount of complexity in deploying such networks. Channel impairments can significantly reduce link and system capacity.

The results from the 70 GHz field trial indicate that the unlicensed spectrum extending from 64 GHz to 71 GHz can potentially be used to deliver Multi-Gbps Fixed Wireless services, without the need to acquire licensed spectrum.

1.5 nLOS/NLOS Advantages in Reach and Setup

Despite modest available spectrum (relative to that liberated for use from 28-71 GHz), sub-6 GHz maintains a key benefit over millimeter wave in terms of wavelength-related propagation improvement (~20 dB lower losses over the same distance), more birefringent than outright scattering behavior with respect to encountered materials in the spatial delivery channel and much lower tendency towards absorption by atmospheric particles or moisture. From POP/Small cell to client, through the air and around or through man-made or natural interdicting materials, sub-6 GHz based signal delivery provides a much more robust probability of successful recovery at the receiving end. Real-world numbers suggest a usable service radius of ~800 meters for the power budget allowed (1W for the client and 50W for POP/Small Cell Category B, in the case of 3.5 CBRS) and may permit indoor use without resort to a piped outside antenna provided a capable enough BS is employed. LOS systems seem bound to perhaps a quarter of that distance (and dare not risk inside-home antenna mounts much past 50-meter range).

In terms of equipment setup, CPE with only a rudimentary MIMO (say, 2x2) ought to be able to pilot, or sound, the channel such that the home user can be iteratively guided to the best in-home location and orientation with an out-of-box onboarding application which levers the CPE's interpretation of the smart antenna tuning parameters relayed by the BS. The upshot here is that new clients for a given nLOS POP ought to be capable of self-installation without the service provider resorting to a truck roll. However, in the event of a technician-assisted install, such should be possible with the complications associated with placing or aligning outside pole-mounted antennae.

1.6 Bandwidth Advantages in mmWave

There is no question that for sheer available bandwidth, the bands above 20 GHz provide instantly scalable, multi-Gbps bitrates at only modest spectral efficiencies. Even with the “narrow” channel bandwidth of 200 MHz assigned to the 39 GHz band (a 33% uptick over *all* the channel BW available at 3.5 CBRS), a pedestrian 7 bps/Hz yields nearly 1.5 Gbps (the equivalent of 35 or so bonded DOCSIS 3.0 channels). And while issues of service radius and link availability (due to atmospheric or topographic masking) give pause, the PHY's raw capability might prove an exploitable contingency in certain use case scenarios where alternate means are either too expensive or not yet ready.

1.7 Spectrum Efficiency: the Power of Massive MIMO

The two-species wireline legacy of cable — coax and fiber — provide a loss gradient and ingress protection which promote E2E spectral efficiency about 10 bps/Hz. A full and fair comparison to the prospect of wireline delivery needs to acknowledge the rather precipitous and (to-date) unappreciated impact of massive MIMO antenna arrays (rather loosely defined as those $T \times R$ matrices in densities larger than 8×8) in achieving an order of magnitude better spectral density over short-haul (≤ 800 m) wireless links.

For example, wireline's DOCSIS 3.1 pride (4K-QAM-based) promises a 12 bps/Hz spectral efficiency (before overheads are invoked). However, in the wireless space, one year ago, academic proof-of-concept (POC) work has yielded data which confirms that a 160-element planar antenna array (its area befitting a mount on a short monopole — see below) could support over 20 fixed and walking-pace mobile, single-antenna clients to the tune of nearly 150 bps/Hz (this, leveraging 256-QAM over a 20 MHz channel at 3.7 GHz)⁶. The extensibility to a 3.5 GHz CBRS POP/Small Cell is obvious and if one were to include channel aggregation of the full available PAL-guaranteed bandwidth of 70 MHz, the result would amount to ~ 10 Gbps. Additionally, if one were to instead apply this to a LOS solution at 37 GHz, the planar array size would collapse to an area of 1% of that used for the 3.7 GHz POC (due to the wavelength shrinking by a factor of 10) and, assuming just the 200 MHz of single-channel spectrum, produce a yield in bitrate of more than 30 Gbps. obviously scalable to whatever channel-aggregation scheme one might apply there and subject to lease considerations). It is noteworthy that Massive MIMO is implementable at both the client and POP/Small Cell ends of a LOS link due to the constrained geometries required (1/2 wavelength spacing, element-to-element).





Figure 26 – Lund University Massive MIMO Array (3.7 GHz)

Microstrip Patch Antenna Calculator

Pasternack's **Microstrip Patch Antenna Calculator** determines the length and width (in millimeters) of a rectangular patch antenna.

Dielectric Constant	4.4	
Dielectric Height:	.062	Inches
Operation Frequency:	3.7	GHz
<input type="button" value="Calculate"/>		



Result:

Width: 24.66 mm
Length: 18.87 mm

Figure 27 – Example of a Single Microstrip Patch Antenna Element @ 3.7 GHz – as Would be Represented in an Array Similar to the Lund Above (courtesy: Pasternack Website)

1.8 Hardening the Delivery

To expand more on the problem of millimeter wave transmission susceptibility to atmospheric absorption or scattering -- the availability past 2 nines (99%) usually implies redundant, orthogonal signaling paths and persistent monitoring of link quality. In the case of a FWA solution, 3.5 GHz and nLOS smart antennas provide both transmit power reach and backup link azimuth tuning which facilitate near-instantaneous switching of the connectivity path when signaled via meshing protocols to the nearest local peer (peer piggyback or peer repeating) in the case of primary loss-of-link. An example of this 3.5 GHz meshing is shown below in Figure 28. This is a potential elegant architecture for sub 6 GHz nLOS solutions.

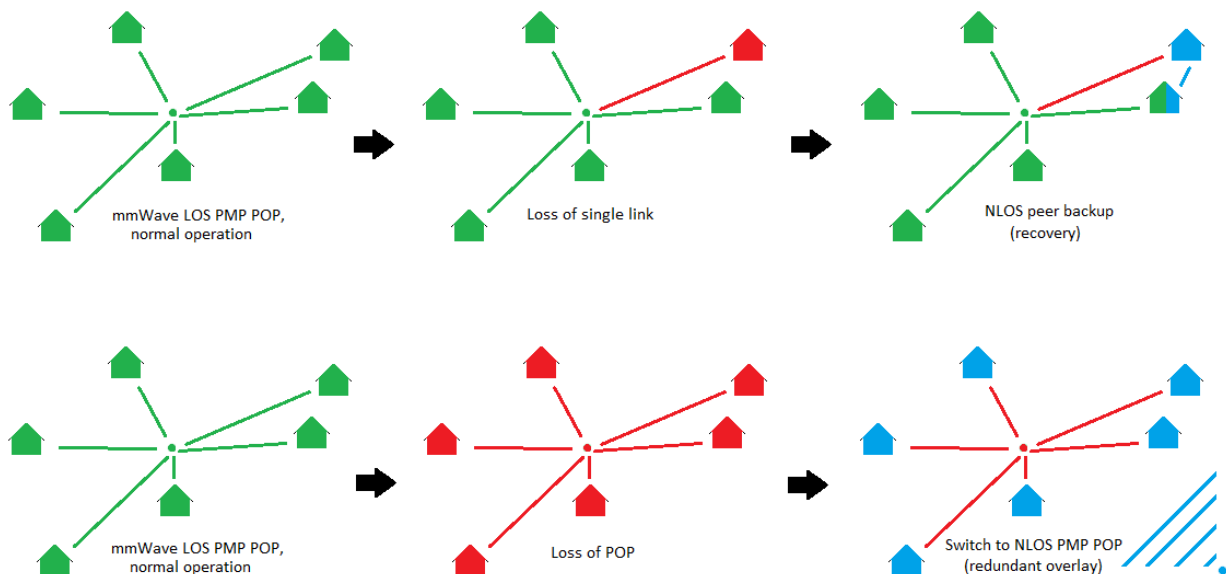


Figure 28 – Example of a 3.5 GHz Backup Signalling Mesh

LOS delivery redundancy is a much more difficult and expensive proposition, however. While it may be possible to establish dual LOS apertures (to separate monopoles) and even allocate emergency front- and backhaul POP/Small Cell wireline resources, redundant antenna solutions at the client sites would be required to affect a purely LOS backup scheme — and this presumes that a convenient LOS aperture to, and a serviceable radius from, a standby LOS base station exists.

A much more sensible scheme — perhaps an evolutionary goal past NLOS delivery at 3.5 GHz — would be to backstop the broad if less reliable LOS downlink with an NLOS “emergency” overlay which advantages itself of the scheme described above for NLOS delivery and provides a reliable uplink (pruning the link failure analysis tree in the process). Broad failure across multiple LOS downlinks in a single POP service group — as would happen for atmospheric interference — would of course tax the backup downlink bitrates; the presumption is that parametric performance reduction is better than complete loss of link.

From a purely qualitative analysis, the availability numbers for LOS wireless delivery versus nLOS should be comparatively worse, given a) its higher susceptibility to service interruption in the first place, b) the possibility that a redundant LOS signaling path is either unavailable or of compromised performance, and c) the potential for thrash or hang in the management of separate frontends (and hence, loss of link) under what may be a common propagation path impairment (rain or fog, for example). Concerns (b) and (c) are obviously mitigated in the case of NLOS downlink redundancy — with the performance hit as noted above.

1.9 HFC Wireless Extension: the New Trunks

Tapping into the HFC plant to supply wireless POPs amounts to an organic extension to the existing network (see below):




Figure 29 – Straightforward Tap of Fiber Trunk for Single Wireless Service Group

Previous analysis has evaluated options for enhancing the HFC plant capacity by pushing fiber deeper into the network [Ulm 16]. The analysis looked at a small number of fiber node service groups in one fiber service group.

In a full-service group overlay the 5G may be deployed to enhance or offload services from the HFC plant for subscribers. Depending on the geography of the service group, this may be a lower cost alternative to the more traditional methods of extending the capacity of the plant - i.e. node splits, fiber deep, etc. However, it should be noted that it is still likely that some costs (in addition to tower construction) will be incurred in the fiber distribution system to support a full-service group overlay.

Figure 29 depicts the same node from [Ulm 16], overlaying with an 800m radius small cell transmitter. Placing the cell at the existing fiber node, the CBRS signal does not cover the entire service group area. In the figure, the tower is placed more central to the service area, approximately 800m or 2600 feet from the fiber node requiring new infrastructure development.

This diagram is intended only as a hypothetical example of how a converged network may look on an existing node. Note that the node diagram has been updated to reflect the geographic distances between elements, however, it still may not reflect the directional depiction accurately.

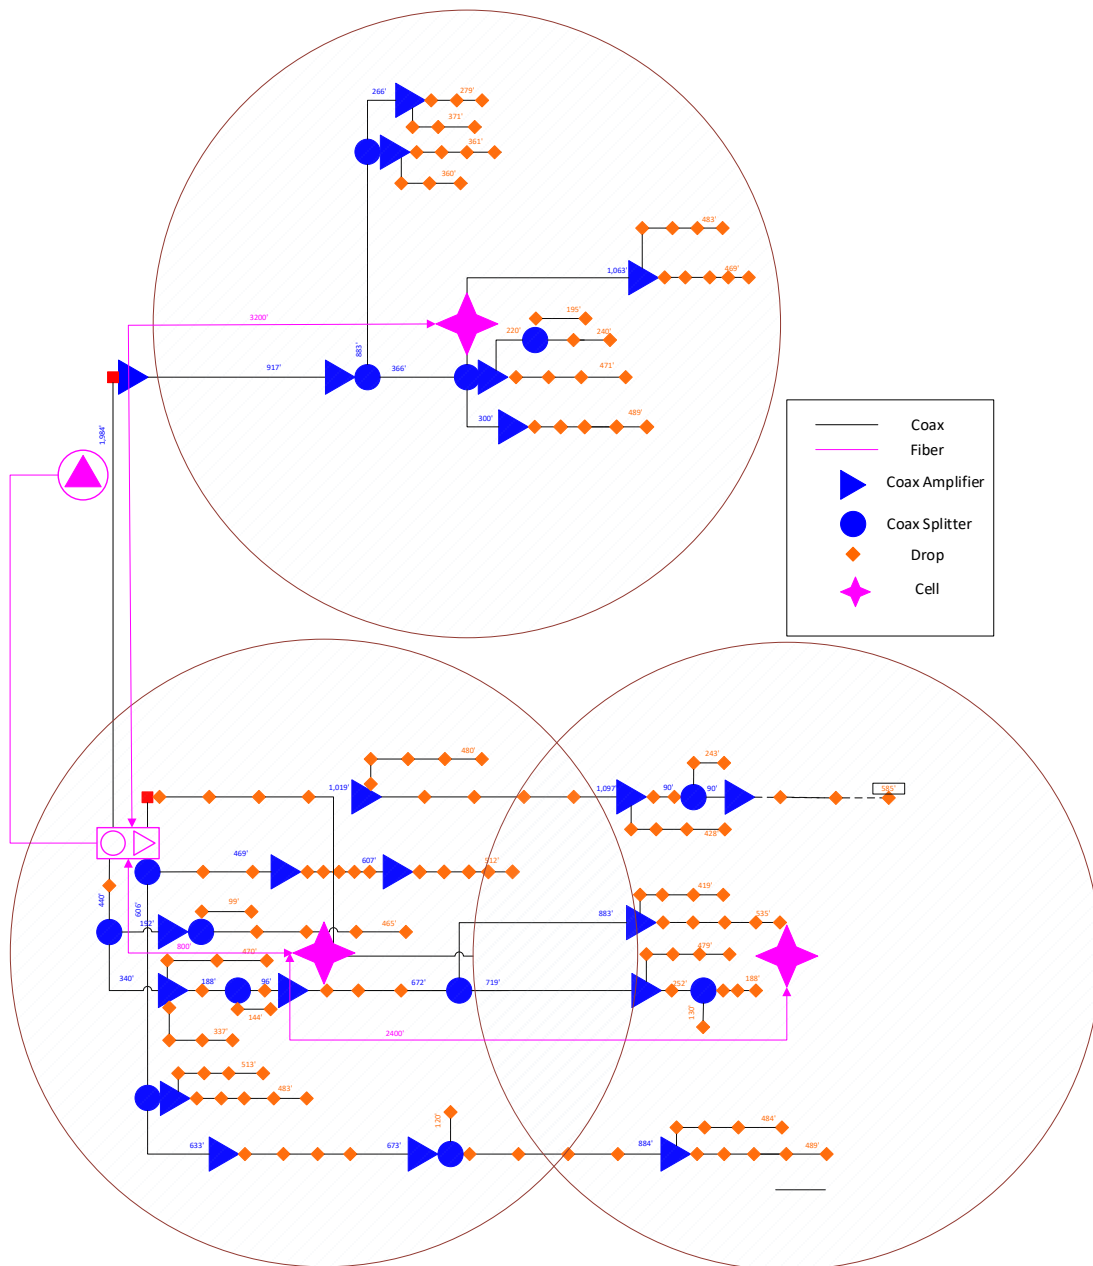


Figure 30 – Feed to Multiple POPs in a Large Overlaid Wireless Service Group

Figure 30 shows the same service group with 200m small-cell towers. The location of the transmitters was chosen to cover the service area with the fewest small cell towers. Alternatively, the location of the small cell transmitters could be placed to allow the reuse of the existing buried or aerial rights within the existing plant as much as possible. Note that in this situation, it may require a larger number of small-cell transmitters to cover the overall node service group, thus trade-off analysis should be done to determine the overall lowest cost configuration. Considerations may include whether the plant modifications will require new aerial or underground distribution feeds and availability of site locations.

2 Cost Considerations

Finding a common reference for the infrastructure investments implied in any of the options listed is difficult due to differentials in regional labor rate, site topography, client population density, relative accessibility of client and tower end points, antennae complexity, back- and front-haul piping to the POP, lease costs, and availability of shared BS infrastructure (like power and tower), the nature and type of the residential CPE and HN interface, and the complexity of installing and setting up the link. The following represents average estimates based upon input from industry professionals, publicly available cost studies and extrapolated CPE costs founded on similar complexity cable devices.

2.1 FTTH Costs

The presumptions applied to, and classifications for service enumerated in, CSMG's 2009 FTTH Deployment Assessment for Corning, have been examined and found no objectors among MSO senior staff solicited for commentary. Of the three company experiences referenced in that paper, a spread of ~35% in cost for connecting homes in "dense" population territories were noted. Applying an aggregated 3%/yr Cost of Living Adjustment (COLA) to those 2009 estimates seems advisable. In that case, for "dense" portions of the client population (~880 households/mi²), the average cost to pass and connect with fiber amounted to \$1,153 (Figure 31 - \$1,460 in adjusted 2017 dollars). When this density fell to ~175 households/mi², there was a cost premium of 40% — which applied to our adjusted for cost of living number above — implies \$2,050 per household. Allowing the client population density to fall to ~72 households/mi² imparts a 71% premium to the "dense" connectivity cost, yielding a "pass and connect" charge of \$2,499 per household.




Figure 31 – FTTH Cost in 2017 to Pass and Connect Per Household

Note that the rapidly rising cost to enfranchise ever-more-rural customers motivated the study to recommend against ever enfranchising the last 20% or so of the customer base. (In fact, the CSMG paper targets a goal of 41.5% of potential clients.) It is noteworthy that the FCC estimate for 100% connectivity in 2009 raised the connection tariff to an *average* of \$3,084 (estimated \$3,907 in current dollars) — so the weighted contribution of that final 20 percentiles is indeed staggering. Figure 32 captures the CAPEX and OPEX considerations.

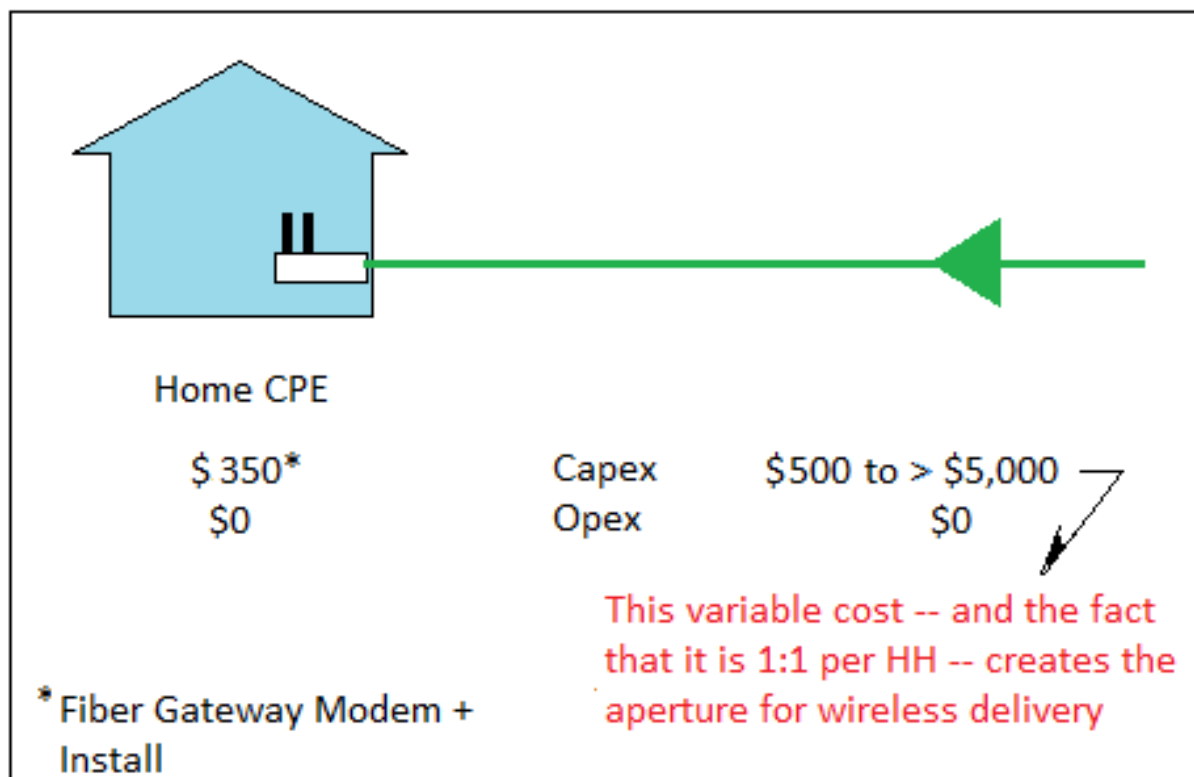


Figure 32 – FTTH Connection Costs at the HH

The upshot of this is that a significant percentage of US data customers could find themselves beyond the reach of high capacity wireline data services and seek an option to connect via Wireless services.

2.2 3.5 GHz CBRS Costs

As mentioned, NLOS 3.5 GHz does not require tall masts to facilitate antenna mounts but can settle for opportunistic 2nd-story types of elevations within 800 meters of its client base. In this way, it is extremely like legacy WLL strand and pedestal plays. Home CPE tends to look very much like existing gateway equipment with a different WAN attachment and the base station consistent with strand product (though with a more sophisticated, massive MIMO antenna array). Costs are anticipated to run as follow in Figure 33.

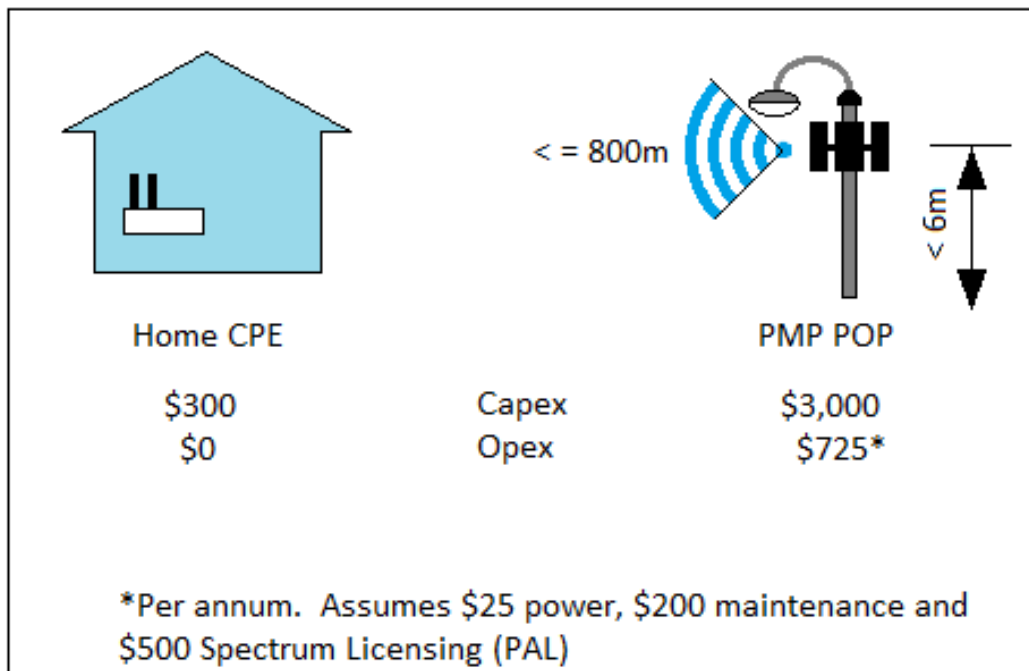


Figure 33 – NLOS/nLOS Sub-6 GHz Infrastructure Costs

Note that the POP CAPEX presumes \$2K in mast costs (hugely variable and depends upon availability of repurposed, available parasitic mounting space), \$500 for POP/Small Cell hardware and another \$500 in network connection costs. It is important to note that the POP costs may be amortized over the client population targeted by the service group (i.e., if the wireless service group population is 50 clients, the \$3,725 annual cost is ~ \$75/client).

2.3 Millimeter Wave costs

LOS wireless infrastructure carries a higher financial ante than NLOS due to higher complexity both on the CPE and BS sides. Additionally, where self-installation is a reasonable consideration for NLOS with its in-home CPE, the fact that LOS systems require an ORU (outdoor receiver unit) and potential placement optimization of this unit, implies that some budgetary citation for installation labor applies (a \$100 bookmark was placed for this). The balance of the \$650 client side costs was assigned to \$300 for the ORU, \$200 for the internal HN router, and \$50 for the interconnecting cabling and grounding aspects.

On the POP CAPEX front, the monopole was tagged with a \$20K cost, the Small Cell and antenna together cost \$2K, and network infrastructure connectivity was set at \$500. The summary looks as follows in Figure 34.

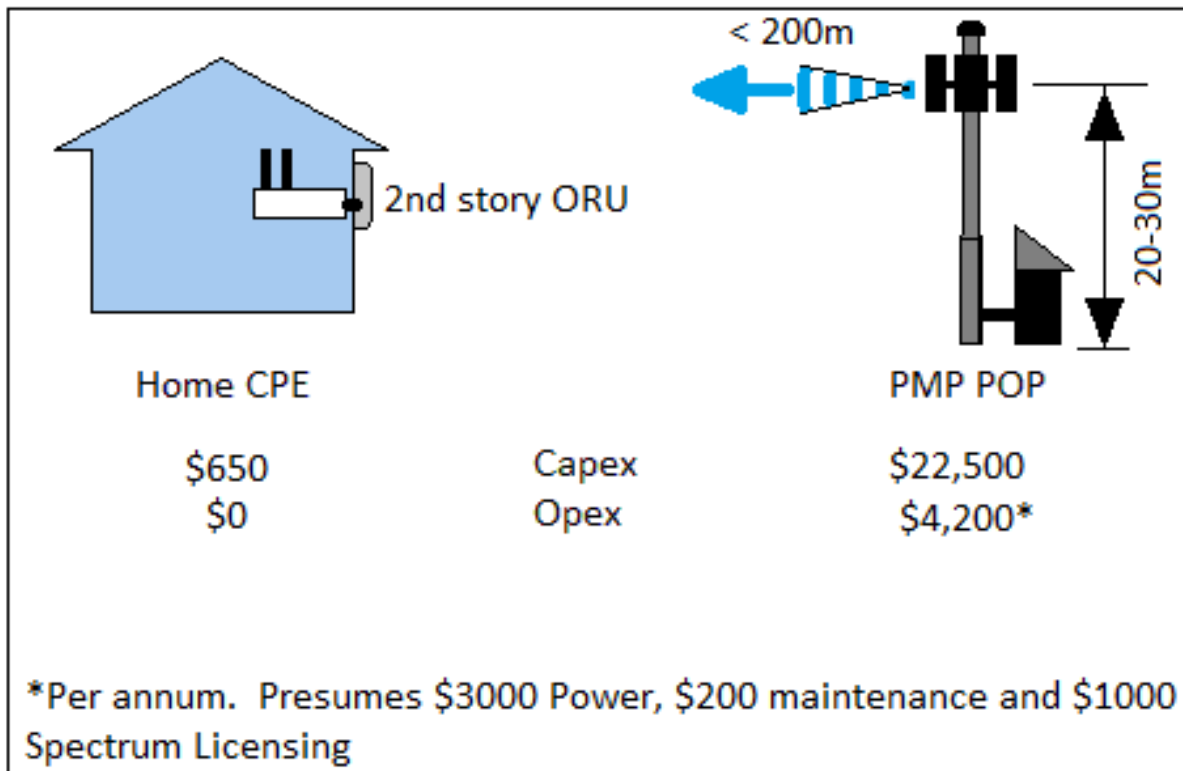


Figure 34 – LOS Infrastructure Costs

As with the NLOS case, amortization of the BS mast, antenna, and electronics may be performed over the client population. The critical differential here is that the service radius for LOS is only ~ 200 meters (which may imply only a couple of dozen clients, or less — and the initial investment is an order of magnitude higher).

3 Intangible Allowances

3.1 Aesthetics

Outdoor wireless AP placement has the potential to be extremely antagonistic from the aesthetic perspective. LOS considerations which must account for potential interference from foliage and interdicting land masses mandate an almost guaranteed level of intrusive environmental presence. Where nLOS APs can be situated in lower-height, disguised locations (and advantage themselves of parasitic placement on street furniture like streetlamps or signage in cases where up to 2nd story height is possible), LOS terminals, for both physical aperture and radiated power considerations, typically are placed well above ground level (to heights of 80 feet and beyond) and intrude on the landscape's skyline. In groomed communities subject to HOA governance on utility visual signature, there are tangible costs associated with potential litigation to obtain the necessary accommodations and intangible costs in public opinion regarding what many communities view as a blight on property values'

In these regards, there is no question that in-ground network connectivity trumps wireless approaches.

3.2 nLOS (3.5 GHz CBRS) Congestion

As the 3.5 GHz band offers an experimentation-friendly 3-tier licensing arrangement which actively promotes low-cost application testing via the General Authorized Access (GAA) entry level, it might behoove cable operators exploiting the band — despite the temptation to surf a GAA tier on accessibility -- to adopt a Priority Access License (PAL), such that explicit bandwidth requirements for geographic areas may be protected from a competing service co-option. And while the GAA tier advertises access to a marginally larger spectrum than PAL (80 MHz versus 70 MHz), there are no air time guarantees relative to other GAA pretenders in the geographic tract such that a QoE can be inferred. The short story then is that, to avoid situations where competing service leverage of the 3.5 GHz CBRS band promotes the accommodation of other services, MSOs' intended use of the band for a geographic area warrants investment in the protection of PAL licenses. (License auctions are expected beginning in 2018).

Conclusions

The implications of providing a future-proofed wireless bitrate capability to all subscribers beyond the reach of wireline in a cable system requires the analysis of wireless delivery options which include LOS, nLOS and NLOS systems — each of which comprises a mixed bag of capability and compromise. The broader bandwidth of millimeter LOS delivery, with its promise of massive MIMO antenna structures on both base station and client endpoints, unfortunately burdens itself with compromises involving client-side signal recovery costs, short signal throw, aesthetic challenges and perhaps too-easily non-deterministic link quality. nLOS and NLOS sub-6 GHz systems can be made to overcome these challenges. However, the available bandwidth puts considerable pressure on massive MIMO and signal processing upgrades on the base station side to create the scalar benefits which effectively multiply spectral efficiency to levels necessary to anticipate user bitrate consumption a mere 4-5 years in the future. The relentless bitrate consumption growth defined in Cloonan's Curve suggest that, ultimately, the facility of sub-6 GHz NLOS will be associated with a redundancy role for more LOS-based delivery — or perhaps in an ad hoc augmentation role for temporal housing arrangements.

There is also an argument which might bear examination (despite the hefty cost of the required redundant infrastructure) which proposes that a hybrid fiber-coax-wireless (HFCW) scheme might be a useful offloading solution in mixed-use cases where separating a few heavy consumers via differential delivery PHY wireless might buy service phase-in time for major upgrades to the legacy wireline business. One could also bind delivery on a flexible basis across both wireless and wireline PHYs and orchestrate a closed-loop, QoE-deterministic multipath delivery scheme which senses delivery impairment per PHY and adjusts link exploits accordingly.

As cable operators move Fiber Deeper going to an all passive coax network, the ability to deliver multiple Gbps of capacity to a single home, seems an easier path than building out a FWA millimeter wave architecture. However, given that 5G POP/Small Cells require wired backhaul, the potential for the MSO to leverage its network for mobile 5G seems to be a more complimentary investment. In discussions with MSOs, who are also MNOs, they struggle now to see a FWA solution to deep residential deployments. They see some potential to use their network to potentially lead out to target MDUs served predominantly by their competitors, and often see value in pulling Fiber. However, they do see the value of adding 5G and CBRS POPS to their HFC and growing Fiber Networks for outside mobility applications.

For MNOs, those that don't own wired broadband networks, the use of FWA is an opportunity to cherry pick areas for a Fixed Wireless overbuild. We have seen some Wireless ISPs already offer millimeter

wave broadband delivery services targeting dense areas with only one incumbent, areas where consumers are deprived of choice of broadband provider. The investment scale which nourishes those shared wireless technical advances applicable to both unlicensed and dynamically licensed space for MSOs (cable, telco, and MNO alike) means that applications of FWA will emerge as we move to mobility on 5G systems. The economics and the size of the optimum cell is still under debate. What is also clear is that the easier direction for FWA is dense MDU environments is targeting a single wireless connection to the outside and using other solutions internally, like Ethernet and Wi-Fi. The Residential 5G deployments will only emerge driven by the rise of mobile 5G devices which will happen in 2021 at scale and will then see the 5G small cell deploy in ever decreasingly small cell sizes.

And lastly, there is still a lot of activity around trying to leverage sub 6 GHz frequencies into the 5G requirements space. 3GPP have added support in New Radio to support sub 6 GHz and even sub 1 GHz to create the overlay capability for millimeter non-determinism to fall back on other transmission frequencies and to allow NB-IOT to run on lower sub GHz frequencies. Additionally, applying Element Arrays to work with sub 6 GHz frequencies offers up potential for paired antenna and spatial streams to reuse spectrum and use spectrum at higher bits/Hz.

We have come a long way in the drive to 5G — but as the saying goes — there is still a long way to go.

Abbreviations

AP	access point
bps	bits per second
BS	Base station
CAGR	Compound annual growth rate
CAPEX	Capital Expenditures
CBRS	Citizens Broadband Radio Service
CBSD	Citizens Broadband Radio Service Device
FEC	forward error correction
FTTH	Fiber-to-the-home
FWA	Fixed Wireless Access
Gbps	Gigabits per second
GHz	Gigahertz
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
LOS	Line-of-sight
nLOS	Near-line-of-sight
NLOS	Non-line-of-sight
OBE	Overcome by events
OPEX	Operating Expenditures
ORU	Outdoor Receiver Unit
QoE	Quality-of-Experience
P2P	Peer-to-peer
PMP	Point-to-Multipoint
POP	Point of Presence

PTP	Point-to-Point
SCTE	Society of Cable Telecommunications Engineers
SG	Service Group
SINR	Signal-to-Ingress and Noise-Ratio
TCO	Total Cost of Ownership
WLL	Wireless Local Loop

Bibliography & References

5G Channel Model for bands up to 100 GHz (2.3 October 2016), 5GCMSIG White Paper

5G NR mmWave, Ozge Koyman, Qualcomm Research

CBRS: New Shared Spectrum Enables Flexible Indoor and Outdoor Mobile Solutions and New Business Models, Kyung Mun, CBRS White Paper, CBRS Alliance, March 2017

Cell Towers and Aesthetics: Blight on the Neighborhood or Sign of the Times?, Paul J. Weinberg, Zoning and Planning Law Report

FCC 15-47, REPORT AND ORDER AND SECOND FURTHER NOTICE OF PROPOSED RULEMAKING, April 21, 2015 (3.5 GHz CBRS band creation)

FTTH Deployment Assessment, Cartesian (formerly CSMG), 2009 (used by permission)

Future Directions for Fiber Deep HFC Deployments, John Ulm and Zoran Maricevic, SCTE•ISBE Expo '16

LuMaMi – A flexible testbed for massive MIMO, Joao Vieira, Steffen Malkowsky, Karl Nieman, Zachary Miers, Nikhil Kundargi, Liang Liu, Ian Wong, Viktor Owall, Ove Edfors, Fredrik Tufvesson, Lund University

Making 5G NR a Reality, Qualcomm Technologies, Inc, Sept 2016

Massive MIMO for Next Generation Wireless Systems, Erik G. Larsson, ISY, Linköping University, Sweden, Ove Edfors, Lund University, Sweden, Fredrik Tufvesson, Lund University, Sweden, Thomas L. Marzetta, Bell Labs, Alcatel-Lucent, USA, arXiv: 1304.6690v3 [cs.IT] 12 Jan 2014

Massive MIMO: Ten Myths and One Critical Question, Emil Björnson, Erik G. Larsson, and Thomas L. Marzetta, arXiv:1503.06854v2 [cs.IT] 18 Aug 2015

Millimeter Wave System Performance Characterization for 5G Data Access, Shirish Nagaraj, Lea Castel, Tommaso Balercia, Bishwarup Mondal, Jong-Kae Fwu, Communications and Devices Group, Intel

Shared Spectrum Market Opportunity for Cable MSOs, Mark Lowenstein, Mobile Ecosystem, April 2017

Fixed Mobile Convergence in the Transition to 5G

A Technical Paper prepared for SCTE•ISBE by

Glenn Laxdal
Ericsson
6300 Legacy Dr. Plano, TX 75002
214 566 2045
Glenn.Laxdal@ericsson.com

Abstract

2017 is witnessing big changes in mobile and fixed industries, from unlimited plans on LTE, to major MSO acquiring 600MHz in the latest auction, to 5G millimeter wave trials, to the heated discussion on adding the 3.7GHz – 4.2GHz adjacent to CBRS band to the trials. Introduction of 5G technologies like higher spectrum bands, wider carriers and massive MIMO creates an opportunity for the convergence of fixed and mobile services both for traditional telecom operators and MSOs.

Both site solutions for small cell grid and high capacity backhaul are of the biggest challenges for Gigabit wireless to the home. MSO have some of the answers to these problems via the extensive Fiber network in residential area to the strand mount right of the way with power and backhaul ready.

Join Anders Svensson, Principal Solution Manager, 5G, Strategic network evolution at Ericsson North America, to learn about technology and deployment strategy of the mix of technologies for Fixed Wireless Access in the convergence of fixed and mobile services.

This paper gives an overview of 5G requirements and capabilities, how 5G technologies may impact the industries and enable convergence of services. The paper covers selected finding around deployment of FWA for both CBRS and 5G at mmw, and the opportunities and challenges with different spectrums.

Overview of the 5G Network

The overall aim of 5G is to provide ubiquitous connectivity with optimal characteristics for any kind of device and any kind of application that may benefit from being connected. It will address high traffic growth and increasing demand for high-bandwidth connectivity. It will also support massive numbers of connected devices with long battery life and low power consumption and meet the real-time, low latency high-reliability communication needs of mission-critical applications. The 5G system is a use case driven network where one network can support multiple use cases via optimized network slicing.

The standardization of 5G is important to build the global eco-system that provide economic of scale and allows the technologies to address the broader use in a Networked Society. The standardization of 5G is aiming to provide capabilities defines by ITU as part of the IMT-2020. The standardization of the implementation and use of technologies are done in 3GPP. The schedule for this standardization has been divided in two phases, with Phase 1 targeting ready June 2018 and Phase 2 target to be ready in second half 2019. Phase 1 is mainly focusing on high capacity/throughput use cases and ultra-reliable / low latency communication(URLLC) use cases, while Phase 2 target the full compliance with the IMT-2020, including massive IoT use cases.

Parallel to the standardization 5G is gaining momentum. It started with trials in 2016 to show case and validate 5G technologies and use cases. Trials as expected to continue through 2018 with early deployments expected in 2018 and 2019. Most interest has been shown in US, China, Korea, Japan. The massive growth of mobile data traffic, primary driven by video consumption, drives the interest in 5G for enhanced Mobile broadband “eMBB” which is gaining tractions in almost all regions (NA, APAC, EMEA). The high capacity provided by 5G also make it attractive for the Fixed Wireless Access(FWA) use case. FWA as an alternative to fiber residential and enterprise is the leading use case in US, while in APC and Europe, a special focus has been seen on the industrial use cases. Examples of industrial use cases are connected factories, deep communication in mines, connected ports and logistics.

The specification of 5G will include the development of a new flexible air interface, NR, which will be directed to extreme mobile broadband deployments. NR will also target high-bandwidth and high-traffic-usage scenarios, as well as new scenarios that involve mission-critical and real-time communications with extreme requirements in terms of latency and reliability. Further capacities deployed in 5G NR will put a broader emphasis on rich fiber networks capable of backhauling this ever-growing volume of traffic but will also allow new players to enter the market to provide mobile and fixed wireless services that leverage their regional or nationwide fixed access network. This would include HFC and fiber networks currently deployed in the MSO community.

LTE is expected to evolve in a way that recognizes its role in providing excellent coverage for mobile users. 5G networks will incorporate LTE access along with new air interface, NR, in a transparent manner toward both the service layer and users. The evolution of LTE to a point where it is a full member of the 5G family of air interfaces is essential, especially since initial deployment of new air interfaces may not operate in the same bands. The 5G network will enable dual-connectivity between LTE operating within bands below 6GHz and the NR air interface in bands within the range 6GHz to 100GHz. NR should also allow for user-plane aggregation, i.e. joint delivery of data via LTE and NR component carriers.

Around 2020, much of the available wireless coverage will continue to be provided by LTE, and it is important that operators with deployed 4G networks can transition some – or all – of their spectrum to newer wireless access technologies. For operators with limited spectrum resources, the possibility of introducing 5G capabilities in an interoperable way – thereby allowing legacy devices to continue to be served on a compatible carrier – is highly beneficial and, in some cases, even vital.

By 2022, it is expected to have more than 500M 5G connections globally. In US the forecast is that 25% of subscription will be 5G enabled according to the “Ericsson Mobility Report”.

5G Requirements and Use Cases

5G is considered to support many use cases beyond the traditional 4G mobile broadband, and thus the requirements on 5G networks are reflecting a wide range of capabilities to enable the many varieties of use cases.

1. 5G use cases

5G will provide wireless connectivity for a wide range of new applications and use cases, including wearables, smart homes, traffic safety/control, critical infrastructure, industry processes and very-high-speed media delivery. As a result, it will also accelerate the development of the Internet of Things.

EXAMPLE USE CASES FOR 5G EXTEND BEYOND 4G



Figure 1 – 5G Use Case examples

Figure above shows some examples use cases that is targeted with 5G. These represents applications that leverage either one or several of the capabilities that 5G will enable in addition to 4G networks that is widely deployed today. The timing of the use cases will be dependent of availability of the technologies as well as the consumer and industry adaption. In North America, the first use cases expected to be widely adopted are fixed wireless broadband, as seen on picture blow, and enhanced mobile broadband eMBB with Gigabit Mobile data.

5G TO THE HOME




Figure 2 – Fixed Wireless Access using 5G

Both these Use Cases has well established business models which allows for a faster adaption as soon as technology becomes available. Other use cases, which require new business models for operators and surrounding industries, are expected to have a slower adaption to reach mass market. However, for operators and MSO that are planning to deploy 5G, it is important to consider the future services when deploying networks.

2. 5G requirements

In order to enable connectivity for a very wide range of applications with new characteristics and requirements, the capabilities of 5G wireless access must extend far beyond those of previous generations of mobile communication.

These capabilities will include massive system capacity, very high data rates everywhere, very low latency, ultra-high reliability and availability, very low device cost and energy consumption, and energy-efficient networks.

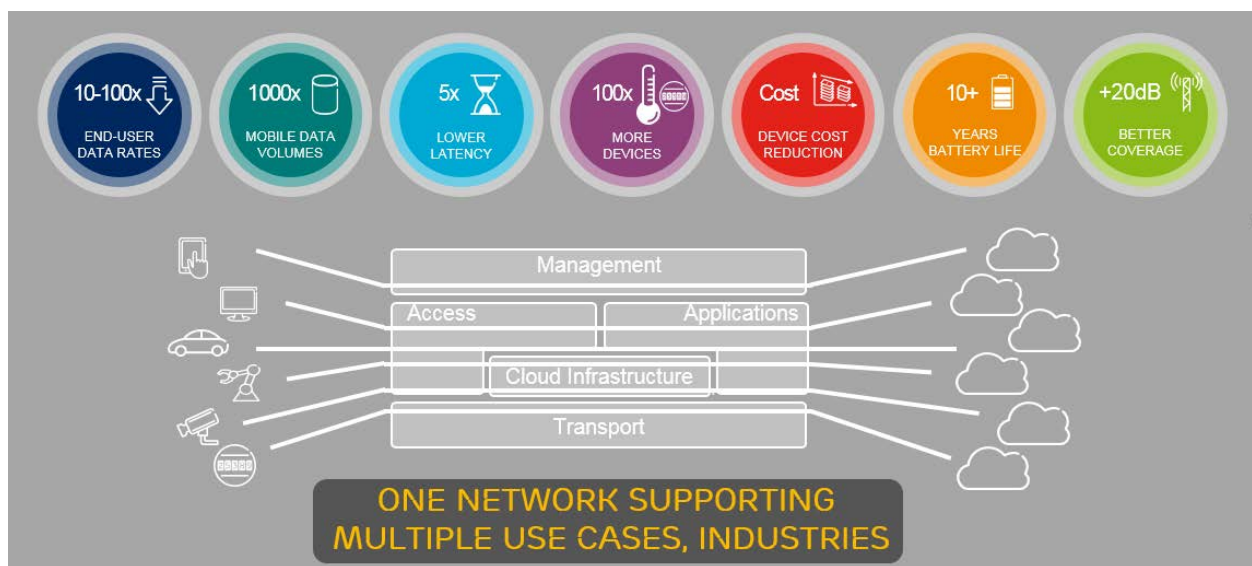


Figure 3 – 5G Capabilities

Figure above shows some of the targets for 5G relative the 4G system that is today widely used for mobile communication.

2.1. MASSIVE SYSTEM CAPACITY

The exponential increase in connected devices, such as the deployment of billions of wirelessly connected sensors, actuators and similar devices for massive machine connectivity, will place demands on the network to support new paradigms in device and connectivity management that do not compromise security. Each device will generate or consume very small amounts of data, to the extent that they will individually, or even jointly, have limited impact on the overall traffic volume. However, the sheer

number of connected devices seriously challenges the ability of the network to provision signaling and manage connections. We expect that a significant amount of traffic and connections will also come from within the household where fixed access networks will see an increase in devices accessing the network. Some of these devices may also collect data while outside the home while using a household internet connection as a connection to report non-real-time sensor data to centralized data collection cloud infrastructure.

2.2. VERY HIGH DATA RATES EVERYWHERE

Every generation of mobile communication has been associated with higher data rates compared with the previous generation. In the past, much of the focus has been on the peak data rate that can be supported by a fixed or wireless-access technology under ideal conditions. However, a more important capability is the data rate that can actually be provided under real-life conditions in different scenarios.

5G should support data rates exceeding 10Gbps in specific scenarios such as indoor and dense outdoor environments.

Data rates of several 100Mbps should generally be achievable in urban and suburban environments.

Data rates of at least 10Mbps should be accessible almost everywhere, including sparsely-populated rural areas in both developed and developing countries.

There are certainly further opportunities for “blended operators” to “fuse” services together to offer data buckets that are commonly shared between mobile and fixed access services.

2.3. VERY LOW LATENCY

Very low latency will be driven by the need to support new applications. Some envisioned 5G use cases, such as traffic safety, control of critical infrastructure, and industry processes, may require much lower latency compared with what is possible with the mobile-communication systems of today.

To support such latency-critical applications, 5G should allow for an application end-to-end latency of 1ms or less, although application-level framing requirements and codec limitations for media may lead to higher latencies in practice. Many services will distribute computational capacity and storage close to the air interface. This will create new capabilities for real-time communication and will allow ultra-high service reliability in a variety of scenarios, ranging from entertainment to industrial process control.

2.4. ULTRA-HIGH RELIABILITY AND AVAILABILITY

In addition to very low latency, 5G should also enable connectivity with ultra-high reliability and ultra-high availability. For critical services, such as control of critical infrastructure and traffic safety, connectivity with certain characteristics, such as a specific maximum latency, should not merely be ‘typically available.’ Rather, loss of connectivity and deviation from quality of service requirements must be extremely rare. For example, some industrial applications might need to guarantee successful packet delivery within 1 ms with a probability higher than 99.9999%.

2.5. VERY LOW DEVICE COST AND ENERGY CONSUMPTION

Low-cost, low-energy mobile devices have been a key market requirement since the early days of mobile communication. The reduction of energy consumption in set top boxes has also generated some voluntary

but effective power consumption best practices within the home. To enable the vision of billions of wirelessly connected sensors, actuators and similar devices, a further step has to be taken in terms of device cost and energy consumption. It should be possible for 5G devices to be available at very low cost and with a battery life of several years without recharging.

2.6. ENERGY-EFFICIENT NETWORKS

The increase in data consumption will result in an increased energy footprint from networks. While device energy consumption has always been prioritized, energy efficiency on the network side has recently emerged as an additional KPI. 5G must therefore consume significantly lower energy per delivered bit than current cellular networks. Much like the Energy 2020 initiative being pursued by the STCE, 5G will focus some significant efforts on how to improve mobile services without creating unsustainable energy footprints.

There are three main drivers for more energy efficient networks: First it is an important component in reducing operational cost, leading to lower total cost of ownership. Second, it enables off-grid network deployments that rely on medium-sized solar panels as power supplies, thereby enabling wireless connectivity to reach even the most remote areas with renewable energy sources. And third, it is essential to realizing operators' ambition of providing wireless access in a sustainable and more resource-efficient way.

The importance of these factors will increase further in the 5G era, and energy efficiency will therefore be an important requirement in the design of 5G wireless access. The SCTE energy 2020 initiative is also focused on many of these general goals as well.

2.7. SUBSCRIBER SPECIFIC SERVICES

New use cases with largely different business cases and business models will require service providers to be able to differentiate the services based on the application and Use Case. 5G will enable this with a functionality called Network Slicing. Network Slicing enables creation of virtual network that run independently on the same infrastructure with different characteristics with regards to peak speed, delays, reliability, device functionality, etc. The operators will use slicing to create new offering addressing the Use Cases in an efficient way.

5G Features

Beyond extending operation to higher frequencies, there are several other key technology components relevant for the evolution to 5G wireless access. These components include access/backhaul integration, device-to-device communication, flexible duplex, flexible spectrum usage, multi-antenna transmission, ultra-lean design, and user/control separation. In the following sections, brief descriptions of these features are introduced.

3. Radio Access Network features

3.1. FLEXIBLE DUPLEX

Frequency Division Duplex (FDD) has been the dominating duplex arrangement since the beginning of the mobile communication era. In the 5G era, FDD will remain the main duplex scheme for lower

frequency bands. However, for higher frequency bands – especially above 10GHz – targeting very dense deployments, Time Division Duplex (TDD) will play a more important role.

In very dense deployments with low-power nodes, the TDD-specific interference scenarios (direct base-station-to-base-station and device-to-device interference) will be similar to the ‘normal’ base-station-to-device and device-to-base-station interference that also occurs for FDD.

Furthermore, for the dynamic traffic variations expected in very dense deployments, the ability to dynamically assign transmission resources (time slots) to different transmission directions may allow more efficient utilization of the available spectrum.

To reach its full potential, 5G will therefore allow for very flexible and dynamic assignment of TDD transmission resources. This is in contrast to current TDD-based mobile technologies, including TD-LTE, for which there are restrictions on the downlink/uplink configurations, and for which there typically exist assumptions about the same configuration for neighbor cells and also between neighbor operators.

3.2. FLEXIBLE SPECTRUM USAGE

Since its inception, mobile communication has relied on spectrum licensed on a per-operator basis within a geographical area. This will remain the foundation for mobile communication in the 5G era, allowing operators to provide high-quality connectivity in a controlled-interference environment.

However, per-operator licensing of spectrum will be complemented by the possibility of sharing spectrum. Such sharing may be between a limited set of operators, or may occur in license-exempt scenarios. The Citizens Band Radio Service in the US in the 3.5GHz band and the 5GHz unlicensed spectrum are examples of managed and unlicensed sharing regimes respectively.

3.3. MULTI-ANTENNA TRANSMISSION

Multi-antenna transmission already plays an important role in current generations of mobile communication and will be even more central in the 5G era, due to the physical limitations of small antennas. Path loss between a transmitter and receiver does not change as a function of frequency, as long as the effective aperture of the transmitting and receiving antennas does not change. The antenna aperture does reduce in proportion to the square of the frequency, and that reduction can be compensated by the use of higher antenna directivity. The 5G radio will employ hundreds of antenna elements to increase antenna aperture beyond what may be possible with current cellular technology.

The use of high degree of beamforming also enable a more efficient use of Multi-user MIMO. The high directivity with beamforming decrease the level of interference in the part of the cell that is not used, which gives the opportunity to sending simultaneously to another user in that part of the cell creating low intra-cell interference. This is an important capacity enabler in 5G, increasing the sector capacity.

In addition, the transmitter and receiver will use beamforming (BF) to track one another and improve energy transfer over an instantaneously configured link. Beamforming will also improve the radio environment by limiting interference to small fractions of the entire space around a transmitter and likewise limiting the impact of interference on a receiver to infrequent stochastic events. The use of Beamforming will also be an important technology to extend coverage and to provide higher data rates over larger area.

3.4. ULTRA-LEAN DESIGN

Ultra-lean radio-access design is important to achieve high efficiency in 5G networks. The basic principle of ultra-lean design can be expressed as: minimize any transmissions not directly related to the delivery of user data. Such transmissions include signals for synchronization, network acquisition and channel estimation, as well as the broadcast of different types of system and control information.

Ultra-lean design is especially important for dense deployments with a large number of network nodes and highly variable traffic conditions. However, lean transmission is beneficial for all kinds of deployments, including macro deployments.

By enabling network nodes to enter low-energy states rapidly when there is no user-data transmission, ultra-lean design is an important component in delivering high network energy performance. Ultra-lean design will also enable higher achievable data rates by reducing interference from non-user-data-related transmissions.

3.5. USER/CONTROL SEPARATION

Another important design principle for 5G is to decouple user data and system control functionality. The latter includes the provisioning of system information; that is, the information and procedures needed for a device to access the system.

Such a decoupling will allow separate scaling of user-plane capacity and basic system control functionality. For example, user data may be delivered by a dense layer of access nodes, while system information is only provided via an overlaid macro layer on which a device also initially accesses the system.

User/control separation is also an important component for future radio-access deployments relying heavily on Beamforming for user data delivery. Combining ultra-lean design with a logical separation of user-plane data delivery and basic system connectivity functionality will enable a much higher degree of device-centric network optimization of the active radio links in the network.

5G Spectrum

The evolution of 4G/LTE as developed to exploit a large number of spectrum to handle the massive growth of mobile data over the last 10 year. With the forecast that this traffic growth will continue, driven by a larger usage of video and other digital technologies, more capacity will be needed. To support the increased traffic capacity and to enable the transmission bandwidths needed to support very high data rates, 5G will extend the range of frequencies used for mobile communication. This includes new spectrum below 6GHz, as well as spectrum in higher frequency bands including millimeter wave spectrum. The 5G standard is both specifying how to leverage these new spectrum band as well as how spectrum can be combined efficiently in a system to provide the diverse 5G services in when deployed.

1. Higher frequency band and mmWave

Specific candidate spectrum for mobile communication in higher frequency bands is yet to be identified by the ITU-R or by individual regulatory bodies. The World Radio Conference (WRC)-15 discussions have resulted in an agreement to include an agenda item for IMT-2020, the designated ITU-R qualifier for 5G, in WRC-19. The conference also reached agreement on a set of bands that will be studied for 5G, with direct applicability to NR. Many of the proposed bands are in the millimeter wave (mmw) above 24 GHz.

In the US, the FCC has allocated 4 primary 5G bands; **28GHz** (27.5GHz-28.35GHz), **37GHz** (37GHz-38.6GHz), **39GHz** (38.6GHz-40GHz) licensed bands and **64-70GHz** as unlicensed band. The 28GHz spectrum is re-farmed to 2 licenses each at 425MHz on county area, the 39GHz will be re-farmed to 7x 200MHz channels on PSA area, and upper 1MHz of the 37GHz will be new band 5x 200MHz auctioned at PSA level. The lower 600MHz part of the 37GHz is under consideration for shared spectrum scheme. There is a proposal for adding more spectrum in 24GHz, 32GHz, 42GHz, 48GHz, 51GHz, 70GHz and 80GHz.

The race to secure these new band allocated by FCC has already started. Verizon has acquired sizable amount of the 28GHz via its acquisitions of XO followed by acquisition of StraightPath. The latter also provides Verizon with good holding in the 39GHz. AT&T acquired Fiber Tower and secured licenses in 39GHz. T-Mobile had already ownership of some 39GHz and 28GHz holding from MetroPCS acquisition. The spectrum that does not have an ownership today is planned to be sold by FCC.

2. Medium high frequency band

The capacity needs of the mobile industry will continue to be served by licensed spectrum, although novel sharing arrangements for spectrum will become progressively more important as restricted opportunities for new spectrum start to impact incumbent services such as satellite communication and radio location. Two examples of sharing arrangements include LSA planned in Europe for the 2.3GHz band and the Citizens Band Radio Service (CBRS) for 3.5GHz in the US. CBRS is 150MHz of spectrum shared between the incumbent FSS and coastal radar as high priority incumbents, then 10MHz channels up to 7 channels is priority access licenses that would be licensed by FCC and will have second level of protection and priority after the incumbent, and finally the General access authorization (GAA) which spans the entire 150MHz with the lowest priority. Environmental sensing system (ECS) that detects the activities of the incumbents, and feeds into the Shared Access system SAS database and control system allow the control and assignment of the shared spectrum among the different tiers. Initially the FCC rules is to have the PAL licensed renewed once for 3 years and on a census track level. Recent proposals are suggesting increase the licensing terms to 10 years and licensing areas to PSA from census tracks.

In the US, cellular and FWA industry is contemplating having 3.7-4.2GHz as NR band. This band has not yet been identified as a 5G band. It is currently used as uplink for the c-band satellite services, but FCC have issued a Notice of Inquiry(NOI) “Expanding Flexible Use in Mid-Band Spectrum Between 3.7 and 24 GHz” which cover the 3.7 – 4.2 GHz. The NOI also explore the used of 5.925-6.425 GHz and 6.425-7.125 GHz, and is due for comment in October 2017

Propagation Characteristics

The propagation channel characteristics is dependent on the frequency band. Figure below summarize the difference between the typical cellular band deployed today in 2.1 GHz and mid-bands like CBRS and mmWave bands.

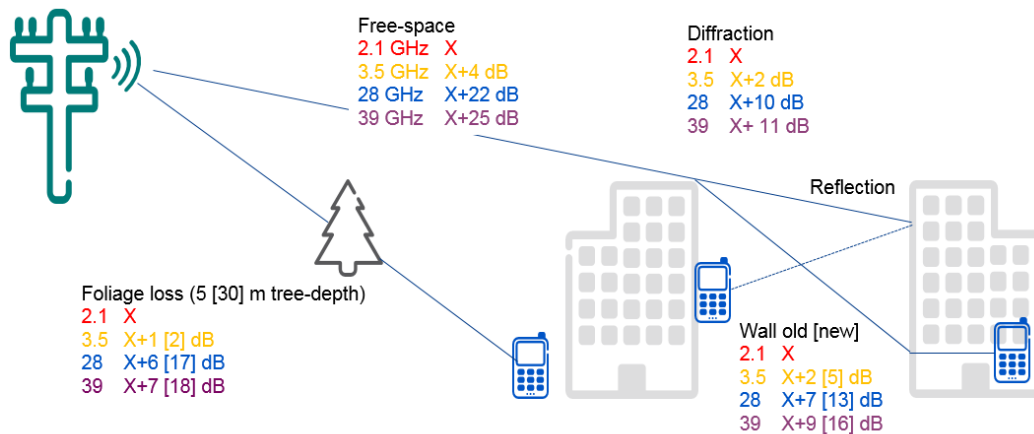


Figure 4 – Path loss comparison between spectrum bands

There is slight difference between the cellular and CBRS, for example it is less than 4dB on free space propagation loss, 2 dB foliage loss and building penetration loss old older buildings. For newer energy efficient building the penetration is even higher. On the other hand, free space propagation loss between mmw and cellular is ranging from 22-25dB, foliage between 6 and 18dB, depending on the tree depth (size) and frequency (28GHz vs 39GHz), diffraction loss is around 10dB more loss, and building penetration loss is more than B1 by 7-16dB depending on type of wall and frequency. It should be noted that reflection losses are rather independent on frequency.

Fixed Wireless Access Performance

The large improvement of capacity access capacity with 5G technologies leveraging much larger channels available in high bands and improving the sector capacity using massive MIMO makes it attractive the address the broadband market to residential and enterprise customers (i.e. Fixed Wireless Access service).

While the wider channels, and a carrier aggregation between them, can be used in any spectrum, the availability of such spectrum is more common in higher mmw bands. The last year there has been significant interest in using these bands, recently made available by FCC, for fixed wireless. Due to mmw limited reach and sensitivity to penetration losses the interest has focused on the possible coverage that can be achieved when offering a FWA service using these bands. Studies shows that the performance that can be expected is dependent on decisions, and viability, to where to deploy both base station antennas and the customer premises equipment (CPE) antennas. The performance of Fixed Wireless Access varies with the deployment scenarios and the spectrum used.

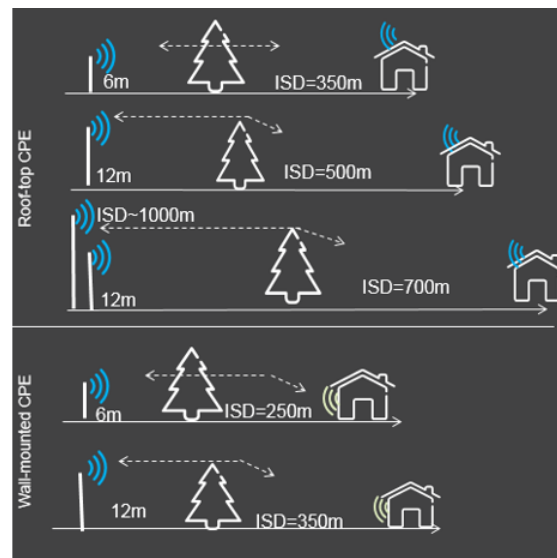



Figure 5 – Performance impacts in mmWave deployment scenarios

The left diagram above shows a result of 28 GHz simulation with different location of the base station antenna and the impact on the SNR received by the CPEs distributed at the houses in a typical single family resident area with mix of single and dual story houses. The graph shows a significant improvement of the cell edge SNR with higher mounted antennas, which translates to large improvement of throughput at the cell edge, which represents the throughput that can be offered to the households in the area.

In addition to the improvement with higher elevated base station antennas, the coverage area is also sensitive to the location of the CPE antenna. The right picture in figure shows the impact on Inter-Site Distance dependent on the different location. All scenarios provide the same cell edge throughput for the user with the most challenges locations, equal to 100 Mbps per 100 MHz of spectrum. While the simulation is done with homogenous deployed CPE, it should be noted that in real scenarios may be a mixed CPE deployment, with roof-top CPE deployed in locations(houses) at challenging RF condition and wall-mounted where RF is less challenging.

For operator and MSO considering deployment in FWA service with mmWave should as part of developing and evaluating the service consider deployment of both the Base Station antennas and CPE antenna. The overall high capacity in mmWave, but limited coverage makes this spectrum most attractive for deployment in urban and suburban single- and multi-family residential area. The high peak speed enabled with the wide channels also make mmWave attractive for enterprise offering.

The other spectrum that has been receiving attraction for FWA service using 5G technologies is the mid-band spectrum 3-4 GHz. While they have potential significant more spectrum than current cellular spectrum, it is expected to be smaller than mmWave. The focus in this spectrum has been more concentrated to the massive MIMO capability with high level of beamforming and MU-MIMO capabilities. These technologies are important to be able to provide the capacity need for the FWA service.

As with mmWave, the important aspect to consider is the cell edge throughput, but in the many cases the mid-band will be primarily capacity limited. Graph below shows the cell edge throughput for different base station antenna configuration under load (Mbps/km²).




Figure 6 – Capacity enhancements using 5G technologies beamforming and MU-MIMO

The baseline is a non-BF system, with 2x2 MIMO. With adding a 8x8 antenna on the base station, the capacity increase both with beamforming and MU-MIMO. For example, with 50 Mbps cell edge at load, the introduction for beamforming with SU-MIMO increase cell capacity from ~80 Mbps to ~130 Mbps (62% increase) and introduction of MU-MIMO increase to further up to ~220 Mbps (total 175% increase). Use higher order of massive MIMO increase the capacity even further.

The better reach and penetration with the mid-band spectrum makes this band attractive for deployment in suburban areas with medium to low density or rural areas. The increased capacity gained with 5G technology enables higher density of subscriber, addressing the high to medium dense suburban areas better. The better penetration through walls also makes this spectrum attractive as a capacity booster for enhance MBB.

Conclusion

The new capabilities enabled by the emerging 5G technology creates an opportunity for a fixed wireless converges. We have started to see the communication industry players in wireless and wireline moving towards each other's domains offering combinations of residential, enterprise, mobile and entertainment/media services.

One area that is seen in US to be an early 5G service is fixed wireless access, offering residential and enterprise customer high speed broadband service in an efficient way. The networks build for FWA are prepared to support full mobility, and will when combined over wireless network currently deployed network, provide a high capacity network to address the increases data demands for mobile users.

Leverage 5G in the convergence of wireless and wireline business will give an operator or a MSO a flexible network that enables the opportunity to offer a wide range of communication services that can be tailored in capability and characteristics based on the end-user and application requirements.

Abbreviations

Table 1 – Abbreviations

3GPP	3 rd Generation partnership project
4G	4 th Generation mobile technology
5G	5 th Generation Mobile technology
AP	access point
MBB	Mobile Broadband
bps	bits per second
BF	Beam forming
CBRS	Citizen Broadband radio services
CPE	Customer Premises Equipment
DL	Downlink
ECS	Environmental sensing system
FEC	forward error correction
GAA	General access authorization
HD	high definition
HFC	hybrid fiber-coax
Hz	hertz
ISBE	International Society of Broadband Experts
LTE	Long term evolution
MIMO	Multiple input, multiple output “ Antenna systems”
mmW	Millimeter Wave frequencies
MNO	Mobile Network Operator
MSO	Multiple System Operator
OFDM	Orthogonal frequency division multiplexing
SCTE	Society of Cable Telecommunications Engineers
UL	Uplink

Bibliography & References

1. Bibliography



ERICSSON

Anders Svensson
Principal Solution Manager, 5G
Ericsson North America

Anders Svensson, Principal Solution Manager, 5G for Ericsson North America, works in the Network Evolution group. Based in Plano, Texas, he is responsible for driving 5G and Network Evolution in North America.

Prior to his current role, Svensson worked as CTO and Principal Solution Manager in one of Ericsson's Tier 1 Customer Units. Svensson has also served as Product Line Manager for Ericsson CDMA solutions. Before moving to the U.S., he worked in Sweden in System Management of Ericsson's packet core products.

Svensson holds a Master's degree in Engineering Physics from the Chalmers University of Technology in Gothenburg, Sweden.

2. References

[1] ICT-317669 METIS project, Updated scenarios, requirements and KPIs for 5G mobile and wireless system with recommendations for future investigations, Deliverable D1.5, April 2015, available at: https://www.metis2020.com/wp-content/uploads/deliverables/METIS_D1.5_v1.pdf

[2] Ericsson, Ericsson Mobility Report, June 2017, available at: <https://www.ericsson.com/en/mobility-report>

[3] FCC, Notice of Inquiry, Expanding Flexible Use in Mid-Band Spectrum Between 3.7 and 24 GHz https://apps.fcc.gov/edocs_public/attachmatch/FCC-17-104A1.pdf

Shaw Communications IPv6 Deployment

Developing Company Momentum

A Technical Paper prepared for SCTE•ISBE by

Darren Gamble
Systems Architect
Shaw Communications
2728 Hopewell Place NE
403-781-4948
Darren.Gamble@sjrb.ca

Introduction

Despite a decade of conferences, papers and seminars dedicated to the subject, IPv6 deployment remains mixed amongst MSOs in 2017.

Most technical staff now have the knowledge, desire, and good reasons for deploying IPv6. However, an organization's culture and processes may remain as obstacles. By both realigning one's IPv6 deployment to meet the immediate needs of their business, and understanding how things are most effectively done in their own organization, they can overcome these hurdles and make better progress.

This document is intended for technical audiences, who benefit most from this information.

Background

Devices on the Internet must be uniquely addressed and use common protocols in order to communicate with each other. Without this, users would not be able to reach some or all other users and sites.

IPv4 is the addressing protocol used on nearly every device on the Internet. Deployed in 1983 [1], it was never intended to be used on a network the size of today's Internet. Problems include:

- The central body for address management has already allocated all available IPv4 addresses [2].
- Large MSOs have already exhausted their internal-only address allocations, and many have resorted to using non-advertised space on the Internet, hoping that it will not be repurposed.
- Address allocation for infrastructure must be carefully allocated. A network that grows outside of its original purpose may need to be painfully readdressed.
- Merging internal networks from two companies is generally impossible without mass-readdressing.

IPv6 was introduced to address these problems, including nearly unlimited address space [3]. Devices can be addressed with both IPv4 and IPv6 address (Dual-stack) which allows compatibility with both IPv6-only and IPv4-only devices. Multiple private networks can be merged without readdressing [4].

While IPv6 is becoming increasingly important with the growth of the Internet, the adoption of IPv6 for both internal and Internet use has been mixed, however. Reasons include:

- Adding IPv6 to one's Internet customers has limited business value in 2017. Websites and Internet services will continue to have IPv4 addresses for the foreseeable future, until the number of Internet users with IPv4-only connections becomes very low.
- Adding IPv6 to one's own website or Internet service also has had limited business value as IPv6-only Internet users are virtually nonexistent. Organizations may still choose to add IPv6 to their sites for altruistic reasons or as a form of positive advertising to other technical users, but only a minority of the top 1000 sites on the Internet have done so as of this paper.

- IPv6 does not solve the IPv4 exhaustion problem that MSOs face, although it may be part of its solution.
- Deploying IPv6 means additional costs, including:
 - Costs and time to replace hardware to support it
 - Costs and time to update firmware and software versions to support it
 - Upgrading resource assurance and resource inventory systems to support it
 - Training for operational and support staff

Organizations such as the Internet Society and Google have tried large-scale events and conferences to encourage adoption [5] [6]. These have been partially successful, however as of January 1 2017, Google reports that only 16.5% of its users are using Native IPv6 connections [7].

IPv6 Adoption

We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.




Figure 1 - Percentage of users that access Google over IPv6 [7]

Internet IP exhaustion aside, there are still many other useful reasons for an organization to start using IPv6 now.

Shaw Communications began its IPv6 deployments in 2009, starting with IPv6 on its backbone and enabling IPv6 transport on its caching and authoritative DNS systems in 2010.

In 2012, Shaw's next project was to deploy IPv6 to its customers. The project was widely supported by its network, broadband, activation and operational teams. However, during a change event in the production field trial, there was an outage with light impact. The project members were asked why they were doing the change event, and the reasons given were "We're running out of IPv4 IPs" and "We need to future-proof our network". It was determined that these reasons were not enough to justify the risk, and the project was postponed.

Shaw instead focused its IPv6 efforts internally, moving its CMs to IPv6 management addresses, and identifying ways that IPv6 would assist its future projects, such as with its BlueSky video product and partner-managed eRouters. Each of these projects had easily quantified benefits for using IPv6, and would also bring Shaw closer to deploying reliable dual-stack Internet service.

Potential IPv6 deployments

IPv6 Customer Internet may be the goal as laid out by the Internet Society and Google, but for most MSOs, this represents a lot of work for very little business value, at least in 2017. Very few residential customers base their choice of Internet over it. It also does not directly solve the problem of IPv4 exhaustion – that requires different work that may or may not use IPv6.

However, technical staff may find greater success in deploying IPv6 in other areas and in smaller projects which have more quantifiable value. Doing so not only makes their infrastructure easier to manage, but will close the gap on the work needed for larger projects like IPv6 Customer Internet, making that decision much easier to justify.

There are many ways that IPv6 can be used in an organization to give value to its customers, enable new architectures, or simplify management. In an organization where IPv6 usage has been slow, or work must be justified with benefits understandable by one's management, one could consider some of the following options:

1. Network support

Having knowledge and support for IPv6 on a MSO's network is a prerequisite for any deployment.

Shaw Communications was able to justify this work due to demand for IPv6 service from some of its non-cable business customers. If an MSO has a direct business need for IPv6 service, then this alone can be justified as a project. Otherwise, it may need to be done as a prerequisite for another project instead.

2. Addressing for MSO-managed CPE equipment

2.1. RFC 1918 Address Exhaustion

For larger MSOs, the number of CMs, CMTAs, digital receivers and other CPE equipment may exhaust the private address space offered in RFC 1918. Even before this happens, the MSO will need to carefully portion out space and perform frequent changes to reallocate it.

The MSO may commandeer public IPv4 address space that is not currently advertised on the Internet, and hope that it does not get used and that its users do not notice it [8]. A safer, simpler solution is to move many of these devices to IPv6. All DOCSIS 3.0+ and some DOCSIS 2.0 modems can be addressed with IPv6. Very old CMs and older embedded devices may not support it, but IPv4 and IPv6 devices can coexist.

Moving all devices to IPv6 is not necessary; only enough that IPv4 address allocation is no longer a problem.

Shaw Communications chose to move all of its non-EOL CMs to IPv6 management addresses. This frees up address space under 10.0.0.0/8, allow rehomes to be done more easily, and to allow modern CMs to be accessed by vendors for future projects.

If a MSO is suffering from this particular problem, IPv6 should be strongly considered as both a solution and also an ideal means to introduce its staff to the technology.

2.2. Network access requirements

Conversely, a MSO may have a need to provide network access to some devices; to a vendor, partner or other entity, which otherwise does not require Internet access.

Doing this with IPv4 may be impractical due to security and/or use of scarce public addresses. IPv6 may be a better solution to these problems, in either a single-stack or dual-stack configuration.

Shaw Communications deployed its BlueSky product single-stack on IPv6. Amongst other benefits, this allows regional IP address definitions to remain static.

3. Infrastructure management

DNS, DHCP, activation and provisioning systems will need to be addressed with IPv6 to have these services.

However, one may also choose to use IPv6 on some or all management networks. This greatly simplifies network management, eliminates the problems with subnetting, and allows any number of virtual servers to be quickly created or destroyed on a given network. While addressing servers is straightforward, changes are needed to the network, resource inventory and resource assurance systems.

Shaw Communications deployed this support in stages, starting with the minimums to support IPv6 servers, and expanding as additional IPv6 systems were required.

4. Caching DNS transport

Addressing an MSO's DNS caches with IPv6 can allow it to reach other IPv6 authoritative servers, and/or allow IPv6 clients to reach it.

IPv6 support in both DNS products and servers is extremely mature [9]. The incremental work needed to implement this change is very small, if the MSO is able to build and support IPv6 servers. By itself, it will increase the reliability and speed of the cache by increasing the number of sources it can get data from. Shaw Communications justified its change based on this benefit. The real value of this change is its ability to serve DNS to IPv6 clients, which is a prerequisite to other deployments.

5. Authoritative DNS record and transport

Like DNS caching, the incremental work to add IPv6 support to authoritative DNS systems is small [9].

This change would be necessary if authoritative DNS systems are needed in IPv6-only networks with IPv6-only caches, or IPv6-only devices that do not use a cache. This is not common but such a need may exist.

A MSO may instead do this for positive perception by the Internet community and its customers. Internet IPv6 authoritative servers are visible to the Internet, and adding IPv6 hosting to one's domains is far easier than adding it to one's Internet customers.

Shaw implemented this in 2010, to both increase the visibility of its servers, and to gather data about how Internet DNS caches use IPv6.

6. Customer Internet through eRouter

In 2017, modern mobile and computer operating systems have very mature support for dual-stack IPv4 and IPv6 configurations. Most services on the Internet with IPv6 enabled are not expected to pose problems.

The reliability of dual-stack IPv4/IPv6 eRouters does vary enormously, however. Each model will need to be tested, and one should expect to stage their deployment by device type. One's testing should include how the device blocks incoming IPv6 traffic to its internal network, what size of prefix(es) the

device requests, how the user can selectively allow traffic in, and how the device gracefully readdresses its internal network when there is a change to its prefix delegation.

One will also need to do detailed testing of how performance will change when a device is made dual-stack, as dual-stack customer devices will generally prefer IPv6 when it is available. APNIC has released data comparing IPv4 to IPv6 performance with loading website images, showing that IPv6 is generally faster [10] [11]. But, a more detailed report showed that YouTube performance was poorer over IPv6. The report cited several reasons, including the time needed for the O/S to determine the appropriate protocol [12]. Shaw's investigations in 2017 on its network revealed that performance does vary from site to site, but on average, IPv6 is slower. Shaw views this as an obstacle to overcome, and not an inherent flaw with IPv6. One should do similar testing on their own network and ensure that all on-site CDN caches are dual-stacked prior to putting customers on IPv6.

V6/V4 RTT Comparison by country (ms)

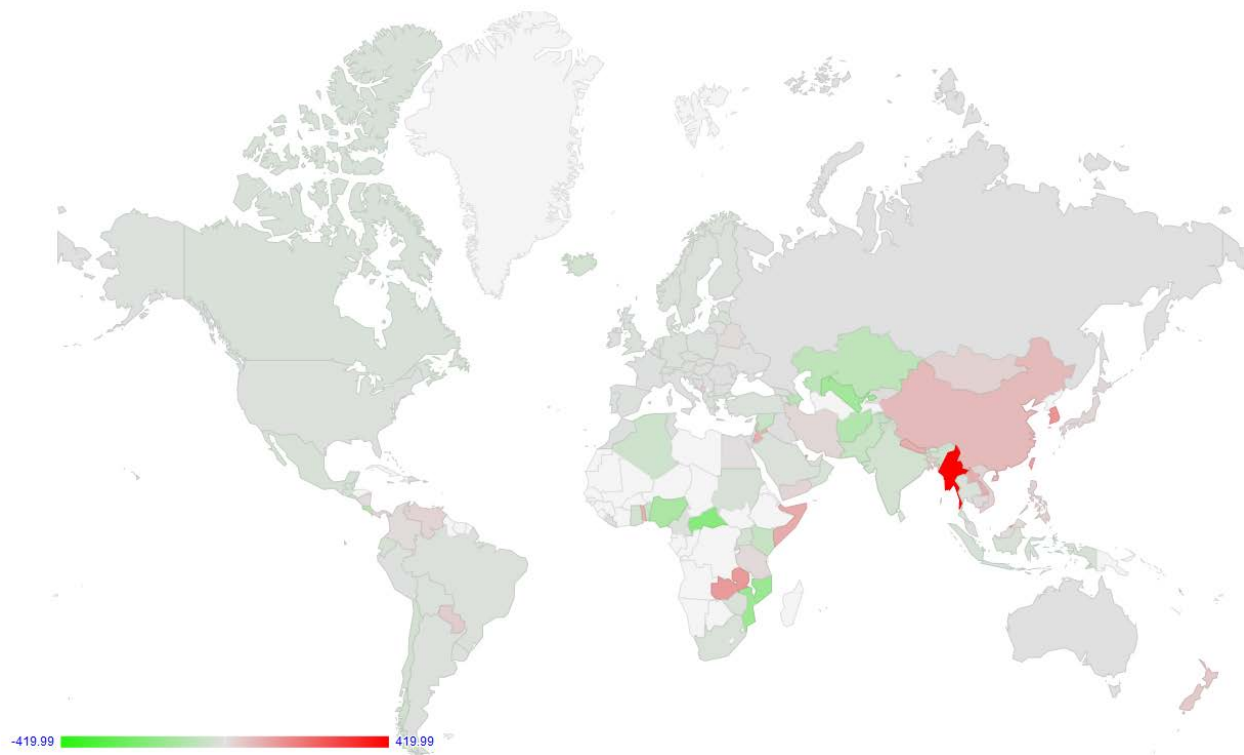


Figure 2 – V6/V4 RTT Comparison by country (ms) [11]

Significant training of frontline staff with IPv6 is also required.

Shaw's provisioning and activation systems can support dual-stack CPE devices as of this report, but deployment has been limited to due concerns with customer experience, eRouter support, and lack of need.

7. Customer Internet without eRouter

A MSO may want to delay supporting 3rd-party devices until it is comfortable with its deployment of dual-stack enabled eRouters.

At the time of this paper, many third-party eRouters do support dual-stack, but not all do it reliably or securely. One's frontline staff must be experienced enough with IPv6 to accurately determine if a problem lies with the customer's device or elsewhere.

Shaw has not extensively tested consumer eRouters.

8. Enterprise network

Addressing one's corporate LANs allow staff members to more easily access and test IPv6 devices.

Moreover, it gives all of the company's staff detailed exposure to the technology.

Shaw has not pursued IPv6 on its Enterprise network.

9. Others

This list is far from exhaustive. IPv6 is not a service- it is a tool. It can be used to make something new easier to build or something existing easier to manage.

Conclusion

An engineer wishing to introduce IPv6 into its network should remember the following:

- 1) **In order to get engagement and momentum in the company, one must understand the motivations and culture(s) in the company.** How important is executive buy-in vs. support of its operational teams? Or do things get done by having many managers come to a consensus? How soon do benefits need to be realized? How much autonomy do staff have?
- 2) **Separate components of IPv6 customer Internet into smaller projects.** They are easier to do, and some have their own benefits, and will bring one closer to their final goals.
- 3) **Justify each of the deployments with quantifiable benefits.** Explaining why IPv6 is needed on a residential Internet network may be difficult or impossible, but adding it to one's infrastructure may be easier.

Abbreviations

DNS	Domain Name System
DHCP	Dynamic Host Configuration Protocol
CM	Cable Modem
eRouter	Embedded Router
DOCSIS	Data Over Cable Service Interface Specification
MSO	Multiple System Operator,
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

- [1] RFC 791, "RFC: 791 INTERNET PROTOCOL," September 1981. [Online]. Available: <https://tools.ietf.org/html/rfc791>.
- [2] L. Smith and I. Lipner, "Free Pool of IPv4 Address Space Depleted," 3 February 2011. [Online]. Available: <https://www.nro.net/ipv4-free-pool-depleted/>.
- [3] RFC 2460, "RFC 2460 - Internet Protocol, Version 6 (IPv6) Specification," December 1998. [Online]. Available: <https://www.ietf.org/rfc/rfc2460.txt>.
- [4] RFC 4193, "RFC 4193 - Unique Local IPv6 Unicast Addresses," October 2005. [Online]. Available: <https://tools.ietf.org/html/rfc4193>.
- [5] Internet Society, "World IPv6 Launch," 2012. [Online]. Available: <http://www.worldipv6launch.org/>.
- [6] Google IPv6 Implementors, "Google IPv6 Implementors Conference," 2010. [Online]. Available: <https://sites.google.com/site/ipv6implementors/>.
- [7] Google, "Google IPv6 Statistics," 1 January 2017. [Online]. Available: <https://www.google.ca/intl/en/ipv6/statistics.html>.
- [8] R. Graham, "DoD address space: it's not a conspiracy," 16 December 2013. [Online]. Available: <http://blog.erratasec.com/2013/12/dod-address-space-its-not-conspiracy.html#.WYSY1VGQxaR>.
- [9] Wikipedia, "Comparison of DNS server software," [Online]. Available: https://en.wikipedia.org/wiki/Comparison_of_DNS_server_software.

- [10] APNIC, "V6/V4 RTT Comparison by country," 16 8 2017. [Online]. Available: <https://stats.labs.apnic.net/v6perf>. [Accessed 16 8 2017].
- [11] APNIC, "Measuring IPv6," 17 8 2017. [Online]. Available: <https://labs.apnic.net/measureipv6/>. [Accessed 17 8 2017].

Addressing IP Video Adaptive Stream Latency and Video Player Synchronization

A Technical Paper prepared for SCTE•ISBE by

Jeffrey Tyre

Distinguished System Engineer
ARRIS
2450 Walsh Ave,
Santa Clara, CA 95051
408-940-2095
jeffrey.tyre@arris.com

Wendell Sun

Member of Technical Staff
Viasat
6155 El Camino Real
Carlsbad, CA 92009
760-893-5577
wendell.sun@viasat.com

Introduction

Since early deployment of IP-based video networks, various technologies have emerged to help cope with the variability associated with delivering video over non-deterministic, best effort IP networking. In managed environments, IPTV operators have traditionally used MPEG-2 TS as the transport mechanism for video over IP networks. Hyper Text Transfer Protocol (HTTP), by virtue of being the content transport protocol for web-based applications, is almost ubiquitously used for video delivery over the Internet.

Traditionally HTTP video was delivered by progressive file download. However, a newer technology called adaptive bit rate (ABR) streaming has become widely used. ABR streaming promises to enable videos to be delivered over unmanaged networks with a very high quality of experience, and is thus applicable to both Internet video environments and managed video networks that are seeking to extend the delivery of premium content to devices other than the television set. ABR streaming has emerged as a technology of choice for many types of video delivery. For managed-networks, Pay TV Operators migration to adaptive streaming fits into an overall strategic objective for a unified, video delivery IP network infrastructure supporting every device screen and all subscriber services using web and cloud-based technologies.

Unlike previous HTTP video technologies, such as progressive download, adaptive streaming introduces the ability to dynamically react to changes in network conditions by switching to a video encoded at a different bit rate. This ability to adapt in real time more accurately reflects the dynamic conditions of today's networks, content, and devices. With users streaming more premium long form content, it is natural to expect that there will be fluctuations in the amount of bandwidth available during a two-hour movie, than say a 3-minute video clip. Adaptive streaming is a recognition of this fact and enables viewers to watch this premium content with a superior quality of experience (QoE).

Adaptive streaming works by leveraging the same content encoded in various bit rates—in a range that reflects the expected quality of the content itself, the network performance, and the screen resolution desired. For example, a video could be encoded in bit rates ranging from 300 kbps (low-quality, online video) up to 6 Mbps or higher (high quality streaming content to the TV). A typical video could be encoded in as many as eight different bit rate profiles, depending on the range of devices and quality desired. Each of these files are then further segmented—or “chunked”—into short segments (typically two to ten seconds long) that are each precisely time-stamped.

As the video is delivered, the HTTP client maintains a communication channel with the adaptive bit rate server. The client (the viewing device) downloads these chunks as individual files, which are buffered by the client, decoded, and played out as a continuous presentation of video and audio to the viewer. During the viewing session, the client player monitors the rate at which the buffer is filling and can thereby infer the performance of the network.

If there is degradation in network performance, the client can request that chunks be delivered from one of the lower bit rate files. This is all seamless to the viewer since each source file is chunked and time-stamped in the same, very precise intervals—so there is no visible interruption or hesitation when switching to a different bit rate. Likewise, if the player detects an improvement in performance, it can request HTTP file segments from one of the higher bit rates.

Adaptive Bit Rate (ABR) video streaming

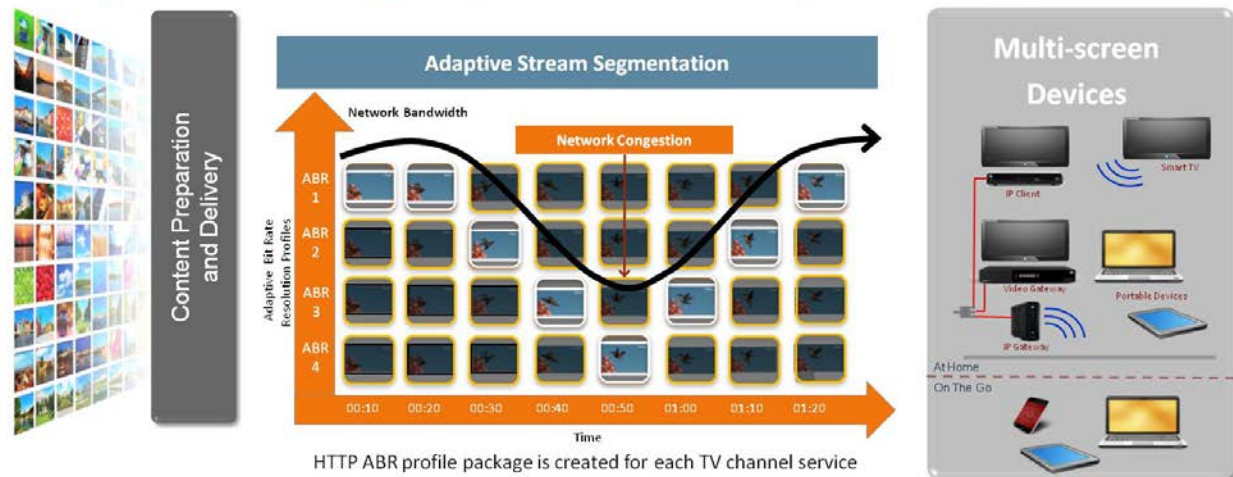


Figure 1 – Adaptive Bit Rate (ABR) Ecosystem in IP Networks

Since adaptive streaming video content is prepared and conditioned into multiple bit rate sources to allow the client player to dynamically select the appropriate bit rate source and seamlessly switch between different bit rate sources per broadband network status, it can be used by almost any type of multiscreen device, provide transport resiliency to a network's condition and gives a much better user viewing experience.

For the video on demand (VOD) type of service, this is near perfect since VOD service is delivered to each end device individually in non-real-time fashion, e.g. IP unicast, and there is no concern for delivery latency and no demand to coordinate viewing experience among end users. The initial buffering delay most likely is acceptable, especially when there is a pre-roll ad play. The commercial application is very successful for Internet-based VOD services, such as Netflix, Hulu, Amazon, etc. However, the HTTP-based transport has major shortcomings when it is applied for live or linear video delivery.

Content

1. Live / Linear TV Service Requirements

The live or linear TV service presents some unique challenges for using adaptive streaming technologies including some of the following characteristics:

- There is a real-time timeline that is referenced by all viewing users at the same time when the content is captured, processed, delivered, and consumed. For example, a sport event is being broadcasted while the event is happening
- Compared with an audience on site, normally there is a constant viewing delay for broadcast viewers. The viewing delay is counted from event happening to being viewed remotely and usually is caused by video content acquisition, editing, processing and delivering, as well as mandatory regulation delay request. The smaller the viewing delay is, the better the user viewing experience will be. See the Figure 2 for reference
- Content is delivered to a large number of viewers simultaneously
- Constant viewing delay is the same to all viewers




Figure 2 - TV Service Viewing Delays

In the ABR case, HTTP is used for content delivery. In general, HTTP is a transaction based protocol designed for file download. The current adaptive bit rate streaming uses small segments to compromise HTTP file transfer request. Instead of bit by bit streaming, content is transported segment by segment.

Normally an HTTP transfer will not start until the whole segment is ready. This will add, at minimum, one segment length of extra time to delivery delay besides other processing delays. This is shown in figure 2 as segmentation delay. The bigger the segment is, the longer the extra latency is. This may force small segments to be defined if low latency live TV service is desired.

However, for seamless switching purposes, a segment is bounded with an instantaneous decoding refresh (IDR) frame, e.g. closed group of pictures (GOP), which requires more coding bandwidth. The smaller

the segment, the worse the video encoding efficiency will be. In addition, each segment delivery corresponds with one HTTP get/reply transaction. The smaller the segment is, the more HTTP protocol overhead will be in network.

Before Internet-based video services were available, IP-based streaming TV services, such as IPTV service, had a long history. IPTV services have been offered by telco operators to compete with cable operators for more than a decade. IPTV services provide subscribers with similar TV viewing experience, such as fast channel change time and low end-to-end transport latency, just like traditional broadcast or linear Pay TV service offers. It uses IP multicast as its primary protocol, which provides true data streaming, minimizes delivery latency, and supports content sharing among multiple clients. However, IPTV services require guaranteed the bit rate to match the bandwidth for smooth service delivery. As a result, an IPTV service is only provided in a managed IP network.

This paper presents how adaptive transport stream (ATS) segment markers, HTTP chunked transfer encoding (CTE), and ABR playlist manifests can be combined into a solution optimized for live video content delivery and become an effective tool for Pay TV Linear Services. It briefly reviews current IP multicast streaming and examines existing or under-developing protocols and standards, such as the HTTP chunked transfer encoding [7] and the CableLabs/SCTE adaptive transport stream [3]. This paper proposes a solution, which combines ABR encoding with content segment markers and HTTP CTE streaming, to minimize the delivery latency for live/linear video service without sacrificing video encoding efficiency. Additionally, a general modification to ABR manifest formats is proposed to address ABR player synchronization challenge.

2. TV Services via Satellite Broadcast and IP Multicast Streaming

Traditional TV services are all in broadcast mode. The content is acquired from source, edited, and processed for transportation over satellite. Satellite is used for large area content distribution. Depending on the business model, there are a couple of ways of receiving content for consumption:

- For Pay TV cable operators, an integrated receiver/decoder (IRD) at the multichannel video programming distributor (MVPD) headend receives and decodes the content, and then the content is re-distributed to subscribers via the cable plant network.
- For Pay TV satellite operators, a satellite set-top box in the subscriber's home receives, decodes, and renders the content.
- For over-the-air operation, it is similar to the Pay TV cable operation in the front line of receiving and decoding, but the last leg of content re-distribution is done by over-the-air radio transmission.

All of these are broadcast one-way from the content source to many content consumption terminals. This usually carries a constant viewing delay, as shown in Figure 2, as regular TV service viewing delay.

Another type of TV service, IPTV, uses broadband Internet as its distribution network. The Internet can support different types of content distributions, such as one-to-one unicast, one-to-many multicast, and one-to-any broadcast. However, IP multicast has significant advantages over IP unicast and broadcast for TV-like services.




Figure 3 - IP Multicast Illustration

IP Multicast offers a one-to-many kind of distribution model, similar to traditional broadcast TV services that are mentioned above with one difference. In the world of the public Internet, IP multicast supports content delivery on per request or per registration basis. To receive content, the end user's device needs to send Internet Group Management Protocol (IGMP) requests to its router to ask for delivery. If the requested content is not available on the end user's device, the router must relay the request all the way back to the content source. To stop receiving content, the end user device can send an IGMP quit command to its router to stop the sending. The user B in Figure 3 is exempted from receiving content. The routers that support IP multicast service only send one copy of same content to the duplicated requests from the next router. For example, if both users C and D ask for the same content, the upstream router just needs to send one copy of content to their router. In this case, even if there is a single user, such as user A, or there are thousand users under the same router, the overall network traffic will look same. Thus it scales very well and is quite efficient for live or linear type of content distribution. Robinson's paper [1] has a lot more details on IP multicast discussion.

More importantly, IP multicast is a true IP streaming protocol. Although content is delivered by IP packets, the protocol is designed for content flowing from point A to point B as long as point B joins the multicast group. IP packets are relatively, very small and do not introduce much extra transportation delay compared with other forms of bit-stream content delivery.

IP multicast has been used extensively by telco operators to provide IPTV service based on delivering MPEG-2 transport streams (TS) to the telco set-top box in the home. It allows the IPTV service to provide a similar user experience as traditional broadcast TV service does. Even for some cable operators, IP multicast may also be used in their backbone for content distribution due to the popularity of Internet and IP network integration.

To supply a satisfied IPTV service or a reliable backbone for content delivery, the only constraint of using IP multicast is the requirement of guaranteed bandwidth – that the bandwidth should be equal or higher than the bit rate of the content stream. For this reason, IP multicast is usually applied in a managed IP network with well-engineered bandwidth allocation.

Whether in a traditional broadcast TV service or in an IPTV service, MPEG-2 TS is used for carrying video and audio content. MPEG-2 TS is the currently the dominant container format for content distribution of most Pay TV services, and several key TS advances recently introduced are providing

enhancements for adaptive content encoding in preparation of adaptive streaming delivery using HTTP as the transport protocol for multi-screen device services.

3. Adaptive Transport Stream and Segment Boundary Points Support

An adaptive transport stream (ATS) is a fully compliant MPEG-2 TS with an embedded segment boundary indicator. The ATS is designed for adaptive streaming, which requires encoding/transcoding of multiple bit rate (MBR) streams, segmentation, and alignment to support seamless bit rate switching. Originally the MBR encoding/transcoding is typically done by encoder or transcoder, while the segmentation and alignment is processed by the adaptive streaming packager. The purpose of ATS is to further exploit the encoding/transcoding process to label the media segment boundary, relieving the effort of parsing and aligning segments in the adaptive streaming packager.

The ATS definition originated in the CableLabs encoding boundary point (EBP) specification [2] and adaptive transport stream specification [3], in which an EBP structure is defined in the private data of the adaptation field in MPEG-2 TS header to indicate the beginning of each segment. The work was promoted to ANSI/SCTE standardization process, in which ANSI/SCTE 223 [4], has been created. Furthermore, it has been contributed to MPEG and an amendment to MPEG-2 TS specification [5] has been generated to define a set of adaptation field (AF) descriptors to serve the same purpose.

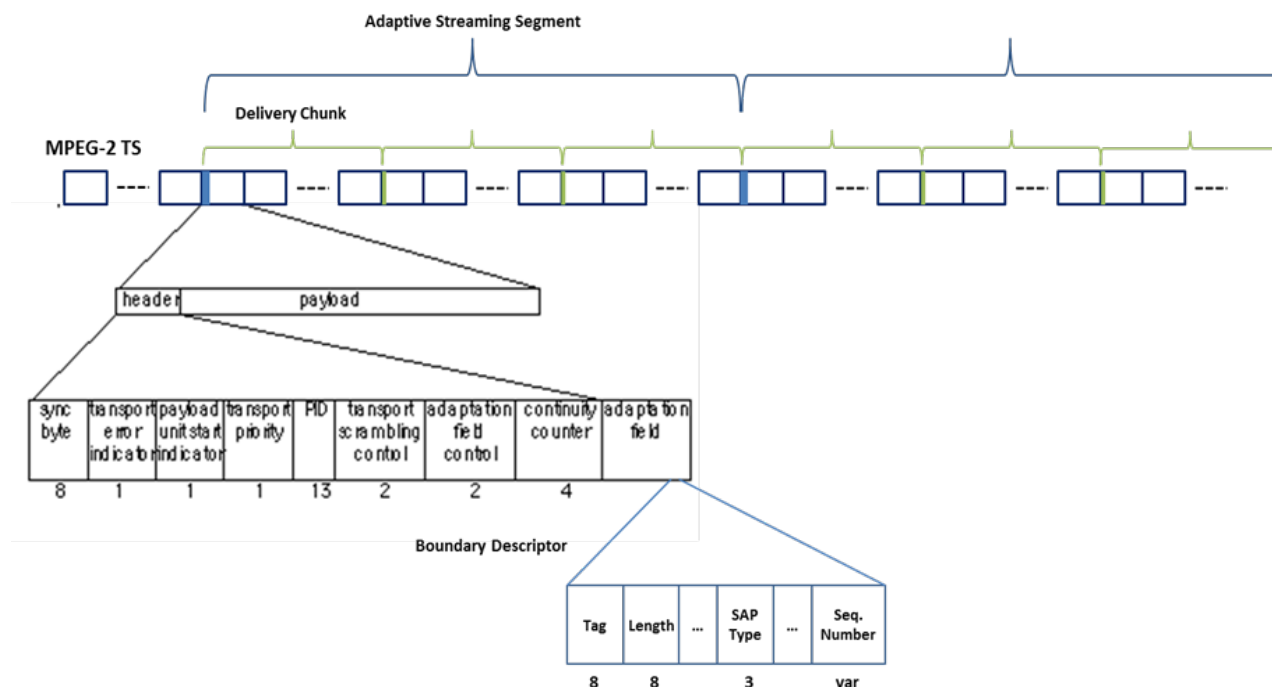


Figure 4 - Structure of Adaptive Transport Stream

As it is shown in Figure 4, the Boundary Descriptor is introduced into the adaptation field of MPEG-2 TS header. The two major parts are the “SAP type” and “sequence number” field. The SAP stands for stream access point and is defined by ISO/IEC 14496-12. A SAP type is a definition of media decoding attribute in that stream point. For example, a SAP type 1 means the media sample at the point can be fully decoded without referring other samples and all samples following it can also be correctly decoded. The sequence

number field has length of 2, 4, or 8 bytes depending on application. It provides a unique identifier for a segment within its context.

To support regular MBR adaptive streaming, the signaling of a boundary descriptor in the segment level is good enough. The packager takes ATS as input stream and needs only to parse the bytes in the transport stream packet headers in order to obtain the boundary information. It does not need to parse any bytes in the packet payload. However, to resolve the extra viewing delay caused by segmentation as it is discussed in the introduction section for live/linear TV service, a smaller segment is desired; and coding efficiency should not be impacted. The segment is designed for bit rate adaptive switching; thus, it requires closed GOP at segment boundary. While the segment structure is maintained, a smaller delivery unit can be designed for low latency delivery.

The MPEG DASH specification [6] introduces a concept of delivery unit media segment to support low latency application. As shown in the Figure 4 - Structure of Adaptive Transport Stream, the delivery chunk is a smaller delivery unit within the adaptive streaming segment. It can be made up by any group of meaningful coding samples, such as a GOP or even a frame, while making it small enough not to cause transportation delay. Similar to each segment, the delivery chunk can also be identified by the boundary descriptor embedded in MPEG-2 TS header.

The insertion of boundary descriptors into MPEG-2 TS headers can be easily achieved as part of the content encoding/transcoding process. It does not add much extra processing to the encoder/transcoder, yet it is a big savings for the packager to not need to look up the coding payload. It is even more essential when the segment durations are not fixed interval - whether that's due to flexible GOP structures in the encoded video, or a content-related change, such as a frame-accurate demarcation in a program or advertisement.

4. Addressing ABR Content Delivery Latency with HTTP Chunked Transfer Encoding Streaming

The HTTP chunked transfer encoding (CTE) is defined by HTTP/1.1: Message Syntax and Routing [7]. It is a data transfer mechanism in HTTP 1.1, in which data can be sent in a series of "chunks" in responding a single HTTP data request. It uses the transfer-encoding header, instead of the content-length header. The HTTP sender does not need to wait for the total size of content being available and can start sending data as small chunks with any amount available while still receiving the content. When the chunk is sent, its size is indicated in the transfer-encoding header. At the end of content, it sends the last chunk with its size set to be zero. For example, the following is a HTTP transaction with chunked transfer encoding:

```
GET /cte-example.html HTTP/1.1
Host: doc.micrium.com
User-Agent: Mozilla/4.61 [en] (Windows 7 6.1)
TE: chunked
```

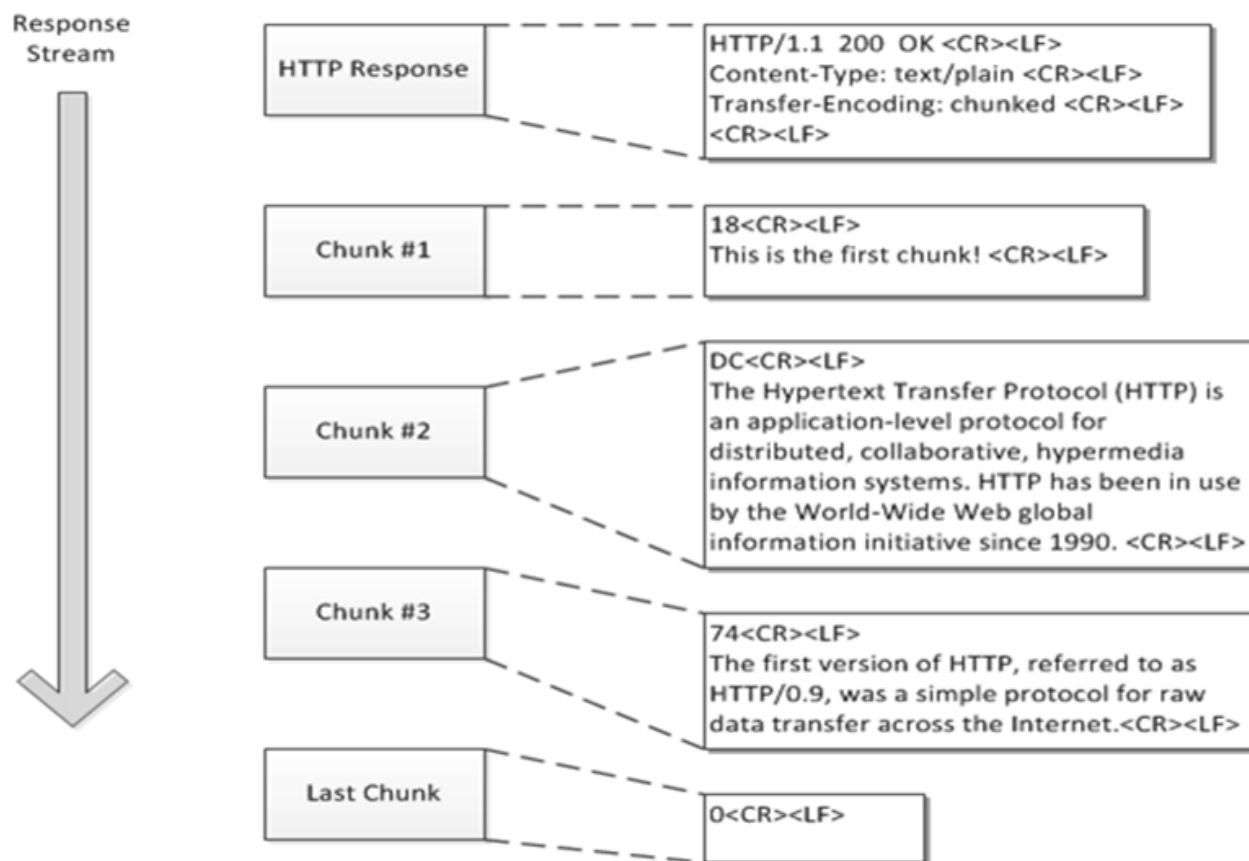


Figure 5 - Example of HTTP Chunked Transfer Encoding Response [8]

As we have discussed in the adaptive transport stream section, ATS are MPEG-2 TS streams with embedded boundary descriptors to virtually divide them into delivery chunks, and further into adaptive streaming segments that are published to ABR players as ABR manifest playlists to select ABR content files. The client side remains with HTTP based ABR streaming, which allows content to pass most access networks and reach almost all types of client devices based on the ubiquity of HTTP content delivery.

The ABR-aware HTTP streaming (A2HS) server is a new component that bridges the two sides of ATS segmentation encoding and adaptive stream delivery [11]. It can essentially be an off-the-shelf HTTP web server with additional key adaptive streaming functions such as ATS parsing, ABR packaging with various ABR format outputs, and HTTP CTE support.

In current ABR content delivery implementations, a content segment is not made available in the manifest published to an ABR client until the segment is ready for delivery. This is the primary cause of content segmentation delay. In the proposed A2HS server solution, especially with the modified manifest playlist (see section below), the content segment(s) will be published in a manifest even before the input ATS is pulled, e.g. IP multicast join. When the A2HS server receives content segment requests from an ABR client and the corresponding IP multicast ATS stream is available (otherwise it initiates IGMP to join to the IP multicast group to get it), the server starts parsing the MPEG-2 TS header of the input ATS seeking the boundary descriptor of delivery chunks and sends the delivery chunks via the HTTP CTE method while it continues to receive the IP multicast ATS input.

The ABR client receives the manifest playlist, and it selects a segment from one bit rate stream per its condition, such as network bandwidth, screen size, etc. to begin pulling the segment from the HTTP server. Each pull is a HTTP file transfer transaction. And among a group of ABR clients, each client acts independently. If bit rate switching is required, the ABR client sends the next segment request of the switch-to ATS representation. Since ABR clients can only switch on an adaptive streaming segment boundary, the proposed A2SH server maintains transport of delivery chunks of current segment until the segment boundary indicated by the boundary descriptor is reached, then it starts transport of the first delivery chunk of the switch-to bit rate segment. In this fashion, the content delivery is in the delivery chunk interval. If the delivery chunk is designed small enough, the HTTP CTE based transport provides a similar function of true video transport streaming.

5. Addressing ABR Player Synchronization Using an Adaptive Bit Rate-Tiered Manifest Playlist

5.1. Current Manifest File Formats and Unsynchronized Video Playout

All HTTP based adaptive streaming approaches currently use a manifest playlist file with segmented content for downloading by ABR video players. The ABR client receives the manifest playlist, and it selects a segment from one of bit rate stream per its condition, such as network bandwidth, screen size, etc. to start pull the segment from HTTP server. Each pull is a HTTP file transfer transaction. And among a group of ABR clients, each client acts independently.

For live/linear TV service, all clients are supposed to pull the same media segment simultaneously, thus giving viewers a synchronized viewing experience. In practice, this does not occur when using adaptive streaming due to the nature of when individual ABR clients initiate their HTTP session requests in reference to the ABR manifest playlist of published ABR segments. The problem may get worse when individual client runs with different segment selection algorithm.

Table 1 - Manifest Example Using Apple HTTP Live Streaming for a Particular Bit Rate Stream

At the moment of 12:00:00	At the moment of 12:00:07	At the moment of 12:00:10
#EXTM3U #EXT-X-TARGETDURATION:10 #EXTINF:10, ./segment3511.ts #EXTINF:10, ./segment3512.ts #EXTINF:10, ./segment3513.ts #EXT-X-ENDLIST	#EXTM3U #EXT-X-TARGETDURATION:10 #EXTINF:10, ./segment3511.ts #EXTINF:10, ./segment3512.ts #EXTINF:10, ./segment3513.ts #EXT-X-ENDLIST	#EXTM3U #EXT-X-TARGETDURATION:10 #EXTINF:10, ./segment3512.ts #EXTINF:10, ./segment3513.ts #EXTINF:10, ./segment3514.ts #EXT-X-ENDLIST

If there are four clients, each starts playback individually, the clients may end up with the following result of segment pulling and playback:

- Client A asks and receives the manifest file at time 12:00:00, and it starts playback of segment 3511

- Client B asks and receives the manifest file at time 12:00:07, and it also starts playback of segment 3511, but comparing with Client A, it is 7 seconds behind in real time
- Client C asks and receives the manifest file at time 12:00:07, and it starts playback of segment 3512, at this moment, it is 3 seconds ahead of Client A and 10 seconds ahead of Client B in real time
- Client D asks and receives the manifest file at time 12:00:10, and it starts playback of segment 3514, at this moment, it is 20 seconds ahead of Client A, 27 seconds ahead of Client B, and 17 seconds ahead of Client C in real time

The time difference of playback among these clients is between 3 – 27 seconds.




Figure 6 - Unsynchronized Playback Based on Current Manifest File Formats

5.2. A Proposed Adaptive Bit Rate-Tiered Manifest Playlist

For a live/linear media service, we are introducing a manifest playlist file that only requires listing different bit rate tier representations. This differs from current ABR manifest formats by not being dependent on individual available segment URLs but instead using a single “virtual” segment URL for the selection of bit rate stream tier by ABR players. The revised manifest file design can optimize the solution by adding one more entry to represent a “zero bit rate” stream, which can be used by the ABR client to signal stopping the HTTP CTE server’s media segment output.

For example:

```
#EXTM3U

#EXT-X-TARGETDURATION:10

./segment-720p-3000kbps.ts

#EXTINF:10,

./segment-480p-800kbps.ts

#EXTINF:10,

./segment-320p-500kbps.ts
```

```
#EXTINF:10,  
./segment-0kbps.ts  
#EXT-X-ENDLIST
```

The target segment duration (e.g. 10 seconds) in the manifest file indicates time interval of virtual segment boundary that can be used for seamless bit rate switching.

5.3. Combining the Adaptive Bit Rate-Tiered Manifest with HTTP CTE Streaming

The HTTP server, which supports HTTP chunked transfer encoding, receives MPEG ATS streams with media segments and maintains the status of available media segments. At any moment, there is only one media segment called the “current segment”, which matches to current media presentation time, for each bit rate stream. This current segment definition moves along with the timeline of media presentation.

Each segment is delivered incrementally in small chunks per chunked transfer encoding, such as one second, half second chunk or even smaller chunk duration to match with video frame, the HTTP server also maintains one current chunk position within the current segment to even closely represent the current moment of media presentation time.

As usual, the playback client initiates a request to receive the manifest file, then it selects one bit rate stream to start for playback, and asks for delivery of media segment via the “virtual” segment URL. Since the HTTP server supports chunked transfer encoding, the client does not need to repeat the media segment request, instead the HTTP server keeps pushing the chunks (and the segments) one after another until either:

- The client elects to switch bit rate streams, then sends a request for the new bit rate stream

or

- The client decides to stop receiving media segment/chunk, and sends a request for zero bit rate stream

When the HTTP server receives the media segment request by the “virtual” segment URL, it simply sends back the current chunk of the current segment in the matched bit rate stream, and keeps doing so for the follow up chunks/segments. If the server receives a request for bit rate switch, it will continue and finish sending chunks in the current bit rate and switch to the chunk of the new bit rate in the next segment boundary. It keeps on until it receives a request of zero bit rate stream, then it stops.

In this way, all playback clients receive the same media chunk if their initial media segment requests fall within media chunk interval. Since the chunk size is designed relatively small, such as one second or half second, it limits the synchronization gap to the minimum degree.

Taking the same example discussed in Table 1 with the current ABR manifest formats presented above:

- Client A asks and receives the manifest file at time 12:00:00, and it starts playback of chunk 3512-0 of segment 3512.
- Client B asks and receives the manifest file at time 12:00:07, and it receives and starts playback of chunk 3512-7 of segment 3512. Even though Client A, started 7 seconds earlier with chunk 3512-0, it now also plays back the same chunk.
- Client C asks and receives the manifest file at time 12:00:07, and it should have same result as Client B does.
- Client D asks and receives the manifest file at time 12:00:10, and it starts playback of chunk 3513-0 of segment 3513. At this moment, Client A, B and C should also start playback of the same chunk.

“*” indicates the current chunk of the current segment being delivered to each ABR player with a chunk size assumed to be a 1 second duration.

At the moment of 12:00:00	At the moment of 12:00:07	At the moment of 12:00:10
#720p-3000bps ./segment3511.ts ./segment3512-0.ts * ./segment3512-1.ts . . ./segment3512-8.ts ./segment3512-9.ts ./segment3513.ts #480p-800bps ./segment3511.ts ./segment3512-0.ts * ./segment3512-1.ts . . ./segment3512-8.ts ./segment3512-9.ts ./segment3513.ts #320p-500bps ./segment3511.ts ./segment3512-0.ts * ./segment3512-1.ts . . ./segment3512-8.ts ./segment3512-9.ts ./segment3513.ts	#720p-3000bps ./segment3511.ts ./segment3512-0.ts . . ./segment3512-7.ts * ./segment3512-8.ts ./segment3512-9.ts ./segment3513.ts #480p-800bps ./segment3511.ts ./segment3512-0.ts . . ./segment3512-7.ts * ./segment3512-8.ts ./segment3512-9.ts ./segment3513.ts #320p-500bps ./segment3511.ts ./segment3512-0.ts . . ./segment3512-7.ts * ./segment3512-8.ts ./segment3512-9.ts ./segment3513.ts	#720p-3000bps ./segment3512.ts ./segment3513-0.ts * ./segment3513-1.ts . . ./segment3513-8.ts ./segment3513-9.ts ./segment3514.ts #480p-800bps ./segment3512.ts ./segment3513-0.ts * ./segment3513-1.ts . . ./segment3513-8.ts ./segment3513-9.ts ./segment3514.ts #320p-500bps ./segment3512.ts ./segment3513-0.ts * ./segment3513-1.ts . . ./segment3513-8.ts ./segment3513-9.ts ./segment3514.ts

Table 2 - Example of CTE Server Output Using Proposed Manifest Format in Delivering Adaptive Bit Rate-tiered Manifests and HTTP CTE Delivered Chunks

The overall solution achieves playback synchronization among all playback clients with low latency delivery. A less than chunk-size time difference of playback should exist among those clients.

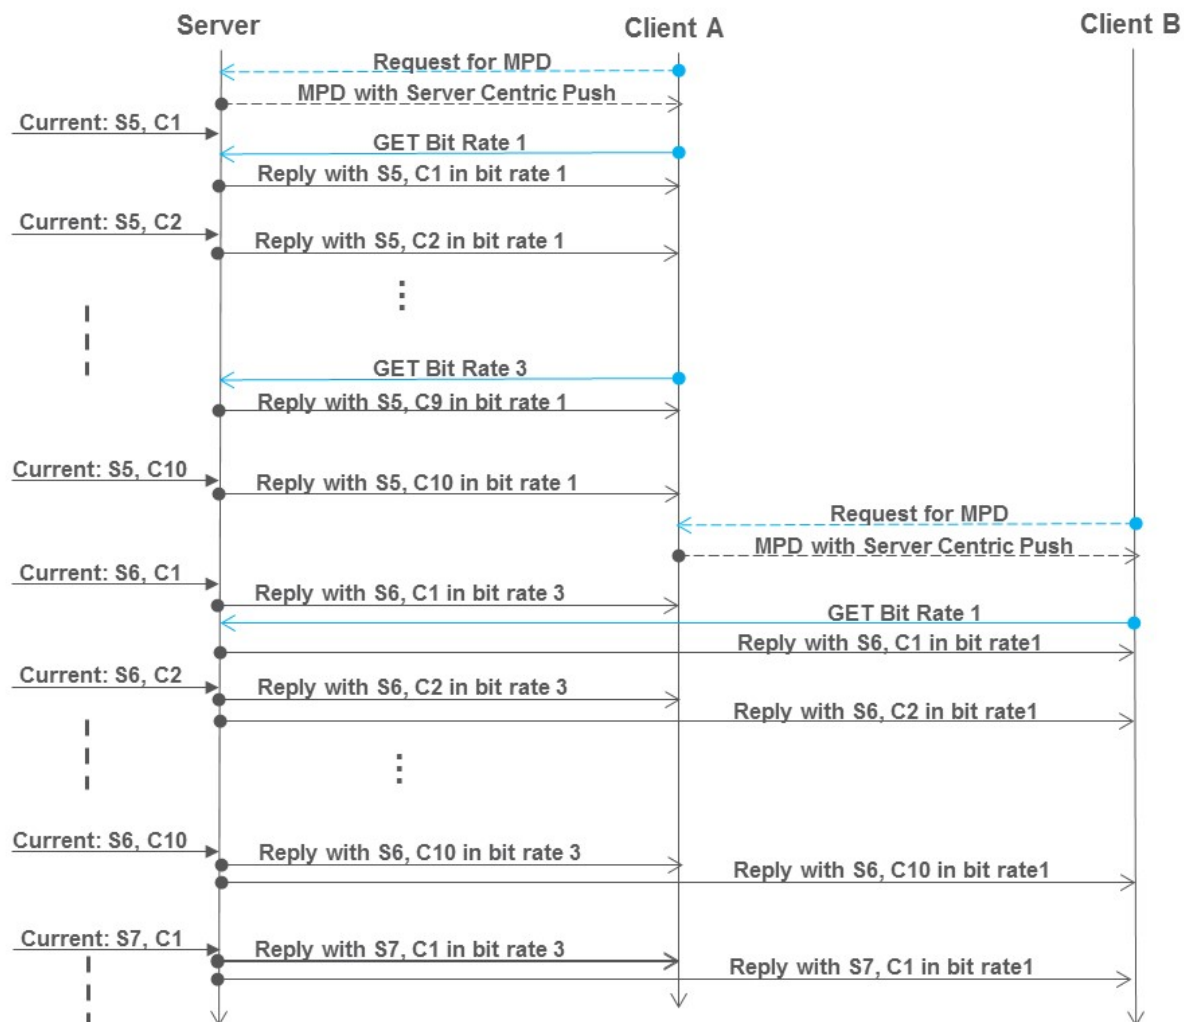


Figure 7 – ABR Client Session Flow with Proposed Adaptive Bit Rate-tiered Manifests

As illustrated in Figure 7, ABR client session synchronization is maintained even when individual clients are selecting content from different adaptive bit rate tier levels. This allows each ABR client to monitor quality of service (QoS) network conditions specific to it, which is one of the cornerstones of adaptive bit rate streaming, and effectively spreads the load of monitoring of IP Video network conditions across the entire population of ABR clients being served.

Conclusion

Compared with existing TV services, especially live and linear TV service, HTTP-based video services add extra viewing delay caused by ABR segmentation. This paper reviewed the advantages of satellite based broadcast TV service and IP multicast based IPTV service, as well as the ANSI/SCTE ATS standard and HTTP CTE protocol.

The paper proposed combining ATS segmentation description for ABR content preparation and HTTP CTE for low latency content delivery to ABR clients with a bit rate-tiered ABR manifest for video / audio playout synchronization between a population of ABR clients. This approach can both minimize the extra viewing delay caused by ABR packaging segmentation and improve live TV services based on HTTP / adaptive streaming technology's quality of experience.

This paper demonstrated that combining multiple advances in the area of HTTP-based, adaptive streaming video delivery can be used to address two of the key challenges currently facing cable operators' migration from legacy to new IP video based, live TV services, potentially paving a path for investment in next generation technologies.

Abbreviations

A2HS	ABR-aware HTTP streaming
ABR	adaptive bit rate
AF	adaptation field
ATS	adaptive transport stream
CTE	chunked transfer encoding
EBP	encoding boundary point
GOP	group of pictures
HTTP	Hyper Text Transfer Protocol
IDR	instantaneous decoding refresh
IGMP	Internet Group Management Protocol
IRD	integrated receiver/decoder
MBR	multiple bit rate
MVPD	multichannel video programming distributor
QoE	quality of experience
QoS	quality of service
SAP	stream access point
SCTE	Society of Cable Telecommunications Engineers
TS	transport stream
TV	television
VOD	video on demand

Bibliography & References

The Return of Multicast: Why it Succeeds in a Live Linear World, D. Robinson,
<http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/The-Return-of-Multicast-Why-it-Succeeds-in-a-Live-Linear-World-108621.aspx>

CableLabs OC-SP-EBP-I01-130118, Encoder Boundary Point Specification, www.cablelabs.com

CableLabs OC-SP-ATS-I01-140214, Adaptive Transport Stream Specification, www.cablelabs.com

ANSI/SCTE 223 2017, Adaptive Transport Stream, www.scte.org

ISO/IEC 13818-1:2015 PDAM 7 Virtual Segmentation

ISO/IEC 23009-1:2014, Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats

IETF RFC7230, Section 4.1, Chunked Transfer Encoding, <http://tools.ietf.org/html/rfc7230#section-4.1>

HTTP Chunked Transfer Encoding,
<https://doc.micrium.com/display/httpref/Chunked+Transfer+Encoding>

Moving Towards the Light:

Migrating MSO FTTP Networks to a Distributed Access Architecture

A Technical Paper prepared for SCTE•ISBE by

Phil Miguelez

Executive Director, Network Architecture

Comcast

Comcast Center

Philadelphia, PA

215-286-1126

phil_miguelez@cable.comcast.com

Introduction

Cable coaxial systems and more recent HFC (Hybrid Fiber-Coax) networks have been expanding across the North American continent for nearly seventy years. MSO coax access links pass more than 85% of the single family and multi-dwelling properties in the US. Cable systems have evolved over the decades from basic video carriage to multi-services transport including voice, high speed internet, and data services for both residential and commercial customers.

In 1982 the landmark United States v AT&T anti-trust settlement broke up the Bell System, which led to the creation of seven independent Regional Bell Operating Companies (RBOCs) just two years later. The Regional Bells no longer had the monopoly protection that the telephone network enjoyed in the past. Competition accelerated in the mid 1990's with the growth of the internet. Cable systems launched DOCSIS® which allowed both digital voice and data services over the MSO coax networks. The twisted pair copper wires of the phone companies reached almost every home in the US, but this basic network of copper lines had been in existence for a hundred years and had limited high speed data capacity. In order to survive, the Regional Bells started to merge and consolidate their coverage areas. They also began the process of migrating their networks from “Plain Old Telephone Service” (POTS) over twisted pair copper lines to fiber optics, in order to compete with the cable and satellite competitors that were rapidly growing by offering expanding video on demand content and faster data.

Verizon initiated FiOS (Fiber Optic Service) in 2005 and built out FTTH networks covering 18 million households, primarily in the northeastern portion of the US. AT&T developed U-verse, a hybrid network consisting of a fiber to the curb transport layer converting to a DSL over twisted pair access link to the home. The cable operator alternative to FTTH PON was RFoG (RF over Glass). RFoG involved very little new technology. The analog modulated laser that had been used to transport the RF channel load to the HFC node was now optically split to feed 32 mini nodes that were located at the subscriber homes. RFoG provided no additional bandwidth or capacity compared to HFC, but it did satisfy the demands of new home developers that insisted on a future proof “fiber to the premise” solution instead of coax. RFoG struggled to gain traction then (as now) due to a number of weaknesses with this technology – limited bandwidth (BW), limited capacity, higher construction costs associated with fiber, and the nagging concerns about optical beat interference (OBI) that was always hard to detect, initially, and had no known cure. The only major benefit of RFoG was that long fiber drops provided a much lower cost solution than HFC in rural serving areas with low homes per mile. The paucity of homes in these locations also provided lower upstream traffic congestion on the network, so the statistical chances of generating OBI was significantly reduced. Many smaller market MSOs serving these communities in the mid-west and southern areas of the country have never had a serious problem with OBI, and swear by RFoG.

The housing boom that generated the increasing interest in fiber to the home ended suddenly as a result of the 2008 Great Recession. New single-family home construction came to a standstill and financial tightening halted all but “business as usual” network maintenance projects.

Post-recession, as the financial and home construction markets recovered, a shift to multi-dwelling unit (MDU) new build projects (instead of single home communities) was clearly evident. But the most significant change in broadband access was the announcement in 2010 that Google Fiber would soon bring Gigabit fiber connectivity to a number of communities across America, starting with Kansas City, KS. Google's entry into the broadband delivery arena was the catalyst for a number of competitive

changes. Google Fiber threatened both telcos and cable operators. Gigabit service was, at that time, out of reach for telco DSL or MSO DOCSIS networks. Cable franchise agreements were now being re-examined as cities lined up to be the next potential Google “Fiberhood.” Emboldened by Google’s negotiating power to extract concessions from municipalities -- such as access to utility poles, and expedited construction permitting -- the lowered barriers to entry encouraged a new wave of fiber overbuilders to enter the market. These new competitors began to target the incumbent operator’s MDU footprint, using the Google template.

“Gigafy America” became the rallying cry from these new competitors.

Cable MSOs now have to plan a fiber strategy -- or watch these competitors take away new greenfield opportunities, as well as existing MDU and bulk contract, gated community subscribers.

Migrating MSO FTTP Networks to DAA

Competitive Pressure

As many as 25% to 70% of US properties passed by MSOs can be classified as MDUs. The broad definition of MDU includes traditional apartment buildings and multi-use buildings containing commercial businesses, plus residential units and a wide range of gated properties -- such as “over 55,” golf course, and waterfront communities. What makes these MDU properties unique is that, for the most part, the developer of the building or community negotiates a long term, bulk contract with a service provider that typically guarantees 100% subscriber penetration. Service provider decision power for these situations is usually concentrated with the property developer or HOA (Home Owners Association). Many of these owner associations quickly learned to hire technology consultants to evaluate the competitive solutions being proposed, so as to force improvements in data delivery rates and other concessions. Gigabit download speeds are an attractive incentive to high-end consumers, who already use lots of internet-ready devices, in almost every room. Plus, “fiber-ready” buildings allow developers to raise the selling price of each condominium as “future proofed” homes. As a result, most developers began to demand fiber-only construction for their new greenfield properties.

New, startup fiber service providers seem to come into existence every week, competing for both new building developments and existing MDU contract renewals. And, for the first time, these upstart overbuilders are not competing with increased video content offerings or HD picture quality, but with offers of Gigabit-speed data -- at extremely aggressive pricing. Figure 1, below, shows the different competitors that have expanded into new cities and in some cases challenged existing Comcast properties.

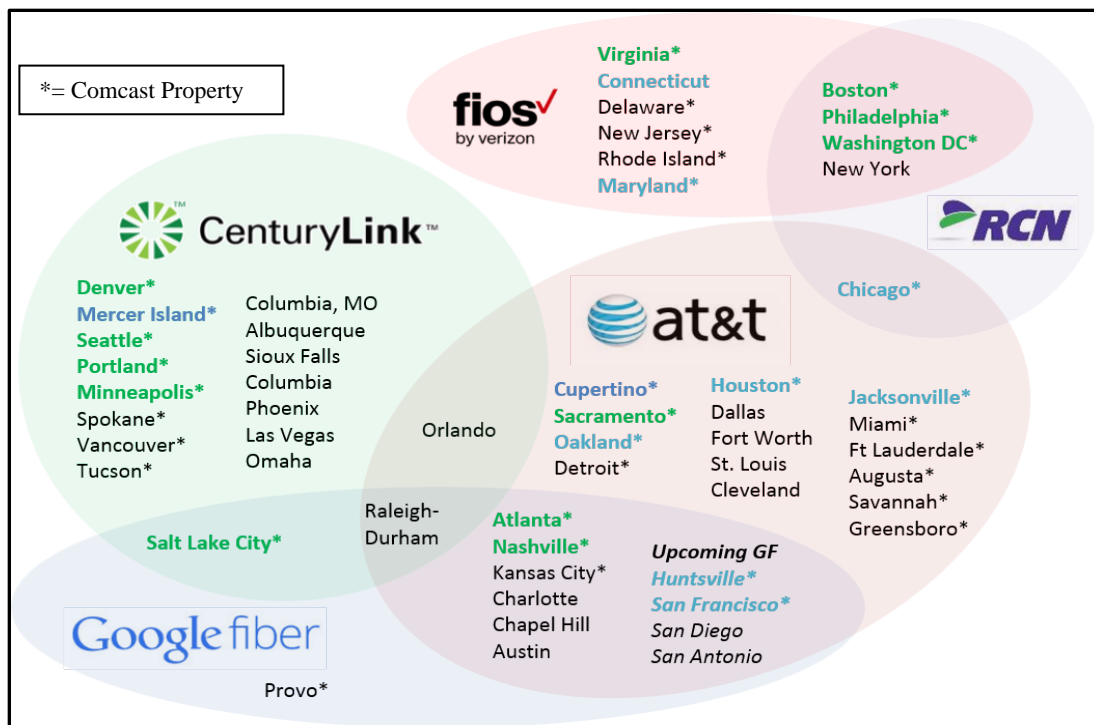


Figure 1 – Fiber Competitors Are Now Challenging Every Greenfield Opportunity

Which PON to Choose?

To meet the demands of greenfield developers for FTTP broadband solutions, Cable MSOs first need to select a PON (Passive Optical Network) operating system. RFoG was created to be the fiber transport equivalent of HFC. RFoG provides the same architecture as PON, while remaining completely transparent to back office systems and subscriber provisioning. The major drawback is that RF over Glass is limited to the same RF capacity as HFC. As a result, RFoG has the same asymmetric DS/US data capacity as HFC and cannot match the symmetric Gigabit speeds that fiber competitors are now offering.

Option 1: Gigabit Passive Optical Network / GPON

Google Fiber and all of the smaller overbuilders selected GPON as their initial architecture. GPON is a mature solution, deployed in volume worldwide, and adopted by all of the major North American telcos, including Verizon and AT&T.

The GPON specification was created by a working group of major telecommunications service providers and system vendors organized as the Full Service Access Network working group (FSAN). The spec was then issued by the International Telecommunications Union (ITU) as the ITU-T G.984 standard. The GPON standard allowed for larger, variable length packets than previous FSAN specifications, to provide 2.488 Gbit/s of downstream bandwidth and 1.244 Gbit/s of upstream bandwidth. This available data capacity is minimally sufficient to support a Gigabit symmetric HSD tier rate, depending on the PON architecture split ratio and the level of congestion on the network.

However, GPON has other limitations that impact MSO adoption. High among these is the lack of a DPoE mediation layer, which is critical for maintaining a common provisioning protocol across all supported services. While it is possible to create an equivalent DOCSIS mediation layer (DML) for GPON, the silicon and software development effort is estimated to be 12 to 18 months. Without support from the majority of MSOs, the cost and risks of investing in and waiting for DPoG were deemed to be too high. Another weakness is the limited interoperability between vendors. Although GPON is a common standard, many vendors customize the OLT and ONU's to achieve optimized performance as a competitive sales tool. An inability to mix and match vendors within the network results in a substantial barrier to lower costs.

Another hurdle in adopting GPON is the makeup of the FSAN / ITU-T standards body that manages the specification. FSAN is a very telco-centric organization. Many MSOs are members of FSAN, but the major contributors are telcos and the vendors that support them. It would be difficult to get substantial traction within FSAN without significant participation and consensus from a block of cable MSOs.

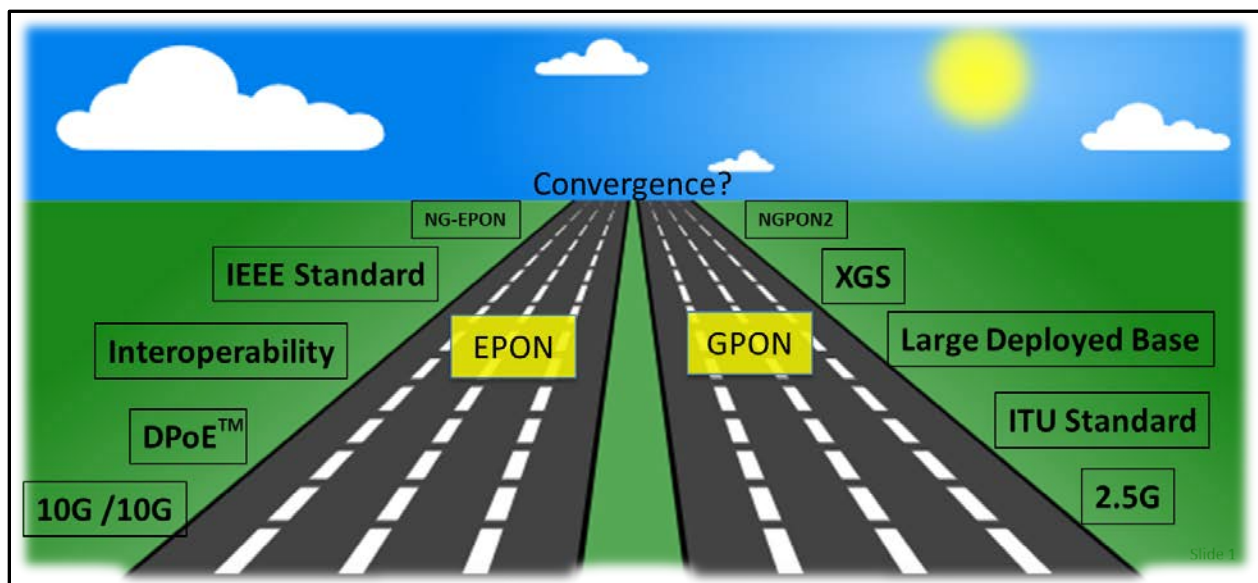


Figure 2 – EPON / GPON Comparisons (Courtesy of Curtis Knittle, CableLabs)

Option 2: Ethernet Passive Optical Network / EPON

Another option is Ethernet PON, developed by the Institute of Electrical and Electronic Engineers (IEEE) -- a leading standards organization that has been very supportive of MSO requirements. The IEEE 802.3ah and 802.3av EPON standards provide 1 Gb symmetric or 10 Gb symmetric and asymmetric options. EPON has also been widely deployed, DOCSIS provisioning is supported via DPoE™, and there is proven interoperability between vendor equipment. The drawbacks regarding the currently available EPON standards are that GEAPON (Gigabit Ethernet Passive Optical Network) is not capable of multi-subscriber gigabit symmetric service, 10G EPON is not widely deployed in North America, and is considered too expensive, primarily due to the higher 10G optics cost.

Selecting which PON technology to deploy is always a balance between cost, operational complexity, and its ability to meet current and future capacity requirements. The rapid progression of high-speed data

consumption, tracked by Nielsen and others for the past 30 years, continues to increase at a compound annual growth rate of close to 50% year over year. At current growth rates, the projected data capacity needed will approach 10 Gb by the early 2020's.

GEAPON and GPON are the lowest cost PON solutions, because of lower data rates and uncooled optics used in optical network units (ONUs). Plus, scale economics play a role, because of annual, worldwide market volumes for ONUs. GPON's downstream data capacity is 2x higher than 1G EPON, giving it a clear advantage. But GPON provides no competitive advantages for MSOs, and based on the previously discussed projections regarding HSD CAGR, the relatively low, 2.5 Gb downstream data capacity will have a limited service lifetime in all but lower tier rate, best effort applications. The GPON standard also specifies a lower, asymmetric data rate for upstream traffic. This significantly reduces the usefulness of GPON in commercial services applications, which usually require symmetric data transport. NGPON2 is a more recent approved ITU standard that is capable of providing 10 Gb symmetric transport, but its use of ONU tunable wavelength filters makes it very expensive, particularly for residential subscriber deployments. As a result, NGPON2 has not been deployed in any operator systems to date.

10G EPON is the most recently deployed IEEE PON technology and has useful capacity to support multi-Gigabit tier rates for both downstream and upstream traffic. 10G is a significant technology leap over GPON, and provides a competitive edge against network overbuilders. It should have a serviceable life of several years based on the current CAGR models.

The current North American market volume for 10G PON lasers is still relatively low. This is one of the reasons for the higher cost of these devices. The announcements that Comcast and other major MSOs are planning to adopt 10G EPON has already produced improved pricing, as optical vendors position their products for the emerging cable operator FTTP business. 10G PON optical transceivers are also mainly available today only as pluggable XFP or SFP+ packaged devices. Conversion to a BOSA-style package, typically used in most 1G EPON and GPON ONU deployments, will help to further lower the cost curve of these transceivers.

Choosing which PON technology Comcast would pursue was a long and arduous process. As detailed above, GPON offered a low cost, mature technology that could immediately defend against the aggressive competitors that were eroding our MDU footprint. 10G EPON had not yet been deployed in a cable MSO network, and other than China, had only been selected for a few commercial installations. On the plus side, 10G EPON's DPoE capability would provide operational similarity with the existing DOCSIS back office, and 4X or higher data capacity than GPON.

In the end, after several months of vigorous internal debate and analysis 10G EPON was selected, based on two primary factors: First, symmetrical 10 Gb allowed the capacity to offer multi-Gigabit HSD for both commercial and residential subscribers, which was not possible with GPON. Second, and even more convincing, was an analysis that took into consideration the growth rate of data in the network. Deploying GPON would result in lower capital costs, but the serviceable lifetime of GPON was limited, due to its lower data rates. The estimated network upgrade costs to eventually replace GPON with a 10 Gb solution was significantly higher than going all-in with 10G EPON on day one.

The Role of PON in Cable MSO Networks

The initial move into FTTP by cable operators was driven by the threat of competition. While some markets will require a fiber to the home solution, the majority of the MSO residential network footprint will be adequately served with DOCSIS and HFC for many years to come. On the business services front, FTTP PON has been discussed for years as the eventual solution, when the available cable plant dark fiber is exhausted.

The introduction of D3.1 has raised new questions regarding the need for FTTP beyond new greenfield and competitive threat applications. MSOs have traditionally met the need for increased capacity by a combination of RF bandwidth expansions and node splitting.

Eliminating analog video channels, digital video compression technology, and higher order QAM modulation formats supported by DOCSIS has allowed operators to keep pace with the BW requirements of HD channel growth and the ever increasing array of new channel content. The arrival of D3.1 and OFDM has the potential of increasing the capacity of the HFC access network to 10 Gb or higher with a corresponding expansion in the RF system bandwidth. Migrating the network to an N+0 architecture paves the way to implementing Full Duplex DOCSIS and the possibility of gigabit symmetric service when it becomes available in a few years. This very real and near term technology will make HFC DOCSIS cable systems the equivalent of a 10G FTTP fiber network.

How long will Gigabit-per-subscriber data rates keep pace with the continuing growth rate of data? Maintaining the viability and lifetime of the HFC network is always the main priority. At what point does the crossover occur, between Fiber Deep and FTTP?

If the growth of data continues at its historical rate, an N+0 DOCSIS 3.1 network, supplemented with FTTP to support high end residential and SMB / home office commercial users, may be sustainable for several years. At some point, though, within the next 10 years, HSD growth will begin to require per-subscriber speeds that exceed the bandwidth limitations of balanced, well designed, N+0 fiber-deep D3.1 HFC networks. Fiber to the tap and similar proposed solutions result in an order of magnitude higher number of active devices -- which will be complicated, power hungry, and need to be provisioned and managed. A Fiber Deep N+0 architecture provides the natural demarcation point for the HFC-to-FTTP transition when needed. Changing the node to a modular virtualized OLT location and adding fiber drops, or possibly wireless access transceivers in place of RG6, could likely be the next evolutionary phase for MSO networks.

The next leap in PON technology is currently being debated at both IEEE and ITU FSAN working group meetings. Most likely the optical device advances that are now driving data center transport rates will be adapted for longer link networks. If so, then data rates of 25, 50, and 100 Gb are not very far off from being a reality.

10G EPON Network Architecture Challenges For MSOs

MSOs face a number of design and operational hurdles on the path to creating a workable FTTP playbook. The list below identifies the most significant issues. Many of these will be expanded on in later sections of this paper.

10G EPON FTTP Architecture Challenges:

- Available dark fiber rapidly decreasing
- Link reach >> 20 km
- Silicon dependency (one vendor provides 90% of all OLT chip sets)
- High cost of fiber construction and home wiring
- New Customer Care, Billing, Tools, Operations, Processes
- Fiber Handling (Field Tech training, tools, diagnostic equipment)
- Homeowner and MDU Consultants
- “Gigabit” speed, speed tests, servers, dependencies
- HSD tier offering in mixed DOCSIS and EPON footprints
- External positioning for FTTx (D3.1 versus fiber comparisons)
- Voice and Video integration (Migration to all IP)

FTTP in Cable Plant Realities

HFC networks are extremely efficient. By relying on active elements in the network, the number of customers that can be reached and supported with triple play services (voice, video, and HSD) in the access last mile link provides the lowest total cost-per-subscriber compared to any competitive solution. As a result, cable networks have been built with link reaches ranging from 20 km to over a 100 km, with the lowest amount of available fiber depending on the homes passed density of the geographic area, and individual operator design practices. Access network fiber availability for MSOs is further challenged by the success of business services. Available dark fiber has been used for adjacent market high revenue applications such as cell tower backhaul and commercial metro Ethernet links, reducing the amount of fiber accessible for a FTTP migration.

Classic PON architectures, whether centralized or distributed, require an extensive amount of fiber. The reference optical distribution network (ODN) design for FTTH is a 20 km passive fiber link to reach 32 subscribers. Higher split ratios -- to increase the number of subscribers per OLT port, connector attenuation loss, and other design margin factors -- reduce the maximum reach of the PON link. Figure 3 illustrates the link reach possible with a completely passive optical network.

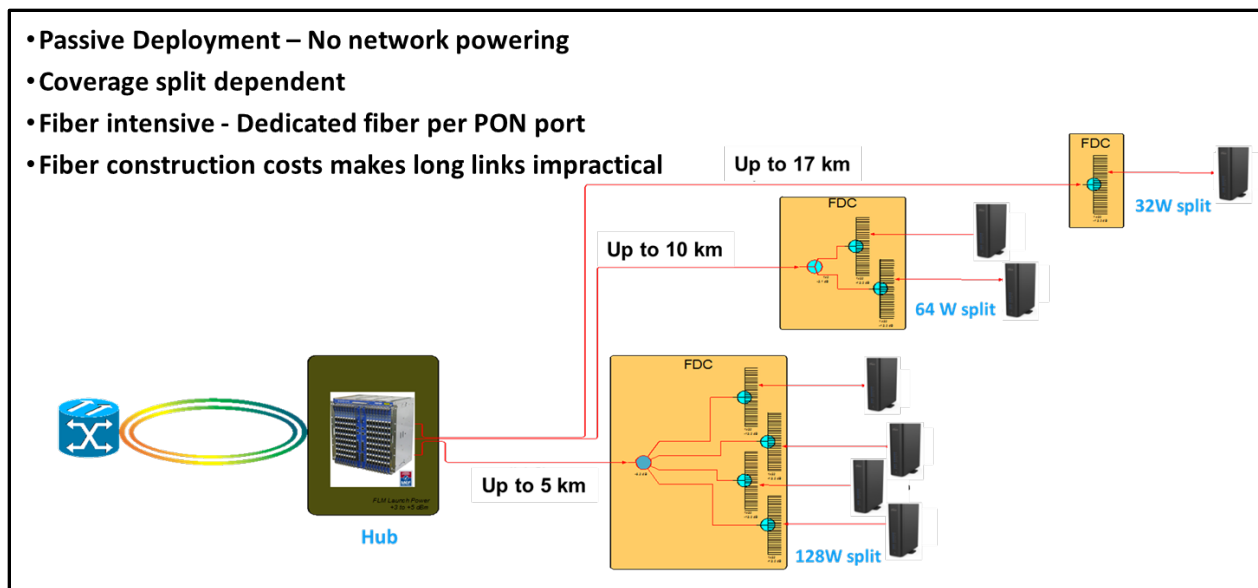


Figure 3 – FTTP Deployment Direct Feed Coverage

Using traditional PON ODN construction, most customers can only be reached by reducing the splitter ratio. This means few subscribers per port, and therefore much higher OLT port counts and cost per connected subscriber. Fiber construction represents one of the highest network deployment costs, particularly in existing legacy brownfield applications. Therefore most cable operators target FTTP primarily for greenfield applications, with the exception of a legacy serving area that is under competitive threat.

While the material and installation costs of fiber and coax are almost identical in new, greenfield construction cases, there are still major differences. Fiber splicing and connectorization is the most significant. Fiber cable is available with several standard fiber counts, from 12 fibers to 288 fibers. Splicing each of these individual fibers requires special equipment and training. Each fusion splice can cost approximately \$25, depending on local construction rates in the build area. In the case of a 288-count fiber bundle, fusion splicing each strand can take several days to complete and verify. Splicing connector pigtails onto a fiber bundle at a distribution cabinet or field splitter is an equally expensive and time-consuming endeavor. Ribbon fiber and multi-fiber MPO connectors are being considered as a means to reduce the time and labor needed by these large-scale optical connections. The drawback is the service impact implications of repairing or replacing a multi-fiber connector.

HFC plant design has typically been 60% aerial construction and 40% underground. In new greenfield properties, 100% of the construction is underground. Underground construction costs can be 2X to 3X higher. This includes not only trenching the fiber, but also installing buried chambers for fiber field connections and pedestal fiber management.

Fiber drop cables to the home represent another significant cost premium to HFC deployments. HFC drop cables are typically limited to 150 feet maximum, due to the high RF attenuation of RG6 coax. As a result, HFC taps are distributed along the access link and sized to provide connections for 2, 4 or 8 homes depending on the home density in the particular serving area. When a new customer requests service, the drop cable usually only has to be trenched across one, or at most 2, building lots. Fiber drop cables, on the

other hand have extremely low optical attenuation, and can span a kilometer or more depending on the location of the field splitter or fiber distribution cabinet (FDC). In this case, the drop cable must be trenching across several lots, which increases the drop's construction cost, not to mention the public relations issues caused by construction crews digging up people's yards.

Home wiring is the next major challenge for FTTP construction. In many new greenfield construction sites, and particularly MDUs, a media panel is usually built-in at a central location within the home. It aggregates all of the fiber and RF distribution connections, and provides a storage location for the ONU / Gateway. Unfortunately, the range of variations from one development to another is wide. There are also significant differences in opinions regarding which party owns or is responsible for installing the media panel and fiber wiring.

The RFoG Overlay

Today, FTTP EPON is primarily used only to deliver high-speed data. In order to provide the same triple play experience over FTTP that subscribers receive with the RF coaxial system, an HFC or RFoG overlay is required. HFC overlays are most likely in cases where a legacy, bulk MDU property is being converted to FTTP, triggered by a competitive contract renewal situation. In the majority of greenfield cases, a fiber-based transport solution is mandated, which means an analog / QAM RFoG overlay must be deployed.

An RFoG overlay is primarily needed to support analog and QAM DOCSIS video transport. The increased complexity and cost of multiple CPE devices needed for each fiber-connected home is a significant deterrent to the growth of FTTP in MSO networks. As shown in Figures 4 and 5, an RFoG overlay implementation for FTTP EPON requires additional fibers and OSP equipment. OBI mitigation is also needed in most cases.

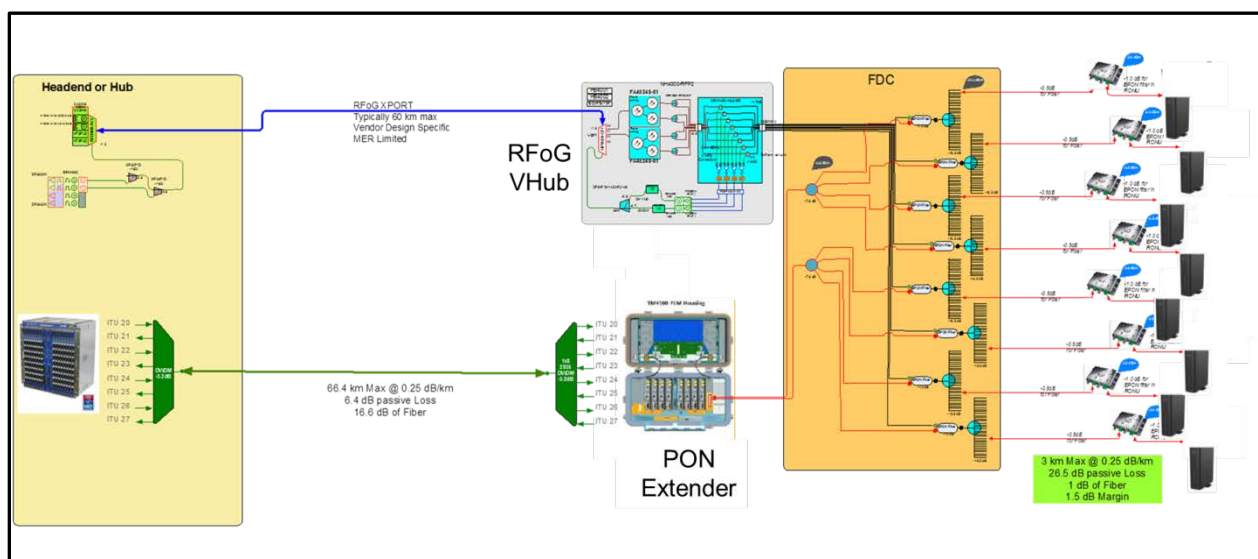


Figure 4 – RFoG / EPON Overlay Design Configuration




Figure 5 – RFoG / EPON Overlay Actual MDU Deployment

An RFoG Vhub is used to amplify the DS wavelengths, in order to compensate for the PON splitting losses. The Vhub also provides for aggregation and OEO analog to digital Ethernet conversion of the RFoG ONU US signals.

In an RFoG overlay, the RFoG and EPON wavelengths are muxed together on the access drop fiber at the fiber distribution cabinet. The RFoG ONU must support an added pass through optical coupler to allow the EPON wavelengths to be coupled from the RFoG ONU to the EPON gateway.

EPON is ultimately capable of duplicating every service flow provided by an all-QAM, DOCSIS provisioned system. The issue is updating, testing, and validating the OLT / ONU software to provide not only the IP packets containing voice, video, and data information, but also a transparent back office operation that is compatible with the existing CMTS. DpoE is only one condition of this compatibility. Other features and service flows within the DOCSIS-provisioned network also need to be replicated for the PON environment. At this time, 10G EPON equipment providers and MSO software developers are working to complete the qualification and field trials that will make a deployable, all IP solution possible by the end of 2017.

Triple play IP transport will not only allow for the elimination of analog RFoG overlays, but is also a necessary step in eventually migrating to a SDN/NFV network for both D3.1 HFC and EPON FTTP.

Cable Plant Designs Are Not PON Friendly

HFC networks consist of a headend, supported by several hub locations, arranged in a star or ring configuration, that serve the local population centers of subscribers. These primary hubs contain the CMTS, local channel content, ad insertion and optical / RF access node links for serving areas ranging from 60K to over 100K homes passed. As new communities developed or expanded, they sometimes exceeded the capacity of the existing primary hub. Smaller, secondary hubs have been added, by leasing space, or building a small facility sized for the new community being served. These secondary hubs

mainly support the access edge for 15K to 30K HP. This system of headends and hubs results in over 85% of all node and subscriber locations being within 20 km of the nearest hub.

While this proximity to a local hub seems to be ideal for PON architectures, the complication is that many secondary hubs are too small to support the rack space needed for a mainframe OLT. Also, because the secondary hub is sized to serve a smaller number of homes, the power and HVAC capabilities of the hub are not adequate to support the added load of the OLT.

Secondary hubs are linked to the main hub or headend via a primary and secondary fiber trunk connection. The problem is that one or both of these link paths, when added to the subscriber access fiber link, may exceed the PON ODN optical budget. This is illustrated in Figure 6, which represents an actual FTTP trial deployment case.




Figure 6 – Secondary Hub Limitations Can Create Extended Link Reach Situations

Although the MDU example in Figure 6 is only 9 km from the local hub, the OLT was located at the primary hub, with access to a larger number of potential FTTP subscriber sites. The additional path distance of the primary and secondary fiber trunks exceeded the PON optical power budget. Additionally, in this RFoG / EPON overlay case, the number of fibers needed for RFoG DS/US links, plus the EPON links for two buildings, exceeded the available dark fiber. This trial proved the need for a distributed access solution, and accelerated the qualification of a 10G capable PON Extender.

PON Extenders – An Interim Distributed Access Solution

A PON Extender is essentially an OEO repeater. Extenders are typically designed as strand-mount, clamshell node housings that can be cable plant-powered. Most extender designs usually support up to 8 OLT output ports. The traditional mainframe OLT PON optics are replaced with ITU grid DWDM optics, which can be muxed onto a single fiber and transported much further than the 20 km limitation of PON optics. DWDM optical transceivers are widely available, with reach capability extending up to 80 km. At the PON Extender side of this link, the DWDM wavelengths are de-muxed, received and then converted to PON wavelengths. The output of the PON Extender is the same as the output of the OLT.

PON Extenders do not manipulate the data or provide any management functions. They can provide basic diagnostic information on the status of the unit, and the optics in particular. The number of ports used can be sized to the property being served.

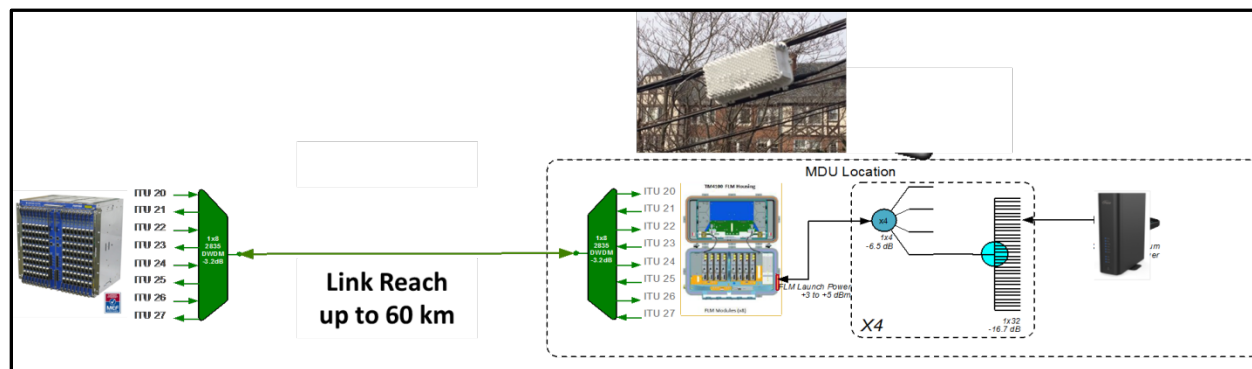


Figure 7 – 10G / 10G EPON Extender Network Configuration

The major advantages of the PON Extender are significantly increased fiber utilization -- up to 8 OLT port links over a single fiber, plus extended reach. The lack of MAC functionality substantially reduces the total power consumption of the unit. The node housing and basic OEO physical layer function make PON Extenders operationally familiar for cable systems engineers and field support technicians.

The main disadvantage of PON Extenders is that they do not eliminate or reduce the OLT equipment requirements in the Hub. The extender is basically the cost equivalent of a second line card needed to reach the same number of subscribers.

Evolution of Distributed Access for FTTP EPON

The driver for distributed access in FTTP EPON networks is identical to the driver for distributed access in HFC networks -- improved scalability. To noticeably increase the data capacity per subscriber in HFC networks, the RF bandwidth available per subscriber must be increased. There are a number of ways to accomplish this: Increasing the raw RF bandwidth, node splitting, or reduced amplifier cascades (i.e. redesigning the network from N+5 to N+0). The consequence of increasing the number of nodes, and thereby reducing the number of subscribers sharing the node bandwidth, is the impact to the hubs supporting those nodes. Each new node requires additional CMTS ports, optical transmission lasers and receivers, plus rack space for signal distribution equipment and connecting cables. Additionally, floor space, power consumption, backup power generators, and HVAC capacity quickly come into play. The alternative -- building new hubs to support the racks of new equipment -- is a long and very expensive undertaking.

The situation with FTTP PON is only slightly different. Mainframe OLT port counts are usually 8 per line card, for a total of between 80 and 112, depending on the vendor. The power consumption of a fully loaded OLT is several kilowatts, and, similar to the HFC scenario previously described, each OLT needs backup power, fiber management and a controlled environment. At a split ratio of 128 subs per port, a

112-port OLT can serve over 14K customers; at 32 subs per port, the number of customers served drops to 3.5K. The impact to the hub facilities and the OSP fiber plant is equally dramatic.

10G EPON has enough data capacity to support multi-Gigabit service to a maximum of 128 subscribers per port, assuming normal contention factors for residential and SMB users. The problem is reaching those subscribers. The majority of PON applications over the next several years will serve non-contiguously located greenfield properties with a typical HHP size of 300 units. Almost all of these applications will be greater than the 5 km link reach limit for a 128-way PON split ratio. Most links will reach beyond 20 km from the primary hub, and almost all will have limited dark fiber available. PON Extenders, as discussed earlier, can bridge the gap -- but at a higher incremental cost per subscriber. Distributing the OLT edge network closer to the subscribers being served solves many of the issues associated with traditional FTTP networks.

The Benefits of FTTP Distributed Access

A Distributed Access Architecture (DAA) for FTTP has many of the same goals and follows the same trajectory as HFC Remote PHY or Remote MAC-PHY. DAA accomplishes the following:

- Dis-aggregation of the data, management, and control planes within the OLT platform
- Reduction of hub rack space and powering requirements
- Distribution of the PHY edge closer to the subscriber
- Provides for flexible deployment sizes
- Increased link reach and fiber utilization
- Interoperability between vendor solutions
- Provides a path to network function virtualization (NFV)

One significant distinction for FTTP distributed access is that it is impractical to completely separate the MAC and PHY functionality of the current OLT silicon. Therefore, a remote OLT design will be similar to MAC-PHY system architectures.

Simply scaling down the number of line cards and rack unit size of a mainframe OLT platform and moving it out into the field is not an efficient or network-safe means of creating a distributed access network. The backplane switching fabric of a mainframe OLT is designed for high port counts, and is not readily scalable for low homes-passed applications, such as serving an MDU with 250 apartments/condos. The existing Ethernet switch silicon is also one of the highest power consumption components within an OLT, so designs based on repackaging an existing line card will typically generate powering requirements of 200 watts AC or more. This is much higher than the cable plant network power grid can support without adding more power supplies to the network, and is definitely beyond the thermal capacity of the largest node housings that can be accommodated on the strand today. Rack-mountable outside cabinets with internal power supplies are a possible option, but with the higher power consumption of these designs, some level of environmental conditioning is typically needed.

The bigger issue with a downsized mainframe OLT platform for remote applications is that all of the provisioning, security, and control functions still reside on board, which dramatically increases operational risks. Maintaining this type of remote design will also require a much higher skill set level and diagnostic equipment tools than the typical field technician has today.




Figure 8 – 10G EPON Remote OLT All IP Triple Play

The target design for a remote OLT is a strand-mountable, node-housed device. Node housings are self-contained environments: Weather proofed, RFI shielded, and capable of being located in the building, in a pedestal cabinet, or on a pole (as originally intended.) R-OLT nodes are designed to operate on cable plant power over an extended ambient temperature range of -40C to +60C.

Creating a node based OLT that is compatible with existing cable plant guidelines and capable of operating in outdoor, uncontrolled environments is a unique challenge for legacy PON equipment suppliers with limited or no HFC plant experience. Newer switch silicon provides lower power consumption and the ability to power down unused ports. The restrictions on total power consumption, combined with the need to meet safe operating temperature limits for 10G optics and PON MAC silicon chip sets, limits the number of 10/10 EPON ports that can be supported to only four, in most cases. While this may appear to be constraining, the ability to place a node-based remote OLT close to the property being served allows 128 subscribers to be connected, per port, or 512 subscribers per node. This is more than enough to cover 90% of typical MDU situations. Larger properties, such as new, single family home developments, would need multiple nodes -- but proximity to the subscribers would reduce the amount of drop fiber construction needed.

By disaggregating the data, management, and control planes of the OLT, these functions can remain at the headend / hub with DpoE / vCM clients at the R-OLT. The R-OLT can then be designed primarily as a layer 2 device, with limited layer 3 functions as required.

Virtualizing the FTTP Network

The ultimate goal for distributed architectures is network function virtualization (NFV) and software defined networking (SDN). Virtualizing the network will allow the use of common, interoperable hardware, with open software, to promote agile development and faster innovation. Under this virtual operating system, all subscribers can receive the same quality of experience -- whether they are connecting through an R-PHY HFC node, or a Remote OLT.

SDN separates data, control and management planes to enable an open architecture with standard interfaces and protocols. NFV decouples network functions from proprietary, built-in hardware, to

provide distributed software functionality that runs over common, multi-sourced hardware. Cloud architectures shift network and operations intelligence into a centralized platform, providing cloud-based applications and service models.

SDN's open architecture provides an end-to-end orchestration and controller platform of both virtual and physical network components, that can be integrated dynamically and programmed automatically, based on service and network demands. Automation combined with programmability enable service agility. In addition, new services and deployments may be introduced with faster time, less investment and reduced cost.

An efficient NFV design for access networks requires breaking software into modular components that correspond to workflows -- that can be distributed at the edge, core or cloud, and coordinated by a centralized SDN orchestration platform. This can be achieved by using containerized software to implement micro-services, as opposed to traditional and monolithic systems. This approach provides smaller and simpler SW subcomponents, which can be created, added and modified independently, using agile development and continuous delivery techniques. These software subcomponents can then be implemented, tested and deployed to introduce new services and network functionality faster, with the help of efficient cross-functional DevOps teamwork.

Cloud-based platforms with distributed functionality provide agility and flexibility of both service and network operations, as wider communication pipes are controlled by a centralized intelligence architecture.

A centralized controller manages both modular and distributed virtual components as well as physical components serving smaller subscriber serving groups. This removes single point of failure issues that can disaffect a large domain, in traditional systems. It reduces CAPEX as the deployment can be scaled over time based on real-time and on-demand requirements. New features are introduced by deploying new software components on the same hardware platform without depending on vendor-specific releases with long lead times. Virtual components can be moved or scaled based on performance, network, and high availability constraints.

This centralized approach, along with powerful data mining techniques, provides end-to-end visibility, better search, and improved correlation of data. Furthermore, virtual probing, together with probing of physical components, may be implemented for troubleshooting, new service integration, and customer QoE assessment. The deployment of a virtualized EPON architecture is flexible and scalable, compared to static, rigid and mass rollouts. New workflows could be introduced as NFV micro-services during trials and deployments to introduce new services, and modify or scale existing services.

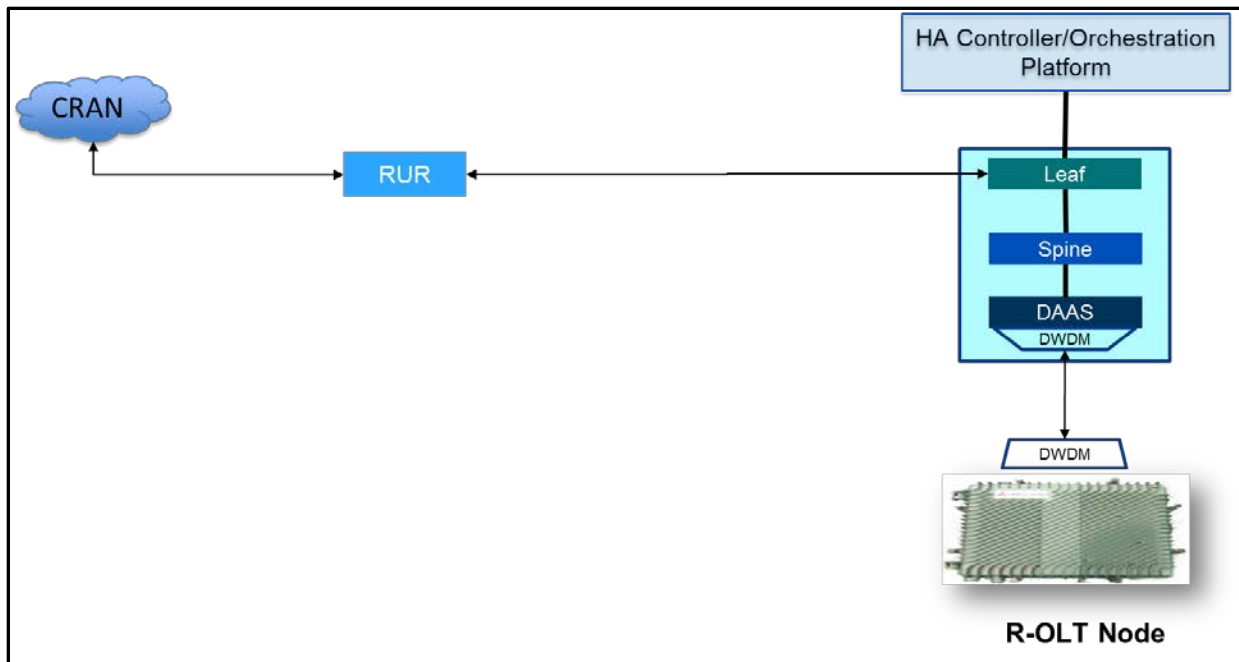


Figure 9 – Virtual OLT Configuration

Figure 9 above illustrates the implementation of a remote OLT in a virtualized network. The remote OLT PON ports are transported over 10 Gb optical links, muxed onto a single fiber and combined at the hub aggregation switch. The controller / orchestration platform directs the flow of data between the regional access router and the R-OLTs connected to the hub, while monitoring performance and policing security, provisioning, and operational policies.

Conclusion

Increased market competition and the continuing growth of data consumption is driving MSOs to deploy networks capable of Gigabit-per-subscriber throughput -- necessary to defend their subscriber footprint against fiber-based overbuilders. Greenfield housing developments, and, in particular, MDUs are being aggressively targeted by PON competitors, prompting cable operators to respond with a fiber to the premise solution.

The applications for FTTP in cable footprints for the next few years will be limited to new greenfield opportunities, entertainment venues, and commercial entities, like small to medium businesses. DOCSIS 3.1, enhanced by FDX, will remain the preferred architecture for legacy, residential, brownfield networks. In order to lower the cost of FTTP construction to a level comparable with HFC, an all-IP, triple-play-capable solution must be deployed to replace the current RFoG overlays and to migrate to a more distributed architecture.

FTTP distributed access architectures provide the same network advantages as remote PHY, in traditional coaxial HFC systems. The current development and eventual deployment of R-OLTs will distribute the

PHY edge closer to the subscriber, enabling increased link reach, higher fiber utilization, and flexible serving group sizes.

Distributed access technology also enables the transition to network function virtualization and a more cloud-based network. Virtualizing the network breaks the long-standing model of bookended, proprietary link equipment, permitting the use of interoperable hardware and open software.

A flexible network that can transparently integrate both HFC and PON subscribers -- while maintaining the exceptional quality of experience of DOCSIS -- may be the best defense against the new generation of fiber competitors.

Abbreviations

BOSA	bidirectional optical sub-assembly
BW	bandwidth
CAGR	compound annual growth rate
CMTS	cable modem termination system
CPE	consumer premise equipment
CRAN	Comcast Regional Access Network
DAA	distributed access architecture
DAAS	distributed architecture aggregation switch
DML	DOCSIS mediation layer
DOCSIS	data over cable system interface specification
DPOE™	DOCSIS provisioning of Ethernet
DPOG	DOCSIS provisioning of GPON
DS	downstream
DSL	digital subscriber line
DWDM	dense wavelength division multiplex
EPON	Ethernet passive optical network
FDC	fiber distribution cabinet
FDX	full duplex DOCSIS
FiOS	Fiber Optic Service
FSAN	Full Service Access Network (working group forum)
FTTH	fiber to the home
FTTP	fiber to the premise
Gb	gigabit
GEAPON	gigabit EPON
GPON	gigabit passive optical network
HA	high availability
HD	high definition
HE	headend

HFC	hybrid fiber-coax
HOA	home owners association
HSD	high speed data
HVAC	heating, ventilation and air conditioning
IEEE	Institute of Electrical and Electronic Engineers
IP	internet protocol
ITU	International Telecommunications Union
ITU-T	ITU - Telecommunications Standardization Sector
km	kilometer
MAC	media access control
MDU	multi-dwelling unit
MPO	multiple fiber push on/pull off (optical connector)
MSO	multiple system operator
NFV	network function virtualization
NGPON2	Next generation passive optical network 2
OBI	optical beat interference
ODN	optical distribution network
OEO	optical-electrical-optical
OFDM	orthogonal frequency division multiplexing
OLT	optical line termination
ONU	optical network unit
OSP	outside plant
PON	passive optical network
POTS	plain old telephone system
QAM	quadrature amplitude modulation
QoE	Quality of Experience
RBOC	regional bell operating companies
RFI	radio frequency interference
RFOG	RF over glass
RUR	Regional Universal Router
SDN	software defined network
SFP	small form factor pluggable
SMB	small - medium business
SW	software
US	upstream
vCM	virtual cable modem
Vhub	Virtual hub
XFP	10 gigabit small form factor pluggable

Bibliography & References

Nielsen's Law of Internet Bandwidth; <https://www.ngroup.com/articles/law-of-bandwidth/>

2016 Spring Technical Forum; Predictions On The Evolution Of Access Networks To The Year 2030 & Beyond, Tom Cloonan, <http://www.nctatechnicalpapers.com/Paper/2016/2016-using-docsis-to-meet-the-larger-bandwidth-demand-of-the>

ITU G.984-1 Gigabit-capable passive optical networks; International Telecommunications Union, <https://www.itu.int/rec/T-REC-G.984.1/en>

IEEE P802.3ca 100G – EPON Task Force; *Institute of Electrical and Electronic Engineers*, www.ieee802.org/3/ca

An Architecture for Distributed EPON Access

A Technical Paper prepared for SCTE•ISBE by

Kevin A. Noll

Sr. Director
Tibit Communications
Petaluma, CA
kevin.noll@tibitcom.com

Steve Burroughs

Lead Architect
Cablelabs
Louisville, CO
s.burroughs@cablelabs.com

Brionna Lopez

Security Engineer
Cablelabs
Louisville, CO
b.lopez@cablelabs.com

Introduction

Distributed Access Architectures (DAA), and Centralized Access Architectures (CAA), have been debated in the industry for many years. Over the last two or three years, though, the advantages of DAA have overshadowed centralized architectures. As a result, cable operators are beginning to deploy DAA for DOCSIS.

The technical aspects of DAA for DOCSIS have been well treated in standards and literature. Products implementing RemotePHY and RemoteMAC/PHY are available in the market today. Some suppliers have included in their designs an SDN-based system that virtualizes portions of the overall DAA-for-DOCSIS solution.

Similarly, the benefits and challenges of FTTx and EPON have been well treated in the literature. The result has been a move by many operators to strategically deploy FTTx and most have chosen to use EPON in those deployments. However, relative to DAA for DOCSIS, little attention has been given to deploying FTTx in a distributed architecture.

This paper will describe a disaggregated architecture for EPON in an MSO network using concepts from the widely discussed distributed access architecture. The architecture will include separation of the management plane and data plane components and describe how they interact. We will discuss required functionality and how SDN and NFV, and breakthroughs in EPON technologies are key enablers of a distributed EPON network.

Reference Architectures in the Industry

Network and system architectures are a dime a dozen in today's industry. Many claim to address network and system architecture in general, but usually we will find that each is focused on solving a specific set of problems. Nevertheless, it is wise to survey those architectures to ensure we don't duplicate prior works and to take advantage of the findings in those previous works. In this section, we survey some architectures that are widely covered in the literature.

1. DAA as a Reference Architecture

Distributed Access Architecture (DAA) is referenced frequently in the literature.

In 2013, (M. Emmendorfer and T. Cloonan, 2013) examined the need to convert from analog to digital modes in the optical portions of the HFC network. The natural conclusion was that a remote PHY or remote MAC/PHY device would be required. This is one of the earliest to note the need and propose the basic architecture for what is now referred to as DAA.

(Emmendorfer, Cloonan, Ulm, & Maricevic, 2014) is one of the first times that the term Distributed Access Architecture is used in the cable industry literature. The authors expanded the analysis of (M. Emmendorfer and T. Cloonan, 2013) by comparing DAA to Centralized Access Architectures (CAA) and further described the architectural foundations of a DAA for an HFC/DOCSIS network. Notably, the authors identify physical locations for various elements (Headend, Node, Subscriber Premises) and decompose the key network elements into their functional layers and components. The CCAP is described as a combination of Upper MAC (Policing Classification, Shaping), Lower MAC and Convergence

(filtering, scheduling, framing), Upper PHY (PCS layer), Lower PHY (PMA layer), and the PMD layer. This decomposition makes possible the authors' proposals to relocate portions of the CCAP system to physical locations outside the traditional MSO headend.

Further, the authors' decomposition of the CCAP/HFC network into its functional components sets the stage for a generalized discussion of a disaggregated access network in the cable industry. Later papers and articles (for example (Bernstein & Ramakrishnan, 2015), (Torbet, Cloonan, & Aftelak, 2016; Torbet, Cloonan, & Aftelak, 2016)), discuss how the component functions of a traditional "big iron" CCAP/HFC network can be split between less complex hardware and software located in data centers, headends, and nodes. This marks the beginning of widespread acceptance of SDN and NFV concepts being applied in the DOCSIS/HFC network.

Even though common roots exist, there is no one definitive architecture that is DAA. In fact, throughout the literature, we find that DAA is more of a concept than a definitive architecture. There are, though, common threads throughout the literature and industry discussions. Key threads are decomposition of the CCAP/DOCSIS elements and relocation of the RF transmit/receive functions much closer to the subscriber than past HFC designs ever thought necessary.

Ultimately, we find that DAA is enabled by several other architectures. Namely, remote PHY, remote MAC/PHY and virtual CCAP.

2. Modular and Distributed Cable Architectures

In (Sundaresan, 2015), the author surveys the various modular and distributed architectures for cable networks. We will not duplicate that work here. Instead we will point out some key re-usable constructs from those architectures.

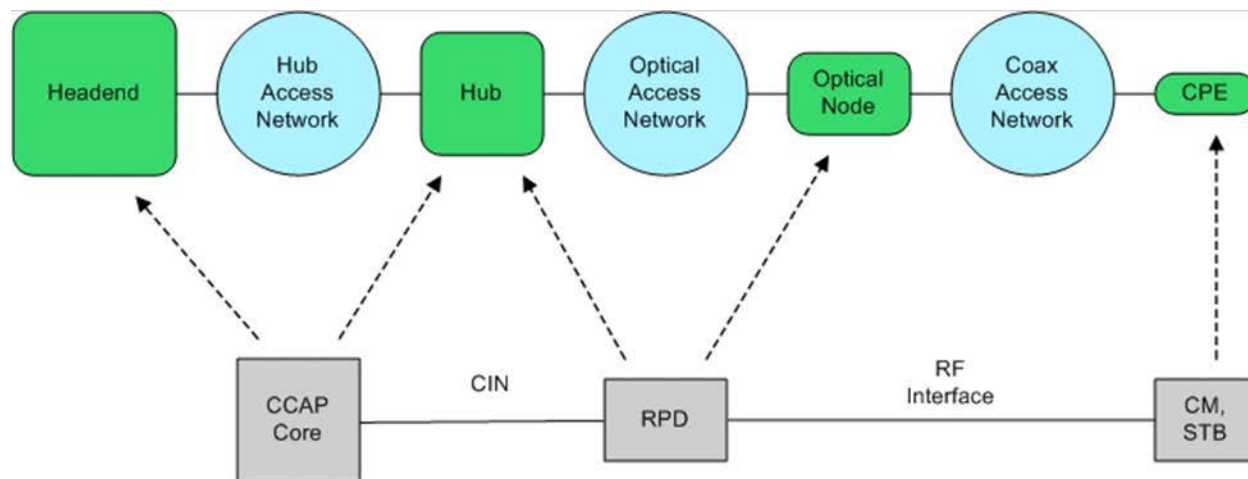


Figure 1 - Remote PHY System Diagram (Remote PHY Specification, 2017)

The R-PHY specification (Remote PHY Specification, 2017) formally describes the structure of a cable operators network and the relative location of the components as they are related to the R-PHY requirements. This foundation for the MHA and R-PHY architecture is shown in Figure 1.

The MHA/R-PHY architecture further describes the interfaces that have many similarities to what is required in the distributed PON network. For example, the Converged Interconnect Network (CIN), which is the network that connects the RPD to the CCAP Core, would need to have an equivalent in the distributed PON architecture.

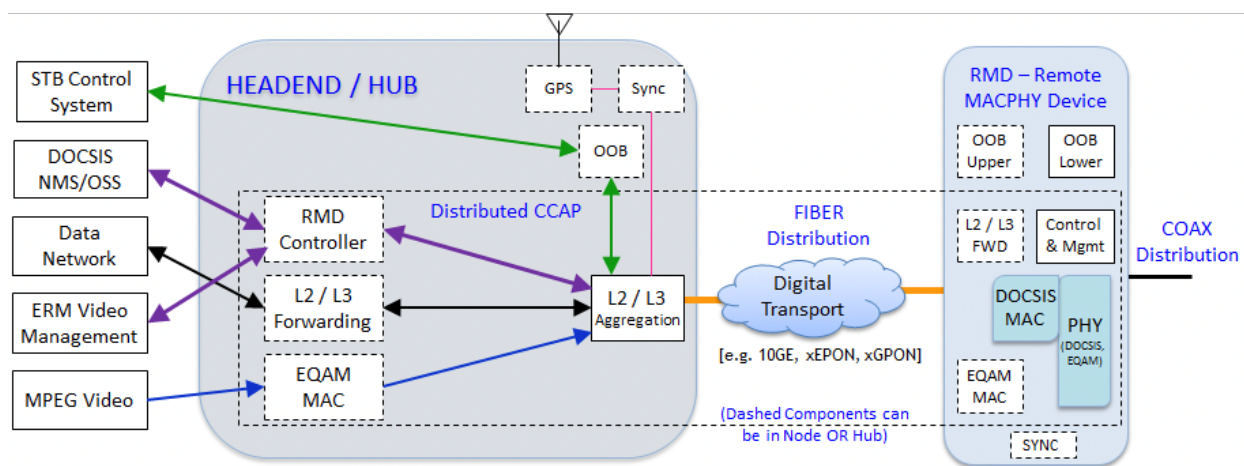


Figure 2 - Remote MAC/PHY System Architecture (Remote MAC-PHY Technical Report, 2015)

The Remote MAC-PHY technical report (Remote MAC-PHY Technical Report, 2015) describes a more complete model, shown in Figure 2, that aligns closely to what might be expected to support a remote PON.

The Remote MAC-PHY model includes elements for the RMD (analogous to a remote OLT), L2 aggregation in the hub/headend, a control function for the RMD, and digital transport (the CIN in R-PHY). The model includes security mechanisms that are critical for operation in a potentially hostile remote location and the model addresses distribution of synchronization and timing which are important for service offerings like wireless backhaul and other commercial and carrier services. This model could be a good example to follow for the physical topology of distributed EPON.

Missing so far, though, is an architecture that describes in sufficient detail the disaggregation and possible virtualization of the control-plane and management-plane functions.

3. DPoE as a Reference Architecture

No discussion of PON in a cable network would be complete without including DOCSIS Provisioning of EPON (DPoE). DPoE v1.0 is specified in a series of nine documents and DPoE v2.0 is specified in a second series of nine documents. Of interest to our analysis is the DPoE v2.0 architecture in (DPoE Architecture Specification, 2016). The DPoE architecture, pictured in Figure 3, is the first time that we find a virtualized network function – namely the vCM – in cable industry specifications.

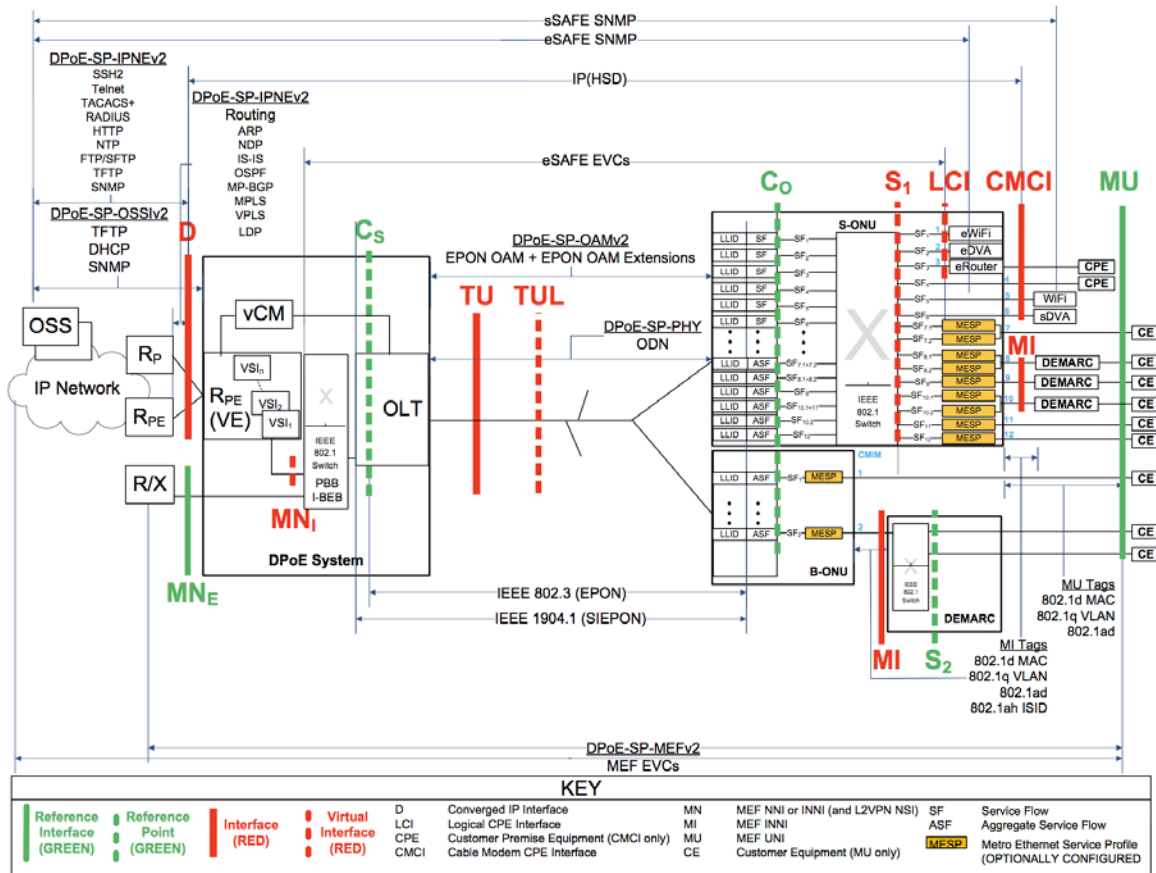


Figure 3 - DPoEv2.0 Reference Architecture, Interfaces, and Reference Points (DPoE Architecture Specification, 2016)

The DPoE specifications, however, do not describe a method by which the OLT can be physically disaggregated from the DPoE System (an extension of the Cs reference point definition). In other words, DPoE practically requires that the DPoE System (analogous to the DOCSIS CMTS) be a full-function system housed in a monolithic chassis. This means that DPoE cannot be an essential basis for the distributed EPON architecture.

It is, however, essential in our proposed architecture that DPoEv2.0 functionality be maintained. Therefore, we cannot dismiss the DPoE architecture from the list of considerations for the distributed EPON architecture.

In fact, a distinct possibility for developing a complete distributed EPON architecture could be to further define the interfaces that would enable the DPoE system to be disaggregated. This was one of the major topics of the Virtual Provisioning Interface Technical Report (VPI) (Virtual Provisioning Interfaces Technical Report, 2017). Section 9 of VPI describes this aspect of the DPoE architecture and calls out several possible solutions, but does not settle on any single solution. Nonetheless, VPI serves as a very good reference for defining our distributed EPON architecture.

4. Non-Cable Reference Architectures

It is important to recognize that the telecommunications and networking industry is a much larger community of which the cable industry is a subset. This means that the cable industry, and any proposal for new architectures, should survey non-cable specifications and standards with the expectation that we will find helpful input there. In this section, we look briefly at three architectures that are frequently referenced in the telecommunications and networking industry.

4.1. IEEE 1904.1

IEEE Std1904.1 (SIEPON) (IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON), 2013) is very familiar to the cable industry because it was originally derived from the DPoE specifications. This adopts the interoperability mission that DPoE sought out in for the cable industry and expands on it by defining two other profiles for OAM messaging to the ONU. SIEPON adds additional requirements and features (such as service availability functions, and PON protection mechanisms) into the EPON system and does not concern itself with the DOCSIS-specific functions that DPoE defines.

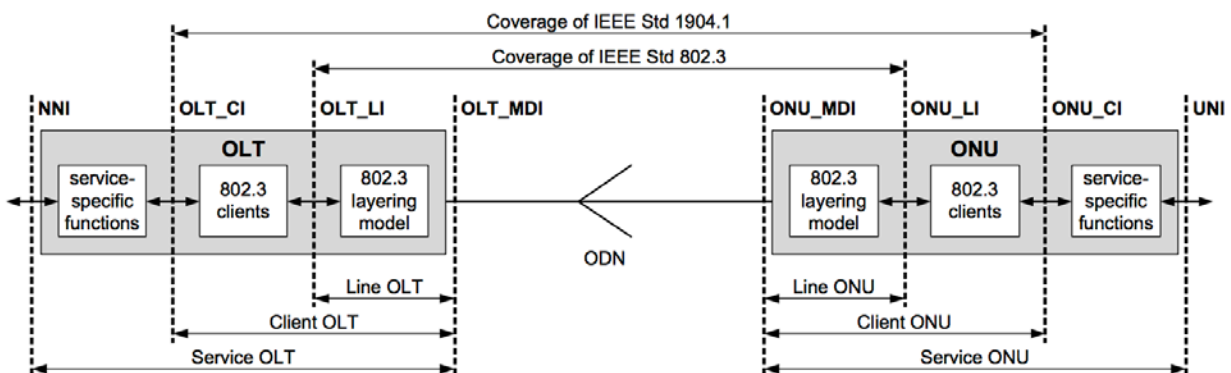


Figure 4 - SIEPON Target System Architecture with Service-Specific Functions (IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON), 2013)

SIEPON also defines features and functions in much finer detail and structure than DPoE. In so doing, SIEPON defines a scope of the system being specified. Shown in Figure 4 is the scope of the SIEPON system architecture.

Note that SIEPON does not define service-specific or system-level functions like DPoE does (e.g. IP routing, DHCP relay, provisioning system interfaces, etc.). As can be seen in Figure 5, SIEPON focuses primarily on the functions and interfaces occurring at the MAC layer and the layers immediately adjacent to the MAC layer (OAM, Data Link, MAC Control, MAC Client).

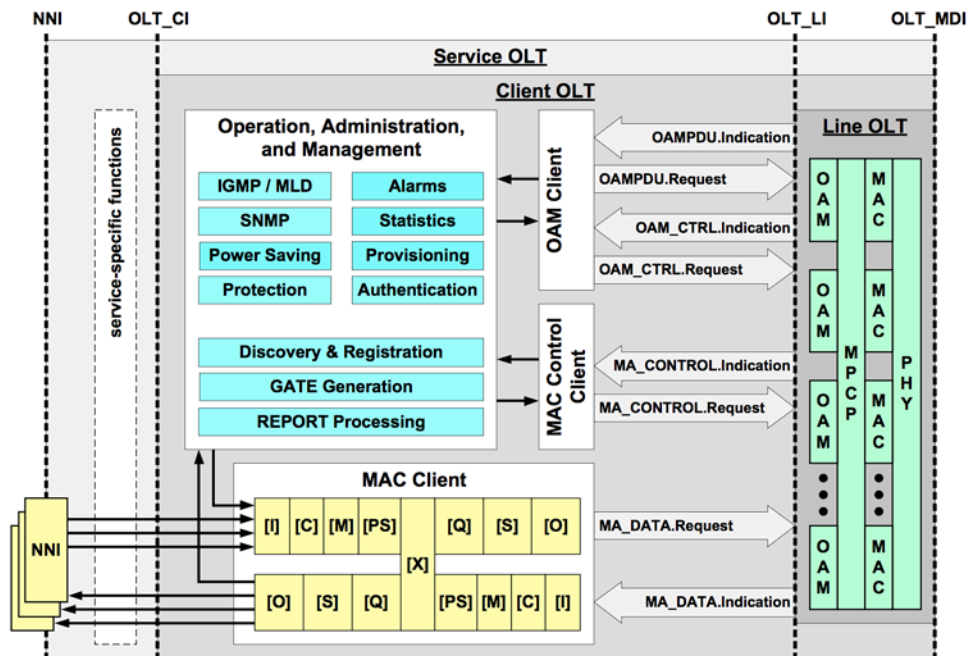


Figure 5 - SIEPON OLT Architecture with single L-OLT (IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON), 2013)

We note from Figure 5 that SIEPON breaks the OLT system down to its basic function blocks required for operation and interoperability. This is of interest to the distributed EPON architecture as it can help guide the distributed EPON architecture in where and how functions can be physically and/or logically separated and still maintain the required functionality.

4.2. Broadband Forum

The Broadband Forum (BBF) has a long history of architecture for access networks. TR-001 (TR-001 ADSL Forum System Reference Model, 1996) the base architecture for an ADSL access network and included abstract definitions of an access node (AN), and the interfaces between each element of the access network, the subscriber and the upstream core networks. BBF documents that followed have adapted this architecture to accommodate changes in technology, additional features and requirements, and the maturing of broadband access networks in general.

What we find now, in TR-200 (TR-200 Using EPON in the Context of TR-101, 2011) is a well-defined architecture that includes EPON in the suite of supported access technologies.

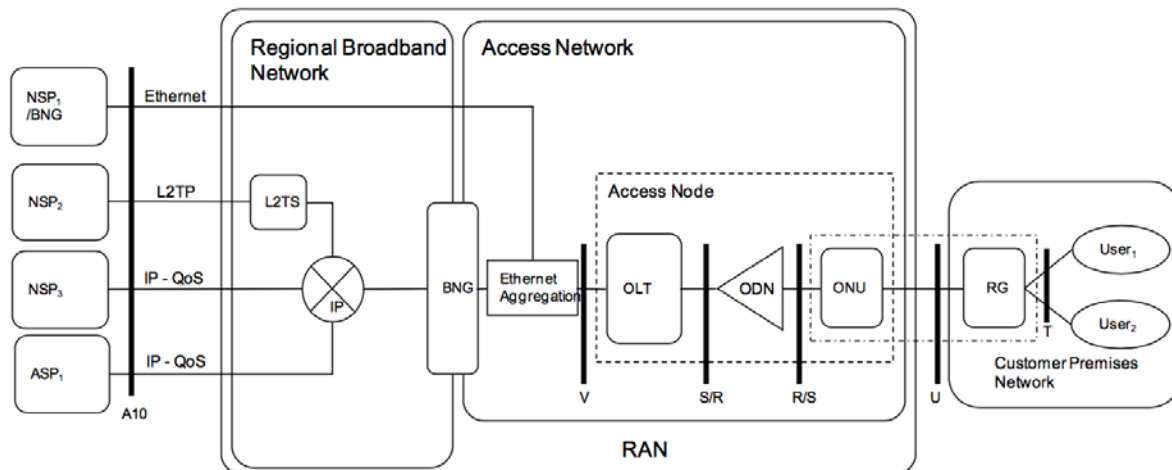


Figure 6 - TR-200 Network Architecture in the case of EPON Access (TR-200 Using EPON in the Context of TR-101, 2011)

The TR-200 architecture for EPON access, shown in Figure 6, inherits from TR-101 (TR-101 Migration to Ethernet-Based DSL Aggregation, 2006), TR-059 (TR-059 DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services, 2003), and their antecedents and successors, the definitions of the network elements like the Access Network, Customer Premises Network, BNG, Access Node, OLT, Ethernet Aggregation, Residential Gateway (RG), and the interfaces between these – the V, S/R, U and T interfaces.

Another line of development in the BBF that continues is a framework for virtualization of network functions in the BBF architectures. TR-345 (TR-345 Broadband Network Gateway and Network Function Virtualization, 2016) develops a framework under which the broadband network gateway (BNG) can be implemented in software as a virtualized network function (VNF) instead of as a monolithic system in hardware. TR-359 (TR-359 A Framework for Virtualization, 2016) moves a far step beyond by developing a framework in which any network function (NF) in the TR-178 (TR-359 A Framework for Virtualization, 2016) architecture can be virtualized. The TR-359 architecture is shown in Figure 7.

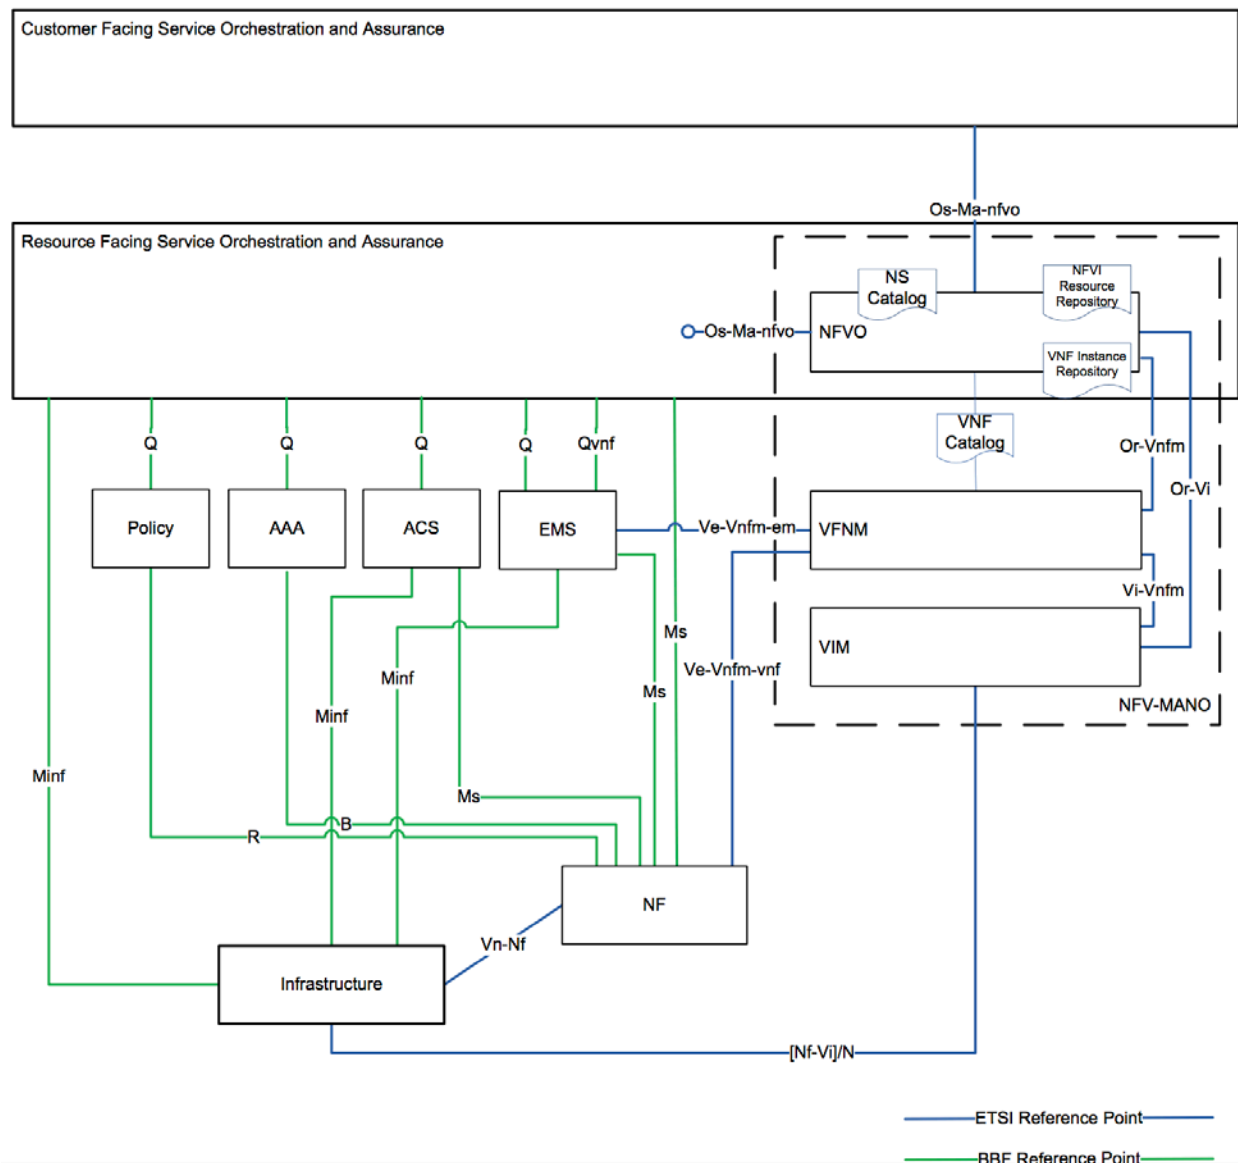


Figure 7 - BBF and ETSI-NFV reference model for service management and control (TR-359 A Framework for Virtualization, 2016)

4.3. Central Office Re-Architected as a Data Center (CORD)

All the reference architectures discussed so far have focused on the physical network elements and network functions required to deliver service to the subscriber. With the notable exception of TR-345 and TR-359, what has not been well covered is the framework for and implementation of software-defined networking (SDN) and network function virtualization (NFV) in a broadband access network.

In 2015, AT&T published a whitepaper entitled *Central Office Re-Architected as a Datacenter* which has been updated in (Central Office Re-architected as a Datacenter, 2016). The whitepaper proposes a major re-think of how a telco's central office (CO) is architected and used. The whitepaper proposes that the CO

be considered a data center that houses general purpose computing hardware that runs in software what previously would be housed in application specific hardware occupying many square feet of space, consuming more energy than necessary, and taking too long to develop and deploy new features. Variations on the CORD architecture have been documented and are under development in the open-source community – namely Residential CORD (R-CORD), Enterprise CORD (E-CORD), Mobile CORD (M-CORD) and others.

The CORD architecture replaces the application specific hardware with virtualized functions running as a service in the cloud infrastructure. An example of this is found in the R-CORD architecture (Figure 8), in which the access network (EPON, GPON, DSL, etc.) is abstracted away from the control and management layers to present a single network architecture to be managed.

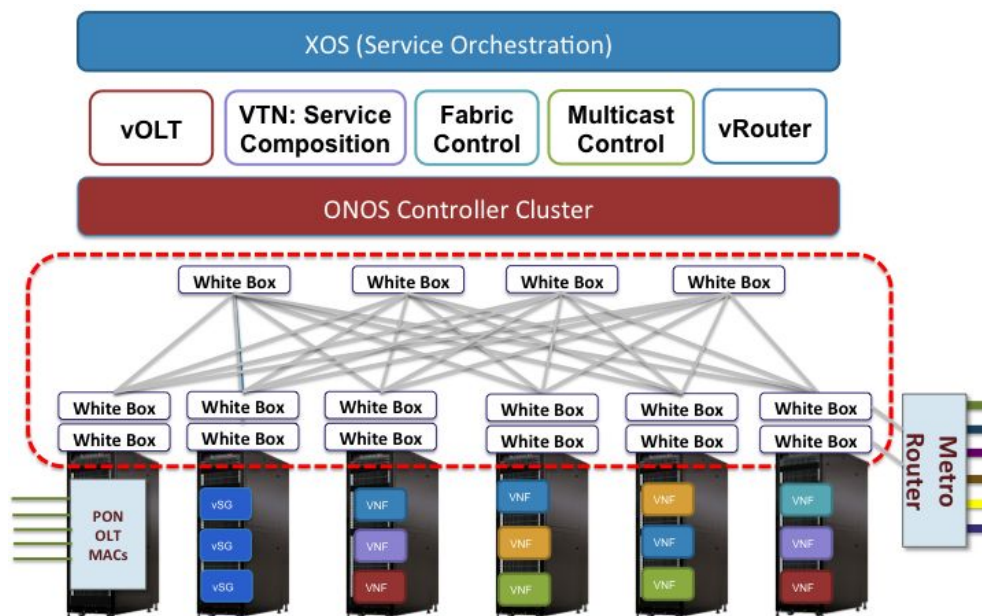


Figure 8 - R-CORD Architecture (Das, et al., 2016)

The virtual OLT (vOLT) which was described in (Al-Shabibi & Hart, 2016) and being developed under the OpenCORD project (VOLTHA Wiki, n.d.), is very relevant to the present distributed PON architecture proposal. The vOLT whitepaper (Al-Shabibi & Hart, 2016) describes an architecture that consists of a simple I/O blade that contains only the PON OLT MAC and a software module running in the cloud that implements all the PON control functions and protocols. Some of the many advantages this architecture brings is the ability to minimize the cost of the hardware by making it a simple media conversion and, by abstracting the control functions, removing many of the interoperability challenges that delay deployment of new PON OLTs and ONU/ONTs.

The CORD architecture is being adopted by telco operators worldwide and has now become a work effort in the BBF. This is exhibited by the sponsors listed on the OpenCORD website (OpenCORD Members, n.d.) and by the efforts going on in the BBF's CloudCO (Broadband Forum Cloud Central Office (CloudCO), 2017)work stream which is seeking to document a formalized architecture.

Architecture

After surveying the industry specifications and standards from which our proposed architecture might benefit, we want to develop the proposed architecture.

Since we are operating in the cable industry, conformance with DPoE is a key requirement for operating EPON. However, we know from the beginning that the distributed EPON architecture cannot fully comply with DPoE because DPoE does not currently support distributed functionality like we are proposing. We offer a compromise in this architecture – this architecture will conform to DPoE by:

1. Exposing the same interfaces to the OSS/BSS;
2. Maintaining the interoperability requirements of DPoE (by using DPoE OAM to the ONU);
3. Exposing the same interfaces to the operator's backbone networks as required by DPoE.

Prior to explaining the proposed architecture, it is necessary to explain how the existing DPoE and related functions can be split apart, or disaggregated. After that is accomplished the proposed architecture can be explained. We, in fact, propose two similar architectures in the interest of showing how an operator might migrate their network over time.

5. Disaggregating DPoE and SIEPON

There are two key areas of required discussion when we decide to create a distributed EPON network based on DPoE. The first is how to disaggregate the DPoE System (the DPoE ONU is not affected by disaggregation) and maintain the same functionality that is specified in DPoE.

According to (DPoE Architecture Specification, 2016), the DPoE System is made up of the following elements and interfaces:

- **vCM** – The virtual cable modem which translates all the operations, administration, maintenance and provisioning (OAM&P) messaging between the DOCSIS-based OSS and the OAM protocol used by DPoE ONU.
- **OLT** – The EPON Optical Line Terminal as defined by IEEE 802.3 and contains the PMD/PHY, MAC, MAC Control, and OAM functions.
- **Ethernet Switch** – The 802.1d Ethernet switch is a connection point and switching matrix for traffic moving between the OLT, R_{PE}, and the M_{NE} interface.
- **IP Router/Provider Edge Router (R_{PE})** – The R_{PE} provides the internal IP routing function, subscriber management, and MPLS PE router functions.
- **TU and TUL Interfaces** – The TU interface is the Optical Distribution Network (ODN); the TUL is a MAC domain that exists on the DPoE network.
- **C_S Interface** – The C_S interface is the downstream classifier that exists inside the OLT.
- **D Interface** – Is the interface between the DPoE System and the operator's network. It has varied functionality that includes bearer traffic (IP only), OAM&P functions and protocols, and MPLS traffic intended for Metro Ethernet (MEF) services connected to the DPoE System.
- **M_{NE} Interface** – Is a raw IEEE 802.3 interface acts as a MEF INNI and/or L2VPN NSI.

The functionality of each of these is detailed in (DPoE Architecture Specification, 2016).

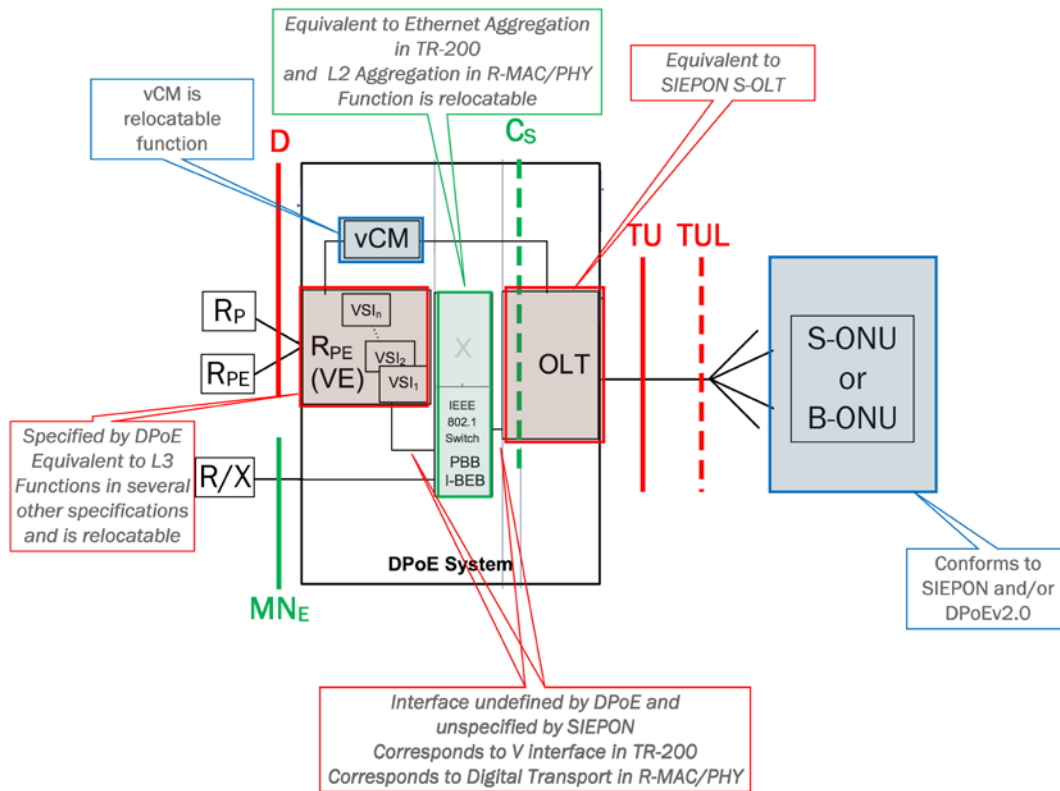


Figure 9 - Disaggregated DPoE System

Each of the elements listed are shown in Figure 9. Missing from the list is the interface between the OLT and the Ethernet switch and the interface between the Ethernet switch and the R_{PE}. These two interfaces must be defined if we are to separate the OLT from the physical DPoE System. These interfaces are similar in function to the V interface in TR-200 (TR-200 Using EPON in the Context of TR-101, 2011) and the Digital Transport described in R-MAC/PHY (Remote MAC-PHY Technical Report, 2015), so our architecture will look to those documents for guidance on their definition.

Figure 9 also shows the mapping between each element and interface and other standards or specifications.

As mentioned earlier, the DPoE ONU (S-ONU or B-ONU) and its functions and requirements are defined in the DPoE v2.0 suite of specifications and in IEEE 1904.1 (SIEPON) (IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON), 2013).

The Ethernet switch is equivalent to the L2 aggregation block shown in R-MAC/PHY (Remote MAC-PHY Technical Report, 2015). It is also equivalent to the Ethernet Aggregation block shown in TR-200 (TR-200 Using EPON in the Context of TR-101, 2011). These two documents will guide our definition of the functions required in this switch.

The vCM is a relocatable function defined in the DPoE v2.0 suite of specifications. *Relocatable* means that the function does not need to be co-resident with its adjacent functional blocks.

The R_{PE} is specified in several industry specification documents. While our architecture will require only the functions defined in the DPoE v2.0 suite of specifications, individual implementations may support additional functions as required by the operator. The R_{PE} is a relocatable function.

In this architecture, we choose to use the SIEPON Service-OLT as the model for the OLT contained in the DPoE System. This choice eases the effort required to define the functions required in the OLT, and which functions can be separated/relocated.

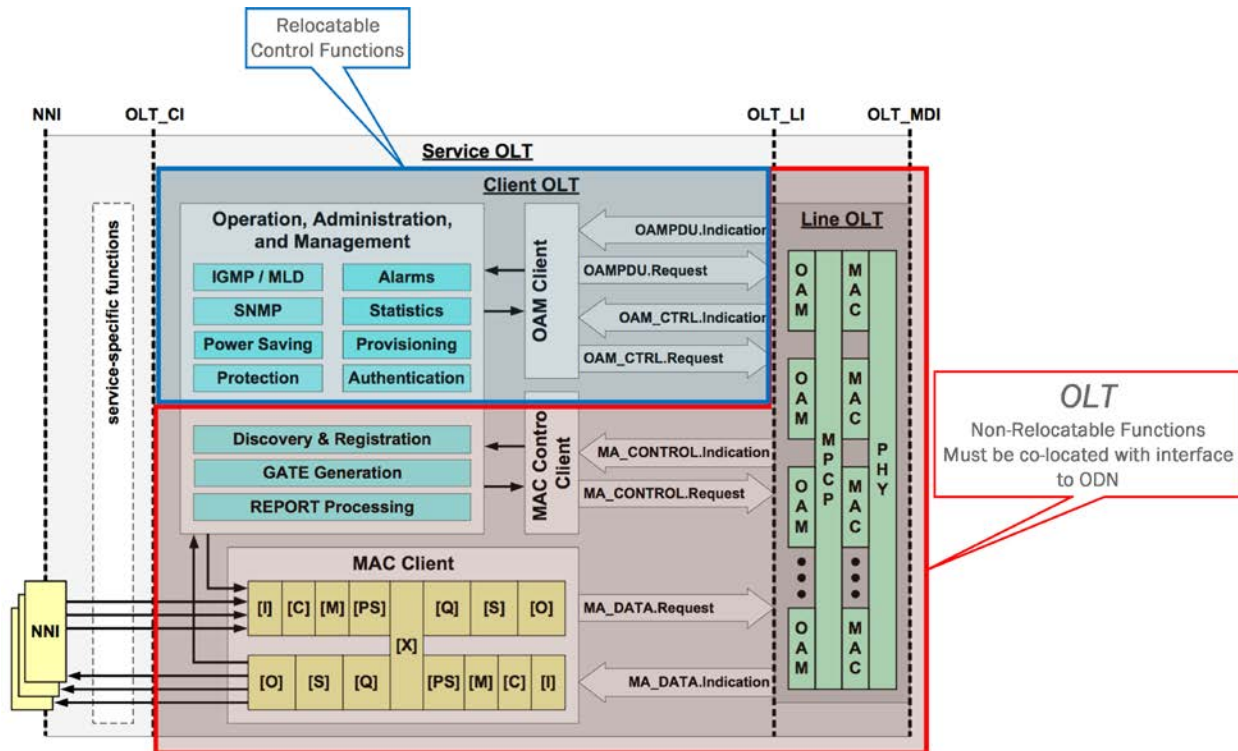


Figure 10 - Disaggregated SIEPON S-OLT

According to the SIEPON architecture, there are three types of OLT:

- Line OLT – is the element that provides the physical connection to the ODN and contains the fundamental functions for MAC, MAC Control (MPCP) and OAM.
- Client OLT – is the element that implements higher-layer functions that control the Line OLT's functions, establishes and manages connectivity with the ONUs, and sends/receives subscriber frames.
- Service OLT – is the element that contains one or more Client OLTs, and provides connectivity to external entities via a Network-Network-Interface (NNI), and may supply additional higher-layer functions such as L3 routing or others as defined in DPoE.

The relationship between these three entities is shown in Figure 10. Annotated in Figure 10 are the elements of the SIEPON architecture that can be separated and/or relocated in a distributed architecture. The Line OLT and a portion of the Client OLT functions cannot be separated because they are intimately involved with the basic functions of the PON and separating them would have negative impacts on performance (performance impacts are explored in (Boyd, Noll, Rahman, Nandiraju, & Villaruel, 2015)).

In the architecture proposals that follow, the OLT element will correspond directly to the combination of Line OLT and the portion of Client OLT functions that are outlined in Figure 10.

6. Reconnecting the Parts – Near Term Architecture

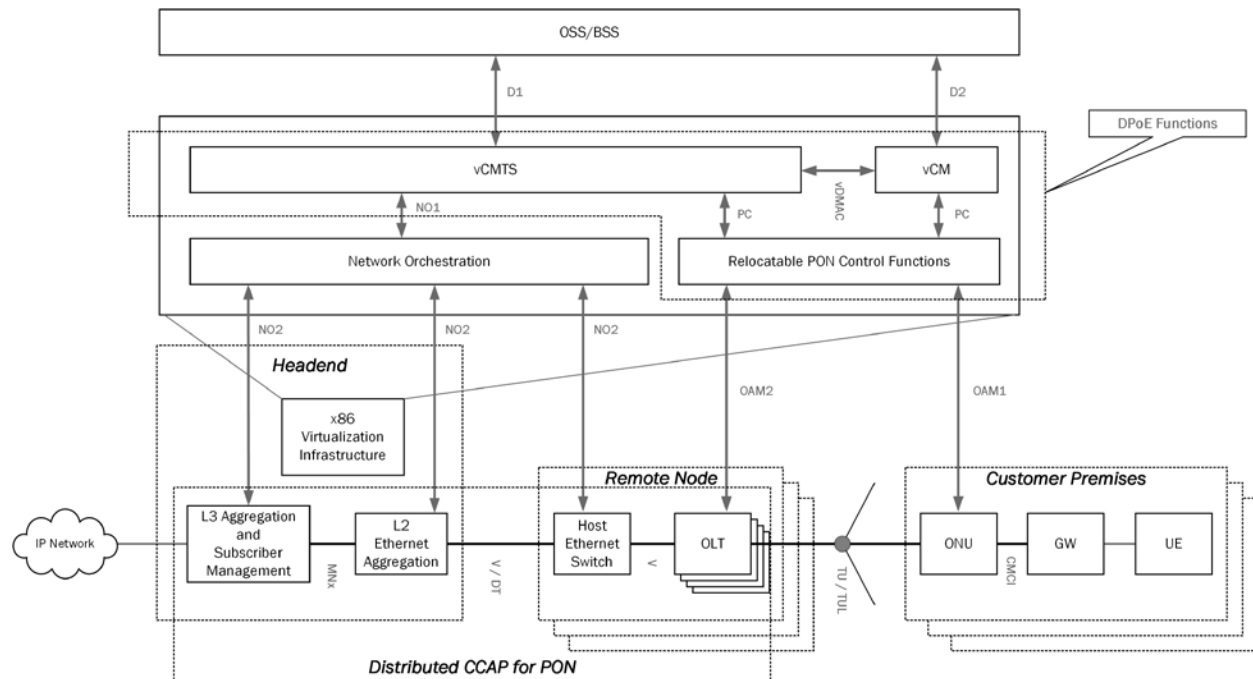


Figure 11 - Near Term Distributed EPON Architecture

The near term distributed EPON architecture is an example of how an operator could deploy distributed EPON today with little disruption to the existing operations. This architecture maintains the existing hardware-based routing and BNG functions, but replaces the OLT system chassis (though it is not necessarily incompatible with this architecture) with an inexpensive OLT that is supported by software elements. Those software elements exist in a virtualized environment and implement the PON control functions and DPoE OAM&P functions.

Key to making this architecture complete, functional, and interoperable is to define the interfaces and elements that are not defined by DPoE. The following sections do this.

6.1. The Elements

- **User Equipment (UE)** – Is the equipment that is used directly by the user and is not involved in providing the network connection. Examples: PC, Set Top Box, WiFi access point. This element is out of scope for this architecture.
- **Gateway (GW)** – Is the element that provides service-specific features and connectivity functions at layers above that provided by the ONU (e.g. Layer 3 and above) at the customer premises. The GW may be operator owned or customer owned.
- **Optical Network Unit (ONU)** – In this architecture the ONU is a SIEPON Client ONU and is equivalent to a DPoE Bridge ONU. This device provides basic layer-2 bridging between the PON

and the GW. The ONU contains all the necessary PON OAM&P functions to operate on the EPON.

- **Optical Line Terminal (OLT)** – In this architecture the OLT is a combination of the SIEPON L-OLT and a portion of the SIEPON C-OLT functions. Specifically, the OLT contains the L-OLT functions, MAC Client functions (bridging, cross-connect queueing, VLAN tagging, etc), and MAC Control functions (GATE generation, REPORT processing, Discovery and Registration).
- **Host Ethernet Switch** – Is the element to which the OLT connects directly through the V interface (defined below). It provides basic layer-2 switching and VLAN tagging functions. It is desirable to minimize the cost of this device; therefore it is desirable for this device not to contain significant queueing and shaping/policing features.
- **L2 Ethernet Aggregation** – Is the element to which one or more Host Ethernet Switches connect for aggregating traffic from many OLTs and many remote nodes. This device implements the functionality describes in R-MAC/PHY for the L2/L3 aggregation (L2 functions only in this architecture) and the Ethernet Aggregation block in TR-200.
- **L3 Aggregation** – Is the element which provides L3 routing and service functions. The functions of this element are well summarized in (Emmendorfer & ZorluOzer, 2016) as IPv4/IPv6 Router (and related forwarding and control-plane functions), DHCP relay, MPLS PE, VPLS, VPWS, etc.
- **Subscriber Management (BNG)** – Is the element at which (borrowing the definition from TR-101 (TR-101 Migration to Ethernet-Based DSL Aggregation, 2006)) bandwidth and QoS policies are applied to subscribers' flows. This function is defined by DPoE as residing inside the DPoE System.
- **Virtualization Infrastructure** – Is an infrastructure of computing platforms (typically x86-based) on which software implementations of network functions can be executed.
- **Relocatable PON Control Functions** – Is the upper half of the set of functions described in SIEPON's Client OLT. These functions are highlighted in blue in Figure 10, and include Authentication, Provisioning, Statistics, Alarms, Power Saving, Protection, IGMP/MLD, SNMP (the latter two are implemented in the vCM and vCMTS for DPoE).
- **Network Control and Orchestration** – Is an abstract function that controls certain lower-layer network elements and coordinates activities between all the various elements so that all required "system" functions are achieved. In other words, Network Control and Orchestration is required to coordinate all the disaggregated functions into a system that completely meets the requirements of the operator.
- **vCM** – Is the virtual cable modem defined in the DPoE v2.0 suite of specifications.
- **vCMTS** – As the name implies, the vCMTS is a virtual CMTS. This element is not currently defined explicitly in any specifications. However, the functions are implied in the DPoE v2.0 suite of specifications. Its functions include CLI access to manage the DPoE System configuration, SNMP, DOCSIS MAC-layer emulation, etc. In this paper, we leave this function vaguely defined due to lack of space. Future development of the proposed architecture could further define the vCMTS.
- **OSS/BSS** – Is the set of Operational Support Systems and Billing Support Systems used by the cable operator. These systems interact with current DPoE Systems via the D interface as defined in the DPoE v2.0 suite of specifications.

6.2. The Interfaces

- **D1** – Is the interface between the vCMTS and OSS/BSS. This interface is a subset of the D interface defined by the DPoE v2.0 suite of specifications. The subset includes SNMP (CMTS-specific MIBs, only), Command Line Interface (defined in the DPOE IPNE specification), policy control (currently undefined in DPoE, but could conform to the PacketCable/COPS suite of specifications for DOCSIS), etc. For lack of space, this paper will not attempt to fully define this interface.
- **D2** – Is the interface between the vCM and OSS/BSS. This interface is a subset of the D interface defined by the DPoE v2.0 suite of specifications. The subset includes TFTP (for vCM configuration file download), DHCP, SNMP (vCM-specific MIBs only), etc. For lack of space, this paper will not attempt to fully define this interface.
- **NO1** – Is the interface between the vCMTS and the Network Orchestrator. This interface is likely to be RESTCONF or NETCONF/YANG as suggested in (Virtual Provisioning Interfaces Technical Report, 2017) and (SDN Architecture for Cable Access Networks Technical Report, 2015). In this architectural model (architecture 1 in this paper), the vCMTS would be responsible for translating function calls from the D1 interface to the necessary function calls on the NO1 interface.
- **PC** – Is the interface to the PON Control Functions element. As exemplified in this architecture, the PC interface exists between the vCM and the PON Control Functions and between the vCMTS and the PON Control Functions. There are two prime candidate protocols for this interface: RESTCONF or NETCONF/YANG and OpenFlow. RESTCONF seems to have more favor in the cable industry, but a combination of RESTCONF and OpenFlow could be a superior solution.
- **vDMAC** – Is the interface between the vCMTS and the vCM. This interface is intended to emulate the minimum set of DOCSIS MAC protocol and messaging necessary to achieve DOCSIS-like functionality on the EPON network. A structured specification of this interface (whether public or proprietary) could enhance the existing DPoE v2.0 functions to bring it closer to full DOCSIS capabilities (for example, addition of dynamic service flows via emulated DSX messaging).
- **NO2** – Is the interface between Network Orchestration and lower-layer network elements – L3 Aggregation, Subscriber Management, L2 Ethernet Aggregation, and the Host Ethernet Switch. This interface is similar in function to NO1 and is likely to be implemented as RESTCONF or NETCONF/YANG. Network Orchestration would be responsible to translate the directives sent by the vCMTS to corresponding directives sent to the lower-layer network elements for implementing the features and functions required by the DPoE v2.0 suite of specifications and any additional or alternate requirements the operator might impose.
- **OAMI** – Is the interface between the PON Control Functions and the ONU. This messaging protocol for this interface is entirely defined by the DPoE v2.0 OAM specification (DPoE OAM Extensions Specification, 2017) and/or the SIEPON OAM specifications (IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON), 2013). What is missing from those specifications is a transport for the OAM messages to be sent across a non-PON network to the OLT and then to the ONU. IEEE 1904.2 (IEEE 1904.2 Task Force, n.d.) is a work in progress that is intended to provide just such a transport.
- **OAM2** – Is the interface between the PON Control Functions and the OLT. This messaging protocol for this interface is not currently defined in an industry standard or specification document. This architecture proposes that this interface adopt the same messaging protocol,

message formats and message content as is defined in DPoE v2.0 OAM specification (DPoE OAM Extensions Specification, 2017) and/or the SIEPON OAM specifications (IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON), 2013). There is much overlap between what is needed to program an ONU and an OLT, but clearly there are instances where the OLT will require new message content. This paper does not attempt to specify the required additional message content, but calls out the need for that work to be done in the industry. Like OAM1, this interface could use the proposed IEEE 1904.2 (IEEE 1904.2 Task Force, n.d.) standard to provide transport between the PON Control Function and the OLT.

- **MN_x** – Is the interface between the L2 Ethernet Aggregation and the L3 Aggregation and/or Subscriber Management elements. This interface should follow the specifications for the MN_E interface define in the DPoE v2.0 suite of specifications. This architecture does not simply call for the MN_E interface because the MN_E is a subset of functionality required. The MN_x interface must also implement the bearer plane functions of the DPoE v2.0 D interface. This can be accomplished by using the functionality of the MN_E interface to move traffic to the L3 aggregation and Subscriber Management functions.
- **V/DT** – Is the interface between the Host Ethernet Switch and the L2 Ethernet Aggregation. It is expected to be a simple fiber-based IEEE 802.3 and IEEE 802.1-based interface and corresponds to the Digital Transport interface called out in R-MAC/PHY and the V interface called out in TR-200. This architecture proposes that the V/DT interface be specified as a combination of requirements from these two industry documents.
- **V** – Is the interface between the Host Ethernet Switch and the OLT. It is expected to be a simple IEEE 802.3 and IEEE 802.1-based interface and corresponds the V interface called out in TR-200. This interface is expected to be short-hop fiber or other IEEE 802.3-compliant physical interfaces. This architecture proposes that the V interface here adopt the TR-200 V interface as a basis for its definition.
- **TU/TUL** – The TU interface is the same TU interface between the OLT and the ONU, as defined in the DPoE v2.0 suite of specifications. The TUL interface is the same TUL interface (a MAC domain) between as defined in the DPoE v2.0 suite of specifications with the subtle difference that the TUL interface will extend across the physical network to the L3 aggregation function.
- **CMCI** – Is the interface between the ONU and GW (or UE). This is the same CMCI interface that is defined in the DPoE v2.0 suite of specifications.

7. Reconnecting the Parts – Long Range Architecture

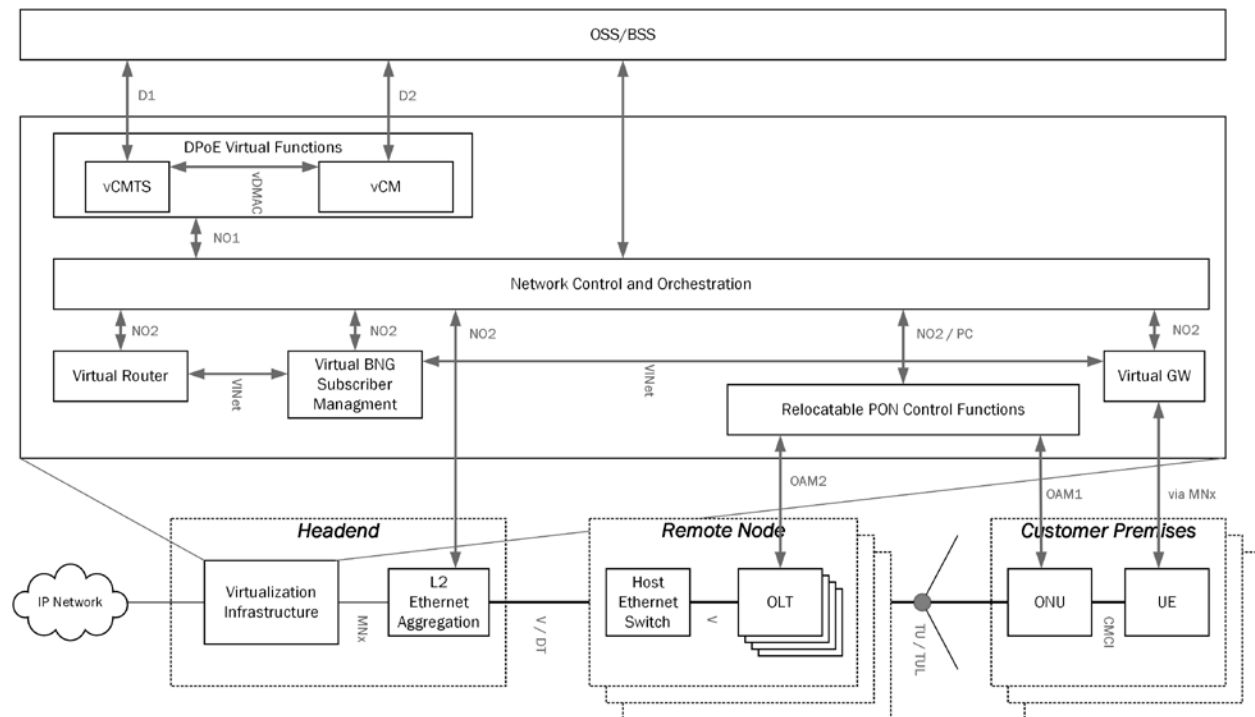


Figure 12 - Long Range Distributed EPON Architecture

The long range distributed EPON architecture is an example of how an operator could deploy distributed EPON with full SDN and NFV control while maintaining a migration path from DOCSIS-based provisioning. This architecture does away with traditional “big iron” application-specific systems and replaces them with virtualized network functions (VNF) running in software on a generic computing platform.

This architecture is a simple extension of the near term distributed EPON architecture proposed above, therefore most of the elements and interfaces are the same. Avoiding duplication, this section will discuss only the elements that are different between the two architectures.

7.1. The Elements

- **Gateway (GW)** – In this version of the architecture, the application-specific hardware GW is removed and replaced by a virtual GW running on the Virtualization Infrastructure.
- **Virtual GW (vGW)** – Is the VNF that implements the functions defined for the GW. It runs as a software module in the Virtualization Infrastructure and benefits from lower cost, due to the lack of application-specific hardware, and faster time to market with new features and functions due to the ability to develop and deploy software faster in the virtualized environment. The vGW must conform to the same feature/functional specifications as the GW, with the clear exceptions that are made for a device running as a VNF (for example, physical interfaces are not required).

- **L3 Aggregation** – In this version of the architecture, the application-specific hardware L3 Aggregation is removed and replaced by a virtual router running on the Virtualization Infrastructure.
- **Virtual Router (VR)** – Is the VNF that implements the functions defined for the L3 Aggregation Function. It runs as a software module in the Virtualization Infrastructure and benefits from lower cost, due to the lack of application-specific hardware, and faster time to market with new features and functions due to the ability to develop and deploy software faster in the virtualized environment. The VR must conform to the same feature/function specifications as the L3 Aggregation Function with the clear exceptions that are made for a device running as a VNF (for example, physical interfaces may not be specified for the VR, but instead would be specified for the Virtualization Infrastructure).
- **Subscriber Management (BNG)** – In this version of the architecture, the application-specific hardware BNG is removed and replaced by a virtual BNG running on the Virtualization Infrastructure.
- **Virtual BNG (vBNG)** – Is the VNF that implements the functions defined for the Subscriber Management function. It runs as a software module in the Virtualization Infrastructure and benefits from lower cost, due to the lack of application-specific hardware, and faster time to market with new features and functions due to the ability to develop and deploy software faster in the virtualized environment. The vBNG must conform to the same feature/function specifications as the Subscriber Management function with the clear exceptions that are made for a device running as a VNF (for example, physical interfaces may not be specified for the VR, but instead would be specified for the Virtualization Infrastructure).

7.2. The Interfaces

The only new or different interface in the long term architecture is the VINet interface.

- **VINet** – Is the interface between virtualized network functions running on the Virtualization Infrastructure. This specification does not attempt to define this interface and leaves it to be implementation specific based on the operator's architecture for the Virtualization Infrastructure.

Conclusion

In this paper, we have developed the framework for Distributed EPON Access in the cable industry. In the interest of reusing as much as possible from existing work, we surveyed the industry for standards, specifications and example implementations. This survey informed much of the proposed architecture and found that all elements and all interfaces, with one exception, have a basis in existing standards and specifications. By reusing the existing body of work, the work required by the industry is reduced significantly if it desires to write a complete specification for a distributed EPON architecture.

The proposed architecture fully supports the cable industry DPoE v2.0 suite of specifications and provides a migration path from current DOCSIS-based provisioning and operational models to SDN-based and NFV-based models that are currently being specified by industry consortiums and standards bodies and being implemented by non-cable telecommunications operators.

It will be important to the cable industry to adopt similar SDN-based and NFV-based models to stay competitive in the telecommunications market. The proposed architecture provides a framework for cable operators to use as they develop their own SDN and NFV architectures.

Bibliography & References

- Al-Shabibi, A., & Hart, J. (2016, February 12). Virtual OLT (vOLT). CORD at <http://www.opencord.org/>.
- Bernstein, A., & Ramakrishnan, S. (2015). SDN as a Matchmaker for Remote PHY Architecture. *Spring Technical Forum*. NCTA.
- Boyd, E., Noll, K. A., Rahman, S., Nandiraju, N., & Villaruel, F. (2015). Remote PON Network Performance. *Spring Technical Forum*.
- Broadband Forum Cloud Central Office (CloudCO)*. (2017, July). Retrieved from <https://www.broadband-forum.org/standards-and-software/major-projects/cloud-central-office>
- Central Office Re-architected as a Datacenter. (2016, March 14). CORD @ <http://www.opencord.org/>.
- Converged Cable Access Platform Architecture Technical Report. (2012). Cable Television Laboratories, Inc.
- Das, S., Al-Shabibi, A., Hart, J., Chan, C., Castro, F., & Moon, H. (2016, March 1). CORD Fabric, Overlay Virtualization, and Service Composition. CORD at <http://www.opencord.org/>.
- Distributed CCAP Architectures Overview Technical Report. (2015). Cable Television Laboratories, Inc.
- DPoE Architecture Specification. (2016). Cable Television Laboratories, Inc.
- DPoE OAM Extensions Specification. (2017). Cable Television Laboratories, Inc.
- Emmendorfer, M. J., Cloonan, T. J., Ulm, J., & Maricevic, Z. (2014). A Side-by-Side Comparison of Centralized vs. Distributed Access Architectures. *Spring Technical Forum*. NCTA.
- Emmendorfer, M., & ZorluOzer, S. (2016). A comparison of Centralized vs. Distributed Architectures for PON. *Spring Technical Forum*. NCTA.
- IEEE 1904.2 Task Force*. (n.d.). Retrieved July 2017, from http://www.ieee1904.org/2/tf2_home.shtml
- IEEE Standard for Service Interoperability for Ethernet Passive Optical Networks (SIEPON). (2013). IEEE Communications Society.
- M. Emmendorfer and T. Cloonan. (2013). Examining the Future Evolution of the Access Network. *Spring Technical Forum*. NCTA.
- Modular Headend Architecture v2 Technical Report. (2015). Cable Television Laboratories, Inc.
- OpenCORD Members*. (n.d.). Retrieved July 2017, from <http://opencord.org/members/>
- Remote MAC-PHY Technical Report. (2015). Cable Television Laboratories, Inc.
- Remote PHY Specification. (2017). Cable Television Labs, Inc.

SDN Architecture for Cable Access Networks Technical Report. (2015). Cable Television Laboratories, Inc.

Sundaresan, K. (2015). Evolution of CMTS/CCAP Architectures. *Spring Technical Forum*. NCTA.

Torbet, D., Cloonan, T., & Aftelak, A. (2016). Service and Management Orchestration in Distributed High Speed Data Networks. *Spring Technical Forum*. NCTA.

TR-001 ADSL Forum System Reference Model. (1996). Assymetric Digital Subscriber Line Forum.

TR-058 Multi-Service Architecture and Framework Requirements. (2003). DSL Forum.

TR-059 DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services. (2003). Digital Subscriber Line Forum.

TR-101 Migration to Ethernet-Based DSL Aggregation. (2006). Digital Subscriber Line Forum.

TR-200 Using EPON in the Context of TR-101. (2011). The Broadband Forum.

TR-345 Broadband Network Gateway and Network Function Virtualization. (2016). The Broadband Forum.

TR-359 A Framework for Virtualization. (2016). The Broadband Forum.

TR-359 A Framework for Virtualization. (2016). The Broadband Forum.

Virtual Provisioning Interfaces Technical Report. (2017). Cable Television Laboratories, Inc.

VOLTHA Wiki. (n.d.). Retrieved July 2017, from <https://wiki.opencord.org/display/CORD/VOLTHA>

Mobile Backhaul Synchronization Architecture

A Technical Paper prepared for SCTE•ISBE by

Jennifer Andreoli-Fang, PhD

Distinguished Technologist
CableLabs
Boulder, CO
303-661-3838
j.fang@cablelabs.com

John T. Chapman

CTO Cable Access and Cisco Fellow
Cisco Systems
San Jose, CA
408-526-7651
jchapman@cisco.com

Introduction

The growth in mobile data consumption has been putting pressure on the mobile network operators (MNOs) to build out small cell networks. All this traffic needs to be backhauled to the mobile core. While traditional choices for backhaul focus on fiber and microwave, hybrid fiber coaxial (HFC) networks have been making advancements. HFC is now being considered as a backhaul contender by the MNOs thanks to its capacity growth, cost efficiency and speed of deployment.

Traditional mobile base stations need to be frequency synchronized to guarantee handover performance, and this service is provided by the backhaul. In the DOCSIS 3.1 specification, the DOCSIS Time Protocol (DTP) was designed into the DOCSIS 3.1 specification to support precision timing from the CMTS to the cable modem (CM). This would allow a CM to provide backhaul services to a mobile base station for backhauling via the DOCSIS link. However, DTP is just one piece of the puzzle, as it needs to work with other elements of the operator network to provide timing to the base stations. This synchronization framework has yet to be defined. Furthermore, each operator network has differing levels of timing support in their existing hardware. This complicates system level designs.

In addition to frequency synchronization, Long-Term Evolution Time-Division Duplex (LTE-TDD) and LTE-Advanced features such as coordinated multipoint (CoMP) and enhanced inter-cell interference coordination (eICIC) all require stringent time and phase synchronization. Supporting these features places additional requirements on the synchronization framework.

In this paper, we review the technologies that can support frequency, time, and phase sync. We propose several architecture options, discuss their corresponding deployment scenarios, and the implications of each option on operations, cost of ownership, and time to market. Finally, we make recommendations on the device requirements and identify optimal designs based on operator deployments.

Drivers for Modern Backhaul Synchronization Requirements

The LTE downlink air interface utilizes orthogonal frequency division multiple access (OFDMA), while the LTE uplink uses single-carrier frequency division multiple access (SC-FDMA). OFDM is attractive for high speed wireless communications mainly due to its ability to combat frequency selective fading in multipath environments without the need for complex equalization techniques. But OFDM also requires orthogonality between the OFDM subcarriers, i.e., that 2 consecutive subcarriers must be non-overlapping in spectrum. Errors in frequency synchronization lead to loss of frequency orthogonality that can cause inter-carrier interference (ICI). In LTE systems, evolved node B (eNB) and user equipment (UE) must be frequency synchronized to 50-250 parts per billion (ppb) to allow the UE to demodulate LTE signals correctly, and to be able to transmit on the uplink.

While traditional macrocell networks only require frequency synchronization, the expected proliferation of small cells poses new challenges on timing distribution both technically and financially. What drives these new challenges is the focus of this section.

Before we begin though, depending on the deployment scenario, it is possible that a small cell deployment does not place further constraints on the synchronization requirements. Table 1 shows a list of deployment scenarios. For example, in rural outdoor deployments where the cell sites will always be able to receive signals from the Global Navigation Satellite System (GNSS) due to having a clear view of the sky, GNSS can be deployed in each cell site. Another scenario is if the traditional LTE FDD (Frequency Division Duplex) mode is deployed, only network frequency synchronization is required using PTP (Precision Time Protocol) which is part of IEEE-1588v2.

Table 1 – Deployment scenarios

	LTE-A Interference Management	Sync Requirements	Sync Methods
Dense urban outdoor (hotspot)	Needed	Frequency, time, phase	GNSS or PTP
Dense urban indoor (venue, MDU, enterprise)	Needed	Frequency, time, phase	No GNSS visibility → PTP
Suburban indoor residential	Not needed	Deployment can be TDD, so frequency, time, phase	No GNSS visibility, need cheap sync method → PTP
Rural	Not needed	Frequency	GNSS

Apart from these particular scenarios, a small cell deployment will introduce additional, and rather stringent requirements on time and phase synchronization. Figure 1 shows the difference between the different kinds of synchronization. If the cable operators want to leverage their DOCSIS networks for wholesale backhaul business, or to backhaul their own small cell networks, they need to have a repertoire of tools to use to solve these technical problems. We will discuss the toolkit in the remainder of this paper after we discuss the drivers.




Figure 1 – Frequency, time, and phase synchronization

1. LTE TDD

With the expected deployment in the 3.5GHz spectrum in the US, the popularity of TDD has grown. LTE TDD is also prevalent in Europe. But TDD requires tight time synchronization.

In LTE TDD, uplink (UL) and downlink (DL) transmissions occur at the same frequency but are separated in time. The eNBs have to easily inform the UEs in the cell whether they should be listening or transmitting. The 3GPP defines 7 TDD subframe configurations so that the UEs know which subframe is for transmit or receive, although most small cells today support subframe configurations 1 and/or 2 only.

The TDD frame structure for subframe configuration 1 is shown in Figure 2. A special subframe denoted as the “S” subframe is defined to include a partial UL and a partial DL subframe, with a “guard period” sandwiched in the middle for switching between the UL and DL transmissions.




Figure 2 – Frame structure for TDD subframe configuration 1

To ensure maximum spectrum reuse, the eNBs operate in the same frequency. Additionally, to minimize interference, a cluster of eNBs are configured to use the same subframe configurations, so that they are either transmitting or receiving at the same time. Consequently, the adjacent eNBs must be synchronized almost perfectly to avoid an UL transmission interfering with a DL transmission in the neighboring cell. The 3GPP has specified in 0 that the neighboring eNBs must be phase aligned to within 3μs.

2. Heterogenous Networks and Dense Deployments

Small cell deployments are intended to address the ever-increasing mobile demands in both indoor and outdoor scenarios. For outdoor deployments, small cells are deployed in the same coverage area as the macrocells to fill in the capacity gaps. It is preferable to deploy small cells in different spectrum from the macrocells, but it is not always possible, due to limited spectrum availability. In co-channel or in-band deployments, small cells operate on the same frequency as the macros to maximize spectrum utilization. Such networks are called heterogeneous networks, or HetNets. The small cells in the HetNets experience inter-cell interference, because the macrocells transmit at significantly higher power levels.

As a large percentage of mobile traffic is consumed indoors, operators need to deploy ultra-dense small cells to fulfill the capacity needs. These co-channel eNBs situated in close proximity cause interference to one another, particularly at the overlapping cell edges.

The operators need to implement interference management techniques to address the interference issues unique to small cell deployments.

Traditional LTE includes simple physical (PHY) layer techniques such as heavy coding or OFDM’s built-in cyclic prefix to combat interference. However, the techniques have all been designed for single cell operation. In case of HetNets and ultra-dense deployments, these methods are not enough.

To address this, a number of LTE Advanced (LTE-A) interference management techniques have been developed. We will now look at 2 techniques: eICIC and CoMP.

Ultimately, these techniques, while improving the small cell system capacity, pose stringent requirements on both synchronization and latency.

2.1. ICIC, eICIC

Suppose we have two neighboring eNBs operating on the same frequency. The UEs situated in the overlapping coverage area will experience high interference. This is because while the eNBs transmit to the UEs situated at the cell center with low power, they must transmit at higher power to the UEs at the cell edge in order to reach them with good enough SINR (signal-to-interference-plus-noise ratio). The situation is depicted in the left side of Figure 3.

Rather than transmitting blindly to the edge UEs with high power that would cause severe interference at the UEs, the two eNBs exchange information about what portion of the frequency spectrum it is planning to transmit with high power. In this way, the interference posed on each UE's data channel (PHY downlink shared channel, or PDSCH) is reduced. The right side of Figure 3 shows an example situation: while eNBs A and B operate on the same frequency resource f1, A will transmit in resource f3 with high power, while B will transmit in resource f2 with high power. This is in essence a way for the eNBs to partition the spectrum, so that they would not be transmitting with high power in the same OFDM subcarriers. This technique, developed in LTE Rel-8, is called inter-cell interference coordination (ICIC).

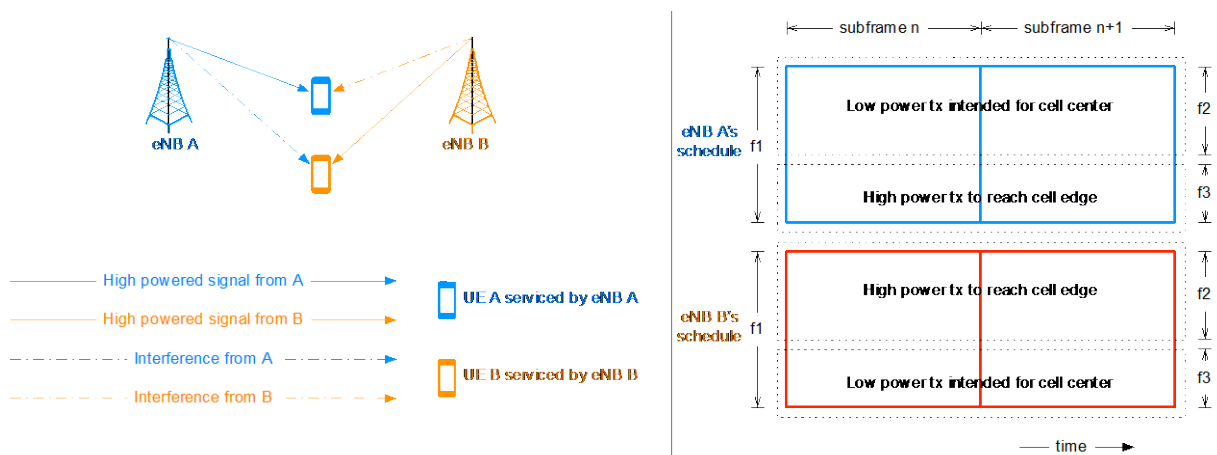


Figure 3 – Overlapping cells (left), and frequency domain inter-cell coordination (right)

While frequency partitioning works well for data channels, it does not solve the interference issue on the control channel. In each LTE subframe, the first 1-3 OFDM symbol(s) includes a broadcast control channel, i.e., the PHY downlink control channel (PDCCH), as shown in Figure 4, that includes subframe format indication and how the subframe is being scheduled to each UE. This channel must be received correctly in order for the UEs to decode the rest of the subframe.

To mitigate interference on the control channel, the LTE Rel-10 defines the eICIC with the concept of “almost blank subframe (ABS).” ABS is essentially subframe muting, and is shown in the left side of Figure 4. One of the eNBs provides information on which subframes in the near future it will mute, and sends this information to the other eNB. The negotiation involves message exchanges and takes place on the X2, which is a point-to-point logical interface between two eNBs.

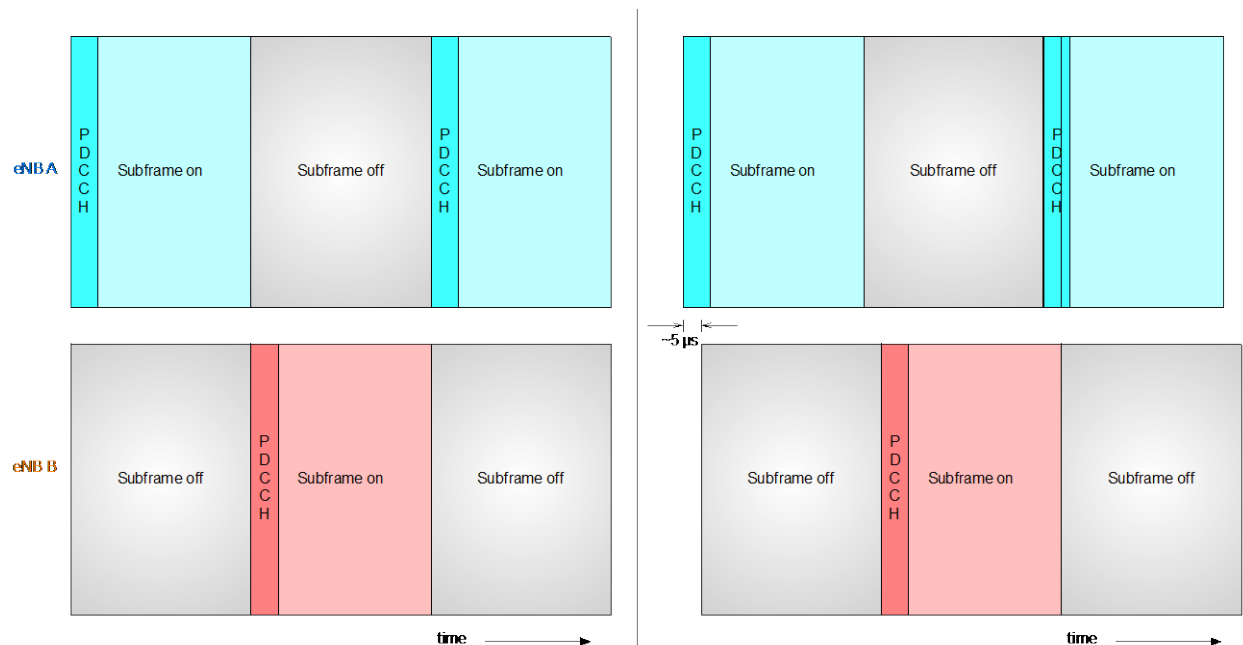


Figure 4 – ABS with perfect phase sync (left), and without (right)

ABS works when the two eNBs are in perfect phase synchronization. It is therefore critical for the clocks of the participating eNBs to be phase aligned, so that the subframes of the overlay eNBs do not overlap when one cell transitions its transmission to on while the other transitions to off. Otherwise, the PDCCH of one of the eNBs will experience severe interference as shown on the right side of Figure 4. The 3GPP does not formally define this phase sync limit. But various eNB vendors have quoted that the participating eNBs must generally be phase aligned within 5 μs, about the size of the cyclic prefix for the first LTE subframe symbol, in order for the technique to result in substantial performance gain.

2.2. CoMP

While eICIC improves the interference level experienced by the UEs at the cell edges, the UE's throughput is limited to what can be achieved in a single cell due to the frequency and time partitioning of the spectrum and airtime. CoMP, featured in LTE Rel-11, enables multiple eNBs to simultaneously serve the UEs residing at the cell edge at the same time, analogous to a MIMO system, to increase the signal level and thereby achieve better edge UE throughput. Furthermore, while eICIC works on a semi-static time frame which is not suited for fast-changing channel conditions, CoMP allows eNBs in the coordinating set to negotiate resources dynamically.

The 3GPP defines several types of CoMP: coordinated scheduling (CS) including beamforming, and joint processing, including dynamic point selection and joint transmission.

CS is in essence a dynamic version of ICIC, but with frequency resource partitioning occurring dynamically at every subframe. The left side of Figure 5 shows a snapshot of signal vs. interference in subframe n after 2 eNBs have coordinated their scheduling. When eNBs A and B operate on the same frequency resource f1, it is possible through CoMP signaling to optimize the bandwidth usage. In this example, eNB A will transmit in resource f3 to reach its edge UE A, while eNB B will transmit in resource f2 to reach its edge UE B at subframe n indicated in the figure. The right side of Figure 5 shows

that the spectrum resources are partitioned and that the scheduling of frequency resources can adapt dynamically on a subframe-by-subframe time scale.

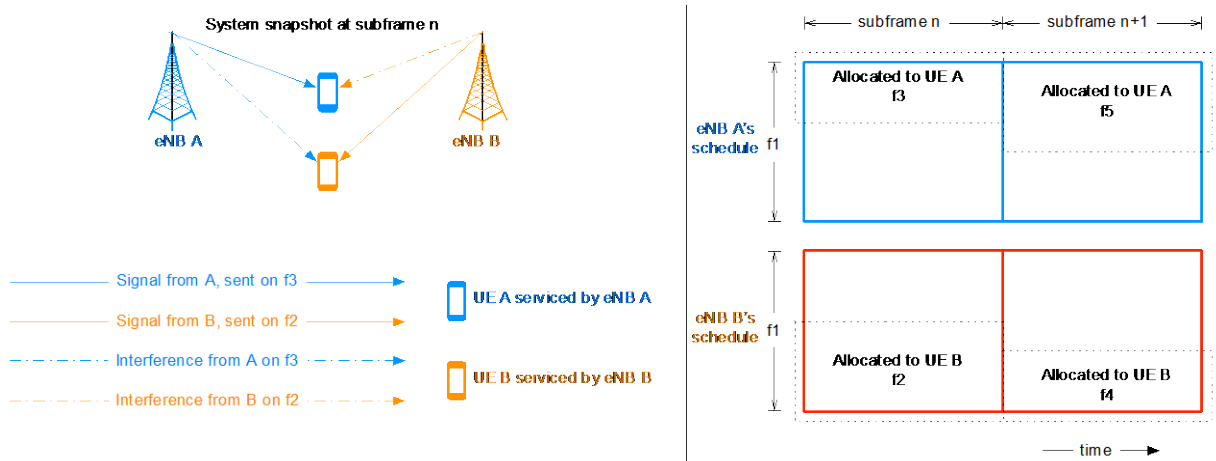


Figure 5 – Coordinated scheduling: snapshot of signal vs. interference for subframe n (left), and subframe scheduling after negotiation (right)

For CS, the data is only available at the UE's serving cell. In our case, eNB A serves as the master eNB for UE A, and eNB B serves as the master eNB for UE B. Scheduling and beamforming decisions are made by using the channel state information (CSI) shared between the eNBs in the coordinating set.

With joint processing, user data is available at multiple eNBs. Dynamic point selection is one type of joint processing. It is similar to CS in that only a single eNB transmits to an edge UE at a given time. The difference is that any eNB can serve an edge UE, compared to just the master eNB in the CS case. Figure 6 shows an example.

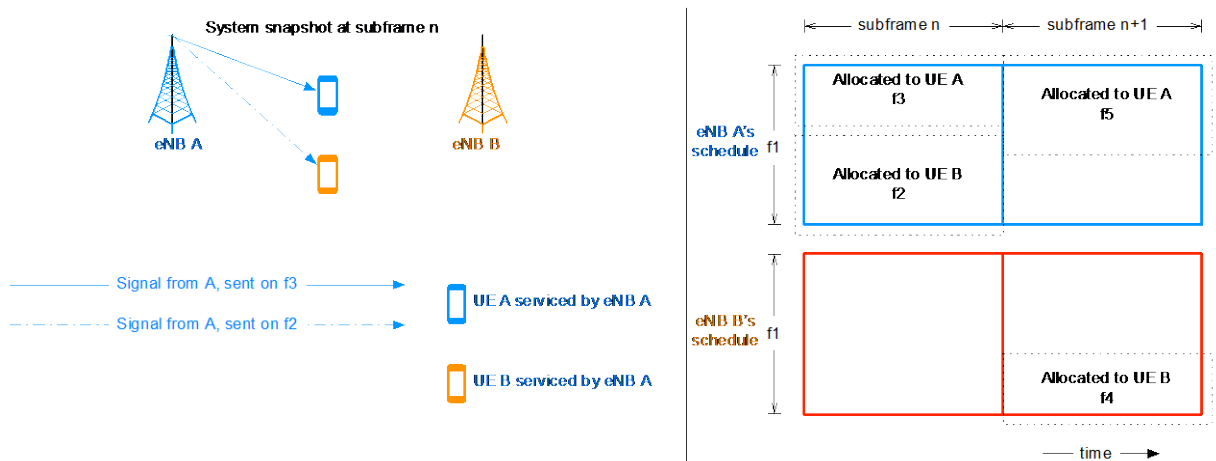


Figure 6 – Dynamic point selection: snapshot of signal vs. interference for subframe n (left), and subframe scheduling after negotiation (right)

With joint transmission, multiple eNBs can send the same data simultaneously to the edge UE at the same time and frequency resource. This improves the received power level at the UE and therefore improves the throughput.

A CoMP resource coordinator coordinates the schedules between multiple eNBs. It can reside in the eNBs in a distributed fashion, or be close to the evolved packet core (EPC) in a centralized fashion. Figure 7 gives a high level view of how CoMP works when a resource coordinator (RC) is located centrally. Referring to the steps in Figure 7: the UE in CoMP mode measures the CSI from all eNBs it can hear and sends CSI feedback to its master eNB (step 1). The eNBs forwards the CSI from the CoMP UEs to the Resource Coordinator (RC, step 2). The RC performs scheduling functions (step 3), and the scheduling information is then conveyed back to each eNB (step 4). If the CSI is delayed on the X2 interface, then the performance gain for the CoMP UEs will degrade.

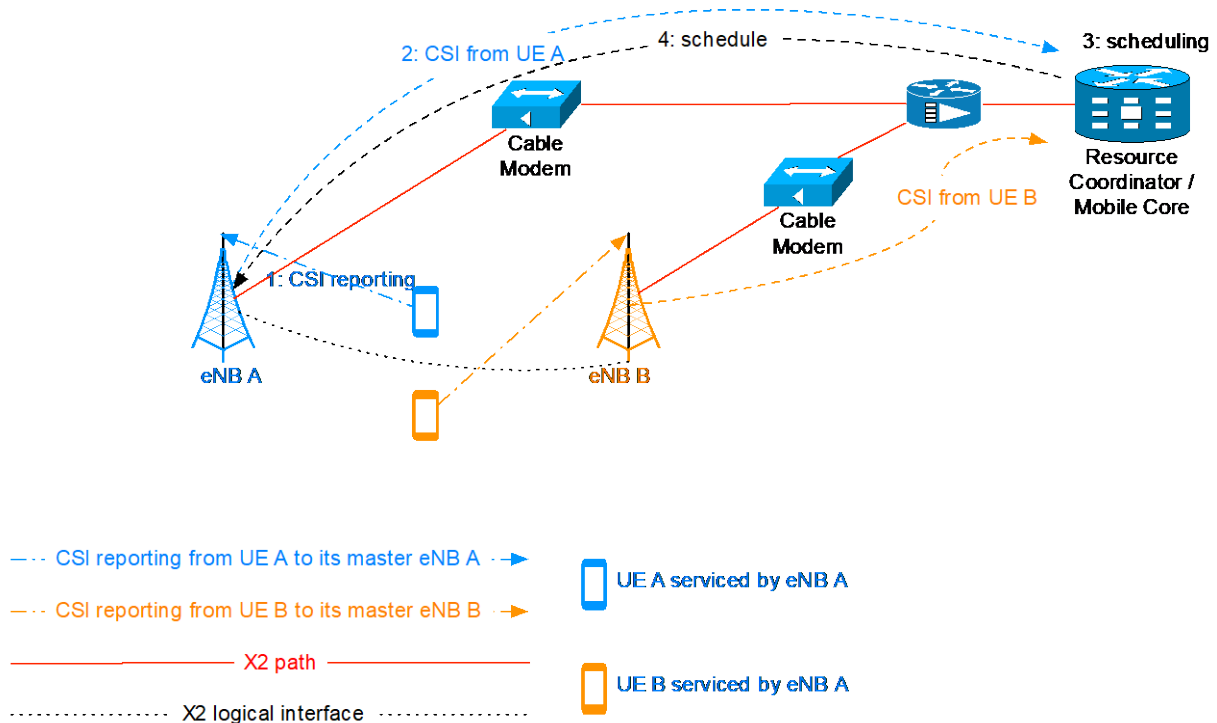


Figure 7 – CoMP operations

So, to support inter-eNB CoMP, the clocks of the neighboring eNBs must be time and phase synchronized to align the radio frames transmitted from different eNBs to the UE. The UE performance degrades with less accurate time and phase synchronization; and the amount of degradation depends on the CoMP technique.

In addition to the phase synchronization requirement, clearly, CSI must be sent expeditiously for the information to stay relevant and allow the cells to coordinate scheduling according to the dynamically

varying channel conditions. This leads to a set of latency requirements (and possible solutions), which is covered by the authors in 0.

3. eMBMS / MBSFN

Another LTE feature is the enhanced multimedia broadcast multicast services (eMBMS), with a common use case being mobile broadcasting of live sporting events. It is supported in LTE over the multicast broadcast single frequency network (MBSFN). MBSFN allows multiple eNBs to transmit identical waveforms at the same time and frequency resources. The UE combines the multiple waveforms as multipath components of a single eNB. So the synchronization requirement is driven by the OFDM cyclic prefix in order to avoid inter-symbol interference.

4. Summary

Table 2 summarizes synchronization requirements for LTE TDD, and LTE-Advanced features such as CoMP, eICIC, eMBMS described in this section.

Table 2 – LTE and LTE-Advanced synchronization requirements

	Frequency	Phase	Notes
LTE FDD	± 50 ppb	None	3GPP TS 36.104 0 §6.5.1
LTE TDD	± 50 ppb (wide area) ± 100 ppb (local area) ± 250 ppb (home)	$10 \mu\text{s}$ (wide: cell radius $>3\text{km}$) $3 \mu\text{s}$ (local: cell radius $<3\text{km}$) $1.33 \mu\text{s} + T_{\text{prop}}$ (home eNB radius $>500\text{m}$) $3 \mu\text{s}$ (home eNB radius $<500\text{m}$)	3GPP TS 36.133 0 §7.4.2
CoMP	None	$\pm 1.5 \mu\text{s}$	
eICIC	None	$\pm 1.5 - 5 \mu\text{s}$	
eMBMS / MBSFN	None	$\pm 10 \mu\text{s}$	

Synchronization Toolkit

There are 3 types of technologies to provide synchronization: physical layer mechanism such as synchronous Ethernet (SyncE), packet-based method such as precision time protocol (PTP) as defined in IEEE-1588v2, and GPS-based method. The latter two can also support time and phase synchronization.

Mobile operators today have various solutions in their repertory to distribute frequency, time, and phase synchronization. One straightforward method shown in Figure 8 uses a distributed architecture where the reference signal is distributed through the satellite signals. The GNSS receiver extracts the signal and is co-located with the end applications, in this case, the cellular radio sites. As such, this method does not require timing support from the network elements. This is an important point, because attaching a GNSS receiver to each cell site versus upgrading all or part of the network elements to support timing distribution poses financial tradeoffs for an operator, a point of consideration with dense small cell

deployments. However, a GNSS-based method is not useable in most indoor deployment scenarios, and is not always reliable in the outdoor scenarios either due to atmospheric effects. So, operators must consider alternate technologies, which will be the focus of the remainder of this section.

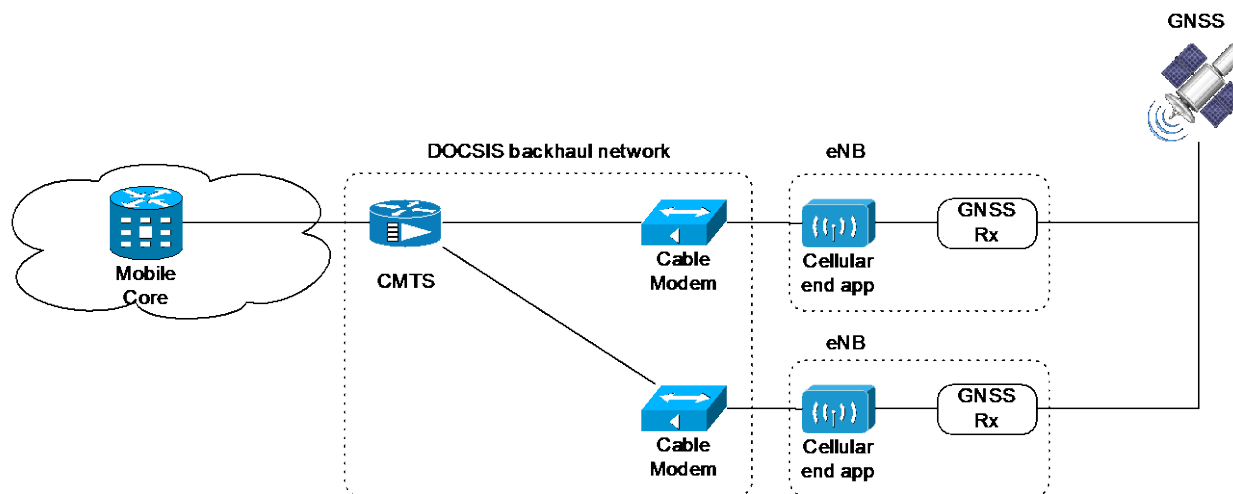


Figure 8 – Distributed GNSS-based synchronization without network support

5. Synchronous Ethernet

Synchronous Ethernet (SyncE) is a physical layer based frequency distribution algorithm. It is based on the Ethernet standards but additionally transmits a PHY transmit clock. In traditional Ethernet, a free running PHY clock is transmitted over the network medium. In SyncE, the Ethernet transmit clock is locked to a reference clock. Because the clock is continuously transmitted, SyncE is not subject to impact of the noise sources as PTP and Network Time Protocol (NTP). As such, SyncE is used as a mechanism to stabilize the frequency clock in case of failure events, i.e., when the timing can no longer be traced to a Primary Reference Time Clock (PRTC).

6. IEEE 1588 Precision Time Protocol

The IEEE 1588 precision time protocol 0, also known as PTP, is a packet-based synchronization technology that provides time, frequency, and phase synchronization. Since numerous tutorials exist on PTP (e.g., 0), this section will only provide a brief overview.

The basic principle is to distribute time sync reference by means of a 2-way timestamp exchange. As the mean path delay is half of round trip delay, the basic assumption is that a symmetric path between the packet master and the packet slave is required.

In contrast to SyncE that defines the timing content of the signals based on the significant edges of a data signal, PTP relies on the transmissions of timing messages. The series of time messages allows a PTP slave to recover the clock by estimating its timing offset from a PTP master.

The level of precision of the recovered clocks is contingent upon the PTP client's ability to filter out the noise sources that will affect the accuracy of the recovered clock. The operator also has the ability to build out the network to minimize the effect of the noise sources, which include but are not limited to:

- Reference clock drift
- Timestamp error at the PTP grandmaster and slave
- Packet delay variation (PDV)
- Network asymmetry

To reduce the PDV that results from the queuing delay of the event messages, IEEE 1588 defines boundary clocks (BC) and transparent clocks (TC). Both BC and TC are switches or routers that participate in the timing protocol but in different ways. The BC terminates the PTP protocol, i.e., all PTP messages, on its slave port, uses the timing message to set its clock, and regenerates the PTP messages on its master port(s). In contrast, the TC does not set its clock based on the event messages, but instead, adjusts the event message timestamp to reflect the propagation time for the message to traverse through the equipment.

7. ITU-T G series recommendations

There has been confusion about the relationship of the PTP and the telecom profiles. The protocol as defined by the IEEE 1588 committee is the protocol that defines a set of message exchanges between two nodes. PTP alone does not guarantee meeting the end application's performance requirements. Equipment that implement the PTP may not interoperate with each other, and may not satisfy any end application performance requirements. So, the International Telecommunication Union (ITU) worked and agreed on a set of architectures, telecom profiles, and performance specifications, all aimed towards meeting the performance requirements for telecom applications. In this section, we will outline the set of ITU-T Recommendations, what they are, what the relationship between them is, and as an operator, which Recommendations should be the focus depending on the deployment scenarios.

The ITU published a comprehensive set of Recommendations for distributing frequency, time, and phase synchronization, particularly geared towards telecom applications, since they have been the ones driving the tightness of the clock requirements. Within the set, the following are for time and phase synchronization:

- G.8260: general definitions and metrics
- G.8271 0: methods for distributing phase and time
- G.8271.1 0: maximum network limits on time errors and requirements on network elements
- G.8272 0: performance requirements for PRTC and T-GM
- G.8273: packet based phase / time clocks
- G.8273.1: performance requirements for T-GM
- G.8273.2 0: performance requirements for T-BC and T-TSC
- G.8273.3: performance metrics for T-TC
- G.8275 0: general architecture for distributing time and phase sync using PTPv2
- G.8275.1 0: PTP profile assuming full timing support from the network
- G.8275.2 0: PTP profile assuming partial timing support from the network

In particular, the operator should first focus on G.8275, where general architecture, along with protection mechanisms (holdover, which we will discuss shortly) are defined. It also specifies telecom-specific clock types, which are more rigorously defined than in the IEEE specs.

7.1. Types of telecom clocks

A primary reference timing clock (PRTC) is capable of providing frequency, time, and phase synchronization for other clocks in a network, by providing reference signals to a telecom grandmaster (T-GM). It is typically traceable to a universal time standard such as the UTC obtained from GNSS. G.8272 specifies the accuracy requirements for the PRTC and the T-GM.

A telecom boundary clock (T-BC) is an IEEE 1588 boundary clock with additional performance requirements defined in G.8273.2.

A telecom transparent clock (T-TC) is an IEEE 1588 transparent clock with additional performance requirements yet to be defined.

A telecom time slave clock (T-TSC) is an IEEE 1588 ordinary clock with only a slave port (i.e., cannot be a grandmaster) with additional performance requirements defined in G.8273.2.

7.2. Telecom profiles

The IEEE 1588-2008 introduced the concept of “profile” which includes a specific set of modes of operations, messages, message rates and attributes designed to satisfy an end application’s requirements.

The G.8275.1 telecom profile requires the network to provide full timing support. That is, boundary clocks must be implemented at every network node on the timing distribution path between the PTP grandmaster and the client. The profile defines a set of PTP parameters used to guarantee interworking between implementations. It specifies aspects such as the PTP messages to be used in the profile, 1-step vs. 2-step masters, message rates, protections, etc.

In contrast, the G.8275.2 telecom profile only requires the network to provide partial timing support. This really means that not every node in the timing distribution chain has to fully participate in the PTP, or to satisfy the performance requirement for T-BC. It is designed for operators who have no full timing support capability and cannot upgrade every switch and router in the timing distribution chain immediately.

The G.8275.2 telecom profile introduces additional clock types such as T-BC-P, T-TC-P, T-TSC-P, and T-TSC-A, where “-P” indicates “partial,” and “-A” indicates “assisted.” But performance characteristics of these clock types have yet to be defined.

Table 3 contains a summary of all the synchronization technologies discussed in this section.

Table 3 – Summary of synchronization technologies

	Pros	Cons
GPS/GNSS	<ul style="list-style-type: none">• Global coverage with great precision	<ul style="list-style-type: none">• Penetration is poor in indoor and dense urban with high rise• Upgrading every client to GNSS capability is expensive• Susceptible to jamming• Can be expensive if every cell site requires a receiver

	Pros	Cons
Packet-based, e.g., PTP	<ul style="list-style-type: none"> • Capable of providing frequency, time, and phase sync • Can be implemented for any deployment locations • Not every cell site must be upgraded (but still needs equipment upgrade on timing distribution chain) 	<ul style="list-style-type: none"> • Performance accuracy subject to noise sources • Operators must do careful testing and measurements of the entire timing distribution chain in order to ensure end application performance requirements can be met
SyncE	<ul style="list-style-type: none"> • PHY layer technology means it is not subject to the noise sources from packet-based distribution mechanisms • Can be used in conjunction with other protocols to increase holdover performance 	<ul style="list-style-type: none"> • Does not support time and phase sync • Point-to-point protocol means if a node in the chain is broken, synchronization for client cannot be achieved • P2P protocol means must upgrade every switch in the chain

Time and Phase Distribution over the DOCSIS Network

As discussed in Section 4, timing and synchronization requirements for small cells are stringent compared to the macrocell, due to interference management techniques and the use of TDD that may be required for small cell deployment. On top of this, DOCSIS is a packet based network. As such, it has the issue of network asymmetry. This makes distributing timing synchronization even more challenging. Luckily, DOCSIS Time Protocol that has been defined as part of the DOCSIS 3.1 specifications 0 several years ago. We will not be discussing the DTP in this paper. But for an excellent tutorial on DTP, see 0.

8. General architecture for phase and ToD distribution over DOCSIS backhaul

The general reference architecture for distributing time and phase using the DOCSIS backhaul network is shown in Figure 9. The PRTC provides timing reference for the timing distribution chain to the end application or the client. The PRTC can get the reference from a GNSS signal. An additional physical layer frequency synchronization signal can be included in the form of a primary reference clock (PRC, used for frequency synchronization only). We will discuss how SyncE can help improve the stability and accuracy during failure events shortly.

The PRTC is attached to a packet master clock known as the T-GM that implements packet-based distribution protocol such as IEEE 1588-2008. From there, the timing distribution chain can include a

series of T-BCs and/or T-TCs. Networks with full timing support will only implement T-BCs compliant with the ITU-T G.8273.2 0 spec. Networks with partial timing support can include T-TCs.

The distribution chain includes the backhaul network serviced by a DOCSIS network. In the DOCSIS portion of the chain, the CMTS and the CM can participate in the IEEE 1588-2008 timing protocol. If so, the CMTS-CM pair can form an IEEE 1588 Boundary Clock that terminates the PTP domain by the CMTS and regenerates the PTP timestamp by the CM.

The timing reference signal will eventually reach packet slave clock(s) in the form of Telecom Time Slave Clock(s) (T-TSCs). The T-TSC may be integrated with the end application, in this case, the eNB.

The general architecture provides requirement flexibilities depending on an operator's use case, performance requirement, and total cost of ownership (TCO) and time-to-market needs.

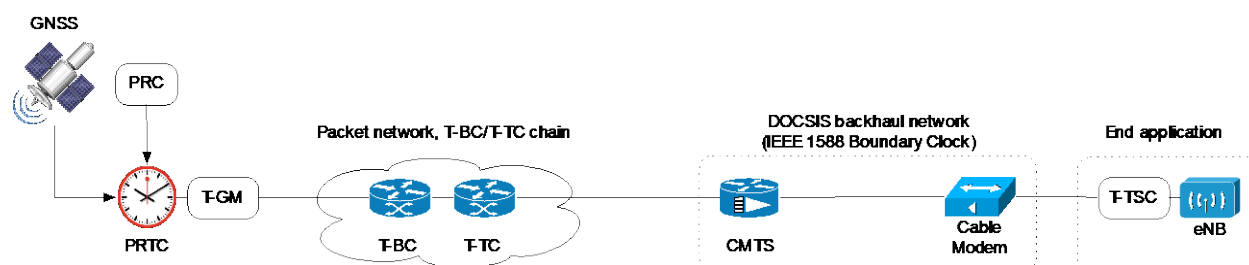


Figure 9 – General architecture to support phase and time distribution

The remainder of this section will cover various technology options for providing synchronization service using DOCSIS as backhaul. As we will see, each option has its merits and deficiencies.

9. Options for networks with full timing support

In the ideal case, each element in the operator network implements ITU-T G.8275.1 0 telecom profile, and is additionally compliant with certain performance requirements as specified in G.8272, G.8273.2.

9.1. G.8275.1 + DTP

This option uses the clocks implementing the G.8275.1 telecom profile to deliver time and phase synchronization throughout the distribution chain. Frequency synchronization is derived through time. The CMTS recovers time and frequency from the PTP messages, and translates the PTP timestamp into the DOCSIS timestamp. The CM regenerates the PTP timestamp based on the DOCSIS timestamp and passes it along to its downstream PTP slave which is the T-TSC. In this way, the CMTS-CM pair acts as a IEEE 1588 BC that terminates the 1588 domain at the CMTS, while the modem acts as a PTP master for the T-TSC in the end application client. Because the network is built with elements that are compliant with G.8275.1 and their corresponding ITU telecom standards (except the DTP domain elements), with time error budgeting, this option guarantees the delivery of frequency, time, and phase synchronization to LTE small cells that implement LTE-Advanced features such as eICIC, CoMP, and for LTE TDD deployments.

Implementing G.8275.1 on every element in the timing chain except the DTP domain provides guarantees that the time error for each link will be within the maximum absolute time error ($\max|TE|$). We will discuss more in the time error budgeting subsection.

The operator has the flexibility of deploying one or more T-BCs between the T-GM and the CMTS. However, this option requires all network elements between the CMTS and the T-GM to be upgraded to be compliant with the G.8275.1 and G.8273.2. This could increase TCO. Alternatively, a T-GM can be collocated with the CMTS to avoid the upgrade.




Figure 10 – G.8275.1 + DTP

Additional requirements on the CMTS and the CM include:

- The CMTS must support the IEEE 1588-2008 slave
- The CM must support the IEEE 1588-2008 master
- The CMTS and the CM must implement DTP

9.2. Time error budget analysis

In Section 4 Table 2, we described a list of LTE and LTE-Advanced features with their corresponding time and phase synchronization requirements. In order to deploy these features in their small cell networks, operators need to carefully design their network to distribute overall maximum absolute time error budget over each network element. In this section, we perform a sample time error budget analysis for the timing distribution chain to show what an operator would need to do to deploy LTE TDD in their networks.

LTE TDD operating mode requires 3 μ s of phase synchronization between the adjacent home eNBs with radius of ≤ 500 meters. The following maximum absolute time error ($|TE|$) has been specified for each clock type by the ITU:

- $|TE| \leq 100$ ns for PRTC 0. This allocation also works for combined PRTC and T-GM function
- A constant time error $|cTE| \leq 50$ ns for Class A T-BC 0
- $|cTE| \leq 50$ ns for Class A T-TSC 0

Additionally, a time error budget of 250 ns is assumed for holdover for the entire distribution chain, and 200 ns is assumed for dynamic time error budget (see 0 Appendix V Note 2). The total time errors incurred in the distribution chain in Figure 11 is:

$$|TE_{PRTC}| + N_{T-BC} \cdot |TE_{T-BC}| + |TE_{DOCSIS}| + |TE_{T-TSC}| + dTE' + TE_{HO} = 1500 \text{ ns},$$

where dTE' denotes filtered dynamic time error (see 0 Appendix IV), and TE_{HO} denotes holdover error. Interested readers are directed to 0 for detailed discussion on constant, dynamic time error, and time error filtering.

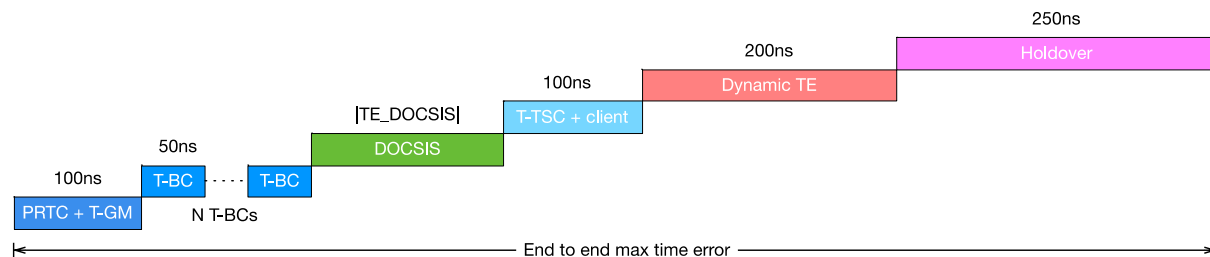


Figure 11 – Time error budgeting for networks with full timing support

DOCSIS 3.1 0 proposes 5 DTP system levels with varying degrees of CM-to-CM skew. Let us suppose an operator's DOCSIS backhaul network satisfies the performance requirement for a Level IV DTP System as defined in Table 10-9 of 0. This means that the operator can allocate a total of only two T-BC in between the T-GM and the DOCSIS network. In other words, the PRTC must be located within 2 hops away from the CMTS.

Alternatively, if the DOCSIS equipment can instead satisfy a Level III DTP System, the operator can allocate 8 T-BCs, or 8 hops between the T-GM and the CMTS. This provides the operator the flexibility in architecting their timing distribution network and reduces the number of grandmaster clocks the operator must deploy.

While the 5 DTP system levels have been defined, further work may be needed to continue to refine the time error budgeting for each of the HFC network elements.

9.3. Protection mechanisms during link failure

As shown in the reference architecture in Figure 9, a series of master-slave clock pairs forms the timing distribution chain that extends from the grandmaster clock to the packet slave clock, or the eventual client which is the eNB. Since packet-based synchronization protocols rely on constant exchange of sync messages between the master and the slave, a disruption on a particular link means timing distribution is interrupted at the eventual client. When this occurs, the T-BC whose upstream link is disrupted will inform the client that the reference signal is no longer traceable to a PRTC.

Two protection mechanisms can allow time and phase to be continuously delivered to the client: redundancy and holdover. A network operator can deploy multiple PRTCs at different locations to provide redundancy. For the DOCSIS backhaul network, this means a CMTS is configured with communication paths to backup PRTCs. The T-BC involved in the failure event will run its PTP best master clock algorithm (BMCA) to look for a new PTP path, and thereby help the client to find a new PRTC and lock to a new traceable reference signal. During this period of network rearrangement, the client's holdover mechanism can kick in to continue to generate clock from the last known traceable timing reference. Since holdover relies on the client's free-running local oscillator, time error can accumulate during this period of network rearrangement. A better way of maintaining time accuracy is to have the client lock to a physical layer frequency reference, i.e., SyncE.

In the absence of a backup reference source, the synchronization stack on the client will enter the holdover state. Since there is no backup plan, the client needs to maintain accurate timing for a longer period compared to the period of network rearrangement until the reference source can be recovered. The 3GPP or the ITU does not define the holdover time and accuracy limits. Instead, typically, an operator

specifies the time period and the accuracy limit for the equipment. Generally, a longer holdover period could be achieved with a higher quality oscillator. However, higher quality oscillators could be costly. While this does not cause an issue with macrocell deployments where cell radius is easily in the 10km range, the dense deployment in the small cell case can cause the cost of small cell equipment to escalate. Once again, having a PHY layer frequency sync support, i.e., SyncE, will improve the holdover performance.

9.4. G.8275.1 + DTP + SyncE

In addition to the requirements discussed in the “G.8275.1 + DTP” option in Section 9.1, each network element implements a Synchronous Ethernet Equipment Clock (EEC), as shown in Figure 12. While SyncE can improve system performance, it is worth noting that since SyncE requires specialized Ethernet hardware support, relying on SyncE will require the operator to replace all of its existing Ethernet equipment in the entire timing chain.




Figure 12 – G.8275.1 + DTP + SyncE

10. What if the network provides partial or no timing support?

In the ideal scenario that requires full timing support from the network, all network equipment must participate in the timing protocol and implement the G.8275.1 telecom profile. The timing distribution chain except the DOCSIS portion includes only T-BCs that are G.8273.2-compliant.

However, replacing and upgrading every network element to be compliant with G.8275.1 and T-BCs specs can become expensive, and the availability of equipment can be an issue. So, for operators who cannot upgrade every clock in their network immediately, an alternative option exists to still enable the delivery of frequency, ToD, and the phase synchronization needed to support LTE small cell deployments that implement LTE-A features and LTE TDD.

10.1. G.8275.2 + DTP (with or without SyncE)

In a network with partial or no timing support, non-participant or non-PTP-aware nodes, as well as T-TC(s) are allowed. In order to enable the delivery of frequency, time, and phase synchronization needed to service LTE-A techniques and LTE TDD, both the CMTS and the CM need to implement the G.8275.2 telecom profile and act as a T-BC-P. Note that the performance requirements for the T-BC-P node have not been formalized by the ITU.

As with the G.8275.1 option, the CMTS recovers time and frequency from the PTP messages, and translates the PTP timestamp into the DOCSIS timestamp. The CM regenerates the PTP timestamp based on the DOCSIS timestamp and passes it along to its downstream PTP slave which is the T-TSC-P. In this way, the CMTS-CM pair acts as a IEEE 1588 BC that terminates the 1588 domain at the CMTS, while the modem acts as a PTP master for the T-TSC-P in the end application client.

Since the performance of non-participant nodes are unknown, and most likely is worse compared to T-BC, the number of non-PTP-aware nodes, especially in a cascade, must be limited in a timing chain. One or more T-BC-P nodes can be placed strategically in the chain to reduce the effect of time error. To ensure required timing accuracy is achieved, the operator needs to perform proper testing, especially when the number of non-PTP-aware hops increases.

As discussed earlier, implementing SyncE at every node will improve frequency stability and holdover performance. However, requiring SyncE will require the operator to replace all of its existing Ethernet equipment with EEC-capable equipment in the entire timing chain – something the operator may have wanted to avoid in the first place by using G.8275.2 rather than G.8275.1. So, the options shown in Figure 13 and Figure 14 are well suited for installing a new Ethernet infrastructure for a new deployment region.




Figure 13 – G.8275.2 + SyncE + DTP




Figure 14 – G.8275.2 + DTP

11. What if DTP is not available

Since new modem silicon may be needed to implement DTP and PTP, none of the options discussed so far is deployable if an operator wishes to deploy the synchronization solution before the next modem silicon cycle. But even with the modem being a non-participant node in the timing distribution chain, depending on the availability of DTP on the CMTS, 2 options exist for distributing phase and ToD reference signals.

11.1. With CMTS participation

If the CMTS implements G.8275.2, it can act as a T-BC-P node to terminate the PTP timing messages from one of its upstream PTP-aware nodes such as a T-BC-P. The CMTS can then regenerate the PTP messages, sending them over the top through the CM directly to the T-TSC-P which could be part of the end application. In this case, the DOCSIS link is PTP and DTP-unaware.

Carrying the timing messages without the support of the network, and in case of DOCSIS, as regular data, will introduce a host of time errors due to PDV and network asymmetry, among other things. At least on the DOCSIS link, the PTP messages should be carried with unsolicited grant service (UGS), real-time polling service (RTPS), or high priority best effort upstream service flows to reduce the time error. Additionally, without DTP, MAC layer asymmetry cannot be determined.

Despite increased time error due to non-participant CMs, this option, shown in Figure 15, reduces the time error accumulated in the chain compared to the next option where the CMTS does not participate in the timing protocol.




Figure 15 – G.8275.2 with CMTS participation

11.2. Without CMTS participation

If the CMTS does not participate in the timing protocol, neither PTP nor DTP, the PTP messages must be sent while both the CMTS and the CM are timing-unaware. Due to the accumulation of time errors from non-participant nodes, this option reduces the number of hops the end application can be placed away from the T-GM.

This option as shown in Figure 16, does have some advantages over other options discussed so far. Since no additional requirement is placed on the CM and the CMTS, the option is deployable today.




Figure 16 – G.8275.2 without CMTS participation

While the options illustrated in this subsection are deployable today, neither provides compensation for the time error introduced by the underlying HFC network. Thorough testing is required to understand the time error budget for the DOCSIS link. Therefore, while it may be possible to provide a coarse level of frequency synchronization with these options, it is highly improbable to provide phase and time synchronization with these options.

12. Summary

Table 4 compares the options discussed in Sections 9 – 11.

Table 4 – Comparison of DOCSIS-based options

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Use Case	<p>Target solution for frequency/time/phase sync to small cells that implement LTE-A features and TDD.</p> <p>Allows install of 1 or more indoor and outdoor small cells on a new Ethernet infrastructure which supports SyncE, when improved holdover performance is required.</p>	<p>For delivery of frequency/time/phase sync to small cells that implement LTE-A features and TDD.</p> <p>Allows install of 1 or more indoor and outdoor small cells on a new Ethernet infrastructure which supports SyncE, when improved holdover performance is required.</p> <p>Allows cheaper T-GM install, if T-GM service is not directly provided by CMTS, or if T-GM is not collocated with CMTS.</p>	<p>Allows install of 1 or more indoor and outdoor small cells on an existing Ethernet infrastructure that does not support SyncE.</p> <p>If time error budget is large, CMTS does not need to be collocated with T-GM/PRTC.</p> <p>Allows cheaper T-GM install, if T-GM service is not directly provided by CMTS, or if T-GM is not collocated with CMTS.</p> <p>8275.2 could be used as a workaround until DTP and 8275.1 are fully supported by entire network.</p>	<p>Allows install of 1 or more indoor and outdoor small cells on an existing Ethernet infrastructure that does not support SyncE.</p> <p>If time error budget is large, CMTS does not need to be collocated with T-GM/PRTC.</p>

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Accuracy	<p>DTP solves DOCSIS asymmetry issue.</p> <p>Frequency sync through SyncE assists and improves time sync.</p> <p>SyncE reduces sync time after first contact, which means quick sync / fast re-sync.</p> <p>Implementing 8275.1 provides guarantees to achieve max time error per hop.</p>	<p>To ensure required accuracy is achieved, proper testing is needed. Time for testing increases when non-PTP-aware hops increase. But, with proper testing and tuning, it is possible to achieve same accuracy as 8275.1.</p> <p>DTP solves DOCSIS asymmetry issue.</p> <p>Frequency sync through SyncE assists and improves time sync.</p> <p>SyncE reduces sync time after first contact, which means quick sync / fast re-sync.</p>	<p>Requires short chain between T-GM and CMTS.</p> <p>Proper testing and optimization is required to keep time error low. Time for testing increases when non-PTP-aware hops increases.</p> <p>DTP solves DOCSIS asymmetry issue.</p>	<p>Requires short chain between T-GM and CMTS.</p> <p>Proper testing and optimization is required to keep time error low. Time for testing increases when non-PTP-aware hops increases.</p> <p>Since only small cells and T-GM implement 8275.2, DOCSIS MAC layer can introduce time error and link asymmetry that cannot be corrected. Use of RTPS or high priority BE services to carry 8275.2 traffic is a must.</p>
Operation	<p>SyncE extends holdover in case PTP fails. By using frequency and time, lower requirements on local oscillators, or can extend holdover time by using same oscillators.</p>	<p>SyncE extends holdover in case PTP fails. By using frequency and time, lower requirements on local oscillators, or can extend holdover time by using same oscillators.</p> <p>8275.2 provides failover and/or better holdover solution for DTP while transport is recovering from service disruption, e.g., after reboot of a CM or CMTS.</p>	<p>8275.2 provides failover and/or better holdover solution for DTP while transport is recovering from service disruption, e.g., after reboot of a CM or CMTS.</p>	<p>In case of link failure, does not need improved holdover time on local oscillators, since sync messages can be transported over IP, as soon as data link is re-established.</p> <p>Reduces the service disruption time after CMTS reboot, since lengthy sync time is not required when carried over IP without DTP.</p>

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Total Cost of Ownership	<p>Either T-GM needs to be collocated with each CMTS, or all elements between T-GM and CMTS must be upgraded, as 8275.1 requires compliance for every hop.</p>	<p>Depending on PDV, asymmetry, and # of hops, same hardware deployed today can be used for 8275.2.</p> <p>Allows for partial upgrade of select nodes, by providing a T-BC in strategic locations to lower time errors.</p> <p>Since SyncE is required, all elements in the timing chain must be upgraded.</p> <p>Also Ethernet switches in the premises must be upgraded to support SyncE.</p>	<p>Backward compatible with existing indoor infrastructure.</p> <p>Support use of very low cost equipment on sites.</p> <p>Works without upgrading any hardware or software in backhaul chain.</p> <p>Depending on PDV, asymmetry, and # of hops, same hardware deployed today can be used for 8275.2.</p> <p>Allows for partial upgrade of select nodes, by providing a T-BC in strategic locations to lower time errors.</p>	<p>Backward compatible with existing indoor infrastructure.</p> <p>Support use of very low cost equipment on sites.</p> <p>Works without upgrading any hardware or software in backhaul chain.</p>

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Time To Market	<p>Radio vendors support 8275.1 today.</p> <p>At least 1 modem silicon vendor supports PTP and DTP in hardware today.</p> <p>May requires new modem silicon to support SyncE master and PTP master.</p>	<p>Depending on radio vendor support of 8275.2, CM support of DTP and PTP, and when Ethernet equipment at the premise and in the timing chain can be replaced with SyncE.</p> <p>At least 1 radio vendor is testing 8275.2, and will support it by EOY 2017.</p> <p>May require new modem silicon to support SyncE master and PTP master.</p>	<p>At least 1 modem silicon vendor supports PTP and DTP in hardware today.</p> <p>CM needs to implement IP stack on CPE facing interface. This is likely a new board design, not a new silicon. But will require additional time.</p>	<p>At least 1 radio vendor is testing 8275.2, and will support it by EOY 2017.</p> <p>Since this solution does not require DTP or PTP support on CM, allows for deployment by EOY 2017.</p>

Conclusion

In this paper, we discussed the drivers for backhaul synchronization requirements needed for modern LTE and LTE-A networks, which are significantly more stringent compared to the traditional macrocell deployments. Although frequency, time, and phase synchronization can already be supported by today's technologies, to guarantee accuracy and holdover performance, operators must architect their network carefully with the right set of equipment, testing, and optimization. With the correct options, DOCSIS-based backhaul networks can support LTE-A features and LTE TDD deployments.

Abbreviations

3GPP	3 rd Generation Partnership Project
ABS	almost blank subframe
BC	boundary clock
BMCA	best master clock algorithm
CM	cable modem
CMTS	cable modem termination system
CoMP	coordinated multipoint

CS	coordinated scheduling
CSI	channel state information
DL	downlink
DTP	DOCSIS time protocol
EEC	Ethernet equipment clock
eICIC	enhanced inter-cell interference coordination
eMBMS	enhanced multimedia broadcast multicast services
eNB	evolved node B
EPC	evolved packet core
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HetNet	heterogeneous network
HFC	hybrid fiber-coax
ICI	inter-carrier interference
ICIC	inter-cell interference coordination
ITU	International Telecommunication Union
ITU-T	International Telecommunication Union-Telecommunication
LTE	long-term evolution
LTE-A	LTE advanced
LTE-FDD	long-term evolution frequency-division duplex
LTE-TDD	long-term evolution time-division duplex
MBSFN	multicast broadcast single frequency network
MNO	mobile network operator
NTP	network time protocol
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
P2P	peer-to-peer
PDCCH	PHY downlink control channel
PDSCH	PHY downlink shared channel
PDV	packet delay variation
PHY	physical
ppb	parts per billion
PRC	primary reference clock
PRTC	primary reference time clock
PTP	precision time protocol
RC	resource coordinator
RTPS	real time polling service
SC-FDMA	single-carrier frequency division multiple access
SINR	signal-to-interference-plus-noise ratio
SyncE	synchronous Ethernet
T-BC	telecom boundary clock
T-GM	telecom grand master, master clock only
T-TC	telecom transparent clock
T-TSC	telecom time slave clock
T-BC-P	telecom boundary clock-partial
T-TC-P	telecom transparent clock-partial

T-TSC-A	telecom time slave clock-assisted
T-TSC-P	telecom time slave clock-partial
TC	transparent clock
TDD	time division duplex
TCO	total cost of ownership
ToD	time of day
UE	user equipment
UGS	unsolicited grant service
UL	uplink
UTC	universal time coordinated

Acknowledgement

The authors would like to thank Yi Tang and Zheng Lu, both of Cisco Systems, for reviewing the white paper and providing valuable comments. The lead author would like to thank the following people involved in the discussions of the DOCSIS-based mobile synchronization architecture: Alvaro Simon of Vodafone Spain, Tilco van Dijk of Liberty Global, Pedro Antão of NOS Portugal, Robert Grimm of Vodafone DE, and Bruno Cornaglia of Vodafone Italy. Additionally, the lead author acknowledges Peter Percosan of Digital Strategy, Volker Leisse and Thomas Nogues of CableLabs for facilitating the work.

Bibliography & References

John T. Chapman, et. al., “The DOCSIS Timing Protocol (DTP), Generating precision timing services from a DOCSIS system,” Proceedings of INTX/SCTE Spring Technical Forum, 2011.

John T. Chapman, Jennifer Andreoli-Fang, “Low Latency Techniques for Mobile Backhaul over DOCSIS,” to appear in *Proc. of SCTE Fall Technical Forum*, October 2017, Denver.

3GPP TS 36.104, “Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception.”

3GPP TS 36.133, “Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management.”

IEEE Std 1588-2008, “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.”

ITU-T G.8271, “Time and phase synchronization aspects of packet networks.”

ITU-T G.8271.1, “Network limits for time synchronization in packet networks.”

ITU-T G.8272, “Timing characteristics of primary reference time clocks.”

ITU-T G.8273.2, “Timing characteristics of telecom boundary clocks and telecom time slave clocks.”

ITU-T G.8275, “Architecture and requirements for packet-based time and phase distribution.”

ITU-T G.8275.1, “Precision time protocol telecom profile for phase/time synchronization with full timing support from the network.”

ITU-T G.8275.2, “Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network.”

“Data-over-cable service interface specifications, DOCSIS® 3.1, MAC and upper layer protocols interface specification,” I-11, issued 2017-05-10, CableLabs.

Mobile Backhaul over DOCSIS

A Technical Paper prepared for SCTE•ISBE by

John T. Chapman
CTO Cable & Fellow
Cisco
San Jose, CA 95124
408-526-7651
jchapman@cisco.com

Jennifer Andreoli-Fang
Distinguished Technologist
CableLabs
Boulder, CO
303-661-3838
j.fang@cablelabs.com

Introduction

One of the next big opportunities for cable operators is mobile. For a long time, cable and mobile have been competitors, both competing for that voice or Internet subscriber, or even a video subscriber. Now, they could become the best of friends. To move into small cell deployment, mobile operators need someone with a plant that can backhaul their user traffic. Cable can be that partner. Cable, meanwhile, has already moved in a direction of wireless as the last hop, and LTE is just another access technology for that last hop.

In this white paper, we examine DOCSIS as a viable option for mobile backhaul technology, particularly for small cells. The business case is that when the small cells are deployed deep in the network, DOCSIS will be there and fiber may not be. Should money be spent on installing new fiber, or can LTE use DOCSIS instead? Judging by all the WiFi deployments being backhauled by cable operators, it makes sense to consider DOCSIS as a backhaul for LTE as well.

There is money to be made and money to be lost. As always, it all depends on how much money it takes to build and operate the network. Keeping the costs down is a great motivation and is something DOCSIS has always taken notice of.

The Cable-Mobile Market

1. Market Opportunity

A Mobile Network Operator (MNO) owns their own plant and sells services. A Mobile Virtual Network Operators (MVNO) rents another operator's plant and sells services. About 50% of CableLabs members are already MNOs or MVNOs. The other 50% are trying to figure out if they should take the plunge as well.

It makes sense. It is difficult to grow revenues significantly in cable when data/voice/video services already have high market penetration and success. Yet, consumers pay similar amounts for cable and mobile services. Most consumers are paying for Internet access from both cable and mobile! By combining cable and mobile services, operators can increase revenue and attract consumers with an attractive combined service at a potentially lower cost.

Today, mobile operators rely on macrocells to provide Long Term Evolution (LTE) services. Macrocells cover areas on the order of one to ten miles in radius. The next step in the evolution of the mobile RF plant is to deploy small cells, and do so in a denser manner. Small cells have an effective radius ranging from 100 meters to 500 meters. Coverage varies and is based upon many factors such as transmitting frequency, antenna height, transmit power and line-of-sight. Small cell footprint also includes both indoor and outdoor coverage.

The movement from macrocell to small cell is analogous to the hybrid fiber-coax (HFC) movement from N+5 (node plus 5 amplifiers) to N+0 (node plus zero amplifiers). Both are segmentations of the physical plant to increase available bandwidth per subscriber and both are happening in the marketplace.

A cable operator could enter the mobile market in a sequence of steps. They rent mobile air time from an incumbent MNO with macrocell-only plant to provide coverage. The cable operator could then begin to

build out a small cell footprint. In doing so, they would decrease the macrocell airtime they need to rent. He could then also rent out the small cell footprint back to the same MNO he is renting the macrocell from. This would reduce the operational expenses. Over time, when the small cell network is built out, the rent-back small cell revenue might equal the rent for the macrocells. Now the business case is cash neutral and the cable operator has built a small cell pant with the help of the rent revenue form the MNO. The next step is for the cable operator to move from being an MVNO to being an MNO.

2. Basic Deployment Requirements for Small Cell

There are at least three major items that are required for a small cell deployment:

1. Location
2. Power
3. Backhaul

The cable operators are in an advantageous position when it comes to offering a solution for all three of these requirements.

1. Location

Cable operators have access to public rights-of-way that can help when placing new infrastructure but that is only true on an overhead plant. Underground plants can be tricky as there is limited antenna access. After that, there is “urban furniture” that works – consumer roof tops, lamp posts, street lights, utility poles, etc.

2. Power

The HFC plant is already powered and can support a small number of small cells. Current plants may not be able to support a large number of small cells but power distribution over coax is flexible and more power can be added if new small cell deployments require it and a business case is made. The responsibility to power the small cell power might even be borne by the consumer, in a residential or business setting where the consumer is looking add service or improve coverage.

3. Backhaul

The HFC plant has two resources: fiber and coax.

Fiber

Today, signals transmit over fiber are analog-based but fiber communications will soon be upgraded to Ethernet (digital fiber) - a requirement for DOCSIS Remote PHY. Wherever there is fiber, small cells can easily be connected.

Coax

Coax is also a very important resource because small cells may not be located near the node. The best location for nodes and small cells don't always coincide because small cells typically need to be placed above roof tops, tree lines and hills to provide optimal LTE signal coverage. Luckily, coax is

nearly everywhere already, especially in residential areas where coax, and thus DOCSIS, is almost always available and fiber may not be.

Imagine a business scenario where the cable team decides it is not cost effective to upgrade a neighborhood from N+5 to N+0 (fiber deep). Then the mobile team from the same operator comes along and deploys fiber to the home to support the small cells. That would be a good example of siloed business economics – where two divisions arrived at opposite business cases instead of a blended business case.

3. DOCSIS vs. Fiber

Throughput

Fiber currently achieves a throughput of 10 Gbps. DOCSIS is competitive with fiber. For example, DOCSIS 3.1 enables throughputs as high as 10 Gbps in the downstream and 5 Gbps in the upstream, using FDX DOCSIS.

With D3.1 & FDX:




Figure 1 – FDX Optical Node

To achieve a comparable DOCSIS throughput to fiber throughput, a node could be deployed as shown in figure 1. Two return paths can be used with FDX DOCSIS on each return path to achieve a bandwidth of 10 Gbps x 10 Gbps - the same bandwidth as fiber!

These DOCSIS throughputs over coax give operators a lot of flexibility. The operators can deploy the small cell anywhere in the HFC footprint and then choose fiber or coax, based on convenience and cost.

Network Timing

Small cells located on the outside plant have the option to receive network timing directly from GPRS. Indoor small cells, for both residential and commercial, don't have direct GPRS access and there would need to receive network timing over DOCSIS. Timing for mobile backhaul over DOCSIS is addressed in 0 0.

Latency

In addition to sufficient bandwidth, a backhaul also needs timing and the ability to offer low latency to high priority traffic. In this paper we will address how to achieve low latency backhaul over DOCSIS.

DOCSIS and LTE Working Together

4. Comparing DOCSIS and LTE Latency

If you are familiar with how DOCSIS works, then you will be able to pick up LTE quickly. Both DOCSIS and LTE are trying to solve very similar problems. Both technologies try to manage a point to multipoint network. Both LTE and DOCSIS use a scheduled upstream and both are managing subscriber traffic.

The main difference between LTE and DOCSIS is that LTE manages a wireless network, while DOCSIS manages a wired network and the protocols are different since they were designed by different committees. The wireless channel and mobility aspect of LTE adds challenges that don't need to be dealt with in DOCSIS but once connected, scheduling and data transfers between customers and the (evolved NodeB) eNodeB and the CMTS follow similar principles.

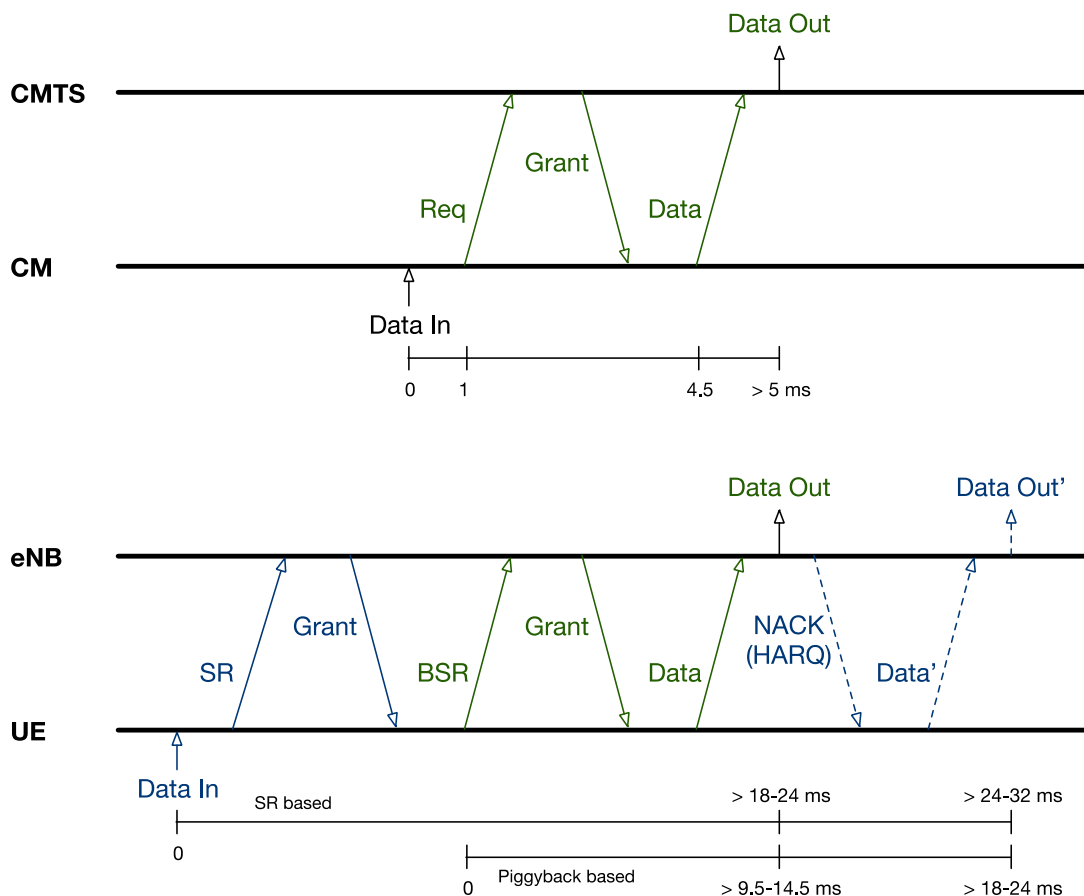


Figure 2 – Comparing DOCSIS and LTE

Duplexing

Another difference between LTE and DOCSIS is that LTE supports frequency division duplex (FDD) and time division duplex (TDD). FDD is similar to HFC in that uplink (UL) & downlink (DL) are segregated by frequency bands. In TDD, essential to massive MIMO antenna operation, the same band is used for UL and DL and the interleaving of the UL and DL transmissions in the same spectrum adds access delay for UL transmissions. Consequently, TDD round trip time (RTT) can be much higher than FDD.

DOCSIS Upstream Scheduling

The DOCSIS upstream uses a request-grant approach as shown in Figure 2. A request message is sent from the cable modem (CM) to the cable modem termination system (CMTS). The CMTS prepares a DOCSIS MAC management message (MMM), called the MAP, and inserts an entry indicating a grant for the CM. The grant entry in the MAP message contains an upstream service identifier (SID) associated with a service flow assigned to the CM, a transmission time signaled as a mini-slot number, and the number of bytes to be transmitted. The CMTS transmits the MAP to the CM, so that the CM can make use of the grant to send its upstream data to the CMTS.

There are many factors that determine the request-grant time in DOCSIS. The DOCSIS specification contains a detailed analysis of the request-grant delay. The minimum delay is basically every second MAP time plus some circuit and queuing delay. The DOCSIS 3.0 CMTSs currently use a 2 ms MAP time. One published set of test results for a Cisco uBR10K CMTS showed a 5 ms minimum request-grant response time.

The CMTS scheduling algorithm can have a significant impact on request-grant delay. For instance, if the CMTS is using a best effort (BE) scheduling algorithm, requests can be made in contention slots where the requests could fail on the first try. If the CMTS is using the real-time polling service (rtPS) scheduling algorithm, the request will be placed in a dedicated slot which ensures that the requests are always successful. BE can perform better than rtPS when DOCSIS is idle because many request contention slots are available. In that case, BE provides lower latency than rtPS. However, in a busy system where there are fewer contention request slots and BE will need to re-request often, rtPS can provide guaranteed latency. Requests can also be sent as a piggyback message with a data packet. Piggybacking is deterministic within a flow and avoids contention.

LTE Uplink Scheduling

At the heart of the LTE uplink scheduler is a similar request-grant mechanism. The request is sent from the user equipment (UE) to an eNodeB. The UE is typically a cellular phone while the eNodeB is a macrocell or a small cell. When data arrives at the UE, the UE first determines if it already has an LTE uplink (UL) grant. If it does not, the UE waits for an opportunity to transmit a scheduling request (SR), where typical SR opportunities can come along once every 1 to 10 ms depending on the configured periodicity of the SR. The purpose of the SR is to keep the upstream signaling overhead low which is critical in wireless where spectrum is costly.

Upon receiving the SR, the eNodeB schedules an UL grant so that the UE can transmit a buffer status report (BSR) to the eNodeB. Once the eNodeB receives the BSR from the UE, it becomes aware of the outstanding UL data present in the UE's upstream queues. The eNodeB then schedules UL grants and transmits the grants to the UE via a downlink control information format 0 (DCI-0) message transmitted in the physical downlink control channel (PDCCH). LTE operates every subframe which is 1 ms. However,

due to processing constraints, the eNodeB's BSR-grant delay is typically 4 ms for both the eNB and the UE. In other words, upon receiving the BSR, the eNB performs scheduling and sends DCI-0 to arrive at the UE in 4 subframes, and the UE is scheduled to transmit 4 subframes later. Assuming a very small SR opportunity periodicity (1 ms), the minimum delay before a UE can transmit UL data is 18 ms which is slightly longer than the minimum DOCSIS request-grant latency. Detailed request-grant latency calculations are shown in Table 1.

Table 1 – LTE REQ-GNT Latency

Latency Components	Latency (ms)	
	SR	Piggyback
Waiting time for SR (assume configured SR period of 5 ms)	0.5 – 5.5	n/a
UE sends SR, eNB decodes SR, eNB generates grant for BSR	4	n/a
eNB sends grant, UE processes grant, UE generates BSR	4	0-4
eNB processes BSR, eNB generates grant for data	4	4
eNB sends grant, UE processes grant, UE sends UL data	4	4
eNB decodes UL data (estimate)	1.5 – 2.5	1.5 – 2.5
Total	18 – 24	9.5-14.5

It is worth noting that both the DOCSIS and LTE system have design efforts underway to decrease the minimum latency, so the latency figures given in Table 1 are for currently deployed systems. It is also worth noting that in a congested network, both DOCSIS and LTE will have longer latencies.

Re-transmissions

The LTE transmissions over the LTE air interface often occur in a harsh wireless environment (eg. cell-edge). Wireless signals are easily degraded due to path loss and interference and transmission errors are far more likely than in DOCSIS. To overcome the transmission errors, LTE has two retransmission mechanisms:

1. Hybrid Automatic Repeat reQuest (HARQ) which operates at the MAC layer.
2. Automatic Repeat reQuest (ARQ) which operates at the radio link control (RLC) layer.

HARQ

HARQ is intended to quickly re-transmit the LTE transport blocks to recover from most errors in conjunction with a good forward error correcting (FEC) algorithm, while RLC ARQ is a higher layer re-transmission mechanism with which has higher overhead but improves the reliability of the link after HARQ.

RLC ARQ

The use of RLC ARQ is optional in LTE and depends on the type of traffic and quality of service (QoS) parameters (error rate, latency, bit rate, etc) desired. For example, for voice over LTE (VoLTE) packet

loss is not that critical, but latency is. The human ear can tolerate small amounts of packet loss but long delays quickly become annoying during a phone conversation. Therefore, for voice traffic, ARQ retransmissions are not required and a LTE EPS bearer for voice can be configured to use the RLC Unacknowledged Mode (UM) entity which does not use ARQ. On the other hand, transmission control protocol (TCP) traffic packet loss triggers slow start, and throughput can suffer. However, longer delays can be tolerated by TCP traffic therefore ARQ is used by configuring the LTE EPS bearer to use the RLC Acknowledged Mode (AM) entity.

The purpose of an additional PHY layer retransmission mechanism is to reduce latency. After transmitting a coded block of bits which is termed a Resource Block (RB), the transmitter, the UE, or the eNB, keeps the RB in the transmission buffer. The receiver PHY layer decodes the RB, and passes the CRC results to the MAC layer. The MAC layer then issues either a HARQ ACK or NACK based on the results. Between the time the RB is received at the receiver and the time a HARQ feedback must be received at the original transmitter of the RB, there is an exact 4 ms of delay.

Overall LTE and DOCSIS Latencies

In summary, if we consider only lightly loaded systems and FDD duplexing for LTE, DOCSIS 3.0 has a minimum request-grant delay of about 5 ms while 4G LTE has a minimum request-grant delay of 18 to 24 ms without re-transmission and 26 to 34 ms with one HARQ retransmission. These latency values will increase under higher loads or if TDD is used in LTE.

Applicability of the Work in This Paper

This paper presents a method to reduce DOCSIS backhaul latency for LTE using pipelining and a new inter-system request message called the bandwidth report (BWR). We focus on DOCSIS 3.0 and 4G LTE since both technologies are well documented and widely deployed and the BWR can be used today. The research covers both LTE FDD and TDD, but we only include FDD timing examples for simplicity. The general principles proposed in this white paper will extend to DOCSIS 3.1 and 5G LTE, although some modifications to the protocols may be required.

5. The Bandwidth Report Concept

LTE request-grant time is much longer than DOCSIS, yet the mobile operators often demand backhaul delays to be on the order of a few milliseconds. There is a hidden opportunity to this high LTE latency. What if we could use the long LTE latency to start DOCSIS processes in parallel and reduce the overall LTE-DOCSIS system latency? This concept is shown in Figure 3 and is made possible using the BWR message.

Current LTE-DOCSIS systems, which do not have BWR, have cumulative latency as shown in the upper diagram of Figure 3. When the BWR is introduced into the LTE-DOCSIS system, the DOCSIS request-schedule-grant loop can be started earlier and in parallel with the LTE request-schedule-grant loop, leading to much lower latency, as shown in the lower diagram of Figure 3. This is accomplished by treating the LTE and DOCSIS systems as one pipelined system rather than the two independent systems they are today. For this pipeline mechanism to work, DOCSIS needs some information from LTE in the form of the BWR message.




Figure 3 – Reducing DOCSIS Effective Latency Under LTE

How does the BWR work?

The BWR allows an external system, such as a LTE eNodeB, to request bandwidth from a DOCSIS system for some specific time *in the future, before* the arrival of the actual traffic. The eNodeB provides a future traffic profile through the BWR message that allows the CMTS to make QoS and granting decisions earlier than it normally would.

The BWR replaces the CM's internal layer 2 request message for LTE uplink data arriving at the CM. The BWR message itself is an external layer 3 IP based message that is transmitted from the eNodeB to the CMTS, through the CM. The content of the BWR is populated with information describing future data that will arrive at the CM.

The BWR is created by the eNodeB's LTE UL scheduler just after the scheduler finishes granting the UE(s) based on the UE's outstanding UL data buffer sizes. The BWR message created after scheduling contains a description of the amount of UL data that is expected to arrive at the eNodeB and be forwarded to the CM. The BWR message is then transmit from the eNodeB to the CMTS, via the CM, in a dedicated upstream service flow earlier than the CM would normally issue its request.

The result of using the BWR is that DOCSIS is made aware of the LTE scheduler's scheduling decisions using the BWR message. Rather than having two separate and sequential latency-additive LTE and DOCSIS request-grant loops, the BWR starts DOCSIS's request-grant loop early and in parallel to LTE's uplink granting. The CMTS scheduler grants the CM just-in-time for the UE's upstream data to arrive at the CM, which the effect of reducing the upstream latency. This is illustrated in Figure 4.




Figure 4 – Pipelining Requests and Grants

Figure 5 depicts an LTE system backhauled by DOCSIS. The resemblance between the LTE and the DOCSIS systems are striking. In both systems, the UE/CM generates a request when they have data to send, and the scheduler in the eNodeB/CMTS responds with a grant.

When the BWR is used, the LTE system's predictor tries to describe the data flow across the Ethernet interface by estimating the data flow across the air interface using the LTE UL scheduler's grants. Unfortunately, the data flow across the air interface will not exactly match the data flow that will occur

later across the Ethernet interface due to segmentation and decoding failures, which will be explained in greater detail in section 6.

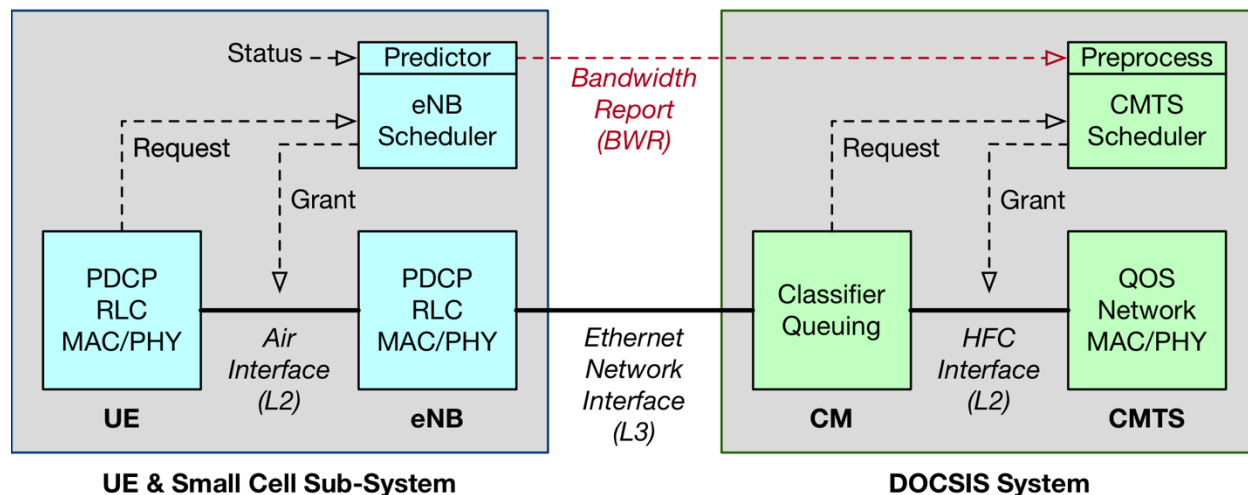


Figure 5 – Linking DOCSIS and LTE Schedulers with BWR

To adapt the BWR message to the traffic flow, a prediction algorithm is used to monitor the operation of the eNodeB scheduler. The predictor looks at the status of the data path to see if there are any changes in behavior and then issues the BWR message. The optimal period for BWR for LTE FDD is 1 ms because the eNodeB scheduler makes scheduling decision once every LTE subframe (every 1 ms). The CMTS uses the information given in the BWR messages, along with other scheduling constraints, and generates grant(s) for the CM.

BWR Example

If a 1000-byte data transfer from the UE will occur in 8 ms and there is an anticipated 1 ms delay in the eNodeB software stack for transport block decoding and packet reassembly before transmission to the CM, a BWR will be transmit to the CMTS to request 1000-bytes in 9 ms. The CMTS could be configured to add an additional 2 ms to the grants generated based on the BWR to ensure they any delayed traffic leaving the eNodeB. Therefore, the CMTS would issue a DOCSIS grant 11 ms after the BWR is created.

That is the basic idea of BWR. Of course, nothing is quite that simple ...

6. BWR Additional Design Considerations

6.1. Synchronization and Timing

If the LTE system is scheduling in the future and the DOCSIS system is granting in the future, both systems must be synchronized in time for the grants to line-up. When high accuracy synchronization, such as GPS is not available or practical, another method to achieve good synchronization like the IEEE 1588 protocol is needed.

For DOCSIS to understand the LTE BWR message and issue grants at the correct time, a common time system must be referenced by the BWR. One proposal is to use a time index (TI) field in the BWR that

stores the time in a similar way as IEEE 1588 timestamps with a reduced resolution of 1 ms to keep the size of the bit field short. With a 1 ms resolution, a 16 bit field could be used for the TI field, for example, with a defined roll-over procedure.

Another proposal is to communicate the LTE frame and subframe to the CMTS. LTE frames are numbered using 10 bits, from 0 to 1023, with rollover beyond 1023, while subframes are numbered 0 to 9 and can be represented using 4 bits. This scheme also requires a total of 16 bits.

The CMTS would translate the TI value to a mini-slot and issue a grant for the upstream data.

6.2. DOCSIS Mini-slot and LTE Subframe Misalignment

LTE and DOCSIS use different framing mechanisms for their transmissions. For instance, the LTE air interface transmissions are based on frames and subframes. A subframe is 1 ms in duration and a frame is 10 ms in duration. Meanwhile, DOCSIS uses a MAP interval of approximately 2 ms and grants are issued on a mini-slot level of granularity. The two framing systems will not be aligned in time. This is illustrated in Figure 6. In essence, the eNB scheduler and the CMTS scheduler are two asynchronous systems that are now trying to communicate scheduling information.

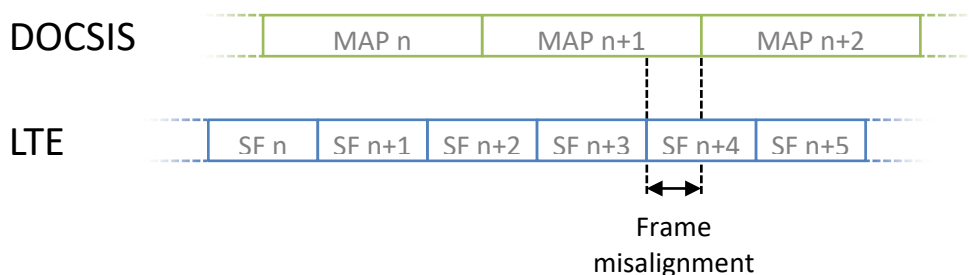


Figure 6 – LTE and DOCSIS Frame Misalignment Despite System Time Synchronization

It is often hard to schedule grants at exact times in a DOCSIS MAP interval because doing so may excessively fragment other flows. Thus, DOCSIS can be configured to allow the CMTS's upstream scheduler some flexibility during scheduling, which results in some grant jitter. This means that the selected mini-slot may not align perfectly with the upstream egress data from the eNodeB.

The CMTS needs to account for all the grant and upstream data egress jitter by adding some engineering margin, otherwise the grant may miss some packets and latency will increase.

6.3. HARQ

If the data from the UE is not received correctly at the eNodeB, the eNodeB will transmit a negative acknowledgement (NACK) to the UE which will trigger data retransmission. This means that the grant requested in the BWR message will not get used for the intended data transfer on the DOCSIS upstream link. The grant can get used by other data, if other data has been delayed, but the grant may also not get used and therefore some DOCSIS bandwidth will be wasted. Refer to 0 for an upper bound on the unused DOCSIS grant.

Since the LTE scheduler knows when a transport block failed to be decoded successfully, the LTE BWR predictor knows when re-transmissions are going to occur. If the LTE transport block is successfully

received when it is re-transmitted to the eNodeB, some packets may egress from the eNodeB. The predictor must then account for the delayed and re-transmit data by adding the byte to a BWR corresponding to the subframe that includes the re-retransmission. More details about this operation and its performance can be found in 0.

6.4. RLC Packet Reassembly

The packets that are transmitted from the UE to the eNodeB are segmented (if required) and placed into LTE transport blocks. If a packet fits entirely within a LTE transport block, the UE can transmit it entirely in one subframe. When the eNodeB correctly receives a complete packet, the packet could get passed immediately up the eNodeB's software stack and out the eNodeB's Ethernet interface.

In the case where not all the segments of a packet have been received, the segments will be stored in a buffer at the RLC layer until all the segments are received. Once all the segments are received, the RLC entity re-assembles the segments into a complete packet and the packet can finally be transmitted out the eNodeB's Ethernet interface.

For an example of the effect of segmentation and re-assembly, let us assume that a packet is segmented into 2 pieces across 2 LTE transport blocks and both are received successfully (no decoding failures). The 1st transport block is transmitted from the UE to the eNodeB but it only contains the 1st segment of the packet. Since the packet is incomplete, the RLC entity handling the EPS bearer will store the segment and wait for the next transmission. 1 ms later in the next subframe, the 2nd transport block is received at the eNodeB physical layer and after decoding the RLC entity receives the 2nd segment of the packet. RLC reassembles the complete packet and the packet is sent out the eNodeB's Ethernet interface.

In our example, the result of the packet fragmentation that occurred on the LTE air interface added an additional 1 ms delay to the time when packet was expected to leave the eNodeB. While this doesn't seem like much, consider that for each additional segment an additional 1 ms delay (eg. $N_{frag} - 1$ milliseconds) is added to the packet egress time and a packet that is late even by a small delay may miss a DOCSIS grant. This is shown in Figure 7.

This packet fragmentation is one of the causes of mismatches between the bandwidth requested by the BWR message and the actual data leaving the eNodeB's upstream Ethernet interface.

Predicting the packet fragmentation and reassembly delay is a difficult problem. The LTE predictor algorithm is predicting the future while the reassembly engine works in the present. If the LTE predictor has packet level visibility when scheduling, then it can address this problem. If it does not then some engineering margin can be added by the predictor to the BWR message time index field to capture segmented packets.

It may appear that the solution is for the BWR predictor to increase the engineering margin in the BWR message. However, there is a tradeoff between increasing the engineering margin to capture more upstream packets and upstream latency gain. The engineering margin that maximizes the upstream latency gain of the BWR feature doesn't necessarily correspond with engineering margin that captures 100% of the delayed packets. So careful BWR predictor and engineering margin selection is needed!




Figure 7 – Packet Delay Due to Packet Segmentation on the LTE Air Interface

6.5. eNodeB Ethernet Queuing and Transmission Delay

All physical devices have speed and bandwidth limitations and queues to help buffer data when the load on the interface is high. In a QoS based system, there will likely be multiple queues serving traffic with different priorities. The queue depth depends on the traffic and the time it takes for a packet to arrive at its destination depends on the time the packet waits in the queue and transmission on the Ethernet interface.

The LTE predictor must allow for the queueing and transmission time. If the predictor knows which QoS level each packet flow is, it can allocate shorter times for the high QoS traffic and longer times for lower QoS traffic to compensate for the time the packet could spend in a queue.

The Ethernet packet transmission time also adds delay. Luckily the transmission delay is small compared to other delays. For example, 1518-byte Ethernet packet on a gigabit interface only takes about 12 microseconds.

6.6. Signaling Traffic

Real-Time Signaling Traffic

LTE signaling could be categorized into the following three categories. This generalization applies to all signaling packets that leave the Ethernet interface of the eNodeB, whether they are intended for another eNodeB or intended for the evolved packet core (EPC).

1. Fully Scheduled Signaling

UE signaling traffic transmit to the network that is accounted for by the eNodeB scheduler.

2. Partly Scheduled Signaling:

eNodeB signaling transactions that have one or more transactions to the UE followed by one or more transactions to the Ethernet interface. This signaling is only partially accounted for by the eNodeB scheduler.

3. Unscheduled Signaling

Signaling traffic originating at the eNodeB, which is not accounted for by the eNodeB scheduler.

The predictor algorithm will have to account for each of these three cases.

Periodic Signaling Traffic

Many signaling messages are time critical and must be transmit immediately. Some signaling messages may be periodic and quite predictable.

Periodic signaling messages are interesting because the grants described by the BWR could be described like LTE describes semi-persistent scheduling (SPS) grants. For example, instead of transmitting 100x 60-byte grants for voice packets arriving in the UE's transmit queue every 20 ms, LTE issues a single SPS grant to the UE for 100x 60-byte grants spaced out by 20 ms. As long as the signaling traffic is predictable and timely, a SPS version of the BWR could be used by the eNodeB to ask the CMTS for periodic grants of a given size. The benefit of the SPS-BWR would be a reduction in BWR messaging overhead.

6.7. BWR Prediction Error

The predictor algorithm that creates the BWR message at the eNodeB is constantly trying to predict the upstream traffic that will egress from the eNodeB several subframes in the future. Despite the BWR predictor's best efforts, we explained in section 6.3 and section 6.4 that HARQ failures and packet segmentation can cause differences between the amount of expected data and the actual data the leaves the eNodeB. In addition, timing errors such as poor synchronization, DOCSIS grant jitter and packet egress jitter can cause the packets to miss the grants. The accuracy of the BWR predictor algorithm and error correction algorithm could be differentiating aspects for eNodeB and CMTS vendors.

For instance, on the eNodeB side, an in depth understanding of the eNodeB's own software processes is a good start to developing a good BWR predictor algorithm that creates accurate BWRs. For example,

knowing timing statistics about the LTE UL stack and Ethernet device can be helpful to predicting when packets will egress the eNodeB, if hard real-time code cannot be written to ensure packets egress at an exact time.

On the DOCSIS side, the CMTS can observe the predicted behavior and actual DOCSIS grant usage. The CMTS can then help to manage the error by either increasing the buffer delay in the CM or by doing active buffer management on the CM and sending extra grants when the buffer increases in size.

When errors do occur, the BWR itself can be used to address inaccuracies. For instance, when transport blocks fail to be decoded successfully, the eNodeB re-requests enough bytes for the future re-transmission.

The eNodeB predictor could be even more proactive and monitor the amount of data leaving the Ethernet device. When the actual upstream data differs from the predicted upstream data, the eNodeB could communicate this difference to the CMTS by subtracting the bytes from the TI it originally requested and add the bytes to a new TI.

6.8. BWR Message Flow

The BWR is a signaling message that gets sent to the Ethernet interface. Depending on how the queuing is done, it may need to be assigned a separate QoS level from data. If the BWR goes to a separate queue on the CM with a rtPS or unsolicited grant service (UGS) flow, then the byte count of the BWR message it does may need to be accounted for in the service flows that carry the EPS bearers and thus in the BWR message itself. If BWR is given high priority and placed on a common service flow, then it does have to be accounted for.

The LTE-DOCSIS system may have to account for the network bandwidth used by the BWR message. This will depend upon the CM queuing configuration.

6.9. BWR Message Delay

If the BWR message is predicting 8 ms in the future, then in theory it has 8 ms to complete its journey from the eNodeB to the CMTS and for the CMTS to issue a grant. In practice, the CMTS will need the BWR message before it runs its upstream scheduler and MAP builder routine. That MAP routine is run early depending on the MAP advance time that the CM requires. If the MAP interval is 2 ms and the MAP advance time is 2 ms, and there is some margin added, then the BWR might need to arrive 5 ms early. If the eNodeB takes 1 ms to generate and transmit the BWR, then that only leaves 2 ms of margin.

This means that it is important that when the BWR message arrives at the CM, the BWR is transmit to the CMTS as quickly as possible. Additional latencies, such as contention in a request slot, will cause the BWR to arrive at the CMTS too late. To avoid contention delays, DOCSIS scheduling algorithms such as rtPS or UGS can be used.

Since the BWR can be a variable length message, UGS may not be the best choice unless the size of the BWR can be fixed or the UGS grant is made larger or equal to the length of the longest BWR message. More advanced scheduling mechanisms may be needed at the CMTS such as waterfall granting which is under evaluation by the DOCSIS committee 0.

6.10. BWR Message Loss

One option for sending BWR messages over Ethernet and DOCSIS is UDP encapsulation. While UDP adds little overhead compared to TCP, UDP should not be considered reliable.

Improving BWR Reliability

To make the LTE-DOCSIS system more robust, the BWR could include a sequence number to help the CMTS detect a dropped BWR message. Note that if the eNodeB has no data to send, it may not send a BWR message to save processing time and DOCSIS bandwidth, for example. Therefore, a TI field alone could not be used to detect dropped BWR messages.

Many communications systems use retransmissions to ensure data is received at the receiver. In the case of the BWR, there is not enough time to detect BWR message loss and re-transmit the original BWR message since a margin of only 2 ms is left in the system as described in section 6.9, therefore re-transmission is not a viable solution.

A better solution may be to duplicate the requested bytes per TI and LCG within a BWR message over “N” BWR messages. The number of duplicate messages “N” can easily be configured at the eNodeB but the proposed value for N is three. The duplication of the requested bytes over three messages keeps the total bandwidth requirements low while helping to increase the reliability of the LTE-DOCSIS system.

6.11. DOCSIS Grant Duplication Due to CM Requests

With the introduction of BWR, there are now two request mechanisms – the external layer three BWR request mechanism and the native internal layer two CM request mechanism.

When data packets from the eNodeB arrive at the CM, the CM may generate a request message (REQ) if the grant from the BWR is not immediately available at CM. When the CM BE flow requests its own grant, two grants (REQ + BWR) may be generated for the same data which could lead to wasted bandwidth if those duplicate DOCSIS grants are not used.

The CM BE flow request behavior is dependent on the CM hardware implementations and may be hard to predict. Based on some observations done during Cisco-CableLabs R&D efforts, it appears that the CM may not issue a duplicate request if the grant from the BWR is scheduled and received by the CM within 1-2 ms after upstream data arrives at the CM.

In theory, the BWR message contains all the request information that is needed for the data path so the additional requests from the CM BE flow could be disabled or ignored. In practice, mismatches between the BWR’s predicted bandwidth request and the actual upstream data from the eNodeB can occur, and the CM’s BE flow’s requests can be useful in managing buffer build-up.

The solution is to keep the CM’s BE flow’s requests enabled and ensure that the BWR is scheduled in a timely manner to prevent the CM from making duplicate requests. This means good synchronization and accurate prediction of when the data will leave the eNodeB and arrive at the CM. In addition, the CMTS may be able to help reduce duplicate messages by detecting the duplicate BE flow requests and ignoring the duplicates, if an intelligent CMTS algorithm can be developed for this purpose.

6.12. Checks and Balances

In order for the BWR concept to work well there needs to be a certain level of trust between the eNodeB and CMTS. For instance, the CMTS doesn't know how much data the eNodeB truly needs to allocate the UE(s) so it must trust the eNodeB - to a certain extent. A rogue eNodeB could intentionally request more bandwidth than it needed if it wanted to try to reduce its latency as much as possible at the expense of wasted DOCSIS grants and potentially other customer's QoS. A poorly designed BWR prediction algorithm could also put an increased burden on the DOCSIS upstream. Whether the BWR inaccuracy is intentional or un-intentional a smart CMTS BWR processing implementation will have checks to ensure the eNodeB is not abusing the BWR.

A simple method to ensure the BWR is not abused is to check if the upstream data flow matches the grants requested. If the eNodeB is using all the grants it asks for there is no reason for concern. On the other hand, a rogue eNodeB that is requesting twice as much bandwidth than it needs would only be using a fraction of the grants it asks for. The CMTS could flag the eNodeB by measuring the eNodeB's grant utilization and enforcing restrictions when poor utilization or other bad behavior is observed. Once an eNodeB is flagged, the CMTS could limit the maximum traffic or scale back the size of the grants dynamically based on the grant utilization of the eNodeB in question.

This puts some responsibility on the shoulders of the eNodeB's BWR prediction algorithm since it must predict the upstream egress data flow with reasonably good accuracy or else it will be bandwidth limited by the CMTS.

7. BWR QoS Considerations

This section will introduce some additional complexities related to the combined data path of the LTE and DOCSIS systems and then discuss how to holistically address these problems.

7.1. A Day in the Life of a Packet

In routing systems, there is a common development exercise called "a day in the life of a packet". The goal is to trace everywhere a packet goes in a system for a given flow. This will help find any data path irregularities that could impact the performance of that data path. Here, we will look at the flow of packets from a source in the UE to a destination just past the CMTS. Figure 8 will be used during this exercise.

The system diagram in Figure 8 depicts an LTE-DOCSIS network. The LTE radio access network (RAN) is connected to the evolved packet core (EPC) using DOCSIS. The LTE system is composed of the UE and the eNodeB which are each composed of their respective sub-systems. The DOCSIS system is composed of the CM and CMTS. For simplicity, this discussion will not go into the detail of the DOCSIS subsystems within the CM and CMTS.

The next thing to notice are the interfaces involved. They are:

- Air
- Ethernet (twisted pair or optical)
- HFC (coax and fiber)

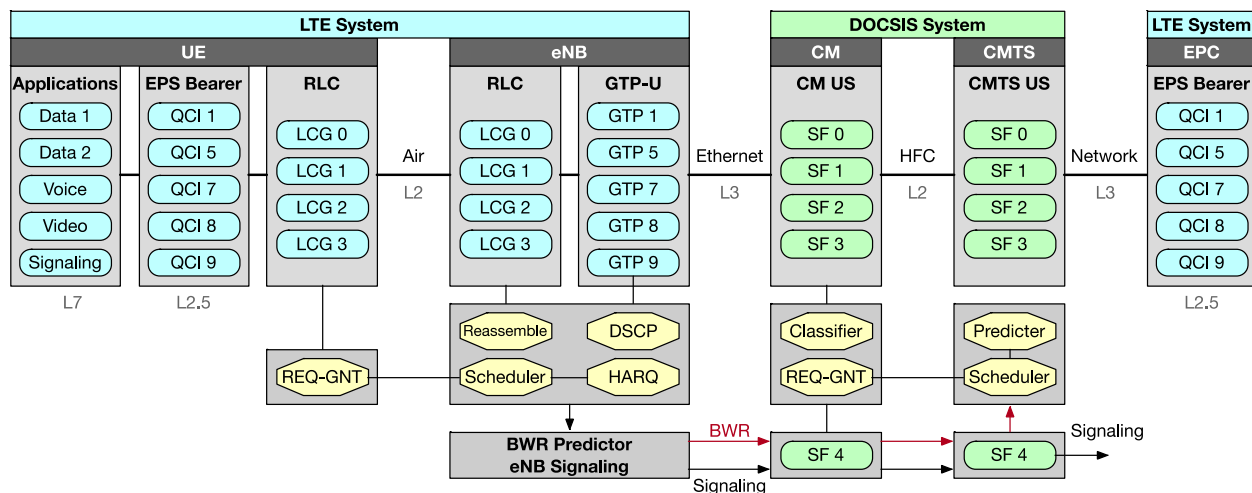


Figure 8 – A Comprehensive Look at the DOCSIS-LTE Interface

The air interface is internal to the LTE system and is a layer 2 interface. The HFC interface is internal to the DOCSIS system and is a layer 2 interface. The Ethernet interface, when combined with IP, is a layer 3 interface. More importantly, the classifiers that the CM uses to classify ingress traffic to a service flow are layer 3 classifiers.

That means that if we have more than one service flow, and hence more than one grant flow, the BWR must describe the request flow *per service flow*. This is a very important point. Service flows are defined by classifiers in the DOCSIS systems. The semantics of the BWR message has to match the semantics of the DOCSIS classifiers.

All this time we have referred to using the scheduling information from the eNB scheduler. That scheduler actually describes flow of segments across a LTE layer 2 construct called logical channel group (LCG). How does this map to something that a DOCSIS classifier can use? Before answering this, let's review the life of a packet in the LTE and DOCSIS data path.

Starting on the left side of Figure 8, there are a number of applications that will initiate data flows. There can be any number of applications that span the classic realms of data, voice, video and signaling. Each of these generate TCP or UDP flows. The LTE system will map these applications into evolved packet system (EPS) bearers and will tag each of these flows with a QoS class indicator (QCI). EPS bearers are end-to-end bearers between the UE, the serving gateway (SGW), and the packet data network gateway (PGW) and are a collection of radio bearer (RB), S1 bearer and S5/S8 bearers. The S1 bearer is carried in an S1 GTP tunnel. Likewise, the S5/S8 bearer is carried in an S5/S8 GTP tunnel. A radio bearer transports the packets of an EPS bearer between the UE and an eNB. Radio bearers (RBs) are in a one-to-one association with logical channels. The SGW, PGW and the mobile management entity (MME) form the basis of the evolved packet core (EPC).

The logical channels are collectively mapped into LCGs to facilitate efficient use of bandwidth for UE bandwidth requests on the radio link. This is important as there is only a maximum of four LCGs but there can be up to 11 EPS bearers per UE. UE signaling is included in one of these LCGs, typically LCG0. The LTE request-grant exchange is based on LCGs, LCs or EPS bearer.

The segments are received at the eNB and assembled back into packets. The eNB will treat each QCI flow differently. This can result in different delays for different QCI streams. The eNB then maps each flow into a S1 GTP tunnel based upon the value of QCI. GTP stands for GPRS Tunneling Protocol and GPRS stands for General Packet Radio Service. Each GTP tunnel will be assigned a DiffServ Code Point (DSCP) for IP quality of service. *GTP is what the CM sees, not the LCG.* This DSCP will usually be copied from the inner DSCP to the outer DSCP. However, the network reserves the right to overwrite the DSCP if it needs to.

The DOCSIS CM will classify the incoming packets into service flows. Each service flow will receive a grant flow. *It is the data into each service flow that the BWR must accurately predict in order to get the right grant flow.*

On the control plane, we can see the eNB prediction algorithm that gathers information from the LTE scheduler, HARQ, the reassembly engine, and other status. It knows about not only LCGs, but also about GTP and DSCP assignments. But it cannot predict based on GTP or DSCP, because it does not have visibility into the LTE requests on the air interface on a per logical channel granularity.

It can also be seen that there are signaling flows for LTE signaling and BWR signaling. In this example, all signaling is going to a dedicated service flow.

7.2. BWR Flow Types

In the previous section, we described the data path from the point of view of a packet that starts at the UE and eventually goes past the CMTS. That packet traverses multiple layer 2 and layer 3 paths. The ultimate realization is that the BWR must talk in a language or syntax that the DOCSIS system can understand. As such, BWR is really an API into the DOCSIS system for requesting bandwidth in the future.

This section will describe four fundamental ways of describing the data flow from the eNB. They are:

1. Based on Bulk
2. Based on LCG
3. Based on GTP
4. Based on DSCP

The methods are ordered to provide increasingly finer granularity of QoS for different types of traffic. We will look at how each method works, where it would be used, its strength, and its weaknesses.

7.2.1. BWR Using Bulk

In this method, all bytes being transferred for a given future point in time are summed together. This would include both data path and signaling bytes.

A variant of the Bulk Method would be to separate BWR out into its own flow.

The advantage of this method is its simplicity. It is useful when there is ample bandwidth and QoS is not required. Or, if the CMTS can accept all traffic on one Service Flow and sort out QoS on its own, say from IP DSCPs, then this becomes a simple interface. The Bulk Flow could be a good match for the DOCSIS Waterfall Granting proposal 0.

The disadvantage of this method is that there is no specific Quality of Service information in the BWR message. As such, the CMTS may not be able to provide lower latency for higher QoS traffic such as signaling.

7.2.2. BWR Using LCG Data Flows

In this method, BWR would be based on the four flows used on the eNB air interface. These are LCG 0 through LCG 3. BWR would describe how many bytes are required for each flow in each future time interval (say 1 ms).

The CM has at least two choices on how to receive these flows.

The first choice uses five service flows – one to match each of the LCGs and one for BWR. Signaling that is not already accounted for in LCG0 could join BWR in its queue. The mapping of resource blocks (RBs) (or LCs) to an LCG is done at the RB setup time by the eNB based on the corresponding QoS attributes of the RBs such as QCI. This mapping is configured by the mobile operator via a provisioning system. If the LTE provisioning system is able to share provisioning information with the DOCSIS system, then it should be possible to know which IP flows are assigned to each LCG. The provisioning system would program into the DOCSIS system the necessary classifiers to recreate four service flows that exactly match the four logical channel groups.

The second choice is two service flows, one for all the LCGs and one for BWR. This second method simplifies the configuration and classification required on the system. It can work well if the CM and CMTS can manage QoS within a Service Flow 0.

A variant of this approach is for the BWR to request per LCG. This will allow the CMTS to make granting decisions based on packet priorities. However, the CMTS would combine all the grant bytes into one service flow and let the CM use the DOCSIS Waterfall Granting proposal 0 to properly allocate the bytes to queues. The UE and eNB use a similar approach for requesting and granting – the UE BSR requests are per LCG, the eNB DCI-0 is a bulk grant, and the UE assignment of packets to LCG flows can be different at grant time than the original BSR intent.

The advantage of this approach is that it is simple and transparent. The external system can see what is really going on.

The disadvantage of this method is that it can require more system configuration for classifiers.

7.2.3. BWR Using GTP Data Flows

From an IP network viewpoint, this method is connection-oriented. That means that the network contains state information about specific IP flows. QoS is achieved by assigning bandwidth to these flows.

In this method, the BWR message expresses the number of bytes sent per GTP flow. The CM would be configured with classifiers that would forward GTP flows onto specific service flows. One or more GTP flows could be combined into a single Service flow, but BWR would still report per GTP. Note that the DOCSIS classifier cannot classify on the QCI field of the GTP flow, so the configuration system would have to create a DOCSIS classifier using the IP tuple (Dest IP, Source IP, Dest Port, Source Port). This information could be derived from the LTE Policy and Charging Rules Function (PCRF) policy system.

Note that LTE signaling messages are sent on a separate protocol such as the stream control transmission protocol (SCTP). BWR would describe the signaling bytes separately.

The advantage of this method is that BWR is describing layer 3 traffic flows instead of internal layer 2 flows. These layer three flows match nicely with DOCSIS classifiers.

The disadvantage of this method is that the provisioning system has to supply precise classifiers. If there is any change in configuration in the LTE system, the DOCSIS system would need updated classifiers. This is achievable, but it does require excellent coordination between the LTE and DOCSIS provisioning systems.

The challenge with this method is that the eNB does not have a proper byte accounting per GTP flow. It only has a byte accounting per LCG flow. Thus, this option may not be implementable in a current LTE system. However, it is included in this proposal in case future mobile protocols or other systems external to a DOCSIS system can use this option.

7.2.4. BWR Using DSCP Data Flow

From an IP network point of view, this method is connectionless. That means that the network does not contain state information about specific IP flows. QoS is achieved by assigning a tag called a DSCP to each packet and using that tag to manage network queuing.

In this method, the BWR message lists the number of bytes per DSCP flow. Usually, each GTP tunnel will have its own DSCP, although one or more GTP tunnels could use the same DSCP. BWR and signaling should have their own DSCP value. This is a very IP centric, connectionless approach, and is a common way of managing QoS on the Internet.

The advantage of this method is that it is very simple. A generic set of classifier rule can be provided to the CMTS at configuration time that is independent of LTE connection information. The LTE system can have any connection it wants. The CMTS will focus on honoring the QoS Policy associated with the DSCP.

The disadvantage of this method is that visibility per GTP flow in the BWR message is lost.

Again, the challenge with this method is that the eNB does not have a proper byte accounting per DSCP flow. It only has a byte accounting per LCG flow. Thus, this option may not be implementable in a current LTE system. However, it is included in this proposal in case future mobile protocols or other systems external to a DOCSIS system can use this option.

7.3. BWR Message Format

This section describes the proposed information elements that will compose a BWR message. These information elements are subject to change as BWR becomes standardized. The actual BWR message format is not specified here as that will be determined by a standards committee.

Table 2 – BWR Message Elements

Information Element	Size	Description
BWR Header		
Version number	3 bits	Version number of BWR message
Sequence number	5 bits	Used for detecting dropped messages
Device Identifier	48 bits	Device unique identifier
1588v2 Reference seconds	48 bits	Value corresponding to the TI reference value.
1588v2 Reference nanoseconds	32 bits	
Time Index (TI) Reference	16 bits	Value corresponding the 1588v2 reference value.
Time Index Increment	32 bits	Total nanoseconds between sequential TI values.
BWR Flow Type	3 bits	Bulk, LCG, GTP, DSCP
BWR Table Size	5 bits	The number of row entries.
BWR Table		
Time Index (TI)	16 bits	Time index that the row refers to
<i>The TI is listed once at the beginning of each row entry. The following entries are repeated as a block within a row entry.</i>		
Last Entry	1 bit	This is the last entry in the current row.
Reserved	5 bits	Not defined at this time.
Operation	2 bits	New – This is a new entry for the bytes Subtract – remove bytes in specified TI Add – add back previously bytes to new TI.
BWR Flow ID	8 bits	For Bulk, this value is 0x00 For LCG, this is 0 through 3 For GTP/QCI, this is 0 through 192 For DSCP, this is the six-bit value 0xFF indicates signaling
Bytes Requested	32 bits	These are the bytes per BWR flow per time interval that cross the Ethernet Interface. The bytes are counted based upon a complete Ethernet frame including the CRC.

The BWR message elements are shown in Table 2. There is a BWR header which contains single entries and a body with the BWR Table.

The BWR message starts with a version field. It is reasonable to expect an updated version for 5G. The sequence field is to allow the receiving entity to detect dropped messages. The sequence number also doubles as a transaction ID. There is no expectation of re-ordering of BWR messages. Note that the eNB may not send a BWR message if there are no traffic updates.

The device identifier field is included for convenience. This allows the CMTS to manage and distinguish among multiple eNB clients and to report statistics with an externally known tag such as the Global eNB identifier which is the combination of PLMN (Public Land Mobile Network) + eNB ID within PLMN. The PLMN is 6 digits: 3 integer digits for mobile country code (MCC), 3 integer digits for a network code (MNC). (24 bits). The eNB identifier is defined as equal to the 20 leftmost bits of the Cell Identity Information Element (IE) as standardized by 3GPP of each cell served by the eNB. An alternate device identifier could be a 48 bit MAC address. A value of all zeros is considered as no identifier supplied.

The BWR message will be sent periodically and will describe repetitive units of bandwidth. Typically, the BWR message will be sent every one millisecond and will describe a one millisecond unit of bandwidth that corresponds to an LTE sub-frame. The Time Index field would count LTE subframes and would correspond to a 1588v2 timestamp. This is done to make BWR more readable and to create a transmission window. It can also help keep BWR message length constrained when the BWR Table gets large. The BWR header contains a cross reference to a 1588v2 timestamp which points to the beginning of the Time Index frame.

In use, the TI at the time the message is created might be X while the TI that bytes are being requested for might be X+8.

There is an entry to indicate what type of flows that BWR is describing – Bulk, LCG, GTP or DSCP. This will dictate what the BWR Flow ID represents. There are some variations on this theme that are possible. For example, bulk and LCG modes could be without much prediction while GTP and DSCP modes would have prediction. Or, prediction could be an option on all modes. This would allow an external prediction mechanism which might be interesting in a cloud implementation.

There is an entry that indicates the number of rows on the BWR Table. Note that the rows can be variable length as flows with nothing to report do not have to be included in the message.

The BWR Table consists of a series of rows. Each row has a time index. If there is no redundancy and no error vectors, there is just the one row entry for the future time index. If the system is configured for redundancy of three, then each row is repeated across three messages.

The BWR operation field allows for New, Subtract, and Add. The New operation is for requesting new bytes for the specified time interval. The Subtract operation is to indicate that the bytes for a predicted TI did not happen. An example of this is a HARQ operation where bytes were supposed to show up but did not. The Add operation is like New except that it is referring to bytes that were previously accounted for. If there are both New bytes and Add bytes for the same flow in the same TI, the bytes would be listed separately. The Subtract and Add operations represent an error vector from the eNB to the CMTS. The CMTS decides what to do with the information.

The BWR Flow ID directly uses the native IDs from LCG, GTP(QCI), and DSCP. A reserved value in this field denotes bandwidth requested for signaling. This would apply for LCG and GTP modes. Bulk and DSCP modes can include signaling within their data flows.

To allow for a variable length row where null flows are not included and there may be separate New and Add entries for a given flow, a last entry marker exists. Then, there is the number of bytes requested for a given flow within a given time interval. A 32-bit field for a one millisecond time interval will support overflows of more than 20 Mbps.

If an eNB uses multiple sectors and/or MIMO channels, each with its own LCG domain, the eNB should collect all byte requests from all sectors and combine it into one BWR message.

A typical BWR message would have an Ethernet/IPV4/UDP overhead of 46 bytes, a BWR header of 24 bytes, a BWR table with three rows (for redundancy) and four flows (78 bytes), for a total of 138 bytes. IPV6 would bring this to 158 bytes.

7.4. Implementation Impacts

If BWR is generated every millisecond, then the eNB has to generate 1000 messages a second that it did not have to do before. This may require additional CPU resources.

To compare actual versus predicted, the eNB would have to monitor its own Ethernet output and timestamp packets in real time. It then could compare off-line the time alignment of the actual packet to the predicted values. This may have some HW impact.

The BWR packet needs an IP destination address (DA) that will cause BWR to be forwarded to the CMTS US Scheduler process. The IP DA could be a well-known IP address, one per CMTS, or one per CMTS scheduler instance. The eNB will need to be configured with this address. This will require some system configuration.

The CMTS has to create a routed path for BWR from an external entity to the scheduler instance.

The CMTS scheduler process will have to receive 1000 BWR message a second *per eNB* that is connected. If there are two eNB per DOCSIS service group (SG) and 32 SG per line card, that would be 64,000 extra messages per second that it is not receiving now. If the CM request flow is disabled for the queues under use, that can help.

Each BWR flow could be on the order of 1.25 Mbps of traffic. If there are a lot of eNBs in a CMTS domain, this bandwidth could be noticeable. There is also the impact of the additional traffic to the control plane in the CMTS.

The CMTS may want to check the percentage of unused grants for a eNB. Unused grants could be a result of improper prediction or over requesting. Results of a grant audit can be combined with the error vector from the eNB to generate statistics.

7.5. Testing out the System

The authors and their respective teams have both simulated and built a system to test out these fundamental theories. Both the MatLab and the physical test environment contain a UE and eNB connected to a CM and CMTS. The test results showed that the desired decrease in latency due to having a BWR interface between the two systems worked.

These test results are being published in a series of IEEE papers 000. The test results showed a decrease in latency from 5 ms minimum to the 2 ms of engineering latency. Also, under heavy load, the overall traffic latency was reduced. Part of that latency reduction under network load can be attributed to BWR not being sent in a contention slot. Figure 9 shows a cumulative distribution function (CDF) of packet delays from the CM to CMTS with and without BWR.




Figure 9 – CDF of CM-CMTS Packet Delays with/without BWR

Extending BWR to Future Platforms

8. Split Small Cell and Cloud CMTS

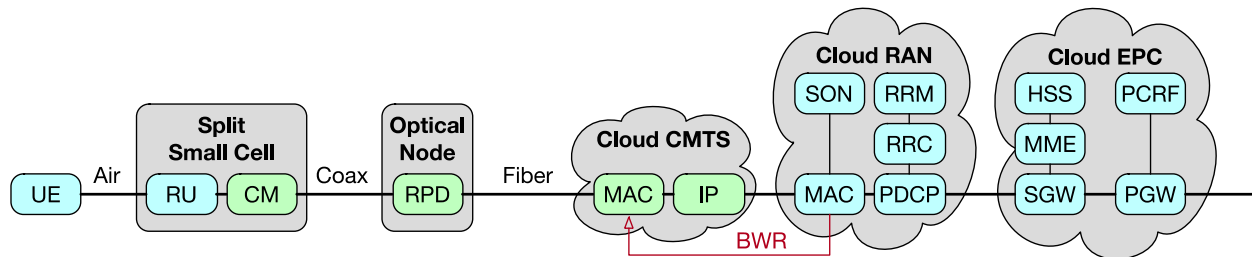


Figure 10 – BWR with vRAN and Cloud CMTS

In the previous sections, the architecture consisted of a physical integrated CMTS and a physical integrated small cell. There are variations of this architecture for which BWR will also work. The CMTS could be an integrated CMTS, a physical CMTS-Core with Remote PHY, or a Cloud CMTS with Remote PHY. The Remote PHY architecture uses a Remote PHY Device (RPD).

The eNB could be integrated or one of eight proposed split configurations 0 0. The split eNB is often referred to as a virtual radio access network (vRAN) or more recently Cloud RAN to reflect updated software architectures. The diagram in Figure 10 illustrates one specific combined scenario known as Network Functional Application Platform Interface (nFAPI) 0, with a Cloud CMTS with a Cloud RAN where the eNB scheduler is located in the cloud. nFAPI is a MAC-PHY split similar in concept to DOCSIS Remote PHY. It can be seen that the BWR message is now passed between software processes within the cloud. The latency for getting BWR from the eNB to the CMTS is greatly reduced now that BWR does not have to traverse up the CMTS access. As a result, the BWR message could be sent less often.

BWR will also not need to include any bandwidth estimation for signaling that would originate and terminate in the cloud. Further, the X2 traffic from eNB to eNB might also be exclusively in the cloud and thus also achieve low latency.

The migration to vRAN does modify the scope of the DOCSIS data path requirements from backhaul to fronthaul. In a backhaul scenario with an integrated eNB, the data packets are fully formed packets and are a tunneled version of the original packets from the original application. In a fronthaul scenario, the content is an interim form that originates from some midpoint in the eNB processing chain.

The ideal fronthaul technology would be a transport that has QoS attributes attached to the data flow. The reason for this is that all transport networks are multi-hop (meaning one or more routers) and bandwidth is typically queued and aggregated. A device like an eNB will have a traffic profile that will contain high priority and lower priority traffic. The task for the network is to maintain that traffic profile through the network.

With a split small cell, where the signaling messages between the UE and the eNB now travel to the cloud, QoS will be required to provide a low latency path for UE signaling. QoS works best with high priority, low bandwidth flows.

From a network QoS viewpoint, a fronthaul like the Common Public Radio Interface (CPRI/eCPRI) 0 where all traffic is equal would not be an optimum choice, whereas a front haul like the network functional application platform interface (nFAPI) 0 would be a good choice.

Conclusion

There were at least three fundamental requirements for DOCSIS to provide mobile backhaul or front haul services: timing, bandwidth and latency. Timing and bandwidth are problems addressed in other white papers. In this white paper, it was described that by coupling the scheduling mechanism from the LTE eNB uplink scheduler to the DOCSIS CMTS upstream scheduler, in effect creating a pipeline that allows for earlier requesting on the DOCSIS system, low latency on the DOCSIS system can be achieved.

The bandwidth report (BWR) message is a proposed API on the DOCSIS system that will allow external systems such as small cells to request bandwidth in advance of when it is needed, thus achieving a low latency transport.

The eNB receives traffic from UEs. In addition, it has its own signaling. The collection of all those bytes and the significance of each byte stream is known as a traffic profile. The eNB needs to place its traffic onto a network. In this case, the network is a DOCSIS system, but even if it was an Ethernet backhaul, that network needs to understand the profile of the traffic from the eNB.

The most popular form of profiling traffic on an IP network today is with IP Differentiated Services (DiffServ). Hence, it is the authors recommendation that the eNB should always properly mark each outgoing IP packet with the appropriate DSCP. The CM classifiers should be setup with DSCP, and thus BWR would be best used by requesting per DSCP. DiffServ techniques also tend to be stateless which means there is less network configuration to worry about. DiffServ is also queuing orientated, and all network ports are queue orientated.

In conclusion, by having LTE and DOCSIS work together, DOCSIS is a viable backhaul mechanism for gigabit-per-second mobile traffic.

Acknowledgements

Many engineers provided software implementation support and technical insights throughout this project. The authors would like to thank Michel Chauvin, Joey Padden, Aaron Quinto, Balkan Kecicioglu, and Vaibhav Singh of CableLabs, and Elias Chavarria Reyes, Zheng Lu, Dantong Liu, Alon Bernstein, and Oliver Bull of Cisco for their contributions.

Abbreviations

ACK	Acknowledgement
AM	Acknowledged Mode
ARQ	Automatic Repeat Request
BE	Best Effort
BSR	Buffer Status Report
BWR	Bandwidth Report
CM	Cable Modem
CMTS	Cable Modem Termination System
CPRI	Common Public Radio Interface
DiffServ	Differentiated Services
DCI-0	Downlink Control Information Format 0
DL	Downlink
DOCSIS	Data over Cable System Interface Specification
DSCP	DiffServ Code Point
eNB	Evolved Node B
EPS	Evolved Packet System
EPC	Evolved Packet Core
FDD	Frequency Division Duplex
FEC	Forward Error Correction
GPRS	General Packet Radio Service
GTP	GPRS Tunneling Protocol
HARQ	Hybrid Automatic Repeat Request
HFC	Hybrid Fiber-Coax
HSS	Home Subscriber Server
IP	Internet Protocol
LCG	Logical Channel Group
LTE	Long Term Evolution
MAC	Media Access Control
MCC	Mobile Country Code
MIMO	Multiple Input Multiple Output
MME	Mobility Management Entity
MMM	MAC Management Message
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
NACK	Negative Acknowledgement
nFAPI	Network Functional Application Platform Interface
PCRF	Policy and Charging Rules Function
PDCCH	Physical Downlink Control Channel
PDCP	Packet Data Convergence Protocol
PDN	Packet Data Network
PGW	PDN Gateway
PHY	Physical Layer
PLMN	Public Land Mobile Network
QCI	QoS Class Indicator

QoS	Quality of Service
RAN	Radio Access Network
RB	Resource Blocks
RLC	Radio Link Control
RPD	Remote PHY Device
RRC	Radio Resource Control
RRM	Radio Resource Management
rtPS	Real Time Polling Service
RU	Radio Unit
SID	Service Identifier
SF	Service Flow
SGW	Serving Gateway
SON	Self-Organizing Network
SPS	Semi-Persistent Scheduling
SR	Scheduling Request
TCP	Transport Control Protocol
TDD	Time Division Duplex
TI	Time Index
UDP	User Datagram Protocol
UE	User Equipment
UGS	Unsolicited Grant Service
UL	Uplink
UM	Unacknowledged Mode
VoLTE	Voice over LTE
vRAN	Virtual Radio Access Network

Bibliography & References

Jennifer Andreoli-Fang, John T Chapman, “Synchronization for Mobile Backhaul over DOCSIS,” Proceedings of SCTE Fall Technical Forum, October, 2017

John T. Chapman, et. al., “The DOCSIS Timing Protocol (DTP), Generating precision timing services from a DOCSIS system,” Proceedings of INTX/SCTE Spring Technical Forum, 2011.

“Data-over-Cable Service Interface Scalable service interface specifications, DOCSIS® 3.1, MAC and upper layer protocols interface specification,” I-11, issued 2017-05-10, CableLabs. < MULPI Specification, Figure on req-gnt delay>.

John T. Chapman, G. White, H. Jin, “Impact of CCAP to CM Distance in a Remote PHY Architecture,” Proceedings of the INTX/NCTA Technical Forum, April 2015.

John T. Chapman, Tong Liu, “Waterfall Granting,” Submitted to the CableLabs DOCSIS MAC committee, June 2017.

Jennifer Andreoli-Fang, John T. Chapman, “Latency Reduction for Mobile Backhaul by Pipelining LTE and DOCSIS,” to appear in *Proc. of IEEE Globecom*, December 2017, Singapore.

Jennifer Andreoli-Fang, John T. Chapman, “Mobile Aware Scheduling for Low Latency Backhaul over DOCSIS,” to appear in *Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, October 2017, Montréal, Québec.

Jennifer Andreoli-Fang, John T. Chapman, et. al “Real-Time Experimentation with Mobile-Aware Scheduling for Low Latency Backhaul over DOCSIS”, submitted to IEEE International Conference on Computer Communications (Infocom) 2018.

R3-161687, Draft TR 38.801 (v030) Study on New Radio Access Technology: Radio Access Architecture and Interfaces, NTT DOCOMO, INC (Rapporteur), 3GPP TSG RAN3, August 2016.

Small Cell Forum Document 159.07.02, Small cell virtualization functional splits and use cases, Small Cell Forum, January 2016.

“Common Public Radio Interface (CPRI); Interface Specification”, Release V7.0, www.CPRI.info, Oct 9, 2015.

“FAPI and nFAPI Specifications”, Small Cell Forum, Release 9.0, Document 082.09.05 , May 2017.

DWDM Access for Remote PHY Networks Integrated Optical Communications Module (OCML)

A Technical Paper prepared for SCTE•ISBE by

Harj Ghuman

Network Architecture and Technology Strategy
Cox Communications
6305 Peachtree Dunwoody Road
Atlanta, GA, 30328
404-449-4711
Harj.ghuman@cox.com

1. Introduction

MSO networks are evolving from traditional Hybrid Fiber Coax (HFC) to remote PHY (RPD) architectures. A typical 500 home node serving area may transition to one with up to 20 RPDs, each requiring a 10G input to provide sufficient bandwidth. An optical access trunk from the headend to the fiber node area will now have to support a high capacity 200G bi-directional link to support 20 RPDs with 10G. Provisions may also have to be made for a PON/10GPON overlay to feed PON networks as well as high capacity 10G/100G for business services from the same fiber trunk.

There are two optical trunk methodologies currently being investigated:

1. A bidirectional 10G DWDM optical trunk
2. A high capacity coherent optical link

A 10G DWDM multi-wavelength optical trunk to each original HFC node area is an efficient low cost solution that makes use of mature, readily available technology. High capacity coherent optical links may eventually be able to fulfill the requirement, but are not yet commercially available in a field hardened, cost effective package. It is likely that networks may have to be designed to allow for the coexistence of both 10G Non-Return to Zero (NRZ) and Coherent networks.

To that end, Cox has developed the OCML (Optical Communications Module Link Extender) concept for cost effectively transporting a mix of DWDM 10GbE, GPON and 10GEPON wavelengths over the same fiber to a typical HFC node serving area.

2. Optimum Optical Access

An optimum access solution for the last mile should be a scalable and technology agnostic solution which is cost, capacity and bandwidth efficient. Today, this calls for a solution which uses mature low cost DWDM 10G NRZ optics to feed multiple RPDs and at the same time allows for a migration to 25G NRZ and 100G coherent for higher capacity requirements. Figure 1 shows a network topology where the backbone and metro core is coherent while the optical access supports both 100G coherent and DWDM 10G NRZ. DWDM 25GHz NRZ for access (40Km) is currently being investigated by many manufacturers and will become available if there is a market need for it, since 25G non-DWDM is already available. However, it should be noted that it is unlikely that RPDs will need more than 10G in the next decade. So, it is important to choose an optical access solution that is technology agnostic and supports a variety of different technologies including 10G, 25G and coherent optics. This will make the last mile truly scalable and serve the needs of the industry for a long time. While it is unlikely that individual RPD's will require 100G anytime in the foreseeable future, it may be beneficial to be able to provide 100G coherent links for high capacity requirements such as enterprises and Multiple Dwelling Units (MDUs).

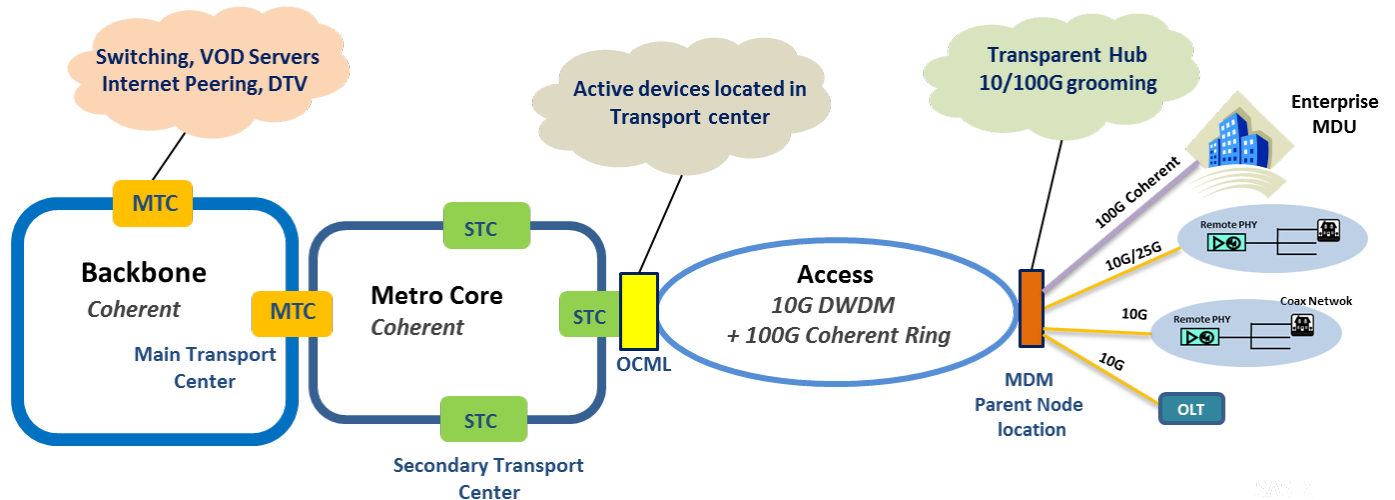


Figure 1 - Optical network: Backbone, Metro Core and Access

3. Access Network Fiber Infrastructure Evolution

The fiber infrastructure between the headend and the fiber node is the most constrained part of the network and the most expensive to upgrade. There is a need to evolve the Access Network fiber infrastructure between the headend and the fiber node to support digital optics thus transitioning away from the current analog (AM) optical connections. This transition is underway to accommodate the impending architectural shift to Remote PHY.

3.1. Optical Access Design Goal

The most important design goal in the transition from analog to digital optics is to preserve the existing fiber network. As all initial Remote PHY Devices will natively support 10G, the use of 10G DWDM optics allows for a direct connection between the RPD and the corresponding aggregation device. This eliminates the need for additional active devices in the field. The placement of active devices in the field are problematic, as these would increase the operational costs to maintain the network and reduce robustness.

3.2. Bandwidth Projections

Cable operators have seen exponential growth in broadband traffic in recent years, with both downstream and upstream traffic growing significantly due to the consumption of video which accounts for a large and growing percentage of bandwidth used in hybrid fiber-coaxial (HFC) networks.

In the access migration to support digital optics, it is projected that each fiber between the headend and the parent node (formerly HFC fiber node) may need to accommodate up to 20 RPDs. Sufficient wavelength capacity remains available with current filter structures to support 80 or more wavelengths. Each RPD can support approximately 64 HHP with a single service group initially shared among up to 4 RPDs. Given current projections it is believed that 10G interconnections between the RPD and aggregation point provides ample capacity for growth. As a 25G DWDM NRZ becomes available, the RPDs could be upgraded to 25G where required, and only the network ends would need to be changed.

3.3. Passive Field Infrastructure

A design goal in the transition to digital optics in the access is to preserve (as much as feasible) a passive outside plant / field infrastructure. As all initial Remote PHY Devices will natively support 10G the use of 10G DWDM allows for a direct connection between the RPD and corresponding hub located aggregation device. Alternative solutions have been suggested that require active aggregation devices in the field but have the disadvantage of increased operational costs to maintain the network and reduced robustness. The OCML concept places all active components in the MTC/STC facilities, with a completely passive Mux/DeMux (MDM) device in the field.

3.4. Time to Market

As the industry begins to implement initial Remote PHY rollouts next year, the OCML provides for the use of available technology, insuring a transport solution is available in time to support the access solution. Finally, although Cox has decided to utilize 10G DWDM NRZ as the initial technology in the transition to digital optics in the access, the underlying infrastructure has been based on ITU standard compliant wavelength plans in anticipation of the requirement to deploy 100G (and beyond) wavelengths.

3.5. 10G DWDM

Figure 2 shows a multi-wavelength 10G DWDM network. Some of the main points are: 10G NRZ is a mature technology and all associated components such as DWDM passives and hardened, pluggable SFP optics are readily available. Some of the main features of an all 10G DWDM access solution are:

- 10G NRZ is low cost, mature and readily available
- Scalable to 25G and 100G coherent
- Large capacity, single fiber pair accommodates 400G (40x10G) – 800G (80X10G)
- Transparent All PASSIVE hub
- Pay-As-You-Grow. Parent node can grow from 2 RPDs to 20 RPDs




Figure 2 - 10G Multiwavelength DWDM Ring

4. Cox Network Transformation

A typical Cox HFC network is shown in Figure 3 where the primary node is connected via dual (redundant) primary and secondary fiber links which can be between 5 to 60 km, with the primary usually the shorter ring. As we migrate to a distributed RPD architecture, these fibers will migrate to a multi-wavelength 10G ring to feed multiple RPDs. Figure 4 shows how a typical HFC node 500 home serving area can be converted to a 20 X 10 DWDM 10G Optical trunk where each RPD can be fed by a 10G. This allows the outside plant to remain passive using a simple 40 channel thin film DWDM.

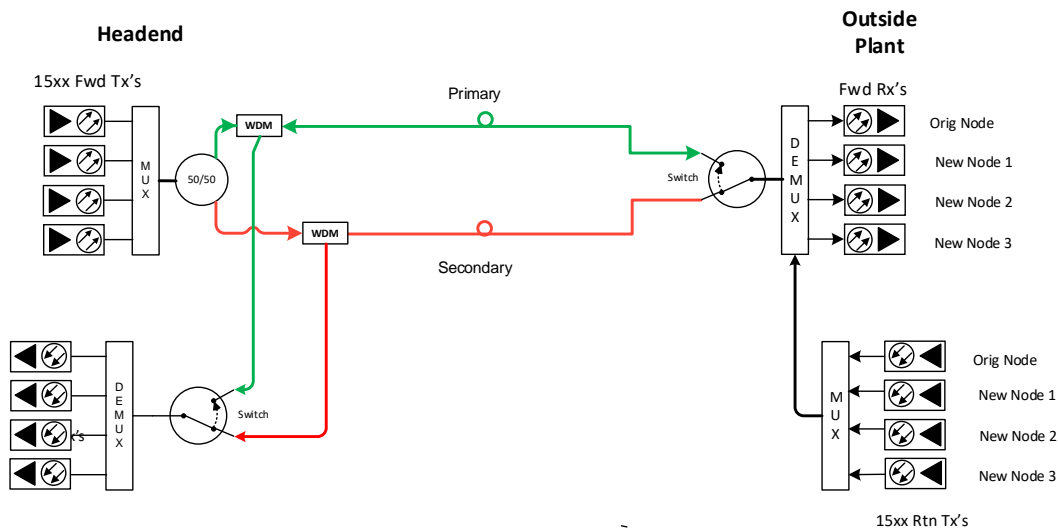


Figure 3 - Current Cox HFC Analog Optical Network




Figure 4 - HFC Node converted to 10G RPD

5. Optical Communication Link Extender (OCML)

Figure 5 below shows a simple 10G DWDM with a GPON/10GEPON overlay, while Table 1 shows the associated loss budget. This shows that the fiber distance associated with such a network is less than 10Km. At Cox, our network supports dual rings of 5 to 60Km and would require field amplifiers as shown in Figure 6.

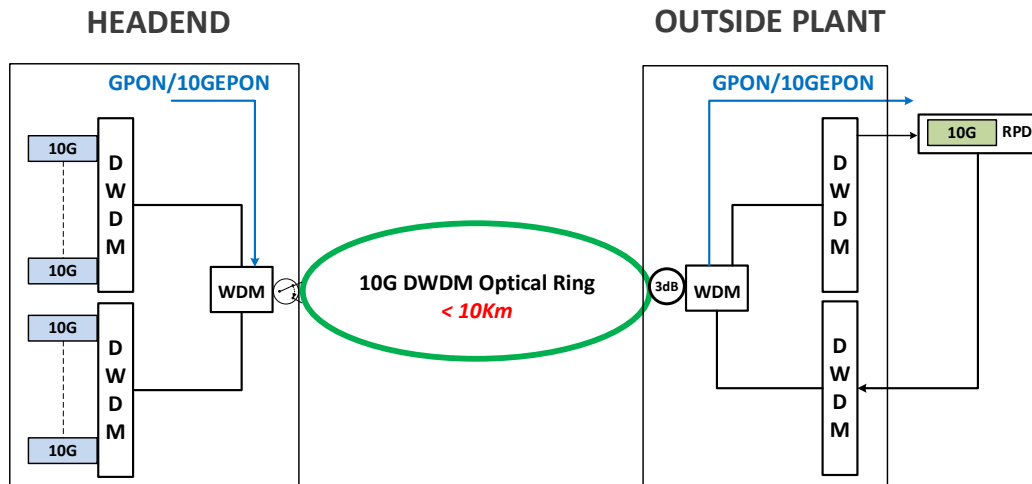


Figure 5 - Multi-wavelength 10G optical trunk without EDFAs

Table 1 - Link loss for 5 km RPD Ring

PARAMETER	LOSS BUDGET		
Headend			
10GbE Txcvr Pwr/WL		0.0	dBm
Fiber Length Headend to MDM	5	1.1	dB
Fiber Length MDM to RPD	2	0.44	dB
Headend DWDM Mux		4	dB
GPON/XGPON WDM		2	dB
Switch		1.5	dB
Outside Plant			
50% Optical Splitter		3.5	dB
WDM		2	dB
DWDM		4	dB
Connectors (2 @0.3dB)		0.6	dB
Safety Margin		3	dB
	Total loss	22.14	dB
	FieldRx I/P	-22.1	dBm

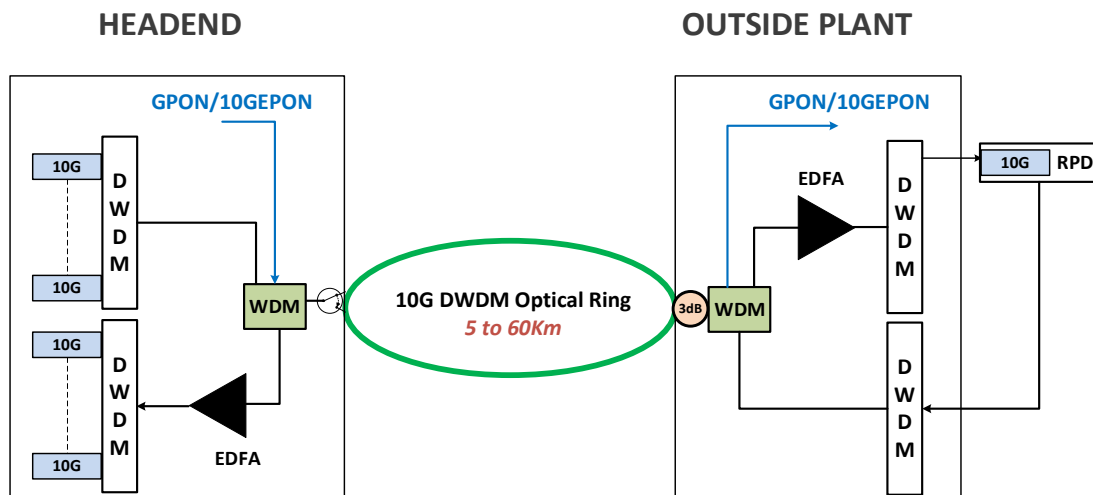


Figure 6 - Multi-wavelength 10G optical trunk with field EDFAs

5.1. The Optical Communications Module Link Extender (OCML)

We have created the Optical Communications Module Link Extender (OCML) concept which allows all the active components of a 10G DWDM access network to be placed in the headend while keeping the outside plant all passive with a Mux/Demux (MDM). The OCML and the remote outside plant MDM unit were required to support next-generation fiber deep DWDM access networks.

Some of the key requirements in designing the OCML were to create an integrated module which could support:

- 5 to 60 Km dual and variable fiber rings -
- 10G DWDM - transport up to 20 X 10G bi-directional wavelengths
- GPON/10GEPON transport over 20Km

The OCML was designed with a view that it could support future requirements for:

- 100GEPON
- 100G Coherent

An added benefit of keeping all the subsystems of the OCML integrated in a box is the estimated cost reduction over individual components.

The Optical Communications Module – Link Extender (OCML) enables DWDM aggregation for optical trunks up to 60 km without any active field devices such as optical amplifiers or optical switches, reducing equipment costs and operating expenses while greatly simplifying installation. The OCML schematics are given in Figure 7 below.

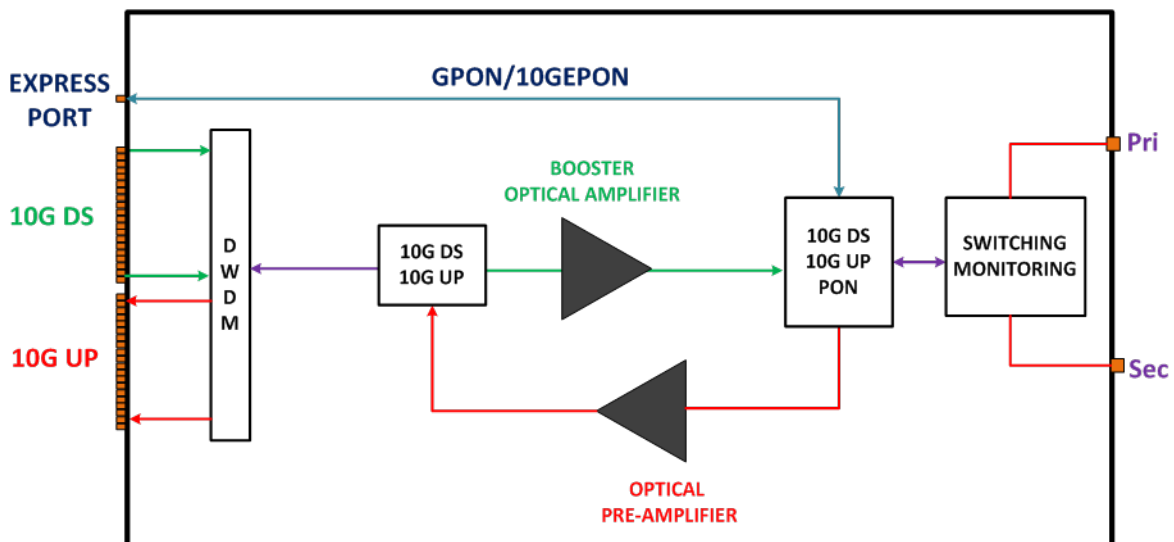



Figure 7 - OCML Block Diagram

The OCML wavelengths to be transported are depicted in Figure 8 below. Future requirements may also include 100GEPON and a 100G/200G/400G coherent to be transported over the same fiber network via the OCML.



5.2. Outside Plant Mux/Demux (MDM) unit

To keep with an all passive outside plant, a Mux/DeMux device is used in the field as shown in Figure 9 below. The MDM is essentially an integrated passive module which reduces cost and keeps with the philosophy of keeping the outside plant simple and passive as much as is possible.

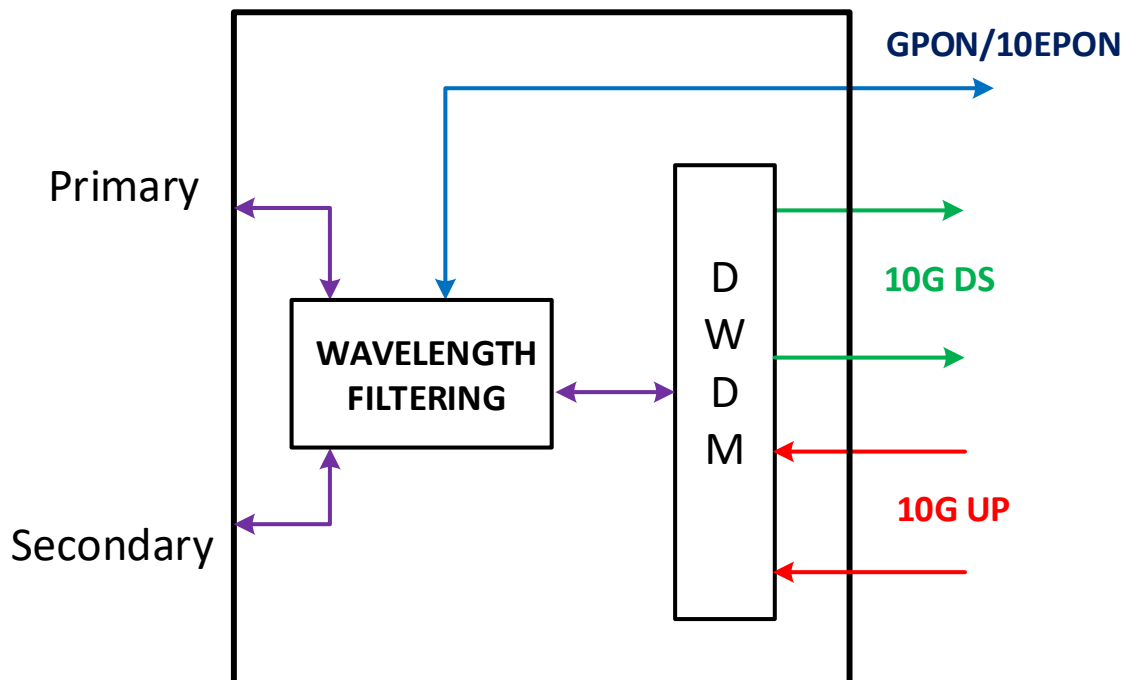


Figure 9 - Outside Plant MuxDemux (MDM)

Figure 10 below shows the numerous applications which can conceivably be fed via a multi-wavelength DWDM ring which also supports other wavelengths outside the C band such as GPON and 10GEPON. Such an architecture can also support a 100G coherent which can be used for high capacity requirements. This is a truly powerful and scalable network architecture which allows various technologies to be transported. In the future, 10G NRZ could be upgraded to a 25G NRZ which is basically a market driven application. If there is a market requirement, 25G NRZ will become available in cost and quantity.




Figure 10 - OCML and MDM Applications

6. Technical Design Considerations of the OCML

DWDM transmission in fiber is always prone to fiber nonlinearities. The potential debilitating impact of Stimulated Raman Scattering (SRS) was considered in the design of the OCML. SRS is a nonlinear process where higher frequency optical channels are depleted and lower frequency optical channels amplified. Figure 11 below shows an illustration and the effect of spontaneous Raman scattering and its impact. If two wavelengths are propagating in fiber and depending on their wavelength separation, polarization states and optical launch power, SRS crosstalk can occur which depletes shorter (pump) wavelength and amplifies the higher (stokes) wavelength. It is sensitive to the power level of the lower wavelength signal, the separation between the two wavelengths and the fiber length. The effect is on the lower RF frequencies carried at the longer wavelength optical carrier. The primary way it manifests itself in FTTH systems is interference from the 1490nm downstream data signal cross talking and causing interference into the 1550nm 10G signals.




Figure 11 - Stimulated Raman Scattering (SRS) Illustration

SRS induced crosstalk is given by the following equation:

$$XT_{SRS,i} = P^2 \sum_{k \neq i} \frac{g^2_{i,k}}{\tilde{A}_{eff}} \left((1 - e^{-\alpha L})^2 + 4e^{-\alpha L} \sin^2 \left(\frac{\Omega d_{i,k} L}{2} \right) \right) / (\alpha^2 + \Omega^2 d_{i,k}^2)$$

This indicates that SRS induced crosstalk is higher for larger channel separation and smaller for higher modulation frequency. One may consider the potential debilitating impact of SRS when coexisting 10GbE with the 1490 nm GPON downstream wavelength since it is only 60nm away for the C band wavelengths of the 10G signals and hence it falls within the Raman Gain profile and can potentially cause crosstalk. This is calculated from the above equation to be about 35 dB and is acceptable.

6.1. Link Loss Budgets

The OCML is a somewhat complex system supporting many variable parameters resulting in careful management of the OCML gains. 10G Link performance is dependent on four main factors:

- Transceiver Rx Power
- Type of receiver technology: PIN or APD
- Dispersion which is dependent on the fiber
- EDFA OSNR: 58dB – NF – Pin
 - NF is the noise figure
 - Pin is the input optical power

In the OCML, the downstream has a high OSNR > 40 due to the relatively high optical input levels of the DS EDFA. On the upstream, DCM are incorporated to reduce fiber dispersion to compensate for the lower OSNR, this creates a robust OCML which can work over a variable 5 to 60Km fiber distance. DCMs will also create negative dispersion over shorter fiber lengths so 10G transceivers should be

specified to work over a wide range of negative to positive dispersion. Figure 12 shows measurements which were done with an APD transceiver, 60Km fiber and DCM (30Km) for different OSNR values. $BER > 10E^{-12}$ can be achieved with a low OSNR of 23dB. In the downstream transceivers need to operate over a lower optical receive power so 80Km APD based receiver diodes are used. In the upstream, both lower cost PIN diode or APD based transceivers can be utilized. All these factors require a careful management of EDFA gains and transceiver input power levels.




Figure 12 - BER v/s OSNR, 60Km fiber, 30Km DCM

Figure 13 below shows the dispersion penalty which is considered by adding a dispersion penalty in the loss budget and/or a dispersion compensating module (DCM) which results in a better upstream link performance since the OSNR could be as low as 23 or 24 dB over longer links. Adding a DCM in the OCML plus a 3-dB margin (for fiber repair, aging of lasers, etc.) makes for a very robust OCML system.




Figure 13 - Dispersion Penalty in Optical Fibers

7. OCML Proof of Concept (POC) Test Results

We have developed a Proof of Concept (POC) for the OCML and while early in the test cycle, it has shown very good optical performance. Figure 14 shows the overall block diagram of the OCML, Figures 15 show the DS OSNR (>40) and figure 16 the UP OSNR > -26 over 60Km of fiber. Figure 17 shows downstream crosstalk as > 35dB. Figure 18 shows upstream crosstalk as >40 dB

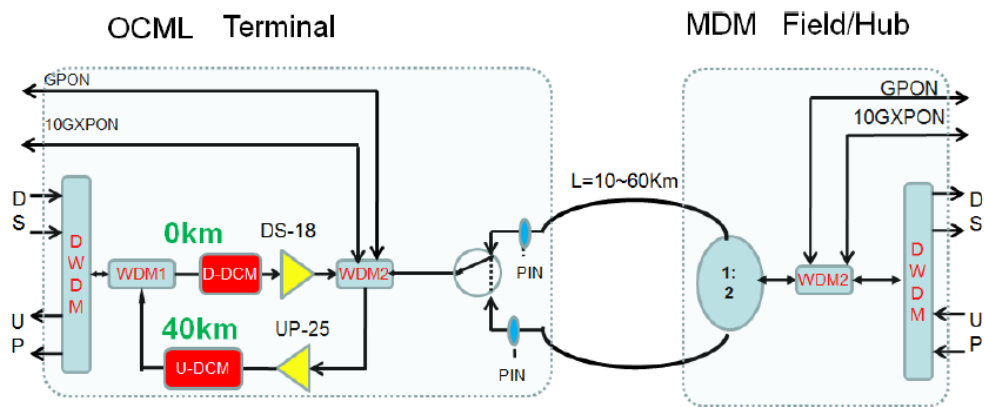


Figure 14 - Block Diagram of Network Set up

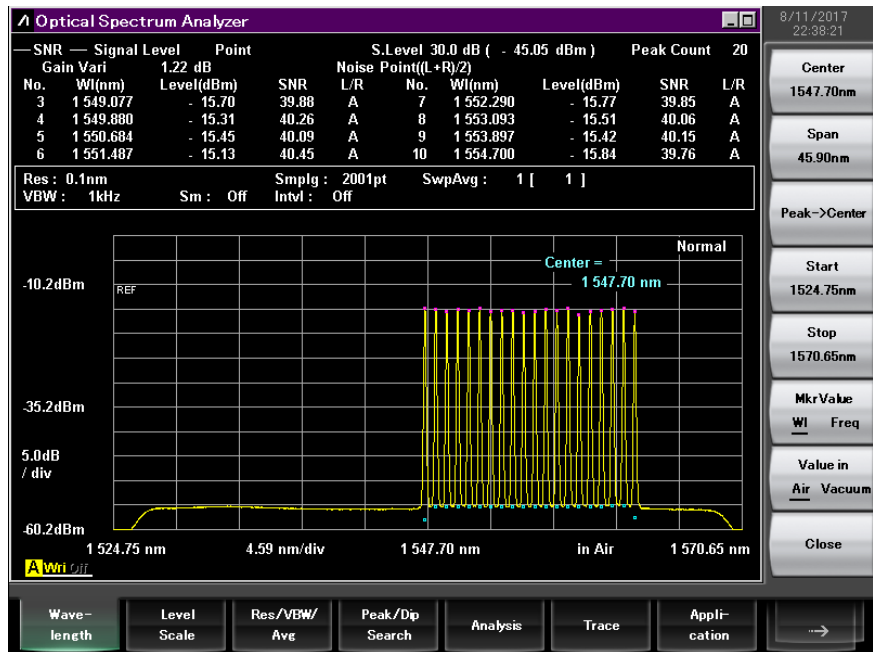


Figure 15 - OCML Downstream OSNR

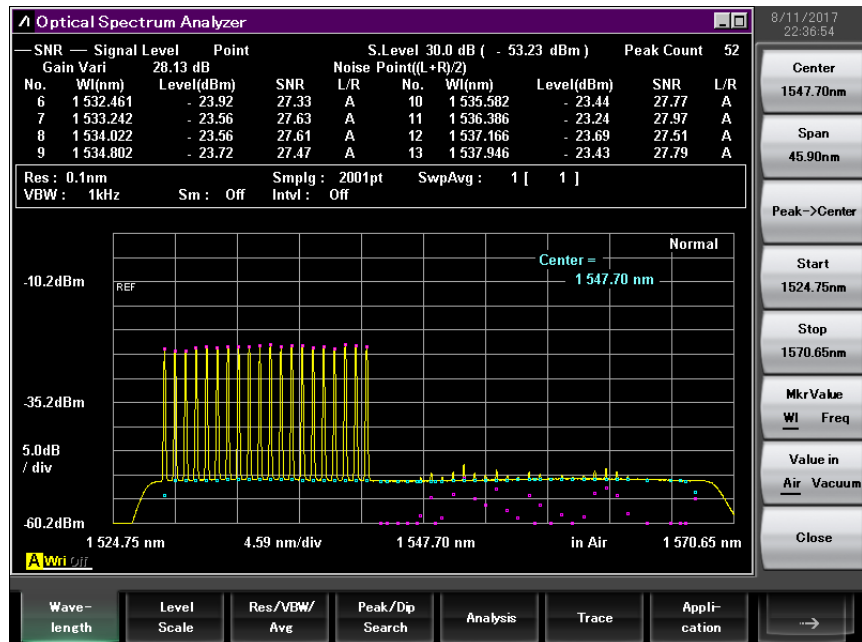


Figure 16 - OCML Upstream OSNR

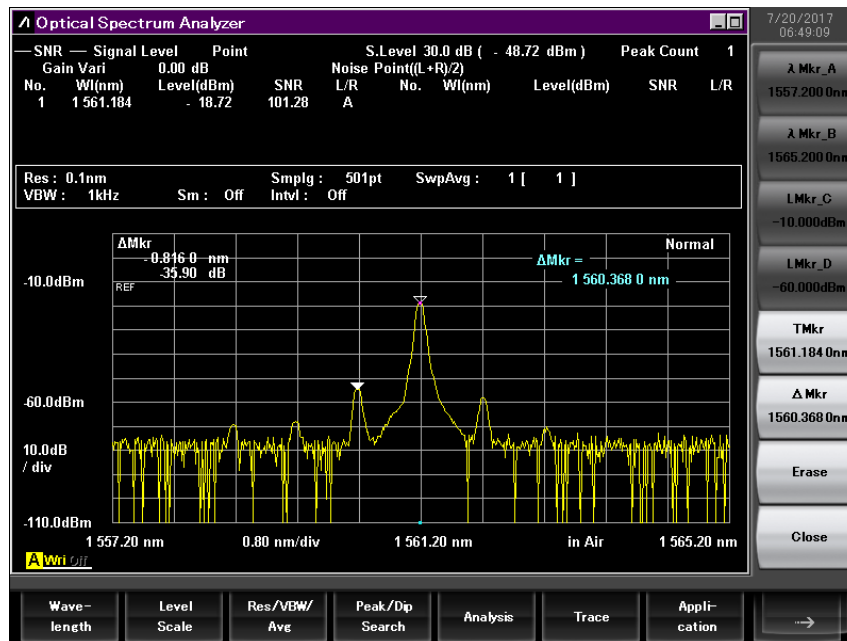


Figure 17 - OCML Downstream Crosstalk

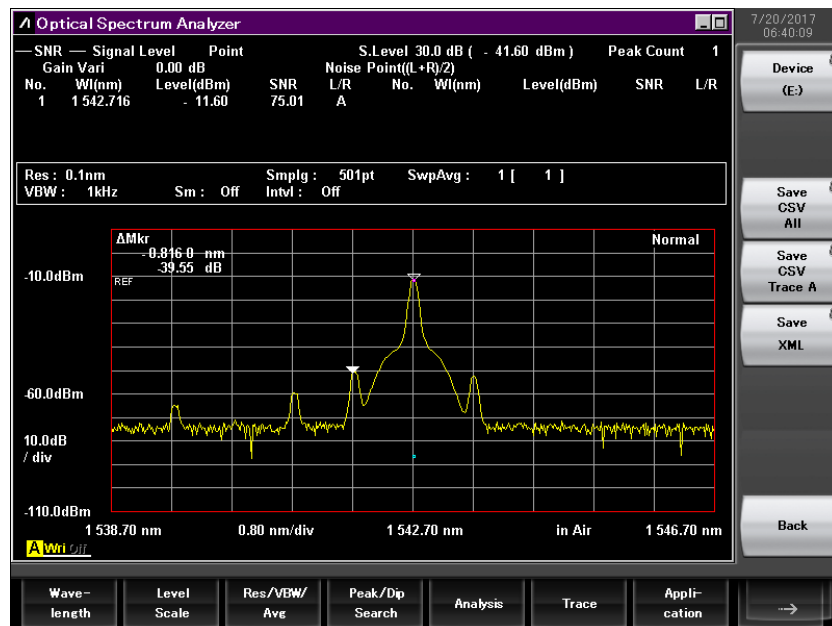


Figure 18 - OCML Upstream Crosstalk

Early testing shows very good results with error free performance over 60Km observed at a low 20dB OSNR and optical receive power of -17 dBm (with PIN 10G transceivers) and -23 dBm

with APD 10G transceivers. These were measured with a 40Km DCM which provided at least 2 dB noise penalty margins.

8. Conclusion

Optical feeds to R-PHY networks can be 10G NRZ DWDM or potentially 100G Coherent. This paper shows that 10G DWDM provides enough capacity for R-PHY networks for the foreseeable future. It is a low cost, mature and scalable technology, plus the network can be all passive from the optical headend to the R-PHY node. A 10G DWDM solution provides 400G of duplex bandwidth over a single fiber (46 wavelengths down and 46 wavelengths up) and is a “Pay-As-You-Grow” network, allowing 10G increments to be added at minimal cost.

The Optical Communications Module Link Extender (OCML) is an innovative low cost solution developed by Cox and will primarily be used for 10G feeds to R-PHY. The OCML is also designed to transport an overlay of GPON and 10GEPON. Future requirements may call for 100GEPON or a coherent optical signal to be transported with the 10G traffic. The combination of the OCML and an all passive outside plant module (MDM) are important building blocks in the evolution from traditional HFC to a Remote PHY architecture. The OCML allows both 100G coherent and 10G DWDM to be transported over the same fiber, making it a very powerful and scalable solution. Optical access networks should be able to coexist with 10G DWDM and coherent technologies to maximize the benefits of both. At the end of the day, it is about building a network that is cost, capacity and bandwidth efficient.

Abbreviations

OCML	Optical Communications Module Link Extender
MDM	Mux DeMux
MTC	Master Terminal Center
STC	Secondary Terminal Center
10G	10Gbps
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

Bibliography & References

1. 40Gbit/s & 100Gbit/s Implementation tradeoffs. Paul Morkel et al. Adva, Internet2 Meeting, April 2009
2. FTTx Networks. Technology Implementation & Operation, Weyl Wang et al, 2017
3. Nonlinear Optical Crosstalk in WDM CATV systems, Dogan Atlas, Antec Corporation
4. Optical Fiber Telecommunications. Ivan P. Maminow. Tingye Li, Alan E. Willner
5. CNR determination at 1550nm. Harj Ghuman, Lightwave 2009

Digital Coherent Transmission for Next-Generation Cable Operators' Optical Access Networks

A Technical Paper prepared for SCTE•ISBE by

Zhensheng (Steve) Jia, Ph.D.

Principal Architect
CableLabs
858 Coal Creek Circle
Louisville, Colorado 80027
303-661-3364
s.jia@cablelabs.com

L. Alberto Campos, Ph.D.

Distinguished Technologist
CableLabs
858 Coal Creek Circle
Louisville, Colorado 80027
303-661-3377
a.campos@cablelabs.com

Chris Stengrim, CableLabs

Jing Wang, Ph.D., CableLabs

Curtis Knittle, Ph.D., CableLabs

Introduction

Cable operators' residential offerings of Gigabit per second service are occurring on a regular basis now, and access bandwidth requirements are expected to grow to multi-Gigabit per second speeds driven by increasing 4K/8K video streaming, proliferation of cloud computing, big data, social media, Internet of Things, and mobile data delivery. Existing Hybrid Fiber Coax (HFC) networks have typically been designed with 6 to 8 fibers connecting the hub to the fiber node; however, many of these fibers have been repurposed for business services, node splits and backhaul services. In many instances, only the two primary fibers remain available for access network transport. This fiber shortage will only intensify as fiber demand for business and wireless backhaul increases and fiber deep architectures become prevalent. Efficient use of optical fiber infrastructure and adoption of innovative technology becomes critical in the evolution towards next-generation cable access networks.

The current analog or direct detection optical schemes face huge challenges because of their low receiver sensitivity and limited options for long-term upgrading, especially in the legacy fiber environment, where operators continue to take advantage of the existing infrastructure to avoid costly fiber re-trenching. Coherent technologies have been recently considered as the most effective future-proof approach for both brown and green field optical access deployments. Thanks to the advancements in digital signal processing (DSP), digital coherent detection enables superior receiver sensitivity that allows an extended power budget and high frequency selectivity enabling dense wave division multiplexing (DWDM) without the need of narrow-band optical filters. Moreover, the multi-dimensional recovered optical signal provides additional benefits to compensate linear transmission impairments such as chromatic dispersion (CD) and polarization mode dispersion (PMD). In the cable access environment, coherent optics allows operators to best leverage the existing fiber infrastructure to withstand the exponential growth in capacity and services. However, there are several engineering challenges of introducing digital coherent technologies into optical access networks. To reduce the power consumption and thereby meet the size and cost requirements for access applications, development of both low-complexity application-specific integrated circuits (ASICs) and optics is essential. In particular, co-design of a DSP ASIC and optics to trade performance against complexity, cost and power consumption is imperative.

In this paper, use cases are explored for near-term and long-term applications, including the deployment for aggregation points in distributed HFC architecture (Remote PHY or Remote MAC and PHY, abbreviated R-PHY/ R-MAC-PHY), remote Passive Optical Network (PON) systems, and eventually coherent optics to the premises. The corresponding economic model for near-term aggregation transmission system will be presented for the comparison with WDM direct detection system. This paper provides an in-depth analysis describing a typical digital coherent optical system, including basic elements of multi-dimensional modulation scheme and a digital coherent receiver structure with fundamental DSP building blocks for both optical transmitter and receiver. The current evolution of coherent optical modules is also introduced.

This paper highlights the motivation for coherent optics in access and potential approaches to re-design and re-engineer the digital coherent concept from long-haul and metro solutions to the access network, leveraging reduction in complexity and cost as well as the benefits of capacity increases and operational improvements. Proof-of-concept experimental results demonstrating multi-wavelength multi-terabit per second within an access environment and the evaluation of coexistence between legacy analog and coherent system are also shown here.

Content

1. Increasing Demand for Higher Bandwidth

Video intensive technologies require the most bandwidth, and currently out of all the video-related applications, Virtual Reality/Augmented Reality (VR/AR) are the most demanding. Current VR applications are little more than 360° video/panoramas. A low quality 360° video requires at least a 30 Mbps connection, HD quality streams easily surpass 100 Mbps, and retina quality(4k+) streams approach Gbps territory. However, there are still many things holding back its use beyond showrooms and proof of concepts, the most glaring problem being our network's capacity.

These days it seems that just about everything is getting smarter, from thermostats to refrigerators, and becoming ever-more connected. Each of these devices -- physical objects with data sensing, analyzing, and recording functions plus the ability to communicate remotely -- collectively form the "Internet of Things" (IoT). Clearly, expansion in the use of smart devices is an unstoppable force, but one thing could hamper this growth -- inadequate bandwidth. While most of the devices that comprise the IoT communicate wirelessly with the world, all the data that they send must be transmitted over a physical wireline network between wireless access points. High throughput, low latency and high reliability networks will be needed for applications such as video analytics in public safety and to support self-driving cars. [1].

In the upcoming 5G era, massive MIMO (Multiple-Input Multiple-Output), Carrier Aggregation, Multi-band support, and radio cell densification are impacting the bandwidth requirements, while the coexistence between macro, micro, pico, and small cells along with a centralized/virtualized processing environment are impacting the flexibility requirements. Fiber and optical access technologies are expected to play more and more important roles in the fronthaul and backhaul services to meet the aggressive performance goals of 5G.

Cable has been undergoing significant changes impacting network speeds which are not only caused by the ever-increasing residential data service tier growth rate but also due to an increasing number of services types being supported, such as business services and architectural changes in the HFC network. With the advent of network function virtualization and more and more applications running from the cloud, cloud computing related capacity is bound to make an impact. Therefore, it is safe to say that not one application or service is driving the increases in capacity but that the aggregate of multiple services has steadily maintained the exponential growth observed in broadband networks. Figure 1 depicts the historical growth of broadband service offerings including a constant growth rate extrapolation up to the year 2030.




Figure 1 - Exponential Growth of Broadband Service Offerings

2. Fiber Shortage Challenge in Cable HFC Network

In cable, the fiber access networks extend from the hub or headend to the fiber node. These fiber links are typically laid out by running a fiber bundle that passes by different nodes. From a splice point near a fiber node, a fiber cable with fewer fiber strands is trenched or strung to the node. In this initial HFC build-out, 6 to 8 fiber strands were typically dedicated to a node. This was ample at the time because cable companies primarily were offering broadcast services. Most fiber distances from node to hub are less than 25 miles, although in a few areas where hubs may have been consolidated, distances may be as long as 100 miles, leveraging the maximum distances allowed in DOCSIS 3.0 (Data Over Cable Service Interface Specification 3.0) and earlier versions. Figure 2 shows a representation of the fiber access network that extends from a hub.




Figure 2 - Fiber Distribution in Access Network

When these HFC networks were built in the mid 1990s, this capacity expansion has been addressed through RF spectrum expansion from 750 MHz to 860 MHz and 1 GHz as well as through node splits, meaning that the original fiber node serving area size of 500 households passed would be split or segmented into smaller portions. These newer smaller nodes consumed some of the spare fibers available.

In addition to spectrum expansion, capacity needs have also been addressed by improvements in transport efficiency. For the past 20 years, DOCSIS technology has been the main transport technology of cable data. After many protocol iterations, in DOCSIS 3.1 today, 4096 QAM (Quadrature Amplitude Modulation) upstream modulation and 16384 QAM downstream modulation are possible. To support these much higher fidelity RF signals, the traditional intensity modulation requires a very high signal to noise ratio (SNR). This SNR, > 47 dB for 16384-QAM and >43 dB for 4096-QAM, requires very high optical power levels. A side effect of using very high optical power levels is that they drive fiber into nonlinear operation and distort the signals within this fiber transmission. WDM is a technique leveraged to make efficient use of fiber resources, however, the nonlinear distortion in fiber at high optical power levels, makes wavelength multiplexing of analog carriers challenging. The end-result is that a single fiber cannot handle many analog optics carriers with high fidelity RF signals.

While the demand for residential data service rates has been steadily increasing, operators also expanded their services portfolio to include business and wireless base station connectivity. As a result of these new

services, some fiber strands in the access were re-purposed to address that need. In high demand business areas as few as 2 fibers per node may be available. Demand for higher upstream rates has been a key driver for a technology called Full Duplex DOCSIS, which requires passive N+0 topology implemented in what is called Distributed Architectures. In a distributed architecture, at least the PHY portion of the CMTS (Cable Modem Termination System) resides at the node. At the remote PHY device at the node, through RF and digital cancellation mechanisms, the strong downstream transmit signal impact on the weak upstream receive signal is cancelled. Figure 3 shows a traditional node serving area of about 500 HHP while Figure 4 shows the N+0 upgrade of that same serving area which could support Full Duplex DOCSIS systems.

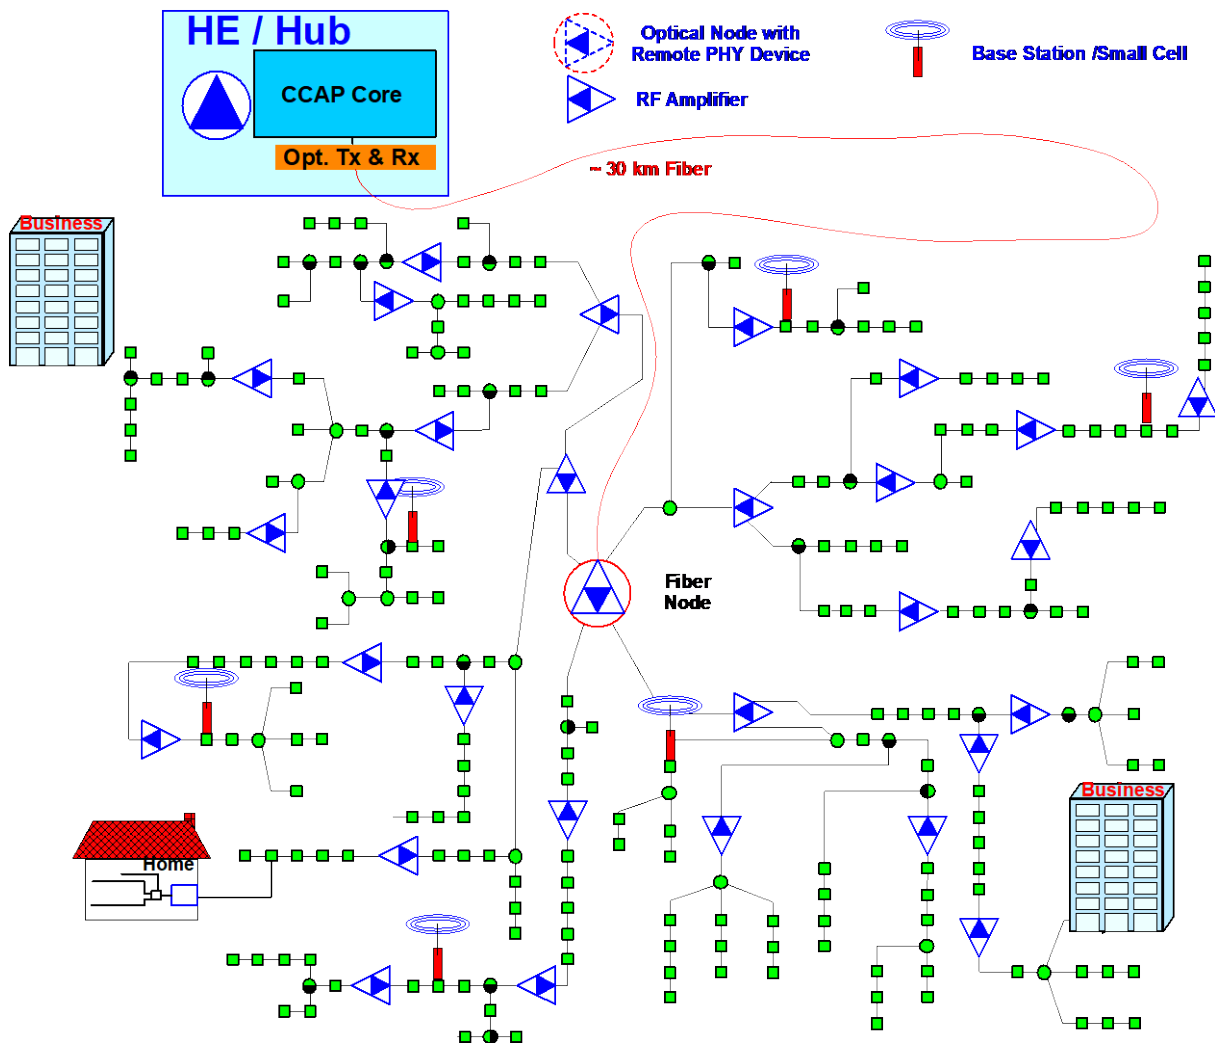


Figure 3 - Traditional Node Serving Area

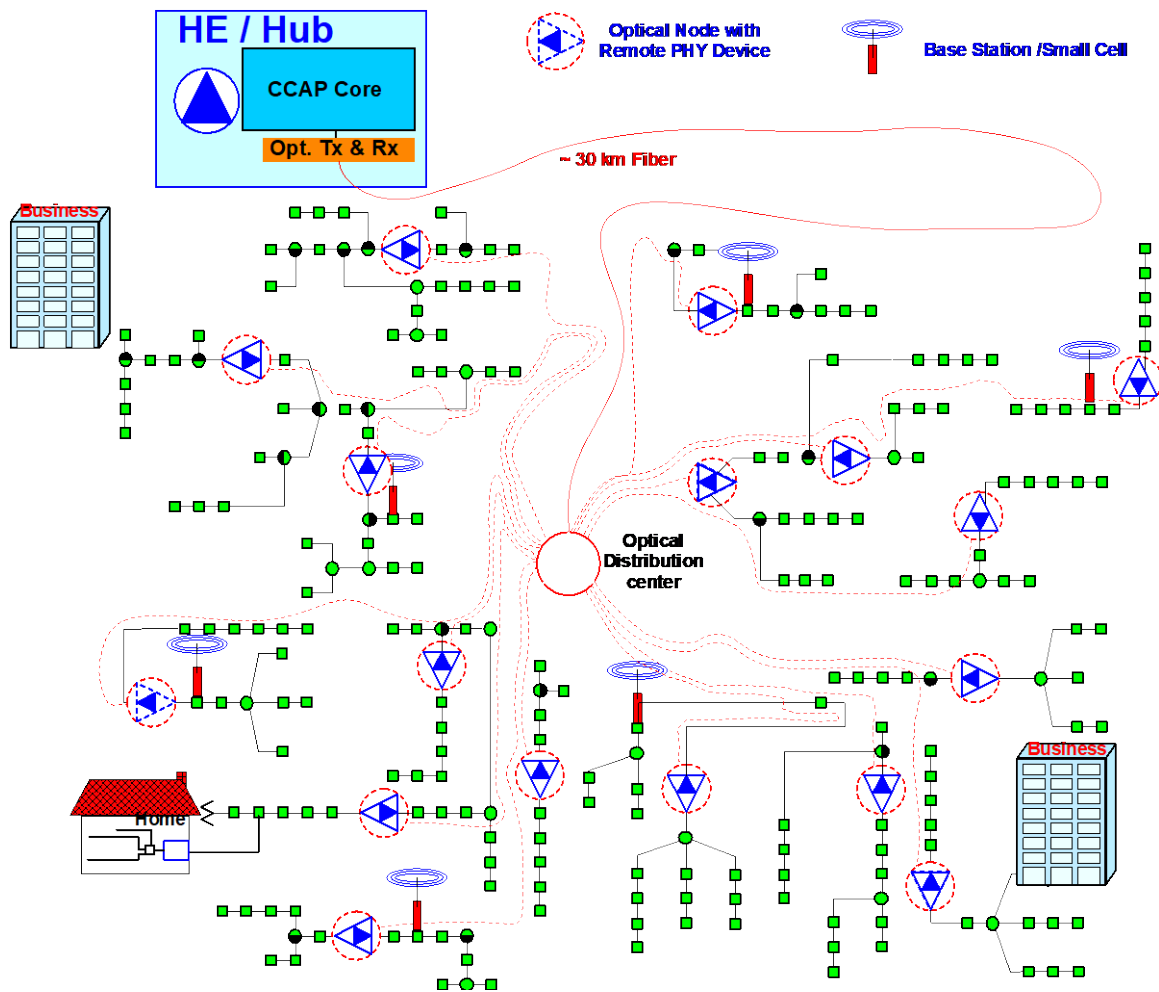


Figure 4 - Original Node Serving Area upgraded to N+0

In Figure 4, some of the amplifiers observed in Figure 3 are not used, while the rest of the amplifiers have been replaced by optical nodes. In N+0 networks like the ones shown in Figure 4, the distance between the node and the subscriber is probably less than 1000 feet [2], [11] through [15].

The cable fiber environment can thus be considered as sparse from a fiber strand count perspective but deep as the fiber is quite close to any potential subscriber. Just driven by the fiber network characteristics and topology, the technologies that the cable industry may consider and select, may not be the appropriate solution for other industries that don't share the same fiber environment characteristics as cable.

The key difference between Figure 3 and Figure 4 is the increase number of optical endpoints in Figure 4. As indicated earlier, transitioning to the architecture in Figure 4, also implies a transition to a Distributed Architecture. This means that it is no longer a centralized architecture, where the CMTS remains at the hub and analog optics carries RF signals to the nodes, but in a Distributed Architecture at least the PHY portion of the CMTS is remotely located in the node. Figure 5 shows the traditional centralized architecture with the CMTS located at the hub.




Figure 5 - Traditional Centralized Architecture

3. Coherent Optics Use Cases in Cable

The combination of the natural evolution of coherent optics technology, along with the 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 strategy.

This demand is not just for residential connectivity but also for business and cellular connectivity. A short-term strategy would not only result in an inefficient use of resources but also in having operators upgrading technology within a few years.

Coherent optics technology can be leveraged in cable following two general approaches. One is when used as a mean of multi-link aggregation and a second one is through direct edge-to-edge connectivity to the desire end-point. Following capacity growth trends, it is obvious that initially the aggregation use cases are going to out-number the direct edge-to-edge connectivity use cases. Within these broad categories, one can identify finer granularity sub-categories of these use cases.

3.1. Aggregation Use Case Scenarios

There are two-dominant aggregation use case scenarios, the first one being the use case of aggregation of multiple remote PHY or remote MAC/PHY devices. Figure 6 shows the distributed architecture example where the CMTS is split between a remote PHY device located at the fiber node and the MAC portion of the CMTS also known as CCAP core, is located at the hub. The hub and the aggregation node is connected through 100G or 200G per wavelength coherent optical transmission system.

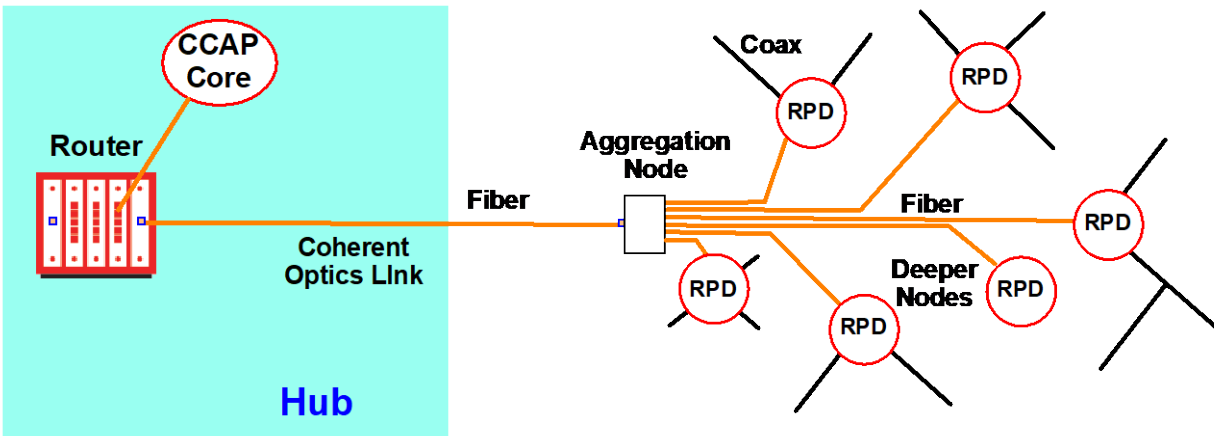


Figure 6 - Distributed Architecture with Traffic Aggregation at Original Node Location

In this distributed architecture, the optical transport is baseband digital optical transport and the DOCSIS RF is generated at the node by the remote PHY or remote MAC/PHY device. The downstream DOCSIS 3.1 RF spectrum is capable of carrying about 10 Gbps worth of data.

While wavelength multiplexing could have been used to carry traffic to and from RPDs or to carry the aggregate traffic from RPDs within a node serving area to the hub, cost and operation complexity of managing a large number of ports and managing wavelengths are important considerations. Section 3.3 covers scenario cost comparisons including some of these considerations.

The second use case is aggregation of PON networks through the connection of multiple optical line terminals (OLTs) (Figure 7).




Figure 7 - PON Aggregation, OLTs may be Collocated with Aggregation Node

In addition, although at a lower scale, there could be aggregation of business service traffic and/or aggregation of traffic from base stations. Cable operators' public announcements towards N+0

architectures over the next decade in part of their footprint make the aggregation of multiple nodes, the most likely aggregation use case where coherent optics will play a role. The advent of Full-Duplex-DOCSIS technology drives the evolution towards N+0 architectures. Cable fiber nodes were designed with 6 to 8 fibers dedicated to an individual fiber node covering around 500 HHP. Since some of these fibers may have been re-purposed to provide business connectivity or for node splitting, a cable service provider can safely rely on having just two fibers available.

Based on traffic demand and traffic growth particularly in the upstream, node demographics and topology, proximity to businesses and other criteria, we could assume that over the next decade a portion of these nodes will gradually migrate to N+0.

In these scenarios, a sustained market size of approaching 100,000 transceiver units after four years has been estimated. Since the need for coherent transceivers supporting N+0 migration is expected to dominate the market, wireless backhaul and Business link aggregation use cases were not included in this estimate.

3.2. Direct Edge-to-Edge Connectivity Use Case Scenarios

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. 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.

As demand for capacity further increases, the capacity required by each RPD will surpass 10 Gbps and operators will likely be required to split that N+0 node and/or extend the RF coaxial spectrum. Such an environment, urges a review of the optical transport approach. This implies that the deeper node instead of requiring 10 Gbps, could require 40 Gbps or higher capacity. At these capacity levels, one has to explore the advantages of an aggregation strategy versus a direct edge-to-edge connectivity strategy. Today a multi sector, multi carrier cellular system requires a 40 Gbps feed. With a capacity compound annual growth rate in cellular greater than 50% [1] and the expected proliferation of 5G access points in the near future, direct edge-to-edge connectivity using coherent optics is quite appropriate.

As is described in the next section, coherent technology, because it operates at the lowest optical power and more efficiently uses the fiber wavelength spectrum, is the best neighbor to other optical transport technologies.

3.3. Economics of Coherent Optics in the Aggregation Use Case Scenario

For the Aggregation Use Case depicted in Section 3.1, other digital optics solutions can achieve similar results in meeting bandwidth demand. Dense wavelength division multiplexing (DWDM) + Direct Detection (DD) allows the combination of incremental 10 Gbps wavelengths to meet the capacity demand of aggregating traffic from multiple end points at a single location (such as shown in Figure 6 and Figure 7). However, DWDM + DD depletes a scarce spectral resource when scaling to meet bandwidth demand, whereas coherent optics can provide sufficient bandwidth on a single wavelength with very high spectral efficiency. The opportunity cost of a wavelength can be quantified in many ways, such as the missed revenue opportunity for leasing high-bandwidth wavelength services to enterprise customers.




Figure 8 - DWDM + Direct Detection for Aggregation Node

The real economic value of coherent optics versus DWDM direct detection (DD) scheme (see Figure 8) can be shown in the scale provided when aggregating multiple end points, whether for Remote PHY Devices or for OLT's. A passive edge-to-edge DWDM + DD solution requires more expensive Fixed DWDM 10 Gbps optical transceivers from the aggregation point to the end points because of longer transmission distance. A coherent optics solution can use less expensive 1310 nm 10 Gbps optical transceivers to connect the aggregation point to the end points with a couple thousand feet distance. According to IHS, the 1310 nm transceiver are 1/3 the cost of Fixed DWDM transceivers today [15], [16], [17]. Further, the Multi-line gearbox (MLG) in the coherent optical transceiver provides flexible bandwidth allocation to end points.

Assuming current coherent optical transceiver (CFP2-DCO) costs for metro networks, coherent optics becomes a more cost-effective solution than wavelength multiplexing when aggregating 14 or more end points as shown in Figure 9. This breakeven point may not seem too impressive until considering the following factors: 1) most Node+0 architecture call for aggregation of 12-18 mini-nodes or Remote PHY Devices; 2) Coherent optics can reach up to 80 km without the need for dispersion compensation; 3) Coherent optics uses a single wavelength up to 200 Gbps, directly connected to the CMTS at a hub without the need for demultiplexing; and 4) the Coherent optics transceivers are not optimized for the Cable Access Network today.




Figure 9 - Breakeven End Points for Coherent Optics vs. Wavelength Multiplexing (2017)

The economics of coherent optics will improve substantially in the coming years. A recent IHS Report [17] projects mature DWDM + DD technology costs will decrease up to 3% annually while coherent optics costs will decrease up to 15% annually. Assuming the 15% annual cost decrease for the coherent Optics transceiver and a 3% annual cost decrease for other optical components, the breakeven shifts to ten end points in the year 2020 as shown in Figure 10. That is, 100 Gbps aggregation will have similar costs whether using wavelength multiplexing or Coherent Optics. This Coherent Optics cost improvement could accelerate with a solution optimized for the Cable Access Network with more benefits from operation.




Figure 10 - Breakeven End Points for Coherent Optics vs. Wavelength Multiplexing (2020)

4. Digital Coherent Optical Transmission Technology

4.1. Optical Transmission Technology Evolution

As shown in Figure 11, optical transmission technologies have evolved over multiple generations. The inventions of semiconductor lasers and low-loss single-mode fiber (SMF) were the major breakthroughs in the 1970s. From the 1980s to early 1990s, electrical time-division multiplexing (ETDM) was the core technology. The invention of the Erbium-doped fiber amplifier (EDFA) in the 1990s and the first commercial use of 8×2.5 -Gbit/s WDM in 1996 were important milestones. The next technological leap occurred during the mid 2000s. The rapid development of silicon-based electronic chips and maturing of DSP technologies reinvigorated coherent detection. This has increased the spectral efficiency of optical signals to 2 bits/s/Hz, and optical transmission has entered the stage of digital coherent transmission with the use of four-dimensional orthogonal signals.




Figure 11 - The Evolution of Optical Transmission Technologies

Besides the optical terminal system evolution, optical fiber media has evolved from early loss-reduction dispersion-managed fiber (DMF), to increased effective area fiber and fiber with further reduced propagation loss. These developments have greatly helped overcome linear and nonlinear impairments in

optical fibers. In the future, space division multiplexing (SDM) is likely to become a technological turning point for addressing optical capacity crunch challenge.

It is noted that the new generation of technology can take full advantage of the technology that has been developed, for example, coherent technology along with polarization multiplexing can easily integrate into existing DWDM systems with much improved SE. In the meantime, new technologies can continue to use the existing deployed optical fiber links. Actually, a standard SMF is more tolerant than non-zero dispersion shifted fiber to nonlinear effects especially for coherent optical systems because of nonlinear phase-matching conditions.

4.2. Rebirth of Coherent Optics

Coherent optical receivers have initially received significant research interest in the 1980s. At that time, no optical pre-amplification was used in front of the receiver. As the local oscillator (LO) signal typically has a much higher power than the received signal, it can be used for coherent amplification gain. This can potentially increase the receiver sensitivity with up to 20 dB in comparison to optically unamplified direct detection.

However, the invention of EDFAs made the shot-noise limited receiver sensitivity of the coherent receiver less significant. This is because the Optical Signal-To-Noise Ratio (OSNR) of the signal transmitted through the amplifier chain is determined from the accumulated amplified spontaneous emission (ASE) rather than the shot noise. Both the surge in EDFA based optical communication solutions as well as the challenges with coherent detection interrupted research and development activities in coherent communications for nearly 20 years. Even though the increase in capacity enabled by EDFAs and WDM has scaled well in the past, a hard limit on capacity exists while non-coherent On/Off Key (OOK) modulation is used. Due to these pressures, advanced modulation formats, capable of transmitting more than one bit per symbol, became a highly desirable technical advancement.

On the other hand, the development and maturity of high-speed digital integrated circuits has offered the possibility of using DSP capability of providing us with simple and efficient means for equalizing fiber transmission impairments, demultiplexing two polarizations, and estimating the carrier phase. Through the adoption of high-order modulation formats, higher spectral efficiencies can be reached through the reduced symbol. Furthermore, only coherent detection permits convergence to the ultimate Shannon limits of spectral efficiency in theory.

4.3. Advanced Modulation Format

As described earlier, the use of coherent detection and DSP enabled system fully leveraged the benefits of advanced modulation formats and provided previously unavailable functionality in systems with direct detection, such as the use phase and the state of polarization as means to convey information. The migration from traditionally used OOK modulation formats to formats with more bits per symbol leads to a reduction of the symbol rate and narrowed spectral widths. Therefore, higher spectral efficiencies, per fiber capacities, and then cost per bit can be realized. This is the main motivation for a system upgrade to higher-order modulation. The constellation diagrams of currently most popular modulation formats are shown in Figure 12, [3] through [7].




Figure 12 - Constellation Diagrams of Various Modulation Formats

From a system perspective, there are multiple optimization parameters when making a decision on these modulation formats, metro access systems generally emphasize the cost, complexity, and receiver sensitivity since the impact of transmission impairments are small compared to long haul systems.




Figure 13 - Theoretical BER Curves for Various Modulation Formats.

It is possible to derive theoretical limits on the performance of different modulation formats and performance being purely impaired by additive white Gaussian noise. The transmitter and receiver are assumed to have ideal matched filters. The noise level is described here using the signal processing convention of E_b/N_0 , where E_b is the mean energy per transmitted bit and N_0 is the mean noise energy per symbol. This metric enables comparison between modulation formats with differing parameters but at identical bit rates, as the SNR is normalized to the number of bits per symbol of the modulation format (unlike, for example, E_s/N_0).

As a result of an increasing number of bits per symbol, noise performance degrades as the Euclidean distances between the symbols become smaller as shown in Figure 12. High-order QAM formats exhibit a significantly better noise performance than high-order phase modulation formats for a certain number of bits per symbol, Square QAM formats in particular, due to the more optimum allocation of symbols on the complex plane. In comparison with 16 PSK as shown in Figure 13, Square 16-QAM has an OSNR performance gain of about 4 dB, for instance. High-order modulation formats offer a way of relaxing the requirements on PMD since a certain group delay difference has a smaller impact on neighboring pulses for reduced symbol rates and longer pulse durations.

4.4. Polarization-Multiplexed (PM) QAM Optical Transmitter




Figure 14 - Architecture of a Standard Mach-Zehnder Modulator

In this section, a most commonly used transmitter architecture is presented and is compatible to arbitrary QAM constellations. As an example, Figure 14 presents an optical nested IQ modulator with dual Mach-Zehnder modulators (MZM) for Quadrature Phase Shift Keying (QPSK) modulation format. 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 the 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.

As illustrated in Figure 15 (a), the recombination of the two Binary Phase-Shift Keying (BPSK) signals with $\pi/2$ phase difference yields a QPSK signal. In a similar way as in QPSK generation, the mechanism for 16-QAM can be considered as two Pulse Amplitude Modulation 4 (PAM-4) signals inference with the phase shifted by $\pi/2$ (see Figure 15 (b)). It is a superposition of vectors on a complex plane. This BPSK or 4-level PAM-4 signals can be obtained by DACs in commercial implementations. One of the most

important parameters of the QAM signal modulation 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

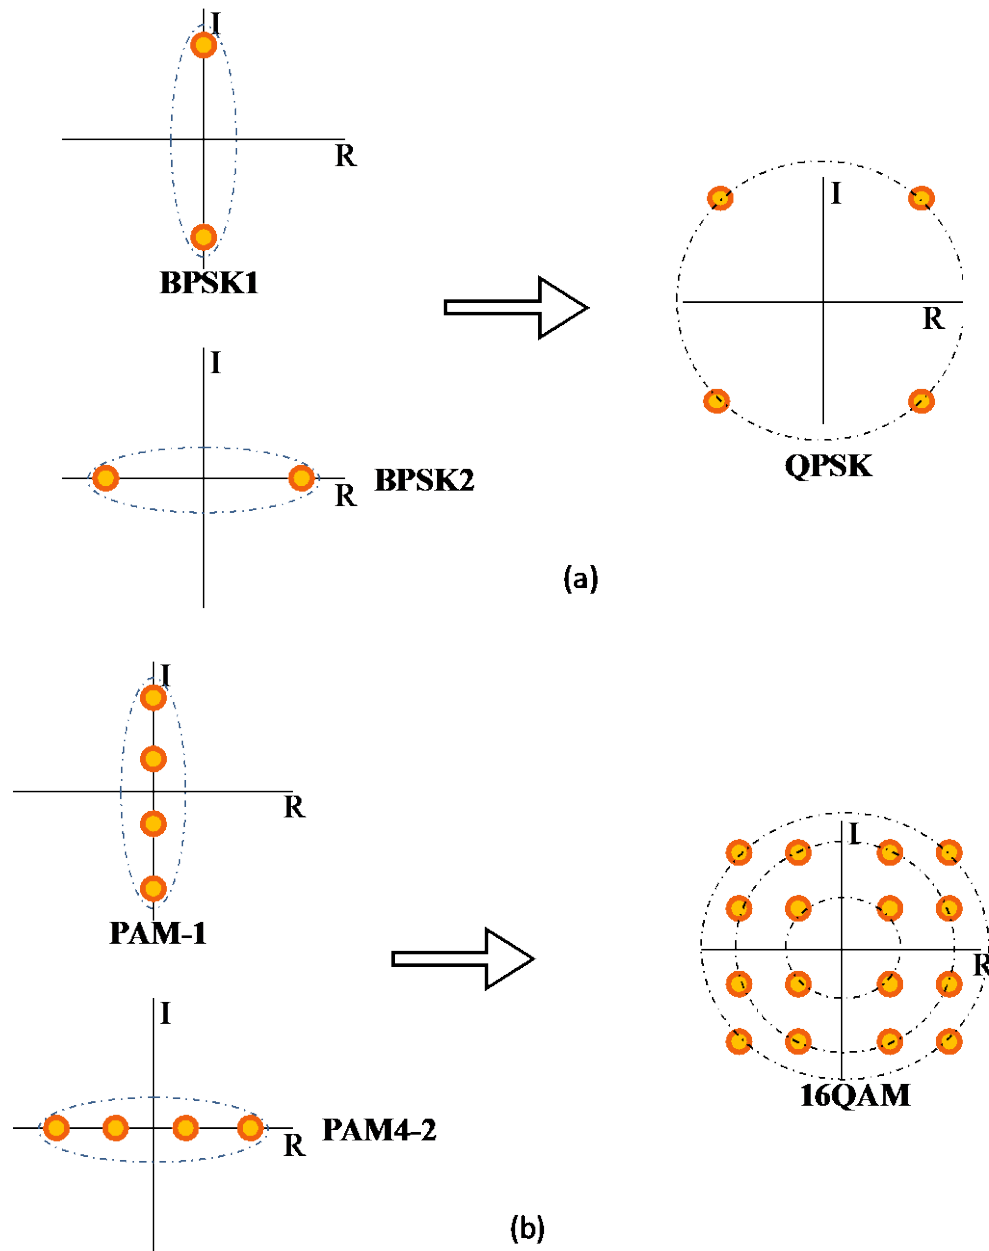


Figure 15 - Illustration of QPSK (a) and 16QAM (b) Signal Generation.

Back to the QPSK example, the transfer function of a single MZM is given by

$$E_{out} = E_{in} \cos(\pi V(t)/V_{\pi})$$

where E_{in} is the input optical signal, E_{out} is the output optical signal of the MZM, $V(t)$ is the driving signal, and V_{π} is a characteristic of the MZM and determines the amplitude required for the driving signal. Thus, the transfer function of this IQ modulator is given by

$$E_{out} = E_{in} \sqrt{\cos^2(\pi V_a(t)/V_{\pi}) + \cos^2(\pi V_b(t)/V_{\pi})} * e^{j \tan^{-1} \frac{\cos(\pi V_b(t)/V_{\pi})}{\cos(\pi V_a(t)/V_{\pi})}}$$

If $V_a(t)$ and $V_b(t)$ take on one of two values $\{0, V_{\pi}\}$, then the phase shift induced on the IQ modulator is one of four values as shown in Table 1 - Phase Shifts in an IQ QPSK Modulator below.

Table 1 - Phase Shifts in an IQ QPSK Modulator

$V_a(t)$	$V_b(t)$	$\cos(\pi V_a(t)/V_{\pi})$	$\cos(\pi V_b(t)/V_{\pi})$	$\tan^{-1} \frac{\cos(\pi V_b(t)/V_{\pi})}{\cos(\pi V_a(t)/V_{\pi})}$
0	0	1	1	$\pi/4$
0	V_{π}	1	-1	$-\pi/4$
V_{π}	0	-1	1	$3\pi/4$
V_{π}	V_{π}	-1	-1	$5\pi/4$

For generating dual-polarization modulation formats, typically two triple MZMs are used in parallel, each modulating an orthogonal polarization (see Figure 16). The two unmodulated carriers come from the same laser and are split into orthogonal linear polarizations with a polarization beam splitter (PBS), before the two-independent polarization modulated signals are multiplexed together with a polarization beam combiner (PBC).

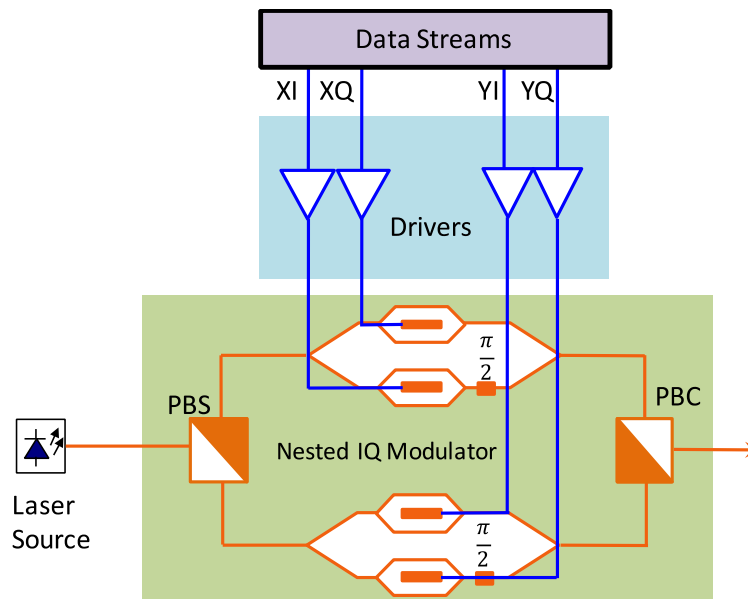


Figure 16 - Optical Transmitter for PM-QAM Modulation Formats

4.5. Digital Coherent Receiver

4.5.1. Homodyne/Intradyne/Heterodyne

In a coherent receiver, an LO is used to down-convert its signal with the incoming signal lightwave for demodulation. Depending on the intermediate frequency f_{IF} defined as $f_{IF} = f_s - f_{LO}$, coherent detection can be realized using a homodyne, intradyne or heterodyne receiver as illustrated in Figure 17, where $Bandwidth_s$ is optical signal bandwidth [8], [9], and [10].

In a homodyne receiver, the LO and transmitter laser have the same frequency and the phase difference should be zero (or a multiple of 2π). Intra-dyne detection is similar to homodyne detection, with the exception that a frequency offset between the LO and transmitter laser exists, but the f_{IF} is chosen to fall within the signal band by roughly aligning the f_{LO} with f_s . In heterodyne detection, the difference between the LO and transmitter laser frequency is chosen higher than the electrical signal bandwidth, the entire optical signal spectrum is directly translated to an electrical bandpass signal centered at the f_{IF} for further electronic processing.




Figure 17 - Three Coherent Detection Schemes: (a)Homodyne, (b)Intradyne, (c)Heterodyne

Compared with heterodyne and homodyne detection, intradyne detection has the advantage that the processing bandwidth is relaxed for the optical and electrical components. It is also noted that both heterodyne and intradyne detection that uses single-ended detection is vulnerable to RIN from LO. This can be solved through balanced detection, which requires a total of 4 photo-diodes and an optical hybrid with 4 outputs, each shifted by 90° degrees.

4.5.2. Coherent Receiver Architecture

The fundamental concept behind coherent detection is to take the beating product of electric fields of the modulated signal light and the continuous-wave LO. To detect both IQ components of the signal light, A 90° optical hybrid is utilized. A key building block of such a hybrid is an 2x2 optical coupler with its property of a 90° phase shift between its direct-pass and cross-coupling outputs via multimode interference (MMI) coupler. By combining such optical couplers into the configuration shown in Figure 18, together with an additional 90° phase shift in one arm, a detection of real and imaginary parts can be achieved. Balanced detection is usually introduced into the coherent receiver as a mean to suppress the DC component and maximize the signal photocurrent.

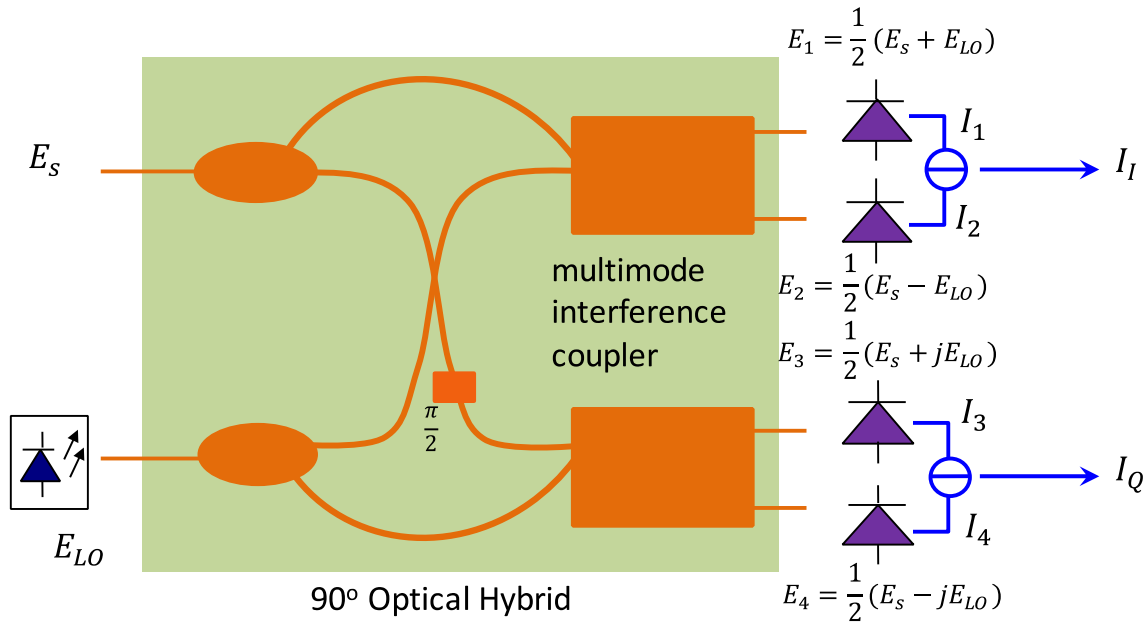


Figure 18 - Configuration of Phase-diversity Coherent Receiver

Output photocurrents from balanced photodetectors are then given as

$$I_I(t) = I_1(t) - I_2(t) = R\sqrt{P_s P_{LO}} \cos\{\varphi_s(t) - \theta_{LO}(t)\}$$

$$I_Q(t) = I_3(t) - I_4(t) = R\sqrt{P_s P_{LO}} \sin\{\varphi_s(t) - \theta_{LO}(t)\}$$

where R is the responsivity of the photodiode, P_s and P_{LO} are the power of the optical fields for incoming and LO signal, respectively. It is possible to estimate the phase noise $\theta_{LO}(t)$ varying with time and restore the phase information $\varphi_s(t)$ through the following DSP on the intradyne detected signal.

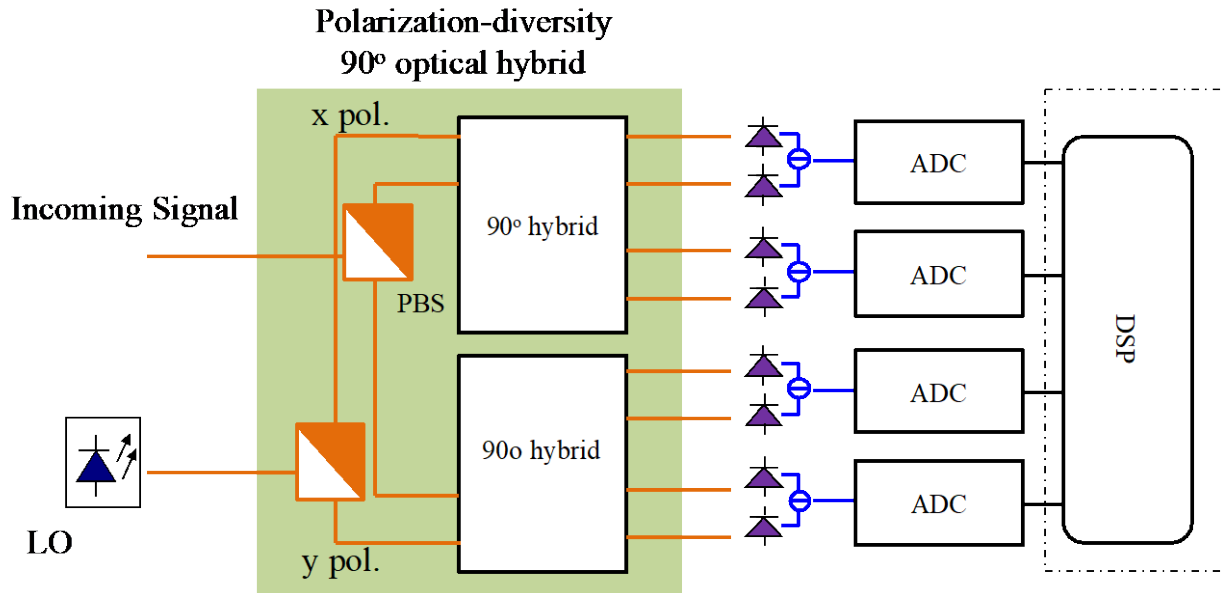


Figure 19 - Configuration of Phase and Polarization Diversity Coherent Receiver Architecture

The schematic diagram of a polarization multiplexed coherent receiver is shown in Figure 19. Both the incoming PM signal and LO are split into two orthogonal polarizations using a PBS, after which the co-polarized signal and the LO are mixed in two 90° optical hybrids to produce an in-phase and quadrature components for each polarization. The four signals are then digitized by four ADCs after which DSP can be performed for signal demodulation.

4.5.3. Digital Equalization Algorithms

Current 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. As shown in Figure 20 in detail, transmitter DSP functions include symbol mapping and signal timing deskew adjustment, optional pre-distortion for dispersion or self-phase modulation, and software-programmable capability of supporting multiple modulation formats and encoding schemes. Transmitter DSP also allows compensating nonlinearities induced by the electrical driver and the optical modulator.

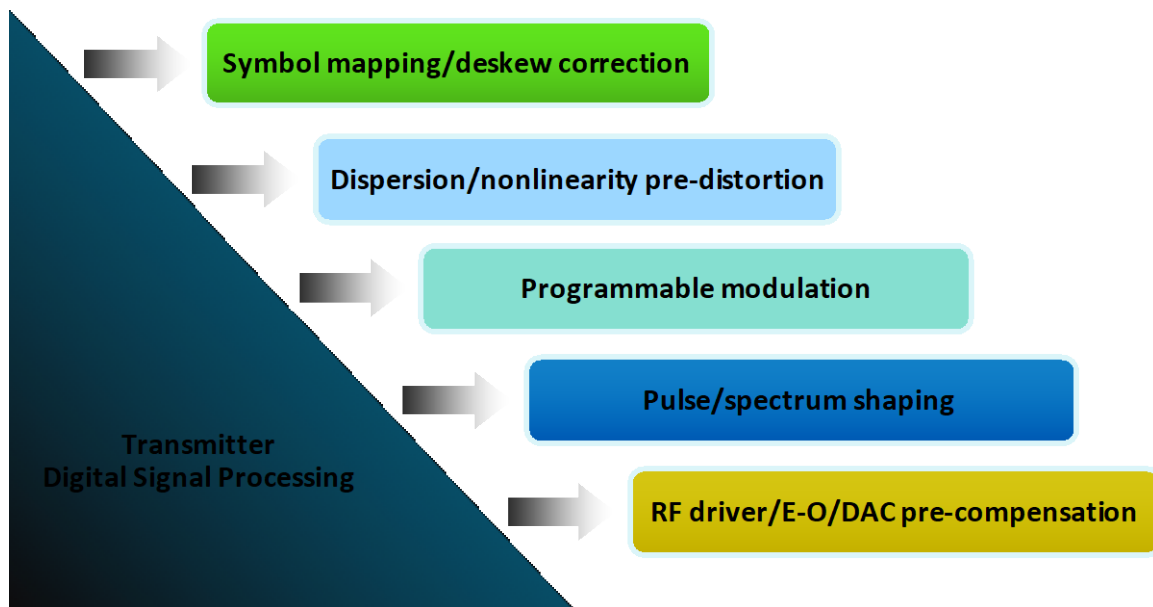


Figure 20 - Transmitter Side DSP Functions

The major benefit of the transmitter DSP is to perform pulse shaping and thus engineer the spectrum of signal, such as generating Nyquist pulse. Pulse shaping controls the spectrum to increase the spectral efficiency or reduce the nonlinear impairments. Basically, there are three types of pulse shaping filters, sinc, raised-cosine, and Gaussian filters. Table 2 summarizes their corresponding transfer function, spectral and pulse shapes, and eye diagrams of QPSK filtered driving signals.

Transfer function in frequency domain (B: bandwidth, T: symbol duration, normalized to 1)	Frequency domain	Impulse response	Eye Diagrams
$H(f) = \begin{cases} 1 & \text{if } f \leq \frac{B}{2} \\ 0 & \text{otherwise} \end{cases}$ <p>(a) Sinc filter</p>			
$H(f) = \begin{cases} T & \text{if } f \leq \frac{1-\beta}{2T} \\ \left[\frac{T}{2} \left(1 + \cos \left[\frac{\pi \cdot T}{\beta} \left(f - \frac{1-\beta}{2T} \right) \right] \right) \right] & \text{if } \frac{1-\beta}{2T} < f \leq \frac{1+\beta}{2T} \\ 0 & \text{otherwise} \end{cases}$ <p>(b) Raised-cosine filter</p>			
$H(f) = e^{-\frac{\ln(2) \cdot f^2}{2 \cdot B^2}}$ <p>(c) Gaussian filter</p>			

Figure 21 - Three Pulse Shaping Filters with Different Transfer Functions

Transmitter-side DSP enables more flexibility not only in channel impairment pre-compensation but software configuration for elastic optical networks. In correspondence to the operation of the transmitter, the major advantage of receiver-side DSP stems from the ability to arbitrarily manipulate the electrical field after the ADC enables the sampling of the signal into digital domain. As shown in Figure 22, the fundamental DSP functionality in a digital coherent receiver for PM-QAM signals can be illustrated by the following flow of steps and their correlations from structural and algorithmic level of details.

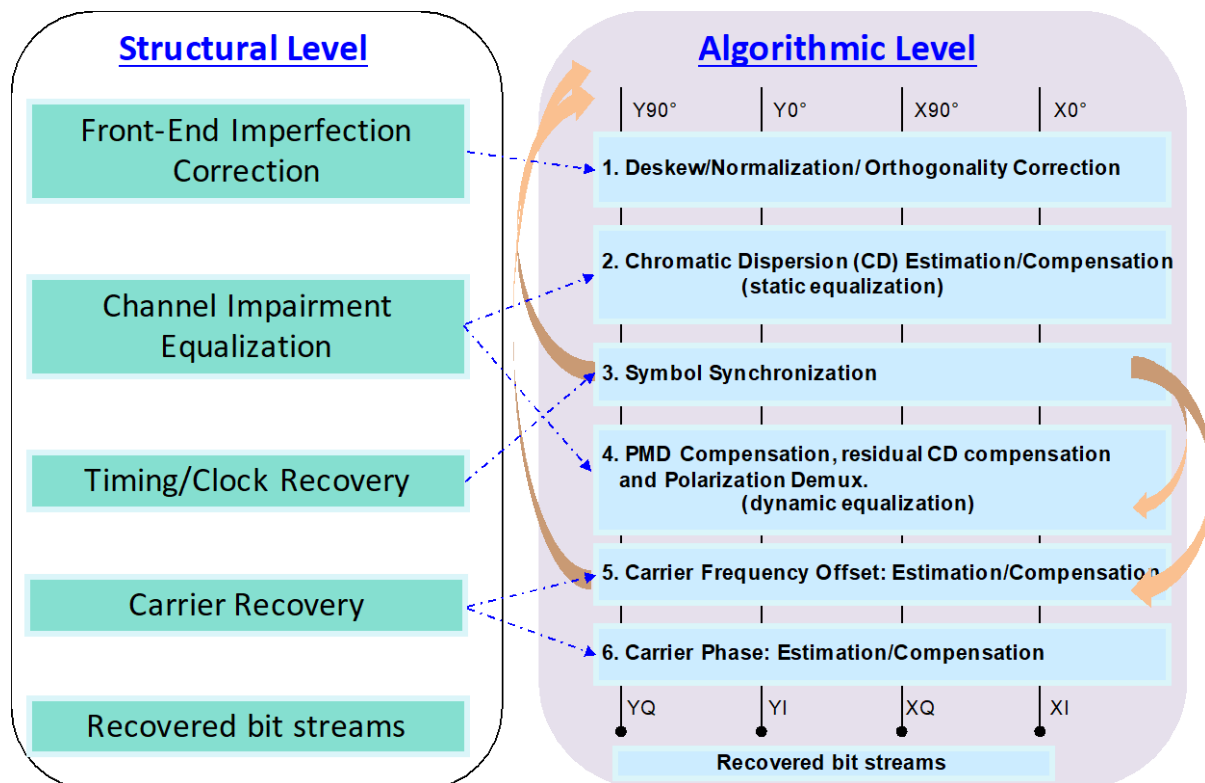


Figure 22 - DSP Flow in a Digital Optical Coherent Receiver

First, the four digitized signals after an ADC are passed through the block for the compensation of front-end imperfections. The imperfections may include timing skew between the four channels due to the difference in both optical and electrical path lengths within a coherent receiver. Other types of front-end imperfections can be 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 90-degree phase shift.

Second, the major channel transmission impairments are compensated through digital filters, in particular, CD and PMD. The static equalization for CD compensation is performed first because of its independence of SOP and modulation format and 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 process between the blocks of clock recovery and polarization demultiplexing for achieving the symbol synchronization (see arrows in Figure 22). 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 the constellation rotation at the intradyne frequency.




Figure 23 - Constellation Evolutions for QPSK/8QAM/16QAM Signals

Finally, the carrier phase noise is estimated and removed from the modulated signal, which is then 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 from those detailed in Figure 22 because of different design choices. Besides feed-forward process, it is possible to perform joint process and feedback among different process blocks such as clock recovery and polarization demultiplexing as mentioned above.

The constellation evolutions show examples (Figure 23) of the received signal after linear transmission over uncompensated SMF link with EDFAs only. Note that the results in this proposal are based on 32-GBaud rate. It is also noted that the impairments of frequency offset of 0.1 GHz, 100-kHz linewidth of Transmitter laser and LO are induced with 20000 ps/nm accumulated CD.

Coherent detection and DSP were the key enabling technologies in the development of 100G optical transmission systems. The next-generation coherent optical systems will continue this trend with DSP playing even more ubiquitous role at both transmitter and receiver. Although the specific algorithms for each process block are typically different because there are various realizations of the same process block in the implementation level, the generic functions in the structural level or function abstractions are similar for all major commercial products.




Figure 24 - Transistor Counts for Both CPUs and ASICs in Optical Communication

Key to advancing the use of DSP in optical communications is the consistent improvement in complementary metal–oxide–semiconductor (CMOS) technology. As shown in Figure 24, DSP ASICs for coherent optics are following the complexity trend with approximately ~1 order of magnitude difference between CPUs. As feature node sizes have shrunk and design tools improved over the years, the maximum complexity (and hence functionality) possible in an ASIC has grown from thousands to several hundred million gates. Today's coherent DSPs use 16/14 nm and tomorrow's will use 10/7 nm [10].

Not only is the feature size in a coherent ASIC decreasing exponentially but the sampling rates of CMOS data converters are also exponentially increasing to support higher symbol rates and hence bit rates. In the meantime, the advancement of high ENOB is very important for higher-order modulation formats. Today's coherent ADC/DACs use 56-64GSamples/s and tomorrow's will use 90-128GSamples/s, as shown in Figure 25, [10].




Figure 25 - ADC/DAC Development in Optical Communication

4.5.4. Coherent Transceiver Module Evolution

First coherent 100 Gb/s interfaces were built using a discrete line card architecture by large network equipment manufacturers. Next to custom line cards, 100G was also defined in a standard line interface module with a total power of less than 80W in 5”x7” OIF compliant Multi-Source Agreement (MSA) modules for QPSK modulation signals. The intention of the 100G MSA module is to support long-haul applications beyond 2,000 km. The second non-pluggable MSA module is 4”x5” with less than 40W of power as shown in Figure 26. To provide high density and low-cost coherent transceivers, hot pluggable modules are being implemented with small form factors. The power consumption for a coherent transceiver mainly includes the power consumed by the ASICs (DACs, ADCs, DSPs), lasers, modulator drivers, and TIAs. In order to pack all the components in a small form factor, the power dissipation need to be reduced significantly.

There are two classes of pluggable transceivers for the optical transmission nodes: one is the line-side transceivers the other one is the client-side transceivers. Most of today’s development of client-side transceivers are focusing on CFP2, CSP, and QSFP+ formats without coherent optics. For the line-side transceiver with coherent optics, there are initial emerging CFP modules available and moving to CFP2 and CFP2 format. The power requirement for a CFP is <28-32W, while a CFP2 has to fulfill a maximum of <9-12W. CFP2 can be implemented with analog host interfaces and the ASIC on an external board or digital CFP2 includes all optics and ASIC.

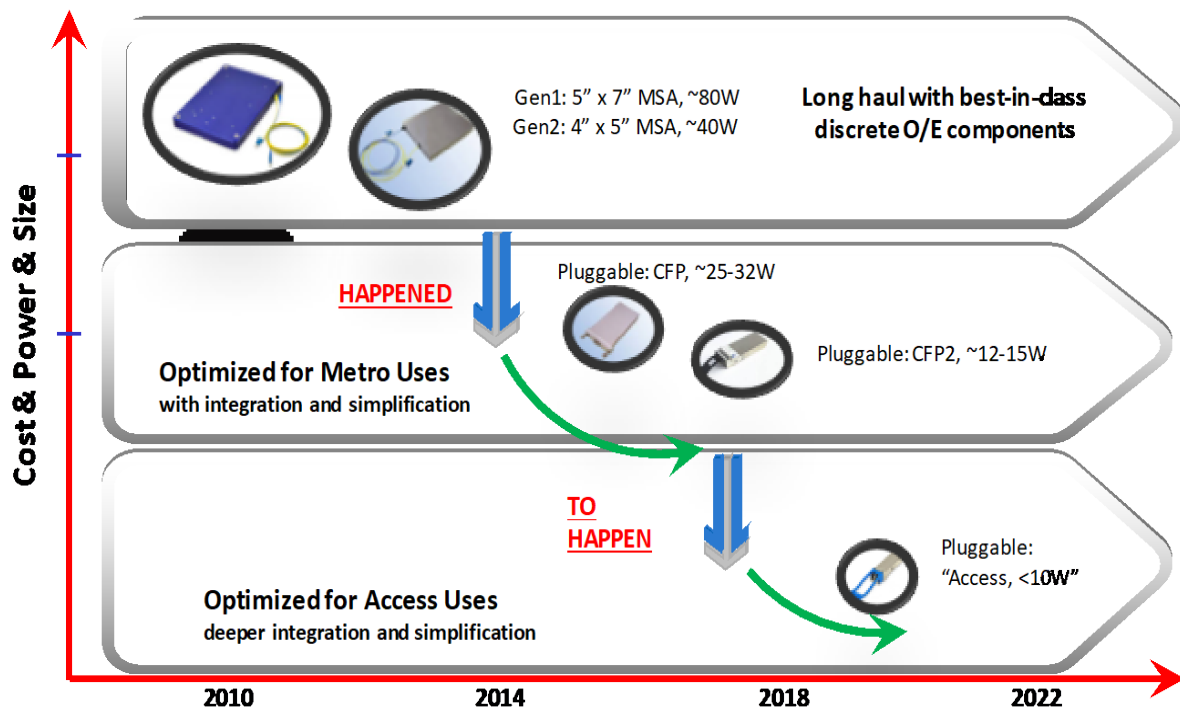


Figure 26 - Evolution of Digital Coherent Optical Module

Photonic integration circuit appears to be a must for CFPx modules instead of separate LiNbO_3 modulators and InP/PLC receivers. To achieve the cost targets, the use of electronic integration and photonic integration needs to be implemented to reduce component count and improve manufacturability. Another dominant cost for the DSP and optics is the packaging, one can further reduce cost, power, and footprint by co-packaging the DSP and optics.

Recent technology advances and ongoing price drop further open the window of opportunity for the application of coherent optics in access networks. It is envisioned that the migration of coherent optics from long haul and metro to access domain will need to further optimization and deeper integration to lower the cost and complexity.

4.6. Coherent Optics for Access

Coherent detection for access networks enables the superior receiver sensitivity that allows for extended power budget, and the high frequency selectivity enabling dense WDM. Moreover, the linearly recovered signal provides additional benefits to compensate for the linear transmission impairments such as CD and PMD, and efficiently utilize the spectral resource and benefiting future network upgrades. In the cable access environment, coherent optics allows operators to best leverage the existing fiber infrastructure to withstand the exponential growth in capacity and services. However, there are several engineering challenges. The coherent technology in 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 process. The coherent pluggable modules for metro solution has gone through CFP to CFP2 via MSA standardization for smaller footprint, lower cost, and lower power dissipation. However, it is still over-

engineered, too expensive, and too power hungry. Therefore, it is not efficient and practical for access applications.

Access network is totally different environment as compared to long haul and metro. To reduce the power consumption and thereby meet the size and cost requirements for access applications, development of both low-complexity ASIC and optics is essential. In particular, co-design of a DSP ASIC and optics to trade performance against complexity, cost and power consumption is imperative.




Figure 27 - Three Major Factors for Impacting Coherent Optics in Access Networks

Besides the natural technology advancement of coherent optics and reuse of existing fiber infrastructure, there are three major factors that impact on the cost and power reduction and eventually lead to successful deployment in access networks. As shown in Figure 27, We know that the access network is the largest component of the network in terms of physical/geographic size and amount invested. Increasing shipment volume for both long haul, metro, and access will drive component and equipment pricing down quickly. The following subsections will illustrate the details on optimization of the digital, optical, and electrical complexity and standardization and interoperability.

4.6.1. DSP Optimization

Naturally, the shorter transmission reach means less distance-dependent signal degradation, requiring less link equalization (i.e., fewer digital filter taps) and less processing in the DSP ASIC for impairment compensation, such as CD compensation, PMD compensation, and tracking the signal state of polarization. A reduction in OSNR performance is acceptable for shorter-reach access applications, which allows for a lower sampling rate and resolution of ADCs/DACs and fewer bits to be carried through the

DSP. Because of shorter distance and less demand on the link budget, soft-decision FEC (SD FEC) encoder and decoder, the major blocks in terms of power dissipation of ASIC, can also be significantly reduced on the complexity by either using hard-decision FEC, less overhead SD FEC, and/or decreasing the number of iterations. Figure 28 shows our analysis on power consumption of typical ASIC from long haul, metro, to access applications. In addition to the percentage change of each constituent element, the total energy consumption is significantly reduced.




Figure 28 - Analysis on Power Consumption of Typical ASICs.

4.6.2. Optical and Electrical Components Simplification

Both optical sources, usually lasers as transmitter and LO, are crucial building blocks to optimize the system cost and performance. Low-cost, small-footprint lasers with relatively large linewidth are preferred over the costly narrow linewidth external cavity lasers in such context, provided acceptable degradation in system performance. As shown in Figure 29, the analysis shows little impact on 16-QAM signals when the linewidth is increased from 100kHz to 1MHz. This will allow the use of cheap laser sources for access coherent systems. It is worth mentioning that higher order modulation formats such as 32/64-QAM are more sensitive to phase noise.




Figure 29 - Impact of Linewidth (Phase Noise) on 16QAM Signal Performance

A miniaturized optical IQ modulator is another key component for coherent optical transceiver. Currently, there are three materials for IQ modulators, i.e., LiNbO₃, Indium-phosphide (InP), and silicon photonics (SiPh). To fit inside a pluggable module like CFP2-ACO/DCO, only InP and SiPh based modulator can be applicable for such small optical modules. A booster optical amplifier, either SOA or EDFA, is often required to compensate coupling and modulation loss in order to meet the power budget of long haul or metro applications. However, this optical amplifier can be avoided to reduce the cost because of the access environment with a less demanding optical link power budget.

Using low-bandwidth, low-cost electrical and opto-electronic components such as driver, modulator, photodetector, TIA, is another consideration to reduce the overall cost for access networks. These components operate beyond their specified bandwidth, the introduced impairments can be equalized in the DSP section to balance the cost of components and the DSP complexity. As shown in Figure 30, two pre-comp schemes at DAC are presented here, including least mean square and constant modulus algorithm (CMA) algorithms. The criterion is to minimize the difference between the received data and convolution result of estimated channel coefficients.




Figure 30 - Pre-Compensation Algorithms for 32-GBaud Signals with 22-GHz System Bandwidth

Figure 31 shows the results of 80-km SMF transmission of PM-QPSK/16-QAM in our proof-of-concept testbed, where error vector magnitude (EVM) is plotted as a function of baud rate. Due to the limited bandwidth of low-cost RF amplifiers (10 GHz RF drivers from Picosecond Pulse Labs 5822B) and the optical IQ MZM (14 GHz IQ-MZM from Covega LN86S-FC), EVM increases with baud rate.




Figure 31 - Transmission Results of Low-Cost PM-QPSK and PM-16QAM Signals.

4.6.3. Standardization and Interoperability

Another important factor to consider is standardization and interoperability. Standardization is driven mainly by short-reach metro/aggregation applications, where optical performance is not a differentiator. Today, around seven DSP solutions from different companies are offered. The standardization will eventually lead to improved interoperability and predictable performance and allow operators to utilize the optical fiber infrastructure more efficiently to meet future bandwidth demand.

In 2016, The Optical Internetworking Forum (OIF) launched a new project related to coherent transmission technology: 400G ZR Interoperability. The OIF expects to develop an implementation agreement (IA) for 400G ZR and short-reach DWDM multi-vendor interoperability. The IA will support single-carrier 400G transmission using coherent detection and advanced DSP/FEC algorithms.

To make the coherent transceiver interoperable to each other, the following main parameters need to be considered:

1. Modulation format.
2. Framing format.
3. DSP algorithms.
4. Symbol mapping.
5. Optical properties.
6. Equalization parameters.

It is believed that the industry as a whole would benefit from a successful standardization of coherent transceiver including both optical performance and DSP functions and capabilities.

4.7. Experimental Results

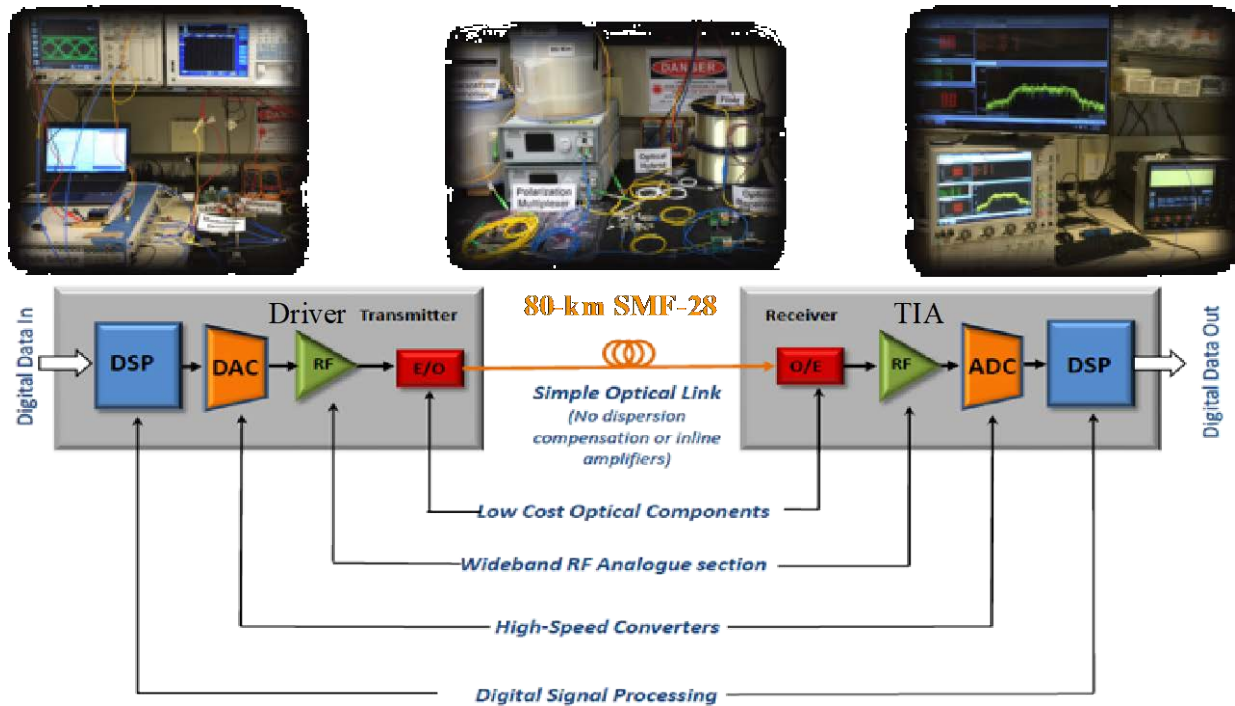


Figure 32 - End to End Coherent Optical Link in the Lab

To demonstrate the feasibility of coherent optics in access networks and verify methods and ideas that have practical potential, CableLabs has established a coherent optical transmission system in an optical laboratory. Figure 32 shows the end-to-end coherent optical link with in-house developed DSP algorithms at both the transmitter and receiver sides. The DAC function is achieved via Keysight M8196A arbitrary waveform generator with up to 92 GSa/s and 32-GHz analog bandwidth. The real-time sampling scope with up to 80 GSa/s and 33-GHz analog bandwidth performs the function of ADC. Different types of off-the-shelf electrical and optical components are acquired to emulate different implementation scenarios, such as low-bandwidth scheme or high-performance solution.

4.7.1. Single Wavelength

In the laboratory, we have achieved 256 Gbps over 80 km on a single wavelength with low-bandwidth optical components and simplified DSP algorithms. That is ~26 times the capacity of what can be achieved over an analog optical carrier fully loaded with 1.2 GHz worth of DOCSIS 3.1 signals. We have achieved that using a symbol rate of 32 GBaud, using 16-QAM modulation over 2 polarizations.

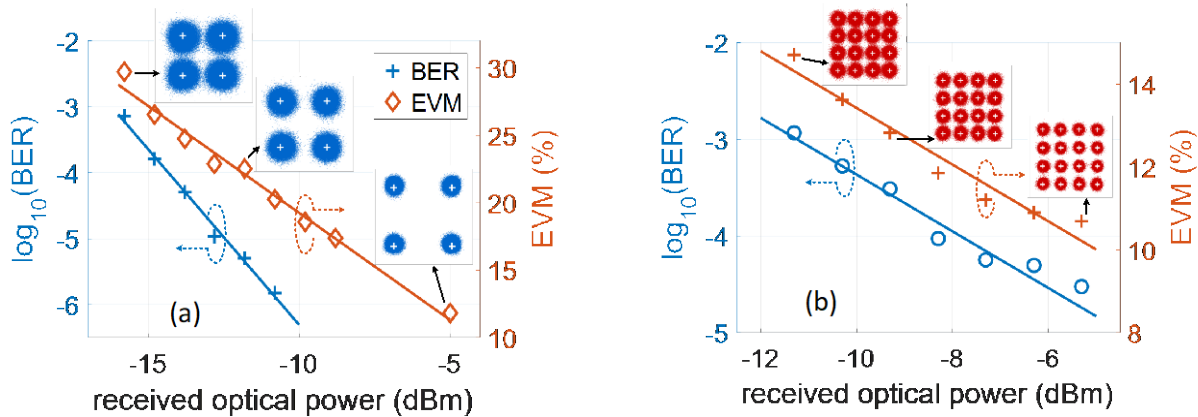


Figure 33 - BER and EVM Performance of Low-Bandwidth Coherent Scheme

BER and EVM performance of coherent 32-Gbaud PM-QPSK transmission are presented in Figure 33 (a), with constellations plotted in the insets. Due to the limited memory of our AWG, minimum measurable BER is 1×10^{-6} , and error free transmission is achieved for received optical power larger than -10 dBm. BER and EVM of 32-Gbaud PM-16-QAM transmission is presented in Figure 33 (b), and the minimum achievable BER is 3×10^{-5} . As shown in Figure 34, high baud rate and high-order modulation formats are also tested based on the high-performance setup.




Figure 34 - Constellations of High Baud Rate and High-Order Modulation Formats

4.7.2. DWDM System

In addition, we have multiplexed eight of these wavelengths to achieve 2048 Gbps. That is 50 times more than what can be achieved over 4 analog optical carriers each with 10 Gbps of DOCSIS 3.1 payload. Each wavelength carries 256 Gbps based on 32GBaud PM16-QAM signals using low-cost components. The

optical spectra of unmodulated CW and modulated coherent signals are shown in the following figure. These eight wavelengths only occupy a very small portion of C-band transmission window, which is shown in (b) of the following figure. In this system, the net SE is 4bits/s/Hz, which is 40 times better than DWDM direct detection scheme.




Figure 35 - Optical Spectra of Eight DWDM Coherent Signals

4.7.3. Coexistence Evaluations (Preliminary Results)

The ideal network for deploying such coherent systems would be a green field deployment on fibers, which is called coherent-only implementation. However, in practice there are many brown field installations, meaning many of these networks are deployed already and have several DOCSIS and or 10G OOK 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 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 in WDM systems.

To provide an option that enables network operators to effectively support 100G on their existing networks, CableLabs has done some preliminary experimental verifications to explore the performance challenges in such coexistence applications.

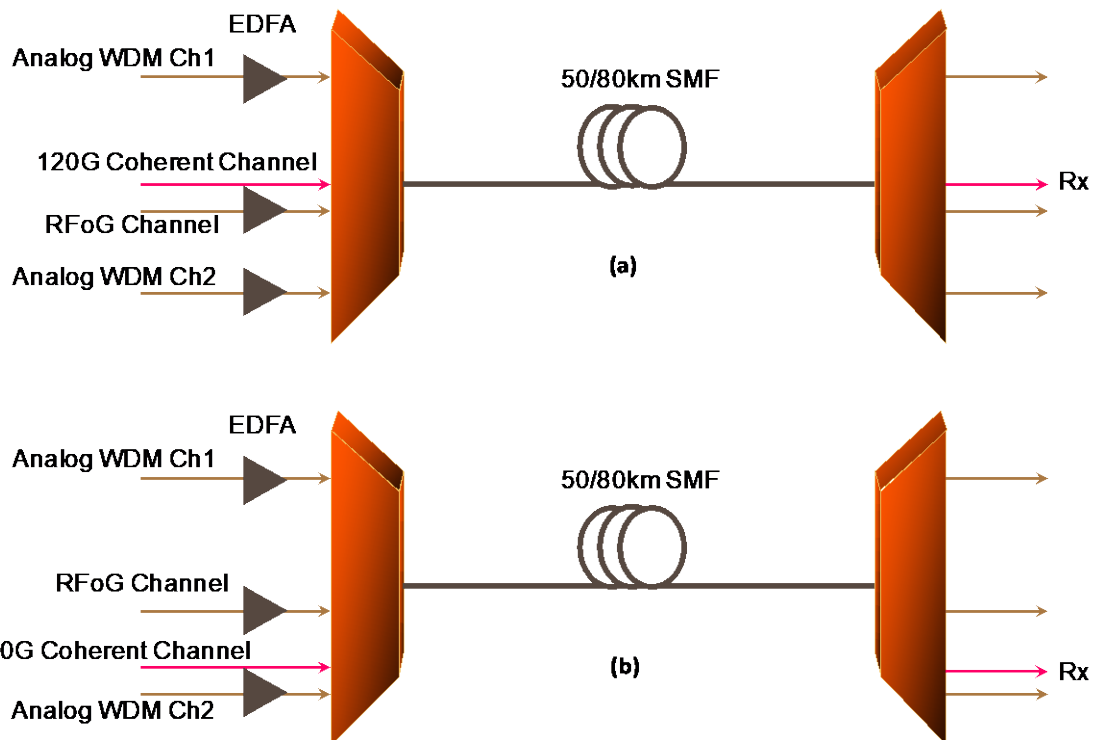


Figure 36 - Experimental Setup for Coexistence Evaluation, (a) Case I and (b) Case II

Because of the limited availability of the number of analog optical channels, two cases are selected for coexistence testing between coherent and analog channels as shown in Figure 36 (a) and (b) for case I and II, respectively. Our experimental work includes a single 120-Gb/s coherent PM-QPSK channel with three co-propagating analog DOCSIS3.1 channels through 50 or 80-km SMF. A CFP module is used to generate the PM-QPSK coherent signal, the two neighboring channels on the 100 GHz ITU-T grid CH28 and CH40 at each side of the test channel are generated by independently modulated commercial 1-GHz DOCSIS 3.1 transmitters. The other channel is based on radio-frequency over glass (RFoG) channel at CH33. The wavelength allocation and launched power of each channel are shown in Table 2. Around 15-dB power difference is set between coherent channel and analog channels.

Table 2 - Wavelength Allocation on ITU-Grid

Channel	ITU-T Grid	Carrier Frequency (THz)	Wavelength (nm)	Input Power (dBm)
Acacia 120G Coherent CFP	CH 34	193.400	1550.097	-1.67
Arris RFoG	CH 33	193.300	1550.967	13.78
Cisco Prisma II	CH28	192.800	1554.92	13.11
Cisco Prisma II	CH40	194.000	1545.309	12.75

Figure 37 shows the optical spectra of all signals before optical transmission, where the coherent wavelength is tuned to CH29.




Figure 37 - Optical Spectra before Fiber Transmission

In this testing condition, the transmission performance of coherent channel with and without neighboring analog channels are shown in Figure 38 (a) and (b). Negligible penalty is observed after 50 or 80-km fiber transmission both testing cases.




Figure 38 - Transmission Performance w/ & w/o Analog Channels

Towards the highly nonlinear regime, a penalty in BER can be observed with the increase launch power of analog channel power (CH 28). As shown in Figure 39, negligible impairment is observed if the analog power is less than 14 dBm in this testing scenario.




Figure 39 - Transmission Performance w/ Different Input Power of Neighboring Analog Channel

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 discussed use cases for near-term and long-term applications, including the deployment for aggregation points in distributed HFC architecture, remote PON systems and eventually coherent optics to the premises. Economics of coherent optics in the aggregation use case scenarios have been analyzed. This economic analysis show that 100 Gbps aggregation by 2020 will have similar link costs using either wavelength multiplexing direct detection or coherent optics, and the breakeven link cost point will occur earlier for a greater capacity demand. In addition, coherent optics systems are advantageous operationally as they consume lower power and require much lower number of ports at the hub than wavelength multiplexed direct-detection alternatives. This coherent optics cost improvement could accelerate further with a solution optimized for the Cable access network.

This paper introduces digital coherent optical system in detail, including advanced modulation formats, architecture of modulation and detection, and DSP flow for both transmitter and receiver. The paper highlights the motivation for coherent optics in access and potential approaches to re-design and re-engineer the digital coherent concept from long-haul and metro solutions to the access network,

leveraging reduction in complexity and cost of electrical and optical components as well as DSP ASIC. While coexistence between coherent signals and analog signals was simulated [2] in a previous paper, in this paper coexistence between these two systems was experimentally demonstrated showing 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.

Standardization and interoperability will play a key role for low-cost implementation for access. Proof-of-concept experimental results demonstrating multi-wavelength multi-terabit per second within an access environment are shown in the laboratory. This also demonstrates coherent optics technology long term scalability leveraging the fiber from hub to node already deployed.

Abbreviations

ADC	analog to digital converter
ASIC	application-specific integrated circuit
BER	bit error rate
BPSK	binary phase-shift keying
CD	chromatic dispersion
CMA	constant modulus algorithm
CMOS	complementary metal-oxide-semiconductor
CMTS	cable modem termination system
DAC	digital to analog converter
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
FEC	forward error correction
FWM	four-wave-mixing
Gbps	gigabit per second
GHz	gigahertz
HD	high definition
HFC	hybrid fiber-coax
HHP	household pass
ISBE	International Society of Broadband Experts
km	kilometer
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
OSNR	optical signal-to-noise ratio
PAM	pulse amplitude modulation
PBS	polarization beam splitter
PHY	physical layer
PM	polarization multiplexing
PMD	polarization mode dispersion
PON	passive optical network
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
SMF	single mode fiber
SNR	signal to noise ratio
SOP	state of polarization
SPM	self-phase modulation
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] Z. Jia, et al, “Performance comparison of dual-carrier 400G with 8/16/32-QAM modulation formats,” IEEE Photon. Technol. Lett., vol.27, no.13, pp. 1414-1417, July, 2015.

[4] Z. Jia, et al, “Performance comparison of spectrum-narrowing equalizations with maximum likelihood sequence estimation and soft-decision output,” Optics Express, vol.22, no. 5, pp. 6047-6059, March 2014.

[5] Z. Jia, et al, “Super-Nyquist shaping and processing technologies for high-spectral-efficiency optical systems,” Proc. SPIE 9009, 90090J, January 2014.

- [6] Z. Jia, et al, “Field transmission of 100G and beyond: multiple baud rates and mixed line rates using Nyquist-WDM technology”, IEEE/OSA J. Lightw. Technol., vol. 30, no.24, pp. 3793-3804, Dec. 2012.
- [7] John G. Proakis, “Digital Communications,” 4th edition McGraw-Hill, 2001.
- [8] I. Kaminow, T. Li, A. Willner, “Optical Fiber Telecommunications Volume VIB,” 6th edition Academic Press, 2013.
- [9] M. Birk, B. Mikkelsen, “100 Gb/s and Beyond Transmission Systems, Design and Design Trade-offs,” Short Course 203, OFC 2017.
- [10] C. Fludger, “Digital Signal Processing for Coherent Optical Systems,” Short Course 393, OFC 2017.
- [11] R. Howald, “Aboard the Technology Wave: Surf Conditions Report,” SCTE Cable-Tec Expo, 2016.
- [12] J. T. Chapman, “Infinite Ports, Infinite Channels, Infinite Bandwidth and Infinite Opportunity,” SCTE Cable-Tec Expo, 2016.
- [13] J. Finkelstein, A. Bernstein, “Remote PHY Deployment Scenarios,” SCTE Cable-Tec Expo, 2016.
- [14] J. D. Salinger, “DOCSIS 3.1 – Experiences from Early Deployments,” SCTE Cable-Tec Expo, 2016.
- [15] S. Druse, J. Miller, “DOCSIS 3.1 Leaves the Lab and Hits the Field with Midco,” SCTE Cable-Tec Expo, 2016.
- [16] Infonetics Research, “HFC Optical Nodes,” July 2015.
- [17] IHS TECHNOLOGY, “100G+ Coherent Optical Equipment Ports,” Dec. 2016.
- [18] IHS TECHNOLOGY, “Telecom Optics and Components,” April 2017.
- [19] Cable Television Laboratories, Inc. “DOCSIS 3.1 Physical Layer Specification - CM-SP-PHYv3.1-I09-1606021”, June 21, 2016.

Competitive Advantages Of HFC Networks for Wireless Convergence

A Technical Paper prepared for SCTE•ISBE by

John Chamberlain

Director, Office of the CTO
CommScope
1100 CommScope Place SE
Hickory, NC 28603
828-431-9800
John.Chamberlain@commscope.com

Mark Alrutz

Director, Field Applications Engineering
CommScope
6519 CommScope Road
Catawba, NC 28609
828-241-6492
malrutz@commscope.com

Introduction

Network convergence has been defined as “The efficient coexistence of telephone, video and data communication within a single network. The use of multiple communication modes on a single network offers convenience and flexibility that are not possible with separate infrastructures.”

By that definition the CATV (Community Antenna Television)/Broadband industry has been a “Converged Network” in and of itself since the deployment of high speed data DOCSIS (Data over Cable System Interface Specification) and VoIP (Voice over Internet Protocol) systems over ten years ago. Although only implied in the statement above, it is the integration of wireless services in a “converged network”, of which the industry has also embraced in the way of millions of single dwelling unit and mobility Wi-Fi hotspots. Now the industry is also embracing the concept of “cellular” mobility on the same network, either through support of 4G/LTE densification or the implementation of 5G.

In all the current definitions of a converged network one feature is rarely mentioned but is of paramount importance the power required to run the network components. In particular, far-edge Access Point (AP) devices require power to operate. This is where the HFC (Hybrid Fiber Coax) architecture deployed by the industry shines. Where other industries have shed the cost of network powering infrastructure for short term gains and reduced OPEX (operations expense), the MSO (Multiple System Operator) industry has maintained network powering to activate signal amplifiers for the coax portion of the HFC network. Power is required every 50-500 meters for the mass deployment of LTE densification and 5G mobility and only HFC networks have the network powering budget required to meet these needs.

A Tale of Two Networks

As an introduction to where the networks are now, a short history of these communication networks is in order.

The Bell System first used -48 volts DC and 90 volts AC to power voice circuitry and the ringer on a phone respectively. Network powering was the only choice as the Bell System required reliability beyond what the power grid could provide. The Bell System pioneered the term “five 9’s” meaning 99.999% reliability. -48 volts DC was chosen as a voltage high enough to allow for voltage drop in a distance defined “local loop” and still power the talk circuit terminal devices. The higher AC voltage was needed intermittently to ring the bell on the phone. In addition, at the time, DC was considered a safer voltage and could be powered with batteries. Negative 48 volts DC to ground was chosen to minimize electrolysis and galvanic corrosion to the conductors and connection points. The power requirements at the terminal equipment were low, on the order of tens of milliamps.

Evolution of DSL

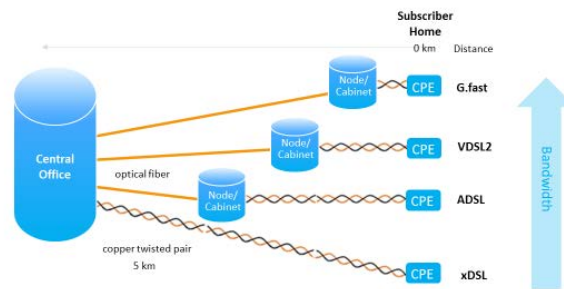


Figure 1 – Evolution of DSL

The twisted pair or “tip and ring” cabling deployed by the legacy telephony carriers has limited bandwidth. The original system was built for 4 KHz voice signals.

Companding, the process of compressing and subsequently decompressing signals, and digital compression and modulation schemes eventually allowed for higher delivered effective bandwidths. First, with telephone modems, and then with the advent of optical or electrical Digital Loop Carriers, and subsequently the shortening of the twisted pair portion of the plant, higher bandwidths and data rates could be achieved through ADSL, VDSL, and G.fast technologies.

As the Bell system matured, the legacy powering system remained in place until the mass deployment of fiber optics and Fiber to the Home (FTTH). FTTH was driven by the consumer need for bandwidth that the legacy twisted pair plant could not deliver. By definition, PON (Passive Optical Networks) are all passive and the only power required in the system is at the terminal unit or in PON vernacular, the ONT. (Optical Network Terminal) This CPE (Consumer Premise Equipment) is powered by the consumer. This created a high bandwidth system with no network powering. No network powering means less OPEX but no access to power in the plant other than through power utility company drops or consumer/enterprise power.

Fiber to the Home (FTTH) / PON

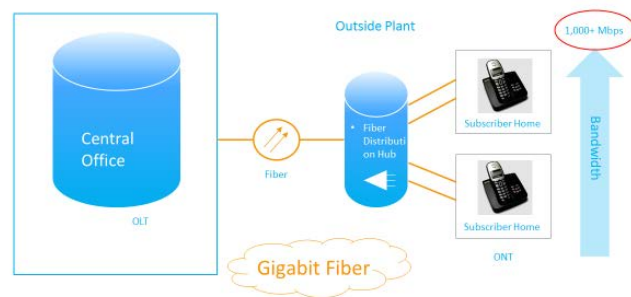


Figure 2 - FTTH PON Network

When legacy CATV systems were being built, 75-ohm coaxial cable was used to transmit AM (Amplitude Modulated) RF (Radio Frequency) signals from a “Head End” to the end users. The

attenuation of the coaxial cable at the RF frequencies in use, is such that the AM signal needed to be amplified approximately every 2500 feet. The RF amplifiers needed to be powered. The most economic and controlled way to do so was network powering. At the peak of coaxial cable deployment, before HFC, cascades of amplifiers as high as 40 were not unheard of. The power system chosen to power these amplifiers along the length of coaxial cable was 60 HZ, trapezoidal AC. Early networks were powered using 60-volt power supplies, but as power demand of network elements increased, standards and regulations were changed to allow 90-volt power supplies. This higher voltage has greatly improved network power utilization and effectively reduced voltage drop and current draw due to more efficient activation of constant power devices. Typically, the power is inserted periodically along the length of the coaxial cable and provides up to 15 amperes of service. There are also designs where power is expressed along a parallel coaxial path, in some cases dedicated for that purpose. This approach allows power supplies to be more conveniently located, and lower resistance cables can be used to minimize voltage drop.

Hybrid Fiber Coax (HFC)

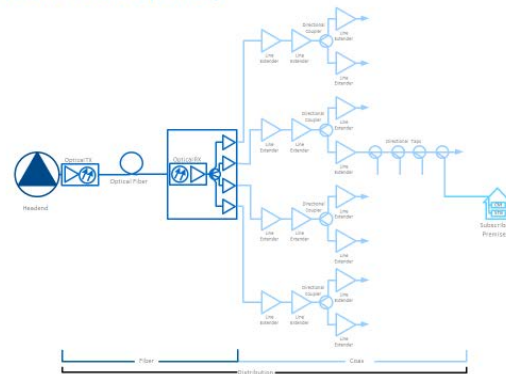


Figure 3 - HFC Network

In the case of the legacy CATV networks fiber was initially installed to reduce amplifier cascades which greatly reduced noise in the system and increased quality of service due to shorter amplifier cascades and the associated reduction in outages. The HFC plant has always had high bandwidth as a requirement for AM and later digital video delivery (1 GHz) and this bandwidth has been maintained or increased as fiber has been deployed deeper into the HFC network. In addition, in an HFC network, network powering has been maintained as the last part of the access/distribution network is still coaxial cable and remains amplified. The HFC architectures designated Node + 1,2,3, etc. refer to the maximum amplifier cascade in each design. This evolution has created a high bandwidth/data rate capable plant with access to power.

Node +0 or “fiber deep” network designs are built around the concept of eliminating amplifiers. In these designs, there is no cascade of amplifiers between the optical node and the subscriber. Power remains a requirement in these networks as well, since the nodes themselves are active devices which require power to operate, typically on the order of 1000 watts.

Network Convergence Small Cell Wireless in HFC

The ubiquitous need for bandwidth in our internet centric world has driven these two networks, as well as more recently deployed wireless networks, to something referred to as “Network Convergence”. In the end, all networks are moving toward delivering high speed digital data from a data or processing source, through a high bandwidth or “broadband” network, to a wireless distribution point and reverse for the upstream.

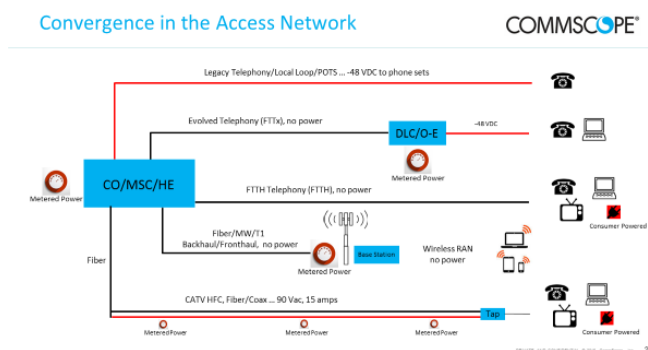


Figure 4 - Convergence in the Access Network

The next network architecture evolution is 4G/LTE densification and 5G wireless. This evolution is an evolution of fixed and wireless access points. Current LTE coverage in North America is relatively ubiquitous. Wireless “coverage” is complete, for all intents and purposes. The issue, as consumer demand for data increases, is capacity. Wireless capacity can be increased in several ways. Better modulation techniques, more spectrum, or spatially. 4G/LTE densification and potentially 5G mobility create more capacity spatially. More small cells closer to each other (250 meters) means that there are less users at each access point which effectively creates more bandwidth per square meter. This is the drive for 4G /LTE densification and potentially 5G. The promise of fixed wireless 5G in the millimeter wave band (i.e. 28GHz) creates more bandwidth with additional spectrum. For instance, in the 28 GHz band there is up to 800 MHz of spectrum to provide fixed broadband as a competitor to HSD from the MSOs.

In the MSO or legacy CATV network there is a headend, connected via router to data centers. The network consists of a high bandwidth HFC network. At the consumer/enterprise end of the network the obvious trend is toward Wi-Fi wireless connectivity in the home, or in the office. In the legacy telephony carrier network the same is true, with the exception of the delivery network described above. The high bandwidth HFC network is instead replaced by a legacy copper network, a version of FTTN (Fiber to the Node) supporting some form of advanced DSL services, or a FTTH network. As in the HFC example, home and business routers and Wi-Fi are typically activated. A traditional cellular network is comprised

of a network of macrocells, each independently powered and interconnected by a backhaul network of varying types, inclusive of fiber, HFC, copper and microwave. Options exist to activate Wi-Fi networks as an additional AP within a cellular wireless network, but these are less common due to the availability of 4G/LTE service with its associated mobility.

These three different networks, 1) legacy telephony carrier TP/FTTx, 2) Legacy CATV broadband HFC and 3) the mobility wireless networks all looked different when deployed due to the need to meet different communications applications, but are now all beginning to resemble each other to the point the networks can, and will merge into one, saving operator OPEX and CAPEX as equipment requirements also become identical This is network convergence. It not only drives the networks toward looking the same, but capex will be reduced as white-box manufactures and NPV/SDN (Network Virtualization / Software Defined Networks) become standard.




Figure 5 - Network Convergence

Requirements for Convergence


In a converged network three things are required, sometimes referred to as PBS:

- 1) Power (for wireless access points, as well as other edge devices)
- 2) Backhaul (data from the edge APs to the central data storage or processing centers)
- 3) Site acquisition


1. Power

The HFC network is well suited for Network Convergence. In fact, an argument could be made the HFC network IS already a converged network as, to date, the number of Wi-Fi hotspots has reached about 500,000.

Estimates are that 80% of HFC plant miles whether coax or fiber have network power availability due to the power on the coax. Coax in many cases runs parallel as a back-feed from an optical node thereby creating power availability even in fiber portions of the plant. In most cases the power availability is more than adequate for Wi-Fi hotspots or 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.






Figure 8 - Access Point Power Requirements

FTTH PON networks simply do not have this access to power, apart from the installation of a utility drop and meter installation. While a utility drop can be acquired at reasonable cost, and perhaps in a

reasonable timeframe, the installation of a utility drop at hundreds or thousands of AP locations will quickly become both an economic burden as well as a project bottleneck.

2. Backhaul

In an existing HFC network, backhaul capacity can be provided in several ways. DOCSIS provides connectivity over either coaxial or fiber links, and current deployments of DOCSIS 3.1 enable interconnect speeds in the Gbps range. Additionally, many HFC installations provide direct fiber connectivity over previously installed dark fiber, or over channels provided on existing fiber using available WDM (wavelength division multiplexing) paths. Equipment to condition power from the HFC plant to activate wireless APs via PoE (power over Ethernet) while providing a DOCSIS channel also exist. For small cells, depending on the technology deployed, tight timing and latency requirements are a concern, but these issues are currently being addressed in standards organizations.

3. Site Availability

Going forward, the majority of wireless mobility to be deployed for LTE densification or 5G will be deployed in urban and suburban areas. Much of the HFC network is aerial in these geographies and lends itself to site acquisition. The 4g/LTE wireless densification effort is about capacity, not coverage and 20-30 feet in the air is adequate for the coverage of a small cell area, estimated at 50-500 meters inter-site distance to create the additional capacity desired. For underground plant, small cabinet based cells are

available, with varying tower heights and a variety of concealment methods. These small calls can be very unobtrusive, and placed along right of way already established for pedestals and cabinets.




Figure 9 - Concealed Access Point in Underground Plant

Conclusion

The HFC network, as deployed, with Power Availability, Backhaul Availability, and Site Availability is very well suited to 4G/LTE densification or 5G convergence. That competitive advantage should not be lost or forfeited as fiber extends deeper into the network or even FTTH architectures are deployed. In order to maintain the competitive advantage that HFC offers to a converged network, coaxial plant should be maintained as power distribution plant. This can be ensured even with the emergence of Node+0 architectures. In situations where coaxial plant does not exist or is impractical, hybrid cable constructions with power conductors placed within or alongside fiber cables can be utilized to retain access to network power.

Abbreviations

AC	alternating current
ADSL	asymmetric digital subscriber line
AM	amplitude modulated
AP	Access Point
CAPEX	capital expenditures
CATV	community antenna television
CPE	consumer premise equipment
DC	direct current
DOCSIS	data over cable system interface specification
DSL	digital subscriber line
RF	radio frequency
FTTH	fiber to the home
FTTN	fiber to the node
GHz	gigahertz
HFC	hybrid fiber coax
HSD	high speed digital
Hz	hertz
LTE	long term evolution
MSO	multi system operator
NPV/SDN	network virtualization/software defined networks
ONT	optical network terminal
OPEX	operating expenses
PoE	power over Ethernet
PON	passive optical networks
TP	twisted pair
VDSL	very-high-bit-rate digital subscriber line
VoIP	voice over internet protocol
WDM	wavelength division multiplexing
xG	x generation

Bibliography & References

Munson, Ben, “Comcast, TWC ‘CableWifi’ shared network deal survives recent M&A hotspot count now 500M.” *Fierce Cable*. Questex, Jul 8, 2017. Web. Jul 17, 2017.

Making More with Less! A Case Study in Converging Wireline and Wireless Network Infrastructures Using Distributed Access Architectures

A Technical Paper prepared for SCTE•ISBE by

Hugo Amaral Ramos
Chief Technologist CALA
ARRIS
hugo.amos@arris.com

John Ulm
Engineer Fellow, CTO – Network Solutions
ARRIS
john.ulm@arris.com

Zoran Maricevic, Ph.D.,
Engineer Fellow, Access Technologies
ARRIS
zoran.maricevic@arris.com

Jose Tavares
Engineer
ARRIS
jose.tavares@arris.com

Claudio Albano
Engineer
ARRIS
claudio.albano@arris.com

Introduction

The maturing of the telecommunications industry has led to a consolidation trend. This consolidated telecom companies often have two completely separate infrastructures to maintain: wireless and wireline. With this, both a challenge and an opportunity emerge. In order to maintain competitiveness and lower both CAPEX and OPEX, the operator must converge infrastructures and associated functions.

At the same time, the evolution of technology is huge. New technologies such as DOCSIS Remote PHY (RPHY), Remote MAC-PHY (RMACPHY), RF over Glass (RfG), and Fiber to the Home (FTTH) will help enable this convergence.

This gives operators the possibility of selecting among an immense quantity of options. However, this can also generate lots of doubts, such as how to make the right decision regarding a future proof infrastructure, while at the same time having the most cost effective, high quality delivery and access network.

A wrong technology decision could be catastrophic for anyone in an industry that is highly competitive. Operators in the Caribbean and Latin America (CALA) region are even more hard pressed since the ARPU and the restriction in CAPEX is often a big burden. The CALA operators may have only one chance to get it right:

“CALA has normally one bullet to shoot the target, aiming right is really crucial.” Author unknown

With the challenges that CALA faces, we analyze a regionally specific case in the CALA market in which the requirements for the planning and deployment of new technologies, such as RPHY, RMACPHY, RfG, and FTTH are somewhat specific to CALA. Requirements specific to typical CALA deployment areas are very high density and lower bandwidth services plans compared to the market in the USA. Consequently, a larger number of subs per service group (SG) are required to make the implementation economically viable.

This case study analyzes some actual underserved cities in the Caribbean and Latin America region, with the goal to profitably enable the delivery of high speed data and other services in a sustainable way by helping the operators to get the best synergy from the wireless and wireline infrastructures.

The paper offers an analysis of various network technology options to serve dense urban areas with High Speed Data and other services while leveraging already deployed mobile infrastructure assets. This is accomplished using new Distributed Access Architectures (DAA) technologies such as Remote PHY.

These options all use the existing mobile backhaul infrastructure, IP Radio Networks (IPRAN) and Node Base locations. They vary in the type of access network that is deployed (e.g. FTTH, RfG, N+0, N+X) and where the Remote DAA elements are deployed (e.g. node or shelf in cabinet). The case study shows a comparison of all options, highlighting a high level normalized total cost of ownership (TCO), technical requirements, benefits, limitations, concerns, considerations, and future proof analysis.

Study of a Convergence Between Wireline and Wireless Infrastructure Using New Technologies

1. Requirements of CALA and Some Difference vs NA Region

Some different challenges and requirements between the Caribbean and Latin America region and North America exist, such as the environment, the economic challenges, and also types of required services. Understanding these differences is important in order to decide which technologies to deploy. In the following paragraphs we can find some of these differences and see how these differences can make an impact.

One important consideration for an adoption of a new technology is the environment in which this technology will be used. A very-dense or ultra-dense environment is typical in the CALA region in the deployment of telecommunication networks today. As noted in Figure 1, some cities in CALA have ultra-dense concentration such as Sao Paulo in Brazil with a density of 7,913.29inh/km² (20,495.3inh/sq. mi) and a huge population of approximately 22 Million persons. Similarly Mexico City is 6,000inh/km² (16,000inh/sq. mi) without as much distribution and spread like the North America (NA) region. These differences show a gap of requirements for deploying new technologies such as DAA.

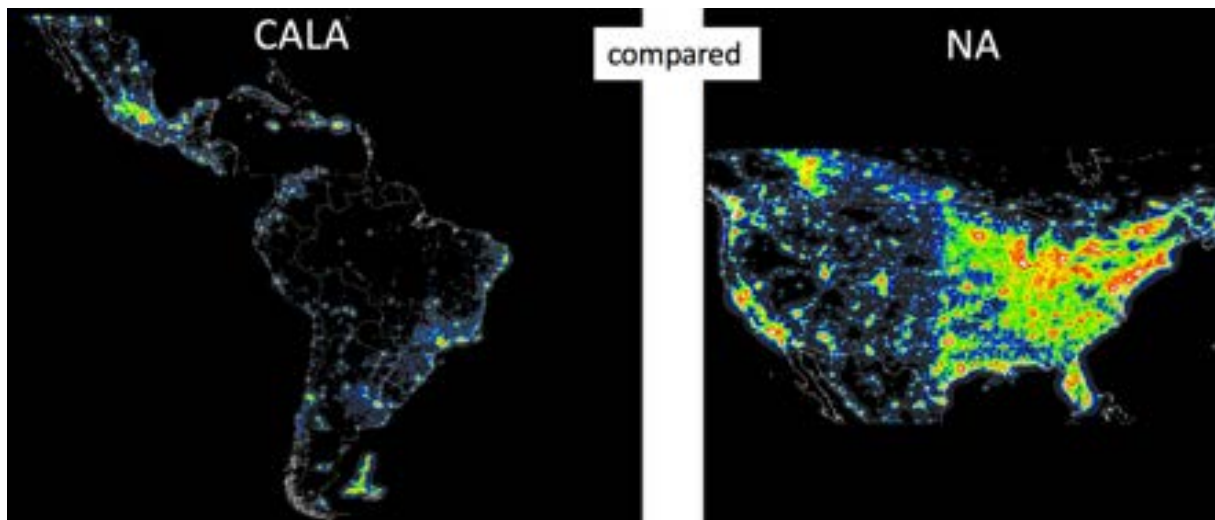


Figure 1 – Light View in Dark Comparing the Concentration between CALA vs NA Region

Another important environmental observation of the region is the type of construction employed. In CALA, the typical implementation is via aerial construction, using poles. One ugly observation identified in the region is the robbery of units and cables in the streets. This sometimes can eventually put some constraining forces against the idea of distributed architectures (DAA), since the fear to put more expensive equipment in the field is significant and justified.



Figure 2 – Photo of a Stolen Optical Unit, Optical and Copper Cables in a Yarn String

Another critical consideration is economical; specifically, how the industry in the region makes revenue, profit, investments, and how the plant is operated and maintained. CAPEX restriction is always strong in the region, mainly due to the FX rate and limited average revenue per unit (ARPU) compared to the North American market. Sometimes the CALA ARPU can be one fifth compared to U. S. market. This is why the CALA region cannot make the wrong selection of standards or technology, nor experiment with hype technologies that are nontraditional initiatives, because there is often no margin for error.

There are a few operators in the CALA region that have a unique position in the market where they have a massive deployment of both wireline and wireless networks and also a dominant market share. These operators are starting to question how convergence opportunities can be deployed, and are beginning to generate initiatives in this direction.

The economic situation in CALA also limits the types of services required in the region. While it varies from country to country, essentially the maximum speed offered in the market is 250 Mbps today. The 1 Gbps speeds that are deployed in U. S. are maybe two years delayed and probably will be deployed for only a small part of only high tier customers.

On the other hand, the demand for data is still projecting very strong growth with the compound annual growth rate (CAGR) in downstream ranging from 45 to 50% Year over Year (YoY). This shows that the requirement for higher data speeds is important to this region. Any solutions considered must be able to grow in a sustainable way, converging infrastructure and gaining synergies.

2. Challenge and Use Case of CALA

Service providers are asking: “What technology should we deploy in greenfield scenarios?” One particular concern in CALA is how to deploy a cost-effective and future proof solution to address the challenges already presented. Should they invest in FTTH and make a revolutionary step? If they do not make this revolutionary step, will they be able to have a future-proof solution? Or will they have a major cost over-run problem in the near/medium term?

Our convergence case study focuses on an under-served city in the region as a typical environment to find a solution to answer these questions. Thinking pragmatically, yet also with a different and more holistic view, how can we solve the environmental and economic constraints while being more effective to maintain sustainability to the service provider?

For this use case, a small under-served city has a density of approximately 6,000inh/km² (16,000inh/sq. mi). While there is currently no wireline broadband infrastructure, there is a 3G mobile infrastructure in place. Our challenge is to use this typical environment and find a solution to deploy video and data services. We focused our analysis on a 25,000 HPs area, which can be replicated and used as a possible “template blueprint” solution over the entire region.

The first option for consideration was the “business as usual” (BAU) HFC implementation. The necessity of deploying new buildings and infrastructure for headend or Hubs was a big challenge. Therefore, it was very hard to make a profitable business case using this type of solution for these cities. The business challenge becomes a civil engineering problem due to time to deploy new buildings, cooling systems installations, and all important considerations that a new equipment rooms require. The implementation cost is also prohibitive to sustain any offering in these small and medium cities. The ongoing costs required to maintain these infrastructures further exaggerate the difficulty of establishing a viable business case.

The next option considered is based on using new Distributed Access Architecture (DAA) technologies. This looked much more promising.

The service offerings considered for the case were up to 100 Mbps, 30% subscriber penetration and 400 kbps average broadband usage in the busiest hour in the peak of consumption. Digital TV was also included, consisting of 100 SD H.264 programs and 72 HD H.264 programs, for the analysis of the growth and capacity models.

With an important consideration that most of these cities have a mobile infrastructure in place and some of these operators already own these infrastructures, our idea is to use and converge these infrastructures, gaining important time in implementation, significant synergy in maintenance, and making this implementation as cost-effective as desired. The typical architecture used in the region to deploy HFC and mobile infrastructures is analyzed in the next section.

3. Technical Implementations of Converged Cable and Mobile Networks

3.1. Typical Fixed Access Infrastructure in CALA

Typically, CALA cable operators use a traditional HFC business-as-usual implementation using 6MHz ITU-T.J83 annex B for video and data, DOCSIS standards and Analog Modulation (AM) based optical nodes in their networks. The average number of subs per service group (SG) in the region is around 1,000, which is much higher than what is typically seen in North America.

It is worth emphasizing the fact that the CALA deployments are typically in a very-dense urban area. Usually the distances between the headend/hub and the node are less than 10 km (6.25 miles) for 99% of the cases.




Figure 3 – Typical Cable Access Infrastructure Implementation

One important consideration to the network evolution towards distributed architectures is that the requirements of the region are different from the big MSOs in North America. As shown in figure 4, the CALA operators need to be cautious about the relative maturity of these technologies before massively deploying anything new.




Figure 4 – Example of Planning and Implementation Differences With RPHY Technology Requirements Between CALA and NA

3.1. Mobile Typical Access Infrastructure in CALA

Figure 5 shows a typical mobile access architecture as utilized in the region. It is comprised of two parts: the IP radio access network (IP-RAN); and the radio access network (RAN). The IP-RAN provides a unified layer of services to deliver important applications, such as synchronization using IEEE1588, QoS, Security and Monitoring to the radio access network (RAN). The second part, the RAN, is the ring of NodeB, eNB, or radio base station (RBS).

These architectures are planned and deployed in a ring topology that provides high availability for the services delivered. That availability is what makes this topology the most recommended and widely adopted IP-RAN/RAN topology in the region.




Figure 5 – Typical Mobile Access Infrastructure

4. Options Analyzed to Converge the Infrastructures with Important Consideration and Synergies

4.1. Options

Understanding how the architectures are deployed, our case study then evaluated four options to converge and use the already in place mobile infrastructure. All of these options use DAA to enable the service offering of fixed High Speed Data (HSD) broadband and Video services.

The major difference between these four options is characterized using two variables: the location of the R-PHY; and the number of RF amplifiers. This is shown as a 2-by-2 matrix in the table below. The location of the R-PHY module could be at the “eNodeB location” or at the “segmentation node”. The number of RF amplifiers is partitioned into either an N+0 passive HFC plant or an N+X active HFC plant with a small number of amplifiers in cascade (e.g. N+1 to N+3).

Table 1 – Options in Two Important Dimensions

	RPD in Field	RPD @ eNodeB Location
N+X	#1	#3
N+0	#2	#4

A high-level architecture diagram of the four options is shown below in figure 6. It illustrates where the Remote PHY device (RPD), switches, and access network equipment are installed.

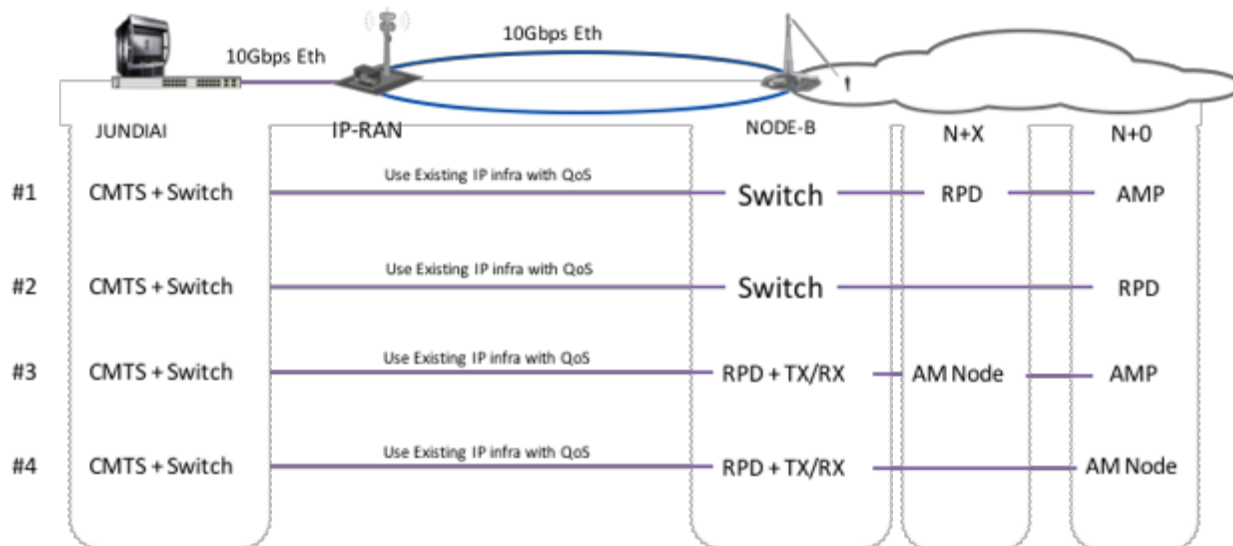


Figure 6 – DAA Technological Options Identified

4.1.1. Option #1 – RPD in the Field with N+X Coaxial Architecture

- A remote CMTS Core contains the DOCSIS MAC function and controls the DOCSIS PHY in the RPD located in a node in the field
- All infrastructure used to get from the remote CMTS Core to the eNB location is Ethernet and is a shared connection with the mobile infrastructure
- From an Ethernet port in the eNB location, an optical SFP with 10km reach is used to connect to the RPD in the field
- From the RPD there is an active N+X coaxial network architecture (e.g. N+1 to N+3)

4.1.2. Option #2 – RPD in the Field with Fiber Deep N+0

- A remote CMTS Core contains the DOCSIS MAC function and controls the DOCSIS PHY in the RPD located in a node in the field
- All infrastructure used to get from the remote CMTS Core to the eNB location is Ethernet and is a shared connection with the mobile infrastructure
- From an Ethernet port in the eNB location, an optical SFP with 10km reach is used to connect to the RPD in the field
- From the RPD there is a passive N+0 coaxial network architecture

4.1.1. Option #3 – RPD in the eNB Location with BaU HFC N+X architecture

- A remote CMTS Core contains the DOCSIS MAC function and controls the DOCSIS PHY in the RPD located in the eNB location
- All infrastructure used to get from the remote CMTS Core to the eNB location is Ethernet and is a shared connection with the mobile infrastructure.
- From the eNB location a BaU HFC TX and RX is used with traditional AM optical nodes in an N+X coaxial network architecture

4.1.2. Option #4 – RPD in the eNB Location with BaU HFC N+0 Fiber Deep architecture

- A remote CMTS Core contains the DOCSIS MAC function and controls the DOCSIS PHY in the RPD located in the eNB location
- All infrastructure used to get from the remote CMTS Core to the eNB location is Ethernet and is a shared connection with the mobile infrastructure
- From the eNB location a BaU HFC TX and RX is used with traditional AM optical nodes in an N+0 coaxial network architecture

4.2. Synergy in Timing and IP Switching Considerations in an RPHY solution

With the evolution of the HFC network to a DAA (Distributed Access Architecture), the IP switching is reviewed with the focus on the future evolution to network function virtualization type of infrastructures. The traditional 3-layer architecture of IP routing and switching is being questioned, changing from the core, aggregation and access topology to Spine Leaf architecture. This new architecture is being used in the data center environment and is viewed by the industry as a future-proof implementation.




Figure 7 – IP Switching Changing from 3-layers Architecture to Spine Leaf Architecture

The Spine Leaf architecture is the IP switching layer between a MAC core that will be instantiated at Headends, hubs, remote locations, and the Remote PHY device, and that will be installed in the field.

These new elements need to be considered during the traffic engineering phase and must support all the traffic required by the end user including signaling information. Also, MSOs will need to think about the traffic demand growth. The selected IP switch design should be able to easily address this traffic demand growth. In addition to traffic needs, MSOs will need to consider redundant topology and security. More and more, IP network engineering and HFC network engineering will be working collaboratively to design new HFC network.

Some specific features will be required and must be considered in the IP switching design. The IEEE1588 timing protocol needs to be supported in the switches. This implies the addition of an IEEE1588 timing server. One interesting aspect to be considered in the convergence is that mobile networks already have this server deployed today to provide timing to the eNodeB. To utilize the server as is in place is a great opportunity of operational synergy. IPv6 support will be required as well.




Figure 8 – Usage of Existing IP Switching to the Leaf Switch and Provide Sync to RPD

4.3. Important Synergy Gains due to the Unification of Two Infrastructure Layers

To meet our cost objectives, synergy gains are exactly what we are looking for in these solutions. One area of synergy is unifying two infrastructure layers together. This is very important. It is critical to converge to a single layer since this analysis is for greenfield areas with new service offerings.

Taking a look in the planning of nodes to provide fixed broadband and video to this city and the deployed mobile infrastructure makes a compelling case to unify the infrastructure as in the next figures 9, 10 and 11:

- Figure 9 is a supposed infrastructure deployment thinking only about the HFC fixed in mind
- Figure 10 uses an example of an existing mobile infrastructure and
- Figure 11 provides the example of the opportunity to unify the infrastructures

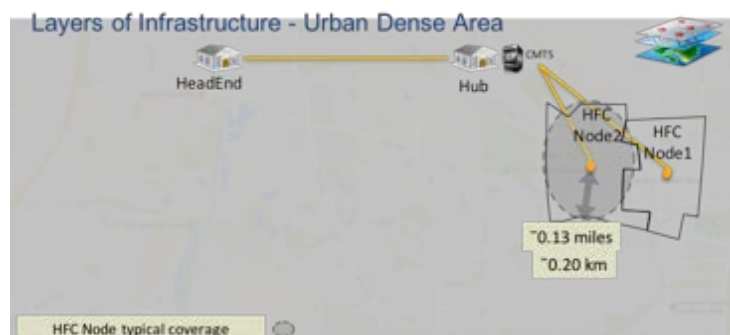


Figure 9 – Example of an HFC Nodes Implementation with a Fixed Services Layer Only

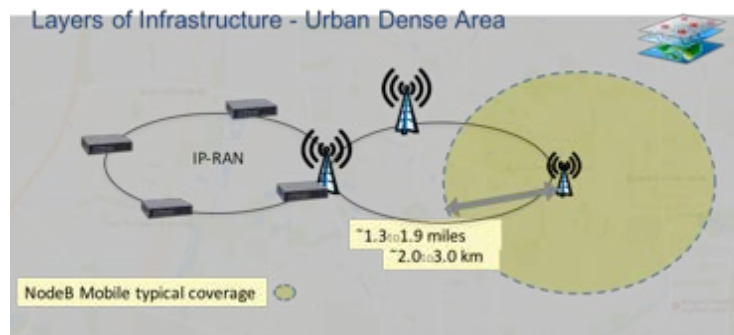


Figure 10 – Example of an Existing Mobile infrastructure

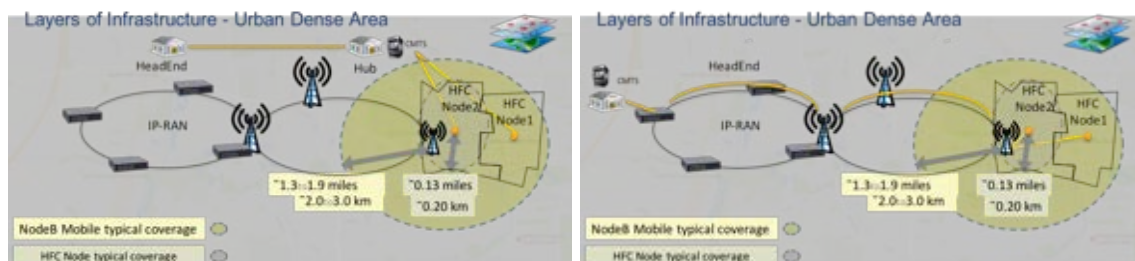


Figure 11 – Opportunity to Unify the Layers

There is a huge opportunity of synergy by unifying the layers of infrastructure. This allows us to:

- Simplify the deployment
- Make a more cost-effective implementation
- Help lower the OPEX and simplify operation
- Help to make the evolutionary step more dynamic

4.4. Network Capacity Planning

Network designs are more cost sensitive in the CALA regions compared to typical North American operator networks. One of the key factors that drives the cost of the access network is the size of the Service Group (SG). The larger the SG, the more an operator can amortize the shared costs. The size of the HFC SG is in turn is determined by the services offered and the available HFC spectrum. This section explores some of the options available for this case study.

4.4.1. Services Offered

It is important to recognize that these areas do not have typical broadband services today. They may currently be limited to just 3G data services. Any broadband services being provided by the new network will be a giant leap forward, even if those broadband service data rates are significantly below other broadband services around the globe.

The majority of users are expected to have a downstream (DS) High Speed Data (HSD) service that is in the 10-25 Mbps range. In some regions, 25 Mbps service has become the minimum acceptable capacity for broadband services. The corresponding Upstream (US) HSD service might be in the 2-5 Mbps range. These data rates are sufficient for the subscribers to have a true broadband experience and stream High

Definition (HD) content into their home. Our network capacity analysis assumes that at least 75% of the subscribers will have a 25 Mbps/5 Mbps (DS/US) service.

An important cost factor is the CPE equipment required. For 25 Mbps/5M bps service, a 16x4 DOCSIS 3.0 modem would be the most cost effective while also providing excellent future growth capabilities. These modems would allow an operator to also offer higher HSD service tiers for premium revenue over time. The 8x4 modems are just slightly less expensive but have one half the capacity, which may limit future growth. For our analysis, we will assume 20% of the subs have 50 Mbps/10 Mbps service and 5% of the subs take the 100 Mbps/20 Mbps service tier.

The above service tiers are the minimum HSD service tiers that we would recommend that the CALA operators support. With newer DOCSIS 3.1 technology, it may also be possible to offer up to 1 Gbps service tiers provided there is sufficient spectrum available. This will allow the operator to offer additional services to businesses, elite residential customers, and/or may be used for wireless backhaul such as 4G/5G &/or WiFi.

As a minimum, the CALA operator will offer an HSD service with sufficient capacity to enable their subscribers to access Over-the-Top (OTT) video services. The operator may also decide to offer their own managed digital video service. They may already have the video infrastructure and set-top boxes (STB) in place for a legacy video service. This kind of offering might support roughly 100 unique video programs and might also support Video on Demand (VOD) services. Since this uses a Distributed Access Architecture over a shared Wireless/Wireline regional network, there is no plan to support any analog video services over the converged infrastructure.

Some progressive operators may decide to offer their video services using IP Video distribution rather than legacy HFC video. This would also allow the video service to be offered over the wireless network and other access networks such as PON. IP Video increases the capacity requirements for DOCSIS. 8x4 modems may not have sufficient capacity for IP Video. This provides another reason for using a 16x4 modem to simultaneously support IP Video.

4.4.2. HFC Spectrum Utilization

Often in brownfield scenarios, an operator is limited by the available HFC spectrum, such as 550 MHz or 750 MHz. However, for this CALA Convergence case study, a greenfield HFC system is being built. This means that the coaxial portion of the HFC will be designed from the beginning with proper components and spacing to optimally support 1002 MHz to 1218 MHz.

The basic HSD service tiers will use DOCSIS 3.0 bonding. Since we are trying to maximize the number of subscribers per SG, the operator should start with 32 bonded 3.0 channels. This consumes 192 MHz of spectrum.

Since there are no older deployed STB in a greenfield, the digital video service is assumed to be delivered using H.264/MPEG-4 encoding technology. This reduces the spectrum in half from older MPEG-2 only STB. A reasonable 100 SD/HD program digital video service with VOD could be offered in 21 QAM channels, or 126 MHz of spectrum.

The HSD + digital video services only consume 318 MHz of spectrum out of a possible 1218 MHz. This means that there is plenty of spectrum for future expansion and potentially other services.

It turns out that digital video STB and DOCSIS 3.0 modems only support up to 1002 MHz. With an HFC designed for 1218 MHz, a DOCSIS 3.1 OFDM downstream channel could be put above 1002 MHz without conflicting with 3.0 HSD or digital video services.

Another interesting decision is choosing the best upstream split. DOCSIS 3.0 only supports up to 85 MHz upstream, while DOCSIS 3.1 can optionally support a 204 MHz upstream that enables a 1 Gbps US service tier. Since the system is not constrained by spectrum, we recommend the 204 MHz US split with the DS starting at 258 MHz. This would allow the operator to also offer a 1G symmetric service over the HFC utilizing existing DOCSIS 3.1 technology.




Figure 12 – Example HFC Spectrum Utilization

In this example, the digital video and DOCSIS 3.0 DS channels have been put in the 684 MHz to 1002 MHz range to compile with a possible future migration to DOCSIS Full-Duplex (FDX). A DOCSIS 3.1 OFDM channel is placed above 1002 MHz to offer 1G DS services on day one.

In the upstream, eight DOCSIS 3.0 ATDMA channels are supported with the rest of the 12-204 MHz upstream being used by a pair of DOCSIS 3.1 OFDMA channels. This also enables 1G US services on day one.

Note that the spectrum from 258 MHz to 684 MHz is not initially used, yet is available for future DOCSIS expansion. This spectrum could be used in several different ways. First, some, or all, of the excess spectrum could be used for DOCSIS 3.0 expansion as user traffic continues to grow. This could help eliminate or defer Service Group splits in the future, saving the operator costs down the road. Alternatively, some of this excess spectrum could be used for additional DOCSIS 3.1 OFDM channels to offer higher service tiers (e.g. 2.5 Gbps DS service).

Finally, some, or all, of this excess spectrum could be used by DOCSIS FDX to offer multi-Gbps symmetric services. Note that use of FDX would also need to coincide with a migration to an N+0 passive plant in the future. This corresponds to Options #2 and #4 above.

4.4.3. Capacity Modeling

For this paper, the network capacity requirements were simulated using the ARRIS Network Capacity modeling tool. To minimize costs, a downstream SG with 1,000 subscribers was paired with two upstream SGs each with 500 subs. It is not desirable to make the US SG any larger to limit the amount of noise funneling in the upstream. If this had been an older brownfield, then the US SG might have had to be even smaller.

Since broadband usage in CALA has been running a couple years behind typical North American usage, the average subscriber usage during peak busy hour (Tavg) is assumed to be 400 Kbps per sub. The higher HSD tiers would have more usage than this and the basic tier would be slightly less than this. Based on other ARRIS research, Tavg is assumed to grow at 40% Compounded Annual Growth Rate (CAGR).

For the modeling a legacy digital video service is considered, as it would use more spectrum than an IP Video service. The digital video VOD service assumes that there would be a peak usage of 5% during peak busy hours.




Figure 13 – Network Capacity Model Results

These results only consider the DOCSIS 3.0 subscribers. By the year 2021, the 32 bonded 3.0 channels have been completely consumed and additional 3.0 capacity is added from the “DOCSIS Futures” spectrum. By the year 2024, this spectrum eventually becomes filled as well.

Initially the model assumes that the SG is split and is now 500 subs per DS, 250 subs per US. By 2026, additional action is needed again. Either another SG split is required, or alternately DOCSIS 3.1 capacity could be used if enough of the subscribers have been migrated to D3.1 modems.

Perhaps the most important conclusion from this section is that it is feasible to have a 1000 sub DS SG that is viable for the next 8-10 years before the access network needs any segmentation. This allows the operator to install the most cost effective access network now, but with segmentation in mind for the next decade. The operator should design their HFC greenfield with the ability to segment easily in the future without requiring a significant fiber or plant investment.

4.5. Trade-offs in Selecting the N+0 vs N+X Implementations

Figure 14 shows topology of a common HFC network [source: HFC wiki]. A regional optical transport ring, in the upper left corner, performs a function of redundant-routes connecting the distribution Hubs, a function similar to that of IP-RAN described above. Fiber links are also the means of connection from distribution Hubs to the optical nodes, with the coaxial cable coming out of the node, through a cascade of RF amplifiers, in order to either extend the reach or to overcome RF splitting losses. The “0” and “X” in N+0 and N+X denote the length of the RF amplifier cascade. For the network of Figure 14, the X = 4, since there are 4 RF amplifiers in a cascade emanating out of the top-most and bottom-most optical nodes.

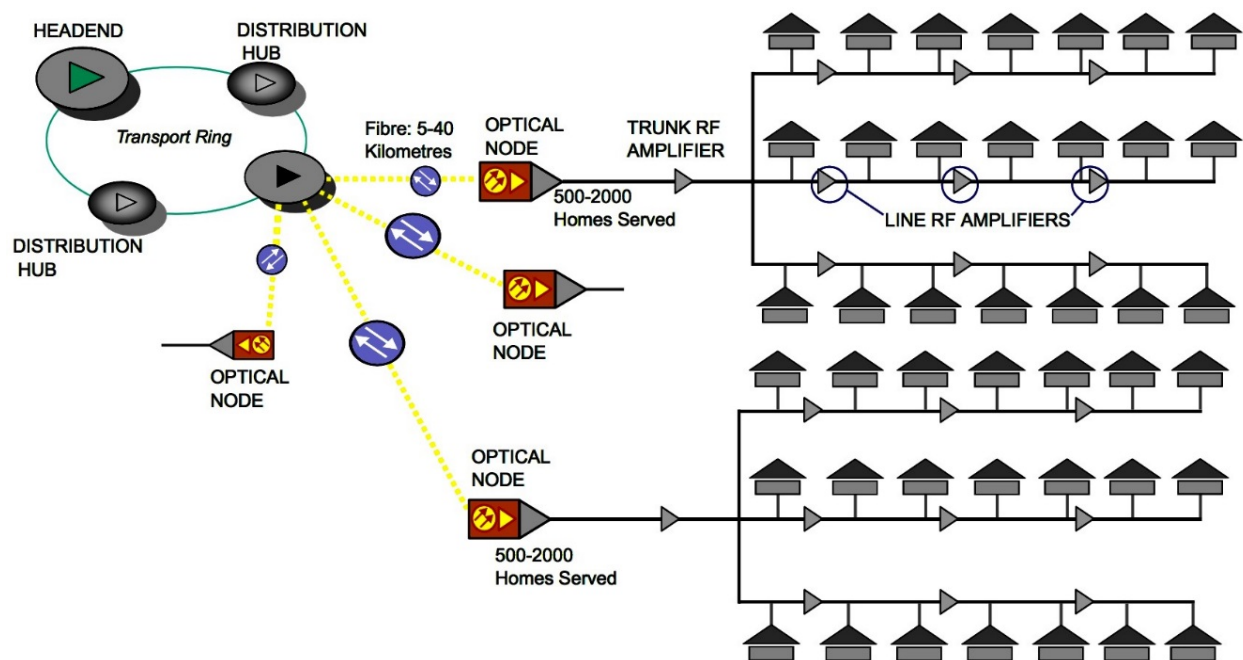


Figure 14 – Topology Overview of Hybrid Fiber Coax (HFC) Network Architecture

True insight into flexibility of HFC can be gleaned from a “pyramid depiction” of the same topology, as shown in Figure 15. Optical nodes typically have 4 coaxial outputs and are followed by no RF amps, as in N+0 case, or with up to ~30 amps, in N+X case. Each RF amp may feed 4-8 taps, and each tap may have 4-8 drop ports. At one extreme, 1 node, x 4 RF outputs, x 4 taps, x 4 drops, results in 64 home-passed coverage, which could be set as a very small service group, with all the capacity to be shared among those

64 homes that have signed up for the services. At the other extreme, 1 node, x ~30 RF amps, x 8 taps, x 8 drops, results in as many as 1,920 homes-passed!

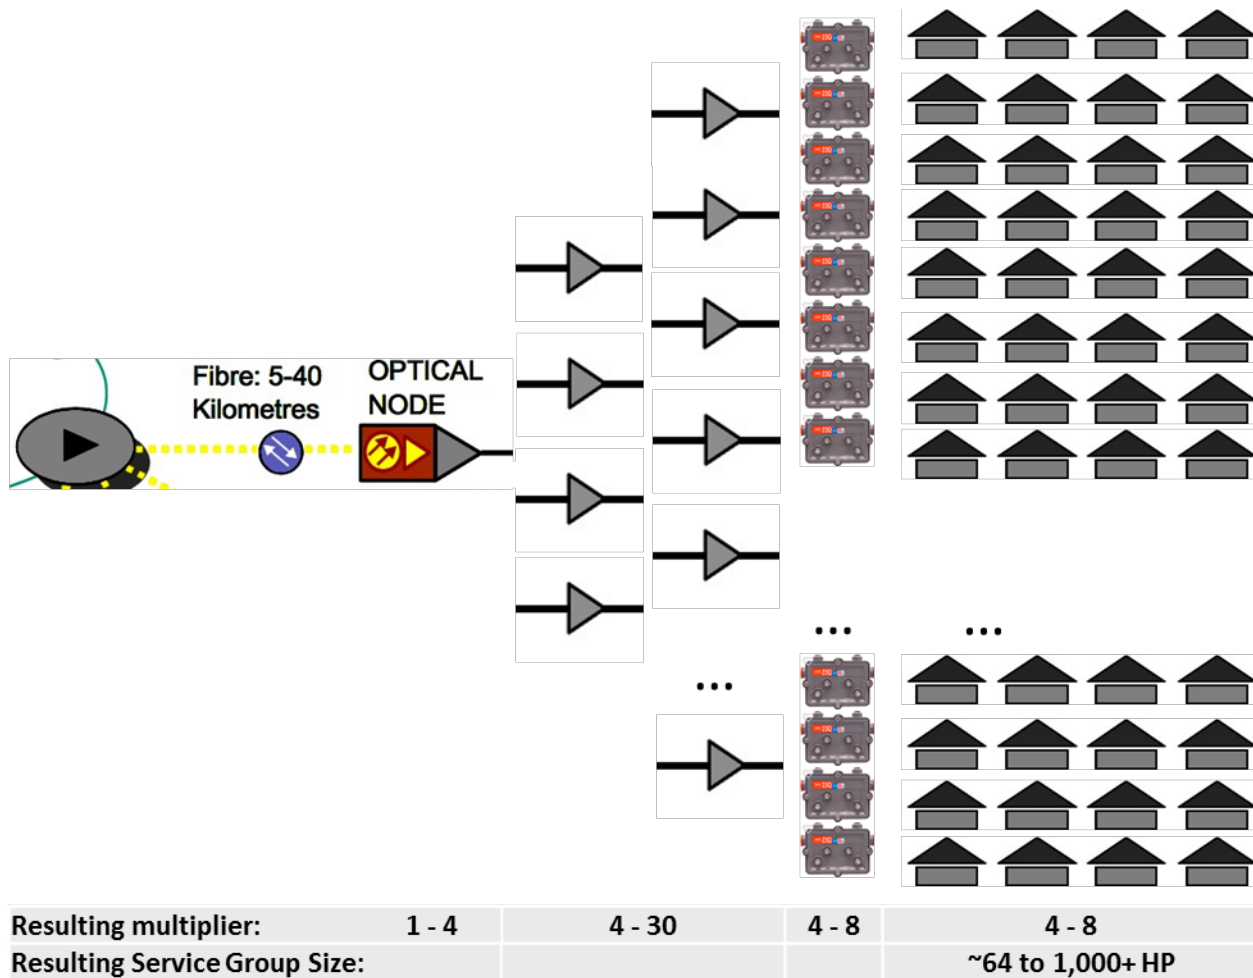


Figure 15 – "Pyramid Depiction" of a Common HFC Network Topology

In figure 15, one fiber link serves 1-4 node segments, followed by 4-30 RF amps, each serving 4-8 taps, each with 4-8 ports, and giving “dynamic range” for the resulting service group of 64 to 1,000+ homes passed.

Distance covered, capacity delivered, and cost per home passed are definitely affected and traded against each other. Nevertheless, the “dynamic range” of this topology is simply powerful. Perhaps the most impactful consequence of this flexibility is network’s unmatched ability to scale as needed, and more importantly for the operator to “pay as you grow” just as the network capacity augmentation is needed.

Those “skilled in the art” of architecting the cable access plant will remember that cable plant was all coax before it became “Hybrid Fiber Coax”. As fiber became a great distance-coverage enabler first, and then cost / power / performance advantageous later, the “N+X” cascade of RF amplifiers following the fiber/coax node kept reducing from as many as 15 down to as few as zero! Furthermore, Deep Fiber

architecture, long advocated by ARRIS for low “total cost of ownership” [Fiber Deep 5 Years Later], establishes a “clean slate” starting point for establishing cost-effective, high-capacity, low power flexible network. Yet the most valuable attribute of N+0 may be its future-proof promise, is that each node can be turned into an RPD / DAA, with further segmentation, if so required, or full-duplex functionality, or even turned into all-fiber network, of either PON or RFoG type.

DOCSIS FDX enables multi-Gbps symmetric services over HFC. This is very powerful. It also requires a passive N+0 HFC plant. In our case study, some options assume that the operator starts with the most cost effective N+X HFC implementation, where x is a small number (e.g. N+3). The FDX services may be well targeted and not required across the entire footprint. If this is the case, then the operator may want to surgically upgrade a particular active component using Fiber to the Last Active (FTTLA) with FDX capable technology while leaving the remainder of the plant alone. An operator may choose to pull extra dark fiber at the same time the coax is built; or simply use conduit that would easily allow the fiber to be pulled to the last active at a later point in time when needed.

4.6. Future Proof Considerations about the Options Identified

We started this paper with a “one bullet” analogy, and how crucial it is for CALA operators to “aim right”. To continue this analogy, why not use a silver bullet as well? Since CALA operators need to get it right first time, they should also consider the ability to offer symmetrical gigabit broadband on day one. Deploying costly Fiber to the Premise (FTTP) network on day one is not economically feasible. However, symmetrical gigabit is achievable by using HFC and choosing the proper downstream / upstream RF frequency split, 204/256 MHz as outlined above, as well as leveraging everything the DOCSIS 3.1 will offer.

By constructing this new Fiber Deep HFC network, CALA operators will be unencumbered by aging plant and historical RF splits. Some of the North-American operators, whose “last mile” cable plant was built long ago, with the RF frequency split already implemented and with a topology that was optimized then, are thinking along the same lines, even though it’s going to be much more work for them to modify the network. Most of them will have to upgrade every active in order to support a new RF splitting ratio. [Stoneback-Slowik].

Thus, the proper RF frequency split selection enables CALA operators to offer and advertise the gigabit symmetrical service availability today. Later, as more and more consumers decide to spring up to gigabit symmetrical service offerings, there will be no need to touch the physical plant topology. The operator only needs to add CMTS / RPD capacity on the headend side and DOCSIS 3.1 modems on the CPE side.

As shown in the capacity plan model, the network will require re-planning to fit the capacity growth required over the years. As any operator manager knows, it is desired for these changes to be as dynamic and as streamlined operationally as possible. The options shown give operators the ability to selectively change the RPD from the eNodeB location to the field in a N+X configuration or change from a N+X configuration to a N+0 design selectively – in a pay as you grow model.

Changing from the eNodeB to the field is already streamlined, since the infrastructure of the optical node is already installed and the change from the optical node to a RPD is simple. Evolving from the RPD installed in the optical node location (N+X) to a N+0 location could also be simplified if the design of the coaxial and amplifier is made using four output amplifiers, as shown in Figure 16 below.

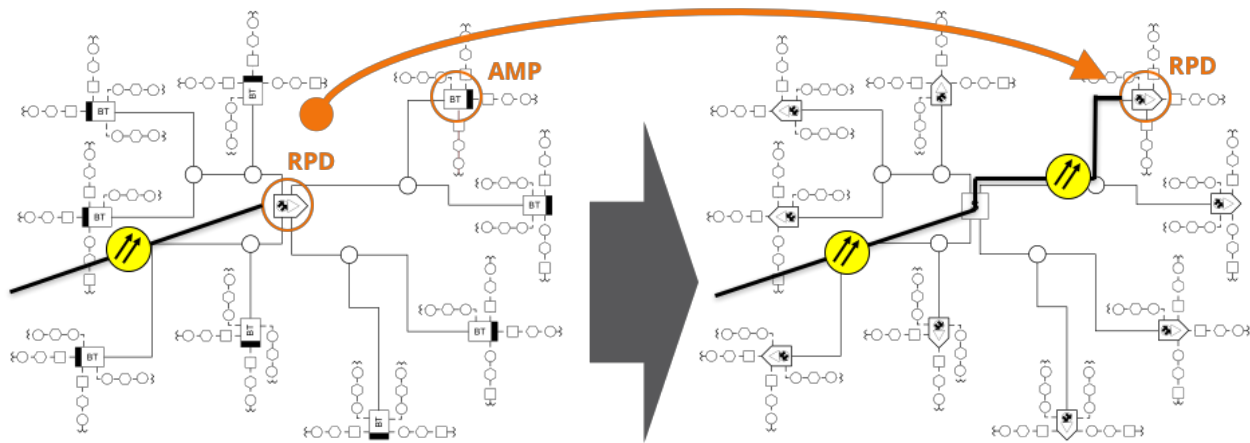


Figure 16 – Usage of a 4-Output Design to Simplify Evolution

Since this will be greenfield construction, it will also be prudent for the operator to consider other future expansions that might be coming within the next five to ten years. Some things to consider include: DOCSIS FDX, FTTH, and wireless backhaul (e.g. 4G/5G, Wi-Fi).

Similarly, an operator may want to provide FTTH services to select users and businesses. As they build out their N+X or N+0 HFC network, they should make sure that it can easily transition to FTTH as needed. As above, this can be done by laying additional dark fibers during the initial construction or by providing conduit that allows fiber to be easily deployed later when needed.

Looking into the future, the need for bandwidth is also growing in the mobile side where the types of service offerings and powerful devices are driving requirements never demanded nor delivered previously. Since mobile devices are powerful handheld computers, new applications with unimagined features to be launched and camera interfaces with more and more pixels will create a need to grow the network even faster, to support capacity, and to pursue higher quality of service.

Today's architecture of mobile macro nodes will not be sufficient for this growth. The need to construct new eNodeB locations will be required and will again generate big constraints. The construction for new antennas placements requires both the premises to support such placements and also power and backhaul links. This takes both time and money to deploy. One of the features that HFC always had is a great way to grow selectively by node splitting, something that cannot be as easily done in the wireless networks of today.

In the future as the RMACPHY evolves as a solution, based on discussions with the industry and standard bodies and becomes more standardized, this could be considered as a trustable option of convergence and can be a good topic to be included in further discussions. Another interesting aspect of the discussion is that 5G is being strongly considered to use mmWave and implementation of small and micro cells will be required. Conquering of the demarcation point (where these small cells will be installed with the place, power, and backhaul) will be a battle for the service operators and will be also a huge factor in providing successful services.

One interesting approach to conquering these demarcation points is to use deploy HFC for the wireless node split. Here the demarcation point uses spaces inside the strand mount equipment (Optical nodes, RPDs and AMPs) which already installed in the aerial deployments where placements can be made in the

street level, power and also backhaul. Also, the usage of street furniture will be an important option where the installation is underground.

The convergence presented here will help to address this opportunity for the future evolution of the mobile network. The usage of the HFC is already unified in terms of infrastructure. The mobile sites will help to connect to new radios to be installed inside the strand mount equipment installed in the HFC network. The usage of the RPHY solution will have an advantage due to the option to include these new radios in the future in this equipment. The technology requirements for the RPHY node to enable this “wireless demarcation point” possibility in the forthcoming products are being strongly considered, and feature prominently in today’s plans.

4.7. Cost, Power and Capacity Comparison

The four options were compared based on relative cost, power and capacity and results shown in Figure 17. The lowest in each category was set to 1.0 and the others are shown proportionately. As a refresher, the four options are shown again in the table below.

Table 2 – Options Remembered

	RPD in Field	RPD @ eNodeB Location
N+X	#1	#3
N+0	#2	#4




Figure 17 – Relative Cost, Power & Capacity for 4 Options

Option #3 has the lowest total overall costs as it uses N+X HFC architecture with a minimal number of RPD. Options #1 and #4 are very close in costs, coming within ~15% of the Option #3 costs. From an energy consumption perspective, Option #4 is the greenest solution. The other solutions require between 25% to 60% more energy. Finally, from a potential capacity perspective, the N+0 HFC architecture for Options #2 and #4 provides more than five times the potential capacity than the N+X HFC architecture of Options #1 and #3. Looking across these combined factors, Options #4 is intriguing as it reduces power consumption and provides the potential capacity of N+0 for only a slight increase in costs over Option #3.

In addition to the above, there are other benefits to the N+0 options. These push the fiber even deeper, so they align better with a wireless micro-cell strategy. This will become more important over time with the introduction of 4G/5G services as well as widespread WiFi hotspot services. The N+0 options also leave the door open to a future DOCSIS FDX and/or Extended Spectrum migration that can support multi-Gbps symmetrical services. This migration can be done as needed selectively on a node by node basis. It does not require the cost of upgrading the entire HFC architecture.

Conclusion

Convergence has been a word overly used in our industry. However, it is the name of the game now that the industry is becoming mature and strongly competitive. In this paper, we attempted to highlight options available to the CALA service providers that are pragmatic solutions ready to deploy today and we took a peek into the next decade, underlining important concerns to take in consideration to make the right decision now to be future proof.

The evolution of DAA, using digital Ethernet communication, provided us the opportunity to demonstrate how unifying the mobile and fixed infrastructure is possible. And the synergy generated by this is huge, not only in terms of CAPEX, but also OPEX and the simplicity of maintaining the network.

Starting with RPD in the eNodeB location and evolving to the field in a N+X and then to N+0 selectively by service groups seems likely to be a logical move. This way of growing networks (node splitting) has proven to be the success factor to the service provider that deployed HFC. The technology and the designs proposed here will help the operators in the region to continue to grow their network capacity selectively.

New technologies being developed today like DOCSIS FDX will require changes in the MAC layer and the PHY layer and deploying R-PHY solution today in a controlled environment such as a eNodeB location could help service providers to streamline operationally the upgrade, when this technology becomes available in the near future.

As the RMACPHY evolves as a solution, based on discussions with the industry and standard bodies and becomes more standardized, this will be considered as a trustable option of convergence and can be a good topic to be included in further discussions.

Finally, mobile services will continue to become more and more important over time. As the operator builds its N+X HFC, it should do so with an eye towards wireless migration and utilizing the HFC access network for 5G & Wi-Fi backhaul. The location of HFC actives and the availability of the HFC power should take into account the needs of both 5G and Wi-Fi wireless distribution.

Abbreviations

3G	Third generation wireless
4G	Fourth generation wireless
5G	Fifth generation wireless
ARPU	average revenue per unit
AMP	Amplifiers
BAU	Business as Usual
bps	bits per second
CAGR	compound Annual Growth Rate
CALA	Caribbean and Latin America
CMTS	Cable modem termination system
DAA	Distributed Access Architecture
eNodeB, eNB	Evolved node B
FX Rate	Foreign exchange rate - A value of two currencies relative to each other
FTTH	Fiber to the home
HFC	Hybrid Fiber-Coax
Hz	Hertz
IP-RAN	Internet Protocol – Radio Access Network
ISBE	International Society of Broadband Experts
N+0	Optical node + zero number of amplifiers in the coaxial
N+X	Optical node + X number of amplifiers in the coaxial
QoS	Quality of service
RAN	Radio Access Network
RPD	Remote Phy Device
RPHY	Remote Physical Interface
RMACPHY	Remote media access control and physical interface
SCTE	Society of Cable Telecommunications Engineers
SG	Service Group
YoY	Year over Year

Bibliography & References

- [CLOONAN_2014] “*Predictions on the Evolution of Access Networks to the Year 2030 & Beyond*”; T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, S. Chari, The Cable Show NCTA/SCTE Technical Sessions Spring 2014
- [HFC wiki] https://en.wikipedia.org/wiki/Hybrid_fibre-coaxial
- [Fiber-Deep 5 Years Later] “*Case Study: Fiber-Deep, 5 Years Later*”; Jeff Gould, Suddenlink and Joseph P. Lanza, Aurora Networks, Communications Technology; October 15, 2008
- [Thinking Green] “Thinking Green Strengthens the Case for Fiber Deep in Cable”; ARRIS/ Aurora White Paper 16, July 2009
- [Stoneback-Slowik] “*Making Rational HFC Upstream Migration Decisions in the Midst of Chaos*”; Dean Stoneback and Fred Slowik, 2013 NCTA/SCTE Spring Technical Forum
- [HFC Green] “*Giving HFC a Green Thumb*”; John Ulm and Zoran Maricevic, SCTE Journal of Energy Management, Volume 1, Number 2, October 2016

The Big Network Changes Coming with 1+ Gbps Service Environments of the Future

A Technical Paper prepared for SCTE•ISBE by

Tom Cloonan

CTO - Network Solutions
ARRIS
2400 Ogden Ave, Suite 180
Lisle, IL 60532
630-281-3050
tom.cloonan@arris.com

Tushar Mathur

Sr. System Engineer
ARRIS
tushar.mathur@arris.com

Ruth Cloonan

Data Analytics Expert
Blue Opus
ruth.cloonan@gmail.com

Ben Widrevitz

Staff System Engineer
ARRIS
ben.widrevitz@arris.com

John Ulm

Engineering Fellow
ARRIS
john.ulm@arris.com

Introduction

Whether operating on HFC or PON or 5G infrastructures, future operator networks will undoubtedly be required to deliver service bandwidths in excess of 1 Gbps. Since these higher-SLA services are not like past services, many network attributes and network operational procedures must be changed to accommodate the new bandwidth levels. This paper will explore many of these changes.

Traffic Engineering and 1+ Gbps Services

Traffic engineering is an important area that will definitely be impacted by the arrival of Gbps services, because many of the traditional rules of thumb and formulae that have worked well for years may no longer be valid.

Consider Figures 1, 2, 3, and 4, which illustrate current consumption (average) and billboard (peak) bandwidth trends and extrapolations into the future for several anonymous, sampled Multiple System Operators (MSOs). The following observations can be made:

- Within Figures 1 & 2, Average Bandwidth growth is still rising exponentially in both the Downstream (36% CAGR) and Upstream (17% CAGR), but the growth rate for the Downstream seems to have slowed in the last couple of years
- In addition, these Average Bandwidth growth rates are much lower than the Billboard bandwidth growth rates (defined by the Nielson Law) shown in Figures 3 & 4
- Even if the Billboard bandwidth growth rates drop to lower rates (as predicted by several operators and illustrated with the green line of Figure 3), it is still expected that the Billboard bandwidths will grow at a much higher rate than the Average Bandwidths

Thus, the distance between maximum bandwidth levels (T_{max}) and Average Bandwidth levels (T_{avg}) will continue to increase, creating an interesting scenario in which T_{max} for a single subscriber will be much higher than T_{avg} for that subscriber.




Figure 1 - Downstream Average Bandwidth Trends


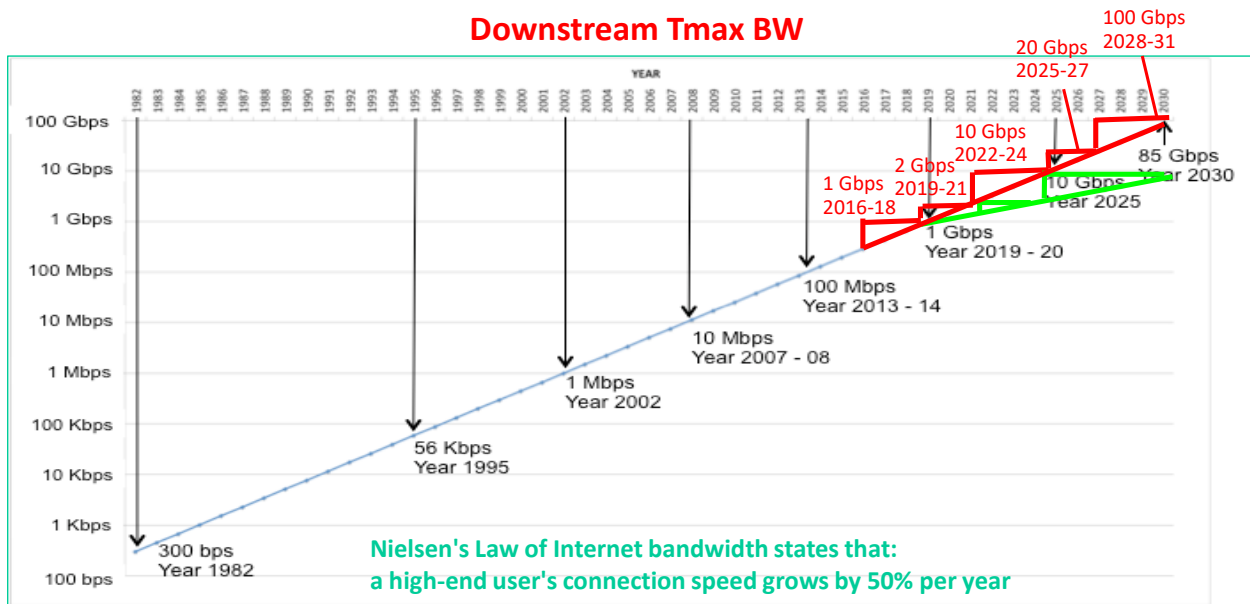


Figure 2 - Upstream Average Bandwidth Trends



Source: <http://www.nngroup.com/articles/law-of-bandwidth/>

Figure 3 - Downstream Billboard Bandwidth Trends




Figure 4 - Downstream & Upstream Billboard Bandwidth Trends

With much higher T_{max} levels (exceeding 1 Gbps) and with T_{avg} levels that are relatively lower (relative to T_{max}), it can be shown that the typical subscriber's transient bursts to maximum bandwidth levels are likely to occur for shorter windows of time and are also likely to occur much less frequently. To illustrate this point, let us consider the following contrived example. Assume that there exists a (somewhat strange) subscriber who only receives downstream traffic at one of two discrete traffic rates. That subscriber either receives bandwidth at a rate of 1 Mbps or 1 Gbps, and they never receive bandwidth at any other traffic rate. Assume also that we know that the subscriber has an Average Bandwidth given by $T_{avg} = 2$ Mbps.

Armed with this simple information about the subscriber, it is interesting to note that we can calculate the probability (the fraction of time) that the subscriber transmits at 1 Mbps and the probability (the fraction of time) that the subscriber transmits at 1 Gbps. This results from the fact that we have two equations in two unknowns if we draw the subscriber's transmission rates on a probability density function (pdf) diagram, as shown in Figure 5. For a pdf, we require that $\int pdf(x)dx = 1$. For a pdf with an average of $T_{max} = 2$ Mbps, we also require that $\int x pdf(x) dx = T_{avg}$. We can solve these two equations to determine that the probability $P(1 \text{ Mbps transmissions}) = 0.999$. We can also solve these two equations to determine that the probability $P(1 \text{ Gbps transmissions}) = 0.001$. Thus, with low T_{avg} values and high T_{max} values, it is clear that the probability of high bandwidth transmissions will be required to be quite small.



Figure 5 - Probability Density Function of Bandwidth of a Subscriber with Low T_{avg} and High T_{max}

Thus, it should be clear that when subscribers do burst to their T_{max} levels, the extremely high bandwidth levels are likely to perform most data transfers in a fairly short timeframe, so they will be on and off the network quite quickly. For example, a typical web page sized at 4 Mbytes = 32 Mbps would download at a 20 Mbps rate within 1.6 seconds, however that same 32 Mbit-sized web page would download at a 1 Gbps rate within 32 milliseconds.

As a result of this effect, it is also likely that the probability of having multiple maximum bandwidth bursts occurring simultaneously for many different subscribers becomes quite low as T_{max} values rise and the probability $P(T_{max})$ values drop. For example, if a T_{max} burst event for subscriber #1 has a low probability P_1 , and if a T_{max} burst event for subscriber #2 has a low probability P_2 , then assuming the two burst events are independent, the probability of the two T_{max} burst events occurring simultaneously has the even lower probability of $P_1 \times P_2$.

Consider the probability density functions shown in Figure 6 for two of our (contrived) example subscribers and for the aggregated bandwidth generated by those two subscribers when they share bandwidth capacity within a Service Group (SG). The aggregated bandwidth can be 2 Mbps (when both

are receiving at 1 Mbps), 1001 Mbps (when one is receiving at 1 Mbps and the other is receiving at 1000 Mbps), or 2000 Mbps (when both are receiving at 1000 Mbps). The probabilities for each of these probabilistic events is also shown within the aggregated bandwidth probability density function, and it can be seen that the probabilities $P(2000 \text{ Mbps}) \ll P(1001 \text{ Mbps}) \ll P(2 \text{ Mbps})$.




Figure 6 - Probability Density Function of Bandwidth for two Example Subscribers & their Aggregate Bandwidth

Carrying this idea even further, if we have N_{sub} subscribers within a Service Group and if each subscriber ' i ' exhibits an Average Bandwidth level of T_{avg} and a maximum bandwidth level of T_{max} with a burst probability given by P_i , then the Average Bandwidth within the Service Group will be given by $N_{sub} \times T_{avg}$. However, the probability of ever seeing a Service Group bandwidth approaching $N_{sub} \times T_{max}$ becomes very small, especially as N_{sub} is increased (since that probability is given by $P_1 \times P_2 \times \dots \times P_{N_{sub}}$). This fact can also be characterized by calculating the Service Group's Average Bandwidth, Standard Deviation, and Coefficient of Variation ($CV = \text{Standard Deviation} / \text{Average Bandwidth}$) for Service Groups of different sizes using data collected from the field today (where $T_{max_max} = 100 \text{ Mbps}$). For illustration, we also show the bandwidth associated with () + (3 ×). These results are shown in Table 1.

Table 1 - Average Bandwidth, Standard Deviation, & Coefficient of Variation vs. Service Group Size

Nsub for SG	Tavg for SG (Mbps)	Std. Dev for SG (Mbps)	Coef. of Var. for SG	(Tavg for SG)+(3*Std Dev for SG) (Mbps)
256	284.16	64.47	0.226879223	477.57
1024	1150.98	128.93	0.112017974	1537.766
4096	4460.54	258.00	0.057840479	5234.544

It can be seen that the Coefficient of Variation tends to be reduced for larger Service Group sizes, implying that the relative spread of bandwidths within the Service Group becomes smaller (relative to the average Service Group bandwidth). In particular, it appears that the Tavg value for the Service Group traffic grows linearly with Nsub, whereas the Standard Deviation for the Service Group traffic grows proportionally to the square root of Nsub (approximately). This is an important point, because it implies that much of the equipment within the headend (ex: CCAP Cores, Routers, Switches, etc.) that supports larger numbers of subscribers will be able to operate quite well even if they only support slightly more bandwidth capacity than $Nsub \times Tavg$. And, they will not typically be required to support anything close to a bandwidth capacity of $Nsub \times Tmax$.

Traffic Engineering studies are currently under way to create better models to determine the actual required bandwidth capacities, and those results will be presented in the future. Due to the complex nature of the new 1+ Gbps services, these studies will likely need to be performed using a “bottom’s up” approach that accounts for the contributions and statistics for both Average Bandwidth levels and burst bandwidth levels for each of the subscribers within the network. Early results of this work led to the realization that the “ARRIS QoE-based Traffic Engineering Formula” developed in [CLO1] is still quite valid for small service group sizes with $Nsub \leq 400$ - even in the era of 1+ Gbps services. That formula is given by:

$$Required\ Bandwidth\ Capacity \geq Nsub \times Tavg + 1.2 \times Tmax_max \quad (1)$$

where Tmax_max is the largest of the Tmax values. Larger Service Group sizes (with $Nsub > 400$) will require new formulae that are being researched. However, using Equation (1), we can find the required bandwidth capacity and DOCSIS 3.0 SC-QAM channel counts DOCSIS 3.1 OFDMA channel counts that typically-sized Service Groups of the future with typical bandwidths might require. It can be seen in Table 2 that DOCSIS 3.1 OFDMA blocks will definitely help provide the required bandwidth capacity, because the use of lower-efficiency SC-QAMs would consume a large portion of the Hybrid Fiber Coaxial (HFC) network spectrum.

Table 2 - Required Bandwidth Capacity, SC-QAMs, & OFDM Blocks for Future Service Group Sizes

Nsub for SG	Tavg for sub (Mbps)	Tmax_max for sub (Mbps)	Required BW for SG (Mbps)	Required # SC-QAM Ch's @ 36 Mbps ea.	Required # OFDM Blocks @ 8 bps/Hz
50	10	1000	1700	47	1.1
100	10	1000	2200	61	1.4
200	10	1000	3200	89	2.1
400	10	1000	5200	144	3.4

Utilization Levels and 1+ Gbps Services

The Required Bandwidth Capacity formula described in Equation (1) is a useful tool for calculating the approximate bandwidth levels required within a Service Group. A related equation can be derived if one considers the general meaning of the terms contained within Equation (1). In particular, the first term of Equation (1) is the Average Bandwidth passing into the Service Group. The second term represents the amount of headroom bandwidth required to support temporary bandwidth bursts in excess of the Average Bandwidth that may occur within the Service Group. The sum of the two terms is the total Required Bandwidth Capacity.

As a result, it should be clear that one can obtain the Average Utilization Level within a Service Group by dividing the first term by the sum of the first and second terms. In equation form, we find that:

$$\text{Average Utilization Level} = (N_{\text{sub}} \times T_{\text{avg}}) / (N_{\text{sub}} \times T_{\text{avg}} + 1.2 \times T_{\text{max_max}}) \quad (2)$$

We can create a table of “typical” values for the future to get a feel for how Average Utilization Levels will vary as a function of time. If we assume that Nsub drops as a function of time and if we assume that both Tmax and Tavg grow by a CAGR of 50%, then we can create the following table showing how Average Utilization Levels may change with time.

Table 3 - Average Utilization Levels as a Function of Time

Year	Nsub for SG	Tavg for sub (Mbps)	Tmax_max for sub (Mbps)	Avg BW for SG (Mbps)	Required BW Capacity for SG (Mbps)	Average Utilization Level
2017	400	1.00	1000	400	1600	25.0%
2019	200	2.25	2250	450	3150	14.3%
2021	100	5.06	5063	506	6581	7.7%
2023	50	11.39	11391	570	14238	4.0%

Table 3 illustrates several interesting trends that will likely occur as operators move into the future. First, it is quite apparent that the Average Bandwidth within the Service Groups of the future will likely grow, but expected Node Splits of the future will likely keep this Average Bandwidth growth to a relatively slow rate.

Required Bandwidth Capacity within the Service Groups of the future will also grow, but rapidly increasing values of Tmax will cause this growth rate to be quite fast.

With a slow growth rate in Average Bandwidths and a fast growth rate in Required Bandwidth Capacities, the Average Utilization Level is expected to drop quite rapidly as operators move into the era of 1+ Gbps services. In particular, for the example shown in Table 3, the Average Utilization Levels may drop from 25% in 2017 to only 4% in 2023.

This low Utilization Level may be viewed negatively by those who are responsible for paying the costs of the ever-increasing Required Bandwidth Capacities, because it appears that the extra Bandwidth Capacity is mostly being used to provide headroom for infrequent bandwidth bursts that exceed the Average Bandwidth levels within the Service Group.

If this low Utilization Level is deemed to be undesirable, then there are several network design techniques that can be utilized by the operator to reduce the severity of this low Utilization problem. These design techniques include:

- Using multiple Remote PHY or Remote MACPHY devices per Service Group (which effectively increases the Nsub value within the Service Group, which increases the Average Utilization level)
- Using MAC-based Nodes with sub-tending Remote PHY nodes (which also increases the Nsub value within the Service Group, which increases the Average Utilization level)
- Using Selective Subscriber Migration (which moves the highest Tmax subscribers off of the HFC plant and gives them Fiber To The Home optical feeds. The resulting subscribers have a much lower Tmax_max value, which increases the Average Utilization level).

These three techniques are clearly illustrated within Figure 7.

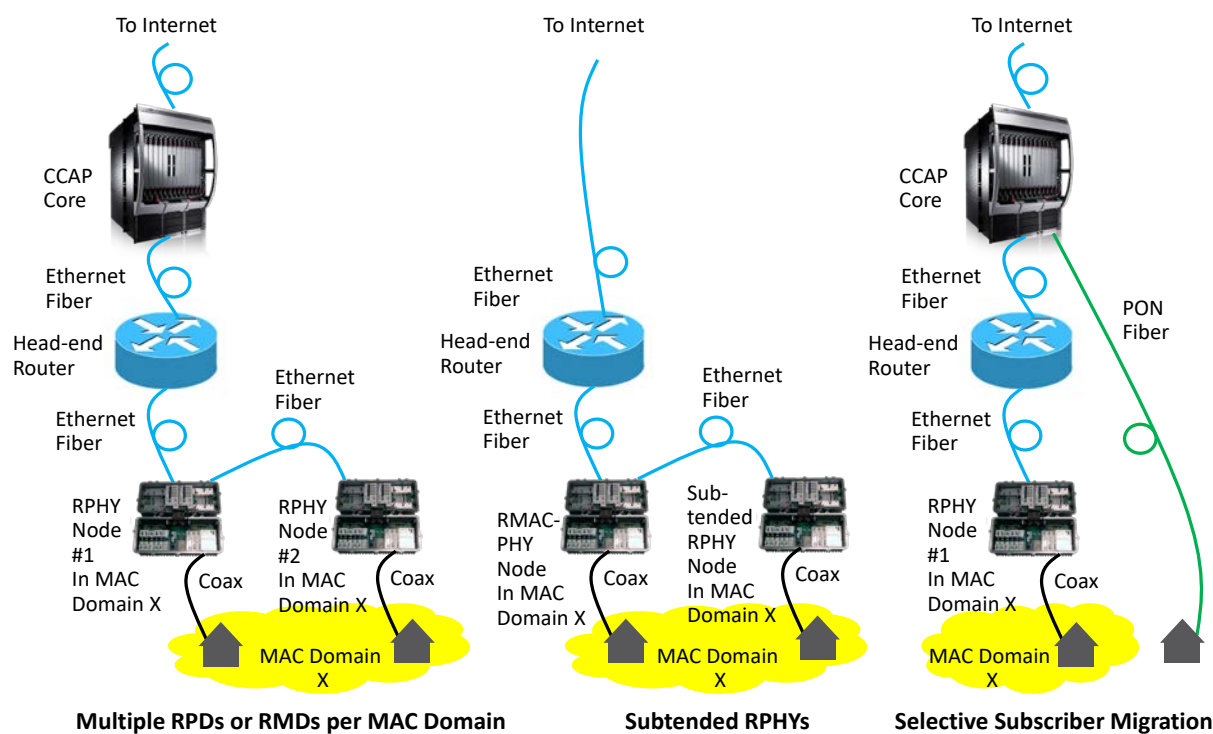


Figure 7 - Various Techniques for Increasing Utilization Levels in MAC Domains

Figure 8 illustrates (using bandwidth candlestick diagrams) how these various techniques help increase the Utilization Levels with example numbers. Within each candlestick diagram, the orange color indicates the magnitude of the first term ($N_{sub} \times T_{avg}$) within the ARRIS QoE-based Traffic Engineering Formula of Equation (1), which is the static bandwidth. The gray color indicates the magnitude of the second term ($1.2 \times T_{max_max}$) within the ARRIS QoE-based Traffic Engineering Formula of Equation (1), which is the QoE headroom bandwidth. As stated above, the Average Utilization Level is given by Equation (2). Within the candlestick diagrams, this is given by the height of the orange color divided by the height of the orange + gray color. In all three scenarios, the changes increase the Average Utilization Levels by about a factor of two. It should be apparent that combinations of these changes (ex: Multiple RPDs per MAC Domain plus Selective Subscriber Migration) can produce even higher Average Utilization Levels.




Figure 8 - Impact of Various Techniques for Increasing Utilization Levels in MAC Domains

TCP Performance Levels and 1+ Gbps Services

Even with very high bandwidth capacities available within the links of the network, the performance of 1+ Gbps services can still be limited by the operation of the Transmission Control Protocol (TCP). TCP is the underlying protocol used to manage the flow of data for most of the data transmissions across the Internet. It includes complex flow control and congestion control algorithms that were put in place to ensure that multiple users can interact and share the Internet bandwidth capacity in a fair and equitable fashion.

Many different flavors of TCP have been proposed and utilized over the years. Each one has a focus on a particular set of performance attributes within high-speed data connections. One of the more popular TCP algorithms in use within many Internet servers today is TCP CUBIC.

The CUBIC algorithm was developed in the 2005 time-frame, and it was optimized to perform well with high-bandwidth connections in long, large-latency networks. The congestion control algorithm within CUBIC was designed to make changes to connection bandwidth levels (as a result of detected packet loss or sudden packet delays) in a much less aggressive fashion than many of the congestion control algorithms that preceded it, such as TCP New Reno, TCP BIC, etc. Generally, TCP's throughput is dictated by packet loss and Round Trip Time (RTT), but CUBIC manages the throughput by only the packet loss. Another important change introduced with CUBIC was the use of timers instead of TCP Acknowledgements to trigger increases in connection bandwidth levels (or congestion window), which resulted in it performing well for both low-latency and long-latency networks.

As the name of the algorithm suggests, the congestion window function of CUBIC is a cubic function:

$$W_{cubic} = C \times (t - K)^3 + W_{max}$$

where C is a window scaling factor, t is the elapsed time from last window reduction, W_{max} is the window size just before the last window reduction, and $K = \sqrt[3]{W_{max}\beta/C}$, where β is a constant multiplicative decrease factor applied for window reduction at time of the packet loss.

β and C are the control knobs in this formula. The two values are also exposed to the end users for TCP CUBIC configuration. An optimal value of the knob variables is described as $\beta = 0.8$ and $C = 0.4$, by the inventors of CUBIC [RHE1].

As can be seen in the congestion window growth formula, CUBIC is independent of RTT and only depends on the packet loss times, unlike many of its predecessor congestion algorithms. Because of these beneficial changes, CUBIC was selected to be used in Linux kernels quite a few years ago. Since many servers within the Internet are based on Linux kernels, this implies that CUBIC is responsible for managing data flows for many of the client-server sessions operating over the Internet. As a result, the authors have opted to focus on CUBIC's TCP implementation within this section.

Two questions related to CUBIC's impact on 1+ Gbps data flows were studied by the authors. These included the impact of typical network delays as well as the impact of typical HFC packet error rates. Each of these is discussed below.

Network Delays during Connection Slow-Start: Once a TCP connection is up and running, network delays are not as problematic for CUBIC as they were for other TCP implementations, because of CUBIC's novel use of timers instead of TCP Acknowledgements.

However, CUBIC still uses Acknowledgements to pace its growth of the Congestion Window and bandwidth levels during the "Slow-Start" operation that occurs when a TCP connection is first established. This early-stage use of Acknowledgements may have an impact on how quickly a TCP connection can achieve its desired 1+ Gbps data rates.

It can be shown that for any TCP connection, the maximum TCP throughput level at any instant in time is given by:

$$\text{Maximum TCP Throughput} = \frac{\min(\text{Congestion Window Size}, \text{Receive Window Size})}{RTT} \quad (3)$$

where RTT is the Round-Trip Time for data transmissions, the Congestion Window Size (cwnd) is an internal state variable maintained by TCP at the source of the transmissions, and that Receive Window Size (rwnd) is an internal state variable maintained by TCP at the destination of the transmissions.

Under normal circumstances, the Receive Window Size is usually quite large, so the Maximum TCP Throughput is limited by *Congestion Window Size/RTT*. In essence, the TCP transmitter cannot send more than a Congestion Window's worth of packets until an ACK arrives for the first packet that was sent. TCP tends to increase the size of the Congestion Window when packets are flowing smoothly (without any packet drops or sudden packet delays), and that increase permits the TCP connection to transmit data at increasing data rates. Longer RTT's tend to reduce the data rates, and shorter RTT's tend to increase the data rates.

During Slow-Start, the Congestion window is increased by one Maximum Segment Size (MSS) every time an ACK arrives. It can thus be noted that the number of Acknowledgements that come back every RTT is double the number that came back during the previous RTT window (assuming no packet loss). If we assume that the MSS is 1500 bytes (12000 bits), then we can develop a formula indicating the amount of time 'T' that it takes during Slow-Start to increase the size of the Congestion Window to a level that permits a particular desired bandwidth level. This formula is given by Equation (4):

$$\text{Desired Bandwidth} = \left(2^{\left(\frac{T}{RTT}\right)}\right) \times \frac{12000 \text{ bits}}{RTT} \quad (4)$$

RTT can vary a lot depending on the geographical locations of the clients and servers, but a "typical" RTT value might be between 20 msec and 100 msec. Using these typical RTT values, we can plug in the Desired Bandwidth of 1 Gbps and calculate the amount of time T that would be required for the Congestion Window to be increased (during the Slow-Start period) to a value that would permit 1 Gbps service to flow. When RTT = 20 msec, the total time T for the Slow-Start Congestion Window to grow to support 1 Gbps is given by 214 msec. When RTT = 100 msec, the total time T for the Slow-Start Congestion Window to grow to support 1 Gbps is given by 1.3 sec.

Thus, it should be clear that for short 20 msec RTTs, the Congestion Window will likely grow to permit 1 Gbps rates in a fairly short window of time (214 msec). However, for longer 100 msec RTTs, the Congestion Window will take a relatively lengthy 1.3 seconds of time to grow to permit 1 Gbps rates, and this lengthy time period may not be ideal for some applications. As a result, future TCP congestion

control algorithms may need to be modified to eliminate this potential problem associated with slow bandwidth growth rates to 1+ Gbps service levels.

Packet Error Rates: Even if there is adequate bandwidth capacity to support high 1+ Gbps TCP connection rates within a network, there is another TCP effect that can cause the actual bandwidth experienced by users to be much lower than the bandwidth capacity. This effect is a result of the actions taken by the congestion avoidance algorithm of TCP.

The congestion avoidance algorithm responds to dropped packets (detected through unacknowledged TCP packets) by assuming that the dropped packets resulted from overloaded and congested buffers within intermediate network elements. This assumption leads the congestion avoidance algorithm to reduce the Congestion Window at the source, which in turn reduces the actual bandwidth of the connection (since bandwidth is roughly given by *Congestion Window/RTT*). These types of algorithms help ensure that the Internet remains operational during periods of congestion.

However, in a lossy network, the packet drops may also be caused by noise that corrupts the packets. Thus, noisy networks will experience packet drops, and those packet drops will also cause the TCP congestion control algorithms to throttle the bandwidth of the connection. As a result, it is important to determine the maximum bandwidth levels that might be permitted by a TCP connection that might be propagating over a noisy HFC plant or over noisy cables within a home.

To perform this analysis, the authors again focused on the operation of the CUBIC TCP algorithm. Simulations were performed using a CUBIC model in the NS2 simulator [WEI1], and various RTTs were simulated and various packet error rates were injected into the simulation to model a lossy HFC plant or lossy home network. In the end, the actual TCP connection performance was monitored for each of the RTT values packet error rate values.

The results are illustrated in Table 4, where the maximum TCP connection bandwidth (in Mbps) is displayed for different combinations of RTT (along the top) and Packet Error Ratio (PER) along the left side.

Table 4 - Maximum TCP Bandwidth (in Mbps) for Various RTT and PER Values

DS PER	RTT →	10ms	20ms	40ms	60ms	80ms	100ms
	DS BER (w/ 1500 byte pkts)						
10^{-6}	8×10^{-11}	1420.95	984.78	788.51	701.19	650.51	606.49
10^{-5}	8×10^{-10}	337.85	214.12	154.67	123.93	101.69	89.49
10^{-4}	8×10^{-9}	116.26	62.86	38.74	29.21	25.52	22.66
10^{-3}	8×10^{-8}	33.33	18.22	10.29	6.95	5.23	4.28
10^{-2}	8×10^{-7}	9.37	5.17	2.87	1.97	1.55	1.29

Within this table, it can be seen that only one of the grid elements (RTT = 10 msec, DS PER = 10^{-6}) actually supports 1+ Gbps service levels. All of the other grid elements are limited to actual TCP bandwidth levels that are lower than 1 Gbps. This illustrates how limiting TCP congestion control algorithms can be in the presence of noise-induced packet errors. The results imply that very low packet error rates and short RTTs must be maintained on the network to ensure high 1+ Gbps TCP throughput levels on any single TCP connection. This can be quite challenging because:

- 1) Reducing packet error rates to this level may require upgrades to the plant to reduce plant noise.
- 2) RTT of 10 msec or less can only be provided by close servers, so TCP connectivity to distant servers would not allow the 1+ Gbps service levels on the TCP connection. Most data would have to be cached in caching servers very close to the subscribers.

There will undoubtedly be pressure on MSOs and content providers to provide improvements like those described above as the world heads toward 1+ Gbps services. However, there are also likely to be a few other changes that will take place in the future. For example, the above results were obtained using CUBIC's TCP algorithm. New TCP algorithms are always being architected and designed to accommodate the new requirements as the Internet evolves, so perhaps a new variant of TCP will be developed that is more amenable to packet errors and longer RTTs and which will permit 1+ Gbps bandwidth levels even in the presence of these challenging environments.

In addition, there is another factor that must be taken into account. The above simulations assumed that the 1+ Gbps service level was going to be offered over a single TCP connection and that the single TCP

connection was required to support the entire 1+ Gbps bandwidth. Fortunately, many applications in the world today do not operate in this fashion, and instead opt to transmit data between servers and a single client using more than one TCP connection. For example, Peer-to-Peer applications such as BitTorrent tend to use many TCP connections to deliver their content. In fact, even traditional web pages are usually delivered using multiple TCP connections to deliver the desired content.

The use of multiple TCP connections to deliver content provides a major benefit to the bandwidth levels for the application. In particular, it can be shown that an application which uses N parallel TCP connections to deliver content to a client will roughly result in an N times speed-up in the bandwidth levels associated with that application (when compared to the bandwidth levels that would have been achieved with a single TCP connection).

This fact was verified in the simulator, where an application was run with 8 parallel TCP connections. The particular network conditions were selected to be $PER = 10^{-4}$ and $RTT = 10$ msec. With a single TCP connection, the maximum bandwidth was found (in Table 3) to be 116.26 Mbps. However, with 8 parallel TCP connections, the total bandwidth was increased to be 917.7 Mbps (which is close to 8 times the bandwidth of the single TCP connection). The actual bandwidth found in each of the 8 individual TCP connections is indicated in Figure 9. The benefits illustrated in Figure 9 may lead more and more applications to take advantage of parallel TCP connections in the future of 1+ Gbps service levels.

FLOW ID	Average Throughput (Mbps)
1	117.61
2	124.09
3	115.02
4	112.56
5	104.58
6	114.21
7	111.31
8	118.28
Aggregate BW	917.70

Figure 9 - Bandwidth in Each of the 8 Individual TCP Connections and the Aggregate Bandwidth

Symmetric Services and 1+ Gbps Services

Operators are continually being presented with new challenges. Over time, there has been more and more marketing pressure for operators to provide so-called “Symmetric Services.” In a Symmetric Services environment, the Upstream T_{max} value and the Downstream T_{max} value must be similar in value. Ideally, the Upstream and Downstream T_{max} values are equal. This type of service offering has become more popular as PON service providers have begun to offer it in recent years.

Providing a Symmetric Services offering has traditionally been a challenge for MSOs, because their traditional HFC plant was originally architected using Frequency-Division Duplex techniques to be asymmetrical, offering much more bandwidth capacity in the Downstream direction (using the 54-1002 MHz range of the spectrum) than in the Upstream direction (which was limited to using only the 5-42 MHz range of the spectrum in North America). Prior to the arrival of DOCSIS 3.1, this yielded a theoretical Downstream bandwidth capacity of ~5.6 Gbps and a theoretical Upstream bandwidth capacity of ~144 Mbps. As can be seen, this is quite asymmetrical.

The arrival of DOCSIS 3.1 has improved the situation to some extent. DOCSIS 3.1 operators can (for example) choose to use a High-Split HFC network with the Upstream spectrum contained within a 5-204 MHz range and with the Downstream spectrum contained within a 258-1218 MHz range. Assuming 10 bps/Hz useable spectral efficiencies during operation, this can yield a theoretical Downstream bandwidth capacity of ~9.6 Gbps and a theoretical Upstream bandwidth capacity of ~2 Gbps. While these are much higher bandwidth capacity levels, they are still asymmetrical.

An asymmetrical HFC network environment of this nature is therefore challenged when asked to provide Symmetrical Services, and the challenge becomes even more difficult in a 1+ Gbps service world where operators may be called upon to provide a 1 Gbps Downstream × 1 Gbps Upstream service offering. In the future, these numbers may grow to be 2 Gbps Downstream × 2 Gbps Upstream, or 5 Gbps Downstream × 5 Gbps Upstream, or even 10 Gbps Downstream × 10 Gbps Upstream. It should be clear that these higher levels of Symmetrical Service could be a challenge for today’s asymmetrical HFC network.

A potential solution to this problem is currently being developed. The resulting solution employs Full Duplex DOCSIS (FDX) capabilities. While the details of this solution are beyond the scope of this paper, the basic idea is to permit the Downstream spectrum and the Upstream spectrum to overlap and utilize the same portion of the spectrum at the same time for both Downstream and Upstream transmissions, as shown in Figure 10. This form of operation will require the use of complex technologies such as Echo Cancellation to permit the simultaneous transmissions of signals in both directions.

Consider a system in which the Downstream spectrum is permitted to operate from 104-1218 MHz and the Upstream spectrum is permitted to operate from 5-684 MHz. Assuming 10 bps/Hz useable spectral efficiencies during operation, this can yield a theoretical Downstream bandwidth capacity of ~11.1 Gbps and a theoretical Upstream bandwidth capacity of ~6.7 Gbps. These much higher bandwidth capacity levels can begin to permit Symmetrical Services- even for a system requiring a 5 Gbps Downstream × 5 Gbps Upstream service offering.

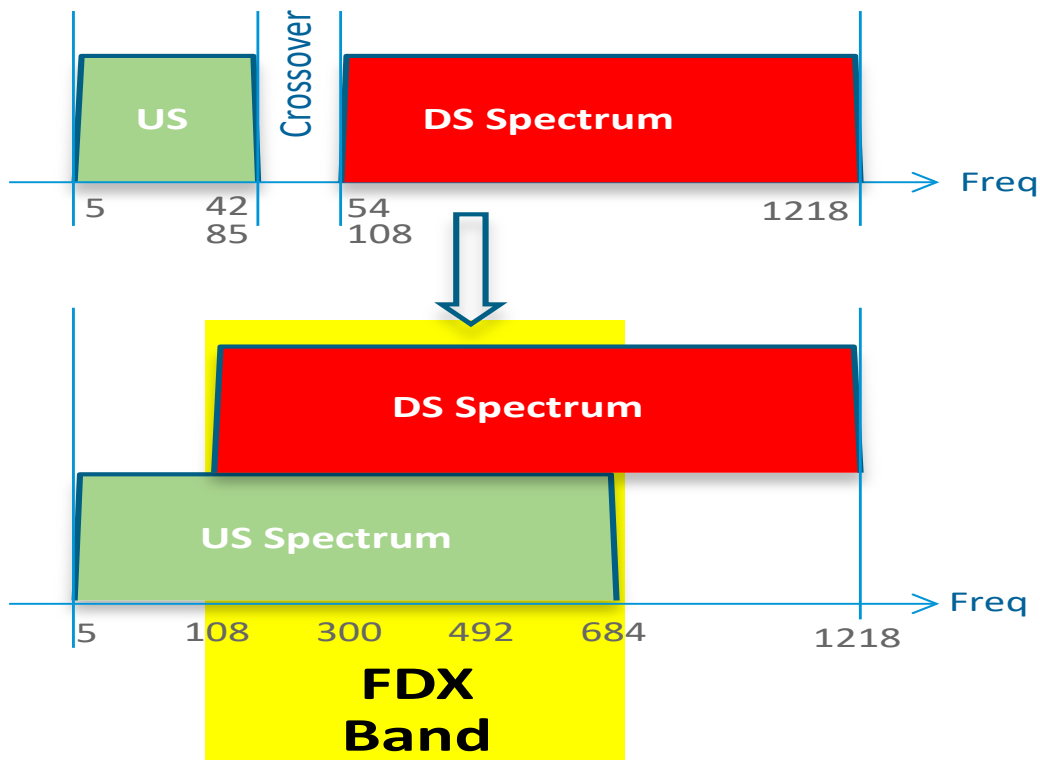


Figure 10 - Full Duplex DOCSIS (FDX) and Use of the HFC Plant Frequency Spectrum

Higher bandwidth Symmetrical Service offerings may require the use of other solutions that are currently being studied. For example, a 20 Gbps Downstream \times 20 Gbps Upstream service offering might require the use of Extended Spectrum DOCSIS techniques, which permit DOCSIS 3.1 and FDX operations to extend beyond 1218 MHz (and even beyond the optional 1794 MHz limit mentioned within the DOCSIS 3.1 specification). These types of Extended Spectrum DOCSIS solutions are illustrated in Figure 11, and they may become practical in the 2020 decade. [CLO2]

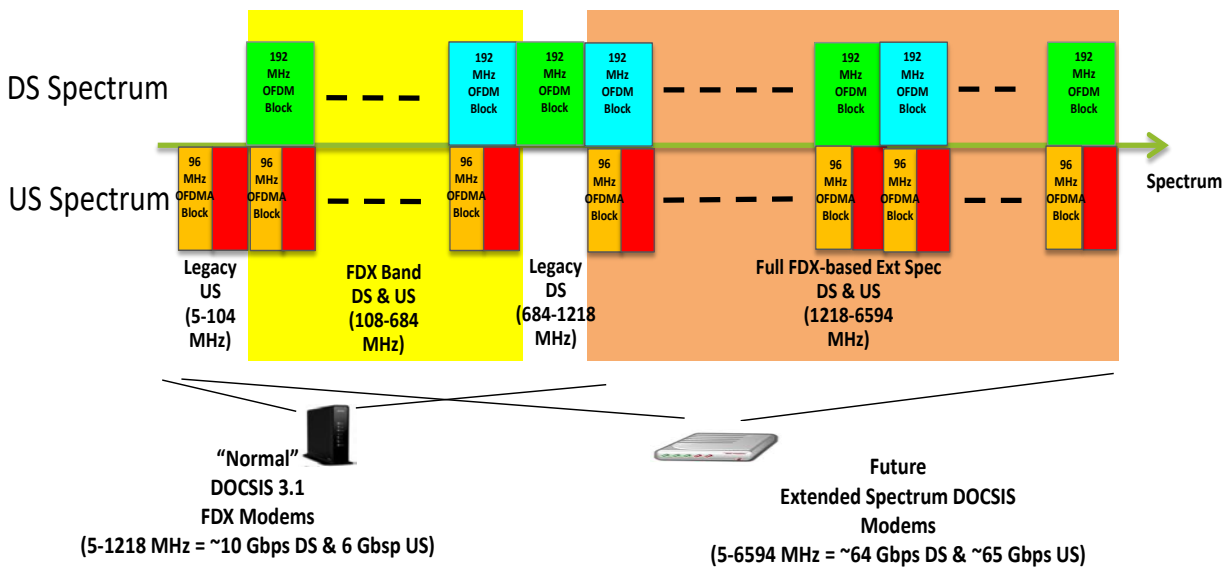


Figure 11- Example of Extended Spectrum DOCSIS Operation

In general, new concepts are continually being developed that should permit Symmetrical Services — even in the 1+ Gbps service environments of the future.

RPHY Linear IP Video Delivery and 1+ Gbps Services

The transition to IP Video is a popular transition that MSOs have been planning for years. Most operators foresee many benefits resulting from this transition, but they have been waiting for the correct timing (technological improvements, bandwidth availability, etc.) to introduce these changes.

IP Video over DOCSIS can yield many benefits, including:

- Convergence of all services over a single DOCSIS infrastructure platform
- The increased number of programs that can be offered due to the Statistical Multiplexing Gains that result from DOCSIS 3.0 Channel Bonding (which are not available when using MPEG video transport)
- The increased number of programs that can be offered due to the higher spectral efficiencies of DOCSIS 3.1 OFDM (which are not available when using MPEG video transport)
- The channel efficiencies that result from Switched Digital Video-like operation provided by the dynamic Service Flows and load-balancing of Service Flows and the use of IP Multicast streams within the CCAPs

There are two general types of IP Video Services that can be offered to IP Video subscribers, and each of these two types can be implemented in at least two different ways:

- Video on Demand (VoD) IP Video Service
 - a. Unicast VoD IP Video Service — using IP Unicast, whereby Unicast DEPI tunnels can be utilized between the CCAP Core and the Remote PHY Nodes in a Remote PHY environment
 - b. Multicast VoD IP Video Service — using IP Multicast- which might be valuable if the VoD Service Group needs to span multiple Remote PHY Nodes in a Remote PHY environment, whereby Multicast DEPI tunnels can be utilized between the CCAP Core and Remote PHY Nodes
- Linear IP Video Service (Linear)
 - a. Nailed-Up, Always-On Linear IP Video Service — Nailed-Up Linear using IP Multicast, whereby Multicast DEPI tunnels can be utilized to send the feed to multiple Remote PHY Nodes in a Remote PHY environment
 - b. Switched Linear IP Video Service — Switched Linear using IP Multicast, whereby Multicast DEPI tunnels can be utilized to send the feed to multiple Remote PHY Nodes in a Remote PHY environment

Thus, it can be seen that IP Multicast and Multicast DEPI tunnels (for Remote PHY environments) may be called into service for at least three out of the four IP Video environments described above. It is therefore instructive to consider some of the side effects of using RPHY Multicast DEPI tunnels for IP Video delivery between the CCAP Core and the Remote PHY Node as the network moves towards the higher DOCSIS 3.1 channel bandwidths associated with the 1+ Gbps services of the future.

According to the Remote PHY Specification [RPH1], the Multicast DEPI support is “meant for the replication of an entire QAM or OFDM channel to multiple RPDs.” Each QAM or OFDM channel is associated with a single unique pseudo-wire within the Multicast DEPI tunnel. It should thus be clear that the QAM channel or OFDM channel associated with a particular pseudo-wire **MUST** carry exactly the same information (i.e.- video programs) to all of the destination RPHY Node endpoints of the Multicast DEPI pseudo-wire.

If there are ANY differences in any of the signals (i.e. video programs) on a QAM or OFDM channel at any of the RPHY Nodes being fed by a particular Multicast DEPI pseudo-wire, then that RPHY Node receiving the different signals cannot receive the same pseudo-wire as the other RPHY Nodes, and a separate pseudo-wire must be set up for the different QAM or OFDM channel feeds to that particular RPHY Node. As will be shown below, this is an important point that may drive many architectural decisions.

In past studies on Service Group bandwidth requirements, it has been shown that the decision on whether it is optimal to transmit a particular set of video program using IP Unicast or Switched IP Multicast or Always-on IP Multicast is a function of several factors. These factors include the popularity of the programs (defined by the alpha value (α) for the assumed Power Law Distribution, where higher alpha values imply a smaller number of very popular programs), the number of viewers, and the number of available programs in the content library.

Note: The Power Law Distribution can be defined as follows. For a list of programs numbered $i = 1$ to $i = N$, the probability that a particular subscriber selects a particular program number (i) is given by:

$$P(i) = \frac{i^{-\alpha}}{\sum_1^N i^{-\alpha}}$$

where the summation is taken from program $i = 1$ to program $i = N$. For the Power Law popularity curve, $\alpha = 0.7$ is a typical value that will be used in this example.

An example scenario for a system with various numbers of available programs in the content library is shown in Figure 12. We will focus on the 250 program case with blocking probabilities of 0.01% (identified by the bottom blue plot within Figure 12).




Figure 12 - Required # of Transmitted Video Programs vs # of Viewers for Switched IP Video

Within this figure, it can be seen that for a small number of Active Viewers ($< \sim 100$), the MSO might as well utilize IP Unicast, because there are little savings from using IP Multicast. If the MSO's IP Multicast solution is robust and low-cost, then there is no harm in using IP Multicast in this range. Within this figure, it can also be seen that for a large number of Active Viewers ($> \sim 900$), the MSO might as well utilize Always-On IP Unicast, because there are little savings from using Switched IP Multicast. If the MSO's IP Multicast solution is robust and low-cost, then there is no harm in using IP Multicast in this

range In the middle region of the Figure (where the number of Active Viewers is between 100 and 900), there is clearly a region where Switched IP Multicast can yield improvements in the bandwidth required by the Service Group over both unicast and Always-On IP Multicast. Thus, when considering the Bandwidth savings within the Service Group, the use of Switched IP Multicast can always yield some improvements.

The above analysis is only focusing on the problem from the point-of-view of the Service Group bandwidth requirements - i.e. it attempts to answer the simple question, “How much spectrum is required to carry the Linear video signals on the RF coaxial connection that runs between the Fiber Nodes and the subscriber homes?”

There is another topic and point-of-view that probably should also be analyzed within an RPHY environment, and that is the topic of the CCAP Core processing and input/output requirements. This alternative analysis would attempt to answer the question, “How much processing power and input/output bandwidth is required on the CCAP Core in the headend or hub to support the Linear video signals?” In actuality, both of these topics (Service Group bandwidth and CCAP Core bandwidth) need to be considered together.

To illustrate the point, let us consider this particular topic area of Remote PHY systems in more detail by using a simple example. Assume that an operator has converted to a Fiber Deep Remote PHY environment, and the resulting headend contains 500 Remote PHY Fiber Nodes with 50 subscriber homes per Remote PHY Fiber Node. This results in a total of 25,000 subscriber homes attached to the headend. Assume also that the operator has 250 Linear video programs within their content library and assume that each video program consumes an average of 9 Mbps of bandwidth in the DOCSIS pipe. Note the 9 Mbps average might result from a blend of Standard Definition and High Definition and 4K content. Obviously, the total Service Group bandwidth required to continuously transmit 250 Linear video programs at 9 Mbps each would be 2250 Mbps (2.25 Gbps). However, there are many ways to architect a system that delivers this 2.25 Gbps stream of Linear bandwidth from the CCAP Core to each of the 500 Remote PHY Fiber Nodes. Techniques include:

- Using “VoD-like” IP Unicasting of a single Linear video program stream from the CCAP Core to each of the active subscribers viewing a Linear video program. This technique is used today by many MSOs to deliver Linear video to many second screen devices (PCs, smart-phones, and tablets) over DOCSIS. There are 25,000 subscriber homes in the example head-end, so if each subscriber home contains ~2.5 people, that results in 62,500 people being provided with service.

If 60% of them (i.e. 37,500 people) are actively watching Linear video content in a particular busy-hour window of time, that implies that there would be $(62,500) \times (60\%) = 37,500$ IP Unicasted Linear video programs (consuming $(37,500) \times (9 \text{ Mbps}) = 337.5$ Gbps of aggregate bandwidth) sent from the CCAP Core to the group of Remote PHY Fiber Nodes. Each of the 500 Fiber Nodes will (on average) receive 1/500th of this aggregate bandwidth, which results in ~675 Mbps of Linear video bandwidth consumed by the subscriber homes in each Service Group. (See Figure 13).

It should be clear that the packets for the Linear programs can be multiplexed in with regular High-Speed Data packets on a channel set going to a particular Fiber Node.

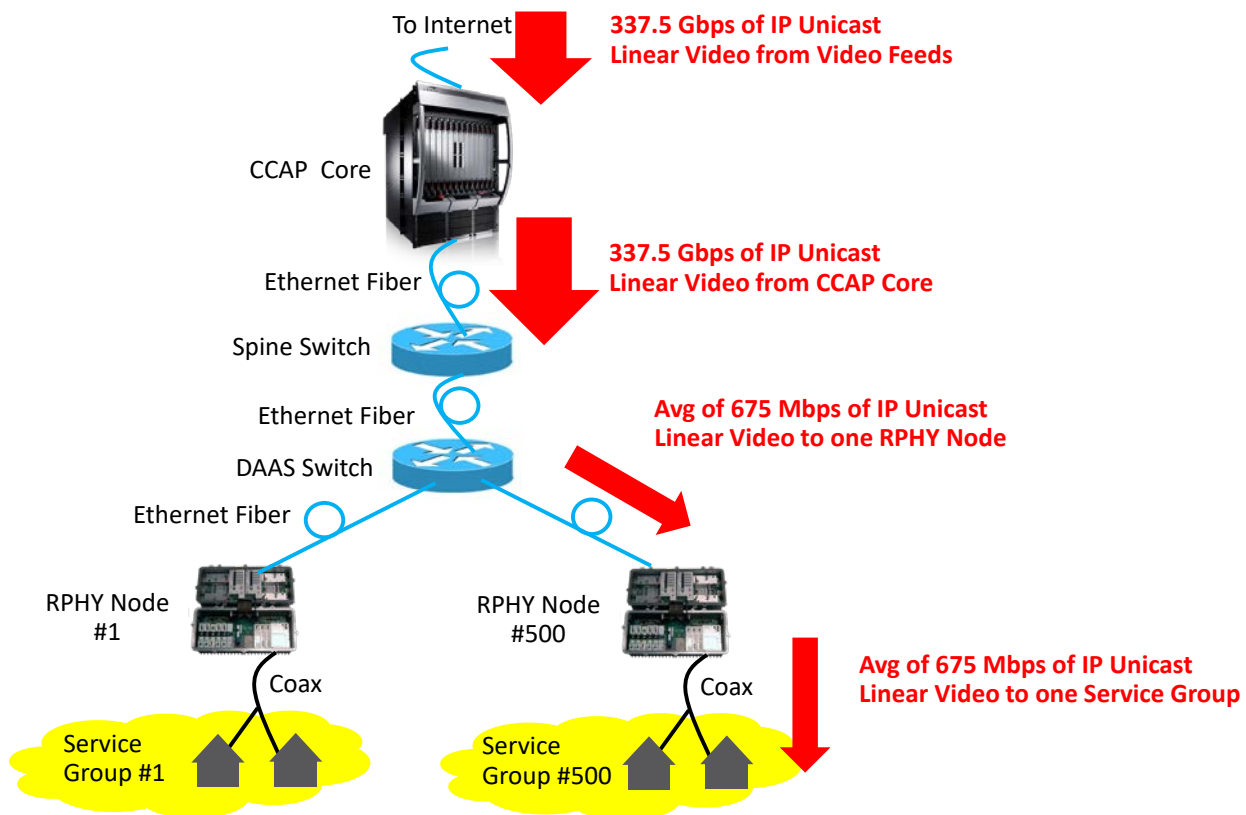


Figure 13 - Linear IP Video Delivery Example using IP Unicasting

- Using Always-On, Nailed-Up IP Multicasting of a single 2.25 Gbps streamed multiplex containing all 250 Linear video programs from the CCAP Core to all of the 500 Remote PHY Fiber Nodes (using the Spine/DAAS Switch Network to replicate the single stream from the CCAP Core to all 500 Remote PHY Fiber Nodes). This implies that only $(250) \times (9 \text{ Mbps}) = 2.25 \text{ Gbps}$ of aggregate bandwidth would be processed and transmitted from the output of the CCAP Core. The Spine/DAAS switch network would replicate this bandwidth to each of the Remote PHY Fiber Nodes.

Each of the 500 Fiber Nodes will continuously receive and transmit this entire 2.25 Gbps streamed multiplex of bandwidth onto the HFC plant. (See Figure 14). It should be clear that the packets for the Always-On, Nailed-Up Linear programs must be isolated to a separate channel set to ensure that the IP packets within that channel set are the same for every Fiber Node. They cannot be multiplexed in with regular High-Speed Data packets going to the Fiber Node.

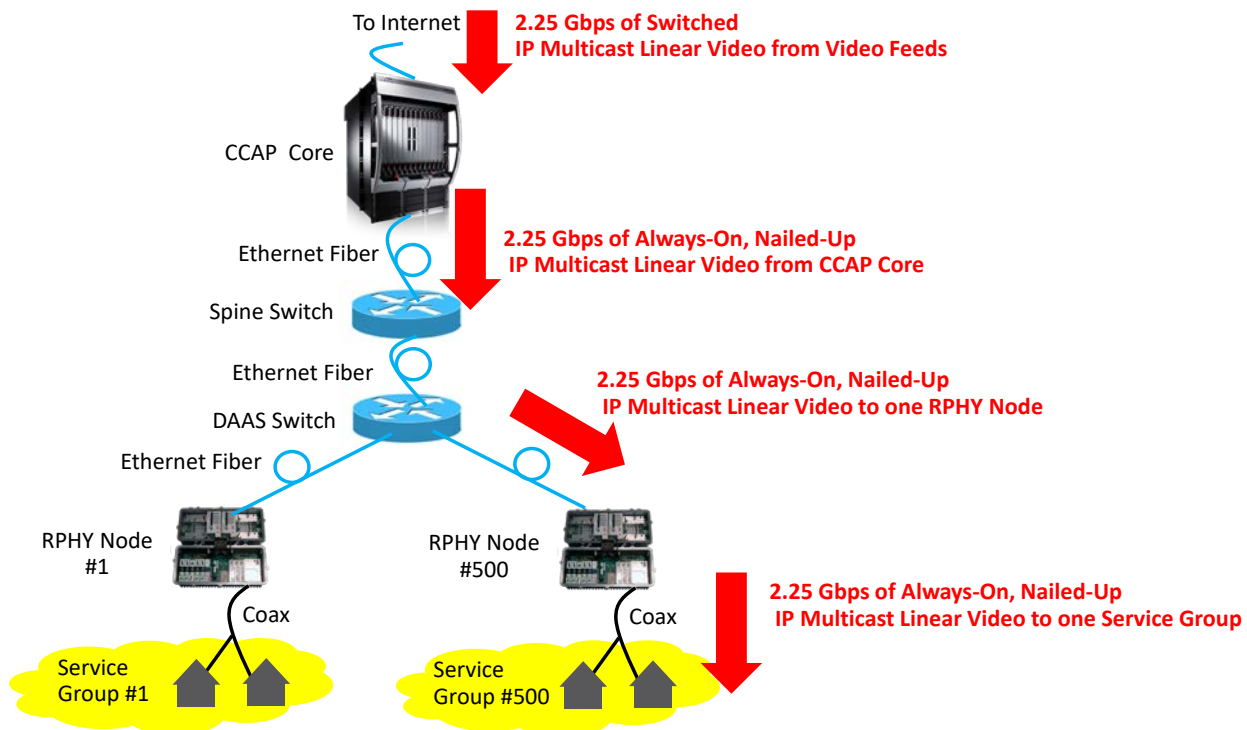


Figure 14 - Linear IP Video Delivery Example using Always-On, Nailed-Up IP Multicasting

- Using Switched IP Multicasting of only the actively-viewed programs from the CCAP Core to the active viewers within all of the 500 Remote PHY Fiber Nodes (using the Spine/DAAS Switch Network to replicate each active video program stream from the CCAP Core to whichever ones of the 500 Remote PHY Fiber Nodes have active viewers for that particular video program stream). With up to 37,500 active viewers supported by the headend, it is practically guaranteed that all 250 programs will need to be processed and transmitted from the CCAP Core, resulting in a total of $(250) \times (9 \text{ Mbps}) = 2.25 \text{ Gbps}$ of aggregate bandwidth transmitted from the CCAP Core.

Assume that each Remote PHY Fiber Node has $(50 \text{ subscriber homes}) \times (2.5 \text{ people per subscriber home}) \times (60\% \text{ viewing activity}) = 75$ actively viewing subscribers during the busy-hour window of time. Based on the blue 250 program curve of Figure 12, it can be seen that a Service Group with 75 active viewers would require ~ 75 programs to be transmitted, which would imply an average Linear video bandwidth of $(75) \times (9 \text{ Mbps}) = 675 \text{ Mbps}$. (See Figure 15).

At first glance, this solution seems ideal, because it keeps the bandwidth low in the CCAP Core and it keeps the bandwidth low in the Remote PHY Fiber Node's Service Group. Unfortunately, this solution is not easily implementable, because the set of active programs being viewed in one Remote PHY Fiber Node would typically be different than the set of active programs being viewed in another Remote PHY Fiber Node. Since the multiplex of IP packets injected into a particular SC-QAM or OFDM channel for a particular Fiber Node have to be created in the MAC, this implies that a different multiplex on a different channel pseudo-wire would have to be uniquely constructed for each Remote PHY Fiber Node by the CCAP Core's MAC processing functions. This ends up implying that a unique and (most likely) different channel set (ex: ~ 675

Mbps/36 Mbps \approx 19 SC-QAM channels on \sim 19 pseudo-wires) containing an average of \sim 675 Mbps must be created for each of the 500 Remote PHY Fiber Nodes.

Note: We could also use of fraction of a single 192 MHz OFDM channel and a single pseudo-wire for carrying the 675 Mbps of Linear video content to each one of the 500 remote PHY Fiber Nodes. Thus, a total aggregate bandwidth of $\sim(500) \times (675 \text{ Mbps}) = 337.5 \text{ Gbps}$ must be processed and transmitted from the CCAP Core to service the Linear program content library containing the 250 programs. (See Figure 16). That unfortunately places a large processing load and interface load on the CCAP Core.

This problem can be alleviated if the MAC processing is moved from the CCAP Core into the Remote PHY Fiber Node (creating a Remote MACPHY solution). In this case, the headend could source 2.25 Gbps of IP Multi-casted Linear video content to each of the Fiber Nodes, and the Fiber Node's MAC processing functions can parse through the video content and only process and forward (via IP Multicast) the \sim 675 Mbps of viewed Linear content onto the HFC plant of the Service Group. (See Figure 17).

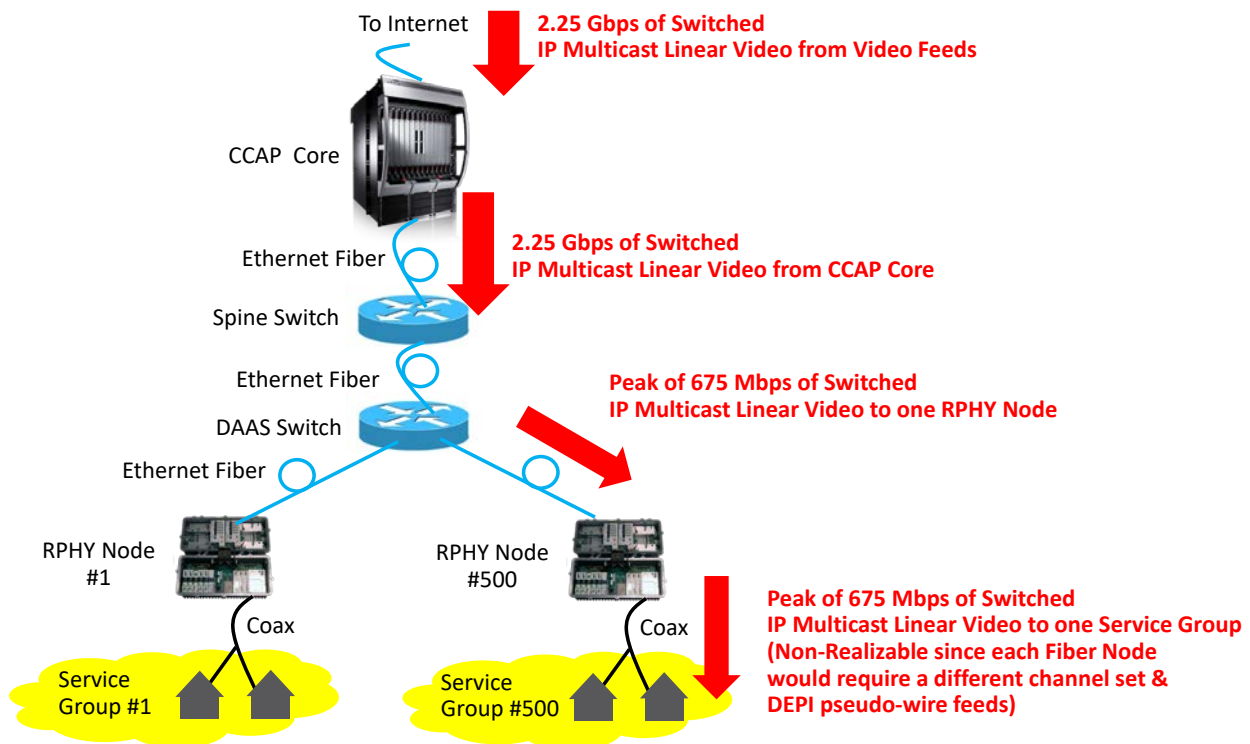


Figure 15 - Linear IP Video Delivery Example using Switched IP Multicasting (Non-Realizable)

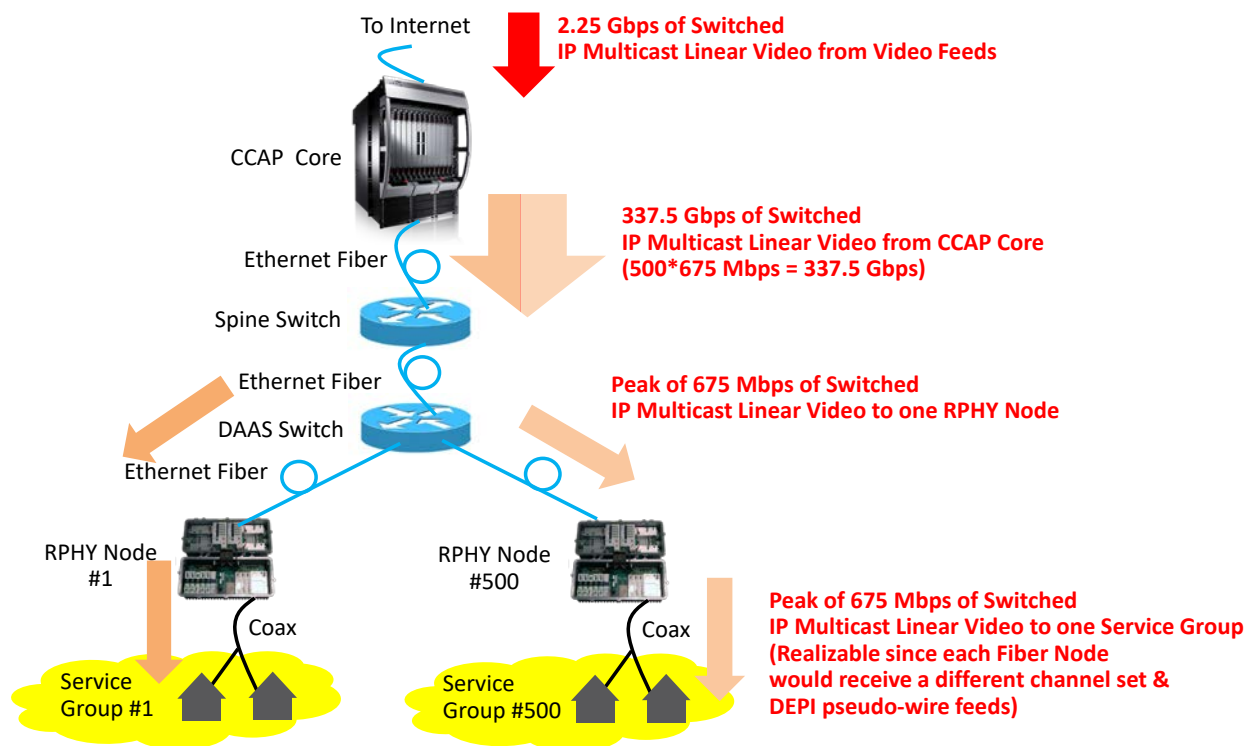


Figure 16 - Linear IP Video Delivery Example using Switched IP Multicasting (Realizable)

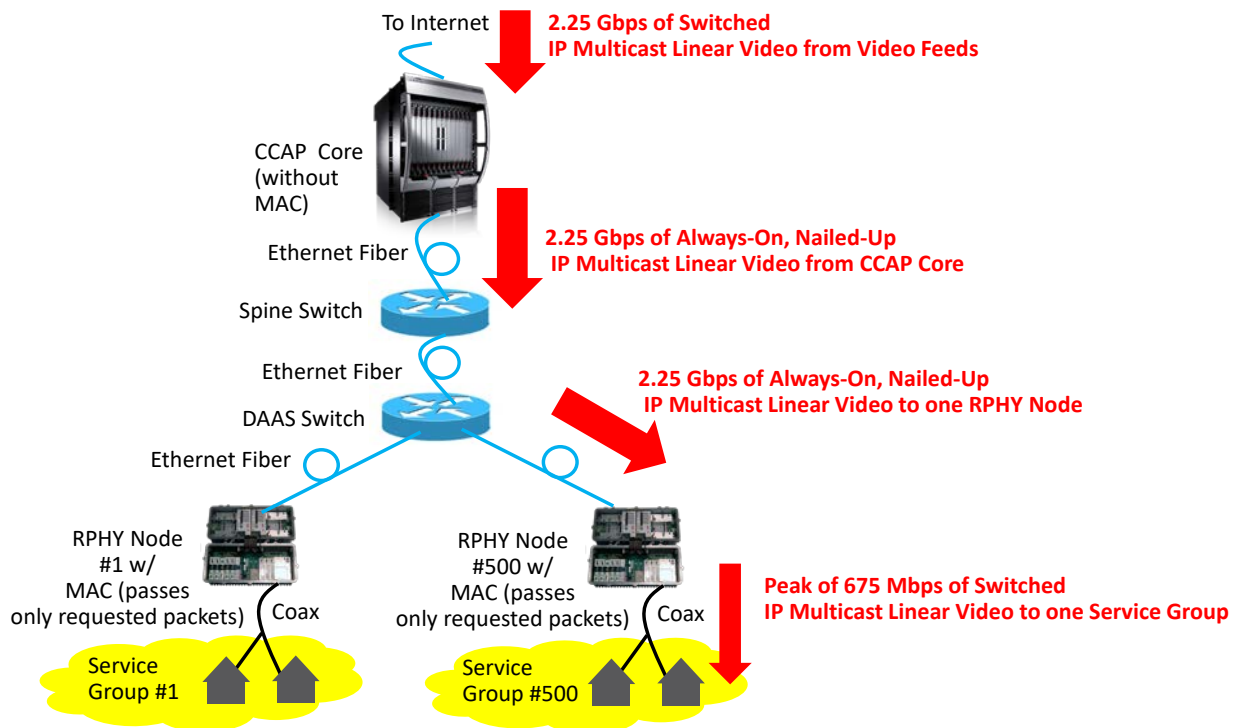


Figure 17 - Linear IP Video Delivery Example using MAC Processing in the Fiber Node

- Using a combination of Always-On, Nailed-Up IP Multicasting for the most popular portion of the Linear video content library along with Switched IP Multicasting for the low-popularity remainder of only the actively-viewed programs. These two sets of Linear video streams would be transmitted from the CCAP Core to each Remote PHY Fiber Node. A single multiplex of the most popular programs would be sent from the CCAP Core and replicated by the Spine/DAAS switch network to go to all 500 of the Remote PHY Fiber Nodes. Then a separate multiplex of low-popularity Switched IP Multicast Linear video programs would be sent separately to each of the 500 Remote PHY Fiber Nodes. This gives a nice compromise solution.

As an example, if the 60 most-popular programs were included within the Always-On, Nailed-Up IP Multicast multiplex, that would imply that the CCAP Core would generate a single multiplex of $(60) \times (9 \text{ Mbps}) = 540 \text{ Mbps}$ containing the 60 most-popular programs. This could be carried in a channel set and pseudo-wires comprising $540 \text{ Mbps} / 36 \text{ Mbps} = 15 \text{ SC-QAM}$ channels. The remainder of the viewed Linear programs associated with a particular Fiber Node would consist of $\sim(675 \text{ Mbps} - 540 \text{ Mbps}) = 135 \text{ Mbps}$ of low-popularity programs that could be sent directly to a particular Fiber Node in a Switched IP Multicast feed on a channel set and pseudo-wires containing $135 \text{ Mbps} / 36 \text{ Mbps} = \sim 4$ channels. Thus, a total of $\sim 675 \text{ Mbps}$ and $\sim 19 \text{ SC-QAM}$ channels and pseudo-wires would be used to transmit the Linear video programs to a particular Remote PHY Fiber Node. OFDM channels could also be used.

At the CCAP Core, a total of $(540 \text{ Mbps} + 500 \times 135 \text{ Mbps}) = 68.04 \text{ Gbps}$ of Linear video content must be processed and transmitted out of the CCAP Core. (See Figure 18). This represents a great reduction in the CCAP Core processing requirements relative to the previous solution. It should be clear that the packets for the low-popularity programs can be multiplexed in with regular High-Speed Data packets on a channel set, but the high-popularity programs must be

isolated to a separate channel set to ensure that the IP packets within that channel set are the same for every Fiber Node.

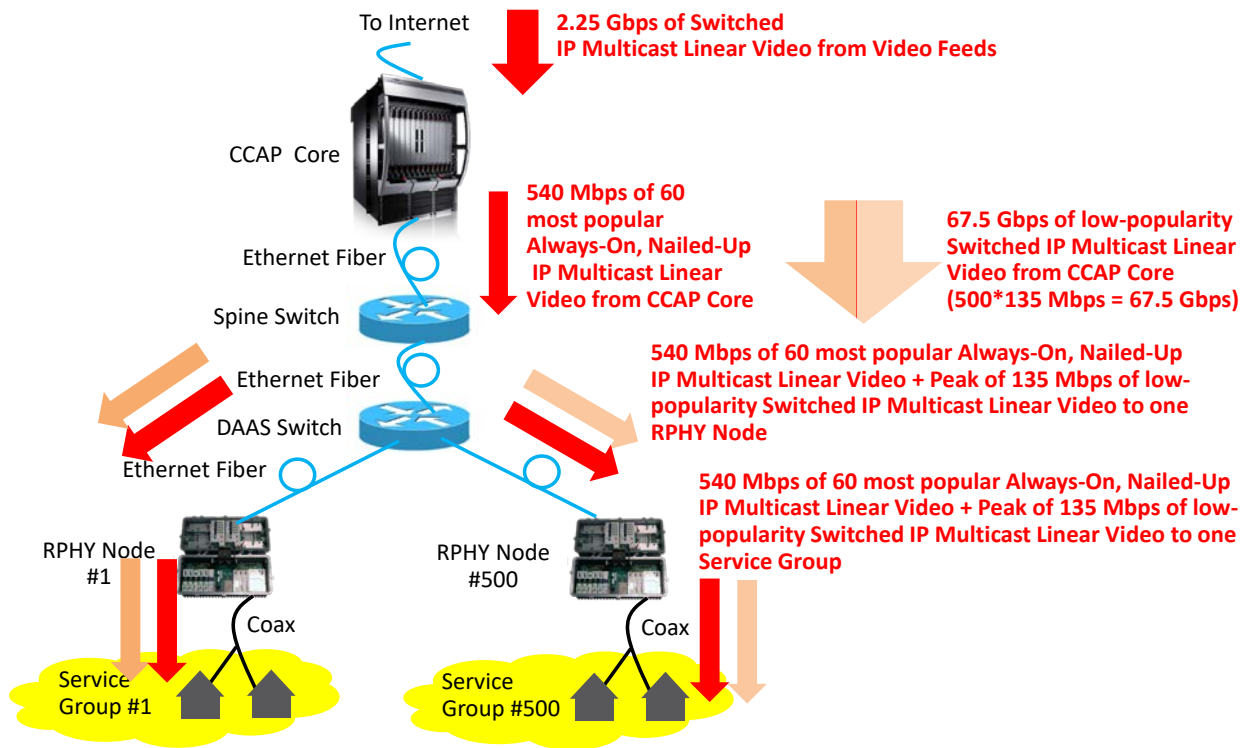


Figure 18 - Linear IP Video Delivery Example using a Combination of Always-On, Nailed-Up IP Multicasting & Switched IP Multicasting

Conclusion

Within this paper, the authors have studied several issues that may need to be addressed as operators move into a future world supporting 1+ Gbps service tiers. The areas that were studied included Traffic Engineering, Utilization Levels, TCP Performance Levels, Symmetrical Services, and Remote PHY IP Linear Video Delivery. Issues were identified, and potential solutions to the issues were proposed. At a high-level, the transition into the 1+ Gbps service environment will require adjustments, but most of the issues seem to have reasonable solutions.

Abbreviations

Avg	Average
BW	Bandwidth
CAGR	Compound Annual Growth Rate
CCAP	Converged Cable Access Architecture
DAAS	Distributed Access Architecture Switch
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable System Interface Specification
DS	Downstream
FDX	Full Duplex DOCSIS
Freq	Frequency
Gbps	Giga bits per second
GHz	Gigahertz (one billion Hertz)
HFC	Hybrid Fiber Coax
HD	High Definition
Hz	Hertz
IP	Internet Protocol
MAC	Media Access Control
Mbps	Megabits per second
MHz	Megahertz (one million Hertz)
MSO	Multiple System Operator
MSS	Maximum Segment Size
OFDM	Orthogonal Frequency Division Multiplexing
PER	Packet Error Ratio
PHY	Physical Interface
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RPHY	Remote PHY
RTT	Round Trip Time
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SD	Standard Definition
SG	Service Group
SLA	Service Level Agreement
Tavg	Average Bandwidth Traffic Rate
TCP	Transmission Control Protocol
Tmax	Maximum Bandwidth Traffic Rate
Tmax_max	Maximum of all Maximum Bandwidth Traffic Rates
UEPI	Upstream External PHY Interface
US	Upstream

Bibliography & References

[CLO1] T. J. Cloonan et. al., “Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements,” in SCTE Cable Tec Expo '14, September 2014, Society of Cable Telecommunications Engineers.

[CLO2] T. J. Cloonan et. al., “Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond,” in SCTE Spring Technical Forum '16, 2016, Society of Cable Telecommunications Engineers.

[RPH1] Remote Downstream External PHY Interface Specification CM-SP-DEPI-I06-170111, CableLabs.

[RHE1] I.Rhee, L.Xu, “Cubic: A New TCP-Friendly High-Speed TCP Variant”,
<http://www4.ncsu.edu/~rhee/export/bitcp/cubic-paper.pdf>.

[WEI1] D. X. Wei, P. Cao, “NS-2 TCP-Linux: An NS-2 TCP Implementation with Congestion Control Algorithms from Linux”, in proceedings of [ValueTool'06 -- Workshop of NS-2](#), Oct, 2006.

Accurately Estimating D3.1 Channel Capacity

A Technical Paper prepared for SCTE•ISBE by

Karthik Sundaresan
Principal Architect
CableLabs
858 Coal Creek Drive
303-661-9100
k.sundaresan@cablelabs.com

Introduction

The DOCSIS 3.1 specification fundamentally changes the nature of information delivery across the cable plant, and how the cable plant will be maintained and managed. The modulation order and FEC can be optimized based on actual plant conditions at individual devices. Devices which receive clean signals will utilize very efficient high-order-modulation across each of the subcarriers, devices with a degraded signal will use more robust modulation, all on the same channel. To manage this optimization the CMTS uses the concept of downstream OFDM profiles and upstream OFDMA profiles. D3.1 allows defining multiple profiles, each tuned to account for plant conditions experienced by a set of CMs.

Estimating downstream and upstream channel capacity was relatively straightforward for SC-QAM channels. But in D3.1, since the modulation orders of each subcarrier could be different and different across profiles, estimating the DOCSIS channel capacity is no longer simple. With multiple modulation profiles in use simultaneously, the capacity of the channel as seen by each CM may change instantaneously. The aggregate channel capacity calculation from a CMTS point of view with the different CMs using different profiles becomes more complicated and varies with which CMs and profiles are in use. How does one account for the NCP, FEC, PHY, MAC layer overhead and other variable factors in determining the effective channel throughput? How does profile definition, number of CMs, and heavy vs light traffic users on the channel affect throughput? The channel capacity affects how many subscribers can be assigned to use the same set of channels. It also affects traffic engineering and when an operator would need to split the node to increase available capacity. What is the reliable method for an operator to get a handle on the network capacity?

This paper will present a framework to calculate the D3.1 downstream and upstream channel capacity accurately and answer the above questions and considerations.

Downstream Capacity calculations

DOCSIS 3.1 introduces Orthogonal Frequency Division Multiplexing (OFDM) downstream signals and Orthogonal Frequency Division Multiple Access (OFDMA) upstream signals to achieve robust operation and provide more efficient use of the spectrum than previous DOCSIS versions.

The DOCSIS 3.1 system will have options of several split configurations that can be exercised based on traffic demand, services offered and the capability of the cable plant.

1. Band edges

DOCSIS 3.1 uses OFDM for downstream modulation. In the downstream direction, the cable system is assumed to have a pass band with a lower edge of either 54 MHz, 87.5 MHz, 108 MHz or 258 MHz, and an upper edge that is implementation-dependent but is typically in the range of 550 to 1002 MHz. Upper frequency edges extending to 1218 MHz, 1794 MHz and others are expected in future migrations of the plants. Within that pass band, analog or digital television signals in 6 MHz channels are assumed present on the standard, as well as other narrowband and wideband digital signals.

The CM supports a minimum of two independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream. The demodulator in the CM supports receiving downstream transmissions up to at least 1.218 GHz and optionally support receiving downstream transmissions up to at least one or more of the following downstream upper band edges: 1.002 GHz, 1.218 GHz, 1.794 GHz.

1.1. Downstream Subcarriers

The OFDM downstream multicarrier system is composed of a large number of subcarriers that have either 25 kHz or 50 kHz spacing. These subcarriers are grouped into independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream, totaling up to 7680 25 kHz subcarriers or up to 3840 50 kHz subcarriers; of which up to 7600 (25 kHz) or up to 3800 (50 kHz) active subcarriers spanning 190 MHz.

The encompassed spectrum of a 192 MHz downstream OFDM channel does not exceed 190 MHz. Therefore, the number of contiguous active subcarriers in a downstream OFDM channel does not exceed 3800 for 4K FFT and 7600 for 8K FFT. When configured for 4K FFT, the CMTS only uses subcarriers in the range $148 \leq k \leq 3947$, where k is the spectral index of the subcarrier in the IDFT equation defining the OFDM signal. When configured for 8K FFT, the CMTS only uses subcarriers in the range $296 \leq k \leq 7895$.

For D3.1 OFDM Channels there is at least 1 MHz of exclusion band between the spectral edge of a legacy SC-QAM channel and the center frequency of the nearest OFDM subcarrier and at least a 2 MHz exclusion band between any two adjacent asynchronous OFDM channels.

1.2. Symbol sizes

Table 1 - D3.1 Downstream OFDM Parameters

Parameter	Value	
Downstream master clock frequency	10.24 MHz	
Downstream Sampling Rate (fs)	204.8 MHz	
Downstream Elementary Period (Tsd)	1/(204.8 MHz)	
Channel bandwidths	24 MHz ... 192 MHz	
IDFT size	4096	8192
Subcarrier spacing	50 kHz	25 kHz
FFT duration (Useful symbol duration) (Tu)	20 μ s	40 μ s
Maximum number of active subcarriers in signal (192 MHz channel -190 MHz used)	3800	7600
Maximum spacing between first and last active subcarrier	190 MHz	

1.3. Cyclic Prefix

A segment at the end of the IDFT output is prepended to the IDFT, and this is referred to as the Cyclic Prefix (CP) of the OFDM symbol. There are five possible values for the length of the CP for the D3.1 Downstream and the choice depends on the delay spread of the channel, a longer delay spread requires a longer cyclic prefix.




Figure 1 - Cyclic Prefix Concept

Table 2 - D3.1 Downstream OFDM Parameters

Cyclic Prefix (μ s)	CP samples
0.9375 μ s	(192 * Tsd)
1.25 μ s	(256 * Tsd)
2.5 μ s	(512 * Tsd)
3.75 μ s	(768 * Tsd)
5 μ s	(1024 * Tsd)

1.4. Modulation Orders

The D3.1 Downstream modulation orders supported are 16-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM and optionally support 8192-QAM and 16384-QAM.

1.5. Pilots & PLC

Continuous pilots are pilots that occur at the same frequency location in every OFDM symbol. The PHY Link Channel (PLC) resides in a contiguous set of subcarriers in the OFDM channel. These subcarriers occupy the same spectral locations in every symbol. Continuous and scattered pilots change from symbol to symbol. Since overlapping pilots are treated as continuous pilots, the number of scattered pilots' changes from symbol to symbol. For bit loading, continuous pilots and the PLC are treated in the same manner as excluded subcarriers

The PLC is comprised of 8 OFDM subcarriers in every OFDM symbol when using 4K FFT OFDM (i.e., a subcarrier spacing of 50 kHz) and 16 OFDM subcarriers when using 8K FFT OFDM (i.e., a subcarrier spacing of 25 kHz).

Using the following notation

N: The total number of subcarriers in the OFDM symbol, equaling either 4096 or 8192

NC: The number of continuous pilots in an OFDM symbol

NS: The number of scattered pilots in an OFDM symbol

NE: The number of excluded subcarriers in an OFDM symbol

NP: The number of PLC subcarriers in an OFDM symbol

ND: The number of data subcarriers in an OFDM symbol

The values of *N*, *NC*, *NE* and *NP* do not change from symbol to symbol for a given OFDM template; the values of *NS* and *ND* change from symbol to symbol. The following equation holds for all symbols:

$$N = NC + NS + NE + NP + ND \quad \text{---eqn(1)} \quad -$$

The value of *N* is 4096 for 50 kHz subcarrier spacing and 8192 for 25 kHz subcarrier spacing. (*NS* + *ND*) is a constant for a given OFDM template. Therefore, although the number of data subcarriers (*ND*) and the number of scattered pilots (*NS*) in an OFDM symbol changes from symbol to symbol, the sum of these two numbers is invariant over all symbols.

There are eight predefined continuous pilots around the PLC. In addition, there are a number of continuous pilots distributed as uniformly as possible over the entire OFDM spectrum. The number of Continuous Pilots is given using the following formula:

$$\text{Number of Continuous Pilots} = \min(\max(8, \text{ceil}(M * (F_{\max} - F_{\min}) / 190), 120) \quad \text{---eqn(2)}$$

F_{max} refers to frequency in MHz of the highest frequency active subcarrier and *F_{min}* refers to frequency in MHz of the lowest frequency active subcarrier of the OFDM channel. Per the equation, the number of continuous pilots is linearly proportional to the frequency range of the OFDM channel. The minimum number of continuous pilots defined using the PLC cannot be less than 8, and the maximum number of continuous pilots defined using the PLC cannot exceed 120. Therefore, the total number of continuous pilots, including the predefined ones(PLC), will be in the range 16 to 128, both inclusive, see eqn(2). The value of *M* in equation as a parameter that can be adjusted by the CMTS with $120 \geq M \geq 48$, the typical value proposed for *M* is 48.

The main purpose for scattered pilots is the estimation of the channel characteristics for the purpose of equalization. There are two scattered pilot patterns, one for 4K FFT and one for 8K FFT. Although these pilots occur at different frequency locations in different OFDM symbols, the patterns repeat after every

128 OFDM symbols; in other words, the scattered pilot pattern has a periodicity of 128 OFDM symbols along the time dimension.

1.6. DS Capacity Calculations

Below are some snippets of Python code which was used in the capacity calculation for the DOCSIS 3.1 channels. The code shows how to calculate the Number of Effective subcarriers in a channel, given certain input settings.

```

DSUpperBandEdge      = DSLowerBandEdge + DSOccupiedSpectrum
DSNumFFTPoints       = (DSSamplingRate * 1000) / DSSubcarrierSpacing
DSSymbolPeriod_usec  = 1000 / DSSubcarrierSpacing
DSCyclicPrefix_usec  = DSCyclicPrefix / DSSamplingRate
DSActualSymbolPeriod_usec = DSSymbolPeriod_usec + DSCyclicPrefix_usec
DSSymbolEfficiency    = 100 * DSSymbolPeriod_usec / DSActualSymbolPeriod_usec
DSModulatedSubcarriers = (DSOccupiedSpectrum - DSGuardBand - DSExcludedBand)
                                                                * 1000 / DSSubcarrierSpacing

if DSSubcarrierSpacing == 50:
    DSNumPLCSubcarriers = 8
else:
    DSNumPLCSubcarriers = 16
DSNumContPilots = min(max(8, ceil(DSPilotDensity_M * DSOccupiedSpectrum / 190)),
                                                                120) + 8

DSNumScatteredPilots = ceil((DSModulatedSubcarriers - DSNumPLCSubcarriers) / 128)
DSEffectiveSubcarriers = DSModulatedSubcarriers - (DSExcludedSubcarriers +
                                                                DSNumPLCSubcarriers * DSNumFFTBlocks +
                                                                DSNumContPilots + DSNumScatteredPilots)

```

2. Downstream FEC

A Downstream codeword (FEC Frame) will be of the size 16,200 bits; this is using the Low Density Parity Codes (LDPC) at a code rate of 8/9. In a D3.1 OFDM channel an FEC Codeword can include mixed modulation. (A mixed-modulation codeword belongs to a profile that does not use the same modulation constellation for all subcarriers of the OFDM symbol.) Codeword shortening is also supported.

Table 3 - FEC coding parameters

LDPC Code Rate	BCH Uncoded Block Size	BCH Coded	LDPC Uncoded	LDPC Coded Block
8/9	14,232	14,400	14,400	16,200

The downstream LDPC codeword shown in Figure 2 is referred to as (16200, 14400). This means that a full codeword is 2025 bytes (16200 bits) that are divided into 225 bytes (1800 bits) of parity and 1800 bytes (14400 bits) of LDPC payload. That payload is further divided into 21 bytes (168 bits) of BCH parity, a 2 byte fixed header, and a variable 1777 byte maximum payload for DOCSIS frames. When the FEC codeword is shortened, only the DOCSIS payload shrinks. All other fields remain the same size.

Codeword shortening is used for creating shortened codewords when there is insufficient data to fill complete codewords or to achieve strong burst noise protection. Codeword shortening is accomplished by

shortening the uncoded block size (14232). Note that the number of parity bits remains the same; there is no shortening of the parity bits either in the BCH or in the LDPC. When a shortened codeword is needed, the CMTS prepends zero bits to the data, do the BCH encoding and LDPC encoding and removes the appended zero bits.




Figure 2 - DS Downstream Codeword header and payload

In the CW Header, the PDU Pointer (11 bits) points to the first byte of the first DOCSIS frame that starts in the payload. A value of zero points to the byte immediately following the codeword header.

2.1. Calculated Number of DS FEC codewords

The CMTS can make a choice to service traffic from each Downstream OFDM Profile in a round-robin fashion. It can do this by sending packets destined to CMs assigned to each profile during a specific window. This window can be on a symbol to symbol basis. i.e. say if there are 4 active profiles: A, B, C, D, the CMTS can choose to send traffic to profile A on a certain symbol and profile B, C, D on each subsequent symbol and cycle through again. It can also decide to make that window larger and choose to service each profile for a longer period of time say 2, 4, or 8 symbols at a time. The CMTS can also send codewords from multiple profiles on the same symbol, this case has not been explicitly calculated in this paper, as the main goal is to try and calculate effective throughput under fully loaded conditions. However, in this case one can surmise that switching between profiles can drop the FEC efficiency further, as there may be present multiple shortened codewords.

The following graphs show the number of full(long) & shortened codewords which will fit into DS symbols, if the CMTS chooses 1, 2, 4, or 8 symbols for packets from each profile. The graphs also show the variation of the number of codewords across increasing Channel Bandwidth. Additionally, it also shows the PHY efficiency of those channel across a symbol.




Figure 3 - Number of DS FEC Codewords per symbol vs Channel BW (1 symbol per profile)




Figure 4 - Number of DS FEC Codewords per symbol, vs Channel BW (2 symbols per profile)




Figure 5 - Number of DS FEC Codewords per symbol, across Channel BW (4 symbols per profile)




Figure 6 - Number of DS FEC Codewords per symbol, across Channel BW (8 symbols per profile)

3. Downstream Rate and efficiency

3.1. Coverage Layer Overhead

Detecting where the next codeword begins in an OFDM symbol can be difficult because more than one codeword may map into one OFDM symbol, the number of codewords per OFDM symbol may not be an integer, a codeword can overflow from one OFDM symbol to another, and the codeword could be shortened. Therefore, the transmitter must convey to the receiver all of the locations where a new codeword begins within an OFDM symbol. When the data codewords are mapped to subcarriers within a symbol, a pointer is needed to identify where a data codeword starts. This is known as the Next Codeword Pointer (NCP).

There are a variable number of NCP message blocks (MBs) on each OFDM symbol. To make sure that all subcarriers are used without reserving empty NCP MBs, the mapping of the NCP occurs in the opposite direction of the mapping for data. The relationship of NCP message blocks to the data channel is shown in Figure 7 (the last NCP MB is always a CRC-MB). These Next Codeword Pointers (NCPs) are encoded using the NCP FEC encoder, as defined in [D3.1 PHY] and are appended to the OFDM symbols. NCP subcarriers are modulated using QPSK, 16-QAM, or 64-QAM and this modulation is signaled by the PLC.

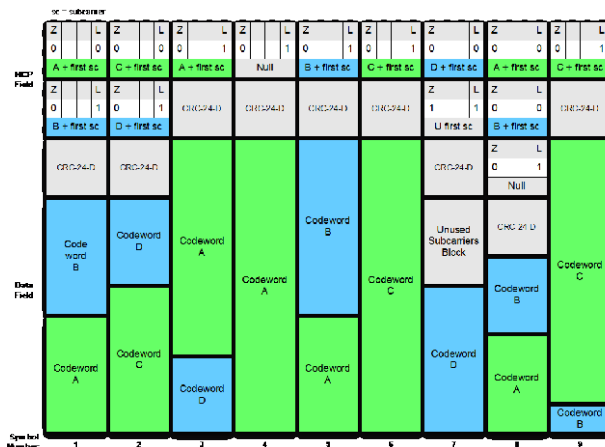


Figure 7 - NCP Examples

3.2. MAC Layer Management Messaging Overhead

The below table is an estimate of the packet sizes and data rate consumed by some of the common periodic MAC Management Messages (MMM): Upstream Channel Descriptor (UCD), MAC Domain Descriptor (MDD), OFDM Channel Descriptor (OCD), Downstream Profile Descriptor (DPD), Upstream Bandwidth Allocation Map (MAP), Time Synchronization (SYNC)

Table 4 - D3.1 Downstream MMM Overhead

	DPD	OCD <i>(On PLC only)</i>	MDD	UCD	MAP	Sync
Total bits (Average msg size * number of messages)	1156 (289*4 profiles)	306 (306 * 1 OFDM channel)	1210	459	594 (66 * 9) (8 SC QAM and 1 OFDMA Channel)	34
Periodicity in milliseconds	500	200	2000	2000	2	200
Average Data rate (Bits per second)	18,496	12,240	4,840	1,836	2,376,000	1,360

Based on this table we can see that MDDs, DPDs, OCDs, UCDs, Syncs, with above average sizes of messages only add up to ~38 kbps. With the addition of MAPs, the MMM overhead adds up to 2.4 Mbps.

Give that in a D3.1 OFDM channel we are talking of channel capacity from ~200 Mbps(24 MHz) to upwards of 1.7 Gbps (192 MHz), the MMM overhead may be considered insignificant in the data Rate Calculation.

3.3. Calculations for Downstream rate & PHY efficiency

Below are some snippets of Python code which was used in the capacity calculation for the DOCSIS 3.1 channels. The code shows how to calculate the data rate of a channel (assuming one profile) and the downstream PHY efficiency.

This step calculates the number of subcarriers consumed by each NCP message block.

```
NCPBitsperMB = 48
SubcarriersPerNCPMB = NCPBitsperMB/NCPModulationOrder
```

This step calculates the number of data bits available in each OFDM DS Symbol. Also if the CMTS services the profile for multiple symbols, then the number of available bits are multiplied by the number of symbols.

```
NumBitsinDataSubcarriers = DSEffectiveSubcarriers * DSAvgModulationOrder
if numSymbolsPerProfile > 1:
    NumBitsinDataSubcarriers = NumBitsinDataSubcarriers * numSymbolsPerProfile
```

This step calculates the number of full codewords available on the DS symbols for that burst of traffic on a particular profile (one or more OFDM DS Symbol, see above step).

```
# DSLDPC_FEC_CW = CWSize, Infobits, Parity, BCH, CWheader
DSLDPFEC_CW = [16200, 14216, 1800, 168, 16]
```

```
NumFullCodewords = math.floor(NumBitsinDataSubcarriers/DSLDPFEC_CW[0])
```

This step calculates the length of a shortened codeword if possible. It does this by finding the number of bits left in the symbol, by taking away the number of bits needed for the NCP MBs for those full Codewords and the CRC NCP MB,

```
NumNCPMBs = NumFullCodewords + ceil(1*numSymbolsPerProfile)
```

```
EstimateShortenedCWSize =  
((numSymbolsPerProfile * DSEffectiveSubcarriers -  
((NumNCPMBs + 1) * SubcarriersPerNCPMB)) * DSAvgModulationOrder) -  
(DSLDPF_FEC_CW[0] * NumFullCodewords)
```

```
if EstimateShortenedCWSize > 0:  
    ShortenedCWData=  
    EstimateShortenedCWSize - (DSLDPF_FEC_CW[2] -DSLDPF_FEC_CW[3]-DSLDPF_FEC_CW[4])  
else:  
    ShortenedCWData = 0
```

This step calculates the number of data bits per symbol and then uses the actual symbol period to calculate the effective data rate of the channel. The Downstream PHY efficiency is calculated by dividing by the Downstream occupied spectrum.

```
totalDataBits = (NumFullCodewords*DSLDPF_FEC_CW[1]) + ShortenedCWData  
RateacrossWholeChannelGbps = totalDataBits/  
    (DSActualSymbolPeriod_usec*numSymbolsPerProfile*1000)  
DSPHYEfficiency = RateacrossWholeChannelGbps*pow(10,3)/(DSOccupiedSpectrum)
```

The graphs in Figure 8, Figure 9 show the effective data rates of the OFDM channel vs. increasing channel bandwidth. Each curve indicates the number of continuous symbols the CMTS uses to send data for a particular profile. The red curve is 1 DS symbol per profile, blue is 4 DS symbols per profile, and green is 8 DS symbols per profile.




Figure 8 – DS Downstream Channel capacity wrt Bandwidth (1 profile)




Figure 9 – DS Downstream PHY Efficiency wrt Bandwidth (1 profile)

The graph in Figure 10, below shows the effective data rates and the PHY Efficiency of a DS OFDM channel (192 MHz) vs increasing Cyclic prefix size. The graph in Figure 11, shows the effective data

rates & the PHY Efficiency of the OFDM channel(192 MHz, 512 CP), vs increasing modulation order for the profile.




Figure 10 - D3.1 Downstream Rates and DS PHY Efficiency vs Cyclic Prefix size




Figure 11 - D3.1 Downstream Rates and DS PHY Efficiency vs Modulation order

The graph in Figure 12 shows the effective data rates of the OFDM channel, vs increasing Modulation order for the profile, and changing cyclic prefix values (each curve is a different CP setting, increasing CP reduces the effective data rate)




Figure 12 - D3.1 Downstream Rates vs Modulation order

4. Multiple profiles

The DOCSIS 3.1 OFDM introduces multiple modulation orders within a channel. Each of the subcarriers in the downstream OFDM channel can be configured to use a different modulation order. This allows the CMTS to optimize the downstream transmissions across the wide frequency band (192 MHz) of the channel. The specific choice of modulation order for each subcarrier is communicated to the CMs in the form of a downstream modulation profile, which allows the CMs to interpret and demodulate the signal.

A modulation profile consists of a vector of bit-loading values, an integer value for each active subcarrier in the downstream channel. Since the modulation orders range from 16-QAM to 16384-QAM, the range of bit-loading values is from 4 to 14 (skipping 5). However, it is expected that very low bit-loading values, 7 or less, will be used very infrequently since most plants support 256 QAM today.

The CMTS uses a “Profile A” that is the lowest common denominator profile, i.e. Profile A is able to be successfully received by all CMs in the Service Group. Profile A is essentially the lowest modulation across all CMs and subcarriers in the channel. It can then generate up to 15 additional modulation profiles, which are communicated to the Service Group. Each CM can be assigned up to four modulation profiles, including Profile A (used for broadcast frames), an optimized profile for the CM’s unicast traffic, and possibly two additional profiles that could be used for multicast traffic. Since the number of CMs in the service group, using that Downstream channel is expected to be larger than 15 in the majority of cases, each profile is expected to be used by a group of CMs that have similar channel characteristics.

Calculating the effective channel capacity when there are multiple profiles active on a CMTS is complicated. The amount of data traffic used by various CMs can vastly differ based on the various subscribers. Each CM can be assigned a different set of profiles that are optimized for it. By combining the profile information and CM data usage one can start estimating the current effective channel capacity. As the data usage of the various CMs changes, the data sent on each of the profiles changes, effectively changing the data rate of channel. This can occur on a symbol to symbol basis. This would probably be best modeled and estimated using a Monte Carlo simulation. That effort could be the topic of another paper.

For simpler estimates which will help us approximate the data rate capacity, one can average out the profile modulation order and weight it across the number of CMs assigned to a particular profile to work out an approximate channel capacity. This assumes each CM is receiving the same amount of traffic. One could also calculate this average with different assumptions on traffic from each CM, their associated profiles and weighting it appropriately.

Using the Profile Management Algorithms defined in [PMA-2016], one can find an optimal set of profiles for use on a channel for a set of CMs. The data one needs to have is the RxMER per subcarrier across the channel for each CM. Based on a set of RxMER Data observed in CMs in a D3.1 trial below are some of the potential profile definitions created by using the [PMA-2016] algorithms.

This figure shows the profile definition of 8 different DS profiles (7 individual profiles + Profile A which is the least common denominator profile), for a group of 121 CMs. It shows the number of subcarriers in each profile that have a specific bit loading e.g. Profile 2 (green) has 27 CMs assigned to it, and has an average Modulation Order (bit-loading) of 11.98, or 4096 QAM across most subcarriers and a few subcarriers at a level below that. Other profiles are also a combination of different modulation orders on different subcarriers.




Figure 13 - Sample set of D3.1 Profile definitions

Now to estimate the D3.1 capacity, assuming all CMs receive equal amounts of traffic, the average modulation order across all the CMs would end up being the weighted average of the profile average modulation order and the number of CMs using it. For the set of CMs/profiles in the Figure 13 that weighted average modulation order (bit-loading) would be 10.75 bits, which will yield a mean channel capacity ~1.6245 Gbps (with smallest CP value). The instantaneous capacity will vary depending on which set of CMs and profiles are active.

Upstream Capacity calculations

DOCSIS 3.1 uses OFDMA (orthogonal frequency-division multiple access) for upstream modulation. OFDMA is a multi-user version of OFDM, and assigns subsets of subcarriers to individual CMs. The upstream OFDMA parameters are derived from the downstream parameters, and described below

5. Band edges

In the upstream direction, the cable system may have a 5-42 MHz, 5-65 MHz, 5-85 MHz, 5-117, 5-204 MHz or pass bands with an upper band edge beyond 204 MHz. Analog and digital television signals in 6 MHz channels may be present, as well as other signals. A D3.1 CM supports one or more of the following upstream upper band edges, (as long as one is 85 MHz or greater): 42 MHz; 65 MHz, 85 MHz, 117 MHz, and/or 204 MHz. DOCSIS 3.1 Network supports a minimum of two independently configurable OFDMA channels each occupying a spectrum of up to 95 MHz in the upstream. A DOCSIS 3.1 network is capable of receiving 192 MHz of upstream active channels when operating with the 204 MHz upstream upper band edge. In DOCSIS 3.1 upstream mode the CM is capable of transmitting OFDMA channels and legacy SC-QAM channels at the same time (as controlled by the CMTS). There are no legacy SC-QAM channels above a frequency of 85 MHz.

5.1. Upstream Subcarriers

The OFDMA upstream multicarrier system is also composed of either 25 kHz or 50 kHz subcarriers. In the upstream, the subcarriers are grouped into independently configurable OFDMA channels each of up to 95 MHz encompassed spectrum, totaling 3800 25 kHz spaced subcarriers or 1900 50 kHz spaced subcarriers or 1920 50 kHz spaced subcarriers. Many parameters of these channels can be independently configured thereby optimizing configuration based on channel conditions

The encompassed spectrum of an upstream OFDMA channel does not exceed 95 MHz. Therefore, the number of contiguous active subcarriers in an upstream OFDMA channel are 1900 for 2K FFT and 3800

for 4K FFT. When configured for 2K FFT, the CMTS only uses subcarriers in the range $74 \leq k \leq 1973$, where k is the spectral index of the subcarrier in the IDFT equation defining the OFDMA signal. When configured for 4K FFT, the CMTS MUST only use subcarriers in the range $148 \leq k \leq 3947$.

5.2. Symbol sizes

Table 5 - D3.1 Upstream OFDMA Parameters

Parameter	Value	
Upstream Sampling Rate (fsu)	102.4 MHz	
Upstream Elementary Period Rate (Tsu)	1/102.4 MHz	
Channel bandwidths	10 MHz, ..., 95 MHz	6.4 MHz, ..., 95 MHz
IDFT size	2048	4096
Subcarrier spacing	50 kHz	25 kHz
FFT duration (Useful symbol duration) (Tu)	20 μ s	40 μ s
Maximum number of active subcarriers in signal (for 95 MHz channel)	1900	3800

5.3. Cyclic Prefix

A segment at the end of the IFFT output is prepended to the IFFT output; this is referred to as the Cyclic Prefix (CP) of the OFDM symbol. The addition of a cyclic prefix enables the receiver to overcome the effects of inter-symbol-interference caused by micro-reflections in the channel.

Table 6 - D3.1 Upstream Cyclic Prefix Parameters

Cyclic Prefix (μ s)	CP samples
0.9375 μ s	$(96 * T_{su})$
1.25 μ s	$(128 * T_{su})$
1.5625 μ s	$(160 * T_{su})$
1.875 μ s	$(192 * T_{su})$
2.1875 μ s	$(224 * T_{su})$
2.5 μ s	$(256 * T_{su})$
2.8125 μ s	$(288 * T_{su})$
3.125 μ s	$(320 * T_{su})$
3.75 μ s	$(384 * T_{su})$
5.0 μ s	$(512 * T_{su})$
6.25 μ s	$(640 * T_{su})$

5.4. Upstream Modulation orders

The DOCSIS 3.1 network supports BPSK, QPSK, 8-QAM, 16-QAM, 32-QAM, 64-QAM, 128-QAM, 256-QAM, 512-QAM, 1024-QAM, 2048-QAM, and 4096-QAM for subcarriers of upstream OFDMA channels. BPSK is used for pilots and complementary pilots only, and not used for data subcarriers.

5.5. OFDMA Frames

Upstream transmission uses OFDMA frames. Each OFDMA frame is comprised of a configurable number of OFDM symbols, K . Several transmitters may share the same OFDMA frame by transmitting data and pilots on allocated subcarriers of the OFDMA frame.

There are several pilot patterns as described below. The structure of an OFDMA frame is depicted in Figure 14. The upstream spectrum is divided into groups of subcarriers called minislots. Minislots have dedicated subcarriers, all with the same modulation order ("bit loading"). A CM is allocated to transmit one or more minislots in a Transmission Burst. The modulation order of a minislot, as well as the pilot pattern to use may change between different transmission bursts (see Figure 15) and are determined by a transmission profile.




Figure 14 - Example Minislot Layout

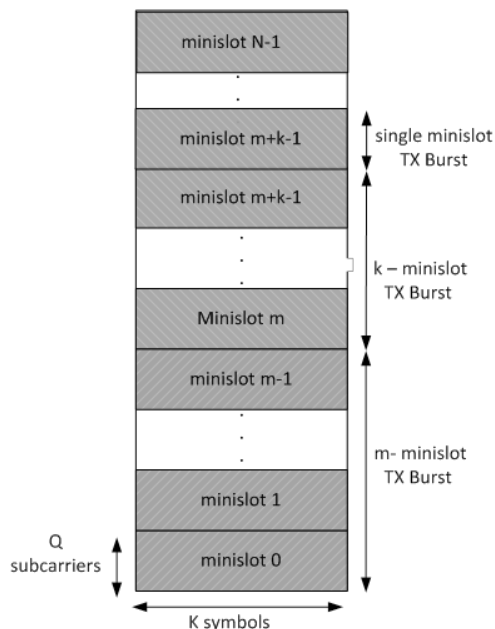


Figure 15 - Grants across Minislots / Tx Burst Layout




Figure 16 - Edge and Body Minislots in a Transmission Burst

5.6. Pilots & Minislots

DOCSIS 3.1 specifies two minislot sizes by specifying the number of subcarriers per minislot. There are 8- and 16-subcarrier minislots. A minislot is always 400 kHz wide (25KHz subcarrier *16, or 50KHz subcarrier *8). Two types of minislots are defined for each minislot size: edge minislots and body minislots. An edge minislot is the first minislot in a transmission burst, and body minislots are used for all other minislots in a transmission burst, see Figure 16. An edge minislot is used for the first minislot of an OFDMA frame that is not a zero valued minislot and also for the first minislot after an exclusion band or after one or more contiguous skipped subcarriers or after a zero valued minislot.

Each minislot is comprised of pilots(P), complementary pilots(CP), and data subcarriers. Pilots are used by the CMTS receiver to adapt to channel conditions and frequency offset. Pilots are subcarriers that do not carry data, Instead, it encodes a pre-defined BPSK symbol known to the receiver. DOCSIS 3.1 also specifies complementary pilots which are subcarriers that carry data, but with a lower modulation order than other data subcarriers in the minislot. If the modulation order used for data in the minislot is M, the complementary pilots are used with modulation order equal to the maximum between M-4 and 1 (BPSK). For example, if the bit loading in a minislot is 12, Complementary Pilots use 8 bits.

For each minislot size, seven pilot patterns are defined. Pilot patterns differ by the number of pilots in a minislot, and by their arrangement within the minislot. The different pilot patterns enable the CMTS to optimize its performance (physical layer rate and pilot overhead) according to different loop conditions and variations of SNR with frequency. Each pilot pattern defines edge and body minislots.

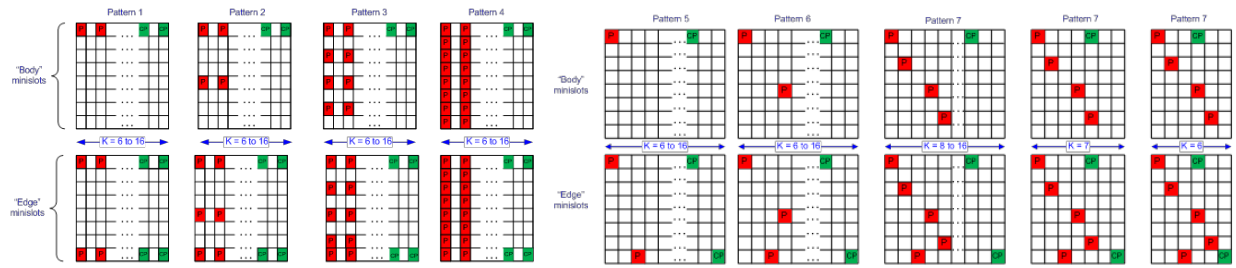


Figure 17 - Pilot Patterns 1-7 (8 subcarrier minislots)

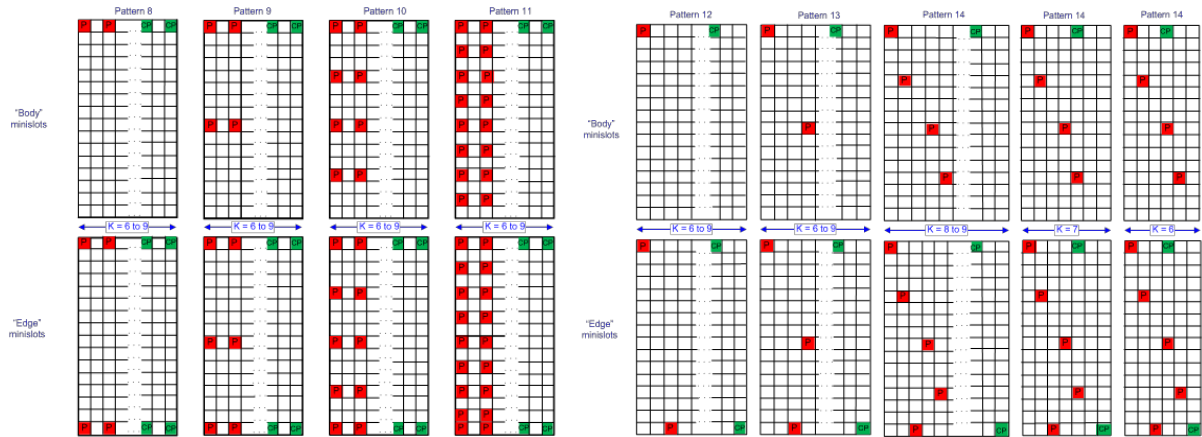


Figure 18 - Pilot Patterns 8-14 (16 subcarrier minislots)

Below are some snippets of Python code which was used in the capacity calculation for the DOCSIS 3.1 channels. The code shows how to calculate the data capacity of a minislot, given certain input settings, such as the OFDMA frame width in symbols (K), num of Subcarriers (Q), body vs Edge Mini slot, and the pilot pattern and complementary pilot pattern.

```
# Pattern Number, minislotsubcarriersQ, Body0/Edge1, Num Pilots, Num CP
minislotpatterns = [[1, 8, 0, 2, 2],
                    [2, 8, 0, 4, 2],
                    [3, 8, 0, 8, 2],
                    [4, 8, 0, 16, 2],
                    [5, 8, 0, 1, 1],
                    [6, 8, 0, 2, 1],
                    [7, 8, 0, 4, 1],
                    [1, 8, 1, 4, 4],
                    [2, 8, 1, 6, 4],
                    [3, 8, 1, 10, 4],
                    [4, 8, 1, 16, 4],
                    [5, 8, 1, 2, 2],
                    [6, 8, 1, 3, 2],
                    [7, 8, 1, 5, 2],
                    [8, 16, 0, 2, 2],
                    [9, 16, 0, 4, 2],
                    [10, 16, 0, 8, 2],
```

```

[11, 16, 0, 16, 2],
[12, 16, 0, 1, 1],
[13, 16, 0, 2, 1],
[14, 16, 0, 4, 1],
[8, 16, 1, 4, 4],
[9, 16, 1, 6, 4],
[10, 16, 1, 10, 4],
[11, 16, 1, 18, 4],
[12, 16, 1, 2, 2],
[13, 16, 1, 3, 2],
[14, 16, 1, 5, 2],
]

```

```

def MinislotCapacity(K, Modulationorder, pattern_index):
    """ function for the Minislot calculation """
    global minislotpatterns

    Q = minislotpatterns[pattern_index][1]

    cp_modulation_order = max(Modulationorder - 4, 1)

    subcarriers = K * Q - minislotpatterns[pattern_index][3]
                    - minislotpatterns[pattern_index][4]
    mscapacity = Modulationorder * subcarriers +
                  cp_modulation_order * minislotpatterns[pattern_index][4]

    return mscapacity

```

The following Figure 19 below show the minislot capacity across various parameters. The first sub-plot shows the variation of minislot capacity ($K=16$), across the different Pilot patterns. It shows 4 lines, two pairs for the 8 & 16 subcarrier minislots, each one denoting body and edge. The edge minislot is a smidge lesser in data capacity compared to the Body minislot

The second sub-plot shows the variation of the Minislot capacity ($Q=16$, pattern 13) as the K increases from 6 to 36 across, and the modulation order of the minislot increases from 2 to 12 bits bottom to top.




Figure 19 - MiniSlot Capacity : Across Pilot Patterns, Minislot Width(K symbols), and Modulation order

6. FEC

The grant indicates which minislots are assigned to a given burst and which upstream profile is to be used. The CM and CMTS use this information to determine the total number of bits in the grant which are available to be used for FEC information or parity.

Table 7 - FEC coding parameters

Code	LDPC Code Rate	Codeword size in bits (Ni)	Information bits (Ki)	Parity bits (Pi)
Long	$8/9 = 89\%$	16200	14400	1800
Medium	$28/33 = 85\%$	5940	5040	900
Short	$3/4 = 75\%$	1120	840	280

6.1. Shortened codewords

Based on the size of the grant obtained from the CMTS, The US FEC codeword creation algorithm follows the procedure as depicted in the Figure 20. The [D3.1 PHY] spec in section 7.4.3.1.1 FEC Codeword Selection Algorithm, shows an implementation of this procedure in Matlab code.

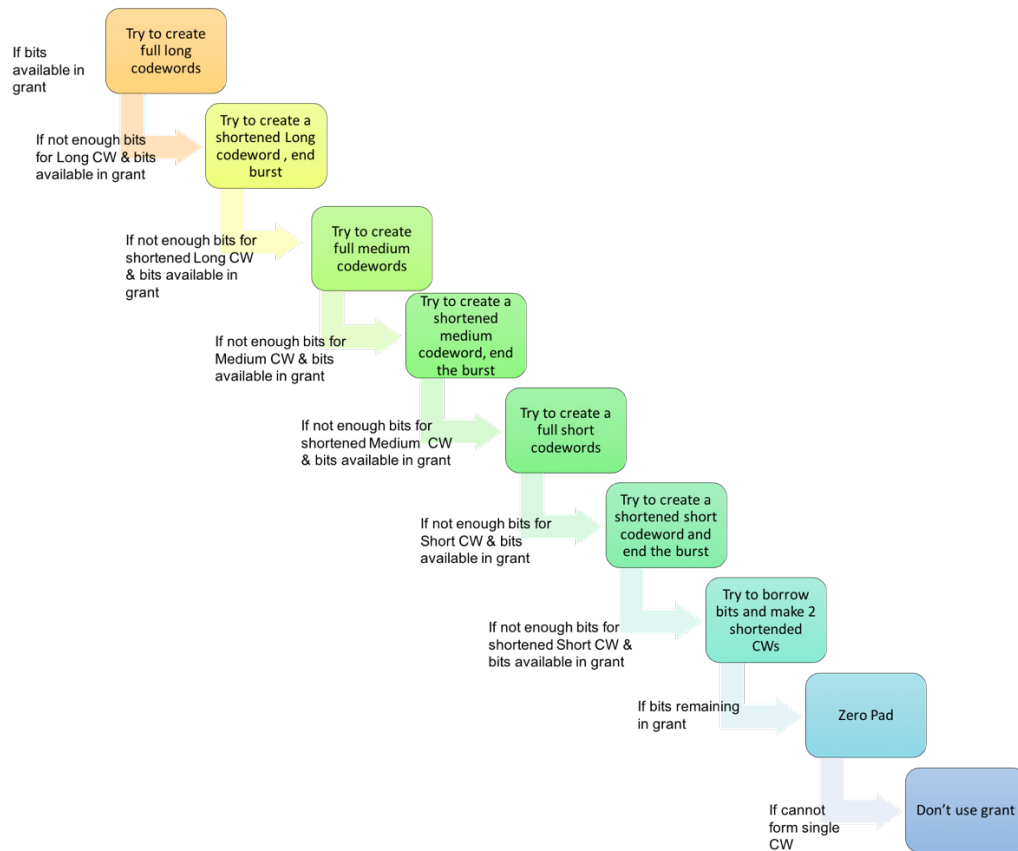


Figure 20 - US CW creation Algorithm

Based on the implementation of the above algorithm, the graphs in Figure 21, show the number of code words of each kind (Short, Shortened Short, Medium, Shortened Medium, Long, Shortened Long) in different colors, as the grant size (in bits) varies. The graph also shows the FEC efficiency of the grant, see Figure 22.




Figure 21 - Number of US FEC Codewords over grant sizes, FEC Efficiency




Figure 22 - Effective US FEC Efficiency vs grant size

7. Calculations for Upstream rate & PHY efficiency

Below are some snippets of Python code which was used in the capacity calculation for the DOCSIS 3.1 channels. The code shows how to calculate the data rate of a channel and the downstream PHY efficiency. The assumption here is an average modulation order calculated across all the minislots for the whole data profile or IUC (interval usage code).

This step calculates the active bandwidth, the symbol period, and the total frame duration.

```
USUpperBandEdge = USLowerBandEdge + USOccupiedSpectrum
USActiveBW      = USOccupiedSpectrum - USExcludedSpectrum

USCyclicPrefix_usec = USCyclicPrefix / USSamplingRate
USNumFFTPoints     = 1000 * USSamplingRate / USSubcarrierSpacing
UsSymbolPeriod_usec = 1000 / USSubcarrierSpacing
USActualSymbolPeriod_usec = UsSymbolPeriod_usec + USCyclicPrefix_usec
USSymbolEfficiency = 100 * UsSymbolPeriod_usec / USActualSymbolPeriod_usec
USFrameDuration_usec = USMinislotSymbolsK * USActualSymbolPeriod_usec
```

This step calculates the number of subcarriers in the Minislot (Q), and prepares to calculate the excluded subcarriers

```
if USSubcarrierSpacing == 25:
    USMinislotSubcarriersQ = 16
    kNBI = 3
else:
    USMinislotSubcarriersQ = 8
    kNBI = 2
```

This step calculates the number of excluded subcarriers, uses the exclusion zones, the number of grants, to figure out the number of minislots

```
UsTotalSubcarriers = 1000 * USOccupiedSpectrum / USSubcarrierSpacing
USExcludedSubcarriers = (1000 * (USExcludedSpectrum + USGuardBand) /
                        USSubcarrierSpacing) + (USExcludedNBI * kNBI)

UsNumofExclSpectrumGaps = USExcludedNBI + USNumContLegacy
USActualSignalSubcarriers = UsTotalSubcarriers - USExcludedSubcarriers

USTempNumMinislots = math.floor(USActualSignalSubcarriers /
                                USMinislotSubcarriersQ)
UsminislotEfficiency = (USActualSignalSubcarriers - UsNumofExclSpectrumGaps * 4) /
                        USActualSignalSubcarriers

USNumMinislots = round(UsminislotEfficiency * USTempNumMinislots)
USNumofEdgeMinislots = USNumGrantsInProfile + USAddnlEdgeMiniSlot
USNumofBodyMinislots = USNumMinislots - USNumofEdgeMinislots
```

This step calculates the appropriate pilot pattern, and then the total US capacity of an US OFDM frame, by calculating the minislot capacity for all the Body and Wdge Minislots.

```

local_USPilotPattern = USPilotPattern-1 # convert frm spec index to array index
if USMinislotSubcarriersQ == 16:
    local_USPilotPattern = local_USPilotPattern + 7

usCapacity = (
    (USNumofBodyMinislots * MinislotCapacity(USMinislotSymbolsK, USProfileModOrder,
        local_USPilotPattern)) +
    (USNumofEdgeMinislots * MinislotCapacity(USMinislotSymbolsK, USProfileModOrder,
        local_USPilotPattern + 7))
)

avgMinislotCapacity = usCapacity / USNumMinislots

```

This step calculates the channel rate using the US minislot capacity number and the actual symbol period, and the PHY efficiency with FEC overhead

```

ProfileRate_Mbps = (usCapacity /
                    (USMinislotSymbolsK * USActualSymbolPeriod_usec))

usPHYEfficiency = ProfileRate_Mbps / USActiveBW

usPHYEfficiency_w_fec_time_overhead = usPHYEfficiency *
    getCWSizesforEfficiency(usCapacity)[0] *
    (USSymbolEfficiency / 100)

```

The graphs in Figure 23, Figure 24, show the effective data rates and PHY Efficiency of the OFDMA channel vs. increasing channel bandwidth.




Figure 23 - US Channel Capacity across Channel BW




Figure 24 - US PHY Efficiency Bits per hz across Channel BW

The graph in Figure 25 shows the effective data rates and the PHY Efficiency of a US OFDM channel (96 MHz) vs increasing Cyclic prefix size.




Figure 25 - US PHY Efficiency Bits/s/Hz & Channel capacity vs CyclicPrefix us

The graph in Figure 26 below shows the effective data rates of the OFDMA channel, vs increasing Modulation order for the profile.




Figure 26 - US PHY Efficiency Bits/s/hz & Channel capacity vs Modulation order

The graph in Figure 27 below shows the effective data rates & the PHY Efficiency of a US OFDMA channel (96 MHz) vs increasing grant size. Each curve is for different values of K ranging from 6 to 36




Figure 27 - US PHY Efficiency Bits/s/hz & Channel capacity vs Number of grants

The graph in Figure 28 shows the effective data rates and the PHY Efficiency of a US OFDMA channel (96 MHz) vs increasing Cyclic prefix size.




Figure 28 - US PHY Efficiency Bits per hz across Channel BW

8. Multiple profiles

Just like in the DS OFDM channel with DS modulation profiles, the US OFDMA channel has support for multiple modulation orders within an IUC (US Profile). All

The DOCSIS 3.1 OFDMA introduces multiple modulation orders within an Upstream channel. Each of the minislots in the upstream OFDMA channel can be configured to use a different modulation order. (A minislot always uses the same modulation order.) This allows the CMTS to optimize the upstream transmissions across the wide frequency band (95 MHz) of the channel. The specific choice of modulation order selected for each minislot is communicated to the CMs in the form of an upstream IUC/modulation profile, which allows them to interpret and modulate the signal.

A modulation profile consists of a vector of bit-loading values, an integer value for each minislot in the upstream channel. Since the modulation orders range from QPSK to 4096-QAM, the range of bit-loading values is from 2 to 12; however, it is expected that very low bit-loading values, 4 or less, will be used very infrequently since most plants support 16 or 64 QAM today.

The CMTS uses IUC 13 which is the lowest common denominator profile, able to be successfully used by all CMs in the Service Group, mainly before registration. CMTS can define additional IUCs (Data Profile IUCs: 5, 6, 9, 10, 11, 12, and 13) which are communicated to the Service Group. Each CM can be assigned up to two IUCs, Since the number of CMs in the Service Group is expected to be larger than 6 in the majority of cases, each IUC is expected to be used by a group of CMs that have similar channel characteristics.

Now calculating the effective channel capacity when there are multiple IUCs active on a CMTS starts getting complicated. The amount of data traffic used by various CMs can vastly differ based on the subscribers. Each CM can be assigned a different set of IUCs, optimized for it. By combining the IUC information and CM data usage one can start estimating the current effective channel capacity. As the data usage of the various CMs changes, the data sent on each of the profiles changes, effectively changing the data rate of channel, on a frame to frame basis. Again, like in the downstream,, this would probably be best modeled and estimated using a Monte Carlo simulation.

For simpler estimates which will help us approximate the data rate capacity, one can average out the IUC modulation order and weight it across the number of CMs assigned to a particular IUC to work out an approximate channel capacity. This assumes each CM is transmitting the same amount of traffic. One could also weight this average with different assumptions on traffic from each CM and their IUCs

Conclusion

Estimating downstream and upstream channel capacity in D3.1 gets complicated since the modulation orders of each subcarrier could be different and different across profiles/IUCs. This paper presents the details on the elements/calculations needed to compute the effective channel capacity of a DS OFDM channel or an US OFDMA channel. For the downstream this includes accounting for pilots, NCPs, FEC overhead of Long or shortened codewords, etc. For the upstream this includes accounting for pilots, complementary pilot patterns, Minislot sizes, cyclic prefix, FEC overhead of short medium or long and their shortened versions etc. The paper presents in graphical form many of the expected capacity numbers for the D3.1 downstream and upstream channel

Default Parameters for calculations

[DownstreamChannel]

DSOccupiedSpectrum = 192
DSLLowerBandEdge = 678
DSGuardBand = 2
DSExcludedBand = 2
DSNumFFTBlocks = 1
DSAvgModulationOrder = 12
DSSamplingRate = 204.8
DSSubcarrierSpacing = 50
DSCyclicPrefix = 512
DSWindowing = 128
DSPilotDensity_M = 48
DSExcludedSubcarriers = 20
NCPModulationOrder = 6

[UpstreamChannel]

USOccupiedSpectrum = 96
USLowerBandEdge = 12
USSamplingRate = 102.4
USSubcarrierSpacing = 25
USPilotPattern = 8
USCyclicPrefix = 192
USWindowing = 128
USMinislotSymbolsK = 36
USNumContlegacy = 1
USExcludedSpectrum = 0
USGuardBand = 1
USExcludedNBI = 0
USAddnlEdgeMiniSlot = 0
USProfileModOrder = 10
USNumGrantsInProfile = 38

Abbreviations

bps	bits per second
CM	Cable Modem
CMTS	Cable Modem Termination System
CP	Cyclic Prefix
DOCSIS	Data-Over-Cable Service Interface Specifications
D3.1	DOCSIS version 3.1 specification, Networks, equipment
DS	Downstream
DPD	Downstream Profile Descriptor
IUC	Interval Usage Code (US Profile)
IDFT	Inverse Discrete Fourier Transform
FEC	forward error correction
FFT	Fast Fourier Transform
HFC	hybrid fiber-coax
Hz	hertz
LDPC	Low-Density Parity Check
MAC	Media Access Control Layer
MAP	Bandwidth Allocation Map, MAC Management Message that the CMTS uses to allocate transmission opportunities to CMs
MB	Message Block.
MDD	MAC Domain Descriptor
MMM	MAC Management Message
NCP	Next Codeword Pointer
OCD	OFDM Channel Descriptor
SC-QAM	Single Carrier Quadrature Amplitude Modulation
PHY	Physical Layer
PLC	PHY Link Channel
QAM	Quadrature Amplitude Modulation
UCD	Upstream Channel Descriptor
US	Upstream

Acknowledgements

Many thanks to Alberto Campos (CableLabs) on his guidance on understanding the complexities of the D3.1 PHY layer, Thanks to Greg White (CableLabs) on discussions on capacity calculations. Thanks to Jay Zhu (CableLabs) on helping with plotting graphs in Python. Many thanks to Paul Schauer (Comcast) for sharing data from their DOCSIS 3.1 field trials.

Bibliography & References

[D3.1 PHY] DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I11-170510, May 10, 2017, <https://apps.cablelabs.com/specification/CM-SP-PHYv3.1>

[D3.1 MULPI] DOCSIS 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I11-170510, May 10, 2017, Cable Television Laboratories, Inc.
<https://apps.cablelabs.com/specification/CM-SP-MULPIv3.1>

[PMA-2016] INTX 2016 DOCSIS 3.1 Profile Management Application and Algorithms (2016) By Greg White and Karthik Sundaresan, Cable Television Laboratories, Inc.
<http://www.nctatechnicalpapers.com/Paper/2016/2016-docsis-3-1-profile-management-application-and-algorithms>

Network Capacity and Machine Learning

A Technical Paper prepared for SCTE•ISBE by

Dr. Claudio Righetti

Chief Scientist & Security
Cablevisión S.A.
Gral. Hornos 690, Buenos Aires, Argentina
Phone: +5411 5530 4468
crighetti@cablevision.com.ar

Emilia Gibellini

Data Scientist
Cablevisión S.A.
egibellini@cablevision.com.ar

Florencia De Arca

Data Scientist
Cablevisión S.A.
fdearca@cablevision.com.ar

Carlos Germán Carreño Romano

Data Scientist
Cablevisión S.A.
caromano@cablevision.com.ar

Mariela Fiorenza

Data Scientist
Cablevisión S.A.
mafiorenza@cablevision.com.ar

Gabriel Carro

VP Engineer and R&D
Cablevisión S.A.
gcarro@cablevision.com.ar

Fernando Rodrigo Ochoa

Security Analyst
Cablevisión S.A.
fochoa@cablevision.com.ar

Abstract

The purpose of this paper is to introduce STEM-ML, an extension of our network-dimensioning tool, which allows us to define the strategy to face the increasing demand, of both our Internet broadband and “Flow”, our Internet Protocol Television (IPTV) services. This tool makes use of machine learning techniques to characterize the optical nodes that integrate our network. Based on such characterization, we can define the technologic and commercial strategy for the access network so that Cablevisión (CVA) is able to afford the short and long-term demand.

Until the development of STEM-ML, characterization was made at hub level, and it was based on the average bandwidth per subscriber parameters. With STEM-ML, the analysis is made at optical node level, and monthly consumption, households passed (HHP), and protocol types, among other variables. Moreover, data from “Flow” our IPTV platform is added. The increasing data volume generates the need for introducing machine learning and multivariate analysis techniques.

Content

1. Introduction

We decided to apply machine learning techniques as an extension of our network dimensioning tool, STEM, presented in Cable-Tec Expo '16 [1]. We called this extension STEM-ML and it makes use of algorithms such as Principal Components Analysis (PCA) and Artificial Neural Networks (ANN). The objective of this work is to characterize the nodes that make up our network in order to define the strategies that Cablevisión will use to meet short and long term demand.

In STEM-ML, we carried out different analysis at node level based on variables such as monthly consumption, households passed, traffic per port and downstream channels distributions, among others. We use a huge volume of data from different sources to obtain examples for the training sets used in the algorithms. As the obtained results are needed for a large number of cases and on a regular basis, it is necessary to automate these processes applying machine learning and multivariate analysis techniques.

1.1. Machine Learning Overview

The main applications of Machine Learning technology in telecommunications and in particular in the cable industry are listed below.

While machine learning has been under research and development for decades, we may wonder why it has just now become Strategic Technology and is in peak expectations of the Gartner's Hype Cycle for Emerging Technologies [2].

The reason why it has become strategic is the great processing power available and the thousands of algorithm developers who have improved the performance of that technology. In addition to the large investments made by companies such as IBM (Watson project), Google (TensorFlow project) or Microsoft (Azure).

Machine learning is the subfield of computer science that, according to Arthur Samuel in 1959, gives "computers the ability to learn without being explicitly programmed" [3].

Machine learning technologies can learn from historical data based on it making predictions or making decisions. That is the fundamental difference between any other applications developed from a program's instructions that are ran in a deterministic manner (Figure 1).

Machine learning is based on algorithms that learn from data without relying on rules-based programming (Figure 2).




Figure 1 - Traditional Programming

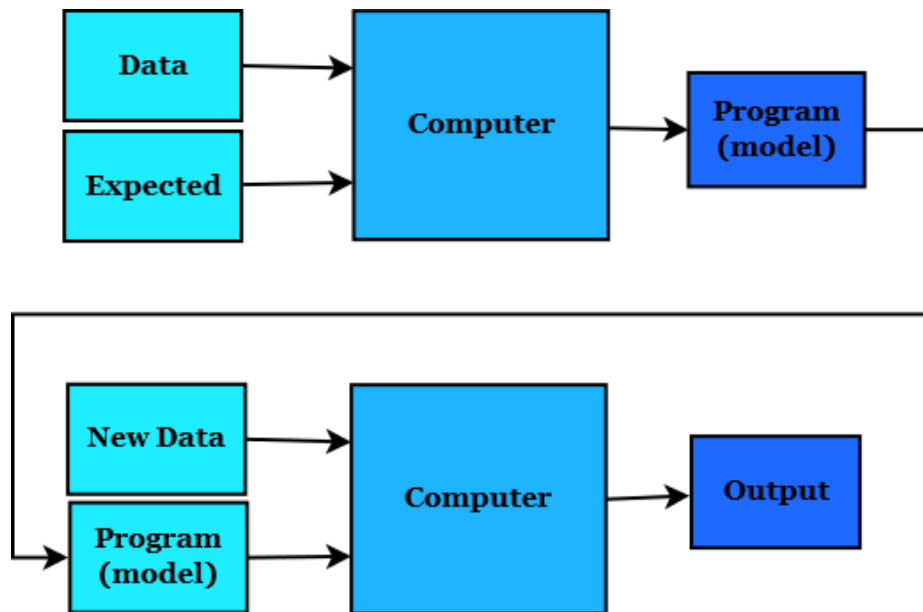


Figure 2 - Machine Learning

At the TM Forum [4] presentations and discussion panels, possible uses and potential applications in the Telecommunications business were presented in the analytics sessions. Among the advantages of the use of Machine Learning, we can mention:

- It allows fast and automatic analysis of large volumes of data that are becoming more and more complex. Getting faster and more accurate results that allow you to make reliable and repeatable decisions.
- Focus on behavioral analysis to detect and predict possible "anomalous" events at an early stage.
- Automate real-time analysis in the orchestration of end-to-end services in a virtualized world.

- Identification and mitigation of security threats in services through predictive analytics and machine learning to detect attacks that escape traditional preventative static defenses.
- Prediction of Churn. Unlike traditional strategies, machine learning allows a multi-class classification of our clients, for example to predict whether they belong to a low, medium or high-risk class.
- Support for automation and management of network orchestration and traceability of end-to-end transactions across the network and OSS / BSS environment.

For the cable industry in particular, Sundaresan et. al. in [5] provides an overview of Machine Learning algorithms, and how their potential applications could be applied:

- Software Defined Networks (SDN) Routing
- Profile Management on DOCSIS 3.1 cable modems [6]
- Proactive Network Maintenance (PNM): for DOCSIS
- HFC's Network Health KPI

Some applications are being implemented in:

- Internet Traffic Characterization
- Network Traffic Engineering
- Wi-Fi Proactive Network Maintenance (PNM)

There are two main paradigms of machine learning, called supervised and unsupervised.

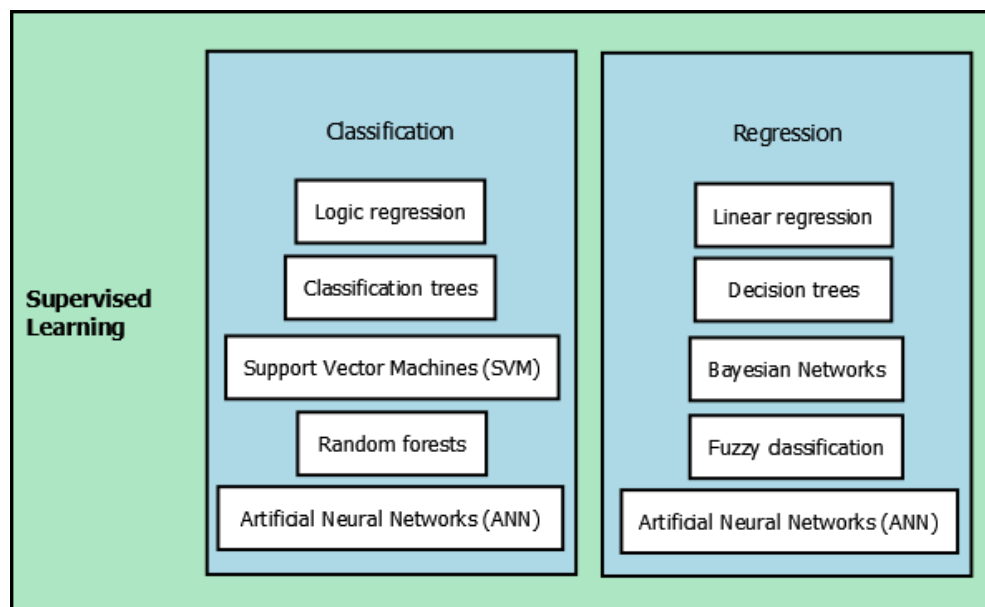


Figure 3 - Most used supervised machine learning techniques.

In supervised learning, your training data consists of some points and a label or target value associated with them. The goal of the algorithms is to figure out some way to estimate that target value. Learning stops when the algorithm achieves an acceptable level of performance.

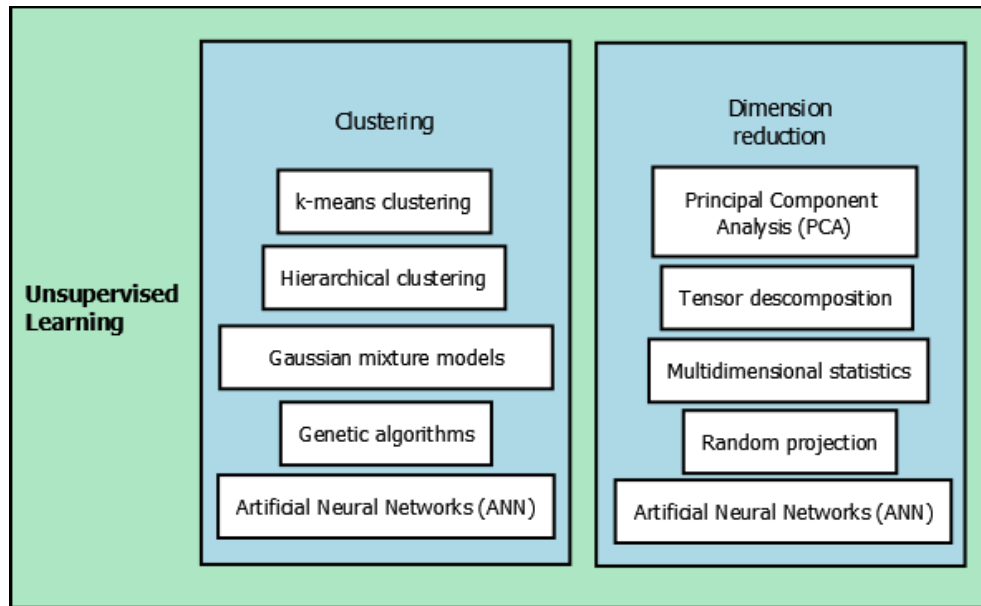


Figure 4 - Most used unsupervised machine learning techniques.

In unsupervised learning, there is just raw data, the output is not known beforehand. Unsupervised algorithms are used for finding hidden underlying structure in the data; there are no correct answers and no teacher.

Supervised learning is somewhat more common in real applications. However, unsupervised learning algorithms are often used as a preprocessing step for extracting meaningful features from a data point, with those features ultimately getting used for supervised learning [7].

1.2. Motivation

In order to define the access network strategy in Cablevisión, last year we carried out a first analysis of optical nodes, in which we characterize them according to their high or low demand, and it was assumed that in all of them the average amount of bandwidth per subscriber during prime time (T_{Avg}) has a 50% Compound Annual Growth Rate (CAGR). Based on this forecast, we suggested that in cases with higher values of this index (T_{Avg}), a new optical node should be installed using architectures N+0 such as Remote MAC-PHY or Fiber Deep, aligned with DOCSIS 3.1 evolution. For those cases with lower values, the suggestion was to enable more QAM channels or to segment the optical node.

At that time, we used historical data about traffic and cable modems volume at HUB and CMTS level. Now we are focused on this same data but at node level so we can make a more detailed and deeper analysis which helps us optimizing the investment plan.

1.3. Datasets Treatment

“Data today is often compared with oil, as in its raw form, its uses are limited. It is through refinement that oil becomes useful as kerosene, gasoline and other goods, and similarly it is through the refinement process of cleansing, validation, de-duplication and ongoing auditing that data can become useful in the kinds of advanced analytics that are starting to shape our world” [8].

Data plays a critical role in the development of smart solutions. Poor data quality puts organizations at risk of making unwise decisions, missing opportunities and undermining the customers' confidence. Rule of thumb: If your human experts struggle to come to conclusions with your existing data, ML will not fix it by itself.

As part of the data treatment phase, some decisions about error treatment based on business guidelines had to be made. Next, an overview of the data processing, main errors found, and ways used to offset them were provided.

As for the terminology used in this paper, a service group (SG) refers to the SCTE definition [9], which is a group of nodes, each node having a number of homes passed (HHP). A headend/hub serves multiple service groups. All nodes in a service group are served by a common switched RF spectrum.

Service group has been borrowed from the video world and has been defined in DOCSIS as the complete set of upstream and downstream channels that can provide service to reach a single subscriber device. Those channels may come from different MAC Domains and even different CMTSs. They could also come from video Edge QAMs.

We add two more concepts: segments and technical zones. Segments are the legs of an optical node that cover a geographical area, they relate to the number of modulators inside the node. Technical zone is the name we use to refer to a specific node/segment combination. This way, we say that a segmented node has an individual modulator for each segment. Therefore, we can say that a node is segmented 1x1, 2x2 or 4x4 depending on the number of downstream and upstream modulators. As an example, a group of technical zones called BON001A, BON001B, BON001C, BON001D refers to the four segments (A, B, C, D) of the node BON001.

Our objective is to estimate traffic at node level (or even segment level in some cases). We will briefly explain how we do this starting from traffic at port level. To do this, it is assumed that the registered traffic at one port is distributed between the nodes or segments in proportion to the CM number in each port.

With this objective in mind, we integrate data that comes from different sources: customer portfolio (cable modem count in each node/segment), traffic in each CMTS port (during prime time periods, Sundays from 18 hrs. to 00 hrs.), and ports/nodes combinations (Figure 5).




Figure 5 - Dataset treatment simplified scheme.

1.3.1. Portfolio, Combinations and Traffic

First, we take CVA data on subscribers and compute the cable modem count in each zone. After unifying the segments' notation, we obtain a data set which combines each CMTS port with its connected nodes, so that they match the one in the portfolio. We merge both data sets and we obtain Table 1.

Table 1 – Example of the resulting database after merging the nodes and CMTS data.

HUB	CMTS_Name	DS_Port	Node	Segment	Zone	CM_Count
ACC	CMT1.ACC1-BSR64K	12/0	ACC001	AB	ACC001A	163
ACC	CMT1.ACC1-BSR64K	12/0	ACC001	AB	ACC001B	164
...

Despite the corrections made, there still are some errors.

- The node is already segmented in the portfolio but not in the combinations data set. **Solution:** copy HUB, CMTS and port information to all segments.
- The node is segmented in the combinations data set but not in the portfolio. **Solution:** divide the cable modem count by the number of segments. This is an arbitrary decision, since the cable modem count is not necessarily evenly distributed among the segments.
- It is not known which port corresponds to each segment. **Solution:** sum up the cable modem count and divide it by the amount of registered ports.
- Some zones have more than one port associated. This is not a data set merge error, but probably a segmented node not yet updated. **Solution:** distribute the cable modem count equally among the ports connected to that zone.

The last dataset that is left is the one that contains traffic data. As we said before, we take this data during prime time and then use the maximum traffic per port.

After joining all the data sets, we still have to decide how to distribute the traffic of each port among the zones that it is connected to. At this point, there are three possibilities:

1. The data we have is complete; we have the cable modem count for each zone connected to the port.
2. The data is partially complete; we have the cable modem count for some zones, but not all.
3. None of the zones associated to the port has the data about cable modem count.

For the first case, we obtain the following weight:

$$weight = \frac{\text{cable modem count in the zone}}{\text{cable modem count in all the zones connected to the port}}$$

Then, each port traffic is distributed using this weight.

For the second case, if one of the segments connected to the port does not have the cable modem count, we consider it a segment with few cable modems, or one that is not yet active. Therefore, we consider this zone has zero cable modems, and the traffic is distributed between the ones that do have this information.

For the third case, another type of estimation is used: the total registered traffic at port level is evenly distributed among all the zones that are connected to it.

Finally, Table 2 shows the errors in the final data set.

Table 2 - Errors in the final database.

Description	%	Cumulative %
No errors.	90.15%	90.15%
Some zones without portfolio information - FIXED.	1.13%	91.28%
No portfolio information - FIXED.	1.31%	92.59%
Ports not associated with any zone.	7.48%	100%

We continue to work towards improving our dataset quality, since it is a common task for data science.

1.4. Variables in this Analysis

This analysis uses traffic data collected every Sunday during prime time, between February 19th and June 4th 2017, for all the ports registered in Cablevisión network.

In particular, one of the variables in these datasets contains the maximum traffic (Kbps) registered in each port. Our approach consists in analyzing two key indicators:

- Average bandwidth traffic per residential subscriber at peak time.
- Ports usage.

The former will provide information about the zones where there is a need for higher bandwidth, and the latter will help us find the optical nodes where ports are operating at almost their full capacity, conditioning the Quality of Service (QoS) and limiting the demand.

To assess the average bandwidth traffic per residential subscriber at peak time, the metric is defined:

$$\text{Average BandWidth per Subscriber [Kbps]} = \frac{\text{Port Traffic}}{\# \text{Subscribers connected}} \quad (1)$$

For measuring ports usage, the maximum utilization was defined:

$$\text{Max utilization [\%]} = \frac{\text{Max Port Traffic}}{\text{Port capacity}} \quad (2)$$

For practical purposes, the maximum utilization metric is also referred to as utilization.

To gather data about how the two key indicators relate to other variables, we also included in our analysis: the count of segments or zones connected to one port, CMTS model, classification of optical nodes according to the region where they are located, 2016 investment plan status, network capacity (1GHz or other), DOCSIS 2.0 and DOCSIS 3.0 cable modems count, HHP, network extension (in Km) and total monthly downstream consumption.

We also included the variable ‘downstream channels available per port’. Depending on the CMTS model, and knowing how many of downstream channels are already being used, we calculate how many more channels the network can support.

2. Exploratory Data Analysis (EDA)

We conducted an extensive exploratory analysis to assess the variables variation ranges, the correlations among them, whether the variables influence the traffic and utilization, and which is the magnitude and trend of that influence to determine its significance in the analysis.

A typical process for data modelling is: Problem → Data → Analysis → Model → Conclusions.

Data visualization is a key aspect of the exploratory phase. A correct visualization should be clear and easy to understand, in order to help any reader detect trends. We used SAS® to obtain a variety of plots and charts, a selection of which are shown in this section.

2.1. Average Bandwidth per Subscriber

Over the surveyed period, there was a slight increase of the average bandwidth per subscriber mean and median. In addition, in Figure 6, the percentile 95 is plotted. This tells us under which values do 95% of the observations lie. The colored area is greater as time passes, so we understand that on every Sunday there are new higher values of the average bandwidth per subscribers, and this means that every Sunday the metric varies among a wider range.




Figure 6 – Evolution of traffic over the surveyed period.




Figure 7 - Traffic at segment level on June, 4th (Mbps) versus subscriber count, colored according to the % of cable modems with DOCSIS 3.0 in each segment.

The diagonal line in Figure 7 divides the segments in the plot. Above it, there are segments in which the average bandwidth per subscriber is higher than 1Mbps. Below, there are the ones for which the metric is lower than 1Mbps. Notice that in most of the zones above the diagonal, the percentage of cable modems with DOCSIS 3.0 is about 60% or higher. On the other hand, the zones below the line tend to have more DOCSIS 2.0 cable modems.




Figure 8 - Evolution of Avg BW per subscriber according to the number of segments in same port.

Figure 8 shows the evolution of the average bandwidth per subscriber mean in the surveyed period. It can be seen that when two or more segments are connected to the same port, the average bandwidth per subscriber drops off. For the cases with two zones connected, the mean is about 20% lower than the same metric for the ones with only one zone. When there are three zones or more, this difference increases another 10%.

**Table 3 - Classification of ports according to the number of segments connected.
Mean and standard deviation of cable modem count for each type of port.**

Segments per port	% of ports	Cable modem count mean	Cable modem count SD
1	65%	208	120
2	30%	322	141
3	5%	472	192

2.2. Correlation between Traffic and Monthly Consumption.

We were interested in studying the correlation between total monthly consumption and the average bandwidth per subscriber. It is well known that the ports with higher traffic are the ones in which there are more cable modems. That relationship actually exists, as is shown in Figure 9, and the correlation between these two variables in May 2017 was 0.85.




Figure 9 - Traffic versus cable modems count at port level.

On the other hand, monthly consumption at segment level is defined as the sum of all consumption that came from the subscribers in that segment. The correlation between monthly consumption and number of cable modems in May 2017, displayed in Figure 10, was 0.61.




Figure 10 – Monthly consumption versus cable modems count at segment level.

If we plot the ports traffic versus monthly consumption values in the segments connected to each port, it may seem as if they correlate. Nevertheless, we already know both variables are associated to a third variable, which is the number of cable modems in each segment.

In order to assess if there is a correlation between maximum traffic during peak hours and total monthly consumption, we tried to remove the effect of cable modems count on both variables. Therefore, we postulated two linear regression models:

Model 1:

$$Y_i = \alpha_0 + \alpha_1 \cdot X_i + \varepsilon_i \quad (3)$$

Where

- Y_i : Maximum traffic registered in port i
- X_i : Sum of the cable modems count of all the segments that are connected to port i
- α_0, α_1 : Fixed coefficients to be estimated
- ε_i : Random error component, $\varepsilon_i \sim N(0, \sigma_1^2)$

Model 2:

$$Z_i = \beta_0 + \beta_1 \cdot X_i + \delta_i \quad (4)$$

Where

- Z_i : Total monthly consumption registered in May 2017 in segment i
- X_i : Cable modems count in segment i
- β_0, β_1 : Fixed coefficients to be estimated
- δ_i : Random error component, $\delta_i \sim N(0, \sigma_2^2)$

In both cases, residuals can be calculated. A plot of the residuals from the first model versus residuals from the second will provide information about correlation between traffic and monthly consumption once the effect of the third variable, cable modems count, is removed. This is equivalent to calculating the partial correlation of the first two variables, conditioned by the third. In Figure 11, it can be seen that there is no association between traffic at peak time and monthly consumption, at least when they are measured at port level. In fact, the estimated partial correlation is -0.06.




Figure 11 - Residuals from Model 1 versus residuals from Model 2.

We evaluated the same correlation at client level. Considering that for the majority, the expected monthly consumption is about 70 GB. We defined heavy users as the ones who download more than 1000 GB per month. It was found that heavy users also produce high traffic during prime time. We can conclude that these clients do have an impact on the ports traffic.

2.3. Principal Component Analysis

Sometimes data is collected on a large number of variables, so it becomes too large to study and interpret properly; there could be too many pairwise correlations between the variables to consider. That is why it may be useful to use a dimension reduction technique to help simplify the analysis. One of these techniques is Principal Component Analysis (PCA).

PCA takes in a collection of d -dimensional vectors and finds a collection of d “principal component” vectors of length 1, called PC1, PC2... and PC d . This means you can get as many main components as original variables. A point x in the data can be expressed as:

$$x = a_1 \cdot PC1 + a_2 \cdot PC2 + \dots + a_d \cdot PCd \quad (5)$$

However, the PCi are chosen so that generally a_1 is much larger than the other a_i , a_2 is larger than all a_i except for a_1 , and so on [7]. Therefore, there will be a subgroup of components explaining a high percentage of the total dispersion of data usually 2 or 3 components.

Principal Components are the underlying structure in the data. Their interpretation will be based on the weights obtained from the original variables. PCA is a way of identifying patterns in data, and expressing it with fewer variables.

We made a PCA for the ports database. The variables included were:

- Maximum traffic per port for each prime time (Sundays from 18 hrs. to 00 hrs.) from 02-19-2017 to 06-04-2017 (Kbps_port1 – Kbps_port16)
- Count of downstream channels in use per port (Channels_used)
- Count of areas connected to each port (Areas_port)
- Households passed (HHP)
- Residential subscribers per port (Subscribers)
- Traffic Management

Table 4 shows the proportion of the total variance of the data that is explained by each component. The first one explains 77.24% of the total dispersion. For the second one, its 5.7%. This means that the first two principal components explain 82.94% of total variance.

Table 4 - Percentage of the total variance explained by each PC.

Principal Component	% of variation	Cumulative %
1	77.24	77.24
2	5.70	82.94
3	4.30	87.24
4	3.66	90.90
...
21	0.13	100

Each principal component is just a weighing of the original variables. Therefore, in order to assign them a name and a meaning, we analyze the weights for the other variables. We concluded there are two main PC, which can be explained as follows:

- PC1: It will take higher values for those ports that registered more traffic during the time surveyed. Channels in use, areas, amount of cable modem and HHP per port also have a positive yet lower impact.
- PC2: This component will take higher values as the amount of areas, cable modems and HHP per port increase, as well as traffic management. On the other hand, it takes lower values as the number of downstream channels in use increases. This variable informs about a port's incapacity to provide a good service in highly populated areas.

We decided to compare PC1 and PC2 to understand the information they provide about the ports.




Figure 12 - PC1 vs PC2 and its relation with traffic management.

Figure 12 shows two groups. The first group where ports connect few subscribers and have many downstream channels. An increment in the subscribers count is associated to a huge increment in the traffic at port level. This increment is possible because these ports have a higher capacity due to the large number of downstream channel. This is not the case for the second group, which presents more subscribers and fewer downstream channels. The increase in subscribers count does not seem to have such an impact in the total traffic at port level. There are few cases where traffic is slightly influenced by traffic management.

In the same figure, we can also see a few ports where the Operations team was performing tasks affecting the service and because of that the traffic management has higher impact on the traffic. These ports connect highly populated areas, meaning, they are associated to high cable modems and HHP count, with many areas connected, but still few downstream channels in use.

We conclude the PCA highlighting that there are two groups of ports. One where the aggregation of more subscribers draws a substantial increment in the traffic, and another where the impact of adding subscribers is lower.

2.1. Ports Usage

In Figure 13, we can observe that the percentage of ports with utilization below 80% decreases over the observed period.




Figure 13 - Evolution of utilization over time.




Figure 14 - Avg BW per Subs (Kbps) versus Utilization (%), colored by number of subscribers.

Figure 14 shows that the traffic per subscriber is generally below 2 Mbps. For obvious reasons the more subscribers the lower average bandwidth per subscriber, the fewer subscribers the higher average bandwidth. We can also interpret that, as mentioned in section 2.2, there is no correlation with utilization. If we pay attention to the coloring in the plot, we can see that when there are more subscribers, the ports utilization is also higher.




Figure 15 - Utilization distribution according to count of downstream channels being used.

Figure 15 shows the distribution of utilization over time according to the downstream channels distribution. Notice that utilization not only is higher for ports with eight channels, but it also has higher variability. In all cases, utilization is lower on February 26th due to a long weekend effect, as on February 27th and 28th it was national holiday in Argentina.




Figure 16 - Utilization in ports with 8 DS channels used, none and 8 or more channels available. The colored area shows the interval between the 25th and 75th percentiles.

We compare the utilization distribution in the ports where there are no channels left to be used against the ones where there is some capacity left. In Figure 16 the blue area shows that half of the ports with eight or more downstream channels available have utilization values that lie between 40% and 80%. Utilization of a quarter of the ports in this group is below 40% and another quarter, is above 80%. On the other hand, the green area shows that percentiles of utilization for the ports with no channels available is generally lower.

3. Construction of an Artificial Neural Network (ANN)

We based our ANN on the following principles [10]:

- Parsimony Principle: the simplest model that fits the data is also the most plausible.
- Sampling Bias Principle: if the data is sampled in a biased way, then learning will produce a similarly biased outcome.
- Data Snooping Principle: if a data set has affected any step of the learning process, its ability to assess the outcome has been compromised.

3.1. Artificial Neural Networks

The inventor of one of the first neurocomputers, Dr. Robert Hecht-Nielsen, provides the simplest definition of an Artificial Neural Network (ANN). He defines a neural network as:

"...a computing system made up of a number of simple, highly interconnected processing elements, which process information by their dynamic state response to external inputs" [11].

The original goal of the ANN approach was to solve problems in the same way that a human brain would. They provide a practical method for learning discrete-valued functions from examples [12].

ANNs are typically organized in layers. Layers are composed of a number of interconnected nodes, which contain an “activation function”. Patterns are presented to the network via the “input layer”, which communicates to one or more “hidden layers” where the actual processing is done via a system of weighted connections. Most ANNs contain some form of “learning rule” which modifies the weights of the connections according to the input patterns that it receives. The hidden layers then link to an “output layer”. If the network generates a “good” output, there is no need to adjust the weights. However, if the network generates a “poor” output, then the system adapts, altering the weights in order to improve subsequent results [13].

A network of many neurons can exhibit incredibly rich and intelligent behaviors. Once a neural network is “trained” to a satisfactory level, it can be used as an analytical tool on other data. They are simple to use and effective classifiers.

3.2. Network Access Strategies

In order to increase the capacity of the network and consequently the access speed for DOCSIS 2.0 and DOCSIS 3.0 HFC networks, we decide which of these four strategies to apply:

1. Chassis upgrade

It consists on an upgrade of the firmware in the CMTS, starting with one or 8 QAMs per physical port, we can upgrade to 16, 20, 24 and 32 QAMs on the same port.

2. Recombination

This process is an upgrade inside the HUB or internal plant. During the deployment of an optical node, and based on the capacity of multiple channels per downstream physical port, a common way to size a service group is to split one downstream port into 1:2 RF combiner. Each one goes to an individual transponder and by the optical Tx goes to two different optical nodes. When we make this kind of upgrade, we remove the RF combiner and reconnect each Tx to a new physical port.

3. Node segmentation

This process is an upgrade in the external plant. We open an optical node, plug a new Tx or Rx (or both) module, and reconnect the segments to the new modules. One optical node may be segmented or not depending on how many modules it has. Typical segmentation schemes are 1x1, 1x2, 2x2, 2x4 and 4x4, where the first and second numbers belong to the number of Tx modules and Rx modules, respectively.

4. Node division

Another task we use to increase the capacity of the network is node division and it can follow two different approaches: the first one is to divide the existent node into two new nodes, reassigning the distribution of subscribers among the two nodes. The other one point to a deep fiber approach. The deep fiber approach requires a node $n+0$ topology. With the existent network it implies that each amplifier will be replaced with a new fiber node reducing the number of HHP per node by the rate of the existing active devices.

Depending on the chosen strategy, in terms of numbers, it implies that we can jump from a node with 500-1000 HHP to two nodes with 250-500 HHP in the case of node division, or $500-1000/x$ where x is the number of the actives above the node in the $n+m$ existing topology. For example, if we have $n+3$, $n+4$, $n+3$, $n+2$ in each segment of a node with 500-1000 HHP, we will increase the number of nodes to at least 4 new deep fiber nodes and the jump goes to 120-250 HHP.

We expect our ANN will classify Cablevisión network nodes into one of these four strategies. The criteria that it will learn is the one given by the expert team.

3.3. Data Sample

We want to get a sample of nodes for the expert team to classify. Then, we use this data to train a neural network to do the same job for all the nodes in the network.

To determine the characteristics of the sample, we classify nodes into three strata: in the first one, we'll have the nodes in which the ports have a mean utilization below 50%; the second stratum contains the ones where mean utilization lies below or is equal to 80%, and in the third one, the nodes for which the mean utilization is above 80%.

The HHP count in each optical node is a determinant factor when it comes to choosing a strategy. Table 5 shows that when utilization is high, there is an increase of households passed mean and standard deviation.

Table 5 – HHP statistics according to the utilization range.

Utilization range	Strata size	HHP Mean	HHP Standard Error
<=50%	2,539	636.10	364.98
50%-80%	1,992	727.73	418.36
>80%	426	804.61	450.16

It is in our interest to have a fair representation of the HHP variable in our sample. Therefore, we will look for a sample size, so that if we had to estimate the mean HHP in each stratum, the estimation would have a certain standard error. As the investment in nodes with higher utilization has a higher priority, we want to be more accurate in these cases. Hence, the desired standard error for the second and third strata will be lower than the desired standard error for the first one.

One way to obtain a sample in a stratified population is to treat each stratum as a “population” and calculate a simple random sampling for each [14]. As the wanted precision varies from one stratum to the other, we decided to calculate separately for each stratum:

$$n_h = \frac{S_h^2}{V_h}$$

Where

- S_h : Squared standard error, which is the same as the variance of stratum h
- V_h : Desired variance for the sample in that population.

The desired standard errors and sample sizes (n_h) are shown in Table 6.

Table 6 - Sample size calculation for each stratum.

Utilization stratum	HHP Desired Standard Error	n_h
<=50%	100	13
50%-80%	80	27
>80%	80	32

If we sum n_h , we obtain a total sample size of 72 nodes, which are randomly selected within each utilization stratum.

3.4. Neural Network Training

It is possible that after training the neural network, which supposedly returns highly accurate classifications, the predicted classifications do not make sense when new data is introduced. This is generally due to overfitting, which means that the neural network adjusts too well to the data used for training but to that dataset only.

In order to prevent overfitting, we split the sample data in three subsets: training, cross-validation and testing. We assign 60% of the cases to the training set, so we use them to estimate the weights in the neural network. Then, another 20% goes to the cross-validation set, which helps us validating the model in terms of the variables and optimization parameters selected. Finally, we use the remaining 20% as testing set. This part of the sample does not participate in the construction of the neural network, it is only used to check whether there is overfitting or not by measuring how well would the network classify ‘new observations’.

The first time we tried to train a neural network, using the software Octave [15], we introduced the variables that usually appear in the process of choosing the strategy. Specifically, these variables were: ports’ utilization, HHP, cable modem count, average bandwidth per subscriber, downstream channels being used, downstream channels available, count of segments and nodes sharing ports, network capacity and segmentation level.

We soon discovered that the existing correlations between these variables were decreasing the network’s accuracy. We used a technique for variable selection known as *stepwise*, to find the optimal combination of input variables. As a result, the first layer of the network contains four inputs: utilization, downstream channels available, count of segments and nodes sharing ports and segmentation level.

At first, accuracy was low and we wanted to improve it. We used the cross-validation set to obtain a learning curve that could help us understanding if we needed some more training observations or to change the selected variables. Figure 17 shows that after 25 observations, even if the sample size increases, the errors in the training set and in the cross-validation set are approximately the same. This indicates that we don’t need a bigger sample to train the network. Actually, we need to reflect a more complex structure.




Figure 17 - Learning curve for the ANN based on four selected variables.

We included quadratic terms to search for higher accuracy. Therefore, as Figure 18 shows, the network's first layer contains eight inputs, the four variables mentioned (represented as x) and the same variables at square (x^2). The second layer, also called hidden layer, contains eight data points too ($a^{(1)}$), and finally the output layer contains five classes ($a^{(2)}$), which refer to the four strategies already detailed and the fifth option 'no action needed' for the nodes where there is no need for investment at the time.

The optimal solution for the weights in the network throws an accuracy level of 96% with the training set. After evaluating the classifications through the testing sample, we found that the accuracy is actually around 90%.




Figure 18 – Our Neural Network scheme.

4. On-going and Future Work

The next step is to incorporate upstream traffic data and improve the ANN accuracy, towards 95% or higher. We will then use it to evaluate the impact of long term actions (in approximately ten years), under different T_{Avg} growth scenarios.

The ANN will be periodically used on all of Cablevisión nodes to guide future investment actions. A Machine Learning tool that involves different scenarios with many variables can be transformed into a very powerful, accurate and scalable resource.

We are using STEM to analyze the impact on a campaign in which we double the speed of access to our customers. It allows us a quick evaluation of the conditions of the nodes and their capacity.

This duplication is complementary, generating a great satisfaction in our customers.

Conclusion

The applications of the technology of Machine Learning have made a very strong advance in the industry of the telecommunications and especially in the Cable industry, providing multiple advantages as detailed above. In addition, it facilitates the operation given the trend of the strong virtualization.

However, we must not forget that the application success is based on the tasks performed by people. In general terms the tasks to be developed to find the learning models are:

- Data: Separate development and validation data. Define instances, classes and attributes.
- Experimentation: Selection of attributes. Performance measures. Cross-validation.
- Validation of the models: Processes intended to verify that models are performing as expected, in line with their design objectives and business uses. It's the most important step in the model building sequence.

In short, applications developed with machine learning technologies are based on human art and science.

Abbreviations

ANN	artificial neural network
Avg	average
BW	bandwidth
CAGR	compound annual growth rate
CM	cable modem
CMTS	cable modem termination system
CVA	Cablevisión S.A.
DOCSIS	data over cable service interface specification
EDA	exploratory data analysis
GHz	giga hertz
HFC	hybrid fiber coaxial
HHP	household passed
IPTV	internet protocol television
Kbps	kilobits per second
Km	kilometers
Mbps	megabits per second
ML	machine learning
OSS/BSS	operation support system/business support system
PCA	principal components analysis
PNM	proactive network maintenance
QAM	quadrature amplitude modulation
QoS	quality of service
SD	standard deviation
SDN	software define network
SG	service group
STEM	Science, Technology, Engineering and Mathematics
Subs	subscribers

Bibliography & References

- [1] M. Fiorenzo, C. Righetti, C. Carreño Romano, G. Carro. IP Traffic Analysis: a Tool for Network Dimensioning. SCTE Expo 2016.
- [2] Gartner's Hype Cycle for Emerging Technologies, 2016,
<http://www.gartner.com/newsroom/id/3412017>
- [3] Arthur Lee Samuel, "Some studies in machine learning using the game of Checkers," IBM Journal of Research and Development 3 (1959).
- [4] TM Forum Live , May 15-18 2017, Nice France
- [5] Karthik Sundaresan, Nicolas Metts, Greg White, Albert Cabellos-Aparicio, Applications of Machine Learning in Cable Access Networks SPRING TECHNICAL FORUM, CableLabs SCTE NCTA 2016 Spring Technical Forum Proceedings.
- [6] Greg White and Karthik Sundaresan, DOCSIS 3.1 Profile Management Application and Algorithms, SCTE NCTA 2016 Spring Technical Forum Proceedings.
- [7] "The Data Science Handbook", Field Cady, 2017, Wiley.
- [8] Pitney Bowes (2017), "The Data Differentiator: How Improving Data Quality Improves Business", Forbes Magazine.
- [9] DOCSIS Engineering Professional, SCTE Course Participant Guide, 2014.
- [10] Professor Yaser Abu-Mostafa, Caltech's Machine Learning Course - CS 156, Lecture 17 - Three Learning Principles. <https://www.youtube.com/watch?v=EZBUDG12Nr0>
- [11] "Neural Network Primer: Part I" by Maureen Caudill, AI Expert, Feb. 1989.
- [12] Machine Learning. Tom Mitchell. 1997. McGraw-Hill.
- [13] Daniel Shiffman, The Nature of Code, Chapter 10: Neural Networks,
<http://natureofcode.com/book/chapter-10-neural-networks/>
- [14] W.G. Cochran, Sampling Techniques, Chapter 5: stratified random sampling (page 65).
- [15] John W. Eaton, David Bateman, Søren Hauberg, Rik Wehbring (2016). GNU Octave version 4.2.0 manual: a high-level interactive language for numerical computations.

Traffic Engineering in a Fiber Deep Gigabit World

A Technical Paper prepared for SCTE•ISBE by

John Ulm

Engineering Fellow, Broadband Systems
CTO - Network Solutions team
ARRIS
john.ulm@arris.com

Tom Cloonan

Chief Technical Officer, Network Solutions
ARRIS
tom.cloonan@arris.com

Introduction

Many new innovations are finding their way into cable operators' plants and many others are on the way. With DOCSIS 3.1 deployments well underway, a new era of Gigabit services in the downstream are being introduced. With DOCSIS Full Duplex (FDX) on the horizon, operators have the promise of symmetric Gigabit services. At the same time, some operators are starting their migration to fiber deep networks and others are looking at Distributed Access Architectures (DAA) such as Remote PHY.

All of this will have significant impact in the way operators manage their traffic engineering and network capacity planning. As an industry, we need to re-evaluate and update our models. This starts with an intimate understanding of subscriber bandwidth behavior. This paper takes a detailed look at a year's worth of live consumer data collected from a single cable site. When sampling during peak busy hours, every packet was tracked to allow for traffic analysis down to the second.

The paper will investigate many different bandwidth trends uncovered. Some of the key variables of interest include traffic consumption based on:

- Differing Service group sizes
- SG to SG variation
- Subscriber service tiers
- Time of day
- Day of week
- Month to month

With these trends in hand, the impacts on existing network capacity models are discussed and how they might morph to provide traffic engineering in a Fiber Deep Gigabit world.

Broadband Bandwidth Trends

The Internet has been growing at a breakneck speed since its inception. And with it, we have seen a corresponding growth in dedicated network capacity. [ULM_2016] provided an overview of these trends which are highlighted and updated below.

1. Nielsen's Law and Cloonan's Curves

While Moore's Law is infamous in silicon realms, Nielsen's Law of Internet Bandwidth has become renowned in the broadband world. It basically states that network connection speeds for high-end home users would increase 50% per year. This law has driven much of the traffic engineering and network capacity planning in the service provider world. It has also led to much research on those topics.

In [CLOONAN_2014, EMM_2014], this research was expanded to also include traffic utilization in addition to the network connection speed. Nielsen's Law is shown in the figure below. Since the Y-axis is a log scale, the 50% Compounded Annual Growth Rate (CAGR) appears as a straight line. An interesting fact is that the graph starts in 1982 with a 300-baud phone modem. The industry is now in the fourth decade of closely following this trend.




Figure 1 - Nielsen's Law – 50% CAGR




Figure 2 - Modified Nielsen's Law – 33% CAGR after 2018

While this trend shows a straight line increase, in reality, internet speeds will make a jump and stay there for a bit. There was recently a jump to 1 Gbps services that happened a couple years ahead of the Nielsen's Law projection. This is shown in the Figure 1. Service Tiers may stay here a couple years before the speed tier climb continues.

While Nielsen's Law focuses primarily on downstream speeds, it has been noted that upstream speeds have generally followed the same growth rate, but at about one-tenth the speed. However, with more Fiber to the Premise (FTTP) deployments and the upcoming introduction of DOCSIS Full Duplex (FDX), it is expected that the highest offered upstream speeds will take a step up as symmetric services become available.

Will Nielsen's Law continue its 50% growth unabated for the next couple decades? In recent years, there has been some suggestions on whether Moore's Law may be slowing down. Will this have a corresponding impact on Nielsen's Law? One could argue that Moore's Law is the fuel behind ever advancing Consumer Premise Equipment (CPE), and these CPE drive the need for Internet bandwidth.

Figure 2 takes a look at a modified Nielsen's Law where the CAGR is reduced to 33% going forward. This stretches the time for 10X growth from the original 5½ years up to 8 years. You can see that from a network capacity planning perspective, the overall impacts are similar. The changes over the next decade are minimal. Longer term, the time it will take to reach the 1 Tbps milestone gets pushed out about five years, from 2035 to 2040. So, even with a slowing in Nielsen's Law, there will be minimal impact in operators' long term network capacity planning.

Earlier work by Cloonan noted that the primetime average subscriber consumption (a.k.a. T_{avg}) has also been following this same basic trend as shown in the Figure 1. For service providers, an important metric is the traffic utilization in a Service Group (SG). The SG traffic utilization is a function of the number of subscribers (N_{sub}) times the average bandwidth per sub (T_{avg}). In [CLOONAN_2014, EMM_2014], this research was expanded to also include traffic utilization in addition to the network connection speed. This was shown in a chart known as Cloonan's Curve, where SG consumption is shown in addition to Nielsen's Law.

In the early DOCSIS days, many nodes were combined together and a SG might consist of thousands of subscribers. At that time, the SG traffic was an order of magnitude higher than the maximum network connection speed (a.k.a. T_{max} after the DOCSIS parameter that dictates max network rates). Over time, the SG size has been shrinking and, with it, the ratio between $N_{sub} * T_{avg}$ to T_{max} . The SG traffic will eventually approach that of T_{max} . As SG sizes dip below 100 subs, then T_{max} starts to dominate the traffic engineering.

2. Broadband Subscriber Traffic Consumption

ARRIS has been monitoring subscriber usage for many years now. The chart below shows T_{avg} , the average subscriber downstream consumption during peak busy hours, for a number of MSOs over an eight year period. At the start of 2017, T_{avg} finally broke the 1 Mbps barrier.

It turns out that the T_{avg} growth rate was higher at the start of this decade and has tailed off a bit in recent years. Over the last 3-4 years, this group of MSOs have an average downstream traffic growth that has been just under 40%. On a yearly basis, traffic growth can be very sporadic. It is not uncommon to see high growth in one year followed by little growth the next. So, the 40% trend should be used as a longer term guideline on downstream traffic consumption. This equates to roughly doubling every other year.




Figure 3 - Tavg, Average Subscriber Downstream Consumption




Figure 4 - Tavg, Average Subscriber Upstream Consumption

Interestingly, the upstream traffic is growing at a significantly slower rate. During the same eight year period, the upstream Tavg only grew at ~20% CAGR. Traffic is also becoming more asymmetric with video applications driving downstream consumption [EMMEN_2014]. While the average DS:US ratio is ~10:1, the MSO with the largest DS:US ratio seemed to have stabilized around a 15:1 ratio. It will be interesting to see what happens with other operators as they reach this point.

3. Selective Subscriber Migration Strategy

At first glance, Nielsen's Law is a scary proposition in that HFC networks might be obsolete in 5-7 years while it may take decades to build out an FTTP infrastructure. However, this is not the full story. As was shown in [ULM_2016, ULM_2014], Nielsen's Law applies to the top speed tiers which is only a very small percentage of the entire subscriber base, perhaps less than 1%. So, the key question then becomes, "What happens to the vast majority of subscribers on HFC who are not in the top speed tiers (a.k.a. billboard tiers) and when?"

The [ULM_2014] case study looked at service tier evolution at a few MSOs. Perhaps the key finding from this study is that the different service tiers are growing at different rates. While the top billboard tier continues to follow Nielsen's Law 50%, each subsequent lower speed tier is growing at a slower rate. Hence, the lower the service tier rate, the lower its CAGR.

The figure below maps out an example of the various service tier growth over the next two decades. While the 1% of subs in the top billboard tier hit 10 Gbps in ~2024, the 14% of subs in the performance tier don't hit that mark until ~2032. Notice that 85% of subscribers in the flagship basic tier and economy tier stay below this mark for several decades. It is important to note that 99% of the subscribers will still be comfortably using today's DOCSIS technology on HFC a decade from now. With a Selective Subscriber Migration strategy, it becomes very important from a traffic engineering perspective to understand the behavior of the individual service tiers. But with this understanding in hand, Selective Subscriber Migration can be used to extend the life of HFC for decades to come.




Figure 5 - Downstream Growth with Multiple Service Tiers

Review of Broadband Traffic Engineering

Previously, [CLOONAN_2014] provided an introduction on Traffic Engineering and Quality of Experience (QoE) for broadband networks. From there, the paper went on to develop a relatively simple traffic engineering formula for cable service groups.

1. The “Simple” Traffic Engineering Formula

The “Simple” formula shown below is a simple two-term equation. Its simplicity is part of its beauty. The first term ($N_{sub} \cdot T_{avg}$) allocates bandwidth capacity to ensure that the aggregate average bandwidth generated by the N_{sub} subscribers can be adequately carried by the service group’s bandwidth capacity. The first term is viewed as the “DC component” of traffic that tends to exist as a continuous flow of traffic during the busy-hour period.

THE “2014” TRAFFIC ENGINEERING FORMULA (BASED ON T_{max_max}):

$$C \geq (N_{sub} \cdot T_{avg}) + (K \cdot T_{max_max}), \quad (1)$$

where:

C is the required bandwidth capacity for the service group

N_{sub} is the total number of subscribers within the service group

T_{avg} is the average bandwidth consumed by a subscriber during the busy-hour

K is the QoE constant (larger values of K yield higher QoE levels)...

where $0 \leq K \leq \text{infinity}$, but typically $1.0 \leq K \leq 1.2$

T_{max_max} is the highest T_{max} offered by the MSO

There are obviously fluctuations that will occur (i.e. the “AC component” of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term ($K \cdot T_{max_max}$) is added to increase the probability that all subscribers, including those with the highest T_{max} values, will experience good QoE levels for most of the fluctuations that go above the DC traffic level.

The second term in the formula ($K \cdot T_{max_max}$) has an adjustable parameter defined by the K value. This parameter allows the MSO to increase the K value and add bandwidth capacity headroom that helps provide better QoE to their subscribers within a service group. In addition, the entire second term is scaled to be proportional to the T_{max_max} value, which is the maximum T_{max} value that is being offered to subscribers. A change in the K value results in a corresponding change within the QoE levels experienced by the subscribers who are sharing the service group bandwidth capacity (C). Lower K values yield lower QoE levels, and higher K values yield higher QoE levels).

In previous papers [CLOONAN_2013, EMM_2014], a similar formula assumed that a K value of ~ 1.0 would yield acceptable and adequate QoE results. [CLOONAN_2014] goes on to provide simulation results that showed a value between $K=1.0$ and 1.2 would provide good QoE results for a service group of 250 subscribers. Larger SGs would need larger values of K while very small SGs might use a K value less than 1.0 .

Using the simple Traffic Eng formula (1), it becomes possible to develop sophisticated network capacity models. Some results from the ARRIS Network Capacity model are shown in Figure 6. It provides an insight into both Tmax and SG Tavg behavior. During the next 5-7 years, the Tmax component dominates traffic engineering as it is driven by Nielsen's Law. The bandwidth needed by the top billboard tier dominates compared to the SG Tavg. But as top tiers are moved off the HFC, then eventually the Tavg component catches up again.




Figure 6 - Network Capacity Model Results

2. Limitations of the “Simple” Traffic Eng Formula

The “Simple” formula has been extremely effective, but it is very important to understand its limitations. The formula was developed for service groups on the order of a couple hundred subscribers, and found that a K value between 1.0 and 1.2 provided good QoE.

However, as the cable world migrates to Fiber Deep HFC designs jointly with Distributed Access Architectures, operators will need to perform traffic engineering on multiple different sized groups. On one extreme, the DOCSIS SG might shrink to less than 100 subscribers, perhaps only a couple dozen. On the other extreme, operators need to engineer the network links in and out of the Routers and CCAP Core with tens of thousands of subscribers. In between may be a multi-tiered Ethernet switching infrastructure where 100G Ethernet links cascade down to 40G links down to 10G links. Every link needs to be managed to make sure it is not a bottleneck to providing acceptable QoE.

The simulations in [CLOONAN_2014] show that the optimum value for K does vary with several different parameters. In reality, finding the optimum value of K becomes very complex and dependent on many variables.

The “DC” component of the formula (i.e. $N_{sub} * T_{avg}$) also has limitations. It appears to be fine for very large sizes but becomes less accurate for smaller SGs where there is much wider SG to SG variation. Going forward, the “Simple” formula will need to evolve to work across these wider ranges of variables.

Subscriber Bandwidth Behaviors

To enhance our Traffic Engineering formula, an intimate understanding of subscriber bandwidth behavior is needed. This paper takes a detailed look at a year's worth of live consumer data collected from a single cable site. A massive amount of data was collected during many peak busy hours. Every packet was tracked to allow for traffic analysis down to the second.

The paper will investigate many different bandwidth trends uncovered. Some of the key variables of interest include traffic consumption based on:

- Month to month
- Day of week
- Time of day
- Differing Service group sizes
- SG to SG variation
- Subscriber service tiers

1. Data Collection Methodology

Data was collected from a live DOCSIS system over the course of a year. Packet monitoring equipment allowed every packet in the system to be captured during each 30-minute sample interval. Typically, multiple measurements were taken during peak busy hours between 6 pm and midnight on a given night. This created massive amounts of raw data that filled disk drives. To make the analysis manageable, the data needed to be parsed into a more usable metric.

Previously in the simple formula, Tavg would typically be calculated across a timespan of many minutes or hours. However, the QoE must be measured on a much finer granularity for the applications people use such as web browsing, OTT video consumption or even running a speed test. All of these events are sensitive to latency on the order of a couple of seconds. We chose to analyze the data at 1 second intervals. This appeared to be the best compromise between observing QoE behavior yet minimizing file sizes to a manageable size.

2. Macro Trends

For a subset of the data, the major trends were reviewed as they varied by month, day and time. This particular data set was collected during the month of June 2016, and then from mid-September through early February 2017. Data was collected across every day of the week and from 6 pm to midnight.

2.1. Month to Month

As discussed earlier, Tavg per subscriber has been rising steadily over the years. The month to month variation for this data, shown in Figure 7 below, confirms this.

As can be seen, the overall trend is higher over time. However, it is not necessarily a smooth linear increase. September saw a big jump and then Tavg dropped a bit for October and November. This was followed by another big jump in December before Tavg slid a bit in January and February.




Figure 7 - Tavg per Sub, Month to Month variations

Tavg was slightly higher than 1100 Kbps in Jan 2017, which puts it a hair higher but in line with other data shown in Figure 3. The Tavg growth from June 2016 to Jan 2017 is just under 20% in 7 months. This equates to ~35% annual CAGR which is also right in line with the Figure 3 data.

2.2. Day of Week

What is the busiest evening for broadband usage? What is the least busy day? It turns out that on average for this data set, Sunday ties Thursday for the honor of busiest evening. This is shown in Figure 8. For least busy day, Saturday barely nudged out Tuesday.




Figure 8 - Tavg per Sub, Day to Day variations

It should be noted that there is only a couple percent swing above and below the overall average, so the relative swings from day to day are not major.

2.3. Time of Day

For traffic engineering, there is often a reference to “peak busy hour”. So exactly when is that? Figure 9 shows Tavg based on Time of Day across this very large data set. Each iteration lasted 30 minutes. The 6 pm bar in the figure represents the average of all iterations that were started between 6 pm and 6:30 pm.




Figure 9 - Tavg per Sub, Hour to Hour variations

On average, the busiest time of day is between 9 pm and 10 pm. Usage is fairly constant during the dinner hours, and gradually increases up to the peak busy hour. The overall increase is ~10%. Usage then drops off as it gets past 11 pm.

3. Micro Trends

When looking across a large data set, the averages shown above provide an interesting data point, but do not have the resolution needed to understand the impacts of data bursts on subscriber QoE.

3.1. Micro view for Time of Day

To get a better understanding of fluctuations on a single day, an example from June is shown in Figure 10. As can be seen, there is ~33% swing from 8 pm to 10 pm which is much higher than the ~5% swing seen when averaged across several months of data. Its peak is about 15% above the monthly average, so this gives a sense for how a daily peak can be higher than the peak when averaged over a month.

The data in Figure 10 shows the variation in Tavg on 30 minute intervals. As was mentioned earlier, applications are concerned with latency impacts on the order of seconds, so even finer granularity is needed.




Figure 10 - Tavg per Sub, Single Day Example

As can be seen in Figure 10, there is a 2½ hour peak busy window from the start of the 8:48 pm interval to the end of the 10:48 pm interval where the peak usage is reasonably consistent, but falls off on either side. A deeper analysis in this window was done at 1 second intervals for more than 1000 subscribers. A histogram of the 1 second Tavg intervals using 10 Mbps bins is shown in Figure 11:




Figure 11 - Single Day Peak Busy Window – BW Histogram, 1 Second Intervals

The average per subscriber bandwidth during the 2½ hour peak busy window is just over 1000 Kbps. This is the highlighted bar near the middle of the chart. The maximum 1-second interval had a Tavg per sub that was just over 1500 Kbps. That's roughly 50% higher than the peak busy window Tavg. The minimum bandwidth seen in a 1 second interval was just under 700 Kbps. Therefore, the ratio of max to min is just over 2:1.

While knowing the max 1 second interval is useful information for understanding burst requirements, that 1 second interval still only represents 0.01% of the peak busy window. A more accurate characterization of the bandwidth distribution is needed for our QoE analysis. Some additional results from this data set are shown in Figure 12.

The standard deviation was calculated to be ~120 Kbps. $T_{avg} + 2$ standard deviations came in just under 1300 Kbps, while $T_{avg} + 3$ standard deviations were about 1400 Kbps.




Figure 12 - Tavg per Sub, Single Day Data Distributions

In addition to standard deviation, an analysis was done using percentiles. The chart shows the percentile levels for 98%, 99%, and 99.9%. The percentile level shows the % of 1 second intervals that are at or below this level. The 98% percentile level was ~1300 Kbps. This means that only 2% of the number of 1 second bandwidth intervals exceeded this level. The 99% percentile was just marginally higher than the 98% level. Finally, the 99.9% level was very close to the ~1500 Kbps maximum. The data points associated with each percentile level are also shown in Figure 11 with the small highlighted bars on the right side of the chart.

4. Service Group Considerations

As discussed previously, the “Simple” traffic engineering equation has considerations as the Service Group sizes become either very small (e.g. Fiber Deep network) or very large (e.g. CCAP Core for Remote PHY). Our research looked at some of the impacts on traffic usage due to varying SG sizes as well as SG to SG variations given the same size and service tier distribution.

4.1. Service Group Sizes

Intuitively, as SG sizes become larger, there is expected to be less relative variation thanks to the benefit of large numbers of samples. Our research tried to understand the extent of this and quantify it. To illustrate this, Figure 13 and Figure 14 show the same data set taken from a single day in Feb 2016 but organized as different sized SG. For Figure 13, the data was organized as one SG with ~1100 subs. For Figure 14, the data was organized as 11 SGs with ~100 subs each. The data provided bandwidth resolution in 1 second intervals. The X-axis in both figures varies from zero to the Max 1-second interval.

Looking first at Figure 13, one can see that it is a much tighter distribution. The maximum 1-second interval is only about 40% higher than the mean value. The minimum 1 second interval is about half of the maximum interval. The coefficient of variation (i.e. standard deviation divided by the mean) is less than 10%. In looking at many other data sets with ~1000 subs, the coefficient of variation ranged from 4% to 10%.




Figure 13 - Bandwidth Distribution for SG with 1100 subs, 1 sec intervals




Figure 14 - Bandwidth Distribution for 11 SGs @ 100 subs, 1 sec intervals

Now looking at the 100 subs per SG data in Figure 14, it is apparent that it is a wider distribution. The maximum 1-second interval is almost triple the value of the mean. The coefficient of variation is much higher, around 35%. In some other data sets, it went over 50%. Figure 14 also indicates the average and 98%, 99% and 99.9% percentile values with shaded bars.

4.2. SG to SG Variations

Even for a given SG size, the traffic engineering must account for variations from SG to SG. Figure 14 shows the aggregated data for 11 unique SGs. But what is happening in each of these SG? Figure 15 helps give us an insight. This is using the same Feb 2016 data set as above.

The heavy blue line in Figure 15 shows the aggregated data from all 11 SGs. The thin lines show the individual SG.




Figure 15 - Tav_g per Sub, Single Day Data Distributions




Figure 16 - Tav_g per Sub, Single Day Data Distributions

Many of the SGs have very similar behavior. In Figure 16, the SGs with the two extremes are isolated: SG 6 with the lightest data traffic and SG 10 with the heaviest data traffic. Tav_g for SG 10 is about twice that of SG 6; despite having virtually the same service tier distribution.

If the aggregated sum had been used for traffic engineering, the 98% percentile level was ~176 Mbps. While this would have been overkill for SG 6, it definitely does not address the needs of SG 10. SG 10 exceeds this value more than 10% of the time. Looking at the 98% level for the individual SG, then SG 6 would be ~124 Mbps while SG 10 would be ~220 Mbps.

5. Subscriber Variations – Heavy Users & Service Tiers

Knowing the average distribution is useful but not enough for understanding SG bandwidth behaviors. This is especially true for smaller SGs. As seen above, two ‘similar’ SG with 100 subs had bandwidth utilization that was different by a factor of two. The following sections explore some of the reasons for the SG to SG variations.

5.1. Active + Heavy Users

Within any given SG, there is a mix of active and idle subscribers, light and heavy users. As SG sizes approach 100 and shrink below that, then the types of user on any single SG can have a significant impact.

Even within a given service tier, the bandwidth usage can vary dramatically. Figure 17 looks at results for the most common service tier, 25M, for the Feb 2016 data set that was analyzed above. This is the mainstream tier and it contains about half of the total subscribers. Almost 62% of the subscribers were predominantly quiet during this 30-minute interval and consumed less than 250 Kbps. Figure 17 shows the Tavag bandwidth distribution for the remaining 200 active subs.




Figure 17 - Tavag BW distribution for 25M Service Tier

The bandwidth average across this entire group was about 1 Mbps. 75% of the total subs were below this level. The highest bandwidth consumer used almost 20 Mbps for an entire 30-minute interval. There were three subs that were over 10Mbps and 21 subs that averaged more than 5 Mbps over that interval. Obviously, if one SG gets a disproportionate number of heavy users (either too many or none at all), this can greatly influence the bandwidth utilization for a given SG.

In addition to analyzing 25M Service Tier, a look at the 100M tier provided a useful insight. SG 6 with the lightest utilization had three 100M subs but they were all relatively quiet. SG 10 had only two 100M subs but one was quiet and the other was extremely active. In fact, the active 100M subscriber averaged

~43 Mbps over the entire 30-minute interval! This was obviously a big factor in SG 10's high bandwidth utilization.

As can be seen for small SG sizes, the number of active and heavy users compared to the number of idle and light users can dramatically affect the SG bandwidth utilization. As SGs become very large, then the laws of statistics tend to even things out.

5.1. Service Tier Impacts

In addition to the active/idle ratio, another important factor in SG variation is the mix of Service Tiers among the various subscribers. Table 1 provides an example mix of Service Tiers with their respective bandwidth utilization.

Table 1 - Example Bandwidth Distribution by Service Tier

Service Tier	% of Subs	Tavg per Sub (Mbps)	Avg Burst Magnitude (Mbps)
6M	8%	0.49	6.9
12M	24%	0.67	7.9
25M	44%	1.01	11.8
50M	11%	1.68	17.6
100M	3%	2.66	26.4
Avg	100%	0.91	10.4

For this data set, the overall Tavg for subscribers was 910 Kbps. As can be seen, the Tavg when measured for each service tier can vary quite a bit. The Tavg for the lowest tier, 6M, came in at just under 0.5 Mbps, while the top 100M tier had Tavg = 2.66 Mbps.

To help manage QoE, our research also investigated the differences in traffic bursts between the service tiers. The bandwidth data was analyzed at 1 second intervals. Figure 18 shows a probability distribution function (pdf) for the bandwidth burst rates at any given second for each of the service tiers.

During relatively idle periods, all the service tiers behave reasonably similarly. This can be seen on the left-hand side of Figure 18. However, once the subscriber becomes very active, then their maximum burst capability is limited by their service tier Tmax value. As can be seen on the right-hand side of Figure 18, the higher service tiers have higher burst rates.

The magnitude of the average burst is also provided in the rightmost column in Table 1. The lowest 6M tier has a burst magnitude of 7 Mbps as its Tmax value (~8 Mbps) is slightly higher than the actual tier level. The 100M tier has a burst magnitude of 26 Mbps, even though its average utilization is only 2.7 Mbps. This gives an insight into the active to idle ratio.




Figure 18 - Bandwidth Transmit Rate Probabilities for a single Subscriber

These differences in service tier bandwidth utilization may be further impacted if an operator has implemented data usage caps for some or all of their tiers.

Updating the Traffic Engineering Formula

The above trends can now be used to evaluate the impacts on existing network capacity models and see how they might morph to provide traffic engineering in a Fiber Deep Gigabit world. The “Simple” formula divided into two parts: a DC component ($N_{sub} * T_{avg}$) that is the average traffic utilization and an AC component ($K * T_{max_max}$) to compensate for traffic fluctuations. However, the AC component has to account for many things including traffic fluctuations.

Stepping back and looking at this slightly differently, our objective is to determine a capacity threshold where applications get good QoE for the network traffic utilization. But how do we measure the QoE component? It turns out that most operators and most subscribers rely on some sort of speed test to determine whether the service is meeting its Service Level Agreement (SLA). Interestingly, it also turns out the speed test is one of the most sensitive applications to increased network latencies and utilizations so it is actually an ideal choice to monitor SLA.

At its most pure form, the traffic engineering requirements are:

$$C \geq T_{burst} + T_{data} \quad (2)$$

where:

C is the required bandwidth capacity for the service group

T_{burst} is the bandwidth target used to meet the SLA test

T_{data} is the overall network bandwidth at the time of T_{burst}

For an operator who wants to have a pure best effort service with no SLA guarantees, then T_{burst} can be set to zero. However, for the many operators under a regulatory microscope, then T_{burst} will be equal to T_{max_max}, the maximum offered service tier. There is a middle ground here as well. Some operators may choose to support a fraction of the advertised service rate. So, for example, an operator might want to guarantee 75% of the service rate, so T_{burst} would equal $0.75 * T_{max_max}$. For the remainder of this paper, it is assumed that T_{burst} equals T_{max_max} as the typical scenario.

The T_{data} component of the above equation is more complex. It must be an estimate of the data utilization during the SLA test. These tests might typically run from 15 seconds to a minute or two. The T_{data} component will obviously vary from time interval to time interval. One can estimate T_{data} by measuring the average bandwidth in the service group and then adding an additional margin to achieve our expected QoE.

1. New Basic Formula

Refining our traffic engineering formula now comes up with this basic one:

$$C \geq T_{max_max} + T_{avg_sg} + QoE_margin \quad (3)$$

where:

C is the required bandwidth capacity for the service group

T_{max_max} is the highest T_{max} offered by the MSO.

T_{avg_sg} is the average bandwidth consumed by a service group during the busy-hour

QoE_margin is additional margin required due to data utilization fluctuations

This becomes our base formula going forward. The first two components, Tmax_max and Tavg_sg are readily available and/or measurable. Our traffic engineering research can now focus on defining the QoE margin component.

But how does this base formula relate to our earlier 2014 “Simple” formula? The answer is quite well. It turns out that the “Simple” formula estimates Tavg_sg using the number of subscribers and the average bandwidth per sub. And it is using a QoE margin of $0.2 * Tmax_max$ when $K=1.2$. See the example below:

$$C \geq Tmax_max + (Nsub * Tavg) + (0.2 * Tmax_max) = (Nsub * Tavg) + (1.2 * Tmax_max) \quad (4)$$

where:

C is the required bandwidth capacity for the service group

Tmax_max is the highest Tmax offered by the MSO.

Nsub is the total number of subscribers within the service group

Tavg is the average bandwidth consumed by a subscriber during the busy-hour

2. Tavg_sg – Operational Considerations

There are multiple considerations for the Tavg_sg component of the new formula. For people in operations, Tavg_sg is more easily attained than either the Tavg or Nsub component in the old formula. The base formula (3) is something that they can measure on a SG by SG basis. Once the QoE margin is established, then this formula can be used as a basis to determine when a SG approaches its maximum capacity before it must be split. So Tavg_sg has operational advantages.

For network capacity planning, it may be desirable to predict required capacity for different SG with different service tier mixes. The generic $(Nsub * Tavg)$ falls short in this respect. By expanding Tavg_sg and breaking out data utilization by service tiers or other groupings, required capacity can be estimated by:

$$C \geq Tmax_max + \sum Nsub(i) * Tavg(i) + QoE_margin \quad (5)$$

where:

C is the required bandwidth capacity for the service group

Tmax_max is the highest Tmax offered by the MSO

Nsub(i) is the number of subscribers on the i^{th} service tier

Tavg(i) is the average BW consumed per sub during the busy-hour on the i^{th} service tier

QoE_margin is additional margin required due to data utilization fluctuations

Referencing back to Table 1 shows an example of different Tavg for different service tiers. When combined with the service tier distribution, then a weighted average can be calculated to find Tavg_sg.

As can be seen, the Traffic Engineering formula can adapt as needed. Either using Tavg_sg when appropriate, or decomposing it down to individual service tier components.

3. QoE Margin – the Magic Delta

Tmax_max and Tavg_sg are relatively straightforward components, so that leaves most of the complexity for future traffic engineering research on the QoE margin. This is the AC component, or the Delta

bandwidth on top of the static bandwidth utilization. Our job is to take the magic out of the Delta bandwidth.

3.1. “Simple” Formula – still valid after all this time

The “Simple” formula is still valid. As discussed above, it maps quite well to the new base formula where $QoE_margin = 0.2 * T_{max_max}$. However, it is just as important to understand its limitations. It works quite well for SGs with a couple hundred subscribers. The “Simple” formula may be a bit of overkill as the SG size shrinks to 100 subs or less.

Its simplicity is its strength. It is well suited for planning and quickly getting a ballpark estimate of capacity needed. In a golf analogy, think of it as the drive that gets you much closer to the hole. Operators should feel comfortable to continue to use this formula.

3.2. Operational Thresholds

As mentioned above, T_{avg_sg} is often easily measured on a SG by SG basis and is an important metric from an operational perspective. But what should be used for the QoE margin component with it?

One thing to consider from our data analysis is standard deviation, or more specifically the coefficient of variation. Depending on how much margin an operator might want to build into their system, they might consider adding two to three standard deviations as their QoE margin. The standard deviation may come from CMTS monitoring tools or research as described earlier in this paper.

Our early research has shown that a 100 sub SG might have a coefficient of variation in the 35% to 50% range. 2-3 standard deviation would then mean that anywhere from 70% to 150% of the measured T_{avg_sg} would be added as the QoE_margin . For example, if this SG measured $T_{avg_sg} = 100$ Mbps, then QoE_margin would be between 70 and 150 Mbps depending on the coefficient of variation used and the number of desired standard deviations.

As another example, consider a Remote PHY CCAP Core port that supports 1000 subscribers. Suppose that it has a measured $T_{avg_sg} = 1.2$ Gbps. From our early research, its coefficient of variation might be 10%, so three standard deviations would require an additional 0.36 Gbps for the QoE_margin .

Instead of using standard deviations, an operator might use a percentile level (e.g. 98%, 99%, 99.9%) to determine the $T_{avg_sg} + QoE_margin$. This might be accomplished with CMTS monitoring tools that provide a histogram for a SG similar to the results in Figure 11. The QoE_margin would equal the difference between the selected percentile level (e.g. 98%) and the measured T_{avg_sg} .

3.3. Big Data Analytics

What our research has found is that there is a massive amount of data and many complicated variables at play here. It turns out that providing sufficient QoE for traffic engineering is a problem that is suited to Big Data Analytics. This work is still in its infancy. Our goal is that Big Data Analytics can be leveraged to not only select optimum QoE margins in existing networks, but become a tool to predict how our networks will morph and the QoE margins of the future.

Conclusion

Many new innovations are finding their way into cable operators' plants and many others are on the way, including DOCSIS 3.1, Remote PHY Distributed Access Architectures, and Fiber Deep networks with DOCSIS FDX. All of these will significantly impact how operators manage their traffic engineering and network capacity planning.

To enhance our Traffic Engineering formula, an intimate understanding of subscriber bandwidth behavior was needed. This paper took a detailed look at a year's worth of live consumer data collected from a single cable site. The massive amount of data took samples during peak busy hour and tracked every packet to allow for traffic analysis down to the second.

Statistics were gathered and many different bandwidth trends uncovered. Some of the key variables of interest include traffic consumption based on:

- Differing Service group sizes – size matters; significantly increased variation for small SG
- SG to SG variation – substantial for small SG
- Subscriber service tiers – you get what you pay for: higher Tavg for higher tiers
- Time of day – peak busy window stretches to 2-3 hours
- Day of week – not much difference but Sunday is the busiest while Saturday is the quietest
- Month to month – erratic growth, but in line with industry's 35% CAGR

With these trends in hand, the impacts on existing network capacity models showed how they might morph traffic engineering in a Fiber Deep Gigabit world. A new basic formula evolved into:

$$C \geq T_{\max_max} + T_{\text{avg_sg}} + QoE_margin \quad (3)$$

where:

C is the required bandwidth capacity for the service group

Tmax_max is the highest Tmax offered by the MSO.

Tavg_sg is the average bandwidth consumed by a service group during the busy-hour

QoE_margin is additional margin required due to data utilization fluctuations

Tmax_max and Tavg_sg are relatively straightforward. Tavg_sg is often easily measured on a SG by SG basis and is an important metric from an operational perspective. Tavg_sg replaces the $N_{\text{sub}} * T_{\text{avg}}$ component from the older "Simple" formula. That leaves most of the complexity for future traffic engineering research on the QoE margin. This is the AC component, or the Delta bandwidth on top of the static bandwidth utilization. Our job is to take the magic out of the Delta bandwidth.

The "Simple" traffic engineering formula is still as valid as ever. It provides an easy method to quickly get bandwidth capacity estimates. However, the newer traffic engineering formulae have been developed to provide more accuracy and handle a wider range of conditions from small Fiber Deep SG to very large CCAP cores.

There are different ways to estimate the QoE margin. This paper discussed several of these. Some use statistical measurements for a SG such as standard deviation or the coefficient of variation as well as percentiles. These statistics might be derived from CMTS monitoring of that SG or from analysis from a very large collection of data over time. Some examples were shown. It is important to note that the

operator can choose how much margin they would like to build in. This may change from region to region based on a particular country's regulatory environment.

What our research has found is that there is a massive amount of data and many complicated variables at play here. It turns out that providing sufficient QoE for traffic engineering is a problem that is suited to Big Data Analytics. This work is still in its infancy. Our goal is that Big Data Analytics can be leveraged to not only select optimum QoE margins in existing networks, but become a tool to predict how our networks will morph and the QoE margins of the future.

Acknowledgements

The authors would like to gratefully acknowledge the assistance of Ben Widrevitz for his countless hours of collecting the live data and processing the raw data. Without that, this paper would not have been possible.

Bibliography & References

[CLOONAN_2013] Tom Cloonan, Jim Allen, Tony Cotter, and Ben Widrevitz, "Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS," Proceedings, The NCTA Cable Show Spring Technical Forum (June, 2013).

[CLOONAN_2014NCTA] Tom Cloonan, Mike Emmendorfer, John Ulm, Ayham Al-Banna, and Santhana Chari, "Predictions on the Evolution of Access Networks to the Year 2030 and Beyond," Proceedings, The NCTA Cable Show Spring Technical Forum (April, 2014).

[CLOONAN_2014EXPO] "Predictions on the Evolution of Access Networks to the Year 2030 & Beyond"; T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, S. Chari, The Cable Show NCTA/SCTE Technical Sessions Spring 2014

[CLOONAN_2015] "Lessons from Telco and Wireless Providers: Extending the Life of the HFC Plant with New Technologies," Tom Cloonan et. al., The NCTA Cable Show Spring Technical Forum, May, INTX 2015

[CLOONAN_2016] "Using DOCSIS to Meet the Larger Bandwidth Demand of the 2020 decade and Beyond" Tom Cloonan, Ayham Al-Banna, Frank O'Keeffe; The NCTA Cable Show Spring Technical Forum, INTX 2016

[EMM_2014] "Nielson's Law vs. Nielson TV Viewership for Network Capacity Planning," Mike Emmendorfer, Tom Cloonan; The NCTA Cable Show Spring Technical Forum, April, 2014

[ULM_2014] "Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning", John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo

[ULM2_2014] "Scaling Traditional CCAP To Meet The Capacity Needs Of The Next Decade" John Ulm, Tom Cloonan; 2014 SCTE Cable-Tec Expo

[ULM_2016] “Giving HFC a Green Thumb: A Case Study on Access Network and Headend Energy & Space Considerations for Today & Future Architectures” John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

[ULM2_2016] “Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies”, John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

Abbreviations

BAU	Business as Usual
Bcast	Broadcast
Bps	Bits Per Second
CAA	Centralized Access Architecture
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
DAA	Distributed Access Architecture
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DWDM	Dense Wave Division Multiplexing
E2E	End to end
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQAM	Edge Quadrature Amplitude Modulator
FD	Fiber Deep
FDX	Full Duplex (i.e. DOCSIS)
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTP	Fiber to the Premise
FTTT	Fiber to the Tap
FTTx	Fiber to the ‘x’ where ‘x’ can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
HFC	Hybrid Fiber-Coax
HP	Homes Passed
HSD	High Speed Data
I-CCAP	Integrated Converged Cable Access Platform
IEEE	Institute of Electrical and Electronics Engineers
IEQ	Integrated Edge QAM
LDPC	Low Density Parity Check FEC Code
MAC	Media Access Control interface
MACPHY	DCA instantiation that places both MAC & PHY in the Node
Mbps	Mega Bits Per Second
MDU	Multiple Dwelling Unit
MHz	Megahertz

MSO	Multiple System Operator
N+0	Node+0 actives
Ncast	Narrowcast
NFV	Network Function Virtualization
NSI	Network Side Interface
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiplexing Access (Upstream)
OLT	Optical Line Termination
ONU	Optical Network Unit
OOB	Out of Band
OPEX	Operating Expense
OTT	Over the Top
PHY	Physical interface
PNM	Proactive Network Maintenance
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RF	Radio frequency
R-OLT	Remote OLT
RPD	Remote PHY Device
R-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RX	Receive
SDN	Software Defined Network
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
TaFDM	Time and Frequency Division Multiplexing
Tavg	Average bandwidth per subscriber
Tmax	Maximum Sustained Traffic Rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
VOD	Video on demand
WDM	Wavelength Division Multiplexing

Are we done yet?

Opportunities in Wi-Fi With 60 GHz

A Technical Paper prepared for SCTE•ISBE by

Carol Ansley

Counsel, Senior Director
ARRIS
3871 Lakefield Dr.
Suwanee GA USA
678-473-2000
carol.ansley@arris.com

Charles Cheevers

CTO CPE Products
ARRIS
3871 Lakefield Dr.
Suwanee GA USA
678-473-2000
Charles.cheevers@arris.com

Introduction

As the 2.4 and 5 GHz spectrum used for Wi-Fi gets more crowded, another unlicensed band that can reliably support multiple Gigabit transmission in the home has promise for home networking and other applications. The 60 GHz band offers a wide bandwidth with little interference from other sources. This paper gives an overview of 60 GHz and WiGig, also known as 802.11ad, and compares testing results with simulations. This paper also covers the upcoming revision still in progress with the IEEE called 11ay, highlighting some new features that will enable outdoor use cases for this versatile technology.

1. Where Do We Stand with Wi-Fi?

Devices using Wi-Fi for data communication encompass every area of technology and many different services. As an example, let's concentrate on using Wi-Fi for video for a moment. Wireless set-top boxes have been on the market for a few years now, yet increasingly people are streaming video to anything with a screen: smartTVs, mobile devices like smartphones or tablets, even a refrigerator. A recent news announcement put the amount of Wi-Fi traffic providing streaming video services at 65%. People also expect to be able to check the weather, stream music, and see who just rang the doorbell, all using almost-ubiquitous Wi-Fi coverage. New uses for Wi-Fi will probably include virtual and augmented reality programming, whether games or scripted entertainment. As was mentioned in our paper last year [1], Virtual Reality (VR) headsets can consume far more bandwidth than any of today's video services.

1.1. Crowded Spectrum, Busy Networks

All this growth has led to two complimentary problems: bandwidth congestion driven by the sheer volume of traffic as well as protocol level congestion caused by the enormous number of devices competing for airtime. Video services in particular demand large amounts of data. Depending upon the device and its distance from its AP, an HD video stream may take up substantial amounts of airtime. For example an 802.11n device streaming a 5 Mb/s video program from an access point (AP) in the same room will take up about 5% of a 20 MHz channel. If the device is a couple rooms over, that percentage could rise to 40 or 50% of the channel.

Aside from the stress of the amount of data, the simple presence of large numbers of devices can place a substantial load on a Wi-Fi network. One paper showed that the presence of more than 25 devices on a single AP can reduce the overall throughput even if the amount of traffic is low.[2] The numbers of Wi-Fi devices are being driven by the increasing number of auxiliary devices incorporating wireless communications. It is useful to note that even if a group of Internet of Things (IoT) devices is not Wi-Fi, but Zigbee or BTLE, traffic on those networks can still contribute to the noise on the 2.4 GHz band reducing the overall channel availability for Wi-Fi. Many IoT devices do not have high bandwidth usage and are battery powered. However, the highest bandwidth devices, such as webcams and other video originating or terminating devices, are commonly set up over Wi-Fi.

1.2. Expanded Competition from Commercial LTE-related Devices

Most of the major telcos in the USA have at least announced MuLTEFire, LTE-U, or LAA trials. These new systems attempt to recapture mobile traffic now commonly redirected to Wi-Fi networks in the home and elsewhere back onto the telco networks. A LTE-U cell typically uses bandwidth within the 5 GHz unlicensed band to augment the current licensed bands. Deployments of this technology are still limited, but if it becomes widespread, the congestion within the 5 GHz band is certain to increase.

2. 60 GHz Wi-Fi Offers Interesting Opportunities

With all of the other WLAN activities in the popular 2.4 and 5 GHz bands, why is the 60 GHz band getting attention? First and foremost, the 60 GHz Industrial Scientific Medical (ISM) band supports unlicensed access across 14 GHz of spectrum in the US, from 57 GHz to 71 GHz. The band from 64 GHz to 71 GHz was just added last year.[3] The FCC has also made a proposal to expand this band even further. The IEEE802.11ad specification supports three 2.16 GHz channels in North America, and four channels in Europe. This is more raw bandwidth than the other unlicensed bands put together. The frequency is very high, so propagation models are more challenging, but modern technology is well able to provide good throughput for many applications. Some early demonstration units with this technology were not very robust with respect to motion or temporary blockages of signal. Testing with the latest units has shown more robust results that we will discuss later.

Two related advantages can be attributed to the higher levels of attenuation in the 60 GHz band. Interference is much reduced, particularly from other active WLAN systems. This advantage compares favorably with the current 2.4 and 5 GHz bands which are often almost unusable in multiple dwelling units (MDUs), because of the large number of closely spaced APs that compete for bandwidth. Security is also increased over the other bands, because stray power from the 60 GHz system is unlikely to make it out of a home. As most users have gotten the message that securing their Wi-Fi network is desirable, this issue is not as concerning as it once was. Even so, since many users use only weak passwords or weak encryption, overall security can be improved if signals from a user's WLAN are less likely to leave their premises.

2.1. 60 GHz Propagation and Antennas

To understand the deployment tradeoffs and advantages of 60 GHz, we need to understand 60 GHz propagation and device antenna characteristics. Propagation characteristics of 60 GHz signals are not the same as the 2.4 or 5 GHz signals that we have all become familiar with. Also, the higher frequencies necessitate a completely different antenna design approach to achieve optimal performance. The next paragraphs will discuss propagation, then antenna design.

Millimeter wave propagation, as 60 GHz is also known, is very different than the 2.4 and 5.5 GHz frequency bands. The high frequency of the radio waves means that a transmission in free space is attenuated more quickly than at lower frequencies. A more challenging aspect is that one band of the 60 GHz spectrum is also absorbed strongly by oxygen molecules. A bit of good news is that the recent extensions enacted and proposed by the FCC are above the band that is most strongly absorbed by oxygen and should provide better performance. Most solid materials tend to reflect or absorb 60 GHz transmissions as opposed to the lower frequencies where transmission through solid materials was less highly attenuated. In a later section we will discuss our testing results characterizing the performance of 60 GHz transmissions in residential environments. Because of the wide bandwidth and power limits, 60 GHz transmissions can still provide acceptable performance within one or two rooms in a residential environment.

A bit of background on the design of millimeter wave antennas is helpful to understand some of 60 GHz strengths and weaknesses. A optimal single antenna supporting the 60 GHz band is very small, less than 2mm on a side for a patch antenna, and does not provide enough directivity or focus to be useful in most applications. To compensate for that fact, 60 GHz systems typically use antenna arrays. The size and configuration of the antenna array determine the performance of the array.




Figure 1 - Examples of antenna arrays, not to scale

The following simulations illustrate a set of results showing the increased focus that results from an increase in array elements. For a 4x4 array, the gain of the antenna is concentrated in a main lobe providing about 10dB of gain over a single antenna, as shown in Figure 2.




Figure 2 - 4x4 Element Array, X-Z and Y-Z Polar Plots

If the size of the array is increased to 8x8, the gain also increased to 16 dB over a single patch antenna. See Figure 3. Note that the tradeoff for these high gain arrays is that the 3dB beamwidth of the main lobe

of the antenna pattern decreases as the gain increases. The 3dB beamwidth of an antenna pattern is defined as the angle of arc within which the antenna pattern's gain declined by 3dB, which is to the angle of arc over which the antenna's transmit power declines by half. The total power transmitted by a device is restricted by FCC regulations, so the increased relative gain thus effectively comes at a cost in the area covered by the beam from that antenna. The behavior is discussed in terms of using the antenna array as a transmitter, yet the same effects also apply when it is used as a receiver.




Figure 3 - 8x8 Element Array, X-Z and Y-Z Polar Plots

While beam steering can compensate by moving the focus of the array in a particular direction, a planar array's beamforming shift is limited typically to 120°.




Figure 4 - Beamforming Example



Figure 5 - Example 11n or 11ac Wi-Fi AP

The beamforming limits mentioned above are significant in the industrial design of a 60 GHz AP or client because the optimal placement of a 60 GHz AP may be quite different from the optimal placement of a traditional Wi-Fi AP in the 2.4 and/or 5 GHz bands. A typical 11n or 11ac Wi-Fi AP may have four or more antennas to support good coverage; together they generally cover 360° in at least one plane. Most

current Wi-Fi APs provide their best coverage when placed in the center of a home so that the Wi-Fi signal can radiate evenly in a more or less spherical fashion.

A 60 GHz device's optimal placement will be heavily influenced by its antenna design. A device with a single antenna array may perform best placed in a corner of a room, or near the corner of a home so that the potential targets of the antenna are within 90° (at least in one plane) of the center of the antenna array.

For a device to be capable of reaching a client anywhere in the plane of the antenna would require three separate fixed antenna arrays to cover 360°. When similar technology has been used in other applications, such as radar, the antenna array is often constructed to spin, so that a single array can track targets spread across a full 360° horizon. This approach is probably not practical in a residential gateway. For a practical 60 GHz deployment, multiple antenna arrays are also probably undesirable since the costs could become prohibitive.

2.2. Residential Testing

We sponsored testing of 60 GHz equipment in a residential environment to determine how useful the technology can be outside of the lab or simple desktop applications. An 802.11ad AP was placed at various locations within a home and the actual throughput to a laptop equipped with 802.11ad was tested. The AP's location was tested in several locations shown on the diagram below.




Figure 6 - Schematic of Test House Showing AP Locations

Location 1 was along an unobstructed wall facing across the living room. The AP was placed on a small table. The table was moved to a corner of the living room for Location 2 as shown in the diagram, and then the AP was tested without any obstructions and again with a large chair placed in front of the table. Also shown on the diagram are the location of a refrigerator and a wall-mounted flat screen television, both of which were found to affect the radiation pattern.

The rooms were divided into grids for testing. Throughput was repeatedly measured from the AP to a laptop as the laptop was placed in grid locations across the rooms. A few locations were also used on the second floor. Rather than show tables of the results, diagrams following show the throughput results from the AP covering the family room and the adjoining kitchen and side room using a color-coded representation of the throughput results for that location. The color coding relates to the throughput measurement and the signal level as shown in Figure 7.




Figure 7 - Legend for Color-coded Diagrams

As an example, a yellow block in the throughput diagrams indicates that the signal level was about -70dB and the throughput was about 900 Mb/s. The test system used for these measurements was 2 years old. A newer system could probably achieve higher throughput than the one used for this testing. The table comparing throughput to signal level and RSSI indicates that at the highest signal levels, the system was limited by its GigE Ethernet port, not by its radio interface. As newer results become available, we might submit the latest results.



Figure 8 - Throughput Measurements for Location 1

The throughput results for Location 1 show that the good coverage was achieved in the living room, with data rates still above 100 Mb/s even in NLOS locations in adjoining rooms. A location that probably fell outside of a direct propagation path can be seen just below the dark red square. As was mentioned earlier, beamforming with a planar array has blind spots. Due to the large amount of reflections, the throughput in that area was still over 700 Mb/s. Follow up simulations showed that 60 GHz signals are reflected strongly by many common in-home building materials leading to good NLOS coverage.

The picture below shows a ray tracing example simulating the coverage due to an AP in Location 1. Testing has shown that sheet rock has a relatively high permeability to 60 GHz transmissions, while brick or cinder block walls tend to reflect most of the energy. That difference can be seen in Figure 9 where energy can cross the sheet rock wall between the living room and the kitchen, yet transmissions hitting the outside walls, which have a brick facing, are strongly reflected.




Figure 9 - Raytracing Example for Location 1

In this image, 2nd floor has been kept invisible to observed the ray propagation in the 1st floor.

The next diagram shows the measurements taken in Location 2, near the corner of the room. Figure 10 shows better coverage in the kitchen, as well as coverage in the upstairs bedroom, which was marginal when the AP was in Location 1. Probably due to the AP's change in angle to the TV, there is now a definite slow spot behind the television. In the test with location 1, the two open doorways on either side of the TV allowed reflected rays to provide good coverage. After the AP shifted to the corner of the room, the farther doorway was not at a good angle for reflections, and the television prevented most through wall transmissions, resulting in the lull behind the TV.




Figure 10 - Throughput Results for Location 2 Without Chair

The results agreed with antenna theory expectations that for minimum blind spots, a corner location for an AP is optimal. With the AP still in Location 2, an overstuffed chair was placed in front of the table holding the AP and the tests were repeated. As shown in Figure 11, the chair absorbed enough energy to lower throughput levels throughout the room, but also drove some reflected energy into the side room.




Figure 11 - Measurement Results for Location 2 With Chair

Summing up the residential test results:

1. Achieving consistently high throughput is very feasible as long as one takes the propagation characteristics of millimeter wave antennas into consideration.
2. The principal obstacles to good transmission are metal or metal backed objects and stone or cement, though they can provide good reflections for NLOS paths.
3. We noted with some surprise that wood and cloth furniture seemed to affect 60 GHz signals more readily than expected.

With these factors in mind, the optimal placement for 60 GHz APs may be as a wall mounted device to give it a better likelihood of being above furniture that might degrade the signals. A creative mind might integrate a 60 GHz AP into a wall mount lighting fixture or shelf designed to be mounted in a corner.

3. Potential Residential Use Cases

The early 11ad demonstrations and products have focused on the elimination of wires, usually in an office or enterprise context. Up until recently, the marketplace for 60 GHz products has been limited as the vendor community struggled to find a niche that fit the strengths of 60 GHz systems. That state of affairs has begun to change with the first general availability AP with 11ad support being released last year. WiGig holds promise in a residential environment for several features or services, with inter-AP and local

access backbones and virtual reality potentially the most significant. These use cases rely on the high bandwidth offered by 60 GHz systems as well as the low added latency of the 11ad MAC.

3.1. 60 GHz Home Network Backbones

As today's consumers bulk up on wireless devices, the traditional single home router is having trouble addressing the expectation that Wi-Fi can be everywhere from the garage to the backyard to every room in the home. Many operators and consumer electronics manufacturers are considering how to provide added coverage that may be needed in other parts of the home. Extender APs and repeaters are popular, but they must still be connected back to the main home router. If a wired connection path is available, that can provide the highest reliability, but often there is not a wired connection can sustain a high bit rate connection, if it's available at all.

Wireless extender APs are showing up on the market, but if they use a standard 5 GHz backbone connection, they are only adding to the congestion already present in the air. A 60 GHz backbone has the advantage of providing a high bit rate connection without impacting the existing services in the home. The testing done to this point has shown that 60 GHz multi-element arrays can provide enough signal to get through at least 2 sheet rock walls and still provide at least 1 Gb/s of service.

It is also important to remember that in-home testing and simulation has also shown that non-line-of-sight connections can be significant for 60 GHz transmissions. Impediments to using a 60 GHz backbone are related to home construction materials. In parts of the world where interior walls are commonly made of cinderblocks, 60 GHz backbones will struggle, just like 5 GHz systems.

Since 60 GHz connections are more dependent on LOS and NLOS reflections than lower frequency Wi-Fi connections, a backup method may be needed for 60 GHz connections that are affected when the home's configuration changes. For example, a bedroom might get enough reflected energy to achieve a useful bit rate when the bedroom door is open, but it might struggle when the door is closed. If the AP(s) can recognize that issue and reconfigure the bedroom's connections to use lower bit rate 5 GHz channels, the end user might experience a lower bit rate connection. However, a change in performance is better than a complete disconnection. Similarly, if there are several 60 GHz APs within the home, they may be able to shift their beamforming to work around changes in the home's physical configuration in real time to continue to provide uninterrupted services.

3.2. Last 100ft Broadband Access over 60 GHz

Another potential use case related to inter-AP backbones is the need for a fixed wireless extension to the home from a local broadband termination facility. That facility might be a fiber node or a strand-mounted DOCSIS 3.1 cable modem. A 60 GHz distribution system could feed high speed connections to outdoor antennas mounted on nearby homes. In particular, the new 11ay specification amendment underway with IEEE 802.11 has new features designed to improve outdoor performance.

Outdoor 60 GHz systems have several challenges with which to contend. An outdoor distribution system has to overcome water as fog, rain, or snow, as well as ice potentially collecting on outdoor antennas. A broadband to wireless distribution node may have enough internal heat generated to avoid some problems with snow or ice, while leaves and tree branches may also block or absorb transmissions. In areas with buried utilities, the other barrier to using 60 GHz services may be the need to get high enough to get a substantially LOS view of the distribution node or nodes to get the best performance. While within a room or in rooms connected by hallways, reflections may be counted on to provide a NLOS path to

hidden nodes, reflections out of doors are as likely to result in energy being reflected into the sky as back in a useful path, making outdoor use cases more challenging than indoor use cases.

To overcome these challenges, complex element arrays that can provide high levels of gain and directivity will be important. Depending upon the network architecture, a distribution node may require multiple antenna arrays to cover different angles. Comprehensive network designs will be needed, similar to the designs of cell sites now, to ensure that nodes with potential overlapping coverage can operate on different channels. An integrated operational management system may be needed to ensure that the network adapts to changing conditions, and that the stations, which may themselves have multiple radios, are kept up to date with their recommended node and channel usage.

3.3. Virtual Reality Needs 60 GHz

Virtual Reality systems require significant amounts of data to generate and maintain the VR illusion. The data must also be provided with low latency. As an example, if a person turns their head, the viewing area must be redrawn within at most 20 milliseconds to prevent the inner ear from disagreeing with the eyes, which can lead to nausea. This requirement is also known as Motion to Photon, MTP. Some types of entertainment with lots of fast motion, such as live sports or gaming, may require even lower latency than 20 milliseconds to avoid the appearance of stuttering video during the active portions of the program, particularly if combined with the user's motion at the same time. The amount of data required to reach that level of performance is still up for debate, but it may be reliably estimated at above 1 Gb/s. VR can require these rates across several steps. If a VR program is streaming from the wide area network, it will drive high data rates over the broadband access facility to the VR controller. The number of cost effective options supporting multi-gigabit throughput is small for home networking solutions. Wireless connectivity using 11ad is an option that gets the high bitrate streaming feeds to the VR controller. Additionally, a wireless VR headset would also need a high bit rate solution. The headset/controller link is needed even if the VR content is from a local source, versus content streaming from the WAN.

If there is more than one user of VR in the same room or if the VR data is passing over two hops (from the WAN to the controller and from the controller to the headset), the rates of course double. Current Wi-Fi systems, while they can reach rates above a gigabit in good conditions, struggle to maintain high throughput levels if other devices are also on the same network. If the systems are in use in an environment with many other overlapping APs with their own demanding clients, it is very unlikely that they can keep up with even one high quality VR experience.

802.11ad systems can provide very low latency and high bandwidth without interference, making them ideal for VR applications. While 11ac systems might be able to keep up with a VR transmission data rate, in many homes they may have trouble achieving the low latency required by VR because of interference from surrounding Wi-Fi systems as well as range and power difficulties serving a very high bit rate to a VR headset that may not be located in the same room as the serving AP. Even if 60 GHz becomes wildly popular, it is unlikely to ever experience the same interference problems currently afflicting conventional Wi-Fi because of the high level of attenuation when a 60 GHz signal tries to traverse an outside wall of brick or stone. Signals from a neighboring building will certainly not get through outside walls with enough energy to disrupt a system in an adjacent building. Even signals from a neighboring apartment are unlikely to line up at just the right angle to interfere with a system in an adjacent apartment since the antenna arrays are highly directional.

Overall, 60 GHz wireless connections have real promise for VR entertainment systems because of their high bandwidth, low latency, and resistance to interference.

4. Next Generation of 60 GHz Wi-Fi – 11ay

In 2015, the IEEE 802 standards group started a new effort to expand the capabilities of the 60 GHz Wi-Fi interface introduced with 11ad. The goal of the group is to increase throughput to at least 20 Gb/s, while maintaining backward compatibility to the current 11ad amendment. The specification is still in progress with work expected to complete in 2019. A new proposed feature provides the ability to transmit to multiple devices on multiple channels simultaneously; channel bonding to a single device is also supported. To support the higher speeds, wider channel definitions are have been proposed along with downlink MU-MIMO.

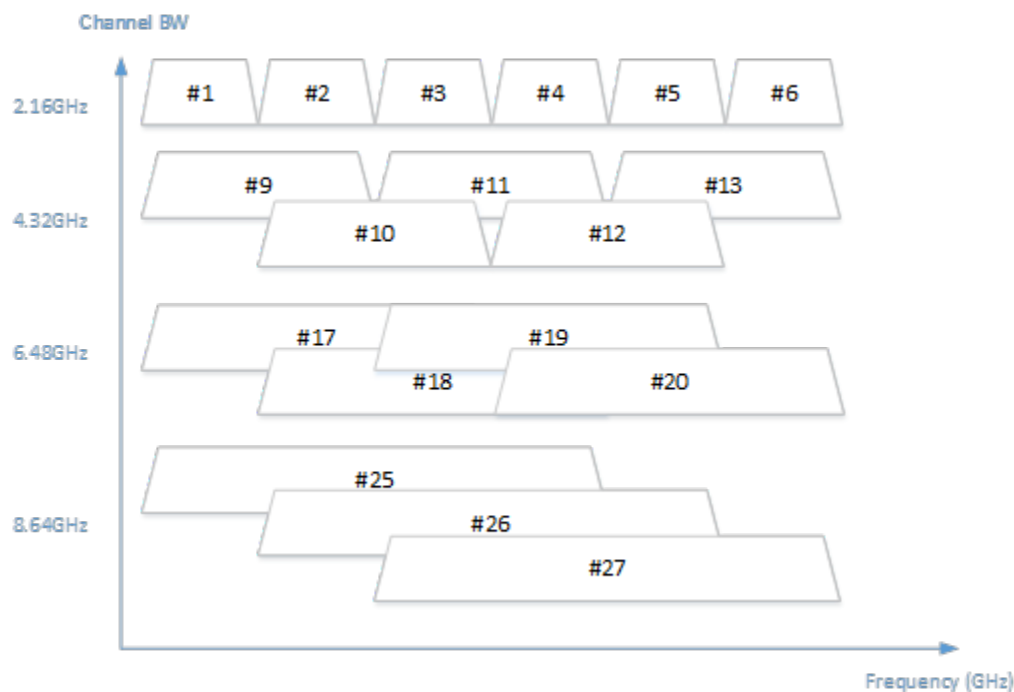


Figure 12 - Proposed Channels for 802.11ay[4]

The high throughput expectations of the 11ay effort have led some to question whether there is a need for such a service. In the timeframes of 11ay, Fiber-Deep and/or DOCSIS 3.1 access technologies may bring to the home data speeds approaching this level. As well in-home services, such as gaming consoles or other entertainment may be supporting VR with its very high throughput demands. Current home wireless networks are already strained with the current service demands, primarily streaming video. The homes of the future are unlikely to require less bandwidth or support fewer services. The ability of 60 GHz services to provide targeted bandwidth to several users could be the linchpin of an in-home VR deployment. No other potential home networking technologies have the mix of high throughput and resistance to interference as 11ay.

Conclusion

60 GHz wireless has many names: WiGig, 11ad, 11ay, and millimeter wave. No matter what label is used for this wireless technology, it can provide real advantages to residential home networking. The current

generation of 11ad systems demonstrates high throughput and good robustness to common home characteristics. The lack of interference means that it can immediately improve the wireless experiences of many consumers in MDUs or other congested areas who are frustrated with the congestion caused by local interference.

As broadband services provide higher multi-Gigabit service rates, higher bandwidth wireless will be needed to provide those high speeds to wireless devices. The 5 GHz band has difficulty providing those high speeds, because of the bandwidth limitations in many regions and the high numbers of other 5 GHz devices. Those multi-Gigabit speeds can be met by 60 GHz solutions, both last 100ft outdoor links, and in home solutions.

A potential driver of multi-Gigabit services, VR, is also a good fit for 60 GHz networking due to the high speeds and low latency of 60 GHz Wi-Fi. The possibilities for 60 GHz Wi-Fi, or WiGig in the residential and home networking space are numerous and compelling.

Abbreviations

AP	Access Point
AR	Augmented Reality
BTLE	Bluetooth Low Energy
GHz	Gigahertz
HD	High Definition
IoT	Internet of Things
ISM	Industrial Scientific Medical
LOS	Line of Sight
MDU	Multiple Dwelling Unit
MU-MIMO	Multi-User Multiple Input Multiple Output
NLOS	Non-Line of Sight
SCTE	Society of Cable Telecommunications Engineers
VR	Virtual Reality
WLAN	Wireless Local Area Network

Bibliography & References

Carol Ansley, Charles Cheevers, “Advanced Wireless Possibilities”, INTX 2016.

Chuck Lukaszewski, Liang Li, “Empirical Measurements of Channel Degradation Under Load”, IEEE 15/0351r02, March 2015.

FCC report and Order and Further Notice of Proposed Rulemaking, FCC 16-89, pages 125-131, July 14, 2016.

Figure 24, Specification Framework for 802.11ay, 11-15-1358-09-00ay, Oct. 8, 2016.

Automation of the Best Practices used to Evaluate 802.11 Access Network

A Technical Paper prepared for SCTE•ISBE by

David Brownell, Shaw Communications

Salman Naqvi, Shaw Communications

Introduction

As the breadth of 802.11 standards increases to meet market demands and convergence with other technologies, the amount of capabilities provided by Customer Premise Equipment and Access Point (AP) devices is increasing dramatically. This coupled with introduction of interpretations by vendors for new and evolving standards places extreme pressure on service providers who endeavor to ensure the highest quality metrics for their network are maintained and enhanced by the new product introduction.

Typically, service providers will rely on the expertise of their engineering teams to vet the new products against the network requirements. The level of test coverage required and the turn-around time to deploy in the market bring in its challenges. Hence, automation of test coverage methodology is necessary to meet these demands.

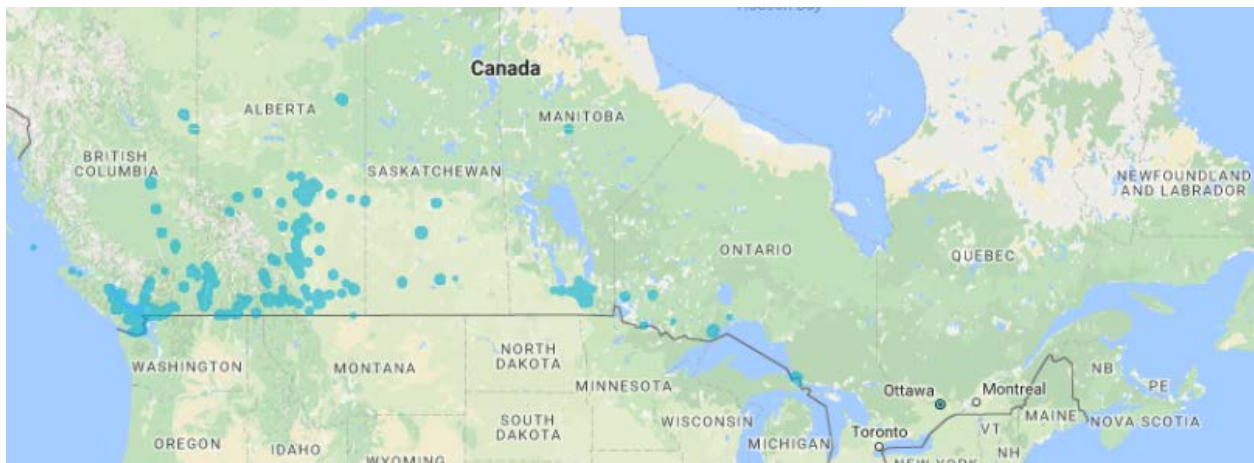
This paper will address the implementation of the process and methodology applied in identifying the Key Performance Indicators to evaluate the 802.11 Access Network. There will be a brief account describing the test cases used and their importance to 802.11 service provider like Shaw Communications. The paper will also describe the challenges and benefits that automation brings to this subject. The test coverage will include the SW/HW tools used to test the full functionality of the network from layer 1 through 7. Based on these results the Quality Assurance (Q) engineering team at Shaw Wireless Lab can provide a set of guidelines to the deployment engineering team, for better deployment of the 802.11 Network.

Content

1. SHAW COMMUNICATIONS

1.1. SHAW WiFi Network

Shaw Communications Inc. is an enhanced connectivity provider. Our Consumer division serves consumers with broadband Internet, Shaw Go WiFi, video and digital phone. Our Wireless division provides wireless voice and data services through an expanding and improving mobile wireless network infrastructure. The Business Network Services division provides business customers with Internet, data, WiFi, telephony and video.




Shaw is traded on the Toronto and New York stock exchanges and is included in the S&P/TSX 60 Index (Symbol: TSX - SJR.B, SJR.PR.A, SJR.PR.B, NYSE – SJR, and TSXV – SJR.A). For more information, please visit www.shaw.ca.

The Shaw network has a more than 80 thousand Shaw Go WiFi Hotspots across Canada.

Shaw offers the following products that utilize WiFi technology:

- Shaw GO WiFi - Launched in 2012 for Shaw Cable and Internet subscribers.
- Managed WiFi - Launched in 2014; Targeting Hospitality
- Smart WiFi - Launched in 2016; Targeting SMB customers and part of Shaw SMART services including Smart Voice and Smart Security.
- Shaw also has a strong presence in Home WiFi products.

1.2. WiFi Technology Roadmap



Shaw Communications network utilizes the latest WiFi technologies in our networks. The latest WiFi deployments for Shaw field the following technology advancements:

- 802.11n
- 802.11ac
- 802.11ac Wave 2 MU-MIMO
- HotSpot 2.0

The 802.11 standard has been ever increasing in scope and has active working groups for 8 different 802.11 standards. One of the more interesting standards expected to be available in the next 2 years is the 802.11ax “High Efficiency WLAN”.

2. Benefit to Shaw for Test Automation

The benefit for Shaw creating an AP test automation implementation is that it fully exercises and measures OSI layer 1 to 4 performance within a finite cycle time. Automation provides consistent repeatable measurements that would not be possible manually and can be run with minimal training.

Automation also has significantly improved the test time from 8 weeks manually testing to 2 weeks for automated testing. In addition, the test coverage has been significantly increased from less than 30% to over 80% with automation.

Some examples of issues found prior to deployment into the Shaw Production Network are as follows:

1. AP displaying high RF power on UNI-1 band exceeding RS245 specification.
2. AP displaying high RF levels of spurious noise in the transmit channel band on AP output.
3. AP displaying Poor EVM modulation performance for high MCS rates at higher RF power settings.
4. AP Beacon modulation rates not aligned with minimum data rates.
5. AP not tuning to some RF channels.
6. AP using UNI -1 frequency range for outdoor model not allowed in Canadian domain.
7. AP no longer forwarding DHCP to clients after several connection cycles.

8. AP candidate firmware revision reducing throughput performance compared to baseline firmware load.

In the case of most issues found with AP performance, we provided detailed feedback and results to the vendor who could address and resolve the issues with firmware releases.

Without the automation capability, these issues may of not been found until the AP was deployed in our production environment and the containment and resolution of the issues would obviously be costlier, time consuming and detrimental to the customer experience.

3. WiFi Network Requirements

The WiFi network requirements are derived from several sources and ultimately place criteria on the technical performance of the AP under test. The test requirements originate from three sources:

- Interpretation of Customer needs into technical requirements
 - Best in class vendor performance specifications
 - System design implementation guidelines
1. Interpretation of customer needs:
 - a. Easy access
 - b. High speeds
 - c. Reliability
 - d. Competitive price
 2. Best in class vendor performance:
 - a. High reliability
 - b. Feature set options
 - c. Latest speeds/spatial streams/performance
 - d. Ease of support/maintainability/fielding configuration
 3. System Design Implementation:
 - a. Overall network design
 - b. 802.11 specifications – ensuring latest technology available
 - c. Access point placement/deployment for optimal coverage/service

Deployment requirements

In addition to 802.11 technical specs, deployment guidelines also provide test requirements:

- Desired throughput – distance selected to 17-18 m between AP and user and expect MCS 5-7 downlink in good conditions based on our link analysis for typical client device performance. The perimeter also defines the typical AP power level settings, as we do not use auto power setting in some network deployments. The question we want to answer is what is the AP RF output power at the downlink MCS rate?
- AP antenna coverage – AP model antenna pattern should support deployment guidelines in directivity, and maximum angle of power. Note that the TRP and TIS measurements provided by

external labs such as CableLabs® quantify performance for TIS (receive uplink) at MCS 7, and TRP (transmit downlink) at MCS 0 only for 802.11n in accordance with the CTIA specification (ref 1). To validate deployment guidelines, we measure RF power at higher MCS for downlink for 11n and 11ac. We also measure AP beacon power as compared to higher MCS power. Beacon power is typically the power measured during site surveys and it helps knowing higher MCS power vs. beacon power to confirm our deployment design intent. The question we want to answer is what is the TRP at our target downlink MCS rate?

- AP Placement/Capacity Planning- the relative spacing deployment numbers of APs for a coverage area. We want to ensure spacing still supports adjacent channel operation between the APs. The question we want to answer is will the transmit RF performance of the AP in adjacent Channels support our AP placement for coverage?

4. Network Test Philosophy

The overall test coverage applied by Shaw in validating a network spans the entire OSI network layers and can be summarized as follows:

1. Component Level Verification of key technical performance metrics (i.e. maximum data rates, standards compliance).
2. Subsystem Level Verification for CPE network performance, example of AP with security appliance and DOCSIS modem.
3. System Level verification through use cases, and mixed traffic tests.
4. System Level verification and soak in pre-production networks. (Where preproduction is an exact copy of the Shaw production network).
5. BETA test trials with customers on the production network.

Overall network performance metrics are validated at higher system integration levels, but we find by measuring the components comprising the network with test results being directly traceable to vendor or industry specifications. This allows Shaw to engage directly with the vendors when non-conformances are found. Verification of the lower layer specifications lays a good foundation for network performance.

Given the coverage, and complexity of the standards, Shaw's approach is to use specialized test equipment and automation to realize the test coverage required. Test coverage is used to perform the first evaluation of equipment, as well as screen changes (firmware updates) throughout the life cycle of the product in the Shaw production environment.

Some of the test equipment Shaw employs for network product verification is shown in Table 1 below.

OSI Model Layers		Examples	Spirent Landslide/I XIA Breaking Point		Spirent Avalanche/IXIA Network	IXIA Veriwave*		Keysight	R&S	Siros
			RF 802.11	Ethernet/802.3	Ethernet/802.3	RF 802.11	Ethernet/802.3	RF 802.11	RF 802.11	802.3
7	Application	NFS, SNMP, Telnet, HTTP, FTP	X	X						
6	Presentation	ASCII, EBCDIC, TIFF, GIF, PICT, JPEG, MPEG, MIDI.	X	X						
5	session	NFS, NetBios names, RPC, SQL.	X	X						
4	transport	TCP/UDP	X	X	X	X	X			
3	network	IP			X	X	X			
2	MAC	802.3			X	X	X			
1	Physical Layer	RF, Ethernet.			X	X	X	X	X	X
*IXIA Veriwave have test coverage for all OSI layers, Shaw uses Veriwave for primarily layers 1 thru 7										

Table 1 - Network Performance Test Tools Used by Shaw

*IXIA IxVeriWave product line does offer test coverage thru all OSI layers, but Shaw uses it primarily for layers 1 through 4.

The general test philosophy applied to WiFi is to perform extensive test coverage at the lower layers 1-4 (channels, MCS rates, frame size etc.). With the foundation components and lower level OSI layers thoroughly tested, higher level test performance (OSI layers 4 through 7) can be validated with less test cases where it does not need to be performed for every possible permutation or channel.

WiFi Test Requirements are defined with the following criteria:

- Must be quantifiable and repeatable.
- Must be traceable to specified requirements. Either 802.11 specification and or vender published specifications.
- Must support overall Shaw requirements and deployment guidelines

For WiFi Access point tests, the direct performance standard is 802.11. Shaw has also augmented this test coverage with derived requirements, other industry standards, and best practices.

The overall WiFi test coverage is summarized in Tables 2 through 4 show traceability to standards where applicable. The automation column uses a color coding of green to indicate the tests selected for automation and currently implemented. As shown in Table 2 through 4, just 7 automated tests implement the test coverage:

- Transmit Test Coverage = 2 automated test scripts
- Receive Test Coverage = 2 automated test scripts
- Link Layer test coverage = 3 automated test scripts.

Transmit Characteristic	Requirement 802.11-2012 a/g	Requirement 802.11-2012 n	Requirement P802.11ac	Requirement, Other	Automation
Transmit Power EIRP Radiated	17.3.9.1	20.3.20.3	20.3.20.3	RSS-247	NA, Radiating TIS TER performed by external lab
Transmit Channel Power - Conducted				RSS-247	RF Characterization vs MCS vs Ordered Power Automated Test
Transmit Power Accuracy			Vendor specification		RF Characterization vs MCS vs Ordered Power Automated Test
Transmit Power Packet to Packet Variation			Characterization only	Characterization only	RF Characterization vs MCS vs Ordered Power Automated Test
Beacon Frame Power				Characterize Only	RF Characterization vs MCS vs Ordered Power Automated Test
SSID Beacon vs MESH Beacon				Characterize Only	Manual
Transmit Spectrum Mask	17.3.9.2	20.3.20.1	22.3.18.1		Conducted Emissions Automated Test
Transmit Occupied Bandwidth	17.3.9.2	20.3.20.1	22.3.18.1		Conducted Emissions Automated Test
Transmit Adjacent Channel Power				Characterize only	Conducted Emissions Automated Test
Transmit Spectral Flatness	17.3.9.6.2	20.3.20.2	22.3.18.2		Manual
Transmit center frequency Accuracy	17.3.9.4	20.3.20.4	22.3.18.3		RF Characterization vs MCS vs Ordered Power Automated Test
Transmit Symbol Clock Frequency Tolerance		20.3.20.6			RF Characterization vs MCS vs Ordered Power Automated Test
Preamble Frequency Error				Characterized only	RF Characterization vs MCS vs Ordered Power Automated Test
Modulation Accuracy – Transmit Constellation Error	17.3.9.6.3	20.3.21.7.4	22.3.18.4.3		RF Characterization vs MCS vs Ordered Power Automated Test
Spurious Noise	17.4.6.9	20.3.16	Not specified		Manual
TX Center Frequency Leakage dB		20.3.20.7.2			RF Characterization vs MCS vs Ordered Power Automated Test
TX Power Peak Excursions dB				US Code of Federal Regulations Title 47, section 15. Para 407	RF Characterization vs MCS vs Ordered Power Automated Test

Table 2 - WiFi Transmit Test Requirements

Receiver Characteristic	Requirement 802.11-2012 a/g	Requirement 802.11-2012 n	Requirement P802.11ac	Automation
Minimum Input Level Sensitivity Radiated	17.3.10.1	20.3.21.1	22.3.19.1	NA, Radiating TIS TER performed by external lab
Minimum Input Level Sensitivity Conducted	17.3.10.1	20.3.21.1	22.3.19.1	Receive Sensitivity Automated Test
Adjacent Channel Rejection	17.3.10.2	20.3.21.2	22.3.19.2	Receive Channel Rejection Automated Test
Nonadjacent Channel Rejection	17.3.10.3	21.3.21.3	22.3.19.3	Receive Channel Rejection Automated Test
Receiver Maximum Input Level	17.3.10.4	20.3.21.4	22.3.19.4	Receive Sensitivity Automated Test

Table 3 - WiFi Receive Test Coverage

Link Layer Testing	Requirement 802.11-2012 a/g	Requirement 802.11-2012 n	Requirement P802.11ac	Automation
UDP Throughput	17.4.6.4	20.6	22.5	UDP Throughput Automated Test
TCP Goodput	17.4.6.4	20.6	22.5	Manual
Rate vs Range	17.4.6.4	20.6	22.5	Rate vs Range Automated Test
MU-MIMO Thruput			22.5	Future
Traffic Stress Test				Traffic Stress Automated Test

Table 4 - WiFi Link Layer Test Coverage

5. Test Automation Architecture

Automation of the test coverage is absolutely required given the complexity and coverage requirements for properly evaluating performance of an Access Point. Manual testing is too cost prohibitive in time and effort and requires a very high skill level.

Automation was realized by combining test equipment products from different companies with custom SW implementation based on industry standard freeware. The automation framework implemented can also be used by other teams within Shaw for any repetitive test tasks, if a suitable ATA interface is available.

Shaw has developed an automation framework that supports the following goals:

1. Open source automation SW
2. Interfaces to all Unit under Test variants
3. Repeatability
4. Reliability
5. Persistence of test data
6. Ease of use
7. Direct interpretation of results to pass/fail criteria
8. Configuration control of test sequences, test SW and test setup conditions

5.1. Test Setup

The test bench hardware setup that supports the Transmit/Receive/Link layer testing is shown in Figure 1 below:

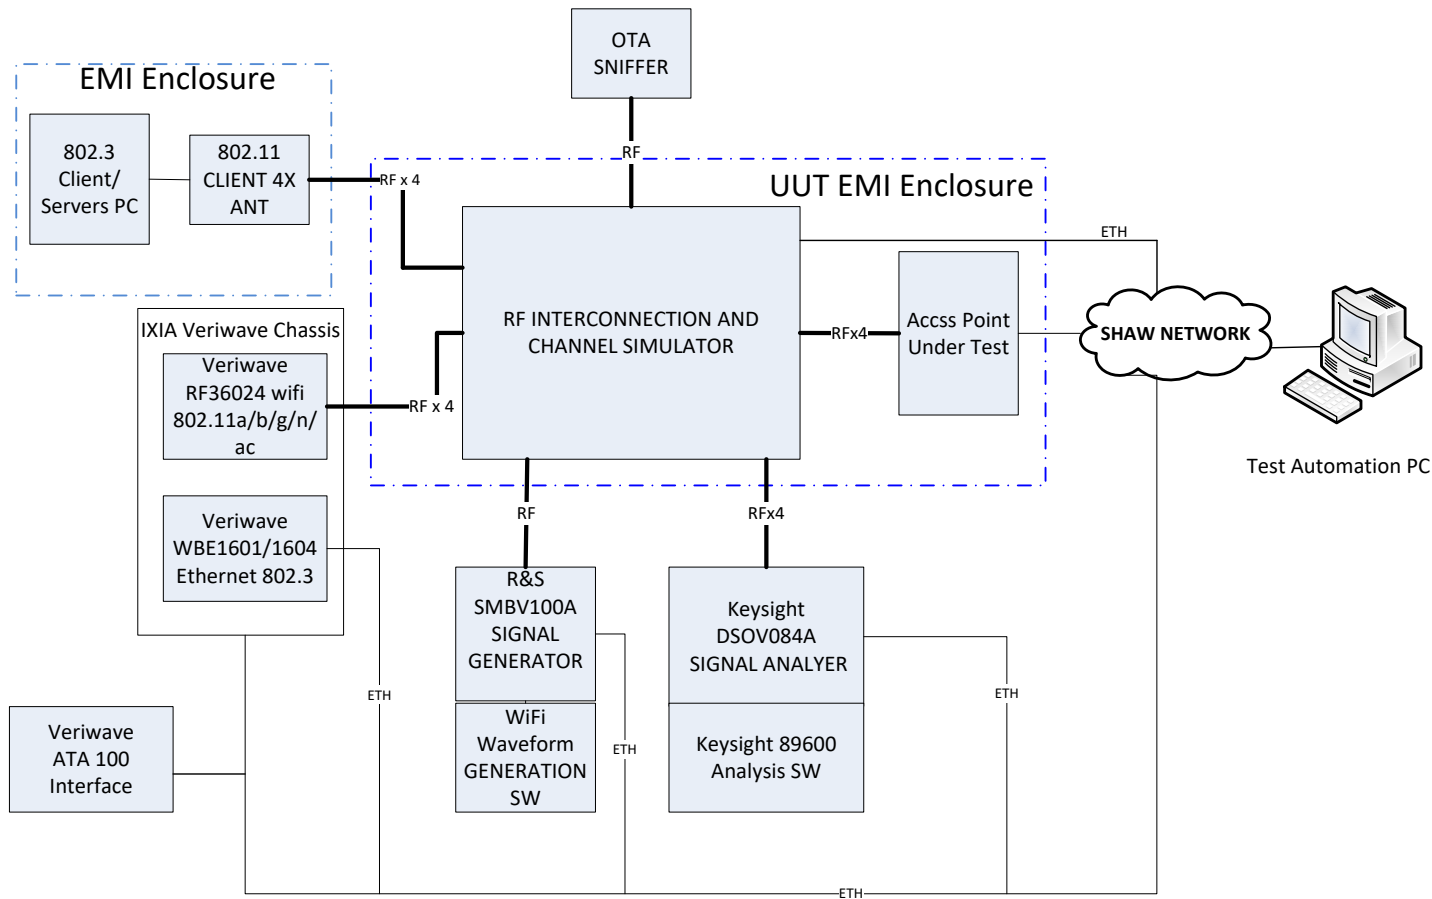


Figure 1 - WiFi Automation Hardware Setup

The block components description of the Automation HW setup is as follows:

5.1.1. 802.3 Client/Server PC

The PC hosts the traffic test tools such as J-perf to perform throughput and packet statistics. It also provides an interface to cloud management for accessing and configuring the AP under test. May also be used to host VOIP, generate video traffic etc.

5.1.2. 802.11 Clients 4x Antenna

The 802.11 clients provide the ability to test the APs via the 802.11 interface standard. The clients support the 802.11 b/g/n/ac to 160 MHz standards. The clients are either ASUS WiFi cards or Octoscope PAL 2.

5.1.3. IXIA IxVeriWave Chassis




Figure 2 - IXVeriwave Chassis

The IXIA IxVeriWave Chassis provides test client capability and test coverage for many of the 802.11 tests. The IxVeriWave RF36024 card supports 802.11a/b/g/n/ac standards for client simulation. The IxVeriWave Ethernet card WBE1601/04 provides the 802.3 client/server interface.

The IxVeriWave ATA 100 interface provides the remote command interface ATA commands via Telnet to the IXIA IxVeriWave chassis. The ATA commands allow full programming capability for configuring clients, generating data flows and running measurements and status queries.

5.1.4. OTA Sniffer

The OTA sniffer provides ability for Wireshark packet capture to analyze the traffic between client and Server.

5.1.5. RF Interconnection and Channel Simulator

The RF interconnection and channel simulator provides the physical RF interconnection of the AP UUT, client and external test equipment. It provides RF switch/coupling paths to support all RF test cases including RF transmit/receive, external interferers, and multiple APs and clients. The components are housed in an EMI chamber to minimize external interference.




Figure 3 - RF Interconnection Views

Channel simulation is realized by injection of “on channel” and “adjacent channel” noise from the external signal generator and is routed via passive splitters/attenuators into the uplink or downlink paths as required.

A 2nd version of the RF Interconnection supports the Rate vs Range automated test (see Figure 24). The channel simulation is implemented with a Butler Matrix device placed between the AP and UUT to ensure samples of each RF path are mixed onto all output ports between client and AP.

5.1.6. 802.11 Signal Generator

Figure 4 Signal Generator R&S SMBV100A



The signal generator used is a Rohde and Schwarz SMBV100A. It is used to transmit 802.11 waveforms with necessary characteristics to support RF test cases such as adjacent channel tests. The signal generator is also used for injecting Gaussian noise to control the C/N ratio of the WiFi channel. The Signal Generator is controlled via Ethernet SCPI command interface.

WiFi Waveform Generation SW

The WiFi Generation SW resides on the signal generator. It provides a tool to generate the waveforms giving access to key parameters within the waveform frame level to change MAC addresses, signaling parameters, duty cycle etc. The waveforms can then be loaded to the signal generator for transmission to the AP.

5.1.7. 802.11 Signal Analyzer



Figure 5 - Keysight Oscilloscope

The signal analyzer used is Keysight oscilloscope DSOV084A 4 channel model running 89600 Analysis SW. This combination provides RF waveform analysis for 802.11 signal physical characteristics such as power, EVM, and in-band and out-of-band channel emissions. The Signal Generator is controlled via Ethernet SCPI command interface.

5.1.8. Test Automation PC

The test automation PC is the host of the test automation SW. It interfaces with all test hardware components via different protocols and the AP under test. The Test automation PC also supports the I-Perf client/server application. The test result data is gathered by the Test Automation PC that interfaces with remote SQL database to store test results.

5.2. RF Interconnection and Channel Simulator Block Diagram

The RF Interconnection is implemented with conducted RF connection cabling so the AP under test is connected-with a conducted RF cable at the antenna input ports. No radiated testing is supported in this configuration. The RF interconnection provides the RF paths for AP to client antenna and RF paths for the signal generator and signal analyzer.

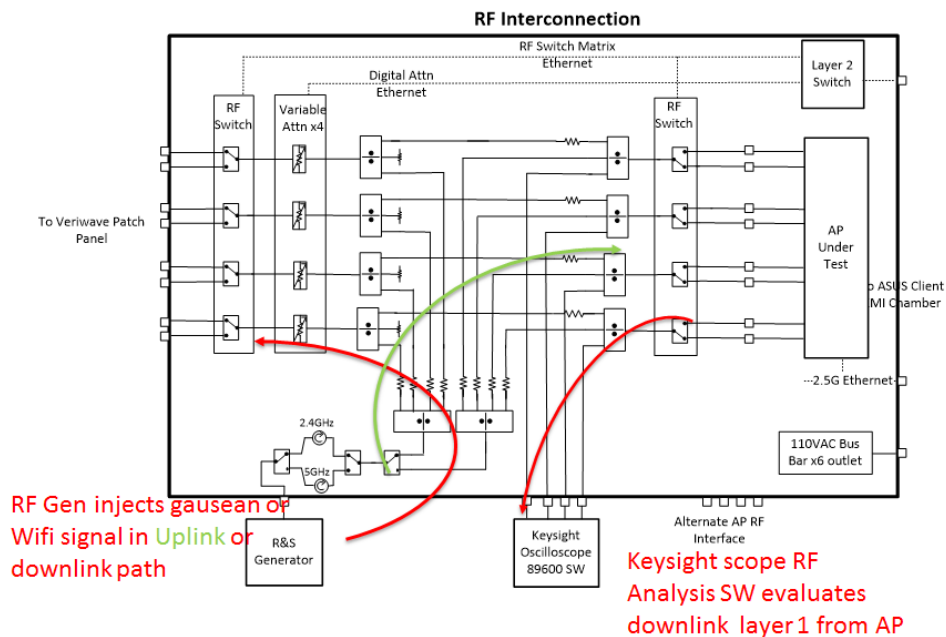


Figure 6 - RF Interconnection Block Diagram

All components are housed within the EMI chamber to minimize interference. Connected RF connections are typically used for testing.

The RF switches and attenuators are controlled externally via Ethernet SCPI command sets.

The RF Variable attenuators provide a dynamic range control of 0- 90dB of in line attenuation. This is used to set the path loss between client and AP RF ports.

The Signal generator is used to inject RF noise on the downlink path to adjust the C/N ratio of the link as shown in RF path in red. The signal generator can also inject noise into the uplink path (shown in green) of the AP under test for receive input co-channel and adjacent channel interference tests.

The signal analyzer receives samples of the RF antenna ports (up to 4) from the AP. The signal analyzer is used to demodulate up to a 4-spatial stream 11ac signal with 160 MHz bandwidth. The signal analyzer is also used to measure the transmit spectral mask, transmit occupied bandwidth and adjacent channel powers.

5.2.1. Example Test Setup for WLC AP

In this configuration, the AP is the Device Under Test (DUT) and creates a CAPWAP tunnel with the WLC via the Shaw Intranet, MPLS Network for the Fiber based Network or DOCSIS 3.1 based Network,

if connected via the Cable Modem. The WLC creates a PMIP/GRE Tunnel with the Local Management Anchor (LMA). The Core Switch interconnects the Access Network with the Core Network components.

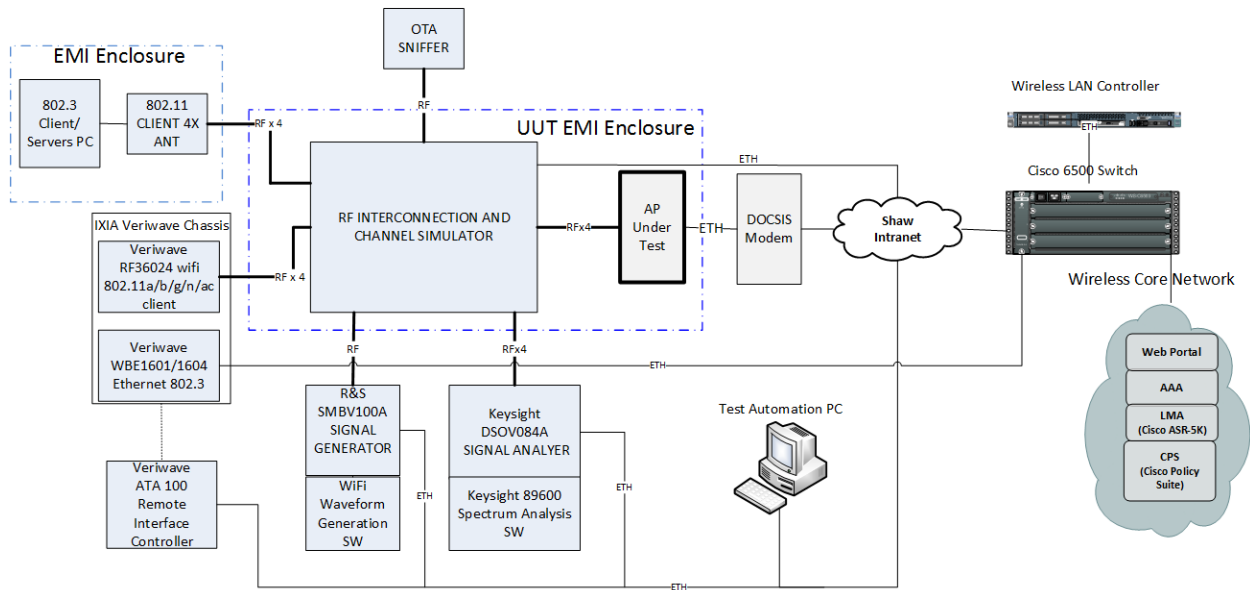


Figure 7 - Automation Example Test Setup for WLC AP

5.2.2. Example Setup for Cloud Managed AP

The cloud managed AP configuration does not require a WLC. All AP management including configuration control is done through remote cloud based applications reducing the CPE requirements. The DOCSIS modem shown provides internet connectivity to the Shaw network.

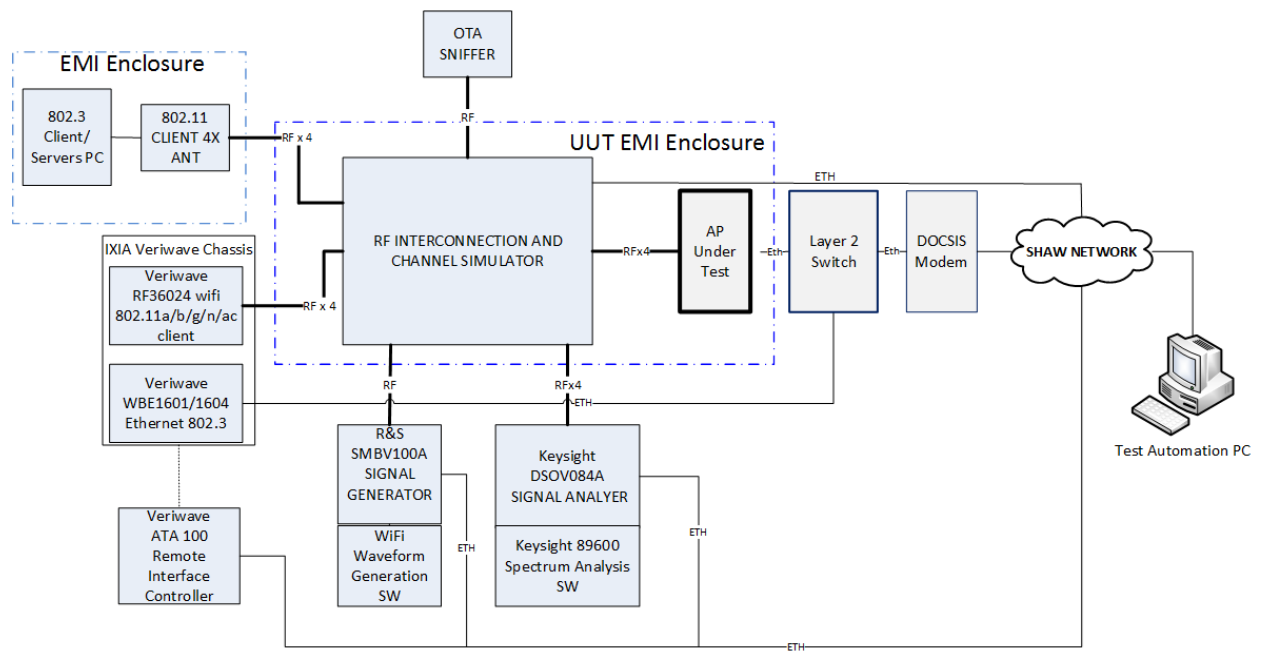


Figure 8 - Automation Example Test Setup for Cloud Based Management AP

5.3. Automation SW Architecture

The Test automation SW architecture is based on Python scripts and RobotFramework with Ride.py GUI interface. Each measurement engine (i.e. Receiver Sensitivity) is written in Python and utilizes common Python subroutines for remote interfaces to equipment, and data gathering. Common python scripts are used to write test setup and test results to the SQL database.

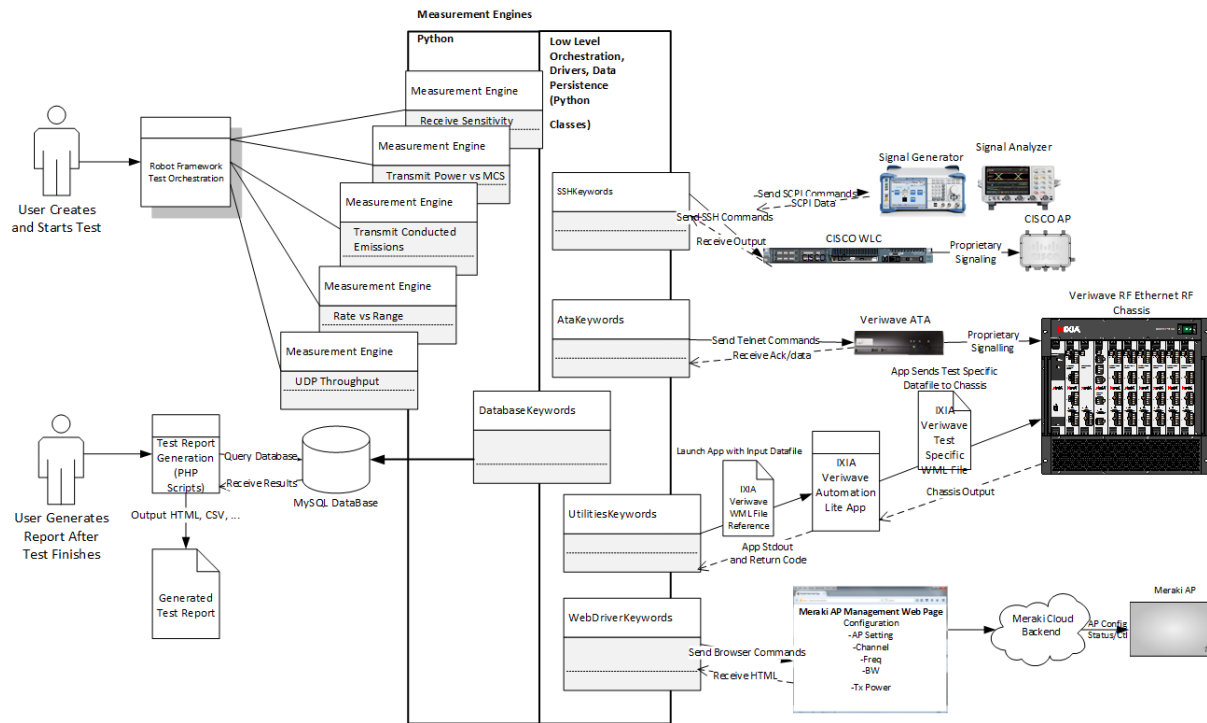


Figure 9 - Automation SW Block Diagram

The User creates the test and tailors it to the AP under test through the RobotFramework Ride.py interface. The user has control of the channel selection, power levels, and MCS rates under test, etc. as inputs to the measurement engine. The measurement engine is common for all APs and is only tailored for the test coverage as selected by the user.

The RobotFramework Ride.py acts as the test sequencer and runs the tests in order as selected by the User. The RobotFramework supports data and error logging of the test sequence results. Many tests can be selected for running sequentially. If a test fails for some reason, RobotFramework continues to the next test until all tests have been completed.

The IxVeriWave chassis is key to running setting up clients and test flows for all tests. In addition, the test Power vs. MCS employs the IxVeriWave WaveAnalyze Test is a SW license that runs on the IxVeriWave RF36014 RF card. The WaveAnalyze SW measures the layer 1 RF performance of the AP downlink under test including RF power, spectral performance, and modulation quality. The test results

from WaveAnalyze are parsed from csv files by python subroutines and the test results sent to the SQL database.

The UDP throughput test uses the IxVeriWave AutoLite IXIA benchmark test SW which provides an automated method to configure and run IxVeriWave UDP throughput benchmark tests. We use this test SW feature to incorporate benchmark tests in the automation framework.

5.4. RobotFramework

Example of the RobotFramework Ride.py GUI is shown below. On the left tab is example of the sequence of tests available to the user. On the right tab is the robot library definitions for the python measurements functions.

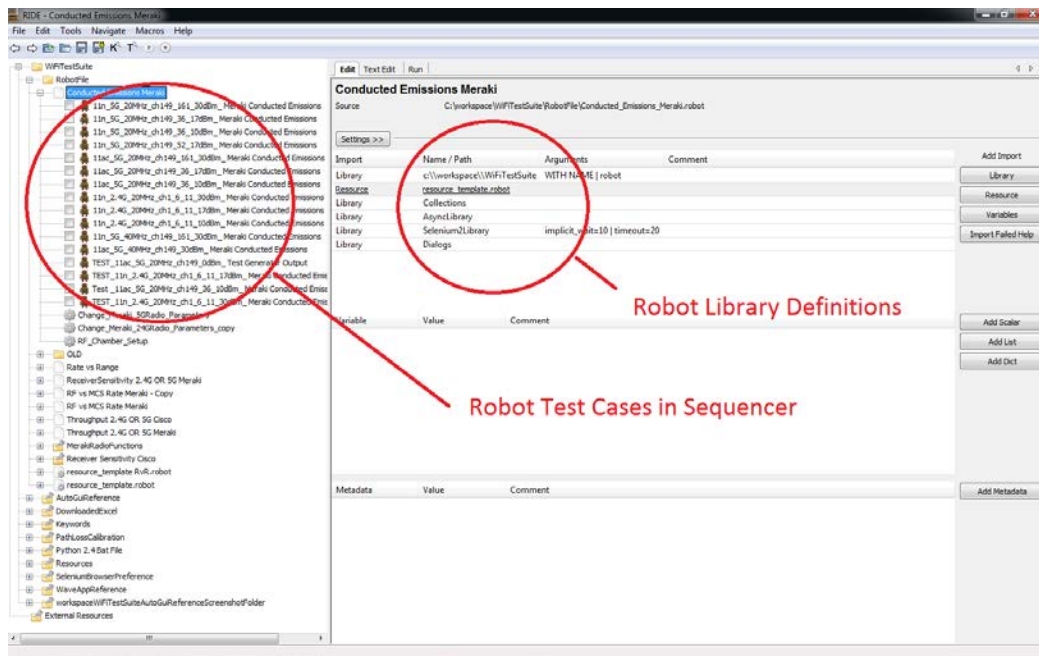


Figure 10 - RobotFramework Sequence

The individual test cases are configurable for the test coverage and input parameters for a AP thorough the GUI interface as shown in figure 11 below. The single test case entitled “11ac_20Mhz_ch149_36_17dBm_Meraki Conducted Emissions” has AP and test variable inputs that are set at the RobotFramework GUI.

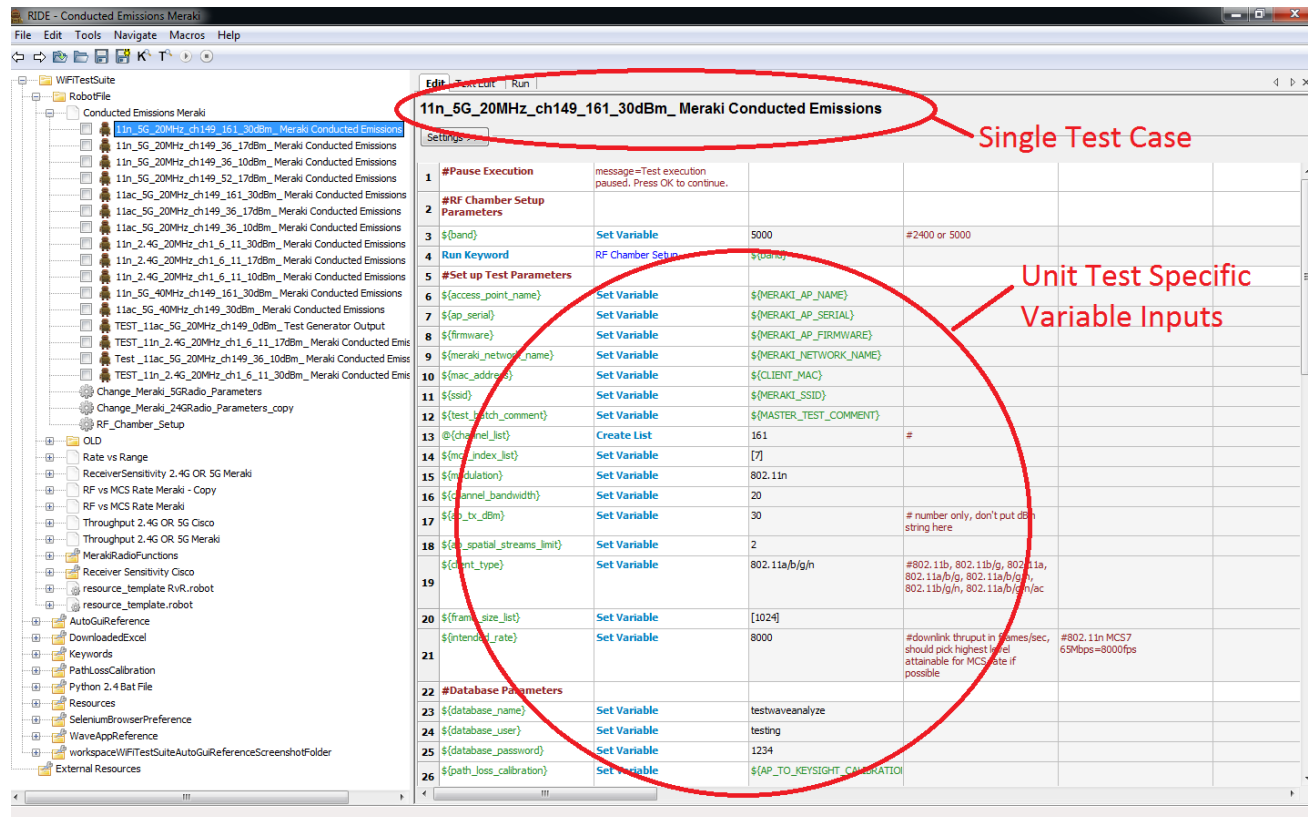


Figure 11 - RobotFramework Measurement Configuration

The RobotFramework library link to the python keywords project allows the measurement engines to be run at the RobotFramework level. RobotFramework sets the variable inputs to the measurement engines for the test.

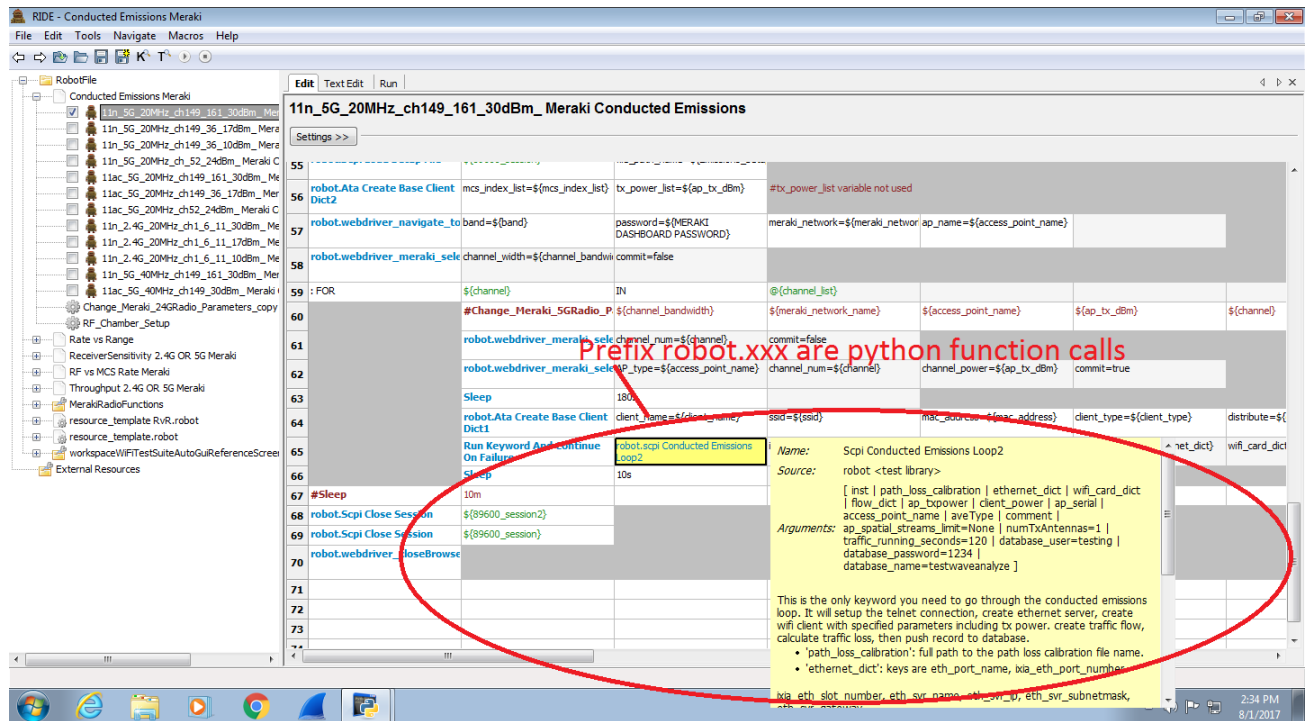


Figure 12 - RobotFramework Measurement Engine Function Call

5.5. Python Keywords Definition

The python keywords implementation is shown in Figure 13 below as viewed using Eclipse SW tool. The keywords defined as external functions can be called by RobotFramework.

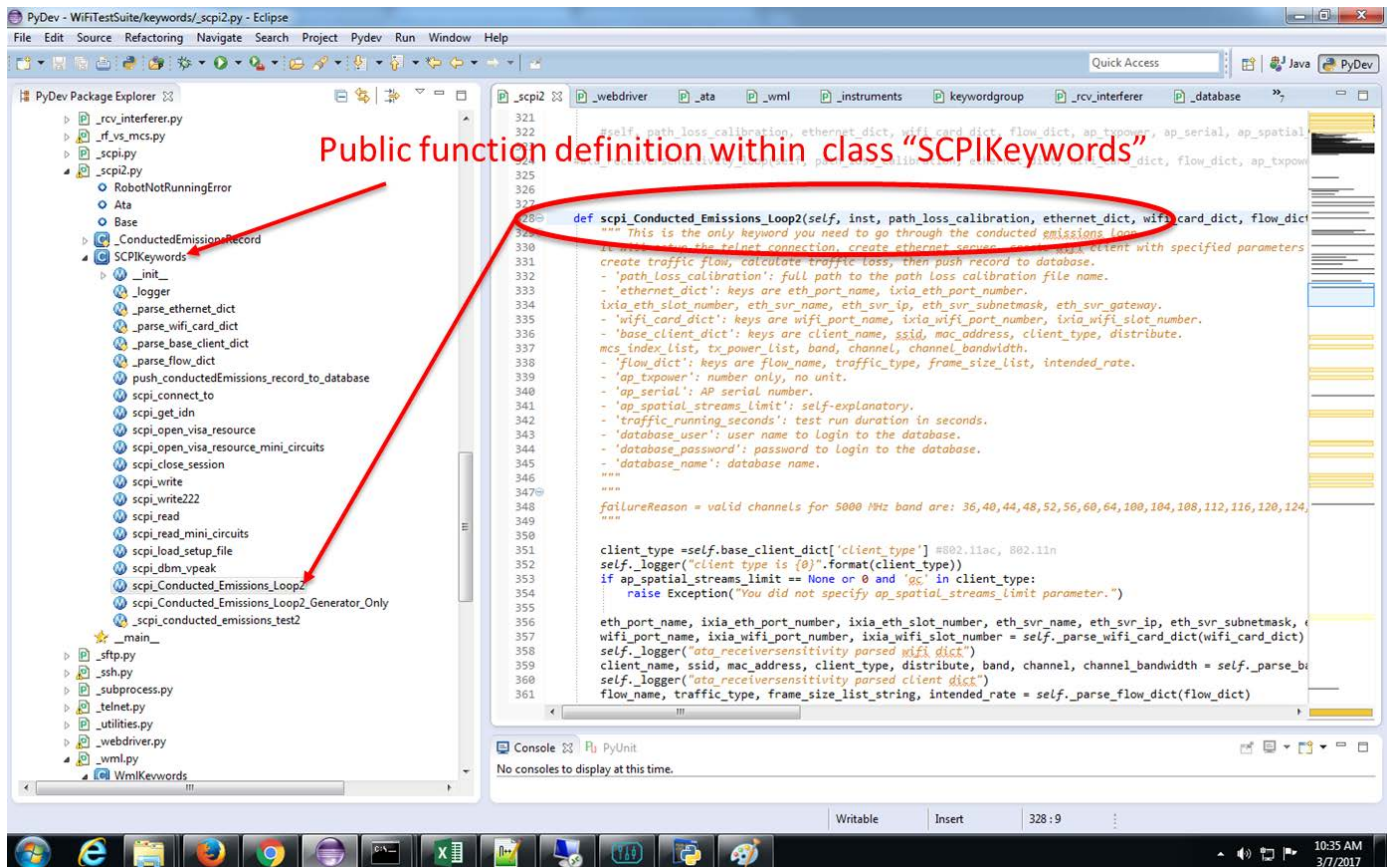


Figure 13 - Python Keywords View

5.6. Test Data Record

The Automation python measurement engines collects the test result data from each test which is then stored on the SQL database.

An example of the SQL database record is shown in figure 14 below for the UDP throughput automated case. The SQL database has two separate tables, one for results, and 2nd for test setup. Both tables are merged to a form a complete test data record.

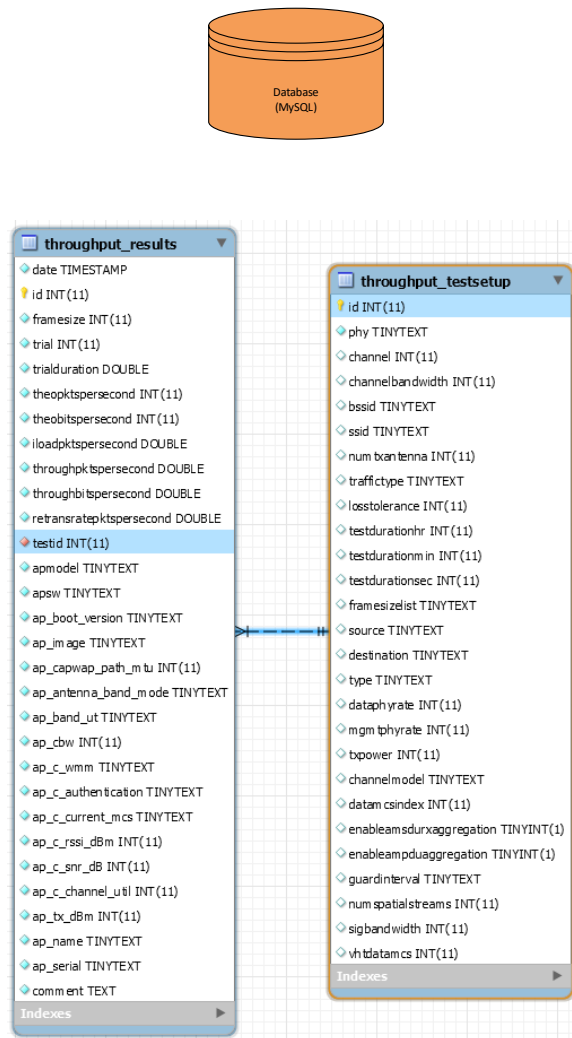


Figure 14 - Test Data Record Example

The SQL database is phpMyAdmin freeware and provides a GUI interface to view the data as shown in figure 15 below. Individual SQL queries can be run on the data, or the entire table exported to a CSV file for post processing.

[illegible]

5.7. Test Data Results and Analysis

The test data stored in the SQL database is quite extensive for each test case. We have written PHP scripts to perform the post processing data analysis steps to present a summary of the test results of interest.

An example is shown in Figure 16 below for the PHP script used to find the best Receive Sensitivity result for each test case (MCS rate, spatial streams).

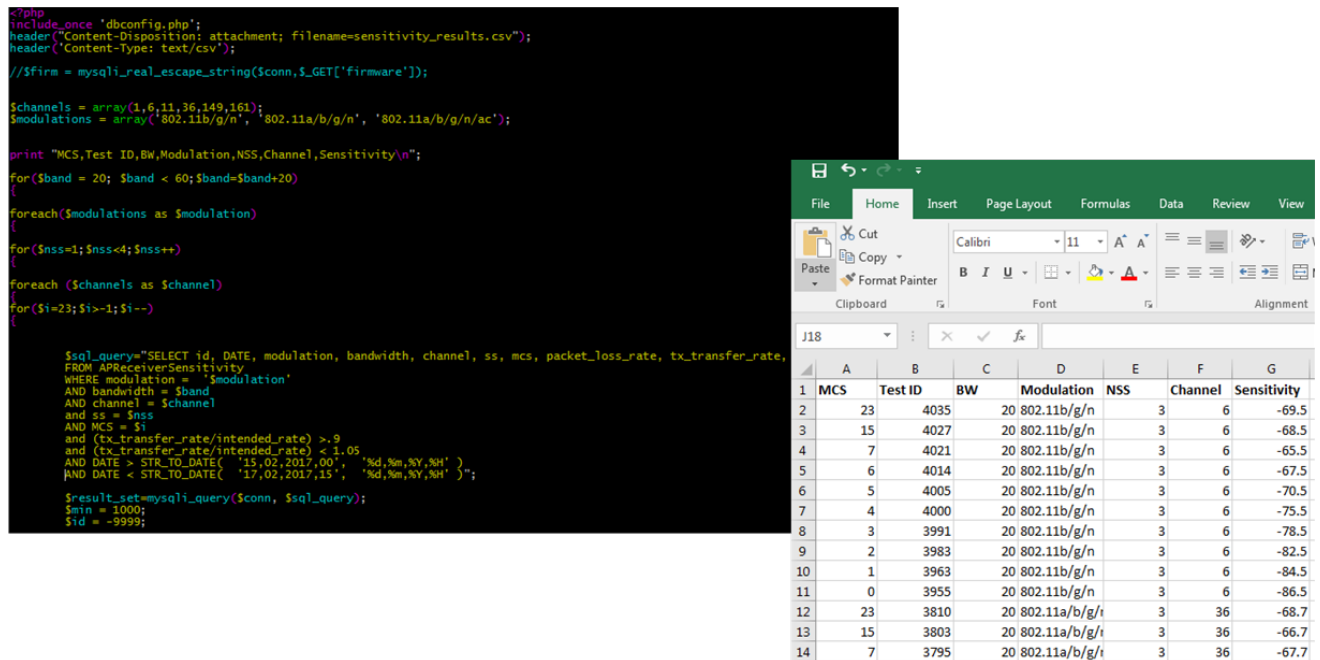


Figure 16 - PHP Script and Result CSV File Example

6. Test Measurements

The test measurement coverage based on automated tests is as follows:

1. RF Characterization Vs Order Power vs MCS Rate
2. Conducted Emissions
3. Receiver Sensitivity
4. UDP Throughput
5. Rate vs Range
6. Traffic Stress Test

The following sections will provide a more detailed overview of how each measurement has been implemented and discussion of typical results attained on the AP under test.

6.1. RF Characterization vs Order Power vs MCS Rate

The purpose of the RF Characterization Vs Ordered Power vs MCS Rate is to measure all RF characteristics for all MCS rates for all modulations over the operational range of output power settings.

The coverage of this test has several dimensions and relies on the IxVeriWave WaveAnalyze SW. WaveAnalyze performs vector signal analysis used to test and qualify 802.11 WiFi transmitters.

WaveAnalyze delivers detailed analysis for every frame in real-time, or in recorded form for future assessment. (The WaveAnalyze SW GUI is shown in figure.) The following measurements are made continuously with output data every five seconds to a CSV file on a per stream/port basis:

- EVM Data RMS, EVM Signal RMS
- Per Subcarrier EVM RMS
- Preamble Frequency Error
- Transmit Symbol Clock Frequency Tolerance
- Transmit Center Frequency Tolerance
- Transmit Average Power
- Transmit Peak Power
- Transmit Peak Power excursion
- Transmit Power Ramp
- Transmit RF Carrier Suppression
- Transmit Constellation per spatial stream
- Transmit Spectral Flatness
- Transmit Spectrum Mask

The WaveAnalyze measurement SW can be run manually via a GUI or called directly from the automation SW. The WaveAnalyze generates a CSV file of results that are parsed and recorded in the SQL database. The example GUI results show the results for power output and EVM measurement.

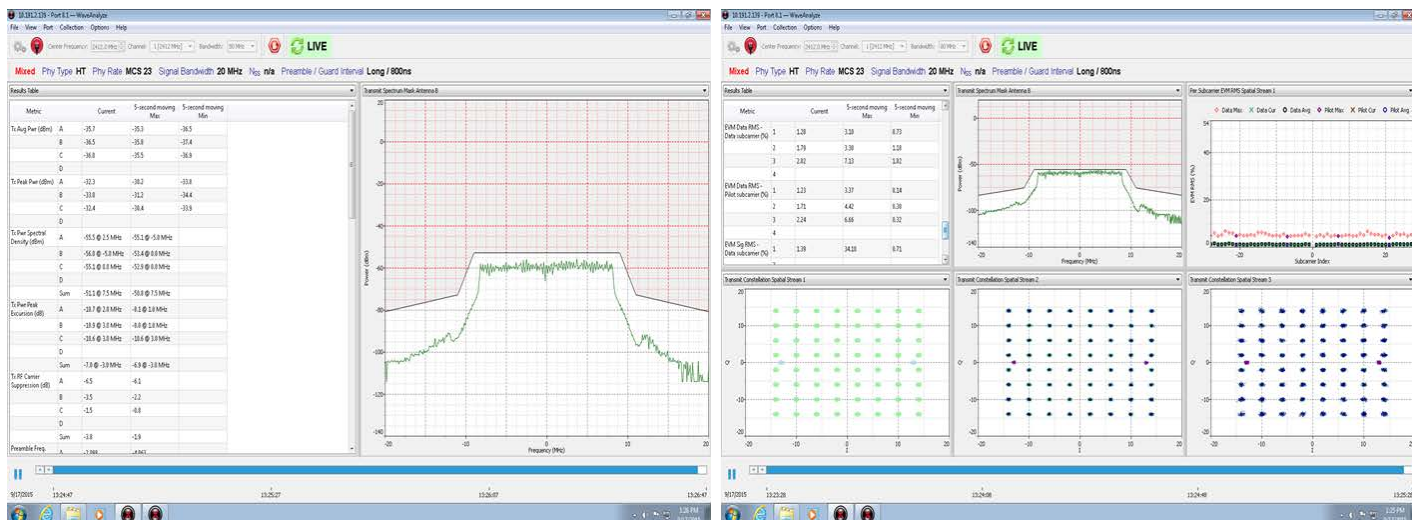


Figure 17 - WaveAnalyze RF Measurement Example

The configurable inputs via RobotFramework GUI to setup the automated test are:

- AP band, channel and modulation type (i.e. 5GHz channel 153, 802.11n)
- Bandwidth 20/40Mhz
- MCS rate of interest (i.e. MCS 7, 15, 23)

- Frame size, data rate (i.e. 1024bytes, 1000fps)
- AP power steps to be measured. (i.e. steps from 12 to 30dBm in 3dB increments)

The test automation then performs the following measurement steps:

- a. Sets AP to the desired channel power level
- b. Sets the IxVeriWave client to advertise the selected band, channel, modulation rate
- c. Connects the IxVeriWave client to the AP
- d. Establish a downlink flow at the desired frame and data rate
- e. Start IxVeriWave WaveAnalyze Analysis SW
- f. Read CSV file to extract measurements results and confirm test results captured for the desired MCS rate
- g. Records results of measurements in SQL database
- h. Repeats measurement at the AP desired power setting
- i. Test duration is approximately 2 minutes for each measurement after initial connection/setup (per MCS under test)

If the target MCS rate is not realized, the test automation will modify the C/N ratio of the test flow by injecting Gaussian noise from external generator in 3dB increments from an initial C/N point. As the C/N is reduced, the AP algorithms will select lower MCS rates to compensate. The test program continues to modify the C/N ratio until the target MCS rate is selected by the AP under test.

We used guidelines (Ref 3) from Andrew Von Nagy shown in Table 5 below as a starting point to set the link SNR when targeting a specific MCS rate.

MCS Value Achieved by Clients at Various Signal to Noise Ratio Levels (SNR)

Protocol	Channel	1	2	3	4	5	6	7	8	9	10	
802.11b	20MHz	None	None	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	Modulation Key None = Grey BPSK = Red QPSK = Orange 16-QAM = Yellow 64-QAM = Blue 256-QAM = Green
802.11a/g	20MHz	None	MCS 0	MCS 0	MCS 1	MCS 2	MCS 2	MCS 2	MCS 2	MCS 3	MCS 3	
802.11n	20MHz	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	
802.11n	40MHz	None	None	None	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	
802.11ac	20MHz	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	
802.11ac	40MHz	None	None	None	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	
802.11ac	80MHz	None	None	None	None	None	None	None	MCS 0	MCS 0	MCS 0	
802.11ac	160MHz	None	None	None	None	None	None	None	None	None	None	
SNR in dB		11	12	13	14	15	16	17	18	19	20	
802.11b	20MHz	MCS 2	MCS 2	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11a/g	20MHz	MCS 4	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 5	MCS 6	MCS 6	MCS 7	
802.11n	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	
802.11n	40MHz	MCS 1	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	
802.11ac	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	
802.11ac	40MHz	MCS 3	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	
802.11ac	80MHz	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	
802.11ac	160MHz	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	MCS 2	MCS 3	
SNR in dB		21	22	23	24	25	26	27	28	29	30	
802.11b	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	20MHz	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	40MHz	MCS 5	MCS 5	MCS 6	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	
802.11ac	20MHz	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	MCS 7	MCS 8	MCS 8	
802.11ac	40MHz	MCS 5	MCS 5	MCS 6	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	
802.11ac	80MHz	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 5	MCS 6	MCS 6	MCS 6	MCS 6	
802.11ac	160MHz	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	MCS 6	
SNR in dB		31	32	33	34	35	36	37	38	39	40	
802.11b	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	40MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11ac	20MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	40MHz	MCS 7	MCS 8	MCS 8	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	80MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 8	MCS 8	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	160MHz	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	MCS 7	MCS 8	MCS 8	MCS 9	
SNR in dB		41	42	43	44	45	46	47	48	49	50	
802.11b	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	40MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11ac	20MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	40MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	80MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	160MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	

Table 5 - MCS vs SNR Estimate

As the MCS rate increases, so does the EVM requirements for the modulation mode used. The EVM is critically important and becoming more difficult to meet for higher MCS rates. This will be even more so with the introduction of 802.11ax. Example EVM results for different candidate APs is shown in Figure 18 below, plotted against 802.11ac MCS 9 EVM requirement of 2.5% for different AP power settings. As shown, AP- C and AP – D suffer from high EVM exceeding the specification at the higher power settings which will result in poorer downlink performance.




Figure 18 - AP EVM Comparison 802.11ac MCS 9

Another important data point is the beacon power vs. MCS data rate power. Typically, the beacon power is the highest power signal from the AP, as this means the beacon is seen at the greatest distance from the AP. It is important to know the relative data MCS power to the beacon power for site survey and deployment purposes. Figure 19 “RF Power vs. MCS Rate 5GHz 11n Product C” shows results of comparing RF power levels per MCS rate. As can be seen there is a power difference between beacon and MCS frame of up to 4dB. “RF Power vs MCS Rate 11ac Product "C" also shows a difference of over 5dB between beacon and MCS frames. This difference of high MCS rate vs beacon power should be considered when determining AP spacing for optimum coverage.

As stated above, the MCS measurements can only be made when injecting noise to adjust the C/N ratio. The relative C/N ratio required to achieve an 802.11ac MCS rate on the downlink is plotted in Figure 20. We do not use this information for evaluation, but it is interesting that for this product MCS 2 could not be invoked when adjusting the C/N ratio.




Figure 19 - RF Power vs MCS Examples




Figure 20 - Average C/N applied to Invoke 11ac MCS Rate

6.2. Conducted Emissions

The purpose of the Conducted Emissions Automated Test is to evaluate the Transmit RF spectrum performance of the AP under test downlink. This test uses the Keysight Oscilloscope with 89600 RF Analysis SW to measure the spectrum performance for the following parameters:

- Occupied Bandwidth
- Adjacent Channel Power
- Spectral Mask

The configurable inputs to the automated test are summarized as follows:

- AP band, channel, and modulation type (i.e. 5GHz channel 153, 802.11n)
- MCS rate of interest (i.e. MCS 7, 15, 23)
- Bandwidth 20/40Mhz
- Frame size, data rate (i.e. 1024bytes, 1000fps)
- AP power steps to be measured

The test automation then performs the following measurement steps:

- a. Sets AP to the desired channel power level via HTML website automation
- b. Sets the IxVeriWave client to advertise the selected band, channel, modulation rate
- c. Connects the IxVeriWave client to the AP
- d. Establish a downlink flow at the desired frame and data rate
- e. Configure the Keysight Analyzer to perform the measurement
- f. Reads back the measurement results from the Keysight analyzer and records results into SQL database
- g. Repeats measurement for next configuration
- h. Test duration is approximately 3 minutes for each measurement after initial connection/setup (per MCS under test)

Part of the challenge with this test is avoiding averaging errors of the frames. IxVeriWave does try to control the periodicity of the downlink frames. The 89600 SW will provide average of the frame spectrum, and not average in any null times. And the test is set at the highest frame rate the downlink can support to maximize channel utilization. We choose to use “peak hold” averaging to evaluate the maximum spectrum density.




Figure 21 - Conducted Emissions Test Result Example

6.3. Receiver Sensitivity

The purpose of the Receiver Sensitivity Automated Test is to determine the minimum sensitivity based on 802.11 specification for conducted sensitivity frame error rate of 10%. This automated test case also tests sensitivity of receiver in adjacent channel and co-channel interference. This test uses the IxVeriWave Client to generate signals at the desired MCS rate for uplink to the AP under test.

The configurable inputs to the automated test are summarized as follows:

- AP band, channel, and modulation type (i.e. 5GHz channel 153, 802.11n)
- MCS rate of interest (i.e. MCS 0-7, 15, 23)
- Bandwidth 20/40Mhz
- Frame size, data rate (i.e. 1024bytes, 1000fps)
- AP input receiver sensitivity range that covers all MCS rates under test.

The test automation then performs the following measurement steps:

- Sets AP to the desired channel power level
- Sets the IxVeriWave client to advertise the selected band, channel, modulation rate
- Connects the IxVeriWave client to the AP
- Establish an uplink flow at the desired frame and data rate
- Perform search algorithm to determine the nominal receiver sensitivity that still supports the required frame error rate in minimum number of steps by adjusting the IxVeriWave Client output power

- f. Records results into SQL database
- g. Repeats measurement for next configuration
- h. Test duration is approximately 7 minutes for each measurement after initial connection/setup (per MCS under test) as up to 7 trials are run to determine the minimum sensitivity point.

One of the challenges for this test is finding the AP uplink receiver sensitivity in as few of steps as possible. We use a simple uniform binary search algorithm that minimizes the number of power levels settings for the data flow to find the desired receiver sensitivity that supports 10% or less frame error rate. Example results for receive sensitivity is shown in figure 22 below.




Figure 22 - Receiver Sensitivity Example

In case of the adjacent channel measurement, the external generator is used to simulate WiFi signal with 50% duty cycle to place on adjacent or co-channel location. The interferer signal is stepped up in power until the receiver sensitivity is degraded to specification limit.

6.4. UDP Throughput

The purpose of this test is to measure the UDP throughput for both Uplink and Downlink and compare results to theoretical rates. This test uses IxVeriWave Benchmark Throughput test and IxVeriWave Wave Automate SW to programmatically configure and run the benchmark test through simple TCL scripts.

The configurable inputs to the automated test are summarized as follows:

- AP band, channel, and modulation type (i.e. 5GHz channel 153, 802.11n)
- MCS rate of interest (i.e. MCS 7, 15, 23)
- Bandwidth 20/40Mhz
- Frame size, data rate (i.e. 1024bytes, 1000fps)

- AP power steps to be measured.

The test automation then performs the following measurement steps:

- Modify the master configuration TCL file for the IxVeriWave Benchmark test.
- Sets AP to the desired channel power level.
- Invokes the TCL file to run IxVeriWave Benchmark Test via Wave Automate SW.
- IxVeriWave Benchmark test runs and generates results CSV file.
- Automation reads CSV file and records results in SQL database.
- Repeats measurement for next configuration.
- Test duration is approximately 5-6 minutes for each measurement after initial connection/setup (per MCS/frame rate under test)

The summary figure 22 “AP UDP Throughput Result Example” is a subset of the information provided by the IXVeriwave Benchmark test report. In this table, uplink/downlink throughput is plotted against the theoretical throughput attainable as calculated by IXVeriwave Benchmark test based on MCS rate, AMPDU/AMSDU settings, guard interval etc. We measure the UDP throughput typically across several frame sizes and modulation rates. We like to see performance above 75% of theoretical attainable given a frame loss tolerance of <10%.




Figure 23 - UDP Throughput Result Example

6.5. Rate vs Range

The purpose of Rate vs Range test is to measure the AP downlink performance to a test client as the relative attenuation representing range is varied simulating a near client to far client.

This test is not performed with IXIA IxVeriWave products. The test is realized using an example client such as ASUS Model PCE-AC68 or Octoscope PAL2 802.11ac client. The data flow is created using I-Perf client/server and the nominal TCP throughput is measured as a function of range.

The test is fully automated within the automation framework, but the RF interconnection is modified to include a butler matrix as shown in figure 24. The butler matrix is necessary to mix samples of all radio antenna outputs from the AP to the client to support spatial stream diversity (See ref 2).

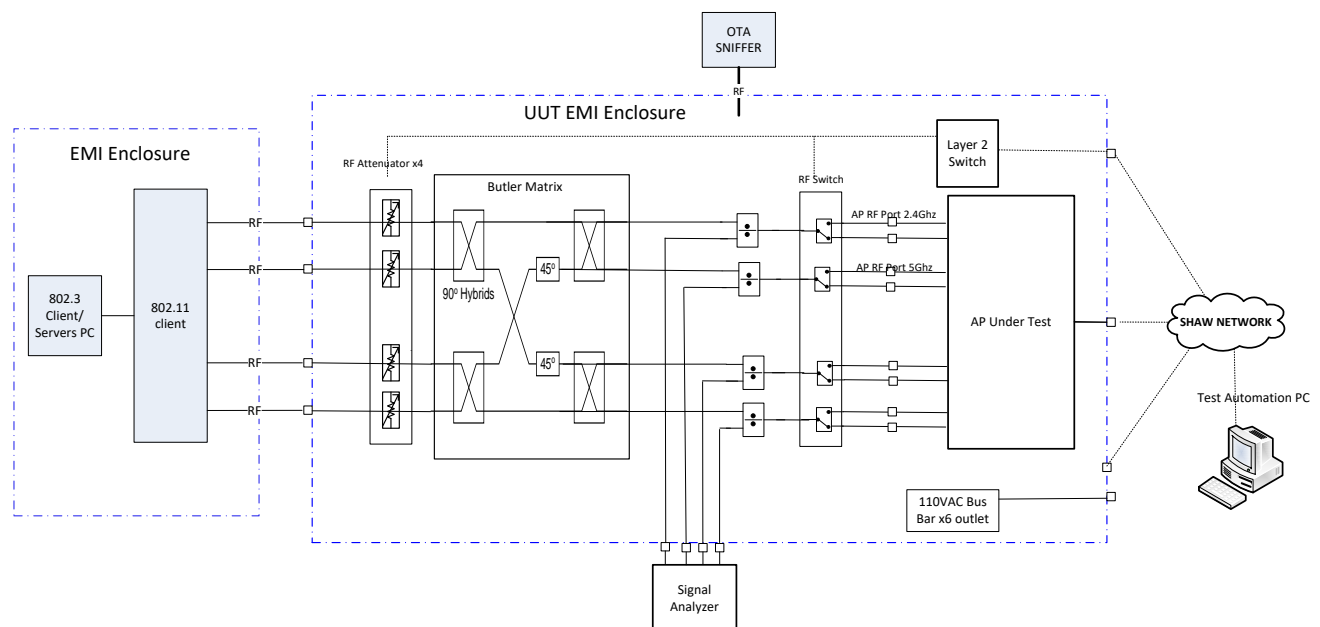


Figure 24 - Rate vs Range Hardware Test Setup

The configurable inputs to the automated test are summarized as follows:

- AP band, channel, and modulation type (i.e. 5GHz channel 153, 802.11n)
- Bandwidth 20/40Mhz
- Frame size, data rate (i.e. 1024bytes, 1000fps)
- AP power steps to be measured

The test automation then performs the following measurement steps:

- a. Sets AP to the desired channel power level
- b. Sets test client (i.e. Octoscope PAL 2) to desired configuration
- c. Initiates I-perf client server

- d. Gathers client server statistics for the attenuation test step
- e. Repeats test for the attenuation steps desired
- f. Repeats measurement for the configuration
- g. Test duration is approximately 15 minutes for each Access Point per channel under test.

The TCP throughput results are written to the SQL database and then plotted as per below.

Rate vs Range test is best performed for comparative purposes between AP or on the same AP for regression test purposes.

In this example in Figure 25 below, product “A” Firmware revision 1.0 is compared to firmware revision 2.0.

The 2nd firmware release was intended to improve throughput at 20 MHz/40MHz. The vendor was successful in improving the 40 MHz case, but new firmware in fact reduced the performance at 20 MHz as shown.




Figure 25 - Rate vs Range Result Example for Different Firmware

Typically Rate vs. Range test results are used as a comparative tool for assessing different models/manufactures of APs. We also wanted to compare the results with theoretical rates attainable for a given power level and link SNR as defined by the relative attenuation setting. An example of comparative testing for different APs is shown in figure 25. In this example, multiple manufacturer product results are compared. The theoretical TCP throughput performance attainable is also estimated and plotted in this example.

The theoretical rate vs. range TCP throughput performance is estimated through the following steps:

1. Assume nominal receiver sensitivity noise floor of -93dBm (allows 9dB receiver NF implementation in 20 MHz vs 802.11 allowance of 15dB.) = RcvSens
2. Measure nominal output power of AP = Pout.
3. Measure total attenuation/pathloss between AP and Client = Attn_dB
4. Determine power level at client receive input= Pout-Attn_dB= Pin_dB
5. SNR = Pin-RcvSens
6. Add 8dB estimate to SNR account for FEC coding gain, receive diversity, beamforming that will improve SNR. SNR_Corrected = SNR+8dB.
7. Compare SNR_Corrected to MCS vs SNR chart (ref 3) to determine the MCS rate that can be supported.
8. For the MCS rate supported, estimate the TCP rate attainable based on PHYrate, UDP throughput at nominal AMPDU setting, and typical TCP rate vs UDP rate. (See table 6). For 20Mhz BW, the estimate is TCP rate is 80% of PHYrate, and for 40Mhz BW, the estimate is TCP rate is 75% of PHYrate.

Mode	Maximum PHY Rate(Mbps)	A-MPDU size	Maximum Throughput(UDP Payload=1500, A-MPDU spacing=0)	% UDP vs PHY (see Note 2)	% TCP vs PHY where tCP = UDP *88% (Note 1)
11n (20 MHz)	72.2	8192	56.3	0.78	0.69
	72.2	16384	62	0.86	0.76
	72.2	32768	65.5	0.91	0.80
	72.2	65536	67.3	0.93	0.82
11n (40 MHz)	150	8192	97.1	0.65	0.57
	150	16384	116.1	0.77	0.68
	150	32768	128.3	0.86	0.75
	150	65536	136	0.91	0.80
11ac (80 MHz)	433	8192	169.5	0.39	0.34
	433	16384	241	0.56	0.49
	433	32768	305.3	0.71	0.62
	433	65536	352.9	0.82	0.72
Note 1:	TCP throughput estimated at 88% of UDP throughput from IPERF test comparison				
Note 2:	UDP vs PHY Reference : http://80211notes.blogspot.ca/2014/03/phy-rate-and-udp-throughput.html				

Table 6 - Estimation of TCP Throughput vs AMPDU

The resulting theoretical Rate vs. Range estimate is plotted on figure 26 below. Also on figure 26 are test results for two AP “Product A” and “Product B”. “Product B” is plotted twice to show performance improvement provided by the vendor updating the firmware to Rev 2.0 based on Shaw test results feedback.




Figure 26 - Rate vs Range Candidate AP Comparison

6.6. Traffic Stress Test

The purpose of the Traffic Stress Test is to simulate many clients connecting to the AP over a long period of time. This simulates a real network case where an Access Point is servicing a Mall or a Train Station.

The example explained here is a test performed in the Pre-production environment the Cisco SP WiFi Network. The hardware Topology of this Network is represented in the Network Diagram shown in the figure 7 “Automation Example for WLC AP Test”.

The generation and control of multiple clients is possible using IxVeriwave chassis and ATA SW interface. The overall test sequencing is performed directly in python and will be incorporated into the RobotFramework architecture in the future.

The python program keeps a list of client MAC addresses that are connected/disconnected with nominal traffic in a controlled fashion. The rate of connection, duration of connection and packet size along with rate-of-transmission of the packets is randomized while keeping the overall aggregate throughput at a nominal rate. The detailed algorithm is shown in Figure 29.

The traffic stress test can be run continuously for a long period to flush out longer term issues such as memory leaks that cause the AP to stop functioning as expected.

Examples of the results are shown in Figure 27/28. Figure 27 is a plot of the overall throughput maintained through the AP as clients are randomly connected, run data flow and dis-associated. Figure 28 is the total client associated/authorized clients over time.




Figure 27 - Soak Test Throughput




Figure 28 - Client Associations vs Time

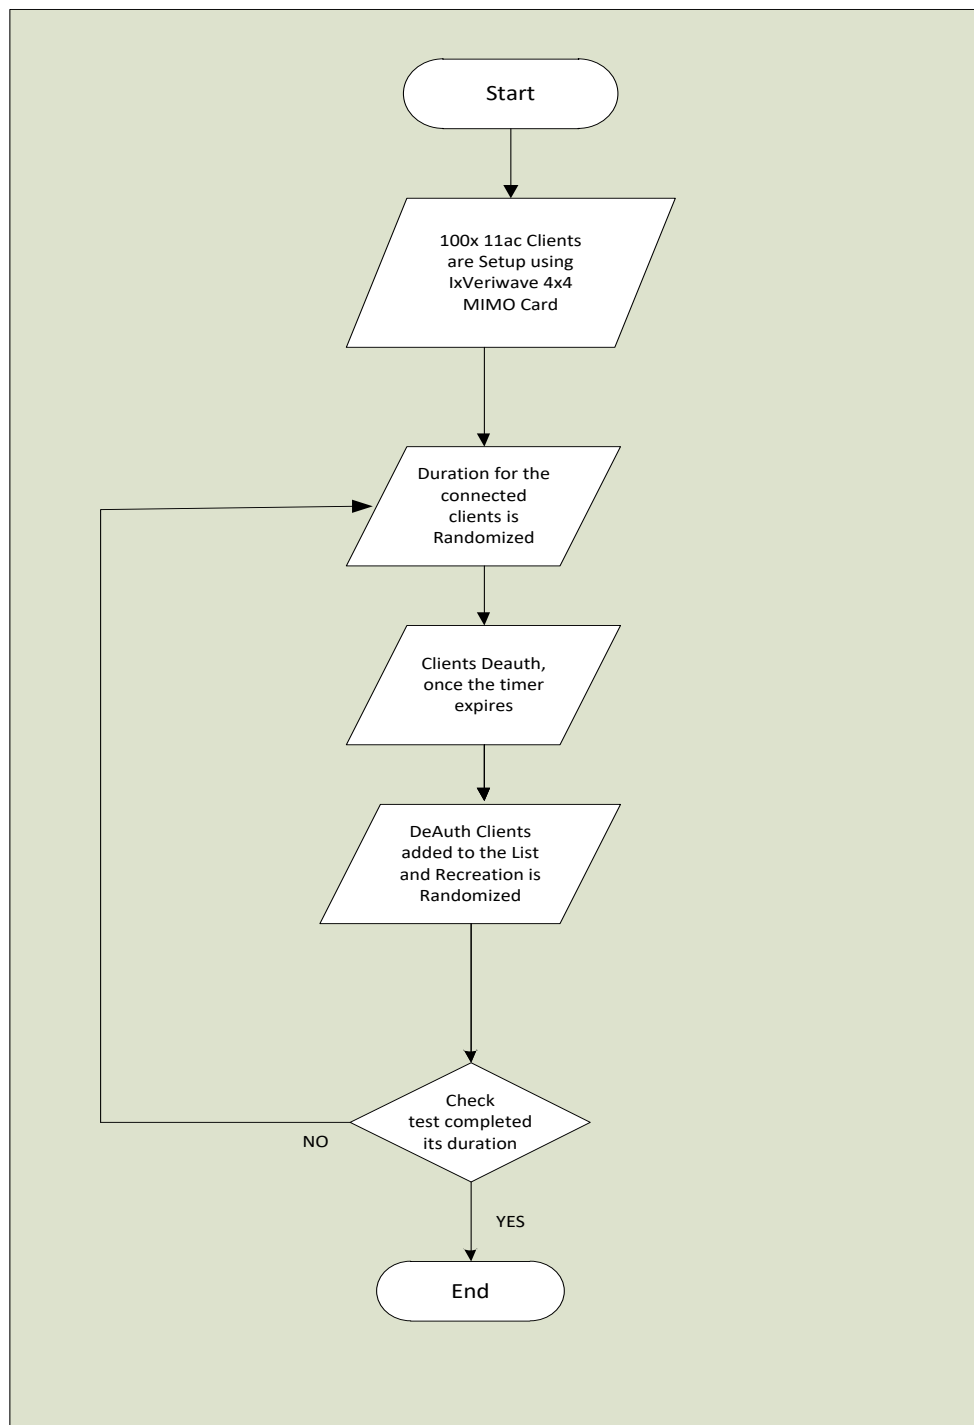


Figure 29 - Soak Test Algorithm

Conclusion

This paper provides an insight into the test philosophy of carrier provider Shaw Communications when evaluating WiFi products for use in the network.

The 802.11 standard is a complex communications channel that supports a multitude of legacy and new products currently in the market.

Shaw has taken a tiered approach in testing of new technology at ever increasing levels of integration. Shaw has found that testing the lower components performance that is traceable to known standards is the best method to engage the vendor when non compliances are found. Given the breadth of the 802.11 standard and the multitude of test cases, Shaw has found it most expedient to develop an automation framework to simplify testing for new products and performing regression testing for product improvements.

This paper has summarized the automation approach using freeware SW that meets the requirements of being a stable test platform. Example test measurements have been discussed showing how the automated framework supports these tests. The automation framework can also be easily expanded to other test requirements for WiFi product or for other unrelated products that require such test coverage.

Abbreviations

AMPDU	Aggregated MAC Protocol Data Unit
AMSDU	Aggregate MAC Service Data Unit
AP	access point
ATA	Agile Test Automation
bps	bits per second
CAPWAP	Control and provisioning of wireless access points
CPE	Customer premises equipment
CSV	Comma separated values
CTIA	Cellular Telecommunications Industry Association
dB	decibel
DHCP	Dynamic host configuration protocol
DOCSIS	Data over cable service interface specification
EMI	Electromagnetic interference
EVM	Error vector magnitude
Fps	Frames per second
GUI	Graphical user interface
GHz	Gigahertz
HTML	Hypertext markup language
Hz	hertz
LAN	Local area network
MAC	Media Access Control
MIMO	multiple-input and multiple-output
MHz	Megahertz
MCS	Modulation coding system
MPLS	Multiprotocol Label Switching
MU-MIMO	Multi-user MIMO
OSI	Open systems interconnection
OTA	Over the air
PHP	Personal home page
QA	Quality assurance
RF	Radio frequency
SP	Service provider
SQL	Structured query language
TCL	Tool command language
SCPI	Standard commands for programmable instruments
SCTE	Society of Cable Telecommunications Engineers
SMB	Small and midsize business
SNR	Signal to noise ratio
SOHO	Small office/home office
SQL	structured query language
SW	Software
TRP	Total radiated power

TCP	Transmission control parameter
TIS	Total Isotropic Sensitivity
Tx	transmit
UDP	User datagram protocol
UNI-1	Unlicensed National Information Infrastructure (band) 1
UUT	Unit under test
VOIP	Voice over IP
WiFi	Not an acronym but is a name used for referencing 802.11 specification compliant devices and networks.
WLC	Wireless LAN controller

Bibliography & References

Ref 1. CTIA Test Plan for Wireless Device Over-the-Air Performance (Method of Measurement for Radiated Power and Receiver Performance)", version 3.2.1 March 2013

Ref 2. IEEE 802.11-06/1839r1 MIMO Testing In A Conducted Environment, 2006-11-10

Ref 3a <http://www.wlanpros.com/mcs-value-achieved-clients-various-snr-levels-andrew-von-nagy/>

Ref3b <http://www.wlanpros.com/wp-content/uploads/2015/06/Revolution-WiFi-MCS-to-SNR-Single-Page.pdf>

Ref 4. RSS-247 Digital Transmission Systems (DTSS), Frequency Hopping Systems (FHSs) and License-Exempt Local Area Network (LE-LAN) Devices Issue 1 , 2015

Pay TV is Not Dead!

Myth Busting 101: It's (NOT) Inferior to OTT Cost and Value Experience

A Technical Paper prepared for SCTE•ISBE by

Charles Cheevers

CTO Customer Premises Solutions

ARRIS

3871 Lakefield Drive

Suwanee, GA 30024

678-473-8507

Charles.Cheevers@arris.com

Michael McCluskey

VP Product Management

Espial

200 Elgin St #1000, Ottawa, ON K2P 1L5, Canada

mmclluskey@espial.com

Introduction

There is not a week that goes by without a member of the media commenting on the way that consumers are consuming video content. Much of this discussion is around the threat to the Pay TV industry, changing viewing habits of consumers, and the increase in choice at different cost points. This threat is defined as coming from several areas:

- The rise of Over-the-Top (OTT) video sources and the potential for choice
- Use of the retail set-top box (STB), Smart TV, and other OTT video source and subscription based services
- Cord Shaving from Pay TV and using OTT sources
- Desire for à la carte video of OTT video services, applications and content
- Overall cost of video entertainment and the growing subscriber pushback on video costs
- New non-linear viewing experiences like time shifted content, binge viewing, and follow me video services
- The quality of the video package and relative value for money against other options
- Re-defining Pay TV to include paid for streaming services and ‘skinny bundles’

Change is certainly happening within many facets of traditional video delivery and video content and consumption. However, this paper will review what is actually happening and which changes are potentially disrupting the Pay TV industry. We will explore how the Pay TV industry is best positioned to still be the aggregation point for the home video experience.

Consumers are always looking for ways to save money. They are much more likely to cord shave if, and only if, the overall Pay TV package is not good enough. It is this “not good enough” equation that this paper will explore to ensure that all the must have factors of the Pay TV package are not only present but strengthened going forward. This “total package” will keep consumers happy (or at least not wanting to change) with their overall home video entertainment experience.

How much is good enough to circumvent the desire for lower cost for all home services? This paper will propose suggestions for the strategy to retain and potentially grow the subscribers for Pay TV services. It will show that this strategy revolves around 5 key tenets and in particular the control point the STB gives the MSO to effect these tenets.

This paper has several sections that explore the following simple 5-point formula for winning the home HDMI input battle as illustrated in Figure 1 below.

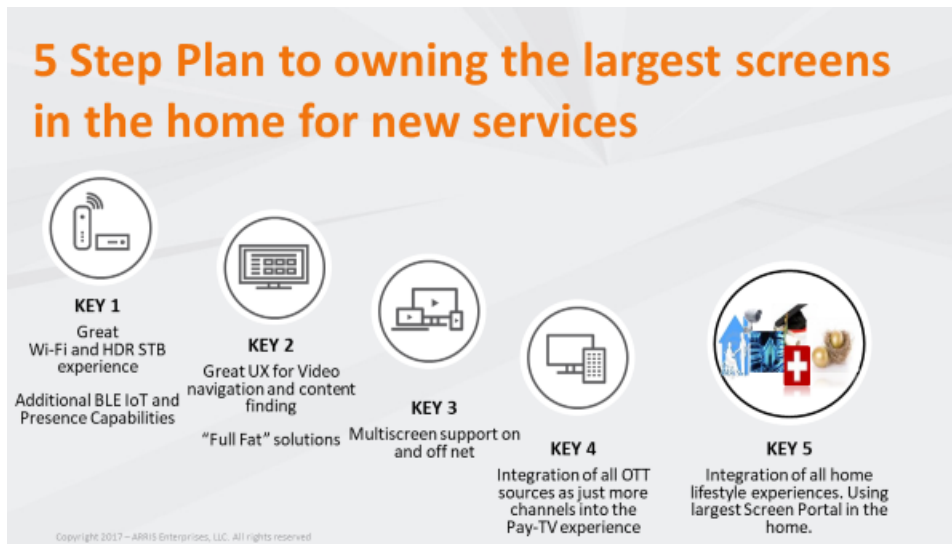


Figure 1 - The Simple 4 Keys to the Pay TV Consumer Retention Experience

Re-enforcing this simple strategy will be analyzing the following key tenets:

- (i) Operator supplied STB device to own the quality of the pixels on the TV screen
- (ii) Focus on in home Wi-Fi distribution for video with operator controlled quality end points
- (iii) Focus on the key user experience (UX) elements of the video experience
- (iv) Integration of key OTT sources into the overall Pay TV service package
- (v) Addition of new services on the TV experience, as part of the Pay TV experience
 - a. In particular, leveraging the TV for new services such as
 - i. Health
 - ii. Energy Management
 - iii. Self Help
 - iv. Education
 - v. Smart Home, Security, and Privacy
 - vi. Aging in Place

Now, read on and explore the following sections in this paper, showing the template for continued success in the Pay TV experience. The paper will first explore the current drivers and satisfaction levels of the Pay TV services available as well as the Cord Shavers OTT packages. It will do this through a combination of public information from providers and analysts and also from an ARRIS and Espial Pay TV and Cord Shaving survey that was conducted for this paper to show the dynamic in the US for existing Pay TV and Cord Shaver consumers. The first sections will outline:

- Current Pay TV solutions
 - Typical Packages and Pricing
 - ARRIS Espial Feedback from Internal Survey of Pay TV consumers
 - Public Analyst commentary as correlation
- Current OTT Video solutions
 - Typical Packages and Pricing
 - ARRIS Espial Feedback from Internal Survey of Pay TV consumers
 - Public Analyst commentary as correlation

Then the paper will look at the data from over 144 thousand sampled homes on their use of the features in their Pay TV solution. This solution based on Espial Elevate solution also has Netflix integrated. It's an ideal platform to see user behavior for what is working in the combined world of Pay TV, Linear content, and the OTT source integrated.

The remainder of the paper will outline suggestions and recommendations for Multichannel Video Programming Distributors (MVPDs) to continue to provide the central entertainment experience for the customer. While cost may always be ever present, the consumer will not churn when all the other factors of user experience and new sticky services are present. If they are not present, then the lure of cheaper video entertainment and a Do It Yourself (DIY) à la carte solution may entice them off an MSO service offering. This section will outline:

- The role of Multiple System Operator (MSO) devices versus customer owned devices versus applications as part of the converged home user experience for video delivery and future other experiential services
- The video bundle and the “thousands of channels to the one you want” problem that seems to always drive the consumer to look for lower cost points for video service
- Taking control of the TV for more than just the TV show and movie experience by leveraging the largest video portal in the house for all digital home experiences and solutions

Content

1. Current Pay TV solutions

Current Pay TV solutions in the USA are widely spread in cost, perceived value for money, user experience, and where they are on their journey to incorporate new OTT sources and new TV experiences.

See below for the current costs for Basic to Premium TV access from a representative group of cable and telco/satellite providers. This information is all taken from their respective websites pricing in June of 2017 and is subject to change.

To just buy TV service is not easy or as attractive as bundling a service with Internet and phone. However, for this introduction of the Pay TV industry to the reader, the figures below show just the cost of TV service where possible for the new subscriber. Some of the offers include features like Digital Video Recorder (DVR) and associated access to streaming services. However, this paper is not focused on trying to define a single value for money estimate.

1.1. Comcast TV Packages

Comcast services for TV range from \$31/pm for Basic Cable TV service to \$109/pm for premium channels like HBO, Cinemax, etc.

One of Comcast's most popular TV + Internet service is the 220+ Channel for TV with additional 75 Mbps broadband service.

Comcast also includes a bundle for Cord Shavers with 300 Mbps Broadband and access to Basic Cable TV with inclusion of HBO package for both TV and streaming.

1.2. Altice

Altice/Optimum TV services range from \$64.95 to \$109.95 as illustrated below with Value TV up to Gold TV packages available.

1.3. Charter

Charters has a \$59.99 TV service for over 110+ channels including DVR support to 4 TVs.

1.4. DIRECTV

DIRECTV has new TV packages ranging from \$50 to \$125 a month.

1.4.1. *The 2015 FCC Report on the Cost and Channels Available to the US Consumer*

The FCC released its yearly report “Statistical Report on Average Rates for Basic Service, Cable Programming Service, and Equipment” in October 2016. It outlines the costs of Basic, Expanded, and the most popular packages as well as Customer Premises Equipment (CPE) costs for 2015. The report also included some competitive analysis for the cable industry, Direct Broadcast Satellite (DBS) providers, and wireless offerings in 2015. With the increase in unlimited data packages and streaming packages from DIRECTV the wireless number is growing faster from the 2015 analysis.

Figure 2 below shows the monthly price of programming across cable and satellite services were an average of \$23.79 for Basic Cable to an average of \$81.75 to the most popular Extended Cable package. That package was also similar with DBS competition average prices. In 2015, the average increase of Cable TV service was about 2.45% over the previous year.

Table 1 Monthly Price of Programming By Status of Effective Competition January 1, 2015								
Cable Service	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
				Second Cable Operator Overbuild Subgroup			DBS	Wireless / Low Penetration
				Incumbent	Rival	Both		
Basic service	\$23.79	\$24.55	\$22.96	\$21.43	\$20.06	\$21.24	\$23.29	\$25.57
Annual change	2.3%	2.5%	2.2%	4.3%	3.3%	4.1%	1.6%	2.2%
Expanded basic	\$69.03	\$67.85	\$70.31	\$69.46	\$74.05	\$70.11	\$70.41	\$69.97
Annual change	2.7%*	3.3%*	2.0%*	3.2%*	10.3%*	4.2%*	1.3%	1.8%
Next most popular	\$81.75	\$81.86	\$81.64	\$78.85	\$86.80	\$79.97	\$82.15	\$82.27
Annual change	2.2%*	2.8%*	1.5%	3.1%*	7.5%*	3.8%*	0.8%	2.0%

Sources: Attachment 2. * Annual change is statistically significant at the 95% confidence level. Expanded basic prices include basic service prices, and next most popular service prices include expanded basic prices.

Figure 2 - Monthly Price of Programming in the US - 2015

Figure 3 below shows the average price per channel in 2015, which overall has continued to decrease as more channels are added. The number of channels raises a lot of discussion amongst consumers as they typically tend to cluster on about 15 channels (ARRIS numbers & Nielsen show about 17 channels) per household for most of their viewing.

Table 2 Average Price per Channel By Status of Effective Competition January 1, 2015								
Cable Service	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
				Second Cable Operator Overbuild Subgroup			DBS	Wireless / Low Penetration
				Incumbent	Rival	Both		
Basic service	\$0.602	\$0.682	\$0.516	\$0.447	\$0.698	\$0.483	\$0.519	\$0.596
Annual change	-2.4%	-1.3%	-3.9%	-0.7%	6.1%	0.8%	-5.3%	-3.9%
Expanded basic	\$0.456	\$0.497	\$0.412	\$0.400	\$0.475	\$0.411	\$0.412	\$0.419
Annual change	-1.8%	-0.6%	-3.3%	-2.4%	3.3%	-1.5%	-3.7%	-4.3%
Next most popular	\$0.359	\$0.392	\$0.326	\$0.328	\$0.351	\$0.331	\$0.323	\$0.339
Annual change	-2.3%	-1.0%	-4.1%	-0.7%	4.7%	0.1%	-5.4%	-4.6%

Source: Attachment 4. * Annual change is statistically significant at the 95% confidence level. Price per channel is the service price divided by the number of viewable channels with that service. Expanded basic prices include basic prices, and prices of the next most popular service include expanded basic prices. Similarly, expanded basic channels include basic channels, and next most popular service channels include expanded basic channels.

Figure 3 - Average Price Per Channel in January of 2015

Figure 4 below shows the 2005-2015 cost of each service and the 10-year compound average rate of change of the cable TV bill is 5.2%. In comparison, the CPI (Consumer Price Index) compounded annual change over the same period was 2.5%. Thus, the consumer feels their cost of TV has increased in cost at 2x+ the rate of other essential consumer products and services. This perception is the basis of much of the Cord Shaver momentum and the cost dissatisfaction amongst existing Pay TV consumers.

Table 4 Historical Averages							
Year	Basic Service Price	Expanded Basic Service			Next Most Popular Service and Equipment	CPI Index	
		Price	Channels	Price per Channel		All Items	Cable
2005	\$14.30	\$43.04	70.5	\$0.620	\$56.03	127.2	169.6
2006	\$14.59	\$45.26	71.0	\$0.650	\$59.09	132.2	174.4
2007	\$15.33	\$47.27	72.6	\$0.670	\$60.27	135.0	179.0
2008	\$16.11	\$49.65	72.8	\$0.680	\$63.66	140.8	183.9
2009	\$17.65	\$52.37	78.2	\$0.710	\$67.92	140.8	186.5
2010	\$17.93	\$54.44	117.0	\$0.560	\$71.39	144.5	191.9
2011	\$19.33	\$57.46	124.2	\$0.569	\$75.37	146.9	192.0
2012	\$20.55	\$61.63	149.9	\$0.505	\$78.91	151.2	199.8
2013	\$22.63	\$64.41	159.6	\$0.484	\$81.64	153.6	206.5
2014	\$22.78	\$66.61	167.3	\$0.496	\$84.65	156.0	212.0
2015	\$23.79	\$69.03	181.3	\$0.456	\$86.83	155.8	216.4
Compound Average Annual Rate of Change							
5-year average	5.8%	4.9%	9.2%	-4.0%	4.0%	1.5%	2.4%
10-year average	5.2%	4.8%	7.1%	-1.4%	4.5%	2.0%	2.5%

Source: 2005-2015 surveys. See Attachment 7 for references. Attachment 7 also shows the series back to 1995.

Figure 4 - 5-10 Year Change in the Pay TV Monthly Bill

The number of channels available in the 3 service tiers has also continued to rise (Figure 5) with the range of increase being 3.2% to 3.9% in added channels, tracking about 2-3 additional channels at the basic tier, and to 8-9 additional at higher tiers. Adding more channels does not necessarily add value for a lot of Pay TV subscribers, and particularly the OTT Cord Shavers, as they believe they need more à la carte services.

Table 5 Number of Video Channels By Status of Effective Competition January 1, 2015								
Cable Service	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
				Second Cable Operator Overbuild Subgroup			DBS	Wireless / Low Penetration
				Incumbent	Rival	Both		
Basic service	58.8	54.3	63.7	72.3	51.0	69.3	62.4	58.0
Annual change	4.9%*	5.1%	4.7%	6.1%	2.6%	5.6%	4.1%	7.6%
Expanded basic	181.3	169.4	194.0	199.8	191.0	198.6	193.3	186.4
Annual change	4.4%*	4.7%*	4.1%*	5.8%	6.7%	5.9%*	3.2%	6.7%
Next most popular	264.4	248.6	280.4	281.1	291.3	282.5	280.3	274.5
Annual change	3.2%*	3.8%*	2.7%	3.9%	2.3%	3.9%	2.0%	6.3%

Source: Attachment 6. * Change is statistically significant at the 95% confidence level. Table refers to viewable channels offered with the service at no extra charge including those requiring equipment to view. Expanded basic channels include basic channels; next most popular service channels include the expanded basic channels.

Figure 5 - Increase in the Number of Channels Added to a Pay TV Package Annually

Within the Basic Tier Cable TV service, the composition of programs is typically similar for Cable, DBS, and Wireless with around 58-69 channels, because the majority are broadcast or other channels.

Table 6 Channel Composition of Basic Cable Service January 1, 2015								
Video Channel Category	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
				Second Cable Operator Overbuild Subgroup			DBS	Wireless / Low Penetration
				Incumbent	Rival	Both		
Broadcast	33.6	30.2	37.3	40.3	36.6	39.8	36.8	33.4
PEG	3.4	3.3	3.5	4.1	3.4	4.0	3.4	2.9
Leased access	1.0	0.9	1.2	1.1	0.7	1.0	1.2	1.9
Regional sports	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Other channels	20.7	19.8	21.6	26.7	10.2	24.4	20.9	19.7
Total	58.8	54.3	63.7	72.3	51.0	69.3	62.4	58.0

Source: 2015 Survey. By individual channel (standard definition, high definition, and multicast).

Figure 6 - Channel Composition across Cable and DBS

Both sports and news channels account for a lot of the thinking around the need for Pay TV and linear viewing. The number of sports networks channels is about 3.4 in the Expanded Basic service with 4.2 in the next most popular tier.

Table 7 Regional Sports Networks By Status of Effective Competition January 1, 2015								
Cable Service	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
				Second Cable Operator Overbuild Subgroup			DBS	Wireless / Low Penetration
				<i>Incumbent</i>	<i>Rival</i>	<i>Both</i>		
Basic service	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Expanded basic	3.4	3.3	3.4	3.5	5.9	3.9	3.4	2.5
Next most popular	4.2	4.2	4.1	3.6	6.3	4.2	4.2	2.7

Source: 2015 Survey. Channels are the number of channels offered at no extra charge including those requiring equipment to view. Expanded basic channels include basic channels, and next most popular service includes expanded basic channels.

Figure 7 - Number of Sports Channels in the Different Pay TV Tiers - 2015

The CPE or the STB is also another discussion topic when it comes to Pay TV. The cost of leasing the device versus the typical one time purchase at retail of an OTT video device also sways opinions on the reasons to cord shave and why a Pay TV customer may not be satisfied. The average cost per month of the single leased STB was about \$8.50 in 2015. Additional costs can be incurred for additional outlet devices.

Table 8 Price for Most Commonly Leased Customer Premises Equipment January 1, 2015								
Cable Service	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
				Second Cable Operator Overbuild Subgroup			DBS	Wireless / Low Penetration
				<i>Incumbent</i>	<i>Rival</i>	<i>Both</i>		
Basic service	\$8.40	\$8.30	\$8.49	\$8.23	\$10.10	\$8.51	\$8.51	\$8.18
Annual change	1.5%	0.7%	2.2%	6.2%	-0.9%	5.1%	0.8%	3.4%
Expanded basic service	\$8.34	\$8.17	\$8.49	\$8.23	\$10.10	\$8.52	\$8.51	\$8.18
Annual change	1.4%	0.5%	2.2%	6.2%	-0.8%	5.1%	0.8%	3.4%
Next most popular svc.	\$8.76	\$8.41	\$9.07	\$8.65	\$10.10	\$8.87	\$9.17	\$9.06
Annual change	1.3%	0.6%	1.8%	5.5%*	-0.8%	4.5%	0.7%	1.8%

Source: Attachment 5. * Annual change is statistically significant at the 95% confidence level. These prices are for a single lease of the most commonly leased equipment and not the average charge per customer or per household, which would depend on the number and type of equipment leases.

Figure 8 - CPE Leasing Costs for STB for Pay TV in 2015

The final table from the 2015 FCC study of the Cable TV sector is the type of features supported by the supplied STB and the software (Figure 9). Some surprising numbers for DVR support were found. 2016 and 2017 saw increased use of DVR as part of the architectures deployed. For example, Comcast X1 and the use of the XG1 DVR gateway (GW) device with only 28% of higher tier services having DVR capability. Additionally, the number of High Definition (HD) channels featured below in each tier is a % of the overall channels offered. The Basic Tier having a higher proportion of HD Channels at 83% versus about 45% of higher tier with more channels offerings. The snapshot showed that many of the additional channels are Standard Definition (SD) only, and shows opportunity to further improve consumer viewing satisfaction with HD upgrade.

Table 9 Features Offered With Most Commonly Leased Customer Premises Equipment January 1, 2015									
Cable Service	Feature	Overall Average	Non-competitive	Effective Competition	Effective Competition Subgroups				
					Second Cable Operator Overbuild Subgroup			DBS	Wireless/ Low Penetration
					Incumbent	Rival	Both		
Basic service	DVR	27%	20%	34%	55%	4%	48%	28%	41%
	HD	83%	80%	87%	92%	93%	92%	87%	66%
	IPG	85%	82%	88%	89%	46%	83%	90%	86%
	RCU	92%	89%	95%	92%	96%	93%	97%	95%
Expanded basic	DVR	28%	22%	34%	60%	9%	53%	27%	44%
	HD	43%	39%	47%	70%	93%	73%	39%	48%
	IPG	93%	93%	94%	93%	100	94%	94%	91%
	RCU	97%	95%	98%	97%	96%	97%	99%	100%
Next most popular	DVR	28%	23%	34%	60%	5%	52%	27%	44%
	HD	47%	42%	53%	74%	87%	76%	44%	55%
	IPG	98%	95%	100%	100%	100	100%	100	100%
	RCU	97%	94%	99%	97%	95%	97%	100	100%

Source: 2015 Survey.

Figure 9 - Features Offered for Pay TV with the Different Tiers of Pay TV Services - 2015

The full report from the Federal Communications Commission (FCC) for the 2015 update on the Average rates paid for Cable Programming can be found at https://apps.fcc.gov/edocs_public/attachmatch/DA-16-1166A1_Rcd.pdf

1.5. Nielsen's Q4 – 2016 Total audience report

Moving on to some of the 2016 Total audience report published by Nielsen:

<http://www.nielsen.com/us/en/insights/reports/2017/the-nielsen-total-audience-report-q4-2016.html>

It is interesting in that it shows that there is a large increase in the use of Smartphone via App or Web with more than 100% growth in the last 2 years. See Figure 10 below where in Q4 2016 consumers spent 2 hours and 16 minutes on average **per day** on their smart phone App/Web.

In the same 2 year period, linear TV viewing **per day dropped only 8 minutes**. Time shifted TV usage grew on average 2 more seconds per day to 33 minutes on average. One of the reasons is that many US homes have TVs left on for ambient noise. Additionally, there has been an 18% increase in the consumption of news from 2015 to 2016, with 44% increase in watching cable news the most significant. While consumers are doing more engagement with phones and tablets they also seem to prefer the larger screen as a bigger part of their at home entertainment. Figure 11 shows the change in consumption per specific medium from 2014 to 2016.

EXHIBIT 1 - BASED ON THE TOTAL US POPULATION




Figure 10 - Consumer Change in Device Usage Time from 2014 to 2016

EXHIBIT 2 - BASED ON USERS OF EACH MEDIUM

	Q4 2014	Q4 2015	Q4 2016
Live+DVR/Time-shifted TV	6:39	6:45	6:24
DVR/Time-shifted TV	2:00	2:01	1:57
AM/FM Radio	2:42	2:43	2:45
DVD/Blu-Ray Device	1:42	1:41	1:51
Game Console	2:48	3:02	3:00
Multimedia Device	2:30	2:23	2:23
Internet on a PC	2:13	2:55	3:11
App/Web on a Smartphone	1:49	2:02	3:17

The data sources in Exhibit 2 should not be added or subtracted; they are based on users of each medium and the bases vary by source. Panel enhancements made in March and August 2016 impacted mobile reporting.

Figure 11 - Specific Time Spent on Each Medium from 2014 to 2016

Just like sports viewing watching the news accounts for one of the main reasons for both linear viewing and the desire for a Pay TV service. There was significant increase in watching the news from 2015 to 2016 mainly can be attributed to the US presidential election and key world events. Cable TV viewing of the news grew 44% Year-over-Year (YoY) with overall news watching growing 18% YoY.




Figure 12 - The Rise in News Watching from 2015 to 2016

1.6. Some of the Analyst Commentary on the Pay TV and Cord Shaver Markets

As there is much analyst commentary on the status of Pay TV solution and the Cord Shaver market the following are some representative dialog based on analysis of Pay TV market trends. MoffettNathanson below show the decline in US Pay TV growth for the last 6 years where the industry has lost 3.9% since Q1 2010 and the trend is downwards. However, the inclusion of Pay TV streaming services like Sling TV and DIRECTV now shows that these services when included reduce the shift by over 1.1%. An additional trend that has not yet fully been realized in this analysis is the inclusion of the primary OTT sources into the traditional Pay TV experience.




Figure 13 - MoffettNathanson Report Showing Pay TV Growth from Q1 2010 to Q2 2017

In a report from Mary Meeker of Kleiner Perkins on Internet Trends 2016 cited the following reasons for the Cord Shavers and highlighted the sensitivity around price and the convenience factor. The first is understandable and the second is the key to target for the Pay TV provider.




Figure 14 - Kleiner Perkins – Q4 2016 Video Trends Report

The Kleiner Perkins report showed the significance of cost in the decision process with 80% of consumers making choices with cost as the primary factor. The ability to cord shave and use an Internet Streaming service drove 48% of people to make the decision. The ability to catch both local and broadcast channels was a factor with 27% of people. A feature, which can be replicated by Pay TV solution, like bingeing full seasons of a series was a factor in 19% of people. 11% of people felt that all their viewing needs were mostly met by Streaming service original content. One surprising factor is that 13% of people will make the change when they move or relocate.

1.7. ARRIS Pay TV Survey Assessment

ARRIS and Espial also conducted its own survey on nearly **900 US respondents across multiple States and job functions**. The survey showed that **78% of those surveyed had a Pay TV service and 22% did not watch TV or had shaved the cord**.

The distribution of USA respondents by state is shown below in Figure 15. Pennsylvania, California, and Georgia had the highest number of respondents to the survey. This survey was conducted in June 2017 and also included another 800 respondents in various countries across the globe.

For the US respondents the survey showed that the split was Pay TV service (78%) and those that did not (22%). This provided a good insight into the drivers, trends, and wants from each tier.




Figure 15 - State Distribution of Respondents in ARRIS Q2 2017 Pay TV and OTT Survey

With 78% of the respondents having a Pay TV service we first looked into this group to see the breakdown of their satisfaction, use, and wants.




Figure 16 - Respondents Who Had Pay TV and Who Did Not Have Pay TV




Figure 17 - Percentage of Respondents with Pay TV by Service Provider

35% of Pay TV respondents were using Verizon. Comcast was second at 31%. Respondents of AT&T and DIRECTV were recorded separately if they did not cite AT&T in their response. DIRECTV was about 7% and AT&T about 5% of overall respondents. Only a few respondents to the Pay TV section cited Netflix or Sling as their primary Pay TV service.

Of the 78% people surveyed who had a Pay TV service, the first big piece of information was that 72% of them were actively looking for or believed they wanted a reduction in their cost for Pay TV. 28% were not actively looking for a reduction in their Pay TV service.

This statistic tracks the Kleiner Perkins Mary Meeker report and is generally a function of most people wanting value for money. However, with competition and choice, the number of people actively looking to reduce Pay TV at 72% shows the importance of other elements of the Pay TV and overall experience for the customer to keep them from switching.




Figure 18 - How Many of the Survey Respondents Were Actively Looking to Reduce Their Pay TV Bill

Let's explore the most important part of this equation, the perceived value for money versus the inertia of changing Pay TV provider versus the inertia of Cord Shaving.

The average overall satisfaction with Pay TV service was 6.3%. This incorporated both the cost and user experience. This is quite low and suggests a fundamental issue or a potential to improve with deliberate steps to focus on the value proposition for Pay TV differently.



Figure 19 - Overall Satisfaction Level with Pay TV by Number of Respondents (1-lowest 10 Highest)

36% of respondents scored 8 or more on the 10-points satisfaction scale.:

34% of respondents scored 5 to 7 on the 10-point satisfaction scale.

20% of respondents scored under 5 on the 10-point satisfaction scale.




Figure 20 - Overall Pay TV Satisfaction Level (1-10) by Overall Percentage

Taking the example of Comcast respondents — they generally tracked to the same trend for overall satisfaction.




Figure 21 - Comcast Overall 1-10 Scoring for Satisfaction Level: 1 = lowest 10 = highest

Taking all the Service Providers who scored an 8-10 on overall survey and then showing their overall % of the 8-10 score against their overall participation level in the survey yields the chart in Figure 22 below. Comcast, DIRECTV (no AT&T), Dish, and Cablevision all scored higher percentage of 8-10 satisfaction scores than their overall participation level in the survey. Interestingly Verizon scored lower % of 8-10 scores to their participation level in the survey. There was a direct correlation between the customers scoring their Service Provider 8-10 level satisfaction and their value on integrated OTT sources as well as the UX services. Over 85% of respondents who rated their SP 8-10 commented on the importance of their OTT service integration, single High-Definition Multimedia Interface (HDMI) port and other features such as multiscreen and other device streaming services.




Figure 22 - SP % of 8 -10 Scores Versus Their Overall Participation Level in Survey

At the other end of the satisfaction scale from 1-5 the following was the respondent's answers. In this case Comcast scored the most 1-5 responses and it was higher than their overall percentage of participation. Verizon had the second highest number of 1-5 responses but was at about 50% the rate of participation and DIRECTV also scored lower 1-5 rate than the participation level. See Figure 23 below. Charter and AT&T also scored higher percentage of 1-5 ratings than their participation level.




Figure 23 - The % of 1-5 Ratings (1 Lowest Score) from The Survey Respondents Against Their Overall Survey Participation

The overall satisfaction level for Cost of Pay TV was the lowest factor and most influential on the overall result — with just 3.9% happy with the cost of their Pay TV service.



Figure 24 - Cost Satisfaction with Pay TV Service




Figure 25 - Cost Satisfaction Level 8-10 by Service Provider

As you can see above Comcast, DIRECTV, Charter, and AT&T scored higher % of 8-10 than their survey participation. Interestingly Netflix did not. Clearly respondents are already trying to get Netflix to reduce costs as well. Figure 26 below shows the lower rating levels for Cost satisfaction. In this one, only Verizon scored a lower number of 1-5 ratings to their overall participation in the survey. Comcast showed about a 50% 1-5 rating per their participation in the survey.




Figure 26 - Cost Ratings 1-5 (5 Lowest Satisfaction) Against Overall Survey Participation Level

The UX satisfaction level of the experience fared better with an overall average satisfaction level of 6.1%. This shows that there is substantial opportunity to improve this — to balance some of the dissatisfaction of cost levels for the service. Figure 27 below showed consumer's satisfaction level with UX from 1-10.



Figure 27 - UX Satisfaction Levels - 1-10 - 1 being lowest - 10 Highest

Comcast scored the largest number of 8-10 satisfaction ratings with over 50% more than their participation levels. DIRECTV and Dish scored higher 8-10 satisfaction levels to their participation level. Verizon had 27% 8-10 ratings of their participation level.




Figure 28 - UX Satisfaction levels (8-10) 10 highest against participation levels

Figure 29 below shows the more dissatisfied level of UX for the respondents — scores for 1-5. Verizon showed a lower number of 1-5 respondents versus their participation level — about 50% better. Comcast scored more 1-5 than their participation level. Interestingly this number is much lower when you review the data for the words XFINITY, X1 — where this is much lower. Charter, AT&T, Cox, Wow, and Cablevision all scored higher 1-5 than their participation level.




Figure 29 - UX Satisfaction ratings (1-5) 1 worst - shown with participation level in survey

The use of a free-to-air (FTA) tuner was also interesting in the poll of consumers. The Pay TV subscribers poll showed that 27% of Pay TV users also used FTA tuner in their home video solution — probably for additional TV outlets rarely used in the home and in conjunction with OTT only TV. Figure 30 below shows this result from survey.




Figure 30 - Pay TV Consumers Using FTA Tuners

As you would expect a higher percentage 63% of Cord Shaver consumers polled were using an FTA source with their video solutions — see Figure 31.




Figure 31 - Cord Shavers using FTA Tuners

As cost remains the main driver for decisions on what the consumer chooses for their entertainment package at home — our survey showed that the US — 65% of people are still paying above \$100 for their total TV service. We found that people included additional to their service provider provided Pay TV solution — their additional subscriptions to Netflix, Amazon, Hulu, and others — as well as including their STB and additional STB costs on a monthly basis.




Figure 32 - How Much are People Paying for Their Pay TV Service

For the people with Service Provider Pay TV service over 87% of people also subscribed to an OTT service as well.




Figure 33 - Percentage of Pay TV Consumers who also Subscribed to OTT Sources

As with the Kleiner Perkins report — ARRIS/Espial survey showed that Netflix was the dominant OTT source. Amazon Video, YouTube/Red, and Hulu made up the top additional top 4 sources. Respondents called out Roku, Apple TV, and Sling TV specifically as their OTT service and also some respondents cited DIRECTV NOW and XFINITY Streaming/Streampix as their OTT source.




Figure 34 - OTT Sources Used by Pay TV Subscribers

ARRIS also independently tracks the Downstream consumption of video traffic on the Hybrid Fiber-Coax (HFC) / Data Over Cable Service Interface Specification (DOCSIS) networks of the US and the following table also tracks consistently with the Kleiner Perkins and ARRIS/Espial survey. Netflix accounting for almost 30% of Downstream video IP traffic, with YouTube at ~17% and Amazon Video at about 3% - with Hulu, Service Provider own streaming services, Sling TV, and others under 2%.




Figure 35 - Downstream Traffic Percentage by Video OTT Source

Kleiner Perkins — Mary Meeker Fixed Access Video Traffic Share Leaders correlates with ARRIS data and survey.




Figure 36 - Kleiner Perkins - Share of Downstream Traffic % NA 2016

The survey response to “What devices consumers use to access OTT content” showed that the Smart TV was most commonly used with 33% of respondents declaring it their preferred OTT device. OTT STB only accounted for 23% of respondents OTT viewing. 27% of respondents responded “other” as their answer in the survey. A subsequent check on this showed that the ‘Others’ were typically smart phones, tablets, and PCs — which were not specifically cited in the question choices but the respondents entered them specifically.




Figure 37 - Most Common Used Devices to Access OTT Services

The device split is shown below in Figure 38 with 77% of homes having both an OTT STB and Smart TV, 25% of homes having all 3 (Smart TV, OTT STB) and Integrated in their Pay TV solution.




Figure 38 - Pay TV Users with Multiple OTT Devices per Home

The most used devices defined as “Other” (27% of respondents viewing devices for OTT content) for watching OTT sources were tablet/iPad at over 70%, second Phone at 48% of people using it consume OTT content, and Game consoles almost 39%.




Figure 39 - The Other Devices That are Used to Access OTT Content




Figure 40 - Would Respondents Like to Have Their OTT Sources Integrated in Pay TV Service

We asked the respondents specifically if they would like to have their OTT sources integrated in their Pay TV UX/user interface (UI) and 61% said that they would — with 21% saying ‘No’ and 19% - not sure/did not know.

For respondents with Pay TV — when we asked if they used “Free to Air Tuners” — 18% of respondents only said, “Yes”, and 49% said explicitly, “No.” The remainder were not sure or did not know (subsequent follow up shows that not everyone is aware what a Free to Air Tuner is). This low use of FTA is expected with people who have local and broadcast channels included in their Pay TV lineup.




Figure 41 - As a Pay TV User Do You use FTA Tuner

Interestingly, when asked if it was of value to integrate FTA tuner into STB– 46% said it was of value with 17% explicitly saying, “No.” 37% of respondents did not answer or did not know.




Figure 42 - Answer to Question if FTA Tuner Added to OTT STB was of Value




Figure 43 - Is changing HDMI port for OTT Source Annoying and Something to Improve

As the HDMI source issue is deemed a constant source of user dissatisfaction we posed the question about respondent's dissatisfaction with changing input sources on HDMI interface. Figure 43 above shows that 36% of respondents replied they did not like switching to other source, 28% of respondents replied — they don't care and 36% of people did not respond or did not know. With a subsequent check of the respondents — the question was not phrased well and it is believed a higher number of respondents would have voted that having a single HDMI source for all video inputs would be a desirable feature.




Figure 44 - Does Your Pay TV Service Offer Multiscreen Support

From Figure 44 above — 58% of respondents said they had a Pay TV service that also supported multiscreen viewing. 9% said explicitly, “No” and 33% were not sure or did not answer. A subsequent check found that many respondents are not fully aware of the streaming features they also have as part of their Pay TV service — for example, respondents who were checked did not realize that Streampix was available to them as part of their XFINITY X1 service.




Figure 45 - Where Does the Multiscreen Service Work?

In trying to determine if multiscreen services only work in the home or when mobile — we asked the question in our survey. From Figure 45 above - 12% said it was ‘in home only’ and 41% stating they could use the service outside the home. 14% of respondents said they ‘did not have multiscreen’ (Higher than the original 9% answer) and 33% did not know or did not answer. Figure 46 below showed that 33% of Pay TV respondents said a multiscreen service was important. 30% interesting said, “No” and a high number 32% ‘Did not answer’.




Figure 46 - Is Having a Multiscreen Service Important to You

In trying to understand the use of linear or time shifted content we asked consumers ‘what percentage of time do they spend watching live versus time shifted content’ and got the following breakdown in Figure 47. The responses showed that 31% of people did not answer the question. 39% of respondents watched it over 20% of the time — and 30% 0-20% of the time. As stated above from the Nielsen and Kleiner Perkins data — this answer does not fully reflect what ARRIS notes from its analysis of channel usage/time on deployed STB where Linear viewing still accounts for 70% of viewing with 15% DVR and 15% Video on Demand (VOD) being generally what is observed on Pay TV systems. We believe the 70% of linear content is skewed with several factors:

- Consumers like to leave TV’s on for ambient background viewing/noise including channels like Music Choice
- Consumers only standby the TV and not the STB — the STB typically remaining locked on Quadrature Amplitude Modulation (QAM) channel




Figure 47 - What % of Time Do You Spend Watching Live Versus Time Shifted Content

Respondents who had Pay TV were asked what would improve their Pay TV offering. Figure 48 below showed the majority of respondents Cost (29%), À la carte (25%), OTT Content integrated (19%), and Skinny Bundles (14%) were the > 10% requests. Significantly, the integration of OTT services and better user experience is something that seems to make a difference in retention of customers.




Figure 48 - What Improvements Would Respondents Like to See in Pay TV Offering

2. Current OTT Video Solutions

2.1. Typical Packages and Pricing

The following sections do a quick analysis of the most relevant OTT packages. As was shown above with the OTT sources that Pay TV subscribers had — the following Figure 49 illustrates the OTT services that they have shaved the cord for The data shows that of the respondents that replied they cited Cost and ‘Having what they needed’ as the major reasons why they shifted to cord shave.




Figure 49 - What Was the Reason You Don't Have a Pay TV Service

Kleiner Perkins report also highlighted cost at 80% as the principle reason why people do shave the cord and cut Pay TV service.

Video = Why Cord-Cutting? Lower Price + Convenience




Figure 50 - Kleiner Perkins - Reasons for Cord Cutting

Of the ‘Other’ reasons cited in the ARRIS survey — there is a high correlation to consumers who are not active TV watchers and have more of an emphasis on using time for other things. 55% of respondents don’t have enough time for TV. 36% of people watch local channels and news only with FTA solution. 36% of people watch local channels and news only with FTA solution.




Figure 51 - Additional Reasons People Cited for Other — When Moving Away from Pay TV

The OTT services that they are using are details below in Figure 52 and shows that 14% of respondents had all 4 services of Netflix, YouTube, Amazon Prime, and Hulu. Netflix was used by 34% of the respondents, YouTube 28% of the respondents, Amazon Video/Prime by 27% of respondents, with 9% watching Hulu.




Figure 52 - What OTT Sources are Primarily Used by Cord Shavers

The respondents had the chance to answer with ‘Other’ sources (which was 20% of the OTT respondents) and the following is the responses recorded. Sling TV (22%) was the highest other source — with DIRECTV Now and Sony PlayStation Vue also registering over 10% the next most significant.




Figure 53 - Definition of Other Sources Outside the Ones Defined for Answer In Survey

2.2. Overview of OTT Main Cord Shaver OTT Source Packages

Based on the respondents OTT sources — the following sections give a quick review of the pricing for those services as of June 2017.



Figure 54 - Main Sources Favored by Cord Shavers

2.3. Netflix

Netflix is the most favored OTT source for Cord Shavers. Its monthly packages range from \$7.99 to \$11.99 per month.

2.4. YouTube

YouTube is essentially a free video service — but with the increase in advertising in YouTube — YouTube Red has now emerged offering ‘uninterrupted’ video watching at \$9.99 per month.

2.5. Hulu

Hulu has several packages starting at \$7.99 per month for access to its streaming library with Limited Ads, for example \$11.99 for no ad interruption. The live TV package is \$39.99 per month. For an additional \$14.99 per month — you can get unlimited DVR and Unlimited screen access.

2.6. Amazon Prime Video

Amazon Prime Video service is integrated into the Amazon ecosystem and offers incentives for people who use Amazon shopping. Using Amazon Prime at \$99 per year — you get free shipping on Prime products and access to Amazon Video.

Amazon Prime is a membership program that gives customers access to streaming video, music, e-books, free shipping, and a variety of other Amazon-specific services and deals. Amazon Prime is a no-brainer for those heavily invested in the Amazon ecosystem. And if \$99 per year seems like too much of a commitment, you can opt for the monthly \$10.99 plan or the Prime Video only plan, which costs \$8.99.

2.7. PlayStation Vue

For the basic \$39.99-per-month Access subscription, users get over 45 popular broadcast, cable, movie, and sports channels in complete HD, including AMC, FX, CNN, BBC America, ABC, FOX, and others. For the Core package at \$44.99 per month, subscribers receive access to have 60 channels, including the NFL Network, ESPN U, and ESPNNews. The \$54.99 Elite Package includes 90 channels, such as Telemundo, Chiller, EPIX Hits, the Esquire Network, E!, and more. Finally, the Ultra package priced at \$74.99, includes everything — all 90 Elite channels, as well as the premium channels like Showtime, HBO, etc.

For the cost of your subscription, you also get the ability to use it on up to five devices (only in your home network), and you also have DVR access with unlimited storage for up to 30 days.

2.8. DIRECTV Now

DIRECTV Now is an all-online streaming service with a starting price of \$35 dollars. The introductory package includes over 60 channels, including FOX News, MSNBC, Comedy Central, ESPN, CNN, SyFy, and others, all being streamed live through the device of your choice. However, the price increases after the introductory period to \$60 per month. There's also the \$50 per month package that includes regional sports networks, as well as all the other channels from the first package. The \$60 package includes over 100 channels that include all the basics, but also MTV, Oxygen, FYI, the NHL network, and Sundance TV. And finally, the last option is \$70 per month, but includes everything from basic to premium channels.

2.9. Sling TV

Sling has gained some market share with its live TV service offering account for somewhere between 1-2% of live viewing in the US. There are 2 basic starting packages — Sling Orange starting at \$20 per month — with à la carte programming that you can add on at \$5 per month (\$10 for sport, \$15 for HBO Go) different genre channel line ups.

For \$25 per month — there is Sling Blue. This option includes regional sports and more. You can add à la carte items to the basic tier at \$5, \$10, or \$15 per month.

2.10. HBO NOW and HBO GO

HBO offers 2 services — NOW and Go.

HBO NOW is standalone streaming only service that you can buy through a subscription provider (Apple, Roku, Amazon, Android, Frontier, Google Fiber, Optimum, Liberty, Service Electric, Verizon, and Samsung). HBO Go is a streaming service — included free — provided you have a HBO subscription through your Cable or Satellite TV package. HBO Now is typically \$14.99 per month.

2.11. ARRIS and Espial Survey – Price Paid for Video Sources by Cord Shavers

2.11.1. Price Paid for OTT Sources by Cord Shavers




Figure 55 - Price Paid for OTT Sources By Cord Shavers

Our respondents showed that over **50% were paying over \$10 per month** with almost **20% paying over \$31 per month on OTT sources.** Only 10% were completely free.

2.11.2. Cord Shavers and Use of FTA Tuner

Not surprisingly 58% of Cord Shavers also used FTA tuner.




Figure 56 - Cord Shavers Use of FTA tuners

When asked if an FTA tuner integrated into OTT STB was of value — 50% said, “Yes” and 25% said, “No”, [they don’t use OTT STB] and 25% did not have an opinion.




Figure 57 - For Cord Shavers is it of Interest to Add the FTA Tuner in the OTT STB

2.11.3. Cord Shavers Preferred Device to Access Content

Cord Shavers preferred to access their OTT content on Smart TV (48%) or separate OTT STB (40%). Other devices accounted for 11% of the cord shavers.




Figure 58 - Cord Shavers - What Device Do You Use to Access Your OTT Sources

Of the 11% of “other” devices used — the respondents showed that Roku (which includes a separate STB OTT experience) was about 17% of the users and Computer/Laptop was the second most used device.




Figure 59 - Cord Shaver - Other Sources

2.11.4. Cord Shavers Time Spent Watching Live Versus Non Live Content

When asked what percentage of time respondents spent watching live content — 65% of respondents watch live content only 0-20% of the time.




Figure 60 - Cord Shavers Time Spent Watching Live Versus Time Shifted Content

3. Market Information on the Consumer Adoption of Pay TV and OTT Video Sources

Kleiner Perkins 2016 report showed the following statistics:

Pay TV household growth was down -1.3% on average for the last 12 Quarters..




Figure 61 - Kleiner Perkins - Pay TV Growth 2016

The number of channels watched was cited as < 10% of the channels received. Kleiner Perkins also showed the typical cost of Comcast, Dish, Charter, and DIRECTV service on a yearly basis below in Figure 62.




Figure 62 - Kleiner Perkins — Channels Watched and Cost of Pay TV Service — Yearly

4. User Experience Features Used — Correlation From STB Usage and UX Usage Data

To test the theories made in this paper — ARRIS and Espial set up a survey that had 1,316 responses with 180 uncompleted — a lot of which you have already seen in the analysis above. Fifty four percent of the respondents were from the US which gave us a large enough statistical sample — approximately 820 respondents to review the landscape in the US for Pay TV and OTT services.




Figure 63 - ARRIS and Espial Pay TV/OTT TV Survey — Size and Country of Reply

Within the US the survey there were 35 States with respondents with 36 respondents not stating their location. The highest participation was in PA and GA.

This data was used above to correlate against the published and public data that is also highlighted above in this paper and is detailed above in the analysis of Pay TV respondents and Cord Shaver respondents. The paper also added in data from the 2016 Kleiner Perkins — Mary Meeker report as another control.

The following sections take a different input direction further review the UX elements of the TV/video experience, Espial also did a monthly analysis of the usage behavior of their Elevate UX platform and the following sections detail the most used and commonly used features of a modern UX solution for Pay TV. Elevate UX experience also includes content and menu access via remote control, tablets, phones and laptop/Web — and this was also factored into the analysis. Elevate also has Netflix integrated into the primary Pay TV interface — already adding the convenience of a single device, single HDMI input for the most popular OTT source. The following sections go through this — as well as correlating some of the watching trends with the 2016 Nielsen report and additional charts from Kleiner Perkins 2016 surveys. The data was taken from ~140,000 households with the Espial Elevate solution.

4.1. The Features of a Modern UX that Consumers Use the Most

Using captured session information over a rolling month period on Espials Elevate user experience and tracking the actual sessions (session being the activation of a UX menu option including remote control, tablets, Laptops, and Phones) — the following was determined on what UX menu options users are using on a modern UX with integrated Netflix and other Applications. As you can see below in Figure 64 — the live TV menu had over 10x more session use on average than the next accessed menu — the Grid Guide menu. Users continue to browse the live TV content with the live TV menu giving fast channel look ahead.




Figure 64 - Frequency of UX Options Chosen on Pay TV Solution with Integrated Netflix

If we remove the live TV access sessions — Figure 65 — then we can see that the Grid Guide remains the most used single user interface solution on Pay TV system — however, the combined use of DVR menu options is higher. Setting DVR recordings is 10% more than DVR Playback usage.




Figure 65 - UX Options Most Frequently Used — Without Live TV Menu

“DVR Recording access” — to view what consumers have recorded- was the next most popular — DVR Recordings and DVR Playbacks and DVR Recorded shows. Combined add up to more access than the Grid Guide usage — almost 2x as much. Netflix sessions were on average 16 per month — however, they were longer in duration than many of the other menu access.

Mobile UX and use of Mobile device — was about 1/6th of the use of the Grid Guide — showing that consumers still like to use the TV or the TV remote more of the time than their tablet or phone devices.

4.2. What Consumers View Live

What people view most is also interesting — with more people still tuning in live to watch sports (expected), news (expected), and late night chat shows (function of browsing for that last thing to watch before bed). News and sports dominate the live viewing most sessions — with the odd movie (placed at the right time) and popular shows like America’s Got Talent. Figure 66 below shows the raw data for the top 100 watched live shows — but the Venn Diagram in Figure 66 below — shows that news accounted for 46% of all live sessions and sports 33% — news and sports accounted for almost 80% of live viewing of the top 100 watched shows. The 100th show has about 20% of the views of the top show — as shown in Figure 67.




Figure 66 - Live Shows Viewed by Genre




Figure 67 - Top 100 Live Shows and the 100+ Tail

The following Table 1 shows the raw data for the top 30 live TV Shows viewed and the household viewership and air date.

Table 1 - Top 30 Shows Viewed by Number of Households and Airtime

Ranking	Households	Title	Episode	Network	Air Date (GMT)
1	78646	2017 NBA Finals	Cleveland Cavaliers at Golden State Warriors	ABC Affiliate	6/13/2017
2	58103	2017 Stanley Cup Final	Pittsburgh Penguins at Nashville Predators	CBC Affiliate	6/12/2017
3	44413	Jimmy Kimmel Live		ABC Affiliate	6/13/2017
4	43174	Anderson Cooper 360		CNN HD	6/15/2017
5	42351	Anderson Cooper 360		CNN HD	6/14/2017
6	41249	Anderson Cooper 360		CNN HD	6/16/2017
7	40703	Anderson Cooper 360		CNN HD	6/13/2017
8	36901	Nightline		ABC Affiliate	6/13/2017
9	36800	CNN Tonight With Don Lemon		CNN HD	6/14/2017
10	36409	CNN Tonight With Don Lemon		CNN HD	6/15/2017
11	35578	Anderson Cooper 360		CNN HD	6/17/2017
12	34418	2017 Stanley Cup Final	Pittsburgh Penguins at Nashville Predators	NBC Affiliate	6/12/2017
13	33452	Sportsnet Central		Sportsnet West HD	6/12/2017
14	32426	2017 Stanley Cup Final	Pittsburgh Penguins at Nashville Predators	Sportsnet West HD	6/12/2017
15	32335	MLB Baseball	Toronto Blue Jays at Seattle Mariners	Sportsnet West HD	6/11/2017
16	30966	CNN Tonight With Don Lemon		CNN HD	6/16/2017
17	30324	2017 U.S. Open Golf Championship	Third Round	Fox Affiliate	6/17/2017
18	30131	Sportsnet Central		Sportsnet (Pacific) HD	6/12/2017
19	30128	CNN Tonight With Don Lemon		CNN HD	6/13/2017
20	28713	2017 Stanley Cup Final	Pittsburgh Penguins at Nashville Predators	Sportsnet (Pacific) HD	6/12/2017
21	27672	The 71st Annual Tony Awards		CBS Affiliate	6/12/2017

Ranking	Households	Title	Episode	Network	Air Date (GMT)
22	27241	America's Got Talent	Auditions 3	NBC Affiliate	6/14/2017
23	26327	Global News Morning		Global BC HD	6/14/2017
24	26327	Morning News		Global BC HD	6/14/2017
25	25925	American Ninja Warrior		NBC Affiliate	6/13/2017
26	25868	Global News Hour at 6		Global BC HD	6/16/2017
27	25868	News Hour		Global BC HD	6/16/2017
28	25823	Transformers: Age of Extinction		Showcase HD	6/17/2017
29	25804	Morning News		Global BC HD	6/16/2017
30	25804	Global News Morning		Global BC HD	6/16/2017

4.3. What Consumers Think They Want to Watch — Scheduled Recordings

It is interesting to review what type of shows people believe they want to watch time shifted or can't miss. The following Figure 68 shows the popularity of scheduled record by program genre. Sport is the highest with 25% of scheduled records set to record (mixture of backup to the live events and also non-live sports shows) and talent shows being number 2 with 19% of the scheduled records. Movies are only at 2% showing that consumers don't generally watch movies recorded from DVR as compared to drama series, variety shows, reality TV, or news.




Figure 68 - Top 100 Shows by Genre That are Scheduled to Record

You can see how the top 100 scheduled recordings tails off quickly on popularity — to 100th show is 9% of recordings scheduled than the most popular one.




Figure 69 - The Frequency of Scheduled Recording for the Top 100 Shows and 100+ Tail

Top 30 Shows scheduled to record in raw form — showing the number of households and air date — Table 2.

Table 2 - Top 30 Scheduled to Record Shows and Air Date

Ranking	Households	Title	Episode	Network	Air Date (GMT)
1	11124	Jillian and Justin	Baby on the Way	W (Women's Television Network)	6/22/2017
2	10886	America's Got Talent	Auditions 4	NBC Affiliate	6/21/2017
3	8709	The Gong Show	Will Arnett; Ken Jeong; Zach Galifianakis	ABC Affiliate	6/23/2017
4	7699	World of Dance	The Duels 1	NBC Affiliate	6/21/2017
5	7406	Grantchester on Masterpiece	4725	PBS Affiliate	6/19/2017
6	7360	2017 U.S. Open Golf Championship	Final Round	Fox Affiliate	6/18/2017
7	6648	NHL Awards		Sportsnet West HD	6/22/2017
8	6372	Little Big Shots: Forever Young	Forever Young	NBC Affiliate	6/22/2017
9	6230	Boy Band	Meet the Boys	ABC Affiliate	6/23/2017
10	5914	The Bachelorette	1304	ABC Affiliate	6/20/2017
11	5431	The Night Shift	Recoil	NBC Affiliate	6/23/2017
12	5285	20/20		ABC Affiliate	6/24/2017
13	4838	Dateline NBC	As Night Fell	NBC Affiliate	6/24/2017
14	4807	Jillian and Justin	Baby on the Way	W (Women's Television Network) (Pacific)	6/22/2017
15	4465	The Story of China	Ancestors; Silk Roads and China Ships	PBS Affiliate	6/21/2017
16	4311	2017 U.S. Open Golf Championship	Final Round	TSN HD	6/18/2017
17	4241	Great Canadian Homes		Home and Garden HD	6/19/2017
18	3949	My Favorite Wedding		Hallmark Channel HD	6/25/2017
19	3829	NHL Awards		Sportsnet (Pacific) HD	6/22/2017
20	3662	2017 NHL Draft	Round 1	Sportsnet (Pacific) HD	6/23/2017
21	3565	2017 NHL Draft	Round 1	Sportsnet West HD	6/23/2017
22	3044	So You Think You Can Dance	Los Angeles Auditions No. 2	Fox Affiliate	6/20/2017

Ranking	Households	Title	Episode	Network	Air Date (GMT)
23	2936	Alone	Hell on Earth	History Television HD	6/23/2017
24	2920	MLB Baseball	Toronto Blue Jays at Texas Rangers	Rogers Sportsnet One HD	6/22/2017
25	2896	My Mother and Other Strangers on Masterpiece	4720	PBS Affiliate	6/19/2017
26	2717	Prime Suspect: Tennison on Masterpiece	4732	PBS Affiliate	6/26/2017
27	2678	MLB Baseball	Toronto Blue Jays at Kansas City Royals	Rogers Sportsnet One HD	6/24/2017
28	2591	Kids Who Kill		A and E HD	6/20/2017

4.4. What Consumers Actually Record — Scheduled and By Viewing Guide or While Watching Content

While there are a proportion of consumers that are using the DVR scheduling to record — there is 75% more recording done for a single show versus the scheduled recording. This more ‘on the spot’ or ‘closer to the event’ recording also during the show itself. This also alters the top show recorded, as well as the genres quite significantly.

If you recall — sports was the highest using the DVR Scheduling capabilities — however — drama series (Game of Thrones, Fear the Walking Dead) and reality TV shows as well as late night chat shows all show a tendency for people to press the record button in the guide or while watching the show. This is done 3x more than scheduled recording — and shows the changes in mood and the on the spot decisions to watch and record things — based on live program schedule. Again, recording movies from live airing is very low.




Figure 70 - The Show Types Recorded at the Time of Airing

The actual recording trend skews more to the top 10 programs — with the long tail at 100th show being 7% of the recordings of the top show. The first 20 shows accounting for 50% of the recordings.




Figure 71 - The Actual Recorded Shows in Popularity and the 100+ Long Tail

Top 30 shows recorded before the air time of the show — including scheduled, select recording before and while viewing the show — Table 3.

Table 3 - The Top 30 Shows Recorded By Household and Airtime

Ranking	Households	Title	Episode	Network	Air Date (GMT)
1	62816	America's Got Talent	Auditions 4	NBC Affiliate	6/21/2017
2	53823	The Bachelorette	1304	ABC Affiliate	6/20/2017
3	30810	World of Dance	The Duels 1	NBC Affiliate	6/21/2017
4	29991	Alone	Hell on Earth	History Television HD	6/23/2017
5	25846	Fear the Walking Dead	100	AMC HD	6/19/2017
6	24526	Dateline NBC	As Night Fell	NBC Affiliate	6/24/2017
7	24389	Better Call Saul	Lantern	AMC HD	6/20/2017
8	23665	MasterChef	Feeding the Lifeguards	Fox Affiliate	6/22/2017
9	22870	Genius	Einstein: Chapters Nine & Ten	National Geographic HD	6/21/2017
10	22318	American Ninja Warrior	San Antonio Qualifiers	NBC Affiliate	6/20/2017

Ranking	Households	Title	Episode	Network	Air Date (GMT)
11	22296	The Handmaid's Tale	The Bridge	Bravo HD	6/19/2017
12	22160	Genius		National Geographic HD	6/21/2017
13	21474	So You Think You Can Dance	Los Angeles Auditions No. 2	Fox Affiliate	6/20/2017
14	20737	Genius		National Geographic HD	6/22/2017
15	19181	Home to Win	Keys to the Competition	Home and Garden HD	6/19/2017
16	17055	Fargo	Somebody to Love	FX HD	6/22/2017
17	16283	Deadliest Catch		The Discovery Channel HD	6/21/2017
18	16194	Nashville	A Change Would Do You Good	W (Women's Television Network)	6/23/2017
19	15730	This Is Us	The Pool	NBC Affiliate	6/22/2017
20	15719	Better Call Saul	Lantern	AMC HD	6/20/2017
21	15613	Fear the Walking Dead	100	AMC HD	6/19/2017
22	14130	Jillian and Justin	Baby on the Way	W (Women's Television Network) -	6/22/2017
23	14076	The Night Shift	Recoil	NBC Affiliate	6/23/2017
24	12658	20/20		ABC Affiliate	6/24/2017
25	11567	Little People, Big World	When It Rains It Pours	The Learning Channel HD	6/21/2017
26	11121	Anthony Bourdain: Parts Unknown	Trinidad	CNN HD	6/19/2017
27	10569	Preacher	Season 2 Greetings From the Set	AMC HD	6/23/2017
28	10521	Nashville	A Change Would Do You Good	Country Music Television HD	6/23/2017
29	10344	The Tonight Show Starring Jimmy Fallon	Will Ferrell; Alison Brie; Shawn Mendes	NBC Affiliate	6/20/2017
30	10322	The Tonight Show Starring Jimmy Fallon	Amy Poehler; Zendaya; Imagine Dragons; Dweezil Zappa	NBC Affiliate	6/21/2017

4.5. Top 30 Shows Actually Played Back on DVR Playback

Yet another way to look at DVR usage is to look at the shows that were actually looked at from the recordings during the same month. This shows that genres like Talent shows tend to be forgotten whereas consumers do make sure more to catch up on drama series and their favorite Reality TV shows. These 2 genres account for 60% of the top 100 shows watched back on DVR.




Figure 72 - The Shows Watched in Playback on DVR

The top 20 shows account for 48% of the DVR views in the top 100 and the tail at 6.8% of the frequency of the most popular first place playback show.




Figure 73 - DVR Playback Frequency and Long Tail of 100+ Top Shows

Table 4 - Top 30 DVR Playbacks by Household and Airtime

Ranking	Households	Title	Episode	Network	Air Date (GMT)
1	41228	America's Got Talent	Auditions 4	NBC Affiliate	6/21/2017
2	36371	The Bachelorette	1304	ABC Affiliate	6/20/2017
3	19762	World of Dance	The Duels 1	NBC Affiliate	6/21/2017
4	15130	Alone	Hell on Earth	History Television HD	6/23/2017
5	13789	Better Call Saul	Lantern	AMC HD	6/20/2017
6	13183	Fear the Walking Dead	100	AMC HD	6/19/2017
7	12636	So You Think You Can Dance	Los Angeles Auditions No. 2	Fox Affiliate	6/20/2017
8	11916	MasterChef	Feeding the Lifeguards	Fox Affiliate	6/22/2017
9	11329	The Handmaid's Tale	The Bridge	Bravo HD	6/19/2017
10	10496	Home to Win	Keys to the Competition	Home and Garden HD	6/19/2017
11	10408	Jillian and Justin	Baby on the Way	W (Women's Television Network)	6/22/2017
12	10182	Dateline NBC		NBC Affiliate	6/17/2017
13	9660	Deadliest Catch		The Discovery Channel HD	6/21/2017
14	9622	American Ninja Warrior	San Antonio Qualifiers	NBC Affiliate	6/20/2017
15	8741	Little People, Big World	When It Rains It Pours	The Learning Channel HD	6/21/2017
16	8692	Better Call Saul	Lantern	AMC HD	6/20/2017
17	8185	Alone	Divide and Conquer	History Television HD	6/16/2017
18	7630	Genius	Einstein: Chapters Nine & Ten	National Geographic HD	6/21/2017
19	7598	Last Week Tonight With John Oliver	105	HBO West HD	6/19/2017
20	7409	The Real Housewives of New York City	912	SLICE HD	6/22/2017
21	7308	Fear the Walking Dead	100	AMC HD	6/19/2017
22	6990	Counting On	Spurgeon's First Birthday	The Learning Channel HD	6/20/2017

Ranking	Households	Title	Episode	Network	Air Date (GMT)
23	6926	Fargo	Somebody to Love	FX HD	6/22/2017
24	6407	Grantchester on Masterpiece	4725	PBS Affiliate	6/19/2017
25	6403	The Night Shift	Recoil	NBC Affiliate	6/23/2017
26	5868	America's Got Talent	Auditions 3	NBC Affiliate	6/14/2017
27	5836	90 Day Fiance	Jorge & Anfisa: Our Journey So Far	The Learning Channel HD	6/19/2017
28	5698	Nashville	A Change Would Do You Good	W (Women's Television Network)	6/23/2017
29	5605	Love It or List It Vancouver	Laila & Dan	Home and Garden HD	6/20/2017
30	5595	Hollywood Medium With Tyler Henry	Khloe Kardashian and Kylie Jenner, Ru Paul, Elisha Cuthbert	E! Entertainment Television	6/19/2017

Finally — to illustrate the changes in viewership across Live, Scheduled Recording, Recorded, and Playback content — the following Figure 74 puts all 4 categories together to see the difference at each different time in the viewing cycle. Some simple analysis reveals:

- Drama Series are watched more on DVR than live significantly — 45% to 1%
- Sport is watched live and recorded by rarely watched again
- Talent shows are scheduled to record but not watched often after recording
- Reality TV is not watched live by majority of people but is watched consistently on DVR playback
- News is watched live and rarely recorded or watched on DVR
- Movies again don't tend to be watched live much or while scheduled don't tend to be watched back on DVR playback




Figure 74 - Comparison of What Consumers Watch By Genre — Most Frequently

Kleiner Perkins — 2016 Mary Meeker report showed the following change in consumption of the top 5 networks — comparison from 2010-2011 to 2015-2016.

Network TV* Minutes Delivered = 2011 Top 5 Networks -10% Average...
Netflix +669% Over 5 Years, USA...




Figure 75 - Kleiner Perkins 2016 report — Monthly Minutes Delivered by Top 5 Networks — 2010-2011 compared to 2015-2016

From the internal survey data, the use of the Elevate Portal to see what consumers are using as part of video user experience and the external public sources of estimating what consumers are doing with Pay TV and new ways to embrace their Video entertainment services in the home — the following section tries to suggest to the cable operator — how to retain subscribers, potentially grow their subscriber base — but at least create the best possible experience to shift the focus from cost — to user experience as part of the home TV experience.

5. A Potential Path Forward for the Cable Operator to Own the TV Experience and Create a New Value Proposition

From the data analyzed to this point the following items stand out to form the strategy for the path forward for the cable operator and the next step in Video, Application, and TV experiences:

1. The value proposition for video sources is broken for Cord Shavers especially — and increasingly for existing Pay TV customers. They are actively looking at reducing the cost of Video and Pay TV services.
2. Both Cord Shavers and current Pay TV subscribers see the path to cost reduction via À la carte and skinny bundle solutions. As they typically favor 15-17 channels that they watch frequently they wonder why they need all the other channels and sources.

3. Consumers see the lease cost of the STB device as part of the overall cost problem. The awareness of the lease cost of STB is not as high as just the overall general cost issue. There is, however, a larger awareness from customers with multiple leased STB and inertia for consumers to add additional outlets to their service from the MSO. Additional outlets at no lease cost increases the value — in the cost value equation — so something to consider going forward and discussed later.
4. Customer care also registered highly in the comments from respondents in the ARRIS and Espial survey. Paying a lot of money and feeling that you are not getting good service (reliability, customer care, technician support, with a number of issues per year) is a relevant factor in the ‘Value to Cost’ equation and something that also needs to improve. This is discussed further in the section. For OTT non-traditional Pay TV Cord Shavers — their customer care issues were less concerning based on the cost spent per month on sources. The more money spent per month on OTT sources, the more likely the Cord Shaver was to be concerned with quality issues and the worry about calling poor support lines.
5. There is a large portion of Cord Shavers who really don’t watch much TV at all and value more time for activities and learning. These consumers may never be converted to traditional Pay TV bundles — but can still be targeted by cable operator creating targeted lifestyle bundles.

The first 5 items are the key items to try and address that are clearly at the forefront of a lot of paying Pay TV consumers — however, the following items also suggest that there is a formula and path forward to continue to retain and potentially grow Pay TV subscribers:

6. The user experience (a really good one) does work for distracting the consumer from the overall cost of video services. Without the modern user experience elements — it gets much tougher to retain customer when cost of their service is an issue. The satisfaction level of the user experience is only 69% overall from the survey conducted so additional work there can offset customers wanting to move.
7. Integrating the popular OTT sources is a must in retention of subscribers. This was clearly shown in the respondent’s answers and comments.
8. Having a single HDMI port for all services to aggregate to the TV is desirable on a number of levels for the cable operator in particular with this control of the consumer to the cable operators’ remote control being key to strategy going forward. The respondents to the survey were not as emphatic on the single HDMI source (many did not answer so could be converted) — but the use of a single remote control, the ergonomic improvement of only one device connected to the TV and future improvements to both the STB and remote control — will add value to the cost proposition of the consumer.
9. There are only 4 or 5 other non-traditional video sources that really count in the USA. These are Netflix, YouTube, Amazon, Hulu, and Sling TV. Integrating these solutions covers most of the desires for the Pay TV user for convenience and also potentially can lure back the Cord Shavers with skinny bundle packages.
10. Live vs. Time Shifted — while the press and analysts want to show that everything is going time shifted — the data recorded shows that Pay TV subscribers still do the majority of their viewing live — with sports and news dominating their live viewing behavior. From the Elevate analysis it was shown that live interaction with the TV was 10x higher than other inputs. This is key because the consumer see’s the live line up as an information portal to ‘what’s going on’ and typically will always start out a TV session with a ‘What’s on Live’ request to the STB.

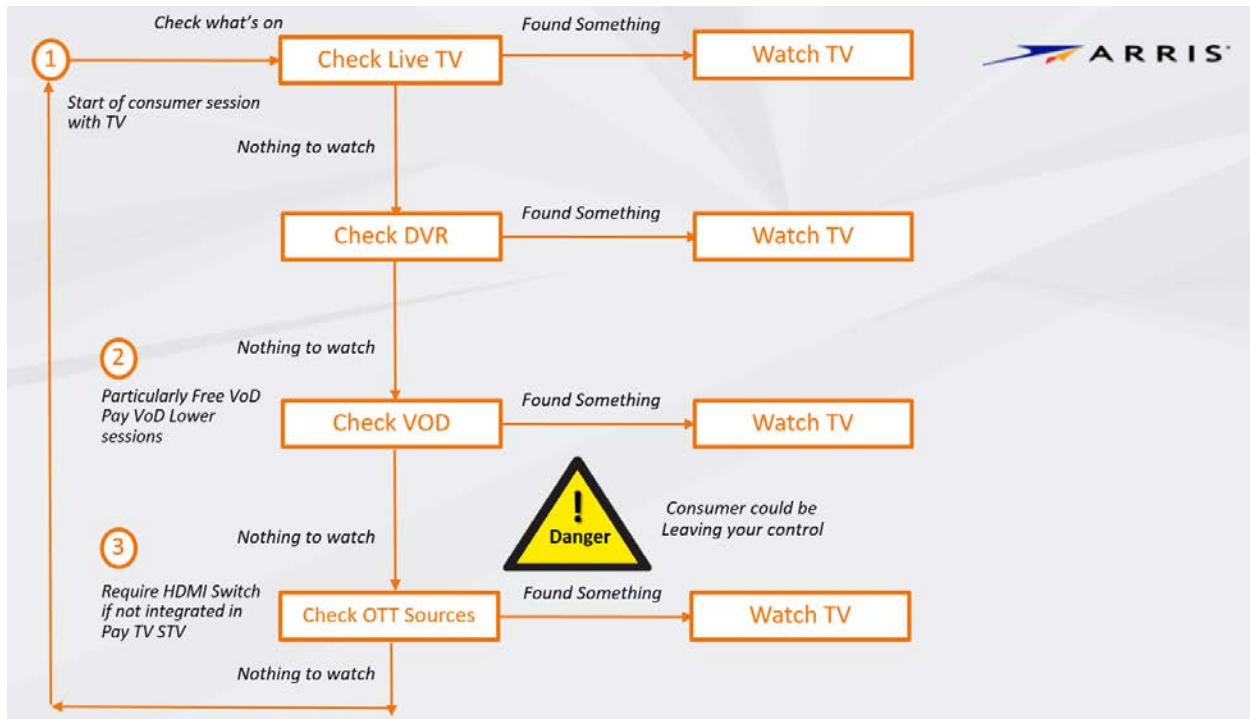


Figure 76 - Most Common Sequence of Engagement with TV on Pay TV system

Figure 76 above illustrates what is the most frequent and important sequence in the consumer engagement with the video experience. As we saw from the analysis of the consumers on Elevate platform — their ‘go to point’ is always to check live TV first. They then drift to DVR when there is nothing to view and then split their time between VoD and OTT Sources integrated — Netflix for example. Consumers also value Free VoD and this retains them more on the VoD system. Consumers will drift to Netflix more for movie’s as we have seen in the data. The key thing however is that because the most popular OTT source is integrated on the same HDMI port — the consumer comes back to the MSO after the OTT session has finished and they typically browse live TV again.

11. The way OTT sources are integrated into the overall UX is also something to consider. The best solutions are when the OTT sources are integrated directly into the guide as well as part of the search and recommendation engines. This requires having access to the metadata for the OTT sources — and has been shown to be a win-win for the MSO, the OTT provider and most importantly the consumer experience. This enhances the value proposition of the Pay TV service. As we saw in the respondent’s survey on cost of Pay TV service — we saw that 14% of Cord Shavers had bought all 4 services — Netflix, Amazon, Hulu, and used YouTube with most of the users having 3 (with YouTube) OTT sources. The average cost of OTT for Cord Shavers is between \$20-\$30 — but if you extract out the customers who do not value TV as much it is above \$30 in monthly payments on average. The challenge then is to get Basic Live TV package plus similar experience or integration of these sources into the Pay TV package to bring back the Cord shavers. With the improved UX elements — the Pay TV provider can charge a higher aggregated cost — how elastic this is — is a function of providing some of the other experiential items discussed below.

12. While respondents were a little damp on the importance of multiscreen services — it was more obvious when correlating their answers with their use of phone and tablet and non-TV viewing devices. In particular, for the OTT Cord Shavers — multiscreen was of higher value — as they strove to get content on other non-TV devices.
13. Last but, not least — is the importance of the remote control and the input device/media to the selection of service on the TV. Whoever controls the remote device — controls the TV and more importantly the content on the TV. Consumers don't want to have to use multiple remotes — and will gravitate to the one that provides access to all their services. This is why good MSO remote solutions try to cover the following:
 - a. Allow the remote to perform the same functions as the Smart TV remote — via a programmed mode.
 - b. Replacing Infrared Radio Frequency solutions with RF4CE and Bluetooth Low energy (BLE) to get to a Non Line of Sight (NLOS) capability for remote control usage and to allow other services like audio streaming and voice streaming from/to the remote control.
 - c. Addition of near field microphones to allow Voice input in a very cost effective manner (Comcast XR11 in Figure 77 is an example).
 - d. The addition of hot key buttons for most prominent or promotional OTT sources (see Roku Remote - with paid for quick access buttons for Netflix, Amazon Video, Sling TV and others like Google Play, Hulu, and RDIO).



Figure 77 - Comcast Voice Capable Remote Control and Roku OTT Services Buttons Remote Control

So, the fundamental questions are:

- i. Is the pursuit of Video services to the subscriber still important for direct ARPU contribution or continued control of the largest screen in the home? **The authors think the answer is, “Yes.”**
- ii. The STB — does it create a continued control point that goes beyond just the plain Video decoder and Graphics generator? **The authors think the answer is, “Yes.”** How important are these other OTT and other TV experiences to the overall strategy and future of the MSO? Figure 78. **The authors think the integration of OTT Video services and the future integration of Home Lifestyle services are the key to customer satisfaction improvements.**

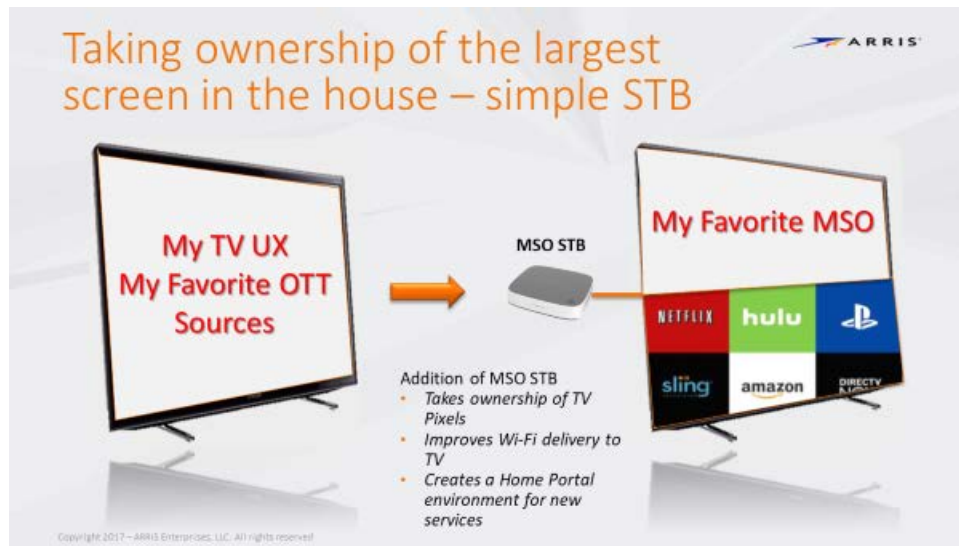


Figure 78 - Addition of MSO STB to Smart TV Keeps Control With the MSO for All Screen Services

- iii. Can the cost/value consumer dis-satisfaction be circumvented by the following 5-point plan as illustrated in Figure 79 below:
1. Hardware - Best Wi-Fi and STB HDR decode experience. Additional support for RF4CE and BLE for more IoT and presence driven capabilities and features to the TV. Future considerations for Voice input as well.
 2. Great UX experience — graphics and application features — with Search and Recommendations across broadest range of sources. The user experience should be ‘Full Fat’ — meaning that there is a correlation between user satisfaction and adding in more and more services to the experience. These have to be relevant, easy to use.
 3. Multiscreen support on and off network.
 4. Integration of all the most popular (and relevant) video sources into one aggregated and one HDMI Port environment.
 5. The extension of the STB control point to bring in new Home Digital Lifestyle and Experiential services around the largest screens in the home.




Figure 79 - The 5 Step Plan to Keep Consumer Demand for Pay TV services

Figure 80 tracks the Epoch evolution in devices and technologies to line up with the video to new service evolution of the home. Central to this is the relationship between the Cloud->GW->STB to deliver these curated Digital Life experiences.

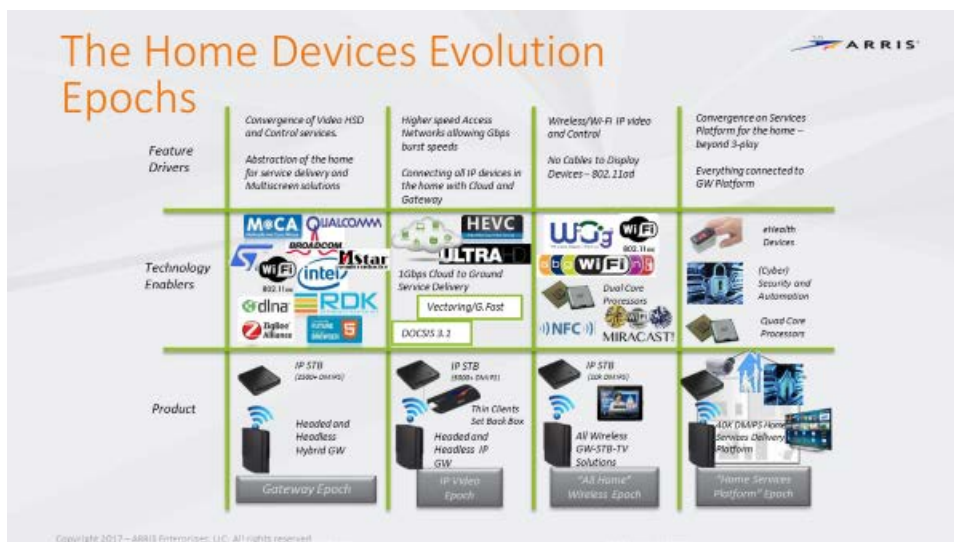


Figure 80 - The Evolution of Home Devices Tracking Evolution of Service Epochs

5.1. The TV as the Home Portal and the Experiences that it Can Drive and Support — All from the Addition of an MSO STB Device

The set-top box is changing and incorporating the changes and capabilities indicated below in Figure 81. The primary function remains to decode video streams — High Efficiency Video Coding (HEVC)->4K->8K and High-dynamic-range (HDR)evolutions — however, it is also now a key demarcation point for the MSO to:

- Improve Wi-Fi — either adding excellent client Wi-Fi at 2x4 or 4x4 or allowing the STB to act as a Wi-Fi extender
- Move to SSD Storage for caching from Cloud DVR or cache points. STB can be extension of CDN
- Change the design to better ergonomics of size and integration with TV or screen
- Higher Graphics processing unit (GPU) capabilities for better Graphics performance. This drives a rich and snappy user Interface including the growing importance of finding content fast for non-Binge viewing. Consider even gaming experiences for casual gamers.
- Add BLE/RF4CE and Wi-Fi for not only Remote Control connection, but also Internet of Things (IoT) Hub capability, presence applications and streaming audio.
- Add Smart Functions — including Voice Input — Smart Assistant — and also presence based applications driven by BLE abilities of the set-top box.
- Add ability to run Home Lifestyle Applications in conjunction with the cloud
 - Home Management Apps — using TV portal as one of the output screens
 - Home Office Apps — using TV portal as one of the output screens
 - Home Health Apps — using TV portal as one of the output screens
 - Home Education Apps — using TV portal as one of the output screens
 - Home Security Apps — using TV portal as one of the output screens
 - Home Self Help Apps using TV portal as one of the output screens

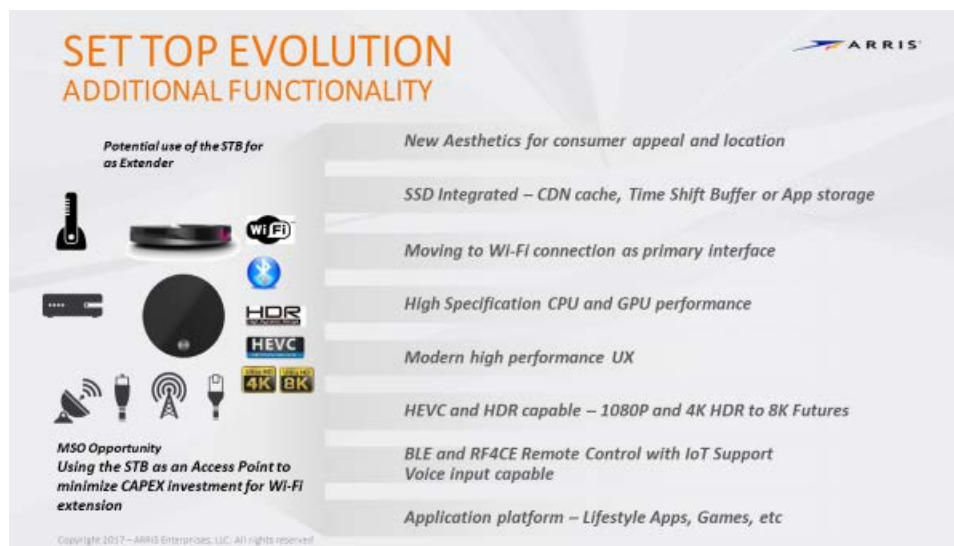


Figure 81 - The Changing Features and Shapes of the STB

The current directions for the STB is illustrated below in Figure 82 — where the industrial design is going more towards a slate/coaster design — and trying to keep the Wi-Fi performance as good as possible for behind TV placement. This drives designs at 76-130mm in size — and the ability to get as thin or even thinner than the size of RJ45 connector. The removal of Ethernet to go to an all Wi-Fi drive with a thinner form factor.

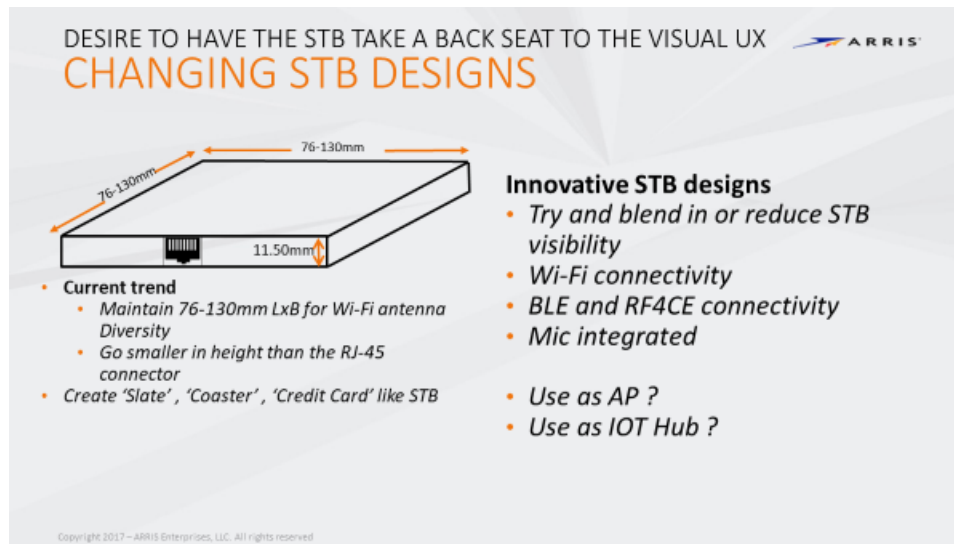


Figure 82 - STB Getting Small and Thinner — Yet Trying to Maintain Best Wi-Fi Performance and Thermals

There is one new interesting increment in the evolution of the STB that we are seeing today. Figure 83 shows LG's advancement in the TV aesthetic and architecture — with the first product leverage of flexible OLED displays. You will see below the separation of the screen itself from the media processing device. Samsung and other TV manufacturers have also started to separate out the media processing from the screen itself. As can be seen the screen function now is purely for display purposes — and the separate Media Box does the following features:

- STB functionality — Decode, Central processing unit (CPU), and GPU
- Wi-Fi Client functionality
- Audio — the LG W7 shown below is an elegant sound bar as well
- BLE/IR
- HDMI and other inputs
- Providing power and control to the screen matrix

It does raise an interesting question whether the screen to Media Box interface should now be standardized like the HDMI inputs were previous standardized on more traditional TV devices something for our industry to consider.



Figure 83 - LG W7 Flexible OLED TV System with Soundbar Media Box

If we peek out a little further from a solution like the LG W7 we can see the natural evolution of mass produced flexible displays to the directions illustrated in Figures 84-85 below. This is the key direction that the MSO needs to track and look differently at the STB and its ability to control these screen portals. There will be primarily one of these all wall devices in the primary living area of the home and it will be almost like the ‘command center’ of the home with the screen offering the ability to display in segments for different services and in full screen mode for entertainment.

Therefore, it is important now to start re-positioning the STB as the anchor point for being able to perform these new services in the future. These additional services — experiential Digital Lifestyle services will.

The next sections will focus on the user experience and what the STB can put on the TV screen — from much better video experience to other Digital Lifestyle services.

5.2. The STB Generated Screen User Experience

It was not long ago we had TV UX like the one displayed below in Figure 84. It’s clear that our surveyed respondents did not favor these guys where they still had them for TV service.

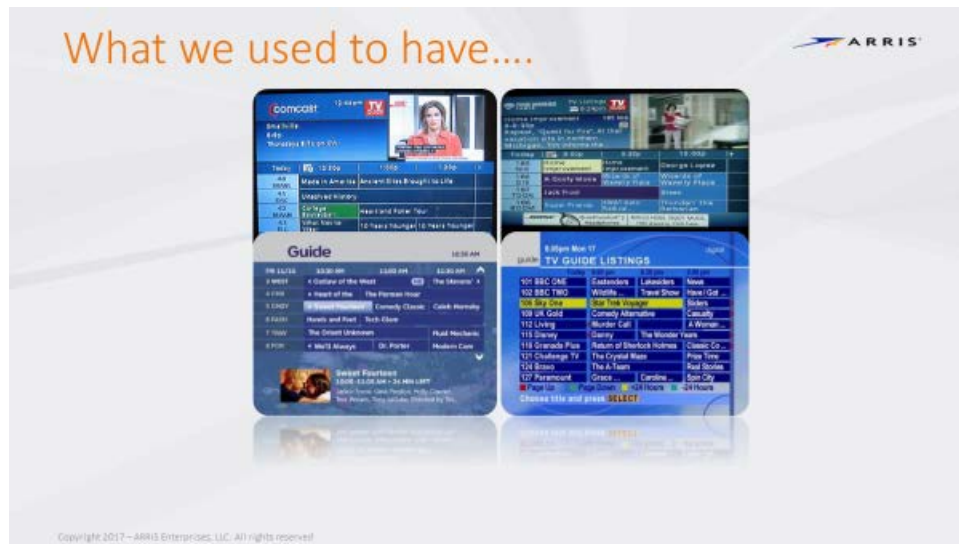


Figure 84 - Old Grid Guide TV UX Solutions

A smooth and focused UX with compelling graphics that make the UX experience quick and also beautiful and efficient — do make a difference. An example of a concept is shown below in Figure 85. Notice the central program and the way related content — either in program additional segments or related content is shown in a new way — taking full advantage of the STB GPU capability and the ability to render graphics to 4K levels.



Figure 85 - Compelling New UX Experience With Compelling Graphic Concepts

Additionally, there can be new paradigms in how to do advertising with a more integrated and less obtrusive experience for advertising — including allowing program sponsorship in new in-line user friendly ways. Figure 86 below is a concept example of this.



Figure 86 - New Ways to Advertise With Compelling Graphics as Part of Overall UX

In the context of user experience (UX) let us consider three aspects the MSO should focus on:

- The screens used by the consumers to access content
- The approaches to integrating a broader range of content
- The ability to change the service rapidly to react to market circumstances

The follow subsections will address each of these in turn.

5.2.1. UX Screens Enabling the Consumer to Access Content

When designing a UX for consumer access to content several key principles need to be considered:

1. The UX should be easy to learn / easy to use
 - a. The user is in lean-back mode, make it very clear what user can do
 - b. Avoid clutter, refine every aspect of design down to minimal information
2. Ensure the user has a consistent experience
 - a. Make actions predictable, re-use similar patterns and behaviors
 - b. If the user can predict what is going to happen, they learn to trust the experience
3. Design for flexibility — changes will occur — plan for these up front
4. Design for usability, strong focus on the users and their user journeys
 - a. Focus on helping the user find and watch content as rapidly as possible
5. Assume users will access content across a range of devices — plan for this
 - a. 10ft experience vs. tablet vs. phone
 - b. Use as a companion device
 - c. iOS and Android expectations (e.g. long touch vs. 3D touch)
 - d. Consider device remote controls — which often have very few buttons e.g. (Figure 87)
 - i. Roku, D-Pad, Back, Options, Play, Pause, FFWD, and REW
 - ii. Apple TV, D-Pad, Back, Search, Play, Pause, Volume



Figure 87 - Variance in Device Remote Control Options

What this means is that the UX design is consistent across multiple devices and is flexible to change.

This tends to drive a clean and simple design like the UX from the Espial Elevate service.

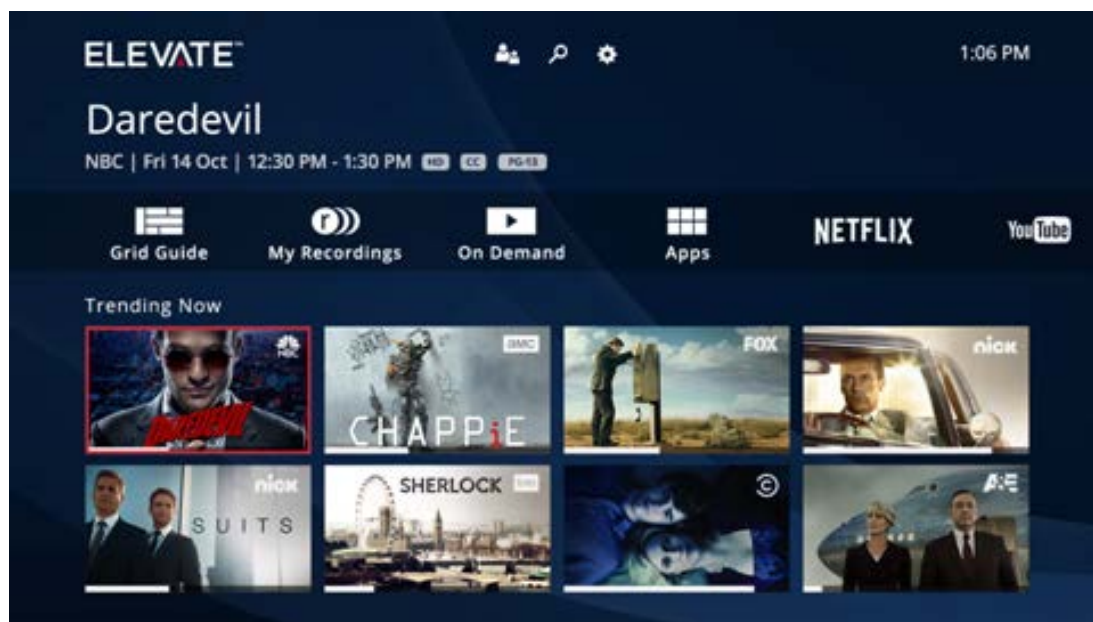


Figure 88 - Example of a Modern UX Design

Some of the key elements of this design are:

- Easy to learn / easy to use / familiar on all screens and across all devices
- Ability to port easily across multiple devices and handle multiple input devices
- Focus on content first — poster centric, recommendations, broad range of content including OTT content.

In addition to this, today's rich metadata sources permit users to get multiple new ways of finding content.



Figure 89 - Offer the User Different Ways to Find Content

In this example, cast and crew information allows the user to easily find information about the cast and then to find additional content with the same actors. This example also shows recommendations where the user is always offered additional relevant options with the objective of ensuring the search for interesting content is successful.

5.3. UX Screens Enabling a Broader Range of Content

In the previous sections, we introduced the concept of integrating OTT content to offer and presenting this to the consumer. Often this is initially performed using an “application” model — meaning the user can access the application through the screen. The benefit of this is that the user can access the content through HDMI 1. The challenge is that finding content across live, on-demand, DVR, and applications is difficult and means the user may have to open and close several applications before finding what they are looking for. Increasingly operators are seeking to offer a much easier access to content for consumers by negotiating access to OTT metadata which supports converged search.

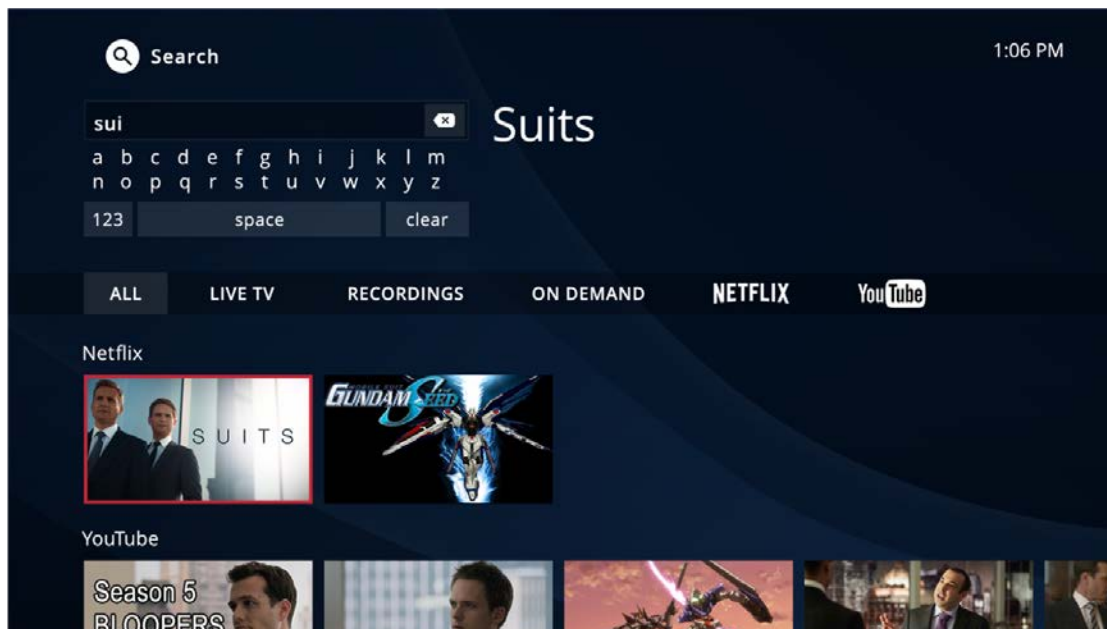


Figure 90 - User Interface with Integrated Search

In this model, the user can much more easily find the content they want across multiple sources and directly access it. This offers the consumer much more value from the service.

5.4. Ability to Change the UX Rapidly

The UX design has been constructed so that it can be rapidly changes from the cloud/back-office. Examples of these changes are in Figures 91 and 92 below.

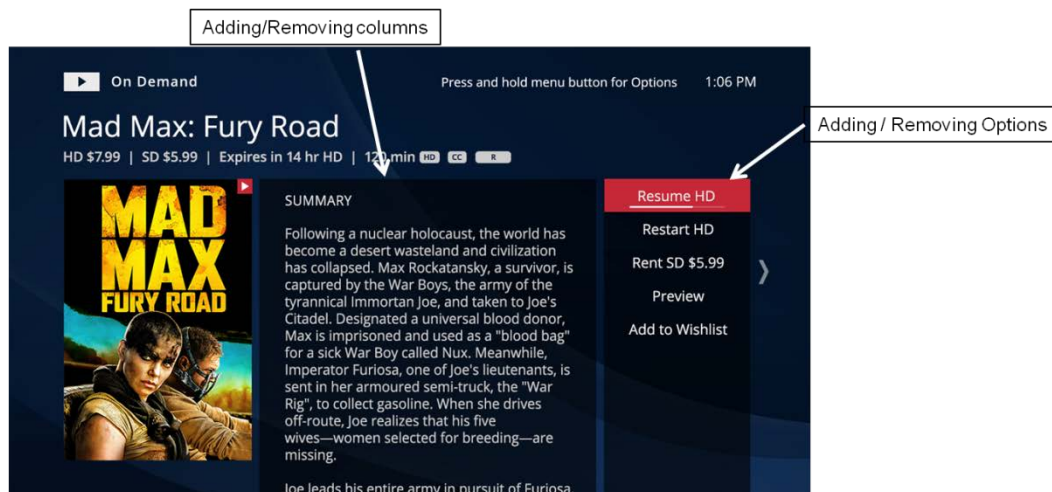


Figure 91 - Cloud / Back-Office Driven Back Office Changes

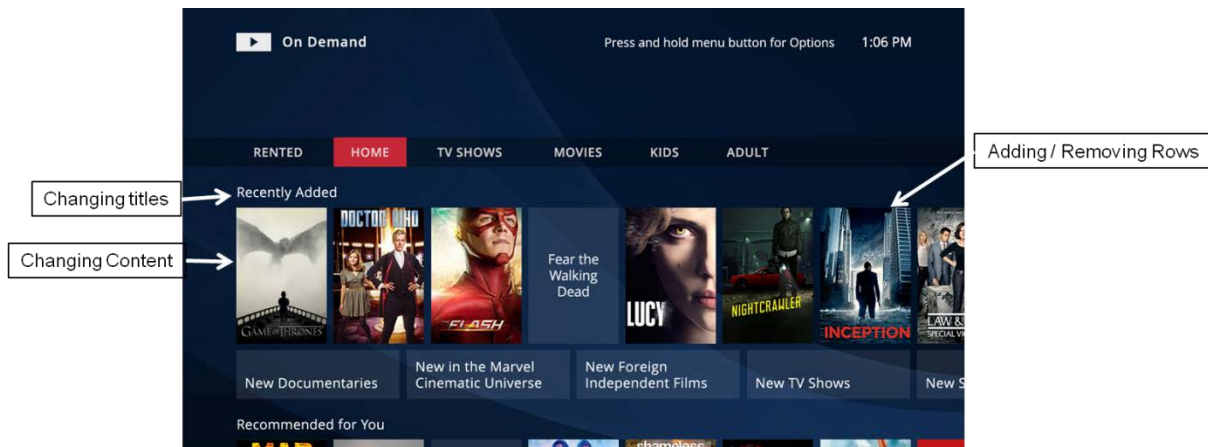


Figure 92 - Cloud/Back-Office Driven UX Changes

Of course the UX is also changing with the size of the TV screen and the real estate you have to play with — we already have 100”+ Flat screen TV’s and as mentioned above we are moving towards larger Wall screens.



Figure 93 - 100" LED Displays to Future all Wall Flexible Displays

It is this flexible screen that will enable multiple view portals and will require a rethink in how to drive them from bandwidth, clean lines, and the media/set-top box to drive them.



Figure 94 - Current - Primary TV use of Flexible OLED — Future — Wall OLED Portal With Split to Multiple Virtual Screens

Today — the mix is something like that illustrated in Figure 95 — where we have multiple generations of screen/TV/STB combinations for different rooms in the home. We have a mixture of HD and growing numbers 4K devices. This combination of cycling TVs to different rooms and the purchase of new TV/screens and STBs will continue to evolve to the first full Wall TV additions.



Figure 95 - Different Rooms, Screens/TVs, and STB

This provides new services and ARPU opportunities for the MSO and some of the ideas and concepts are illustrated below.



Figure 96 - There is an Opportunity to Generate New Experiences and Revenue Opportunities Across All Screens

New area of Home Digital Lifestyle can be mapped to the screens in the home particularly the largest screens with the addition of MSO STB.



Figure 97 - The STB Converging New Services from Video to Health Energy Management Wealth

Figure 98 below shows that the STB is the aggregator of experience and quality delivery for many services, not just video consumption. Allowing (when appropriate for lean back experience) services like commerce (select a dress over your avatar), watch a concert live or in Virtual Reality (VR), select parallel feeds to view the house that you saw advertised, etc.



Figure 98 - STB Providing All New Experiences on the Screens of the Home

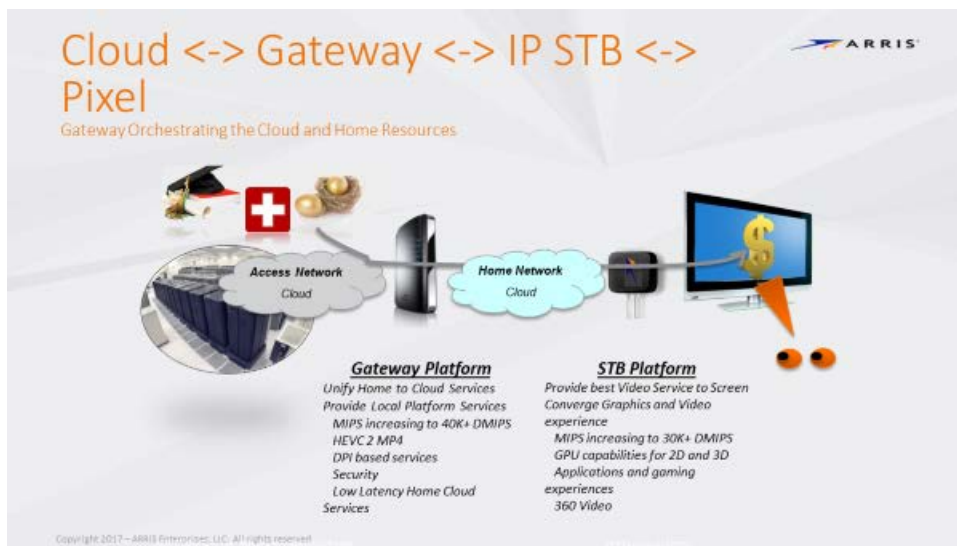


Figure 99 - The STB and GW Enable a Whole New Set of Services to be Leveraged to the Largest Screens in the Home

The synchronization between the 1:1 devices (phone and tablet) and M:1 of the large TV — is also an opportunity for leverage and favor the provider of the STB. Figure 100 below shows the use of tablet/smart phone to provide offers for the big screen as well as notifications across both devices. We have already seen how video casting selections from the small screen device to the large screen device is another part of the convergence of screen usage in the home.



Figure 100 - Use the Small Screen to Send Recommendations and Offers to Show on Large TV Screen

New additions to recommendations and search can be implemented to allow consumers to find not only full show content, but also specific items within a show or across multiple shows. Not only the metadata, but also the actual audio and video content can be digested and added to improved navigation and viewing experience.

The example below shows a framework that has been put in place to allow the ability to index video/audio content to the point where you can ask for only content in AMC walking dead series that contains zombie scenes! This is a new powerful way to correlate content sources and even to leverage social media trending items.

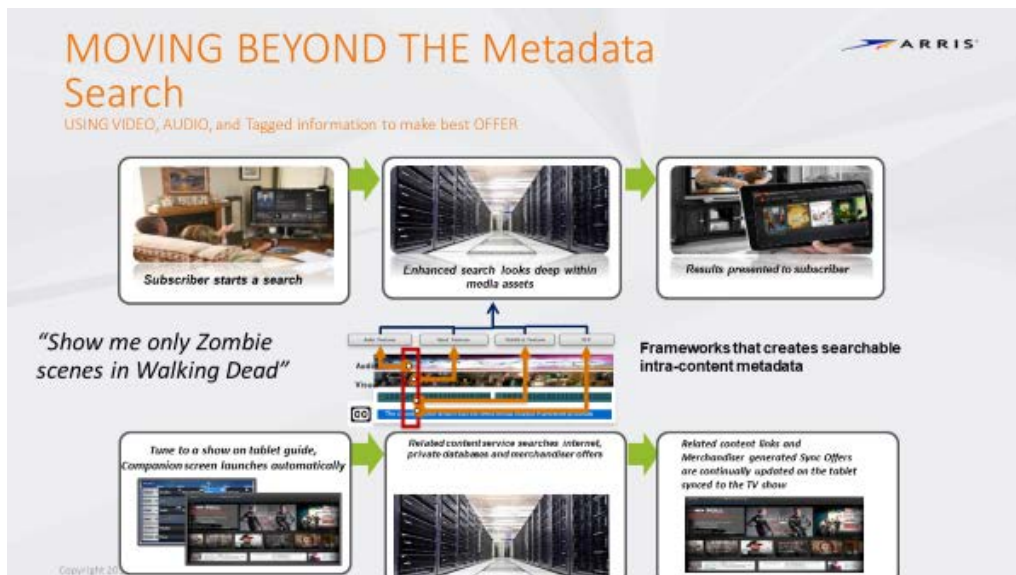


Figure 101 - Media Analysis Framework — Allowing Very Granular Search on Words, Video Scenes and Other Tags

As we showed above, even advertising can be enhanced to be a positive source of information for the consumer vs. something that people actively seek solutions to skip. It is clear that the correlation between small screen (personal) and TV screen (communal) is a key part of new advertising directions that users can see as beneficial to their commerce experience.

Figure 102 below illustrates some of the concepts in the advertising space to synchronize small and large screen. It can leverage new technologies such as parallel feeds and control between small and large screen — where the small screen (mobile/tablet) as a personal device can react with the context of what is happening on the larger screen in communal multi-person mode. This is very powerful and an extension of the STB presence capabilities.

The Cloud can recognize the content being displayed on the large screen at the same time using presence based features that could be added to the STB — also control the content on the tablets or phones present in the home. The cloud also knows if MSO application is running on the companion device and use this to send notifications — update context of display etc. This can be extended to other services — examples include

- Recommendations and next content to watch — can be pushed to the secondary device
- IoT applications can be synchronized across the STB/TV and the companion small screen device(s)

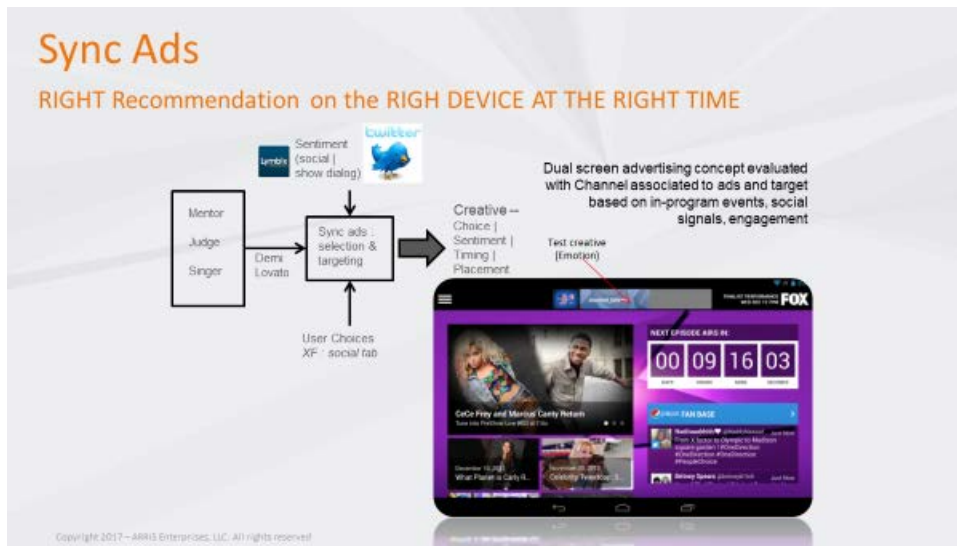


Figure 102 - Multiple Screen Synchronized Advertising



Figure 103 - New Media Control Plane to Allow Synchronized Parallel Activity Across Main TV and Companion Devices

New services can be built up from the parallel feeds architecture — where for example friends living in different locations with the same MSO — can order a movie together as a group — using a “Together TV” concept — where they can also open audio connections and also share social media comments. This service may appeal to Cord Shavers as they share a lower cost when group sharing the watching of video content.



Figure 104 - Group — Together TV — Offers New Social Ways of Watching TV

Integrating social media into the TV experience has been somewhat clunky to date. We know that typing and using lean forward applications like Facebook does not work and small screen devices are better. We do know that Twitter tags and feeds are now part of trending content that people want to share and participate in — so finding ways for users to link segments of video content they see on screen — with the ability to post and comment quickly seems to be of value as people feel compelled to give their opinion on a specific short form segment of video or comment made. There is still work to do to find the best way to integrate social media on the large TV screen (notification for sure) but again using parallel feeds to sync to both large screen and companion device may be the answer.



Figure 105 - Social Media Applications Need to be Split to Work on Large and Small Screen in Parallel

Personalization of the UX with the ability for the consumer to set their home screen feeds — and allow background images and streamed feeds — will become more important as we get closer to the all wall TV. There is already a lot of investment going into adding status (weather, stocks, your orders) to smart assistants and routers — however, there is already a TV driven by a STB in the home — so an

opportunity to make the existing TV through STB drive these information feeds for the home owner. See Figures 106 & 107 that show personalization tools concepts for the UX widgets including background and streaming feeds and menus.



Figure 106 - Personalization Screens and Menus for Dad



Figure 107 - Personalization Themes for Halloween

Another simple example is to leverage the STB (with BLE) for health applications. Most medical devices and wearables have BLE connectivity and they connect to BLE hubs to send patient readings to a secure cloud care portal. The STB acting as this secure hub seems like a logical extension of potential STB capabilities — for aging in place applications as many elderly or infirmed people spend so much time in front of the TV. This allows simple but powerful applications to be built — to allow reminders to patients to take readings, prompt multiple times, and even pause TV use until vital daily readings are taken. See

Figure 108 below — where a simple notification to remind stay at home patient to take their pills (smart pill box has not been opened) using TV notifications.




Figure 108 - Health Notification Example of Patient not Taking Their Daily Pill

Finally — a quick comment on the potential new ‘other’ STB — the VR Renderer — and the associated Head Mounted Display (HMD). There are new experiences the HMD offers particularly for immersion. However, there does seem to be value in also synchronizing the largest flat screen with VR and AR experiences.



Figure 109 - The Technical Parameters that Must be Met for Full Immersive 4K VR

The blend of live TV and HMD experience — switching back and forward in the HMD (watch the live event and the specific thing you want — and then switch to live for replay or zoom in) — is an experience that enhances the VR live viewing experience. One consideration that will be played out is whether in a

home with one VR HMD live viewer — is it a worthwhile application — to also have the viewer's HMD view synchronized to the TV screen using the 360 degree videos capabilities of the current generation of STB SOC's. This may be an application that gives some shared experience in the home — allowing a single HMD wearer to also navigate the 360 experience for others — Figure 110/Figure 111.



Figure 110 - Will Live VR Bring the Best Seat In the Sporting Event to our Living Room Chair?

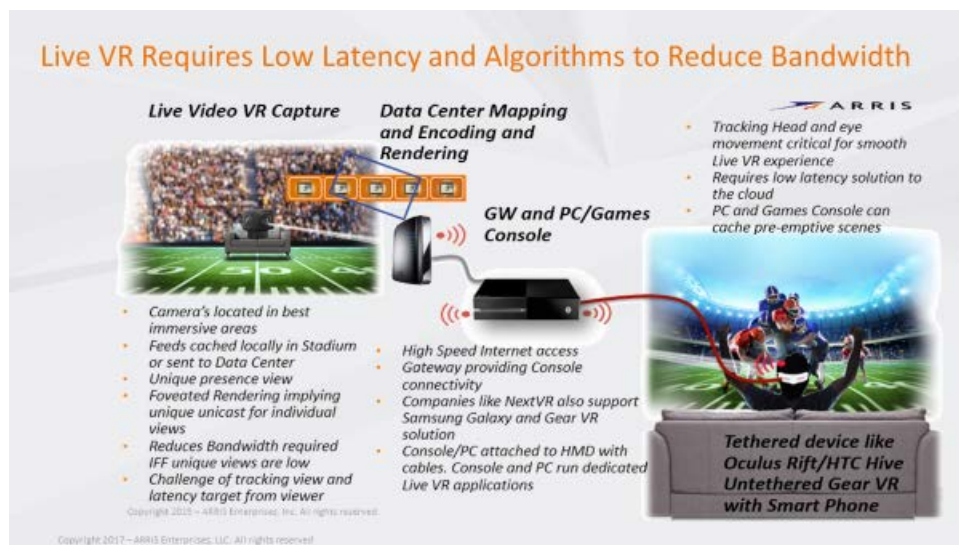


Figure 111 - Low Latency Networks and Low Latency Connectivity in the Home via Gateway and Rendering Device Key for Live VR

Should the MSO consider being the VR aggregator and renderer for a new class of entertainment that is both HMD and TV based? A potential VR home architecture is shown below — in Figure 112 — where future STB add to their ability to do 360 video but also perform full stereoscopic VR rendering to HMD as well.



Figure 112 - Could the Gateway and Future VR Rendering Capable STB Using Wi-Fi and 60 GHz be a Future Solution?

Conclusion

The Pay TV sector has a bright future ahead of it with consumer desire for entertainment at the right bundle and price and the right user experience. While consumers will always be looking for a better deal and value for money — they will stay with the provider who can offer them a path to a convergence of not just “TV”, but a new convergence of all home experiential services. As the screens change to get larger in the home — and as the interaction with the large screen home portal and personal small screen devices — starts to integrate and find its killer applications — the device that converges these experiences to the pixels remains a key part of the MSO architecture and devices.

It is clear that the home gateway and Wi-Fi (and possibly 60 GHz) remain the key transport to connect cloud services and experiences to the home — and it is also clear that the device that terminates them and also aggregates them into a simple and cohesive user experience on the largest screens in the home — is also key. This device is the set-top box, though its name does not fit its future role in the home. Inserting this device in front of the screen — ensures ownership by the MSO of all the experiences, Wi-Fi quality, and the overall quality of the consumer experience. Letting this end point go to a third party device or application on a third party device — seems to give the control up — before the real convergence happens.

It is also key that unless there is an investment in cloud services to bring more and more of the new services to the larger screen pixels then the users — will continue to see the cost of their service — as the video experience only. The fundamental message of this paper and the surveys done show the path forward for the MSO as

- Continue to invest in your own physical set-top device
- Expand its role as a termination of all services to present to the largest screens in the home
- Develop the next level of new services around the leverage of the largest pixel real estate, easiest to use interface — the remote control — and even add smart assistant voice inputs

- Look to drive things like:
 - Customer Self-Healing — why not use TV to take people through problem resolution since 50% of all customer calls require customer education?
 - Health, Commerce, Security and Home Automation, Energy Management, and other services — all of which have reasons to be on the largest screen — even if for notifications
 - Advertising — rethink how it's done with large screen and synced to small personal screens
 - New video experiences like Together TV, with live streaming from new live streaming sources
 - Live VR synchronized to the largest 360 degree viewing capable device
 - Presence based applications — detecting consumer's presence to simplify and improve automated decisions and context for decisions

The set-top box and the user experience it drives remain the key element for retaining customers and growing with them as we see screen technology evolve with flexible wall screens being added to the home. While it may not be a “set-top” box anymore — maybe one day it will look like the image below where it's integrated into every room as the wireless connectivity micro-node serving all video and wireless and experiences to that room. This is the control point to ensure that these services get delivered e2e by the MSO and not by some third party.



Figure 113 – Future in Room MicroNode Media Device

Abbreviations

BLE	Bluetooth Low energy
CPE	Customer Premises Equipment
CPI	Consumer Price Index
CPU	Central Processing Unit
DBS	Direct Broadcast Satellite
DIY	Do It Yourself

DOCSIS	Data Over Cable Service Interface Specification
DVR	Digital Video Recorder
FCC	Federal Communications Commission
FTA	free-to-air
GPU	Graphics processing unit
GW	Gateway
HD	High Definition
HDMI	High-Definition Multimedia Interface
HDR	High-dynamic-range
HEVC	High Efficiency Video Coding
HFC	Hybrid Fiber-Coax
HMD	Head Mounted Display
IoT	Internet of Things
MSO	Multiple System Operator
MVPD	Multichannel Video Programming Distributor
NLOS	Non Line of Sight
OTT	Over the top
QAM	Quadrature Amplitude Modulation
RF4CE	Radio Frequency for Consumer Electronics
S/W	software
SD	Standard Definition
STB	Set top box
UI	User Interface
UX	User Experience
VoD	Video on Demand
VR	Virtual Reality
YoY	Year-over-Year

Bibliography & References

“Statistical Report on Average Rates for Basic Service, Cable Programming Service, and Equipment” in October 2016, https://apps.fcc.gov/edocs_public/attachmatch/DA-16-1166A1_Rcd.pdf

Kleiner Perkins – Mary Meeker report – Internet Trends 2017 Code Conference
<http://www.kpcb.com/internet-trendshttp://dq756f9pzlyr3.cloudfront.net/file/Internet+Trends+2017+Report.pdf>

Nielsen Q4 2016 – Total Audience Report 2016, <http://www.nielsen.com/us/en/insights/reports/2017/the-nielsen-total-audience-report-q4-2016.html>

<http://www.fiercecable.com/cable/top-cable-satellite-and-telco-pay-tv-operators-q2-ranking-comcast-to-directv-to-charter-to> - MoffettNathanson Pay TV subscribers report

Smart Recordings

Dynamic Search and Record of Live TV

A Technical Paper prepared for SCTE by

Chris Lintz
Sr. Principal Architect
Comcast VIPER
1515 Wynkoop St.
Denver, CO 80218

Introduction

X1 – Comcast’s advanced video platform and the X1’s Cloud DVR (cDVR) are used by millions of customers across the country. They expect uninterrupted service, seamless access to thousands of programs, and world-class product features. Comcast VIPER designs and develops IP video solutions supporting X1, cDVR, and mobile technologies.

Users searching for linear programming on their set-top box or mobile device typically enter static metadata such as the program title or series they are interested in. This synchronous search will immediately return results highlighting channels and scheduling for any programming results. Channel surfers may simply find programming by continuously paging the on-screen guide or by rotating through their favorite set of channels.

A user’s interest in content cannot always be defined by searchable static programming metadata. The dialogue within a program can often be a much richer description of the content, but today’s services offer no way to search the dialogue in near real-time and inform a user that their interest is appearing on a live program.

Traditional DVR search and record functionality is also based on static metadata such as show title, program series and genre. Users can press the record button and instantly start recording a program or alternatively, schedule future recordings. While these are key functionalities for DVR use, automated recordings triggered from continuous search across linear dialogue offers a broader content discovery.

Searching and recording based on linear dialogue opens up more value and provides a leap forward in linear and DVR services. Imagine coming home and finding multiple video clips on your DVR containing in-depth interviews of your favorite sports stars appearing on various programs. Channel surfing becomes more interactive when you receive a guide notification informing you that a topic of interest is appearing on a channel you rarely watch.

Shorter, more relevant videos of interest offer a convenience for both at-home and mobile users always on the go. Investors can stay on top of every important conversation about their stocks. Fantasy football fans can capture news and interviews with their roster of players. Researchers and entertainment fans will no longer miss content that brings value to their lives. The possibilities are endless when given the ability to continuously search live TV dialogue 24/7 across all available channels.

This future smart search and record product is running in Comcast’s lab and planning is underway for a customer trial launch, which will be in full compliance with all applicable laws and contractual obligations. The innovation was invented and pushed forward by engineers with the passion to bring more value and content discovery to X1 and mobile device customers. This document focuses on the architecture and technologies supporting the product functionality.

Logical Architecture Overview

Searching linear program dialogue and recording content of interest is achieved through multiple Comcast systems (Figure 1). Video Analysis is a system that provides real-time video, audio, and Closed Captioning analysis. Cloud DVR (cDVR) is a system running in regional Comcast data centers, which supports scheduling, recording, and playback of video for various devices. Linear Search consumes output from a Video Analysis pipeline and user queries, performing search as dialogue streams in a program.

Successful search matches result in notifications sent to the Notification System. Entitlements are checked before recording content – purposely not before linear search. This provides a broad based search across all available linear streams and allows customers to be notified that content of interest is appearing on a stream – even if they are not yet entitled to the stream. Customers entitled to the linear stream result in scheduling requests to the cDVR system. The Smart Search Service provides application interfaces for publish and retrieval of search results, video clip metadata, and notifications.

Future customer notification methods will be supported, including SMS, email, and push notifications to mobile devices. This gives customers the ability to tune instantly to the program and at a position in front of the content of interest.

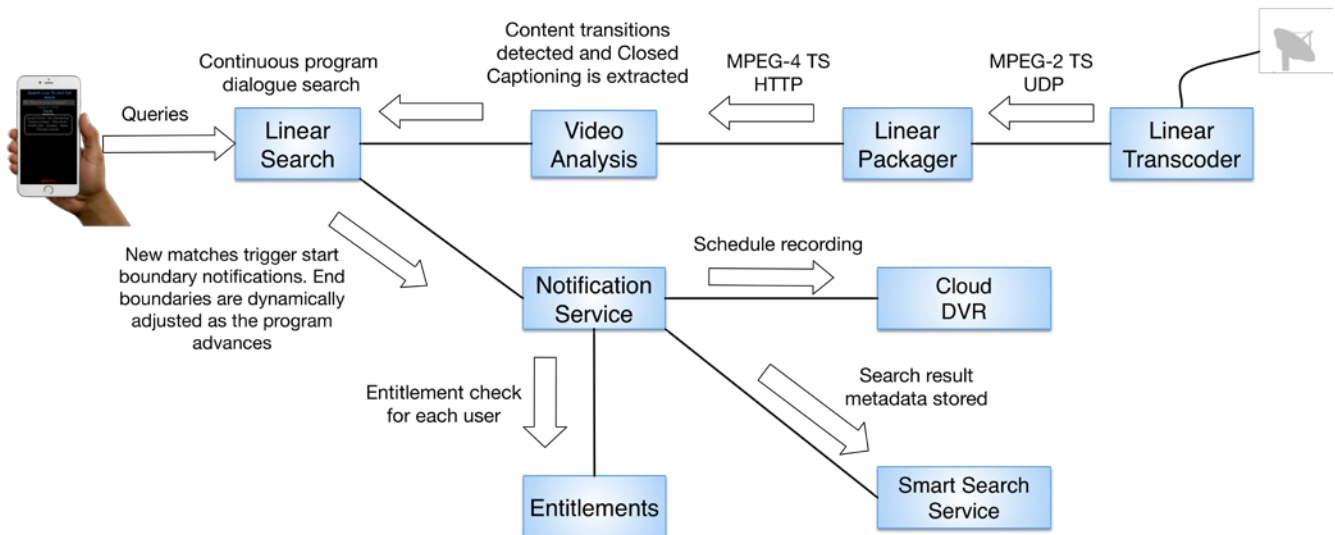


Figure 1 – Linear Search and Record Workflow Detailing System Components

Closed Captioning and Video Analysis

Video segments carry all the necessary metadata needed for searching linear dialogue. Linear textual dialogue and identified content transitions are what downstream linear search components utilize when executing user queries and processing results. Dialogue in the form of Closed Captioning (CEA-608/708) is carried in the picture user data on the transport stream. I-frames contain image data, which is used when analyzing the video to determine any content transitions.

Transitions within the content of a program provide opportunities to define boundaries around content when a match occurs from linear dialogue search. A shot change (Figure 2) is a slightly different camera perspective within the same scene of content. Scene changes occur when an entirely different camera perspective occurs within the same program.



Figure 2 - Visual Example of a Shot Change

SCTE-35 signaling in the manifest can be relied on for local ad spots, identifying a scene change. When this signal is not available, shot and scene detection algorithms help identify content changes in the video. This involves decoding image packets for color and edge information and applying mathematical formulas to detect movement from one frame to the next.

While these algorithms can be computationally expensive and require storing previous computations from frame analysis, they are achievable in real-time and provide a high detection accuracy. Detecting shot changes is less computationally expensive and requires less storage of previous frame analysis. These algorithms are the first step in a video analysis pipeline (Figure 3).

A process called a Stream Reader monitors linear manifests from a Linear Packager. Stream Readers scale horizontally consuming in aggregate over 10,000 local and national streams. The Linear Packager is an internal video packager supporting at rest encryption and an intermediate format derived from DASH. Each time the monitored manifest is updated, video segments are pulled and video frames are analyzed for shot and scene changes.

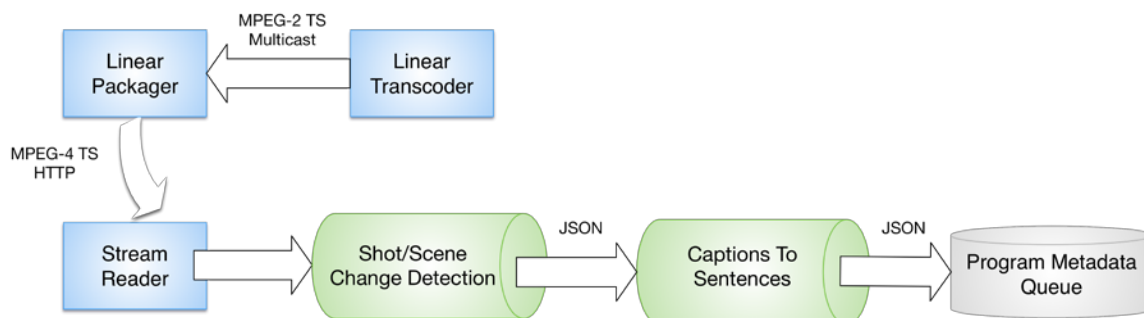


Figure 3 - Video Analysis Pipeline

Each segment carries an encoder boundary point (EBP) containing a sequential timestamp relative to the transcoder. These EBP timestamps are extracted along with the textual Closed Captioning data. Sentence formation is constructed if there is a partial phrase. A series of phrases, which ultimately form a sentence, may be spread over multiple segments. Multiple segments may result in more than one shot or scene change. All shot and scene change times are reflected as an array of EBP times in the program metadata document (Figure 4).

```

{
  "program": "Mad Money",
  "streamName": "CNBC HD",
  "streamId": "8951620516683951163",
  "ebpTime": 1496438204,
  "text": " >> Tesla passes Ford in market value now up over $350 per share
           as it continues an unreal run ahead of the Model 3",
  "shotChange": [1496438204, 1496438208],
  "sceneChange": [1496438212]
}
  
```

Figure 4 - Program Metadata Example

Once a sentence is formed it is also included in the resulting program metadata document, which is then pushed onto the queue making it immediately available for search.

Linear Search and Record

Typical search systems store static documents, inverted indexes are built, and queries are executed against the indices. When a large amount of queries are searched over the same document, optimizations are made when inverting this concept. As streaming documents arrive they are tokenized and searched against query indices. Candidate query matches are returned, requiring a document search in order to resolve search hits and relevancy. This inverted search concept is commonly known as stream search, or reverse search, and can greatly reduce the number of queries executed.

Searching linear dialogue at Comcast scale can equal hundreds of thousands of queries running across hundreds of linear streams in large regions. Linear packagers output two-second segments resulting in dialogue changes received at that cadence per linear stream. In a typical Comcast region with hundreds of linear streams, the Video Analysis Pipeline produces streaming text documents at a rate of 200-300 per

second. The volume of queries combined with the influx of text documents makes streaming search a desirable technique for near real-time dialogue search.

Querying dialogue between the program start and live point of the program also provides value to users. A query added during a live program will search the past dialogue allowing for back-in-time notifications and recordings if the video segments are still available in the Linear Packager.

1. User Queries

Queries added to the system exist as live searches until removed by the user. Query Parsers (Figure 5) receive the submitted user queries and are responsible for filtering and expanding queries. Supported user query types are simple terms and phrases with limited conjunction and disjunction clauses.




Figure 5 –Query Parser and User Queries

Within a program scene, a close distance between two phrases in Closed Captioning text represents conversationally related phrases. Matches resulting from proximity queries on textual dialogue identify conversational relevance. Proximity queries are used internally when expanding some multi-phrase queries.

Editorialized synonyms help expand popular queries into broader meanings. As one example, the two queries “Donald Trump” and “President Trump” would result in the same query “President Trump OR Donald Trump”. Queries are then normalized into an internal query representation and submitted to the queue.

2. Query Partitions

The Query Queue is partitioned so that each partition holds a subset of user queries. Using simple hashing on expanded queries provides a common routing technique resulting in the application writing identical queries to the same partition: $Partition\ ID = Hash(Query) \% Total\ Partitions$

This approach ensures that the same Query Engine handles identical queries. This allows us to create one-to-many relationships of queries to users so that only a single query is executed for multiple users. The reduction in the amount of queries can be drastic for popular queries. It also provides optimizations resulting from being able to batch notification messages and cDVR recordings.

Queue partitioning (Figure 6) is a pattern that can be implemented with any persistent store. We utilize Kafka as our queue primarily for the built-in partitioning, consumer groups, and log compaction features. Log compaction maintains at least the latest version of a key. Each query partition is essentially a

persistent store for user queries. This allows the Query Engines to glue onto their assigned partitions and handle re-start and failure scenarios by again consuming the set of queries.




Figure 6 - Query Partitioning

3. Query Engines

Two popular streaming search libraries are Lucene based Elasticsearch-Percolators and Luwak. Pre-filtering techniques are used in Luwak, which eliminates queries that are not a possible match and presents candidate queries that may be matches. We have chosen to embed Luwak in Query Engines for stream search functionality because of the performance gains proven in lab testing.

Query Engines scale horizontally and perform in-memory search for both stream search and full program dialogue search. Figure 7 below depicts high-level workflow between components.

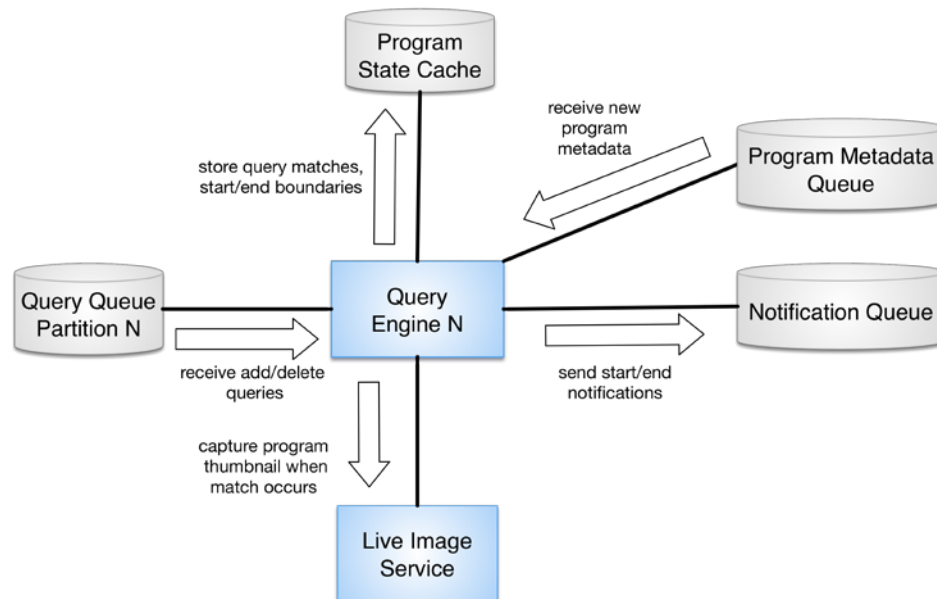


Figure 7 - Query Engine High Level Workflow

3.1. Query Filters

Queries can contain preferences such as stream or program filters for fine-grained search over desired programming. For example, a user may choose a broad search across all available linear streams rather than filtering on a single program.

Users with the same query are combined and any stream or program filtering is applied. In the example below (Figure 8), Chris and Tony are interested in “Comcast OR Netflix” but only if it is discussed on CNBC’s “Squawk on the Street” or “The Closing Bell”. Jen is interested in “Climate Change” if it appears on any Weather Channel program.

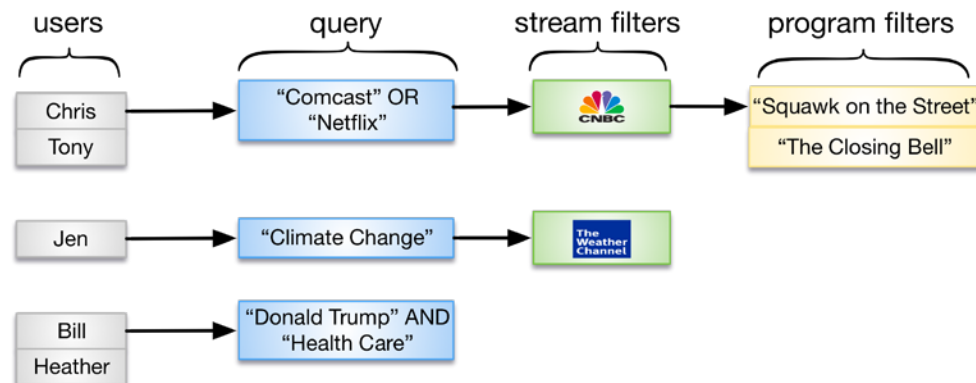


Figure 8 - Stream and Program Filters

Bill and Heather want the broadest search across any linear stream for discussions of “Donald Trump” AND Health Care”.

3.2. Stream Search and Program Transcript Search

Query Engines are assigned with a query partition ID and a list of linear stream IDs. Documents for each live program are retrieved from the queue for each stream ID. Luwak monitors queries and matches new program metadata documents against them. The documents are also used to form a program transcript document for each stream.

These program transcript documents are not part of stream search, but they are searched directly with Lucene for two use cases. The first is when a query is added during a program. This allows matches to be found behind the live point of programs. Searching program transcripts are also needed for some complex queries. For example, a proximity query use distance between words or phrases, which may require a search into past dialogue.

Searching the transcript document not only provides opportunities to trigger recordings with a start boundary back-in-time, but notifications from matches can result in other non-recording actions. For example, a customer can be presented an option to tune to a point in time behind the live point where their interest appears. Tuning back-in-time is possible through Instant VOD (iVOD), a Comcast service supporting live program rewind. A user can also be presented an option to set a scheduled recording for the program’s next airdates.

Just a couple minutes into a typical chatty hour-long news analysis program such as MSNBC's Hardball with Chris Mathews, an average transcript document contains less than a couple hundred words. By program end the transcript document can be over 12K words not including commercial dialogue. This can produce a roughly 20KB document size - resulting in just 10MB of RAM for 500 one-hour programs. These documents are maintained locally in-memory for the duration of the Query Engines runtime.

It should be no surprise that both stream search and transcript document search are compute bound. The frequency of program transcript document updates, which requires re-indexing, is an added burden. Both stream search and traditional search techniques have different performance considerations. We address these considerations by adjusting two key parameters in the system. (1) Total query partitions in the system and (2) total list of linear streams consumed by each Query Engine. This also allows for a great deal of flexibility for tuning deployments for different regions with different numbers of local and national streams, running on different hardware.

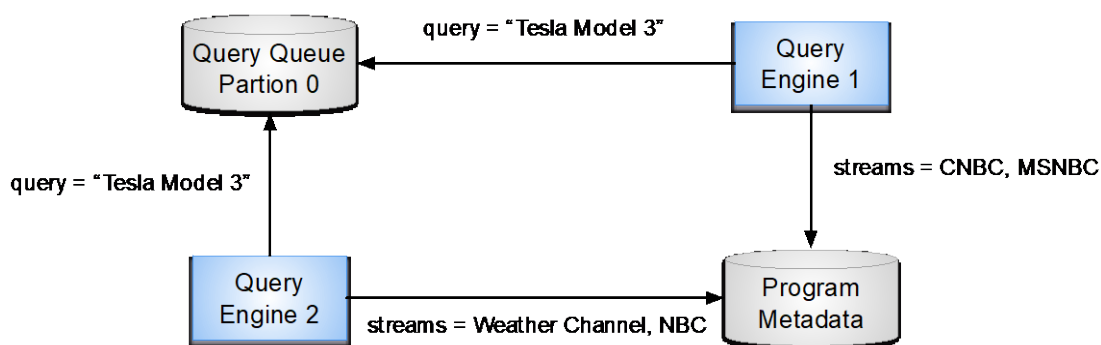


Figure 9 - Query Engines Consuming the Same Queue Partition but Different Linear Streams

3.3. Utilizing Content Transitions for Video Clip Boundaries

A content transition timeline for each program is maintained by extracting shot and scene change times from program metadata documents. Closed Captioning drifts in varying durations on all linear streams all day long. Frame accuracy search matches are not of great value due to this drift. Even if the EBP time near the matched sentence were to be used as the start boundary for the video clip, it likely would be in mid-dialogue or in the middle of scene. While this will capture content relevant to the query, it likely will not result in a great user experience.

A better quality video clip can be achieved by using content transitions before and after the time of the query match. For example, the moment of a commercial end can be used as the start boundary of the video clip – in front of the match. A scene change that occurs at some time after the match can represent the end boundary of the video clip.

The Query Engine utilizes the content transition timeline in attempt to capture a better start boundary and end boundary for a desired video clip. In Figure 10, if a search match occurs at the live point, the EBP time of the nearest scene change (walking backwards in time) is used as the start boundary for the video clip.

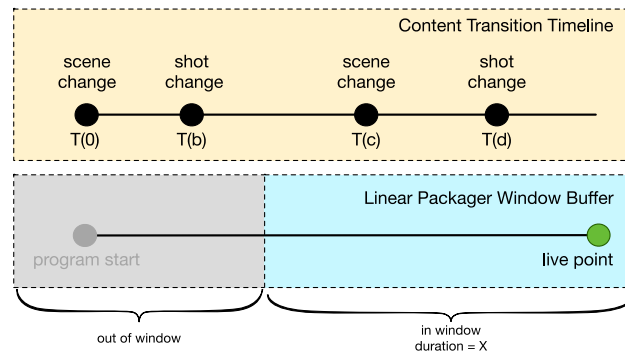


Figure 10 - Content Transition Timeline and Linear Packager Window Buffer

Scene changes are preferred over shot changes and if neither transition change is available in the past, the related EBP time of the match is used. In the above example the scene change EBP time at T(c) is selected as the start time of the recording. T(c) is within the Linear Packager window buffer allowing for a successful back-in-time start of a recording.

A more detailed example describing how basic start and end boundaries work around a content of interest is depicted in Figure 11 below.

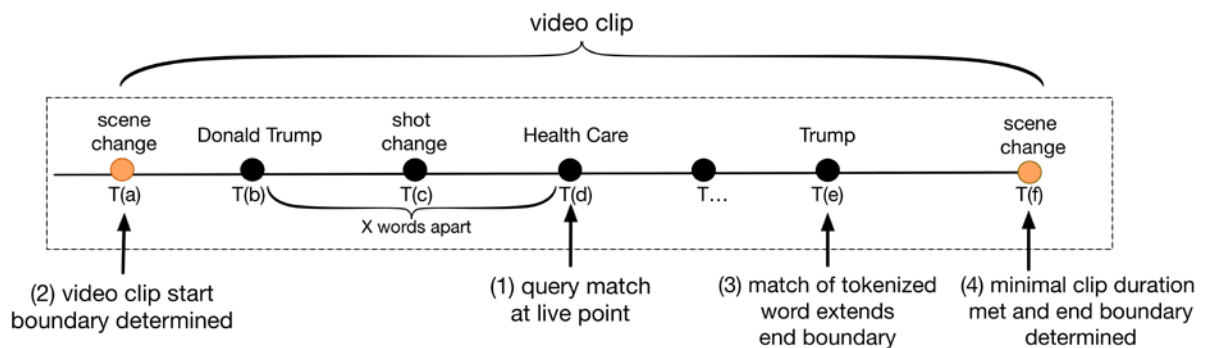


Figure 11 - Video Clip Timeline Utilizing Content Transitions

“Donald Trump” AND “Heath Care”{X} is a conversationally relevant query that will only match if the two phrases are at most X words apart. At the live point T(d), the tokenized phrase “Health Care” appears in the current sentence triggering the proximity search against the program transcript document. A match is found and the EBP time before Donald Trump appears is used to find a prior scene change EBP time in the content transition timeline cache, which results in T(a). The desire to be in front of the matched phrases eliminates shot change at T(c) because it falls within the proximity query.

If there is a preference to record the remaining program, the end time of the program is determined and set as the end boundary for the video clip. Otherwise the end boundary is set as a fixed duration and adjusted dynamically as the program progresses. At this point a notification can be fired off, resulting in the start of a recording.

Desired video clip durations are attempted and end boundaries are extended if tokenized words from the query are found in new sentence dialogue. Time progresses and at T(e) the tokenized word “Trump” is in the current dialogue and the end boundary is extended by a fixed duration. At T(f) the desired duration of

the video clip has been met and a scene change has occurred. This triggers the end boundary to be set to T(f).

Searches against the program transcript document also utilize the content transition timeline. EBP times embedded in the transcript document provide timestamps needed for content transition time lookups. For example, Figure 12 below represents a Lucene highlighter result from the query “Tesla” against a program transcript document.

[1498423117]Ford’s market value has just been passed by Tesla as the stock breaks through \$350 a share.[1498423121] The first deliveries are expected in late July for the Model 3.

Figure 12 - Example Program Transcript Snippet w/EBP Times (Lucene Highlighter Result)

Similar to stream search matches occurring at the live point, The EBP time in-front of “Ford’s” will be used to find the nearest scene change time in the content transition timeline.

4. cDVR Supprt for Smart Recordings

A Scheduler, external to cDVR regions, manages and schedules recordings via a Recording Manager at located at each local region. Scheduling logic was not changed, but a more condensed down scheduling logic was built into the Notification system. Modifications were introduced to cDVR Manifest Agents in order to support back-in-time recordings.

Linear Packagers contain a window of content behind the live point. Manifest Agents are continuously monitoring manifests for updates and maintain a cache of manifests within a rolling window (Figure 13).




Figure 13 - Cached Manifests for Back-in-time Recordings

Maintaining a cache of manifests allows for recording video segments behind the live point that fall within the Linear Packager’s buffer. Requests to record segments outside of this buffer will default to the oldest segment available in the window.

A match from a single query representing multiple users generates a batched notification. This results in a batched recording request. Copies are unique per subscriber, but batched recordings result in optimizations to the underlying network and storage system by generating a fan-out request to persist the unique video segments per user. While batched recordings are not new to cDVR, it is important to note that combining of users with the same query not only optimizes search, but also cDVR resources.

Future Considerations

The total tuners available limit concurrent recordings for a set-top box. Concurrent recording limitations also apply to cloud-based recordings. Popular topics and current news event queries can quickly reach the concurrent recording limit. As an obvious example, a simple query like “Donald Trump” is likely to appear on many channels simultaneously and can quickly consume the total active recording limit – particularly if stream and program filters are not set for each query.

In order to optimally capture the most content for a user while maximizing the concurrent recordings, many options are being considered:

1. The avoidance of recording video clips for programs already scheduled to record. For example, if a program is set as a recurring recording it is clear a customer is interested in the full program. Rather than capturing video clips of the already scheduled program, metadata can be added to the recording or notifications can inform the customer that matches appear within the recording.
2. Multi-weighted algorithms that include customer priority channel, program, and query rankings.
3. Suggested program and/or stream filters for popular queries that are likely to consume multiple concurrent recordings.

Conclusion

Identifying content transitions combined with Closed Captioning are key to dynamic searching and recording video of interest. Small modifications to cDVR that allow for back-in-time recordings provide opportunities for smooth recording starts, which also capture more context in front of a user’s matched interest. These back-in-time recording starts combined with dynamically extending end times and ending a video clip on a content transition produce a better quality video clip and user experience.

The value of content discovery on linear streams in near real-time is made possible by a combination of streaming search techniques on current sentence dialogue and search over active program transcript documents. Notifications providing opportunities for users to tune to a program back-in-time or to set scheduled recordings for the next airdate are also value adds from the system.

cDVR revolutionized how customers record and view video content. Expanding search and record capabilities using linear dialogue and video analysis is the next leap in offering additional features and value, allowing customers to never miss their interests appearing in any stream with a new personalized experience.

Abbreviations

EBP	encoder boundary point
cDVR	cloud digital video recorder
VIPER	Video Internet Protocol Engineering and Research
DASH	Dynamic Adaptive Streaming over HTTP
MPEG	Motion Pictures Expert Group
iVOD	instant VOD

Bibliography & References

Zheyun Feng, Jan Neumann, “Real Time Commercial Detection in Videos”

Apache Kafka, <https://kafka.apache.org/>

Luwak, <https://github.com/flaxsearch/luwak>

Elastic Percolator, <https://www.elastic.co/guide/en/elasticsearch/reference/current/search-percolate.html>

SCTE-35, Digital Program Insertion Cueing Message for Cable, ANSI/SCTE-35 2013

Echo Cancellation Techniques for Supporting Full Duplex DOCSIS

A Technical Paper Prepared for SCTE•ISBE by

Hang Jin

Distinguished Engineer
Cisco Systems
2200 E Pres. George Bush Turnpike, Richardson, TX75082
469-255-2666
hangjin@cisco.com

John Chapman

CTO, Fellow
Cisco Systems
3700 Cisco Way, SAN JOSE, CA 95134
408-526-7651
jchapman@cisco.com

Introduction

Full duplex (FDX) DOCSIS® allows the downstream and upstream to use the same radio frequency (RF) spectrum at the same time, leading to ~100% increase of spectral efficiency. With FDX DOCSIS, the upper band edge of the upstream spectrum can be extended beyond 204 MHz, leading to five to 10 times increase in upstream throughput. The downstream throughput is also increased as the use of FDX DOCSIS can eliminate the crossover band of current frequency division duplex (FDD) systems and push the low band edge of the downstream below 258 MHz. Using the 10 MHz to 1.2 GHz spectrum, full duplex DOCSIS has the capability to provide 10 Gbps throughput for the downstream and 5 Gbps throughput for the upstream.

As the downstream and upstream spectrums overlap in FDX DOCSIS, interference occurs between transmission and reception. Thus, interference mitigation is a key enabler for supporting FDX DOCSIS. Echo cancellation (EC) is required in FDX systems to suppress the interference that is coupled or leaked from the transmitter to the receiver as they operate on the same frequencies. Cisco invented and prototyped FDX DOCSIS echo cancellation algorithms, and demonstrated them in August 2016 at the CableLabs Summer Conference.

This paper explains the types of interference that occurs in FDX DOCSIS operation and the corresponding echo cancellation techniques required. The paper is organized as follows. Section 1 explains the basics of FDX DOCSIS operation. Section 2 explains the challenges with FDX DOCSIS operation: the interference from transmitter to receivers and the interference among cable modems (CMs). Interference cancellation must be implemented for supporting FDX DOCSIS operation. Section 3 explains the network topology for supporting FDX DOCSIS. The details on the interference types and the corresponding echo cancellation techniques are given in 4-7. Section 8 explains the echo cancellation lab prototype system and test results, and Section 9 explains the live FDX DOCSIS proof of concept (PoC) demonstration system.

Content

1. FDX DOCSIS: The Continuing Innovation

Today cable access (DOCSIS) employs frequency division duplex. With FDD, the usable frequency spectrum is divided into non-overlapping downstream (DS) and upstream (US) spectrums and a crossover spectrum in between (Fig. 1).




Figure 1 - FDD frequency split

There are three frequency divisions, so called frequency splits, in use today: low split, mid split and high split (see Table 1). DS traffic from the cable modem termination system (CMTS) to CMs is sent in the DS spectrum, and US traffic from CMs to the CMTS is sent in the US spectrum. This FDD DS and US frequency division is completely overthrown in FDX DOCSIS: Both DS and US traffic can use the same spectrum at the same time, resulting in doubling the spectrum usage efficiency as the same spectrum is used simultaneously for DS and US traffic. (Fig. 2).

Table 1 - Frequency splits

	US Spectrum	crossover	DS spectrum
Low split	5 MHz-42 MHz	42 MHz-54 MHz	54 MHz-1218 MHz
Mid split	5 MHz-85 MHz	85 MHz-108 MHz	108 MHz-1218 MHz
High split	5 MHz-204 MHz	204 MHz-258 MHz	258 MHz-1218 MHz

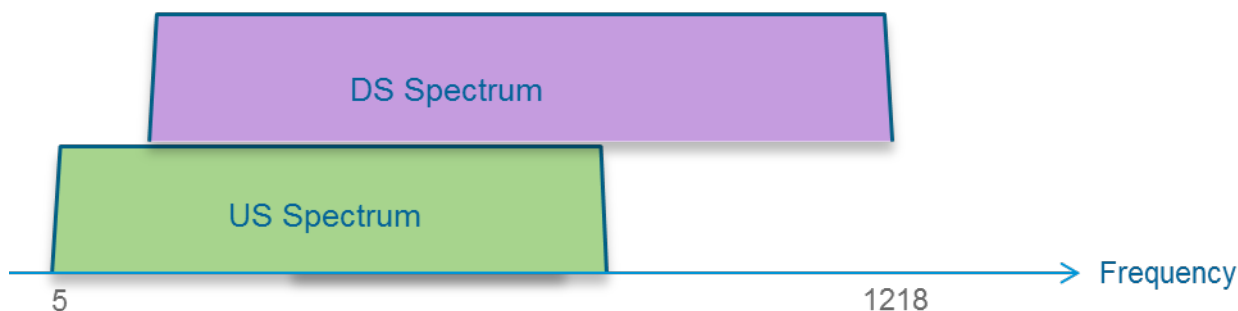


Figure 2 - FDX concept

Compared to today's FDD DOCSIS, FDX DOCSIS has a clear advantage: It 'creates' frequency spectrum for US traffic without sacrificing the frequency spectrum for DS traffic. Today's coaxial network has roughly 1.2 GHz of usable spectrum (limited by the attenuation of the taps installed in the field, which have a sharp roll off around 1.2 GHz). Given this fixed 1.2 GHz of usable spectrum in the coax network, increasing the US spectrum will reduce the DS spectrum if FDD is used. FDX DOCSIS allows DS and US traffic overlap on frequencies, effectively doubling the usable spectrum of the coax network to 2.4 GHz.

Cable operators have been working on increasing the US spectrum in order to keep up with user data demands and stay ahead of the competition. Most of the networks deployed today are either low split, which has 37 MHz total US spectrum, or mid split, which has total 80 MHz US spectrum. 37 MHz or 80 MHz US spectrum is definitely not enough to provide acceptable user experiences given the US spectrum needs be shared among dozens or even hundreds of users. Migrating the frequency plan to high split will increase the US spectrum to ~200 MHz, which may ease some of the US data congestion, but migrating to high split not only results in CAPEX, but also the obstacles of 75 MHz/107 MHz out-of-band (OOB) signals: one needs to create a "jumper" in the middle of the US spectrum at every active device to allow the DS OOB signals to propagate from the headend to customer premises equipment (CPE) (Fig. 3).




Figure 3 - OOB issue with high split

FDX DOCSIS has a completely new paradigm to solve the US spectrum shortage: it extends the upper edge of US spectrum to as high as 1.2 GHz without taking away any spectrum from the DS. As FDX DOCSIS still operates in the frequency range of 5 MHz to 1.2 GHz, no taps in the field need be replaced. US spectrum in FDX operation can be allocated as high as 1.2 GHz, the OOB signal can still be present at 75 MHz/107 MHz as a part of a guard-band, so no changes are required for supporting OOB signal.

2. Challenges with FDX DOCSIS

While it provides all the benefits, FDX DOCSIS presents many implications and design challenges. The biggest challenge among them is the interference from transmitter to receiver at the CMTS (RPD node) and among CMs.

2.1. Interference from transmitter to receiver at RPD node

Since both DS and US signals use the same spectrum in full duplex operation, the transmitted and received signals will overlap in frequency and time. The transmitted DS signal has a much higher signal level than the received US signal, and the leakage or coupling from the transmitter to the receiver will become co-channel interference and may completely wipe out the received US signal if there is not sufficient isolation between the transmitter and the receiver. The co-channel interference from transmitter to receiver is one of the hurdles that needs be overcome to make full duplex work in the coaxial network.




Figure 4 - Interference from transmitter to receiver

Interference from transmitter to receiver at RPD-equipped nodes needs be cancelled out through echo cancellation for supporting FDX operation.

2.2. Interference at CM

In FDX DOCSIS, the CM still works in FDD mode. Although the CM transmits and receives on different channels, there is still interference occurring at the CM. There are two types of interference at the CM: self-interference and interference among CMs (neighboring CMs).

2.2.1. CM self-interference

Although the CM transmits and receives on difference channels, CM self-interference can occur when a CM transmits and receives at the same time. The out-of-band spurious emissions of the transmitted channel may couple and leak into the received channel and become co-channel interference, thus increasing the noise floor of the received signal. Also, the transmitted signal coupled into the receiver may have much higher power than the desired DS signal and saturate the receiver front end.

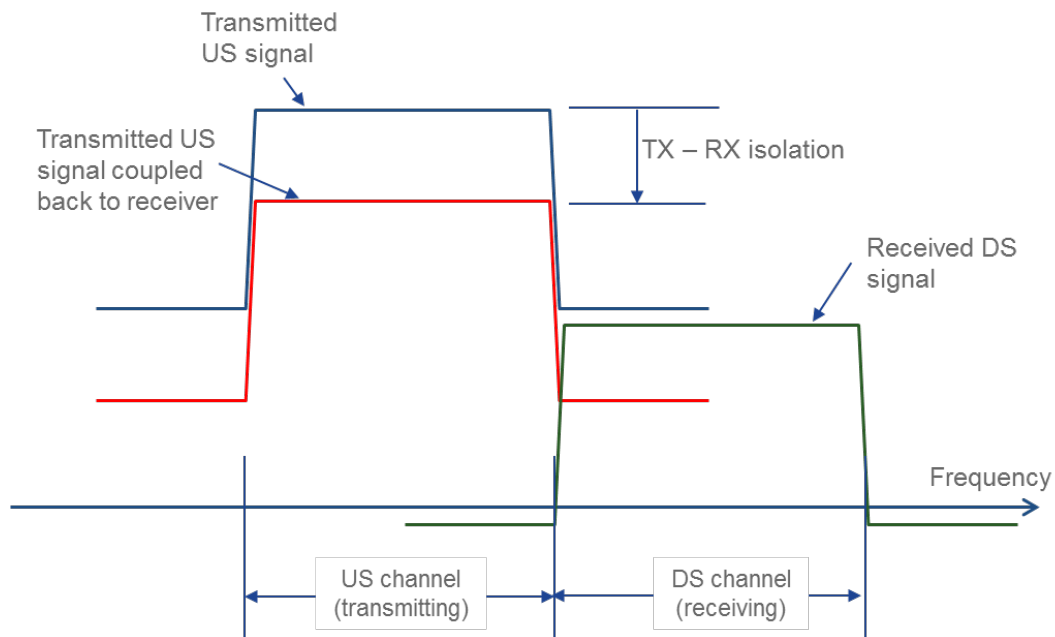


Figure 5 - Interference at CM (CM self-interference)

CM self-interference needs be cancelled out through echo cancellation.

2.2.2. Interference among CMs

Interference could occur among CMs. For example, CM1 is on tap 1 and transmitting on channel 1, and CM2 is on tap 2, receiving on the same channel (channel 1). The transmitted signal from CM1 on channel 1 may leak into CM2 and impair CM2's reception on channel 1, if there is not sufficient isolation between CM1 and CM2.




Figure 6 - Interference among CMs

Interference among CMs is mitigated through smart scheduling: schedule the DS and US channel allocations among CMs in such a way to better leverage the isolation among CMs so there is no or little interference among CMs that are transmitting and receiving on the same channel. The nutshell of this smart scheduling is to avoid allocating overlapping DS and US channels to CMs that may interfere with each other. For details on this interference avoidance scheduling, please see reference [1].

3. Network topology for FDX DOCSIS

Technically, FDX DOCSIS could work with any network topology as long as one could develop FDX nodes as well as FDX bi-directional amplifiers. Bi-directional amplification of FDX signals results in bi-directional interference: the output of the DS transmitter will interfere with the reception of the US receiver, and similarly, the output of the US transmitter will interfere with the reception of the DS receiver. This will require bi-directional echo cancellation on both DS and US paths, and ensure the total interference cancellation of the complete loop will be greater than the closed loop gain of the amplifier. Designing FDX amplifiers presents a great challenge on an echo cancellation scheme and RF design.

The assumption today is that N+0 network topology is required for supporting FDX operation to avoid the implications related to FDX amplifiers (design challenges, active amplifiers in the field).

3.1. N+0 network topology

N+0 means there are no active amplifiers between the R-PHY node and CMs. The coax and taps between the R-PHY node and CMs are all passive RF components and can support bi-directional RF signal transmission according to Lorentz reciprocity theorem. The only changes required for supporting FDX are in the R-PHY node and CM. This will avoid expensive upgrade or replacement of the coax and taps already deployed in the field.

4. Interference in FDX operation

In this section, we list and discuss all the interference sources and their power levels in a FDX DOCSIS system.

The interference sources and their power levels are different in the RPD node and CM.

4.1. Interference sources at RPD node

With FDX operation, the DS and US are overlapped in time and frequency, so at the RPD node, the DS traffic may couple from the transmitter into receiver and become co-channel interference to the received signal. This co-channel interference may come from multiple sources (Fig. 7)

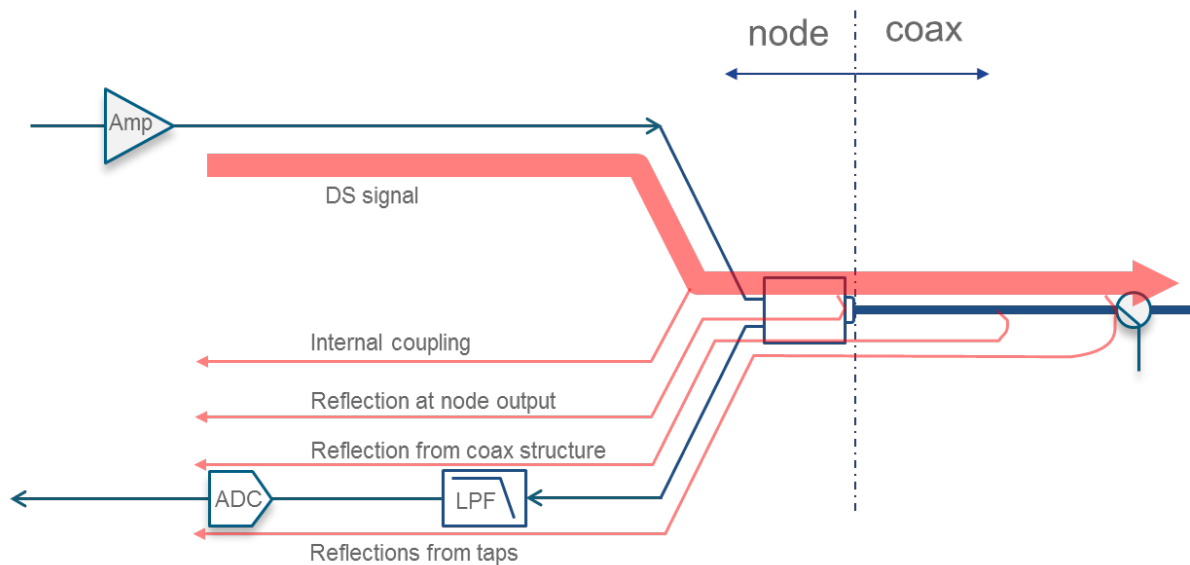


Figure 7 - Interferences in FDX node

4.1.1. RPD internal coupling

The internal coupling is caused by limited isolation between the transmitter and receiver paths. This also includes the coupling between the transmitter and receiver ports of the three port device used to connect the transmitter and receiver to the coax (the common port). Typically this is a directional coupler with limited isolation between the input and coupled ports.

4.1.2. Reflection at node output

The node output has a limited return loss, depending on the type of connector and its quality. Any reflection of the DS signal at the node output will become co-channel interference at the receiver.

4.1.3. Reflection from the taps

Each tap has a limited return loss. All the taps deployed in the field have >23 dB return loss, per published tap specs. Reflections of the DS signal at taps will go toward the US direction, experience the path loss between the taps where the reflections occur and the node input port, and become co-channel interference to the receiver. A long cable between the tap and the node will help reduce the power level of the tap reflection received at the node receiver.

4.1.4. Reflection from coax discontinuities (coax structure reflection)

There are reflections caused by discontinuities in the coaxial cable itself: the coaxial cable structure imperfections cause reflections. The structure imperfections could be the inhomogeneous nature of the dielectric or the inner conductors and outside shields. The reflections resulting from the coax structure imperfections are small in power compared to the reflections from other sources, but they spread over much large time intervals.

4.2. Interference power levels

The interference power levels are different for different sources. The reference point used for interference level is the node input port (node interface D)

4.2.1. RPD internal coupling

This is the internal coupling within RPD. The general design guidance is to reduce the interference due to internal coupling to a negligible level compared to that of the interference from other sources. The dominant interference occurring within the RPD may result from coupling between the transmitter and receiver ports of the three port device used to connect the transmitter and receiver to the coax (the common port). A minimum of 40 dB isolation between the transmitter and receiver ports is preferred.

4.2.2. Reflection at node output

This is the internal reflection at the node output. The same principle is applied here: to reduce the internal reflection from the output port to a negligible level compared to that of interference from other sources.

One may sum up all the interference from all the internal sources (internal coupling and internal reflection) and specify its level to be X dB below the DS output power. With 72 dBmV total composite power (TCP) node output power, the interference level from internal coupling and reflection will be $72 - X$ dBmV (208 MHz to 1218 MHz). $X > 35$ is preferred.

4.2.3. Reflection from the taps

Assuming 23 dB tap return loss, the reflection from the tap will be 23 dB below the DS transmit signal in power. Assuming 2 dB loss each way in the initial feeder (~100 ft express cable), and 72 dBmV TCP node output power, the reflection at the node input will be $72 - 23 - 2 \times 2 = 45$ dBmV (108 MHz to 1218 MHz).

4.2.4. Reflection from coax discontinuities (coax structure reflection)

The reflection from coax discontinuities (structure reflections) are much lower than other sources. Its magnitude is much lower than the reflection from the tap as illustrated in Fig. 8. Fig. 8 is the measurement data of all the reflection sources from an actual N+0 network.




Figure 8 - Lab measured results on the reflection from passive coax network

Although the reflections from the coax discontinuities are lower in power, they can't be ignored. As indicated in Fig. 8, the reflections caused by coax discontinuities are ~20 dB below the tap reflections in magnitude, equivalently having a return loss 43 dB. Given 72 dBmV TCP node output power, with -43 dBc reflection, the reflections at the node input is about 29 dBmV (108 MHz to 1218 MHz), significantly above the thermal noise level (thermal noise power within 108 MHz – 1218 MHz \approx -30 dBmV, assuming 5 dB noise figure). The reflections from coax discontinuities must be cancelled out to ensure proper FDX operation.

4.3. Interference power level in FDX frequency spectrum

The interference levels computed in the previous sections are for the frequency range of 108 MHz to 1218 MHz. As the FDX frequency spectrum covers only 108 MHz to 684 MHz, a filter with a cutoff frequency at 684 MHz will be put in the receiver path to suppress the interference power above 684 MHz. Because the DS signal has a 21 dB up tilt from 108 MHz to 1218 MHz, filtering out the power in the frequencies above 684 MHz will significantly reduce the total power of the interference seen by the receiver. There is roughly 10 dB power reduction if the interference above 684 MHz is filtered out. Using the same

assumptions in 4.2.3, the total power of the reflection received at the node input in the frequency range 108 MHz to 684 MHz will be 35 dBmV.

4.4. Impacts of DS up tilt

The actual reflection power level varies with the frequency as the DS transmitted signal power level varies with frequency. The power density of DS transmitted signal is not flat over the frequency spectrum, it has 21 dB linear up tilt from 108 MHz to 1218 MHz. This is because the loss of the plant at 1.218 GHz is roughly 21 dB higher than that at 108 MHz, so in order to compensate for this un-equal loss of the plant, the power density of the DS signal is pre-emphasized by uptilting the power density of the DS signal at the transmitter to make the received DS signal power density at the CM relatively flat over the spectrum.

This 21 dB up tilt of DS power density implicates the power density of the reflections received at the RPD-equipped node receiver. As in most of the cases, the dominate reflections come from the node output and first a couple of taps, the total loss does not have 21 dB of down tilt, thus the reflection received at the node receiver is still largely up tilt. This means most of the interference power is concentrated at higher frequencies. Per the FDX DOCSIS 3.1 specification, the FDX operating spectrum is from 108 MHz to 684 MHz, and three DS OFDM channels overlap with six US OFDMA channels. The interference level observed on the last OFDMA channel will be much higher than that on the first OFDMA channel, and in an extreme case (most of the reflection comes from the node output), the interference observed on the last FDX OFDMA channel could be 10 dB higher than that of the first FDX OFDMA channel as indicated in Fig. 9. To achieve the same post-EC performance, the EC must achieve 10 dB better performance on the last FDX OFDMA channel than the first FDX OFDMA channel.




Figure 9 - The interference power density. FDX US channel 6 could see the interference power density 10 dB higher than channel 1 due to the DS signal up tilt.

4.5. Impact of interference on node receiver

The impact of the interference on the node receiver are twofold, as discussed in the next two sections.

4.5.1. Receiver dynamic range

The receiver has a limited dynamic range. The dynamic range of the receiver is mainly limited by the analog-to-digital converter (ADC). For a FDX RPD node, the frequency range of the US is from 108 MHz to 684 MHz. For this range of frequencies, a state-of-the-art ADC can achieve ~50 dB in-band modulation error ratio (MER) for an OFDMA signal with a flat power density. As indicated in Section 4.3, the power level of the reflections in the FDX frequency spectrum 108 MHz to 684 MHz can be as high as 35 dBmV. In FDX operation, the desired US signal power density is around 0 dBmV/6.4 MHz, or ~20 dBmV in the frequency range of 108 MHz to 684 MHz (all six OFDMA channels), indicating the reflection can be 15 dB higher than desired US signal in power. The receiver ADC needs to leave sufficient head room to accommodate the high reflection power level. Leaving head room for reflections will directly impact the effective MER that can be achieved with the ADC, and the achievable MER is reduced dB by dB with the head room reserved, as indicated in Fig. 10




Figure 10 - ADC dynamic range, and the impact of the interference on achievable SNR

One technique to reduce the impact of the reflection on ADC dynamic range is to incorporate analog echo cancellation before the ADC. More details on this are given in Section 5.1

4.5.2. Co-channel interference

Another impact of the reflection on US receiver performance is co-channel interference. As DS and US share the same spectrum at the same time, the reflection of the DS will become co-channel interference to the US. Co-channel interference must be suppressed to a level that the targeted US MER can be met. Co-channel interference will be cancelled through a combination of analog EC and digital EC.

5. Echo cancellation

To ensure proper operation of FDX, the interference resulting from FDX operation must be suppressed/cancelled. As the interference always occur in the form of reflection, or echoes, the interference in FDX operation will be considered as echoes, and interference suppression/cancellation will be called echo cancellation.

Two types of EC techniques can be implemented in a RPD node to cancel or suppress the echoes.

5.1. Analog EC

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 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 (node output, first tap, second tap, etc.), multiple cancelling signals need be generated, one for each echo. All these need to be done in the analog domain.

As the echoes may come from taps that are located a few hundred feet away from the node, delays are needed to add into the cancelling signals. In coax, 1 ft. distance corresponds to ~1.2 ns delay in time, to cover the actual delay of all the echoes, a delay line with a variable delay of 1 ns~500 ns is required. Also, the bandwidth of this delay line needs to be >684 MHz. Such a delay line doesn't exist today.

The analog EC used in FDX DOCSIS 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) (Fig. 11). All the delays and magnitudes are computed and set in the digital domain through EC digital signal precessing (DSP).

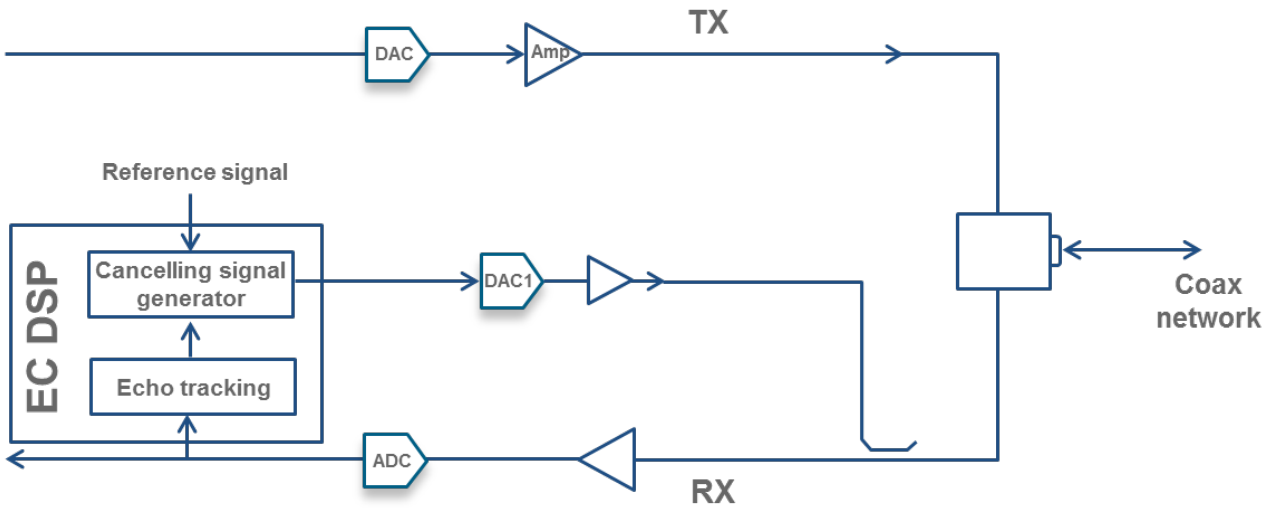


Figure 11 - Analog EC hardware architecture

5.2. 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 DS 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.

5.3. Reference signal for EC

The cancelling signal is generated from the DS signal with the proper magnitude and phase computed from the echoes embedded in the received signal. The DS 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 DS signal, just with various magnitudes and phases, depending on how the echoes are generated. The reference signal can be taken from DS data path in the digital domain, or taken from the output of the last stage amplifier if the noise generated from node launch amplifiers needs to be taken into account in the EC algorithm.

5.4. EC performance target

Analog EC and digital EC complement each other. One can partition the total EC performance target between analog EC and digital EC. For example, the analog EC can be designed to cancel out the echoes to the extent that the ADC dynamic range is not impaired, that is, the echoes are suppressed to a level that is below the desired US signal level, and the digital EC will cancel out all the echo residue to meet the final requirements on in-band MER. As indicated in Section 4.5.1, the echo could be 15 dB higher than the US desired signal in power, one could target the analog EC to have 15 dB echo cancellation, and so that the echo after analog EC will have the same power level as the desired US signal to minimize the

impact of the echo on receiver dynamic range. The digital EC then has 40 dB echo cancellation to further suppress the echo to 40 dB below the desired US signal to reach 40 dB inband MER.

5.5. EC coefficient training

To cancel out the echo, the canceling signal is generated from the reference signal. The cancelling signal needs to have proper magnitude and phase. The magnitude and phases of the cancelling signal are called EC coefficients. The EC coefficients are computed over a time period by comparing/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 which the EC is trained is called EC training period. There are two type of EC training: explicit training and implicit training.

5.5.1. Explicit training

Explicit training means there is a dedicated period of time when the EC is training. Normal system operation may be altered to facilitate the EC training. For example, the US traffic may be halted so the received signals are 100% echoes. This will help the EC coefficient computation/tracking algorithm to better compute the magnitude and phase differences between the reference signal and received echo without ‘interference’ by the US desired signal. Generally speaking, explicit training could lead to a simpler EC training algorithm and more accurate EC coefficients, but may impair normal system operation (for example, halt US traffic from all CMs)

5.5.2. Implicit training

Implicit training means the training is carried out without any explicit training period. With implicit training, the EC coefficient is computed/tracking in the background without impacting system operation. Implicit EC training has no impact on system operation. Implicit EC training needs to deal with the ‘interference’ of US traffic and usually involves a more sophisticated EC training algorithm and long training time to achieve the accuracy required.

6. Echo cancellation at CM

While most of the EC techniques explained in the previous sections can be used at both the RPD and CM, some changes or improvements to the EC techniques may be required when they are used at CM. At the node, all the echoes present as co-channel interference, while at the CM all the echoes present in a form of adjacent channel interference. This is because the RPD node the true FDX operation, that is, DS and US are completely overlapped at the RPD input/output port, but the CM still operates with frequency division duplex: its DS and US are not overlapped in frequency (ref [1]). The main reason behind this is to reduce the complexity of the CM. To support true full duplex operation, the CM receiver would need to have very high dynamic range, much higher than the RPD’s, as the echo would be much higher in power at the CM than RPD. The return loss of a CM F-connector is only 6 dB per specs vs. 23 dB return loss at the tap. Supporting such a high dynamic range results in a very expensive receiver and is cost prohibitive for CMs.

Although the CM supports only FDD, there is still interference that the CM needs to cancel out. It is the interference coming from adjacent channels. More specifically, there are two types of interference at the CM: adjacent leakage interference (ALI) and adjacent channel interference (ACI).

6.1. Adjacent leakage interference

ALI occurs when a CM transmits on one channel and receives on an adjacent channel. The out-of-band spurious emission of the transmitting channel leaks into the receiving channel and cause degradation of the received signal quality (Fig. 12).

6.2. Adjacent channel interference

ACI occurs when a CM transmits on one channel and receives on an adjacent channel. The power of the transmitting channel coupled back to the receiver and causes overload/saturation of the receiver front end.

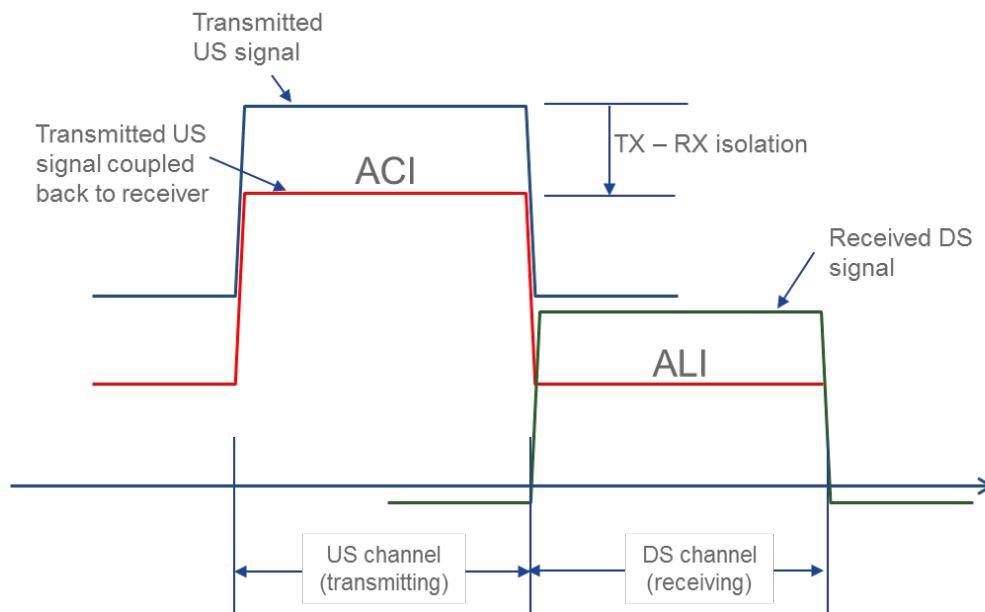


Figure 12 - Adjacent leakage interference and adjacent channel interference at CM

6.3. EC for ALI and ACI

As explained in Section 5.3, the theoretical base for EC is that all the echoes are true copies of a reference signal. The echoes can be expressed by the reference with EC coefficients in a linear space. The core algorithm of the EC is to find these coefficients, which is the same for EC used at the RPD and CM. However, as the echoes at CM all come in a form of adjacent channel interference, there are some unique requirements for EC at CM:

6.3.1. EC for ALI

To cancel out ALI at the CM, the reference used in EC must be the spurious noise generated by the transmitted signal on adjacent channels. This requires either using a reference from the output of the transmitter or predict the noise through an amplifier nonlinear model.

6.3.2. EC for ACI

As the ACI impacts on the receiver front end, the EC for ACI needs be implemented in the analog domain before the front end receiver AGC.

7. Summary on EC

In summary, EC is required at the RPD node to cancel out co-channel interference, and the EC can be implemented in the analog or digital domain or combination of both. While the analog EC must be implemented in the time domain, the digital EC can be implemented in the time domain or frequency domain or combination of both. EC is also required at the CM to cancel out ALI and ACI. The EC technique for ALI is similar to that used at the RPD node, but the difference is that the reference used to cancel out ALI at CM is the spurious emission of CM transmitter. EC for ACI needs be implement in the analog domain before receiver front end AGC.

8. Full duplex EC lab prototype system

A full duplex EC lab prototype system was built to validate the EC algorithm at the RPD node for supporting full duplex operation.

The system consists of three main components (Fig. 13):

1. The transceiver that emulates the RPD node:
 - a) It transmits three OFDM channels in the frequency range 108 MHz to 684 MHz
 - b) It receives six OFDMA channels in the frequency range 108 MHz to 684 MHz (overlap with the three DS OFDM channels)
 - c) It contains an EC function block sitting in front of the receiver demodulator
2. The transmitters that emulate CMs (CM1 and CM2, transmitting only)
 - a) CM1 transmits three OFDMA channels in the frequency range 108 MHz to 396 MHz
 - b) CM2 transmits three OFDMA channels in the frequency range 396 MHz to 684 MHz
3. The coax network between the emulated RPD and CMs.
 - a) Three sections of Series 6 cables, each section 100 ft. long.
 - b) The tap values are 26 dB (first tap), 14 dB (second tap) and 8 dB (last tap)
 - c) The CMs are connected to the last tap
 - d) The total loss between node and CM is ~35 dB.

DS traffic conforms to DOCSIS 3.1 DS frame, but the payload is random data. US traffic conforms to DOCSIS 3.1 US frame but the payload is random data.

The RPD node continuously transmits DS traffic with output TCP = 73.8 dBmV (63 dBmV between 108 MHz and 684 MHz), and at the same time receives US traffic from the CMs. US MER is computed from the constellation after the EC and channel estimation. For all six US OFDMA channels, >37 dB post-EC MER is achieved.

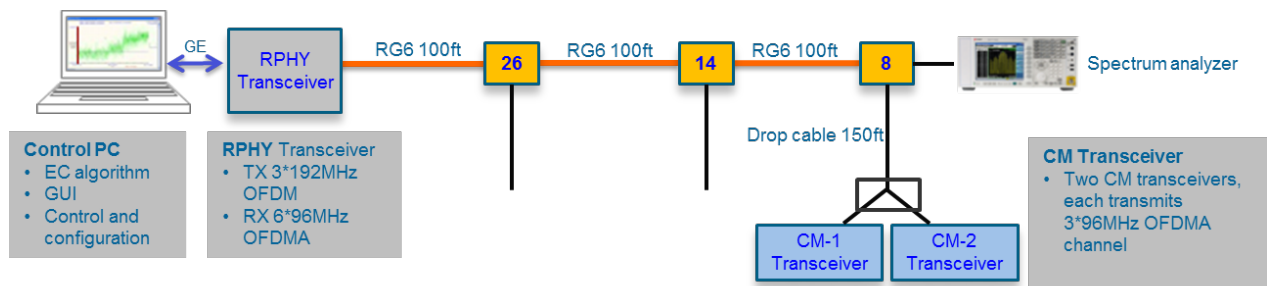


Figure 13 - FDX EC lab prototype system (block diagram)

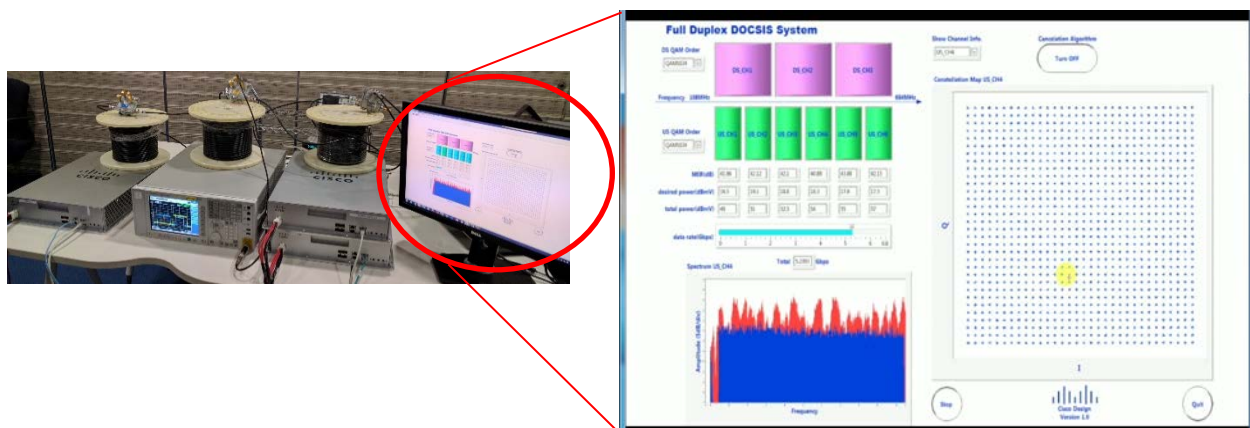


Figure 14 - FDX EC lab prototype system (Actual hardware)

9. FDX DOCSIS live demo system

Cisco partners with Intel and conducted a live demo of FDX DOCSIS at ANGAearlier this year. The system consisted of all essential boxes required in a real system: CMTS (cBR-8), R-PHY node, CMs and coax network to emulate N+0 network. The FDX operation used frequencies between 108 MHz and 204 MHz (one 96 MHz channel). CMs used in the demo were normal DOCSIS 3.1 CMs, and their US transmission frequency was limited to 204 MHz. One of the CMs was configured as low split, receiving DS traffic on 108 MHz to 204 MHz, and the other was configured as high split, transmitting US traffic on 108 MHz to 204 MHz. There was also one DS QAM channel at 300 MHz as the primary DS channel, and one US QAM channel at 8 MHz as the primary US channel.

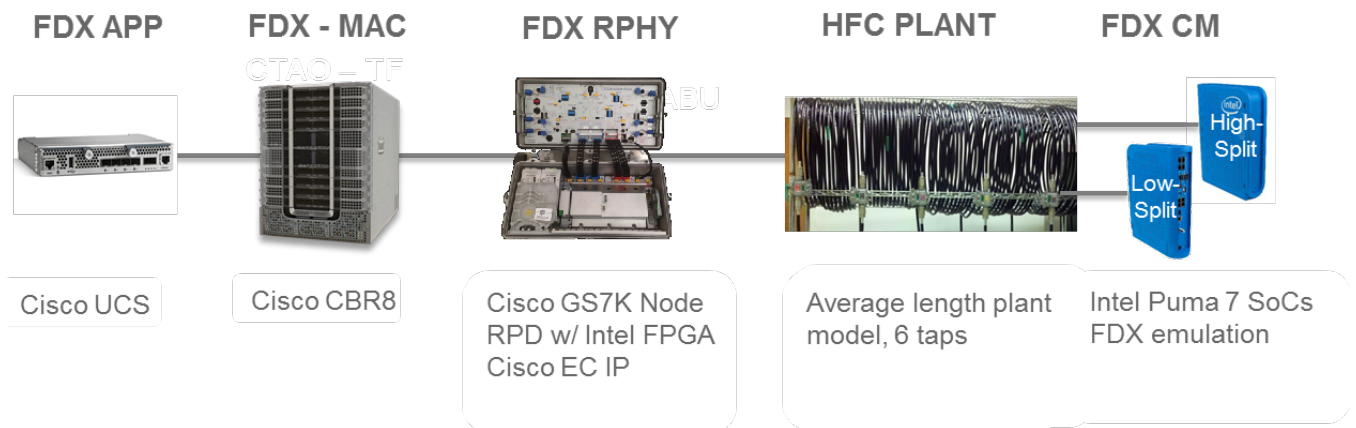


Figure 15 - FDX DOCSIS live demonstration system

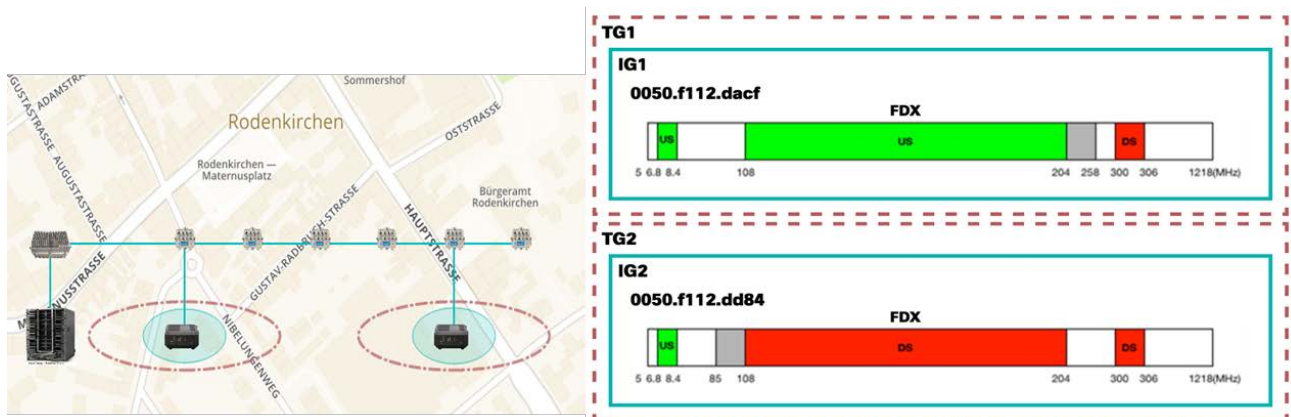



Figure 16 - FDX DOCSIS live demonstration system – network and channel configurations



L1=80ft; L2=80ft

Tap type: 4 ports, Cisco taps (Surge Gap taps, no equalizers)

All cables use RG6
(Comm/Scope RG-6, F6SSV
Super Shield drop cable)

Figure 17 - FDX DOCSIS live demonstration system – plant model

The DS supported 4096-QAM and the US supported 1024-QAM. The US pre-EC MER is -2 dB, and post-EC MER is 37 dB. With a 96 MHz FDX channel, DS throughput achieved ~940 Mbps, and US throughput achieved ~620 Mbps.




Figure 18 - FDX DOCSIS live demonstration system - measured echoes and post-EC constellation and MER.

Conclusion

Full duplex operation allows DS and US traffic to use the same RF spectrum at the same time, leading to ~100% increase in spectral efficiency and five to 10 times increase in US capacity. The key enabler for full duplex DOCSIS is echo cancellation. Echo cancellation can be implemented in time and/or frequency domains, and in analog and/or digital domain, depending on the system performance targets, hardware limitations and vendor preferences.

A full duplex EC lab prototype system was built in Cisco's lab to validate the EC algorithm at the RPD node for supporting full duplex operation. The system supported full duplex operation of three DS OFDM channels and six US OFDMA channels, for a total of 576 MHz of overlapping FDX spectrum. Over 45 dB echo suppression and >35 dB post-EC were achieved.

A live FDX DCOSIS PoC system was demonstrated at ANGA. The live demo system consisted of all essential boxes required in a real system: CMTS (cBR-8), R-PHY node, CMs and a coax network to emulate N+0 network. The FDX operation was running between 108 MHz and 204 MHz (one 96 MHz channel). 37 dB post-EC MER was achieved. The DS supported 4096-QAM and the US supported 1024-QAM. With a 96 MHz FDX channel, DS throughput achieved ~940 Mbps, and US throughput achieved ~620 Mbps.

Abbreviations

ACI	adjacent channel interference
ADC	analog-to-digital converter
ALI	adjacent leakage interference
CM	cable modem
CMTS	cable modem termination system
DAC	digital-to-analog converter
dB	decibel
dBmV	decibel millivolt
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	downstream
DSP	digital signal processing
EC	echo cancellation
FDD	frequency division duplex
FDX	full duplex
Gbps	gigabits per second
GHz	gigahertz
HFC	hybrid fiber/coax
ISBE	International Society of Broadband Experts
LPF	low pass filter
Mbps	megabits per second
MER	modulation error ratio
MHz	megahertz
ns	nanosecond

OFDM	orthogonal frequency division multiplex
OFDMA	orthogonal frequency division multiple access
OOB	out-of-band
PoC	proof-of-concept
QAM	quadrature amplitude modulation
RPD	remote PHY device
R-PHY	remote PHY
RX	1) receiver; 2) receive
SCTE	Society of Cable Telecommunications Engineers
SNR	signal-to-noise ratio
TCP	total composite power
TX	1) transmitter; 2) transmit
US	upstream

Bibliography & References

Full Duplex DOCSIS, John Chapman and Hang Jin, Spring Technical Forum, May 16-20, 2016, Internet and Television Expo (INTX), Cable Labs/NCTA/SCTE.

Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture

A Technical Paper prepared for SCTE•ISBE by

Richard S. Prodan, Ph.D.
Broadcom Fellow
Set-Top Box and Cable Modem Group
Broadcom Limited
1811 Pike Road, Suite 2C
Longmont, CO 80501
720-864-4227
rich.prodan@broadcom.com

Introduction

DOCSIS 3.1 has significantly increased the bandwidth utilization and capacity flexibility of cable data service through the use of new physical layer modulation based on OFDM/OFDMA, as well as extending bandwidth in both upstream and downstream transmission bands. DOCSIS 3.1 and prior versions were concerned with forward and reverse path loss and minimum power levels to achieve acceptable signal-to-noise ratio (SNR) for good spectral efficiency. Only unidirectional frequency response limits need be considered and reflected signals back toward the source can be reasonably ignored.

The addition of full duplex (FDX) simultaneous transmission and reception within the same spectrum introduces numerous additional considerations that can be neglected in prior DOCSIS frequency division duplex (FDD) versions. Such concurrent full duplex transmit/receive operation in the same spectrum introduces interference into and reflections back toward the transmission source that must be effectively canceled for simultaneous reception of signals traveling in the opposite direction within the same frequency band.

This paper analyzes the additional considerations for full duplex operation within a fiber deep (passive node plus zero amplifier) required cable system architecture and the resulting expected SNR and spectral efficiency impacts both at the node and the cable modem. These include:

- Signal levels and path loss over the fiber deep coax plant and within the node and modem
- Tap signal conditioning (equalization) over wide bandwidths (up to 1218 MHz)
- Noise and interference sources and levels both across different taps and across ports on the same tap including co-channel interference, adjacent channel interference, and adjacent channel leakage interference from transmitter noise and spurious emissions
- Transmitted signal reflection levels impacting received signals within the same or adjacent frequency bands
- Interference due to limited isolation between transmitted and received signals both at the fiber node and the cable modem
- The segmentation of cable modems into “interference groups”
- SNR and spectral efficiency (bit-loading) achievable within such cable modem groupings.

The Fiber Deep Cable Architecture

1. Cable Network Characteristics and Design

Traditional hybrid fiber/coax (HFC) cable distribution networks have been built as tree and branch networks consisting of a fiber node connecting multiple cascaded amplifier coax cable sections. Each section connects to a series of multiport taps transmitting signal to and receiving signals from drop cables to customer premise equipment. An example of one coax branch of a conventional Node + N HFC architecture shown in Figure 1. The node span contains multiple amplifier spans, each with multiple taps between amplifiers.




Figure 1 - Conventional node plus N network architecture

This conventional architecture provides two-way signal transmission on separate spectral bands using frequency division duplex operation. Each amplifier section uses diplex filtering to separate upstream transmissions toward the node in the narrower lower frequency band (typically 85 MHz or less) from downstream transmissions from the node in the much wider upper frequency band (up to 1 GHz). Such diplex filtering prevents two-way transmission within the same bandwidth. Each multiport tap contains a directional coupler that diverts a portion of the downstream signal to the drops connected to the tap ports and injects the upstream signals present on the tap ports toward the node. The directivity of the directional coupler prevents upstream signals from propagating in the downstream direction or from diverting to other drops upstream from that tap port.

The full duplex architecture provides two-way signal transmission within the same spectral band. This requires a passive architecture without amplifiers. In this case, the fiber node connects to a single series of multiport taps. Without any amplifiers that require diplex filtering, both upstream and downstream signals can share the part of the same spectrum but with the same directivity of the conventional architecture. The bandwidth supported in this passive fiber deep architecture can be wider than the conventional cascaded amplifier architecture (1218 MHz or more). An example of one coaxial branch of a fiber deep node + 0 architecture is shown in Figure 2.

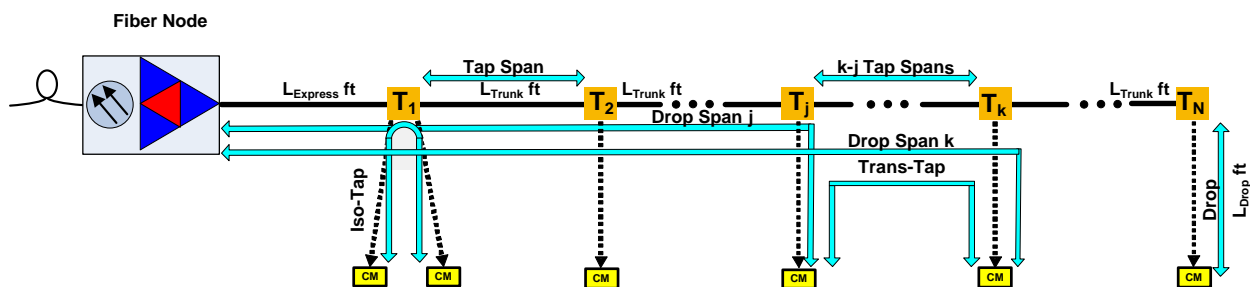


Figure 2 - The fiber deep node + 0 architecture

The design of a conventional or fiber deep network requires specifying the individual components and their values to obtain consistent signal levels across all the subscriber drops. These include:

- hardline cable (trunk and express feeder) and drop cable type and parameters
- tap parameters (insertion loss, tap loss, return loss, isolation)
- downstream and upstream node and tap equalization type, frequency range, and tilt value
- tap spacing
- drop cable length
- amplifier output (level and tilt)

Typical hardline and drop cable specifications include attenuation per length versus frequency and velocity of propagation are shown in Table 1.

Frequency (MHz)	Cable Attenuation (dB/100 ft)
5	0.14
55	0.48
83	0.58
85	0.59
204	0.93
211	0.95
250	1.03
300	1.13
350	1.23
400	1.32
450	1.40
500	1.49
550	1.56
600	1.64
750	1.85
865	2.00
1000	2.17
1002	2.16
1218	2.41

Velocity of Propagation:	0.87
Cable Span Delay (μsec):	0.2045093
Linear Regression Parameters	
Slope	0.1219285
Intercept	-0.0991351




Table 1 - QR 540 hardline cable parameters

Note the linear attenuation versus square root frequency characteristic typical of coaxial cable. This useful property simplifies modeling the behavior of signal attenuation of the interconnecting hardline cable between taps or between a tap port and the terminating device on the drop cable (see Appendix 0).

Tap parameters including insertion loss, tap loss, return loss, and tap-to-output isolation are also specified as attenuation versus frequency. Table 2 shows an example of such tap specifications.

2-Way Maximum Insertion Loss (dB)											
Number of Ports (#-Way), Tap Value (dB)											
Frequency ≤(MHz)	2,4	2,7	2,10	2,12	2,14	2,17	2,20	2,23	2,26	2,29	
5		3.9	2	1.8	1.3	1.2	1	0.7	0.7	0.7	
10		3.6	1.8	1.6	1.2	1.1	0.9	0.6	0.6	0.6	
50		3.5	1.8	1.6	1.2	1.1	0.9	0.7	0.7	0.7	
100		3.6	2	1.8	1.3	1.2	1	0.8	0.8	0.8	
450		4.2	2.4	2.1	1.6	1.5	1.3	1.1	1.1	1.1	
550		4.2	2.4	2.3	1.8	1.5	1.4	1.2	1.2	1.2	
750		4.2	2.4	2.4	1.8	1.6	1.5	1.3	1.3	1.3	
870		4.3	2.7	2.7	2	1.6	1.6	1.4	1.4	1.4	
1000		4.3	3	3	2.3	1.6	1.6	1.6	1.6	1.6	
1218		5.2	3.6	3.6	2.7	2.2	2	2	2	2	

Minimum Return Loss (dB)			
For ALL Tap Values			
Frequency ≤(MHz)	Minimum Return Loss In-Out (dB)	Minimum Return Loss Tap (dB)	
5	16	16	16
10	16	16	17
750	16	16	17
1000	16	16	17
1218	16	16	16

2-Way Nominal Tap Value (dB)											
Number of Ports (#-Way), Tap Value (dB)											
Frequency ≤(MHz)	2,4	2,7	2,10	2,12	2,14	2,17	2,20	2,23	2,26	2,29	
5	4	7	10	12	14	17	20	23	26	29	
1000	4	7	10	12	14	17	20	23	26	29	
1218	4	7	10	12	14	17	20	23	26	29	

Tap Value Tolerance ±(dB)			
Number of Ports (#-Way)			
2	4	8	
2	2	2	3
2	2	2	3
2.5	2.5	3.5	

2-Way Tap-to-Output Isolation (dB)											
Number of Ports (#-Way), Tap Value (dB)											
Frequency ≤(MHz)	2,4	2,7	2,10	2,12	2,14	2,17	2,20	2,23	2,26	2,29	
5		20	20	22	22	26	29	32	35	38	
10		20	20	22	22	26	29	32	35	38	
85		25	25	25	26	30	33	36	38	40	
300		21	22	23	26	30	33	36	38	40	
750		22	22	23	26	30	31	34	36	39	
900		20	20	22	23	28	30	33	35	37	
1218		20	20	20	22	25	29	31	33	35	

Tap-to-Tap Isolation (dB)	
For ALL Tap Values	
Frequency ≤(MHz)	Isolation (dB)
5	20
10	25
85	27
300	27
750	23
1218	20

Table 2 - Tap specifications for insertion loss, tap value loss, return loss, and tap-to-output isolation versus frequency for each tap value

An important consideration for supporting the wider bandwidth on the passive fiber deep plant is the use of tap plug-in equalizers. Such equalizers provide either upward or downward attenuation (tilt) with increasing frequency. The amount for each tap is determined to provide approximately the same output levels versus frequency at all tap ports. An example of tap equalizer magnitude frequency (amplitude) responses is shown in Figure 3.




Figure 3 - Tap equalization (“signal conditioning”) plug-in characteristics

The design of the fiber deep coaxial cable plant requires specification of signal levels versus frequency, selection of cable lengths and types, and tap values including equalization appropriate for the inter-tap spacing and drop cable length to provide uniform signal levels across all tap ports. An example of fiber deep signal levels and limits is given in Table 3.

Node Signal Levels (21 dB linear up-tilt from 111 to 1215 MHz)		Tap Specifications	
Transmit Power/6 MHz	Receive Power/6.4 MHz	0' to 50' RG-6 short drop model for 256-QAM	
37 dBmV (111 MHz)	8 dBmV (5 to 85 MHz@node;	Minimum Level Low Frequency	4 dBmV@111 MHz
58 dBmV (1215 MHz)	≤32 dB loss@85 MHz to tap port)	Minimum Level High Frequency	12 dBmV@1218 MHz
		Maximum Tilt	10 dB Minimum Tilt 6 dB
		Maximum Return at Tap Port	40 dBmV@85 MHz
Subscriber Signal Levels		Tap Specifications	
Transmit Power/6.4 MHz	Receive Power/6 MHz	51' to 100' RG-6 standard drop model for 256-QAM	
40 dBmV (5 to 85 MHz@tap max)	-6 dBmV (min@4 outlets)	Minimum Level Low Frequency	5 dBmV@111 MHz
	4 dBmV (111 MHz@tap min)*	Minimum Level High Frequency	15.5 dBmV@1218 MHz
	6 dBmV (111 MHz@tap max)*	Maximum Tilt	12 dB Minimum Tilt 8 dB
	12 dBmV (1218 MHz@tap min)*	Maximum Return at Tap Port	40 dBmV@85 MHz
	19 dBmV (1218 MHz@tap max)*	Tap Specifications	
		101' to 150' RG-6 long drop model for 256-QAM	
		Minimum Level Low Frequency	6 dBmV@111 MHz
		Minimum Level High Frequency	19 dBmV@1218 MHz
		Maximum Tilt	15 dB Minimum Tilt 11 dB
		Maximum Return at Tap Port	40 dBmV@85 MHz
		Tap Specifications	
		151' to 200' RG-11 extra-long drop model for 256-QAM	
		Minimum Level Low Frequency	6 dBmV@111 MHz
		Minimum Level High Frequency	18 dBmV@1218 MHz
		Maximum Tilt	14 dB Minimum Tilt 10 dB
		Maximum Return at Tap Port	40 dBmV@85 MHz

*Absolute limit but actual tap port level depends on drop length

Table 3 - Fiber deep node, tap, and subscriber RF signal levels and limits

Analysis of Fiber Deep Networks

2. Fiber Deep Transmission Model and Analysis

An approach to modeling signal transmission through the various elements of the fiber deep network is discussed in this section. Analysis of the frequency response including signal levels transmitted from the node and received by the cable modem or vice versa is described. This approach allows the use of log magnitude values versus frequency from component spec sheets to model the transfer function of each section of transmission line (i.e., cable) terminated at either end with an impedance (i.e., a tap, node, or cable modem) characterized by its magnitude return loss versus frequency.

Consider the fiber deep architecture model depicted in Figure 4. The fundamental transmission transfer function block is denoted by the tap span shown in this figure.




Figure 4 - Fiber Deep model for Full Duplex operation

It is shown in Appendix 1 for a length of trunk cable with cable propagation delay T , cable amplitude response $A(f)$, and tap input/output port return loss $RL = -10 \log(\rho)$ where ρ = reflection coefficient, that the transfer function $H(f)$ for the tap span cable transmission line is given by:

$$H(f) = \frac{A(f)}{1 - A^2(f) 10^{-\frac{(RL_i + RL_o)}{20}} e^{-j4\pi f T}}$$

Note that this formulation provides the complex frequency response with only the scalar amplitude versus frequency of the cable transmission line and the magnitude return loss versus frequency of the tap terminating impedances. This avoids the need to measure complex valued s-parameters versus frequency for each component and convert to t-parameters for transmission frequency responses that can be cascaded to compute the end-to-end transfer function between any two points in the network.

The same analysis applies to a drop cable section between tap port and the cable modem where the attenuation model and propagation delay are specified for the drop cable instead of the hardline cable and the input and output return losses are specified for the tap port and cable modem F-connector port respectively.

Denote the drop frequency response as $H_{drop}(f)$ and the hardline trunk feeder section frequency response by $H_{trunk}(f)$. The return path frequency response $H_n(f)$ for a drop on the n^{th} tap after the amplifier with tap loss value T_n and insertion losses $I_{n-1}, I_{n-2}, \dots, I_1$ (all tap loss amplitudes are a nearly constant function of frequency by design) is simply the product of the frequency responses of each section given by:

$$H_n(f) = T_n I_{n-1} I_{n-2} \dots I_1 H_{drop}(f) [H_{trunk}(f)]^{n-1}$$

This approach is used to model the transfer function between the node and any cable modem or between cable modems on the same or different taps.

3. Fiber Deep Network Example Design

An example design of suburban cable plant with standard 100 foot length drops is analyzed next using the approach described in the previous section. The plant model (Model 1) is shown in Figure 5.

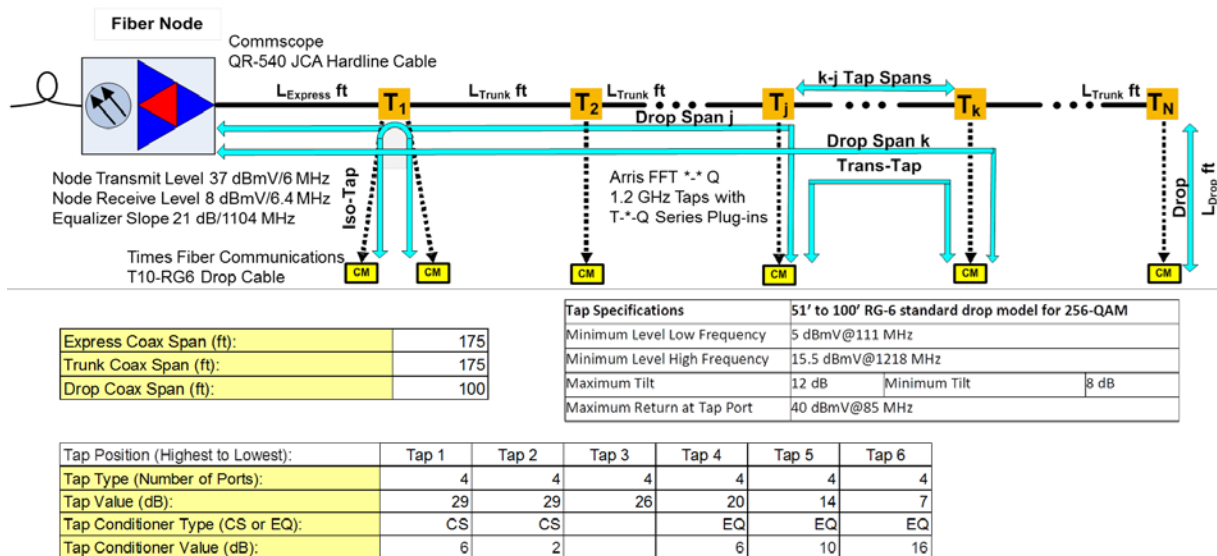


Figure 5 - Model 1 – suburban plant and standard drops with fiber deep signal levels

An analysis of received downstream signal levels at the tap ports and cable modems attached to the drops is provided in this section with the methodology described previously using the node signal launch levels and up-tilt, tap and plug-in equalizer characteristics, and cable attenuation versus frequency for the hardline and drop cable types. The resulting tap port received levels are plotted in Figure 6.




Figure 6 - Receive power levels for all taps in Model 1

Note that the tap port levels for each tap are approximately the same and compliant within a dB with the fiber deep signal levels and limits given in Table 3 for standard length 100 foot drops. This is the result of design selection of appropriate tap values and tap plug-in equalizers for the cable lengths and types.

The resulting cable modem received levels at the end of the drops from each tap are plotted in Figure 7. An expanded view of the received power of the cable modem on tap 1 shows the detailed amplitude response due to micro-reflections from finite return losses in the tap ports and cable modem F-connector.




Figure 7 - Receive power levels for modems on each tap drop

Due to the symmetry of the passive cable plant, the same calculated frequency responses from the node to each tap port can be used to evaluate the cable modem transmit levels at each tap drop. However, the received level at the node port is specified to be a flat power spectral density across the 108 MHz to 684 MHz full-duplex band. The value of that power spectral density level is determined such that the total composite power transmitted in the full-duplex band by any cable modem in the network does not exceed 64.5 dBmV. In this design example, it will be shown later in the node port analysis that received upstream power spectral density level is approximately 5 dBmV/6.4 MHz. The resulting cable modem transmit levels for each tap are shown in Figure 8.

Note that the transmitted signal levels are approximately the same for all cable modems across all taps. As will be shown later for the node, the tilt in the upstream transmissions from each cable modem results in the received flat power spectral density level at the node port noted previously.




Figure 8 - Modem port transmit levels at the drop for each tap

Full Duplex Network Operation

4. Full Duplex Interference Groups

Simultaneous true full duplex transmission and reception within the same frequency band only occurs at the node. Each cable modem utilizes FDD in a dynamic fashion across multiple channels within the full-duplex band. This is necessary to prevent upstream transmission of one cable modem in a given channel from corrupting downstream reception of another cable modem in that same channel. Unlike the node where the transmitter and receiver are co-located, upstream transmitting modem and downstream receiving modem are located on different tap ports of the same tap or different taps. Thus, the downstream receiving cable modem has no reference for the upstream transmitting cable modem signal making cancellation within the same frequency band impossible.

It will be shown that cable modems on the same tap in this situation can only either transmit or receive in a given channel, but never simultaneously. However, if the transmitting and receiving modems are separated across different taps, then simultaneous transmission by one modem and reception by another modem on a different tap with sufficient separation to provide isolation between the signals sharing the same frequencies in a given channel is possible. This segregating of each modem into a group where modems in each individual group mutually interfere while other modems in a different group do not mutually interfere significantly is known as an interference group (IG).

Figure 9 shows the dynamic FDD upstream transmit and downstream receive channel allocations within a single IG.




Figure 9 - Time variability of FDD transmission and reception in the FDX spectrum for a single interference group

All modems in the same IG see the same up/down resource block assignments (RBA) which are the fixed channel assignments to upstream transmission or downstream reception. Modems in different IG's may get the same or (usually) different RBAs. The collection of all IG's within the shared full-duplex spectrum is known as a transmission group (TG).

The TG at the node receives signals from all modems across all channels in every IG resulting in true full-duplex spectral utilization at the node. However, each modem uses dynamic FDD as shown in the example of Figure 9. Several interference mechanisms and their impact on spectral efficiency (bit-loading) is examined next.

5. Full Duplex Interference Group Analysis of Co-channel Interference

Referring to the fiber deep architecture of Model 1 in Figure 5, modem-to-modem co-channel interference (CCI) is examined in the following. Consider a modem on tap j which transmits in each FDX channel. The worst case interference seen at the receiver of a modem on tap k relative to its downstream receive signal level on each FDX channel is shown in the green line of Figure 10.




Figure 10 - Modem-to-modem co-channel interference

The interference level depicted in this figure is representative of modems on the same tap or when j equals k . Modems that are more widely separated across different taps will experience considerably lower levels of CCI. The downstream signal to upstream CCI SNR versus frequency between two modems with the widest separation of taps is shown in Figure 11 using the calculation method of Appendix 3: SNR Calculations of CCI.




Figure 11 - Example of modem CCI SNR across interference groups

Note the symmetry in either transmission direction due to the symmetric frequency response of the passive coaxial network. The SNR at each frequency can be compared to the threshold value for DOCSIS 3.1 OFDM modulation efficiency (bits/subcarrier) to determine the QAM order that can be supported in a given FDX channel as shown in Table 4.

QAM Order	Modulation Efficiency (bits/subc)	Spectral Efficiency (bits/subc)	FEC SNR Threshold (dB)	CNR Threshold (dB)	Channel Capacity (bits/subc)
0	0.0	0.00	-100.0	-100.0	0.00
QPSK	2.0	1.76	7.5	9.0	2.73
16-QAM	4.0	3.51	13.0	15.0	4.39
64-QAM	6.0	5.27	18.6	21.0	6.20
64/128-QAM	6.5	5.71	20.4	22.5	6.79
128-QAM	7.0	6.15	21.4	24.0	7.12
128/256-QAM	7.5	6.59	23.3	25.5	7.75
256-QAM	8.0	7.03	24.2	27.0	8.04
256/512-QAM	8.5	7.47	26.0	28.7	8.64
512-QAM	9.0	7.91	26.9	30.5	8.94
512/1024-QAM	9.5	8.35	28.7	32.2	9.54
1024-QAM	10.0	8.79	29.7	34.0	9.87
1024/2048-QAM	10.5	9.22	31.6	35.5	10.50
2048-QAM	11.0	9.66	32.4	37.0	10.76
2048/4096-QAM	11.5	10.10	34.2	39.0	11.36
4096-QAM	12.0	10.54	35.2	41.0	11.69

Table 4 - Downstream bit-loading table for calculation of modulation efficiency

Modulation efficiency represents the OFDM subcarrier bit-loading as a function of CNR threshold. Calculate bit-loading using Table 4 of modulation efficiency vs. CNR at each frequency and average bit-loading per frequency over a frequency band for calculating average bit-loading in a channel. The results for average CCI SNR and bit loading for Model 1 is shown in Table 5. Note that a modem transmitting on the same tap as other receiving modems on the same tap results in a negative CCI SNR. Downstream reception in the same channel as the upstream transmission is not possible in this case.

R E C E I V E	SNR (dB)	SNR in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	-4.7	40.4	40.4	40.4	40.4	40.4
	Tap 2	40.4	-7.2	36.2	36.2	36.2	36.2
	Tap 3	40.4	36.2	-5.4	32.2	32.2	32.2
	Tap 4	40.4	36.2	32.2	-7.3	24.0	24.0
	Tap 5	40.4	36.2	32.2	24.0	-5.7	13.7
	Tap 6	40.4	36.2	32.2	24.0	13.7	-7.2

R E C E I V E	Bit-Loading (bits/subc)	Bit-Loading in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	0.0	11.6	11.6	11.6	11.6	11.6
	Tap 2	11.6	0.0	10.5	10.5	10.5	10.5
	Tap 3	11.6	10.5	0.0	9.3	9.3	9.3
	Tap 4	11.6	10.5	9.3	0.0	6.8	6.8
	Tap 5	11.6	10.5	9.3	6.8	0.0	2.4
	Tap 6	11.6	10.5	9.3	6.8	2.4	0.0

Table 5 - Average CCI SNR and bit loading for Model 1

The SNR table values are estimated in the “sounding” process. A modem transmits a reference signal (CW tones or an upstream data profile testing burst within an OFDMA Upstream Data Profile (OUDP) interval usage code (IUC) grant). All other modems receive the reference signal plus the downstream signal containing zero bit-loaded (ZBL) subcarriers or symbols depending on the reference signal chosen. The modulation error ratio (MER) is measured over a sufficient number of OFDM symbols. Each modem is assigned to an IG determined by the sounding MER results (“IG discovery”).

An example of using this sounding table to assign modems on different taps to interference groups is shown in Table 6.

R E C E I V E	SNR (dB)	SNR in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
Tap 1	-4.7	40.4	40.4	40.4	40.4	40.4	40.4
Tap 2	40.4	-7.2	36.2	36.2	36.2	36.2	36.2
Tap 3	40.4	36.2	-5.4	32.2	32.2	32.2	32.2
Tap 4	40.4	36.2	32.2	-7.3	24.0	24.0	24.0
Tap 5	40.4	36.2	32.2	24.0	-5.7	13.7	13.7
Tap 6	40.4	36.2	32.2	24.0	13.7	-7.2	-7.2

R E C E I V E	Bit-Loading (bits/subc)	Bit-Loading in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
Tap 1	0.0	11.6	11.6	11.6	11.6	11.6	11.6
Tap 2	11.6	0.0	10.5	10.5	10.5	10.5	10.5
Tap 3	11.6	10.5	0.0	9.3	9.3	9.3	9.3
Tap 4	11.6	10.5	9.3	0.0	6.8	6.8	6.8
Tap 5	11.6	10.5	9.3	6.8	0.0	2.4	2.4
Tap 6	11.6	10.5	9.3	6.8	2.4	0.0	0.0

OFDM Modulation:	4096-QAM	2048-QAM	1024-QAM	512-QAM
SNR Threshold:	41 dB	37 dB	34 dB	30.5 dB

Interference Group partitioning:

- **IG 0** SNR ≥ 41 dB \leftrightarrow (no Taps)
- **IG 1** SNR > 37 dB \leftrightarrow (Tap 1)
- **IG 2** SNR > 34 dB \leftrightarrow (Tap 2)
- **IG 3** SNR > 30.5 dB \leftrightarrow (Tap 3)
- **IG 4** SNR ≤ 30.5 dB \leftrightarrow (Taps 4, 5, 6)

Interference Group resulting bit-loading:

- IG 0** (no Taps) receive 4096-QAM (12 bits/subcarrier)
- IG 1** (Tap 1) receives 2048-QAM (11 bits/subcarrier) when IG 2, 3, or 4 transmit
- IG 2** (Tap 2) receives 1024-QAM (10 bits/subcarrier) when IG 1, 3, or 4 transmit
- IG 3** (Tap 3) receives 512-QAM (9 bits/subcarrier) when IG 1, 2, or 4 transmit
- IG 4** (Taps 4, 5, 6) receives 512-QAM (9 bits/subcarrier) when IG 1, 2, or 3 transmit

Table 6 - CCI interference group possible assignments for Model 1

Assignment of a modem to an IG can be done by successively checking to see if the measured MER is above a threshold SNR for a given bit loading starting at the highest SNR threshold. If so, then the modem is assigned to this highest IG associated with that threshold SNR (IG 5). If not, then the measured MER is checked to be above the next lower threshold SNR. If so, then the modem is assigned to this next highest IG associated with that next lower threshold. If not, this process continues until the MER is above the given threshold for an IG. Finally if the MER is below the SNR for the lowest IG, then the CM is assigned to this lowest IG.

The results of this sorting process of cable modems into IG's is shown in Table 6. No modems were found to be above the threshold for the 4096-QAM IG. The modems on tap 1 are above the threshold for 2048-QAM when modems on taps 2 through 6 transmit and are therefore assigned to this next highest IG 4. The modems on tap 2 are above the threshold for 1024-QAM when modems on taps 3 through 6 transmit and are therefore assigned to this next highest IG 3. The modems on tap 3 are above the threshold for 512-QAM when modems on taps 4 through 6 transmit and are therefore assigned to this next highest IG 2. The modems on taps 4 through 6 are below the lowest threshold for 512-QAM when modems on taps 4 through 6 transmit and are therefore assigned to this lowest IG.

In summary, modem-to-modem CCI is the primary impairment that limits spectral efficiency (i.e., maximum throughput) for a full duplex modem. CCI between different taps increases with increasing distance from the node and decreasing proximity between transmitting and receiving modems. SNR is reduced for modems further from the node. SNR is reduced with decreasing distance (fewer number of intervening taps) between transmitting and receiving modems. CMs on the same tap will have negative CCI SNR and will be in the same interference group.

6. Full Duplex Interference Group Analysis of Adjacent Leakage Interference

Again, referring to the fiber deep architecture of Model 1 in Figure 5, modem-to-modem adjacent leakage interference (ALI) will be examined in the following. Consider a modem on tap j which transmits in the two highest frequency FDX channels. The worst case interference due to out-of-band noise and spurious emissions seen at the receiver of a modem on tap k relative to its downstream receive signal level on the lowest frequency FDX channel is shown in the red line of Figure 12.




Figure 12 - Cable modem-to-modem adjacent leakage interference

The interference level depicted in this figure is representative of modems on the same tap or when j equals k . Modems that are more widely separated across different taps will experience considerably lower levels of ALI. It will be shown that ALI from modems across different taps is insignificant.

The modem spurious emissions transmit mask for the FDX spectrum is shown in Figure 13.




Figure 13 - Modem spurious emissions with 570 MHz FDX active spectrum

The upstream transmit signal is shown for the maximum 576 MHz grant and 10 dB up-tilt with the maximum 64.5 dBmV total composite power. The spurious emissions limit has a linear taper from -39 dBr at 108 MHz to -44 dBr at 684 MHz. Spurious emissions level at modem port with this transmit up-tilt is adjusted by $10 \cdot \log_{10}(\text{upstream grant power} / \text{total composite power})$ for grants less than the full FDX band down to the under grant hold bandwidth of $1/6^{\text{th}}$ of the modulated spectrum width.

The above spurious emissions mask is used to calculate ALI power adjusted by the grant size. The modem SNR is then calculated with the ALI power that migrates across ports of the same tap using the calculation method of Appendix 4: SNR Calculations of ALI. The results are shown in Figure 14 for the first and last tap.




Figure 14 - ALI SNR at modems across ports of the same tap

R E C E I V E	SNR (dB)	SNR in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
R E C E I V E	Tap 1	43.3	87.8	86.6	88.0	87.2	88.2
	Tap 2	86.5	42.0	83.0	84.4	83.6	84.6
	Tap 3	86.5	84.2	43.1	81.2	80.4	81.3
	Tap 4	86.6	84.3	79.9	41.8	72.4	73.3
	Tap 5	86.6	84.3	79.9	73.2	42.6	63.6
	Tap 6	86.6	84.3	80.0	73.3	62.7	41.7

R E C E I V E	Bit-Loading (bits/subc)	Bit-Loading in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
R E C E I V E	Tap 1	11.9	12.0	12.0	12.0	12.0	12.0
	Tap 2	12.0	11.8	12.0	12.0	12.0	12.0
	Tap 3	12.0	12.0	11.9	12.0	12.0	12.0
	Tap 4	12.0	12.0	12.0	11.7	12.0	12.0
	Tap 5	12.0	12.0	12.0	12.0	11.8	12.0
	Tap 6	12.0	12.0	12.0	12.0	12.0	11.7

Table 7 - average ALI SNR and bit loading for Model 1

The average ALI SNR and bit loading for model one is tabulated in Table 7. Comparing to the average CCI SNR and bit loading in Table 5, it can be seen that iso-tap ALI degradation of SNR given along the diagonal of the table is comparable in the lowest frequency channel and several dB lower on average over all FDX channels. Also note that ALI SNR across different taps given by the off diagonal table entries is negligible. The SNR impact on combining these impairments is calculated in the next section.

7. Full Duplex Interference Group Analysis of Combined Interference

The combined trans-tap CCI plus iso-tap ALI SNR at modems across interference groups is plotted in Figure 15.




Figure 15 - Trans-tap CCI plus iso-tap ALI SNR that modems across interference groups

The average trans-tap CCI plus iso-tap ALI SNR and bit loading for Model 1 are tabulated in Table 8.

R E C E I V E	SNR (dB)	SNR in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
R E C E I V E	Tap 1	-4.7	38.6	38.6	38.6	38.6	38.6
	Tap 2	38.1	-7.2	35.2	35.2	35.2	35.2
	Tap 3	38.5	35.4	-5.4	31.9	31.9	31.9
	Tap 4	38.0	35.2	31.8	-7.3	23.9	23.9
	Tap 5	38.3	35.3	31.8	23.9	-5.7	13.7
	Tap 6	38.0	35.1	31.7	23.9	13.7	-7.2

R E C E I V E	Bit-Loading (bits/subc)	Bit-Loading in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
R E C E I V E	Tap 1	0.0	11.2	11.2	11.2	11.2	11.2
	Tap 2	11.1	0.0	10.2	10.2	10.2	10.2
	Tap 3	11.2	10.2	0.0	9.1	9.1	9.1
	Tap 4	11.1	10.2	9.1	0.0	6.8	6.8
	Tap 5	11.2	10.2	9.1	6.8	0.0	2.4
	Tap 6	11.1	10.2	9.1	6.8	2.4	0.0

Table 8 - Average trans-tap CCI plus iso-tap ALI SNR and bit loading for Model 1

This table of combined interference shows about a 2 dB degradation in SNR. This is less than 1 bit/subcarrier spectral efficiency loss if the resulting SNR crosses the CNR threshold for an interference group. The combined average trans-tap CCI plus iso-tap ALI interference groups are recalculated in Table 9. Note that the additional impact of iso-tap ALI on interference group SNR is typically less than 2 dB or less than 0.5 bits/subcarrier bit loading.

R E C E I V E	SNR (dB)	SNR in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	-4.7	38.6	38.6	38.6	38.6	38.6
	Tap 2	38.1	-7.2	35.2	35.2	35.2	35.2
	Tap 3	38.5	35.4	-5.4	31.9	31.9	31.9
	Tap 4	38.0	35.2	31.8	-7.3	23.9	23.9
	Tap 5	38.3	35.3	31.8	23.9	-5.7	13.7
	Tap 6	38.0	35.1	31.7	23.9	13.7	-7.2

R E C E I V E	Bit-Loading (bits/subc)	Bit-Loading in Total Channel (108 MHz to 684 MHz)					
		T R A N S M I T					
		Tap 1	Tap 2	Tap 3	Tap 4	Tap 5	Tap 6
	Tap 1	0.0	11.2	11.2	11.2	11.2	11.2
	Tap 2	11.1	0.0	10.2	10.2	10.2	10.2
	Tap 3	11.2	10.2	0.0	9.1	9.1	9.1
	Tap 4	11.1	10.2	9.1	0.0	6.8	6.8
	Tap 5	11.2	10.2	9.1	6.8	0.0	2.4
	Tap 6	11.1	10.2	9.1	6.8	2.4	0.0

OFDM Modulation:	4096-QAM	2048-QAM	1024-QAM	512-QAM
SNR Threshold:	41 dB	37 dB	34 dB	30.5 dB

Interference Group partitioning:

- **IG 0** SNR ≥ 41 dB \leftrightarrow (no Taps)
- **IG 1** SNR > 37 dB \leftrightarrow (Tap 1)
- **IG 2** SNR > 34 dB \leftrightarrow (Tap 2)
- **IG 3** SNR > 30.5 dB \leftrightarrow (Tap 3)
- **IG 4** SNR ≤ 30.5 dB \leftrightarrow (Taps 4, 5, 6)

Interference Group resulting bit-loading:

- **IG 0** (no Taps) receive 4096-QAM (12 bits/subcarrier)
- **IG 1** (Tap 1) receives 2048-QAM (11 bits/subcarrier) when IG 2, 3, or 4 transmit
- **IG 2** (Tap 2) receives 1024-QAM (10 bits/subcarrier) when IG 1, 3, or 4 transmit
- **IG 3** (Tap 3) receives 512-QAM (9 bits/subcarrier) when IG 1, 2, or 4 transmit
- **IG 4** (Taps 4, 5, 6) receives 512-QAM (9 bits/subcarrier) when IG 1, 2, or 3 transmit

Table 9 - Combined trans-tap CCI plus iso-tap ALI interference groups for Model 1

In summary, analysis for the increased degradation of trans-tap average CCI by iso-tap average ALI shows the following trends:

- Iso-tap ALI across ports of same tap adds slight degradation to trans-tap CCI across IGs
- Iso-tap ALI across ports of same tap adds slight degradation to trans-tap CCI across IGs

For the highest capacity interference group (Tap 1) receiving and the lowest capacity interference group (Taps 4, 5, and 6) transmitting: ALI bit-loading impact is ~ 0.5 bits/subcarrier in across the FDX band (108 MHz to 684 MHz).

For the next highest capacity interference group (Tap 2) receiving and the lowest capacity interference group (Taps 4, 5, and 6) transmitting: ALI bit-loading impact is about half the previous case.

For the next highest capacity interference group (Tap 3) receiving and the lowest capacity interference group (Taps 4, 5, and 6) transmitting: ALI bit-loading impact is again about half the previous case.

The ALI impact in further division into more interference groups would be negligible. For example (Tap 4) receiving and the lowest capacity interference group (Taps 5 and 6) transmitting.

Full Duplex Echo and Self-Interference Analysis

8. Full Duplex Echo Model and Analysis

An approach to modeling signal reflection from the various elements of the fiber deep network is discussed in this section. Analysis of the echo response including transmitted signal levels reflected back toward the node and received by the node and similarly for the cable modem is described. This approach allows the use of log magnitude values versus frequency from component spec sheets to model the reflections of each section of transmission line (i.e., cable) terminated at either end with an impedance (i.e., a tap, node, or cable modem) characterized by its magnitude return loss versus frequency.

Consider the fiber deep architecture model depicted previously in Figure 4. The fundamental reflection model function block is again denoted by the tap span shown in this figure. It is shown in Appendix 2 for a length of trunk cable with cable propagation delay T , cable amplitude response $A(f)$, and tap input/output port return loss $RL = -10 \text{ Log}(\rho)$ where ρ = reflection coefficient that the echo response $E(f)$ for the tap span cable transmission line is given by:

$$E(f) = \frac{A^2(f) 10^{-\frac{(RL_i)}{20}} e^{-j4\pi f T}}{1 - A^2(f) 10^{-\frac{(RL_i+RL_o)}{20}} e^{-j4\pi f T}}$$

Note that this formulation provides the complex frequency response with only the scalar amplitude versus frequency of the cable transmission line and the magnitude return loss versus frequency of the tap terminating impedances. This avoids the need to measure complex valued s-parameters versus frequency for each component in the network.

The same analysis applies to a drop cable section between tap port and the cable modem where the attenuation model and propagation delay are specified for the drop cable instead of the hardline cable and the input and output return losses are specified for the tap port and cable modem F-connector port respectively.

The echo response $E_n(f)$ for a reflection from the n^{th} tap from the node port with tap insertion losses $I_{n-1}, I_{n-2}, \dots, I_1$ (all tap loss amplitudes in linear magnitude as a function of frequency) is given as the echo response $E(f)$ above with the length of the echo path delay from the node to the n^{th} tap being n times the tap span delay T and with the square of the linear magnitude cable attenuation $A^2(f)$ being multiplied by product of all tap insertion loss responses in the echo path squared (two passes through the tap for each reflection).

The echo response from the node to the n^{th} tap in the echo path is given by:

$$E_n(f) = \frac{A^2(f) (I_{n-1} I_{n-2} \dots I_1)^2 10^{-\frac{(RL_i)}{20}} e^{-j4\pi f n T}}{1 - A^2(f) (I_{n-1} I_{n-2} \dots I_1)^2 10^{-\frac{(RL_i+RL_o)}{20}} e^{-j4\pi f n T}}$$

Denote the node total echo response as $E_{\text{node}}(f)$. The sum of all node echoes from all N taps is the sum of each echo path depicted in Figure 16 and given by:

$$E_{node}(f) = \sum_{i=1}^N E_i(f)$$




Figure 16 - Node principal echo paths

This approach is used to model the echo response between the node and all taps. The Fourier transform of the echo response in the frequency domain yields the impulse response in the time domain. The node port impulse response so obtained is shown in Figure 17.




Figure 17 - Node port echo impulse response

The principal echo from each equidistant tap is evident in the delayed peaks of the impulse response at multiples of twice the electrical delay path length to each tap which is equal to $2 \times 175' \times 1' / \text{ns} / 0.87 = 0.41 \mu\text{s}$ for a 175 foot inter-tap cable spacing with a velocity of propagation of 0.87 times the speed of light.

A similar formulation is used to model the echo response from any cable modem attached to the n^{th} tap with two echo paths, one through the tap toward the node and the other through the tap-to-output isolation toward the last tap as shown in Figure 18.




Figure 18 - Modem principal echo paths on the drop from tap j

This similar approach is used to model the echo response between the modem and all taps both upstream and downstream from the transmitting modem. The Fourier transform of the echo response in the frequency domain yields the impulse response in the time domain and is shown in Figure 19 for modems connected to the first tap 1 and the last tap 6 from the node.




Figure 19 - Modem port echo impulse responses connected at the first and last tap

Note that the impulse response of a modem connected with 100 feet of drop cable to the first tap exhibits a single reflection from that tap. Twice the electrical delay drop length to the tap is equal to $2 \times 100' \times 1' / \text{ns} / 0.85 = 0.24 \mu\text{s}$ for a 100 foot Series 6 (“RG-6”) cable. The very small echo at $0.41 + 0.24 = 0.65 \mu\text{s}$ is the primary reflection from the node.

Also note that the impulse response of a modem connected with 100 feet of drop cable to the last tap exhibits a multiple reflections from the drop plus all other taps and the node. The echo path delays are given by the drop delay plus multiple inter-tap delays at $0.24 + N \times 0.41 \mu\text{s}$ for $N = 0, 1, \dots, 6$.

9. Full Duplex Echo, Adjacent Leakage Interference, and Adjacent Channel Interference

Consider first the simultaneous full-duplex transmission and reception at the node. A high level functional block diagram of the node is shown in Figure 20.

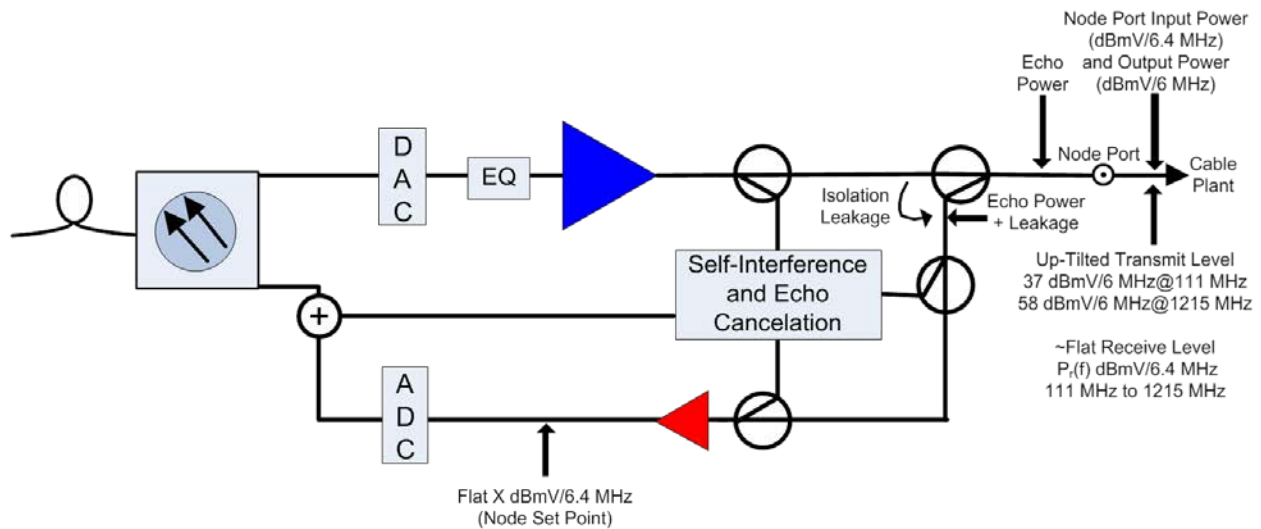


Figure 20 - Node functional block diagram for self-interference and echo cancellation

A directional coupler connected to the node port replaces the diplex filter in a conventional DOCSIS 3.1 FDD system. The directional coupler is needed to separate the upstream received signal from the downstream transmitted signal respectively entering and leaving the node port and occupying the same true full-duplex spectrum. Two sources of interference corrupt the reception of the upstream signal. The first is self-leakage of the downstream high power transmitted signal from the directional coupler output port to the tap port of the where the upstream signal is received. The second is the downstream high power echo returning from the cable plant through the node port and into the received signal path through the directional coupler tap port.




Figure 21 - Node echo plus leakage interference and signal-to-interference ratio

The downstream transmit power out of the node port has a 21 dB up tilt from 37 dBmV/6 MHz at 111 MHz to 58 dBmV/6 MHz at 1215 MHz as shown in Figure 21a. The received level at the node port is specified to be a flat power spectral density across the 108 MHz to 684 MHz full-duplex band. The value of that power spectral density level is determined such that the total composite power transmitted in the full-duplex band by any cable modem in the network does not exceed 64.5 dBmV. In this design example,

the received upstream power spectral density level meeting this modem transmit power limit is approximately 5 dBmV/6.4 MHz.

The node downstream echo power increases from an average 20 dBmV/6 MHz to 25 dBmV/6 MHz in the FDX band. The echo power level closely tracks the transmit power level attenuated by approximately 20 dB. The variation in echo power is over 10 dB peak-to-peak about the average. This is due to the multiple echoes with different path lengths that add on a voltage basis with rapidly varying group delay versus frequency causing both maximum constructive and destructive interference at various frequencies.

The self-leakage power from the directional coupler through port isolation to the tap port is seen to be higher average power than the echo average power. Again comprising attenuated and delayed versions of the same signal, the self-leakage power and echo power add coherently on a voltage basis at the directional coupler tap port as shown in Figure 21a. The node upstream signal to total echo plus self-leakage interference ratio is highly negative from -32 dB to -40 dB as shown in Figure 21b. Thus at least 60 dB to 70 dB of self-interference plus echo cancellation and/or suppression is required to obtain a positive signal to interference ratio that can support 1024-QAM upstream signal reception in the node.

Consider next the simultaneous FDD transmission and reception at the modem. A high level functional block diagram of the modem is shown in Figure 22.

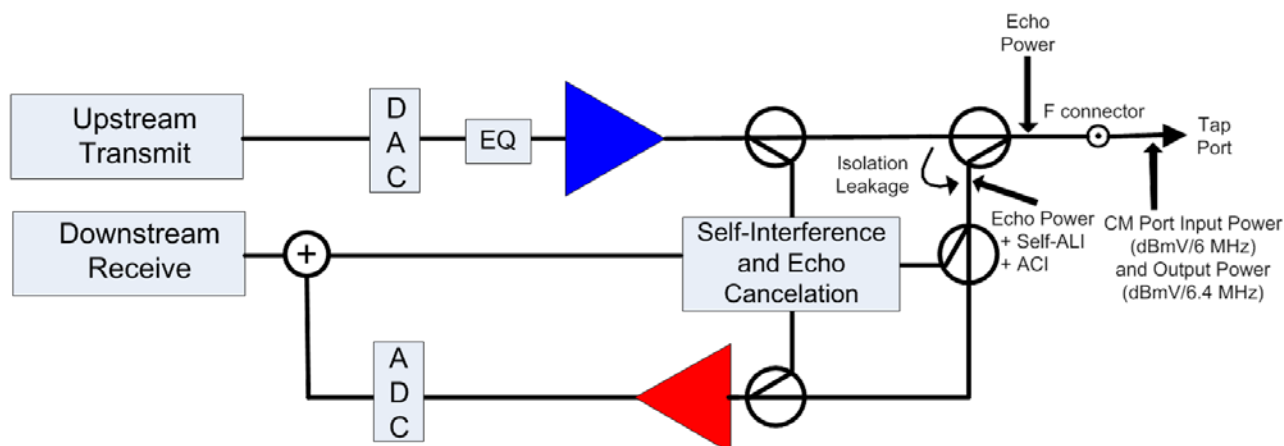


Figure 22 - Modem functional block diagram for self-interference and echo cancellation

A directional coupler connected to the modem port replaces the fixed diplex filter in a conventional DOCSIS 3.1 FDD modem. The directional coupler is needed to both separate the downstream received signal entering the modem port and combine the dynamically allocated spectrum of the upstream transmitted signal leaving the modem port. However, the coupler does not offer the isolation of the fixed-frequency diplex filter between the upstream transmitted and downstream received frequency bands.

The reason for separating transmission and reception within different channels is to obviate the need for echo and leakage interference cancellation in the same received downstream channel as the upstream transmission in a modem. If the modem transmits in a given channel, then both the echo and the leakage of the transmitted upstream signal can be canceled in that modem since the transmitted signal is known. Other modems in the same IG (e.g., on the same tap) will experience large co-channel interference from the transmitting modem. Unlike the situation at the node, these modems do not have knowledge of the transmitted signal and therefore have no reference with which to cancel the transmission leakage and

echoes from a different modem. Thus true full-duplex operation within the same channel in the modem is precluded.

Even by separating transmission and reception dynamically within different channels in the modem, two sources of in-band interference can still corrupt the reception of the downstream signal in a transmitting modem. Referring to Figure 22, the first is isolation leakage of self-adjacent channel interference from the upstream high power transmitted signal (i.e., self-ALI from transmitter out-of-band spurious emissions) across the directional coupler output port into the tap port where the downstream signal is received. The second is the echo of this self-ALI returning from the cable plant through the modem port and into the received signal path through the directional coupler tap port.




Figure 23 - Modem echo plus leakage interference and signal-to-interference ratio

The echo plus leakage interference and signal-to-interference ratio for a modem attached to a drop from tap 6 is shown in Figure 23a and Figure 23b respectively. This figure shows the interference across the entire downstream band. However, the upstream and downstream signals must occupy non-overlapping spectrum as explained previously. This is depicted in Figure 24.




Figure 24 - Modem adjacent channel echo and self-adjacent leakage interference

The received level at the node port is specified to be a flat power spectral density across the 108 MHz to 684 MHz full-duplex band. The value of that power spectral density level is determined such that the total composite power transmitted in the full-duplex band by any cable modem in the network does not exceed 64.5 dBmV. In this design example, the received upstream power spectral density level meeting this modem transmit power limit is approximately 5 dBmV/6.4 MHz as shown previously in Figure 21a.

The modem upstream echo power in the adjacent channels to the downstream receive channel is approximately 20 dB higher than the received downstream signal both measured at the modem port as shown in Figure 24. This echo power, although not in the downstream receive band, is a significant adjacent channel interference source that may impair the operation of the downstream receiver.

The high level upstream transmitted signal in the 40 to 50 dBmV/6.4 MHz range will contribute self-adjacent leakage interference due to the limited output-to-tap isolation of the directional coupler.

Consider the situation for dynamic self-interference mechanisms in the modem. Figure 25 shows time varying self-interference of FDD transmission and reception in the FDX spectrum for a single interference group. Each modem in the same IG receives in all channels except the transmitting CM channel in that IG. As shown in the figure, the transmitting CM in an IG is subject to both self-ACI and iso-tap ACI (from modems on other ports of the same tap) adjacent to its received bands, and self-ALI within all received bands.

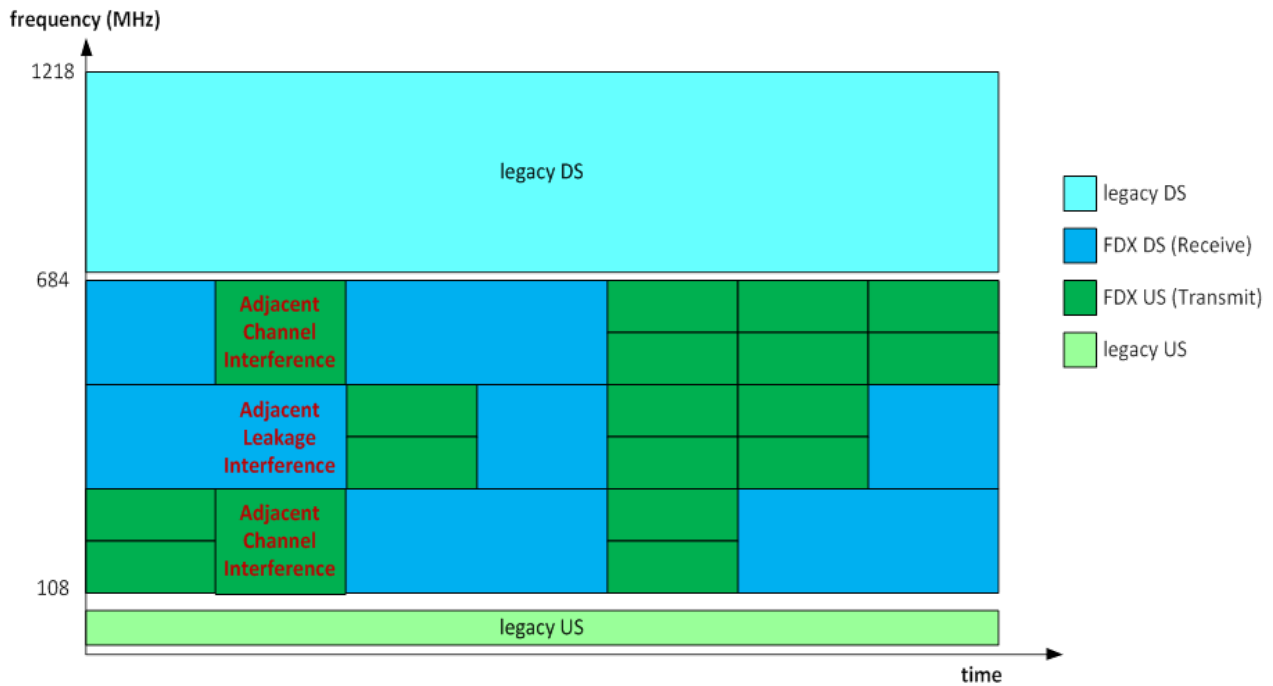


Figure 25 - Time varying self-interference of FDD transmission and reception in the FDX spectrum for a single interference group

Figure 26 shows one possible scenario for modem adjacent channel interference. The modem transmits upstream in the two highest FDX channels. The green line shows worst-case self-adjacent channel interference seen at the receiver of the same modem relative to its downstream receive level on the lowest FDX channel.




Figure 26 - Modem adjacent channel interference

Iso-tap ACI power at the receiving modem port is calculated as the power at the transmitting modem port less the drop losses and tap port-to-port isolation. The iso-tap signal-to-adjacent channel interference ratio is calculated for a modem on tap 6 of Model 1 using the calculation method of Appendix 5: SIR Calculation of ACI and plotted in Figure 27. Note that the ACI arising from a different modem on the

same tap cannot be canceled since the receiving modem has no transmitted signal reference from a different modem.

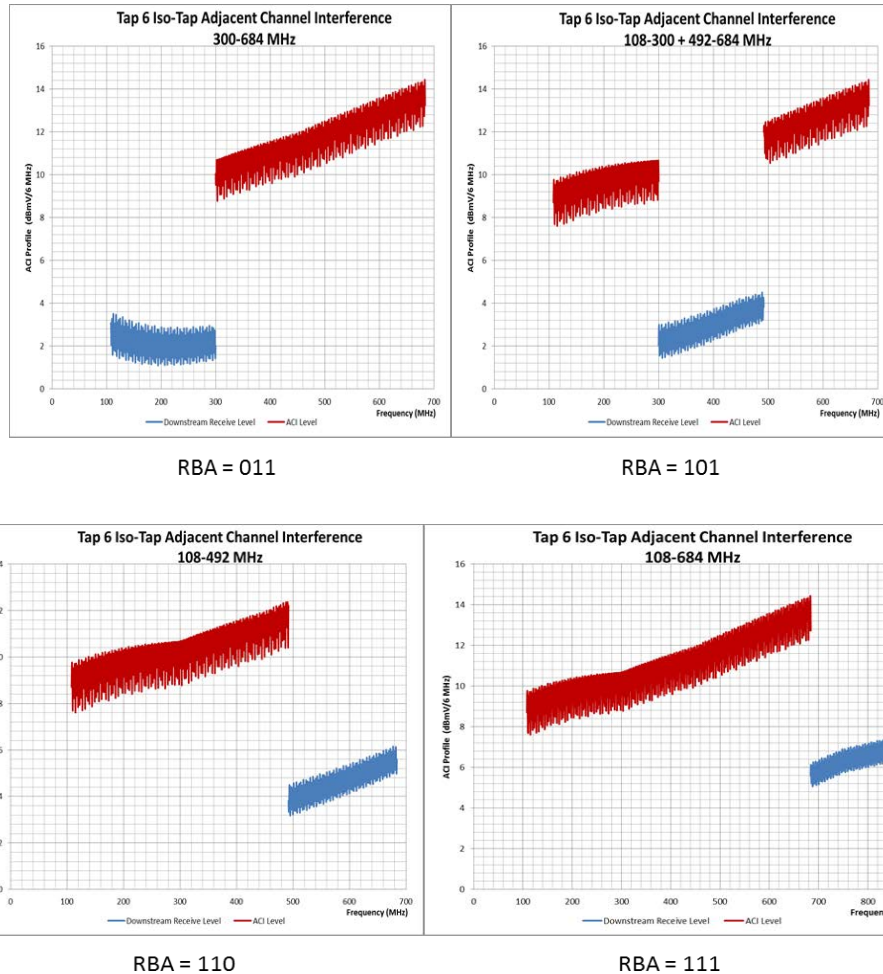


Figure 27 - Modem iso-tap adjacent channel interference for Model 1, tap 6

Self-ACI power is calculated as the transmit power plus the coupler insertion loss less the coupler isolation. The signal-to-self- adjacent channel interference ratio is calculated for a modem on tap 6 of Model 1 and plotted in Figure 28. This power will vary and may be mitigated in the transmitting modem based on the performance of the directional coupler chosen and the addition of self-ACI cancellation.

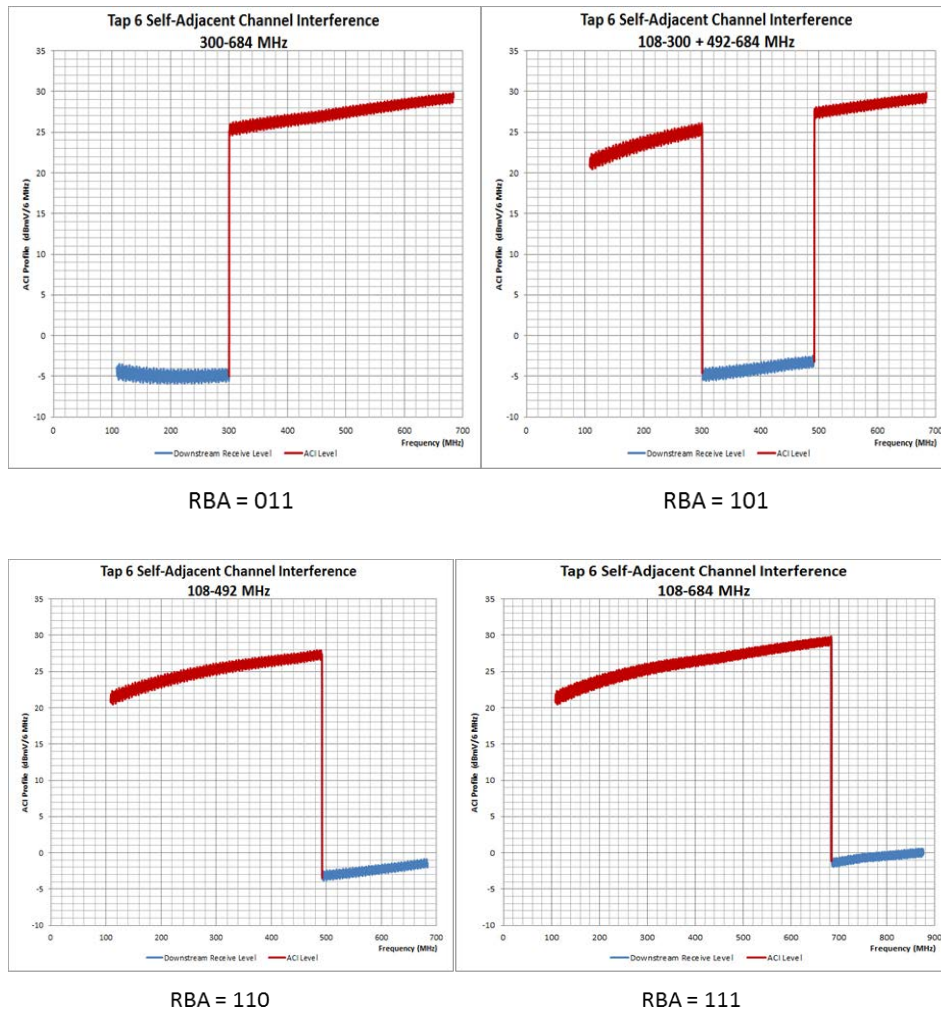


Figure 28 - Modem self-adjacent channel interference for Model 1, tap 6

Modem self-adjacent leakage interference will be considered next. Figure 29 shows one possible scenario for modem self-adjacent leakage interference. The modem transmits upstream in the two highest FDX channels. The red line shows worst-case self-adjacent leakage interference seen at the receiver of the same modem relative to its downstream receive level on the lowest FDX channel.

The self-adjacent leakage interference at the modem port in each sub-band with all other sub-bands transmitting is calculated for a modem on tap 6 of Model 1 and plotted in Figure 30. Self-ALI power in an FDX band receive channel arises from both other channels transmitting in two of the three sub-band channels that contribute interference. Self-ALI power in a receive channel in the legacy downstream band above the FDX band arises from all three FDX channels transmitting. Such high self-interference levels will require self-ALI mitigation to maintain high bit-loading in the FDX band. A traditional low pass filter plus a transition band would be an alternative approach instead of self-ALI cancelation above 684 MHz in the legacy downstream band.




Figure 29 - Modem self-adjacent leakage interference




Figure 30 - Modem self-adjacent leakage interference for Model 1, tap 6

Conclusion

The fiber deep architecture is designed for uniform spectral efficiency both downstream and upstream in frequency division duplex systems such as DOCSIS 3.0 and 3.1. Fiber deep specified transmitted and received levels are designed with enough margin to provide 4096-QAM spectral efficiency downstream with up to a 200 foot drop cable plus a 4-way splitter and additional in-home wiring to the customer premise equipment.

Full duplex systems based on DOCSIS 3.1 over the fiber deep architecture lower the spectral efficiency of the downstream progressively in interference groups as a function of tap distance from the node. Simultaneous transmission and reception within the same frequencies in different interference groups introduces co-channel interference lowering spectral efficiency. Simultaneous transmission and reception within the different frequencies in the same interference group introduces adjacent leakage interference lowering spectral efficiency. Both these interferences can occur concurrently and add destructively impairing downstream reception.

Impaired receivers due to self-interference and transmitter echoes in both the node and the cable modem require self-interference and echo cancellation.

Operation in the full duplex band lowers overall downstream capacity while greatly increasing overall upstream capacity (which is the primary objective) at the expense of substantially increased complexity.

Appendix 0: Cable Attenuation Model

It will be shown that the signal level attenuation solely due to frequency dependent passive cable attenuation (or “tilt”) can be reliably calculated from the measurement of actual cable attenuation characteristics such that the minimum mean squared error of the measured amplitude variations is achieved. This results in a simple model for the frequency dependent attenuation characteristic of a coaxial cable.

The coaxial cable produces signal attenuation per unit length that is frequency dependent whereby the higher frequency signals are subjected to greater attenuation than the lower frequency signals over the same length of cable. The attenuation in dB is proportional to the square root of frequency.

Specifically, for the cable attenuation A_1 dB at frequency f_1 and A_2 dB at frequency f_2 , the ratio of attenuations is equal to the cable loss ratio:

$$A_1/A_2 = \sqrt{f_1/f_2}$$

Define:

$A_H \equiv$ signal attenuation at the highest carrier frequency f_H

$A_L \equiv$ signal attenuation at the lowest carrier frequency f_L

Tilt \equiv the difference in signal attenuation between f_L and f_H is given by

$$Tilt = A_H - A_L = A_H - A_H \sqrt{f_L/f_H} \text{ (dB)}$$

Therefore the signal attenuation at the highest carrier frequency is given by

$$A_H = \frac{Tilt}{1 - \sqrt{f_L/f_H}} \text{ (dB)}$$

The signal attenuation at the lowest carrier frequency is given by

$$A_L = A_H \sqrt{f_L/f_H} \text{ (dB)}$$

Therefore, the *relative* signal attenuation at a frequency f where $f_L \leq f \leq f_H$ is given by

$$A(f) = A_H \sqrt{f/f_H} - A_L \text{ (dB)}, \text{ where } 0 \leq A(f) \leq \text{Tilt}.$$

Substituting $x = \sqrt{f}$ above yields

$$A(x) = \frac{A_H}{\sqrt{f_H}} x - A_L$$

which is a linear function of attenuation versus the square root of frequency.

Suppose one wants to estimate the cable attenuation at any frequency using the above attenuation model derived from some measured frequency vs. attenuation data pairs.

The attenuation model at any frequency could be derived from the measured data pairs in a least squares fit using linear regression (i.e., the line $F(x) = mx + b$ with slope m and intercept b) on the set of attenuation levels versus the square root of frequency.

Using the method of least squares for determining the best linear fit for the attenuation y vs. the square root of frequency x yields

$$m = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^N (x_i - \bar{x})^2}$$

$$b = \bar{y} - m\bar{x}$$

where $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$ and $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$ are the means of x and y respectively.

An example of this method for determining the attenuation vs. frequency characteristics from measured Series 11 (“RG-11”) drop cable attenuation at multiple frequencies is shown in the following figures. **Figure 31** shows the linear fit as a function of the square root of the frequency and **Figure 32** shows the attenuation as a function of frequency by straightforward independent variable substitution. The actual measured data points from the manufacturer’s data sheet are shown in blue and the estimated cable attenuation model calculated with a least squares linear regression are shown in yellow. Note the close correlation between the actual and the estimated attenuation levels (typically within less than 1 dB).




Figure 31 - Series 11 cable attenuation as a (straight line) function of \sqrt{f}




Figure 32 - Series 11 cable attenuation as a function of frequency

Appendix 1: Transmission Transfer Function

Consider a signal transmitted downstream from a tap output to the adjacent tap input as shown in Figure 33 with amplitude response $A(f)$ and linear phase response which has the corresponding impulse response of the cable denoted by $a(t)$. The transmitter at the signal source has (nearly) matched impedance to the drop cable but with a return loss RL_o (dB). The signal traverses the cable to the tap with propagation delay T which has a (nearly) matched impedance to the cable with return loss RL_i (dB). A portion of the signal equal to the reflection coefficient $10^{-RL_i/20}$ is reflected back to the source, which in turn a portion of the reflected signal equal to the reflection coefficient $10^{-RL_o/20}$ is re-reflected back toward the tap, and so on ad infinitum. This can be represented as a sum of the incident signal $x(t)$ and the infinite series of reflections each delayed by the round trip (i.e. twice) the propagation delay T of the cable.

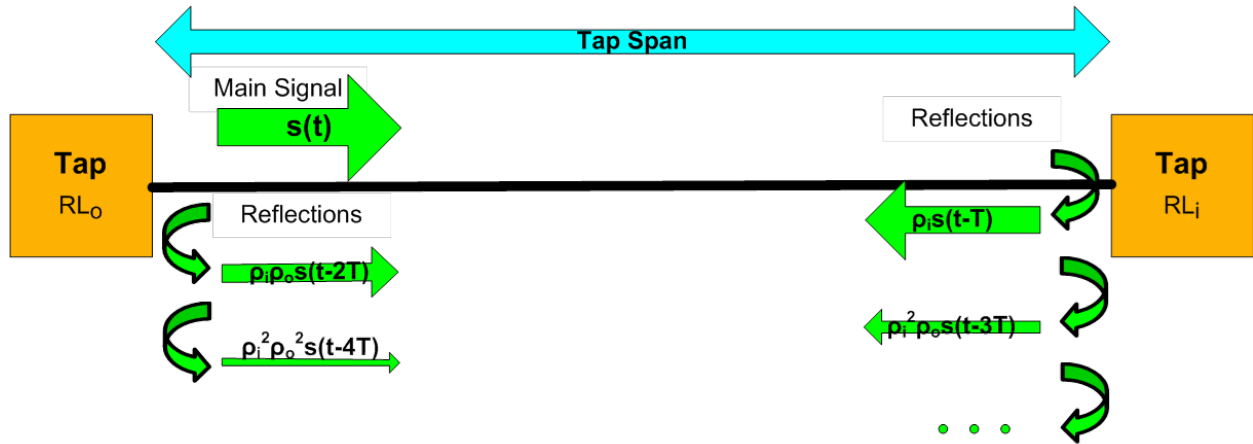


Figure 33 - Signal reflections in a cable between adjacent taps

Thus, the tap input consisting of the incident signal plus the discrete delays of each round trip reflected signal is given by:

$$y(t) = x(t) \otimes a(t) + 10^{-\frac{(RL_i+RL_o)}{20}} x(t-2T) \otimes a(t) \otimes a(t) \otimes a(t) + 10^{-\frac{2(RL_i+RL_o)}{20}} x(t-4T) \otimes a(t) \otimes a(t) \otimes a(t) \otimes a(t) \otimes a(t) + \dots$$

where \otimes denotes convolution. Taking the Fourier transform results in:

$$Y(f) = A(f)X(f) + A^3(f) 10^{-\frac{(RL_i+RL_o)}{20}} e^{-j4\pi fT} X(f) + A^5(f) 10^{-\frac{2(RL_i+RL_o)}{20}} e^{-j8\pi fT} X(f) + \dots$$

The transmitted signal transfer function $H(f)$ is given by:

$$H(f) = Y(f)/X(f)$$

$$= A(f) [1 + A^2(f) 10^{-\frac{(RL_i+RL_o)}{20}} e^{-j4\pi fT} + A^4(f) 10^{-\frac{2(RL_i+RL_o)}{20}} e^{-j8\pi fT} + \dots]$$

or

$$H(f) = A(f) \sum_{k=0}^{\infty} A^{2k}(f) 10^{-\frac{k(RLi+RLo)}{20}} e^{-j4\pi fTk}$$

Using the relationship

$$1 + r + r^2 + r^3 + \dots = \frac{1}{1-r}; |r| < 1$$

yields the closed form transfer function $H(f)$ as

$$H(f) = \frac{A(f)}{1 - A^2(f) 10^{-\frac{(RLi+RLo)}{20}} e^{-j4\pi fT}}$$

The same analysis applies to a feeder cable section between taps where the attenuation model and propagation delay are specified for the hard-line cable instead of the drop cable and the input and output return losses are specified for the through ports of the taps.

Appendix 2: Echo (Reflection) Transfer Function

Consider a signal transmitted downstream from a tap output to the adjacent tap input as shown in Figure 33 with amplitude response $A(f)$ and linear phase response which has the corresponding impulse response of the cable denoted by $a(t)$. The transmitter at the signal source has (nearly) matched impedance to the drop cable but with a return loss RL_o (dB). The signal traverses the cable to the tap with propagation delay T which has a (nearly) matched impedance to the cable with return loss RL_i (dB). A portion of the signal equal to the reflection coefficient $10^{-RL_i/20}$ is reflected back to the source, which in turn a portion of the reflected signal equal to the reflection coefficient $10^{-RL_o/20}$ is re-reflected back toward the tap, and so on ad infinitum. This reflected signal can be represented as a sum of the infinite series of reflections each delayed by the propagation delay T plus multiples of the round trip time (i.e. twice the propagation delay or $2T$) of the cable.

Thus, the tap output reflections consisting of the discrete delays of each round trip reflected signal is given by:

$$y(t) = 10^{-\frac{(RL_i)}{20}} x(t - 2T) \otimes a(t) \otimes a(t) + 10^{-\frac{2(RL_i)+(RL_o)}{20}} x(t - 4T) \otimes a(t) \otimes a(t) \otimes a(t) \otimes a(t) + \dots$$

where \otimes denotes convolution. Taking the Fourier transform results in:

$$Y(f) = A^2(f) 10^{-\frac{(RL_i)}{20}} e^{-j4\pi fT} X(f) + A^4(f) 10^{-\frac{2(RL_i)+(RL_o)}{20}} e^{-j8\pi fT} X(f) + \dots$$

The reflected signal transfer function $E(f)$ is given by:

$$E(f) = Y(f)/X(f) \\ = A^2(f) 10^{-\frac{(RL_i)}{20}} e^{-j4\pi fT} [1 + A^2(f) 10^{-\frac{(RL_i+RL_o)}{20}} e^{-j4\pi fT} + A^4(f) 10^{-\frac{2(RL_i+RL_o)}{20}} e^{-j8\pi fT} + \dots]$$

or

$$E(f) = A^2(f) 10^{-\frac{(RL_i)}{20}} e^{-j4\pi fT} \sum_{k=0}^{\infty} A^{2k}(f) 10^{-\frac{k(RL_i+RL_o)}{20}} e^{-j4\pi fTk}$$

Using the relationship

$$1 + r + r^2 + r^3 + \dots = \frac{1}{1-r}; |r| < 1$$

yields the closed form transfer function $E(f)$ as

$$E(f) = \frac{A^2(f) 10^{-\frac{(RL_i)}{20}} e^{-j4\pi fT}}{1 - A^2(f) 10^{-\frac{(RL_i+RL_o)}{20}} e^{-j4\pi fT}}$$

Appendix 3: SNR Calculations of CCI

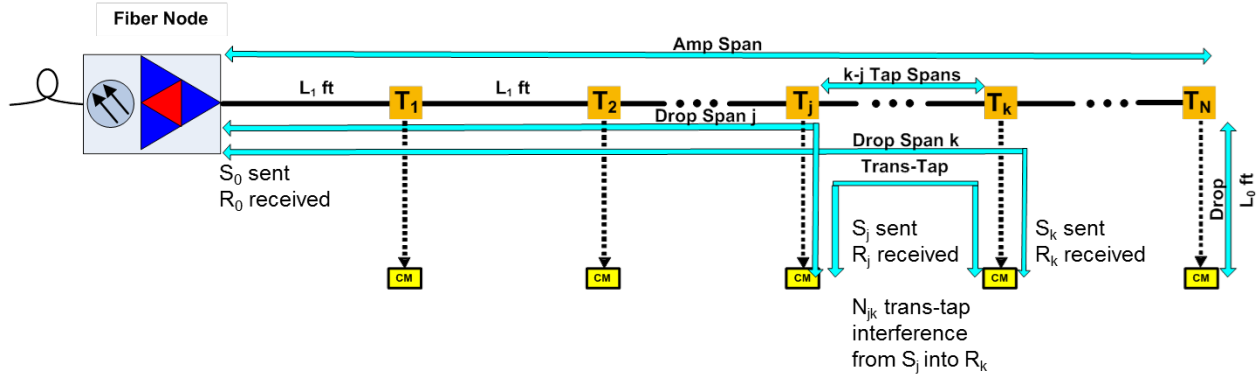


Figure 34 - Modem trans-tap interference in a coax cable system

Denote the transfer function $H(f)$ between a modem on tap j and a modem on tap k in Figure 34 as:

$$H_{jk} \equiv H(f) \text{ from tap } j \text{ to tap } k$$

where the node is defined as position 0. The received signal R_k is related to the node transmitted signal S_0 by:

$$R_k = S_0 H_{0k}$$

The trans-tap CCI N_{jk} by the transmitted signal from tap j to the received signal into tap k is:

$$N_{jk} = \{S_j | R_0\} H_{jk}$$

Noting that

$$R_0 = S_j H_{0j}$$

$$\text{so } \{S_j | R_0\} = R_0 / H_{0j}$$

Therefore

$$N_{jk} = R_0 H_{jk} / H_{0j}$$

and the SNR caused by the transmitted signal from tap j to the received signal into tap k is:

$$SNR_{jk} \equiv 20 \log \frac{R_k}{N_{jk}} = 20 \log \frac{S_0 H_{0j} H_{0k}}{R_0 H_{jk}} (dB)$$

SNR_{jk} is a function of frequency f_i

Average over f_i ; $0 \leq i < n$ for n frequency points to obtain:

$$\overline{SNR}_{jk} = S_{0dB} - R_{0dB} - 20 \log \frac{1}{n} \sum_{i=0}^{n-1} \frac{H_{jk}(i)}{H_{0j}(i)H_{k0}(i)} (dB)$$

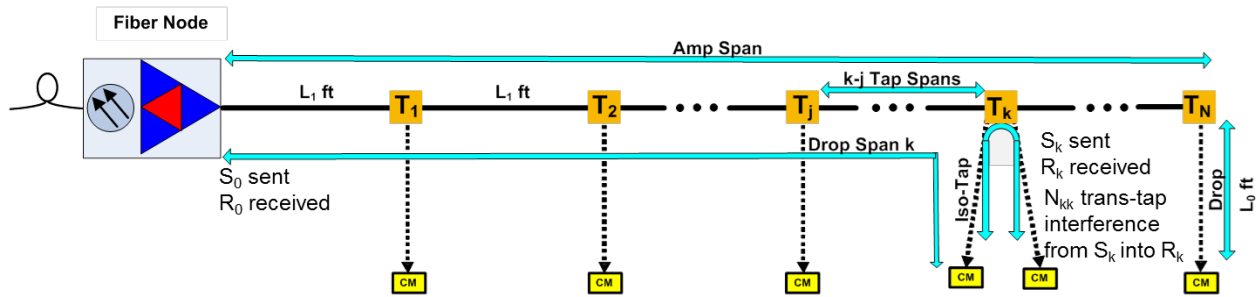


Figure 35 - Modem iso-tap interference in a coax cable system

Denote the transfer function $H(f)$ between a modem on a drop of tap k and a modem on another drop of the same tap k in Figure 35 as:

$H_{kk} \equiv H(f)$ between different drops on tap k

where the node is defined as position 0. The received signal R_k is related to the node transmitted signal S_0 by:

$$R_k = S_0 H_{0k}$$

The iso-tap interference N_{kk} by the transmitted signal between drops on tap k is:

$$N_{kk} = \{S_k | R_0\} H_{kk}$$

Noting that

$$R_0 = S_k H_{0k}$$

$$\text{so } \{S_k | R_0\} = R_0 / H_{0k}$$

Therefore

$$N_{kk} = R_0 H_{kk} / H_{0k}$$

and the SNR caused by the transmitted signal on a drop from tap k to the received signal into another drop on the same tap k is:

$$SNR_{kk} \equiv 20 \log \frac{R_k}{N_{kk}} = 20 \log \frac{S_0 H_{0k} H_{0k}}{R_0 H_{kk}} (dB)$$

SNR_{kk} is a function of frequency f_i

Average over f_i ; $0 \leq i < n$ for n frequency points to obtain:

$$\overline{SNR}_{kk} = S_{0dB} - R_{0dB} - 20 \log \frac{1}{n} \sum_{i=0}^{n-1} \frac{H_{kk}(i)}{H_{0k}(i)H_{k0}(i)} (dB)$$

Appendix 4: SNR Calculations of ALI

Denote the transfer function $H(f)$ between a modem on tap j and a modem on tap k in Figure 34 as:

$$H_{jk} \equiv H(f) \text{ from tap j to tap k}$$

where the node is defined as position 0. The ALI from the modem port into tap j is defined as A_j which is a function $f(\cdot)$ of the modem transmitted signal S_j . Noting that

$$R_0 = S_j H_{0j}$$

$$\text{so } \{S_j | R_0\} = R_0 / H_{0j}$$

Therefore

$$\{A_j | S_j\} = f(S_j) = f(R_0 / H_{0j}) = A_j$$

The trans-tap ALI N_{jk} by the transmitted signal from tap j to the received signal into tap k is:

$$N_{jk} = A_j H_{jk}$$

and the SNR caused by the transmitted signal from tap j to the received signal into tap k is:

$$SNR_{jk} \equiv 20 \log \frac{R_k}{N_{jk}} = 20 \log \frac{S_0}{A_j} \frac{H_{0k}}{H_{jk}} (dB)$$

SNR_{jk} is a function of frequency f_i

Average over f_i ; $0 \leq i < n$ for n frequency points to obtain:

$$\overline{SNR}_{jk} = S_{0dB} - A_{jdB} - 20 \log \frac{1}{n} \sum_{i=0}^{n-1} \frac{H_{jk}(i)}{H_{0k}(i)} (dB)$$

Denote the transfer function $H(f)$ between a modem on a drop of tap k and a modem on another drop of the same tap k in Figure 35 as:

$$H_{kk} \equiv H(f) \text{ between different drops on tap k}$$

where the node is defined as position 0. The ALI from the modem port into tap k is defined as A_k which is a function $f(\cdot)$ of the modem transmitted signal S_k . Noting that

$$R_0 = S_k H_{0k}$$

$$\text{so } \{S_k | R_0\} = R_0 / H_{0k}$$

Therefore

$$\{A_k|S_k\} = f(S_k) = f(R_0/H_{0k}) = A_k$$

The iso-tap ALI N_{kk} by the transmitted signal on a drop of tap k into the received signal of a modem on another drop of the same tap k is:

$$N_{kk} = A_k H_{kk}$$

and the SNR caused by the transmitted signal on a drop from tap k to the received signal into another drop on the same tap k is:

$$SNR_{kk} \equiv 20 \log \frac{R_k}{N_{kk}} = 20 \log \frac{S_0}{A_k} \frac{H_{0k}}{H_{kk}} (dB)$$

SNR_{kk} is a function of frequency f_i

Average over f_i ; $0 \leq i < n$ for n frequency points to obtain:

$$\overline{SNR}_{kk} = S_{0dB} - A_{kdB} - 20 \log \frac{1}{n} \sum_{i=0}^{n-1} \frac{H_{kk}(i)}{H_{0k}(i)} (dB)$$

Appendix 5: SIR Calculation of ACI

Denote the transfer function $H(f)$ between a modem on tap j and a modem on tap k in Figure 34 as:

$$H_{jk} \equiv H(f) \text{ from tap j to tap k}$$

where the node is defined as position 0. The received signal R_k is related to the node transmitted signal S_0 by:

$$R_k = S_0 H_{0k}$$

The trans-tap ACI N_{jk} by the transmitted signal in adjacent bands from tap j to the received signal into tap k is:

$$N_{jk} = \{S_j | R_0\} H_{jk}$$

Noting that

$$R_0 = S_j H_{0j}$$

$$\text{so } \{S_j | R_0\} = R_0 / H_{0j}$$

Therefore

$$N_{jk} = R_0 H_{jk} / H_{0j}$$

and the signal-to-interference ratio (SIR) from the ACI of the adjacent channel transmitted signal from tap j to the received signal into tap k is:

$$S/ACI_{jk} \equiv 20 \log \frac{R_k}{N_{jk}} = 20 \log \frac{S_0 H_{0j} H_{0k}}{R_0 H_{jk}} (dB)$$

S/ACI_{jk} is a function of frequency f_i

Average over f_i ; $0 \leq i < n$ for n frequency points to obtain:

$$\overline{S/ACI}_{jk} = S_{0dB} - R_{0dB} - 20 \log \left[\frac{\frac{1}{m} \sum_{i=0}^{m-1} H_{jk}(i) / H_{0k}(i)}{\frac{1}{n} \sum_{i=0}^{n-1} H_{0j}(i)} \right] (dB)$$

with n points in-band and m points in the interfering adjacent band.

Denote the transfer function $H(f)$ between a modem on a drop of tap k and a modem on another drop of the same tap k in Figure 35 as:

$$H_{kk} \equiv H(f) \text{ between different drops on tap k}$$

where the node is defined as position 0. The received signal R_k is related to the node transmitted signal S_0 by:

$$R_k = S_0 H_{0k}$$

The iso-tap ACI N_{kk} by the transmitted signal in adjacent bands to the received signal between drops on tap k is:

$$N_{kk} = \{S_k | R_0\} H_{kk}$$

Noting that

$$R_0 = S_k H_{0k}$$

$$\text{so } \{S_k | R_0\} = R_0 / H_{0k}$$

Therefore

$$N_{kk} = R_0 H_{kk} / H_{0k}$$

and the signal-to-interference ratio from the ACI of the adjacent channel transmitted signal from a drop on tap k to the received signal on another drop of the same tap k is:

$$S/ACI_{kk} \equiv 20 \log \frac{R_k}{N_{kk}} = 20 \log \frac{S_0 H_{0k} H_{0k}}{R_0 H_{kk}} (dB)$$

S/ACI_{kk} is a function of frequency f_i

Average over f_i ; $0 \leq i < n$ for n frequency points to obtain:

$$\overline{S/ACI_{kk}} = S_{0dB} - R_{0dB} - 20 \log \left[\frac{\frac{1}{m} \sum_{i=0}^{m-1} H_{kk}(i) / H_{0k}(i)}{\frac{1}{n} \sum_{i=0}^{n-1} H_{0k}(i)} \right] (dB)$$

with n points in-band and m points in the interfering adjacent band.

Abbreviations

ACI	adjacent channel interference
ALI	adjacent leakage interference
CCI	co-channel interference
CNR	carrier-to-noise ratio
CW	continuous wave
dB	decibel
dBmV	decibel millivolt
dBr	decibel reference (sometimes decibel relative)
DOCSIS	Data-Over-Cable Service Interface Specifications
FDD	frequency division duplex
FDX	full duplex
FEC	forward error correction
GHz	gigahertz
HFC	hybrid fiber-coax
IG	interference group
ISBE	International Society of Broadband Experts
IUC	interval usage code
MER	modulation error ratio
MHz	megahertz
OFDM	orthogonal frequency division multiplex
OFDMA	orthogonal frequency division multiple access
ODUP	OFDMA upstream data profile
QAM	quadrature amplitude modulation
QPSK	quadrature frequency shift keying
RBA	resource block assignment
RL	return loss
SCTE	Society of Cable Telecommunications Engineers
SIR	signal-to-interference ratio
SNR	signal-to-noise ratio
TG	transmission group
ZBL	zero bit loaded

Future Proofing Access Networks Through Wireless/Wireline Convergence

A Technical Paper prepared for SCTE•ISBE by

Martin J. Glapa

Partner and Bell Labs Fellow
Nokia, Bell Labs Consulting
366 Monte Vista Rd.
Golden, CO. 80401
303-517-1273

Martin.Glapa@bell-labs-consulting.com

Hungkei (Keith) Chow

Head of Department
Access Platform Research, Fixed Network Lab, Nokia Bell Labs
600 Mountain Ave., New Providence, NJ 07974
908-582-7685
Hungkei.Chow@nokia.com

Werner Coomans

Member of Technical Staff
Fixed Networks Lab, Nokia Bell Labs
Copernicuslaan 50, 2018 Antwerp
werner.coomans@nokia.com

R. J. Vale

Leader of Strategic Techno-Economic Analysis – Future Networks R&D
Nokia, Bell Labs Consulting
972-477-8674
Rj.vale@bell-labs-consulting.com

Enrique Hernandez-Valencia

Partner and Bell Labs Fellow
Nokia, Bell Labs Consulting
600 Mountain Avenue, Room 7E-318
Murray Hill, New Jersey, 07974
908-582-3144

Introduction

Multi-Gbps wireless and wireline services are pushing fiber deployments deeper into the outside plant and setting the stage for converged access networks. On the cable front, Distributed Access Architectures (DAA) using higher modulation orders and symmetrical Full Duplex (FDX) DOCSIS® will leverage existing drop coax, but will require MSOs to upgrade their HFC access networks to an N+0 deep-fiber architecture. In the special cases of MSO MDU (multi-dwelling unit) copper access through Fiber-to-the Distribution Point (FTTdp), copper assets will be leveraged using next-generation DSL technologies such as G.fast or G.mgfast. On the wireless front, licensed and unlicensed spectrum services (i.e., 5G, LTE, MulteFire, CBRS, etc.) will require small cells located close to the end-user, and require fiber all the way to the radio headend for backhaul. These deep-fiber-based architectures will all require that fiber be brought in very close-proximity to the end-users, creating an ideal convergence point to integrate short-reach distance-sensitive wireless and wireline technologies into a single converged access platform. Such a converged access platform will: 1) offer seamless service and content movement across any form of available access media, 2) provide dynamic resource shifting across multiple media access forms active at any time, 3) enable dynamic slicing, adaptation and coordination of network resources under the control of centralized virtual network orchestration in edge cloud/cable hub data centers, 4) be optimized in cost, space and power consumption, and 5) enable operational savings by reusing a converged infrastructure.

Technical, architectural and economic studies and insights on how MSOs can leverage their deep-fiber DAA/FDX DOCSIS networks to become leaders in deploying cost-effective future-proof converged wireline/wireless access networks and services, and become leaders in the deployment of mmWave technologies, are presented in this paper.

1. Future Access Network Vision

The future of access networks is defined by four major pillars that result in both architectural and technological shifts away from current access network design paradigms. While these pillars apply to all types of access networks, the focus in this paper is on cable and fixed wireless access networks, and their convergence on a common platform. Figure 1 illustrates these pillars which are:

1. Deep fiber – resulting in a future-proof network (to enable multiple MSO deployment options).
2. High capacity, low latency - enabling a hyper-scalable (subscriber and bandwidth tier scaling), hyper-capacity (symmetrical multi-Gbps bandwidth), hyper-performance (millisecond delays) network to serve a variety of MSO subscribers and bandwidth needs.
3. Convergence - of fixed wireless access networks on a common access platform – to enable hyper-flexible (dynamic support of multiple access mechanisms) based on need, demand, and competition.
4. SDN/NFV enablement – resulting in a programmable virtualized network.

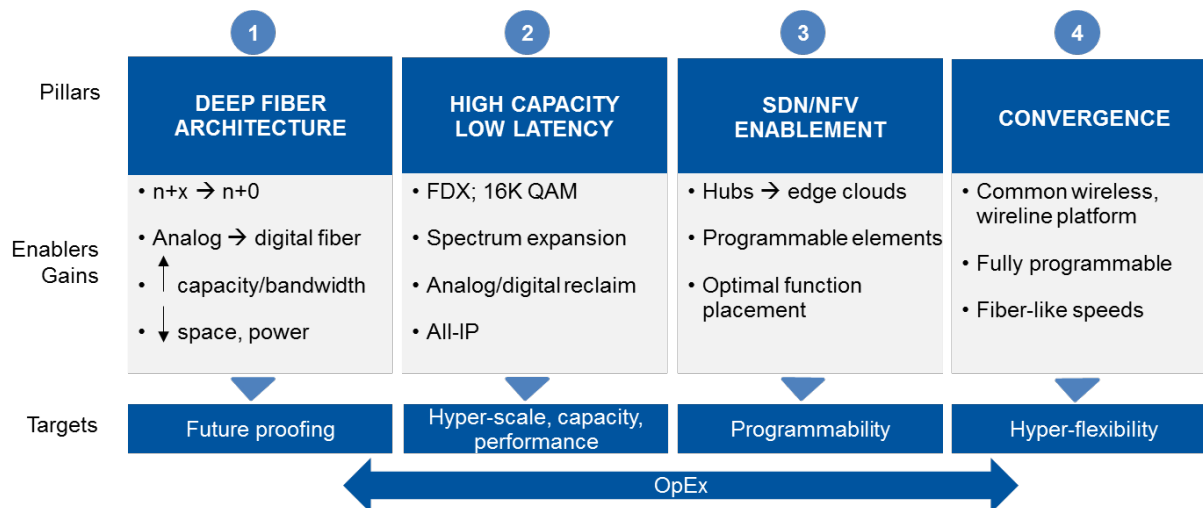


Figure 1 - Pillars of Future Access Networks

The following sections describe how these pillars enable future proofing of network investments through a common converged wireless/wireline access platform.

1.1. Deep Fiber Access Evolution

Access network architectures of all types: coax, copper, and fiber, have evolved and will continue to evolve dramatically as access speeds continue to increase. Generally, achieving high signal rates on access networks is distance-limited due to physical propagation characteristics. The use of higher spectrum bands that goes with the higher signal rate requires shortening the distance between the access node and the end user. In many deployment scenarios, fiber is brought to within 50m to 100m of end-users to achieve such multi-Gbps rates.

Cable access networks have been progressively pushing fiber deeper toward end users for many years. Early days of Hybrid-Fiber Coax (HFC) deployments used for video distribution saw the number of Households Passed per fiber node (HHP) in the several thousands, with many amplifiers on the coax between the fiber node location and the last user. As the need for voice and data services were added over time, and as the demand for high-speed data has continued to increase, fiber node serving areas have been made smaller by means of pushing fiber nodes (and hence fiber) closer to the end users. Thus, the number of subscribers served per fiber node have been reduced while providing greater bandwidth to users. In this process the average number of amplifiers on the coax cable between the fiber node and the last user has been reduced from $N+6$ (fiber node + the number of amplifiers) to $N+x$, ($x=5,4,3,2$) in many deployments. These architectures enable 1Gbps, and potentially even higher, downstream service access rates. Deep-fiber cable architectures, including deployment of Remote-MAC/PHY (RMD) and Remote-PHY (RPD) equipment at $N+x$ locations (typically where $x=2$ or 0), are referred to as Distributed Access Architectures (DAA). The introduction of Full Duplex (FDX) DOCSIS® will require an $N+0$ (i.e. fiber node deployed at the last amplifier location) architecture as FDX signals cannot be passed through existing amplifiers. In this point-to-multipoint architecture, fiber will be brought to within approximately 100m to 300m of end users, on average, and a maximum of several tens of users will be served from each node. Using 1.2 GHz of spectrum, FDX can theoretically provide 10Gbps symmetrical bandwidth. However, FDX DOCSIS will initially only target FDX operation over a reduced spectrum, enabling multi-gigabit upstream service rates. A longer-term alternative is to drive fiber all the way to the last tap,

bringing fiber to within tens of meters of about four to six users, and exploit signal bandwidths larger than 1.2 GHz.

In traditional twisted pair copper access, there is a long history of continuously increasing the data rates by increasing the fiber penetration depth. This evolution is continuing today with the advent of the G.fast standard, and the recent start of a new G.mgfast ITU-T project, which stands for multi-gigabit capable G.fast.

A similar evolution is occurring in wireless access networks. Macro cells, typically with a coverage radius of a few Km in a dense urban environment (depending on many factors such as radio frequency propagation, traffic, population density, etc.), are being supplemented with small cells with a much shorter radius. For example, 4G small cells may have a 100-300m radius, and 5G mmWave fixed wireless access will be limited to about 100m. These small cells will support peak rates of 1Gbps and sustained rates of up to 100Mbps¹. Deep fiber deployment will be required to backhaul the end-user traffic of such Radio Access Networks (RANs) to the wireless core sites.

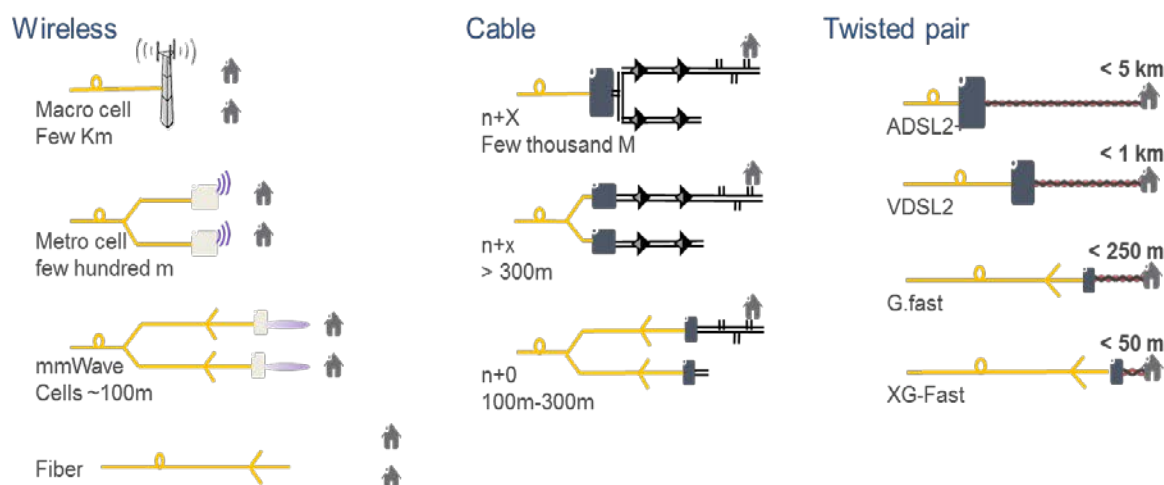


Figure 2 - Deep Fiber Access Networks

Figure 2 illustrates the migration to deep fiber in cable and wireless networks. All-of-the access architectures discussed, have at this point, one thing in common, fiber to within about 100m-300m of the end user. This enables an ideal convergence and colocation point from which to serve both fixed wireline and fixed wireless customers, and to enable future-proof networks.

1.2. High-Capacity and Low-Latency Cable Access

In parallel with deep fiber deployment, cable access networks are evolving along spatial, spectral efficiency, and signal spectrum dimensions as illustrated in Figure 3.

¹ Peak and sustained bandwidths are dependent on many factors such as the number of simultaneous users, the distance from a user to a wireless access point, and many others.

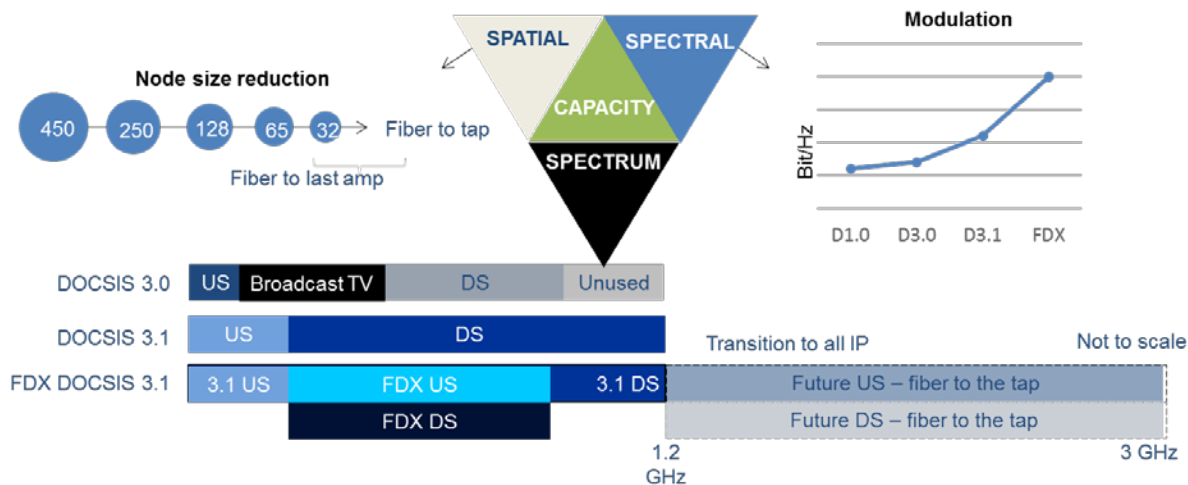


Figure 3 - Evolution to High Capacity Low Latency Networks

Spatial evolution is the progression of node, amplifier and service area reduction in the N+x architecture, facilitated by deep fiber deployment, as discussed in Section 1.1. Many MSOs are also evolving their outside plant spectrum to 1GHz or higher, enabling additional bandwidth to be used to support new high-bandwidth services. Future Fiber-to-the-Tap deployments may enable exploitation of broader signal spectrums on the order of a couple of GHz. In addition, spectral efficiency improvements over time, building up to DOCSIS 3.1 and FDX DOCSIS, will enable higher utilization of the available spectrum. Finally, the shift of analog towards all-IP further contributes to increasing the utilization of available spectrum resources.

DAA enables significant power and space savings at the hub or headend sites. These savings will be very useful for future MSO growth as it lowers the barrier to house general purpose servers and create edge clouds. Edge clouds are needed to support the new high-bandwidth and low-latency applications such as connected cars, augmented/virtual reality-based applications, and even virtual access nodes (such as virtual RANs). The longer reach provided by digital optics in DAA, combined with edge clouds, also enables consolidation of hub sites, saving further operational costs. Service-related functions can also be virtualized and deployed at these edge clouds locations, as needed, to derive additional benefits of increased performance, service agility, CapEx and OpEx reductions (over comparable monolithic approaches). Clearly, there is even greater opportunity to share the costs of facility, compute/storage resources and administration/management staff when wireless and wireline access networks' components are integrated (converged) at the same site, whether it is at a hub or a headend. Such convergence of access network components opens the opportunity to share services components such as firewalls, CDNs, DVRs, etc. across wireline as well as wireless customers using the same underlying hardware and software resources, and the same operational platforms and processes.

By bringing services closer to users, edge clouds not only enable better quality of experience for latency-sensitive services, but also permit elastic capacity scaling and help reduce transport network costs by avoiding backhauling user traffic to core locations. This is consistent with trends from leading Cloud Service Providers that are also collocating many of their service enabling capabilities closer to customers. Figure 4 shows examples of classes of applications requiring different low latency.

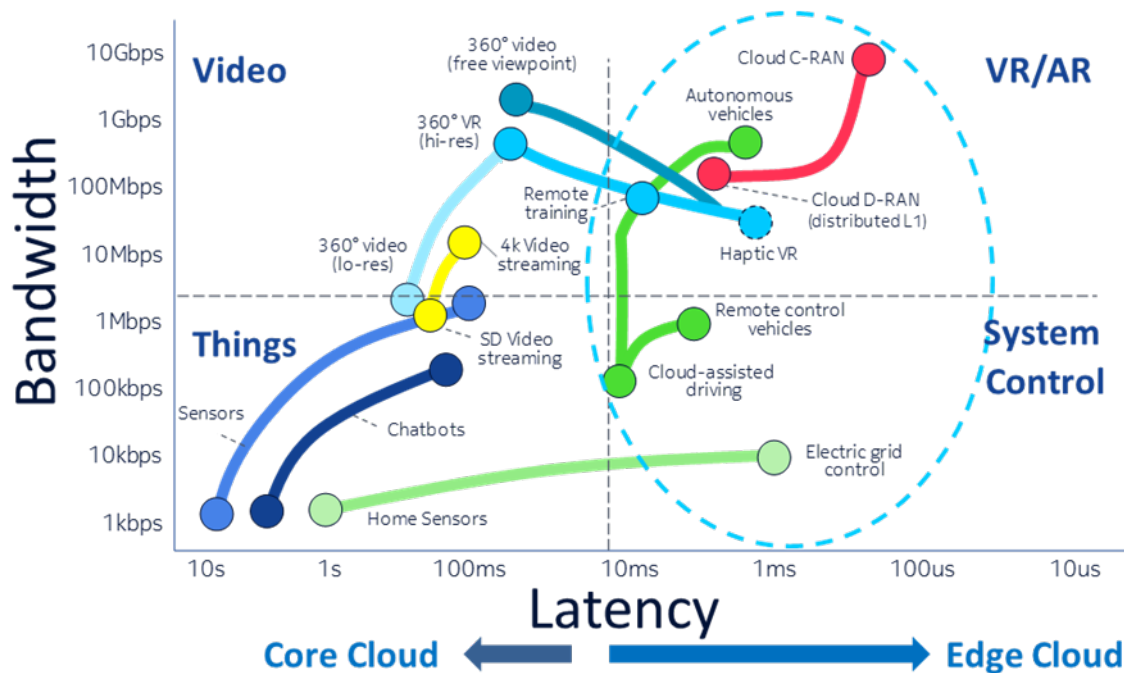


Figure 4 - New Applications Redefine Network Capacity and Latency Requirements

Hyper-scalable, hyper-capacity, hyper-performance networks thus become possible.

1.3. SDN/NFV Enablement

An SDN-oriented framework to control resources and automate the management of such converged networks is critical to achieve the envisioned economic benefits. A hierarchy of SDN controllers will be needed to disaggregate the management of the end-to-end services from the network resources across the multiple network domains (access, transport, hubs/headends and data centers). This programmability enables independent dynamic adaptations in the network resources according-to changing traffic patterns, chaining of new service functions on-demand and agility to instantiate, or develop, new revenue-generating services. Such modular abstraction of the network capabilities is needed to support a massively scalable and agile network architecture and operations that can handle a 100x expected increase in the number of service sessions from IoT and other digital home and enterprise applications.

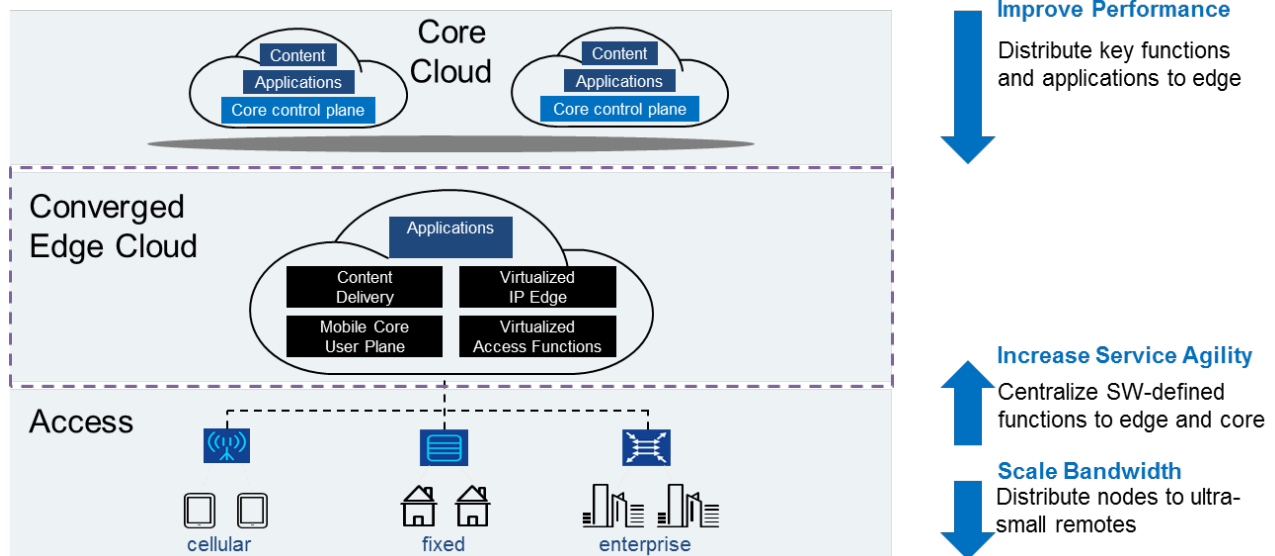


Figure 5 - SDN/NFV Enablement

1.4. Convergence

As described earlier, future high-speed wireless and wireline access technologies will be located within 100m-300m of end-users. As such, other access network characteristics, such as common backhaul, and co-resident physical layer (PHY) and media access control (MAC) layer further enable the development of a convergence access platform. Figure 4 illustrates these and other characteristics of future access nodes.

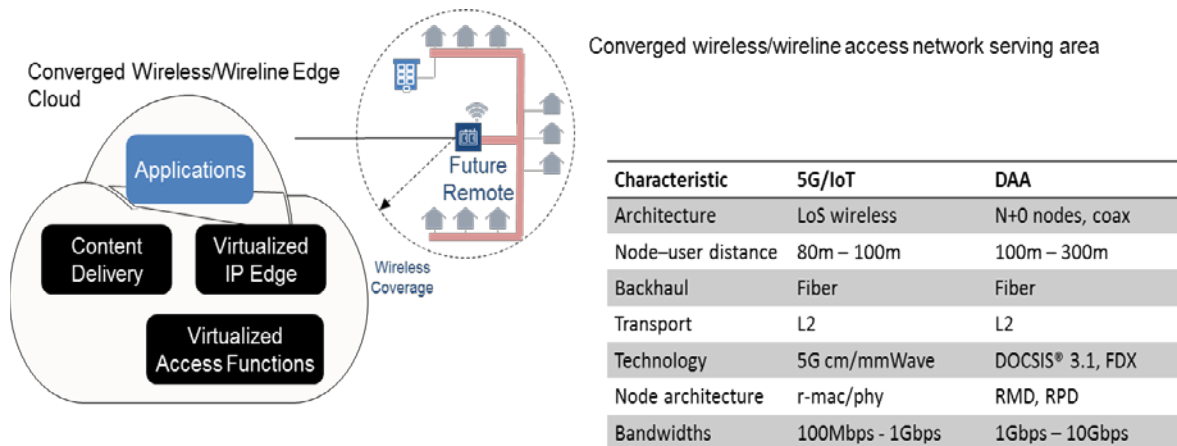


Figure 6 - Access Network Convergence Characteristics

Collocating common functions enables flexible and cost-effective deployment options. Variations in geography, morphologies and competition will dictate the specific configuration needs. A programmable converged access platform will offer flexible deployment options to meet such diverse needs. Longer

term, as wireless traffic is projected to grow at a faster pace than native wireline traffic, the platform can be reprogrammed to meet shifting consumer needs. Near term, as wireless and wireline usage patterns shift by time of day (refer to the use case example in Section 3), platform resources can be dynamically allocated to flexibly accommodate these needs. The next section discusses the technical details.

2. A Common Access Network Convergence Platform

The advent of NFV and SDN mentioned in Section 1.3 has given rise to the disaggregation of functionality that typically resided in bespoke hardware and software platforms. This approach, which also decouples hardware and software development, can be used to optimize the scaling and performance of each individual network function, thereby improving the overall performance of the network in various scenarios. The next step beyond NFV and SDN is to make the underlying hardware platform programmable so that it can be reconfigured and reused for multiple purposes. We have evaluated the potential for such a solution in the access network that is not only programmable to support a single technology dimension, but also across multiple key technology dimensions, e.g., wireline (Cable and Telco) and wireless (Cellular and Wi-Fi).

The anticipated future converged access network architecture is shown in Figure 6. In this architecture, a pool of common off-the-shelf computing and storage resources located at the future edge cloud hub can be configured and programmed to act as any combination of access technologies. Some common ones are shown as examples in Figure 6. More importantly, such a split/remote network architecture is already foreseen in the telecom service provider space as FTTN (Fiber-to-the-Node) and FTTdp (Fiber-to-the-distribution point), in the wireless service provider space as vRAN (virtual Radio Access Network) and in the MSO space as vCCAP with Remote-MAC/PHY. Such a converged edge cloud implements various Layer-2 and above networking functions as well as control and management interfaces.

Another degree of convergence in this architecture is the underlying feeder/interconnection network. Passive Optical Networking (PON) is becoming a technology of choice given its low equipment cost and flexibility in fiber distribution topology. PON has seen a significant increase in access capacity, for example, the evolution from EPON or GPON (1Gbps line rate) to 10G-EPON or XG-PON (10Gbps line rate). Recently, both the ITU-T and IEEE have begun the specification of higher data rate PON, noticeably as XGS-PON, NG-PON2 and 25G/100G TWDM PON.

Yet another less obvious degree of convergence is in the remote nodes, depicted as Future Remotes (FRs) in Figure 6, which form the basis of the future converged access solution. The FR implements the Layer-1 (physical layer) function of the supported access technologies. Operating as a combined wireless and wireline remote node in any combination, the FR can offer simultaneous converged access services that leverage both wireline and wireless connectivity to home/enterprise subscribers at the same time. Such an approach will provide better user experience that can be delivered at any point in time, using all available access resources in concert. Furthermore, given the common functional blocks in the PHY layers of future wireline and wireless technologies, economic benefits can be achieved by system vendors, chipset vendors and operators because of the multi-purpose usability of both these functional blocks and the resulting systems.

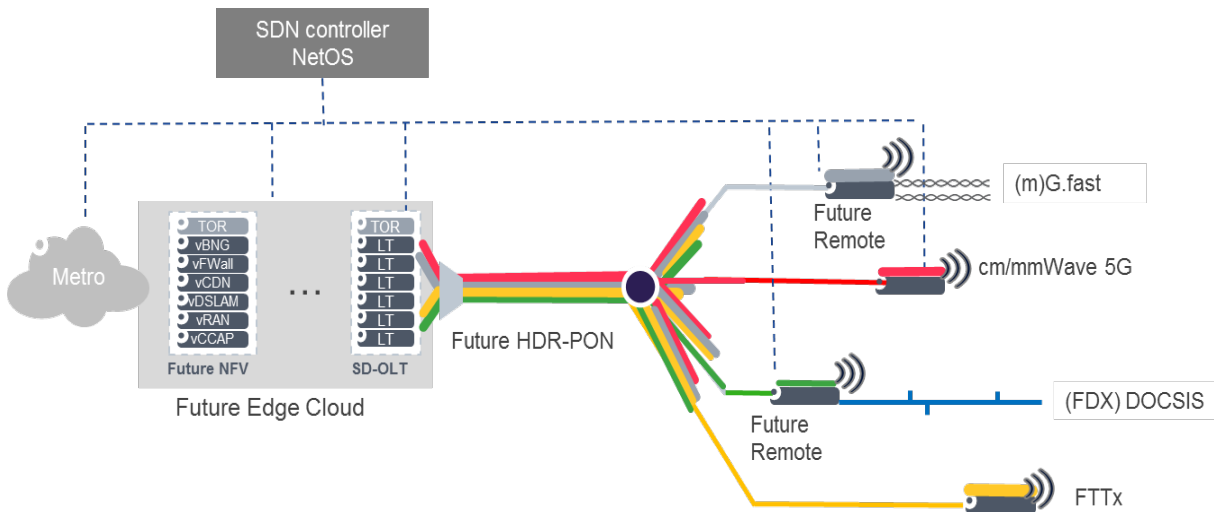


Figure 7 - Future Converged Access Network Architecture

2.1. Future Remote (FR)

Figure 7 shows a simplified block diagram of a FR node. On the subscriber line side, a FR is equipped with various analog front-end (AFE) modules. An AFE module implements a set of typical analog functions such as power amplification, RF mixing and filtering for a particular-application. Each supported access technology will have a corresponding AFE module that can be built as a pluggable module to suit the specific deployment scenario. Beside the AFE modules, the core of a FR is a universal PHY System-on-Chip (uPHY-SoC) which is realized as a programmable and reconfigurable ASIC. A uPHY-SoC implements the network side interface, e.g., a PON Optical Network Termination (ONT) or an Ethernet interface, a variety of digital baseband physical layer signal processing functions and/or MAC protocols. The operations, administration and management (OAM) as well as control functions of a FR, are performed over the SDN control and management interface.

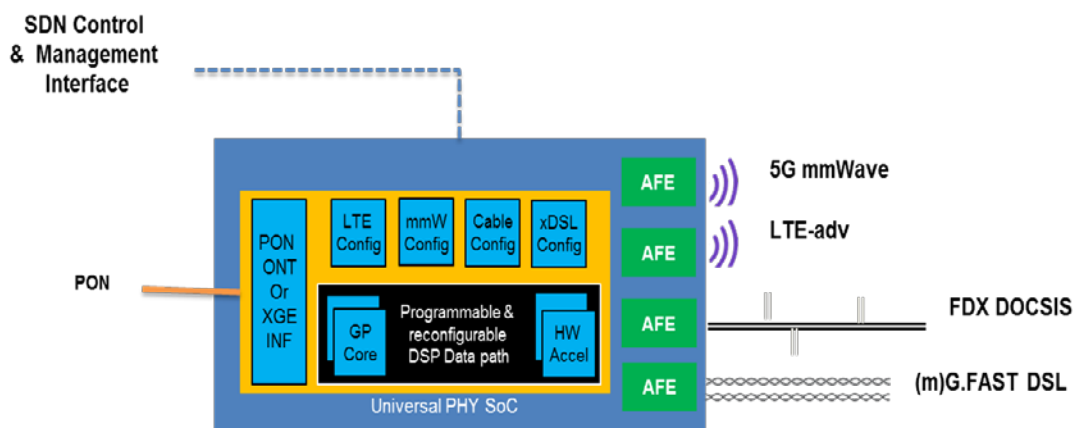


Figure 8 - Simplified Functional Block Diagram of a Future Remote

The future remote is in proof-of-concept stage and has been demonstrated by Bell Labs.

2.2. Universal PHY System-on-Chip

To realize the uPHY-SoC, we consider the digital baseband signal processing pipelines for all target access applications, such as the ones shown in **Error! Reference source not found.** The design objective of the uPHY-SoC is to achieve flexibility through programmability and reconfigurability, while maintaining the target throughput and efficiency, so that any combination of access technologies can be realized on a single uPHY-SoC. Even though the exact implementation of a particular-processing function may vary, our studies have shown that certain processing functions are better suited to be realized by either a general-purpose instruction processor (GPP) core, an array signal processor (ASP) or a specific function hardware accelerator (SFA). As indicated in 9, for example, the (I)FFT function, MIMO precoder function or pre-distortion filter function can be realized efficiently in an array signal processor structure, or more generally a Single-Instruction Stream-Multiple-Data stream (SIMD) processor structure, because of their regular and parallelizable data and instruction structure. Forward Error Correction (FEC) functions such as LDPC, turbo and Trellis/Viterbi code, however, are more suitable to be implemented as a reconfigurable special FEC hardware accelerator, given the massive internal connectivity of its signal flow. Other functions, such as scheduler, channel estimator and/or protocol processor, are good candidates for GPP cores, because of their irregular data and instruction flow implementation.

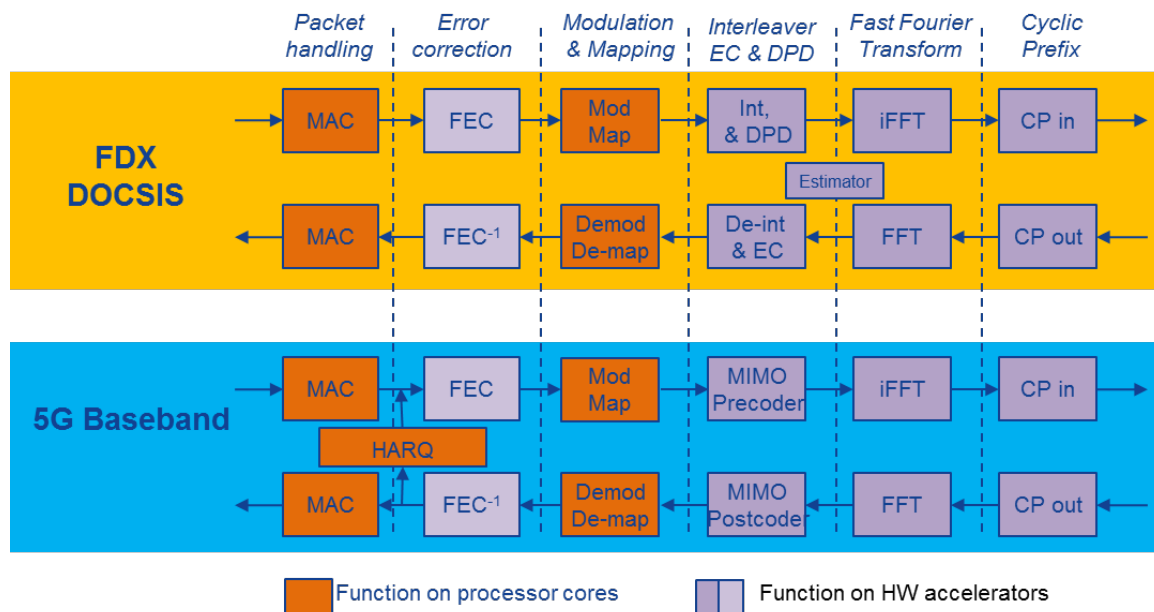


Figure 9 - Digital Baseband Processing Pipeline for FDX DOCSIS and mmWave Wireless

Given a mix of access applications under a target deployment scenario, one can determine the number of GPP cores, ASPs and SFA units required in a uPHY-SoC to support such a use case. Figure 9 illustrates the internal architecture of a uPHY-SoC. Various processing elements, GPP core, ASP cores and SFA units are connected on a configurable interconnection fabric. Even though **Error! Reference source not**

found. depicts a single interconnect fabric, actual implementation may be split into several dedicated fabrics with different throughputs and dimensions. On-chip SRAM may also be placed tightly-coupled with a certain set of processing elements or be placed as a shared memory pool. The high-speed input/output interfaces are typically being used to realize the network side interfaces, while the subscriber-side I/Os are connected directly via the DAC or ADC interfaces.

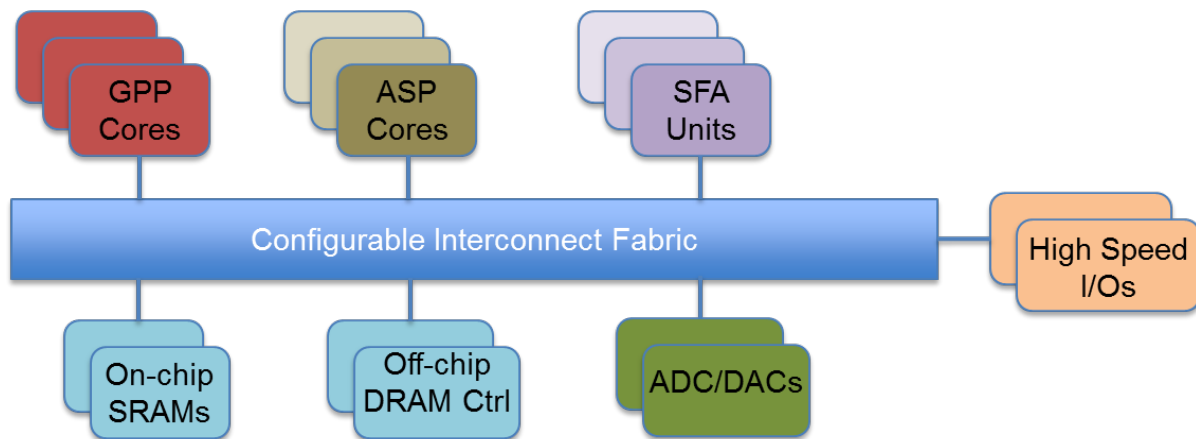


Figure 10 - Generic Universal PHY SoC Architecture

3. Use Case Example

3.1. Dynamic capacity

This use case represents typical user behavior pattern in terms of bandwidth consumption by access type varying by time-of-day. Refer to Figure 11. The use case assumes that the converged access platform supports both FDX DOCSIS and 5G fixed wireless access. By using a common platform and dynamically allocating capacity across the two networks, a converged operator can optimize fixed wireline and fixed wireless access deployment cost and provide flexibility to deal with dynamically changing traffic patterns.

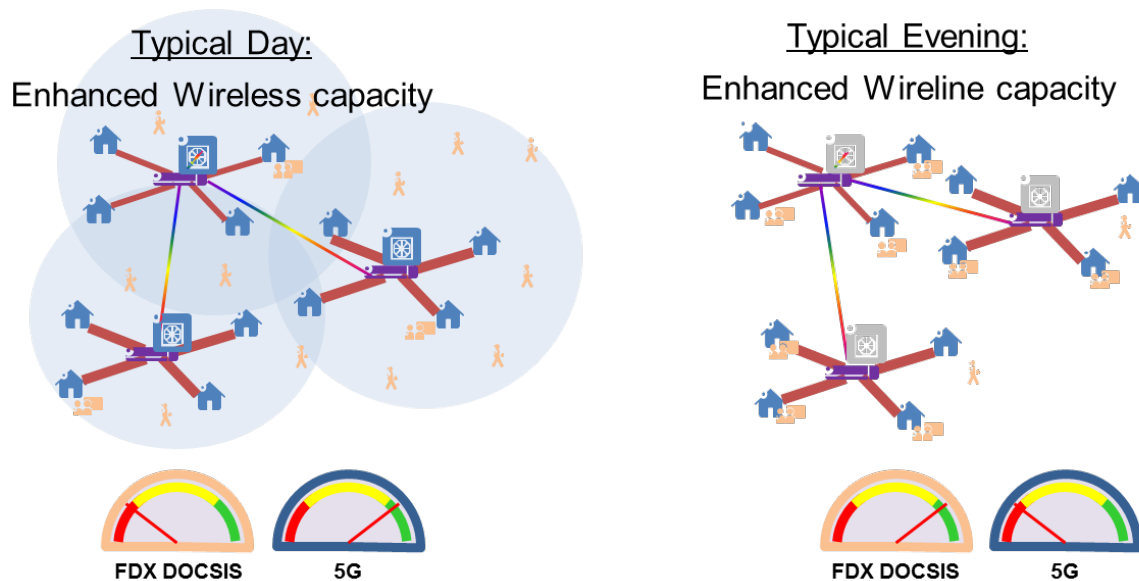


Figure 11 - Dynamic capacity use case

The left side of the figure illustrates that the primary day-time use is fixed wireless or nomadic access, while the right side of the figure illustrates the evening shift to primary wireline use of FDX DOCSIS as consumers settle in front of their PCs and TVs. The authentication for nomadic usage could happen through the standard mobile infrastructure enabling nomadic users to migrate to a mobile network when they move out of range. As an example, users at a bus station on a nomadic 5G fixed wireless connection switch over to mobile network as their bus moves out of range.

3.2. Economics / Comparative TCO analysis

A detailed total cost of ownership (TCO) analysis was performed to compare deployment models with discrete wireline and wireless systems against models using a converged system housing both wireline and wireless hardware. These were designed to provide dynamic capacity coverage as described in section 3.1 in three different morphologies – urban, suburban and dense urban. The wireless systems coverage was line-of-sight (LOS) with nomadic application. The results indicate that collocating wireless and wireline functions on the same converged access platform enables significant CapEx and OpEx savings over building discrete wireless and wireline access networks. Benefits are derived primarily from shared site acquisition, common fiber backhaul, powering, real estate and maintenance.

Dynamic wireline/wireless capacity use case

- Platform serves fixed wireline and nomadic 5G wireless users
- Day-time usage primarily nomadic wireless evening usage primarily 5G wireline
- FR resources dynamically shift between wireless and wireline service needs
- Up to 40% savings vs. building traditional-discrete wireline, wireless access networks




Figure 12 - Economic benefits of convergence

Figure 12 illustrates TCO savings of up to 40% over discrete networks in dense urban environments. Figure 13 illustrates the CapEx and OpEx savings details for the dense urban morphology.

TCO savings derived from:

- Site development: survey, acquisition, backhaul, power, enclosure
- Deployment: RF/network design, engineering optimization, install & commission, testing & integration, project management
- Network electronics: common baseband processors and Digital Signal Processors (DSPs)
- OpEx: network trouble management, preventive maintenance, RF performance optimization, capacity management, power

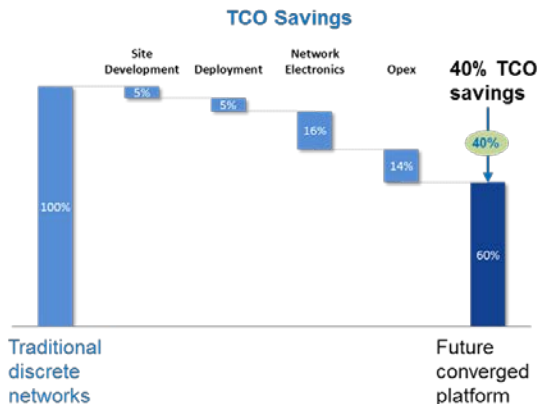


Figure 13 - Economic benefits details

The cost model is based on a converged network electronics when possible (e.g., uPHY-SoC, GPP etc.) that is shared between 5G fixed wireless and FDX systems. This architecture, in turn allows sharing many other common components such as power supply and mechanical systems on the same physical platform. Since both the platforms are contained in the same physical unit, it also results in considerable savings in site development, installation and operational expenses.

Conclusion

The future of access networks is defined by four major pillars that result in both architectural and technological shifts away from current access network design paradigms. These pillars are:

1. Deep fiber
2. High capacity, low latency
3. SDN/NFV enablement
4. Convergence

Deep fiber will bring fiber to within a few hundred meters of end-users. High capacity and low latency networks – enabled by deep fiber and the move to edge clouds, will enable MSOs to offer new applications. SDN/NFV will enable highly programmable MSO networks. Convergence of wireless and wireline technologies will be done at deep fiber locations, within a few hundred meters of end-users, on a common programmable access platform, and leverage the programmability of SDN/NFV. Such a platform is in the proof of concept stage at Bell Labs. Our studies have shown that up to 40% TCO savings can be achieved through the deployment of a converged multi-access (wireline and wireless) future remote platform over building discrete wireless and wireline access networks.

Abbreviations

ADC	analog digital conversion
AFE	analog front end
ASIC	application specific integrated circuit
ASP	array signal processor
Capex	capital expenditures
CBRS	citizens broadband radio service
CDN	content delivery network
DAA	distributed access architecture
DAC	digital analog conversion
DSL	digital subscriber line
DVR	digital video recorder
HFC	hybrid fiber-coax
EPON	ethernet passive optical network
FDX	full duplex
FEC	forward error correction
FFT	fast Fourier transform
FTTdp	fiber to the distribution point
FTTN	fiber to the node
FR	future remote
HHP	households passed
Gbps	gigabits
GHz	giga-hertz
GPON	gigabit passive optical network
GPP	general purpose processor
ISBE	International Society of Broadband Experts
I/O	input/output
IoT	Internet of Things
L1	layer 1
LDPC	Low density parity check
MAC	media access control (layer)

MDU	multi-dwelling unit
MIMO	multiple input, multiple output
MSO	multi-system operator
NFV	network function virtualization
OAM	operations, administration, management
PHY	physical (layer)
PON	passive optical network
RAN	radio access network
RMD	remote mac device
RPD	remote phy device
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networks
SFA	specific function hardware accelerator
SIMD	single-instruction stream-multiple-data stream
SoC	system on a chip
vCCAP	virtual common converged access platform
vCDN	virtual content delivery network
vRAN	virtual radio access network

Bibliography & References

The Future X Network A Bell Labs Perspective Book
Version Date: 20150910

How Integrated Photonics Enhances Capacity and Scalability for Fiber Deep Networks

A Technical Paper prepared for SCTE•ISBE by

Wayne Hickey

Advisor, Product Marketing
Ciena Corporation
1300 E. Lookout Dr. Suite 240
Richardson, TX, USA, 75080
469-771-9483
whickey@ciena.com

Joseph Shapiro

Director, Product Line Management
Ciena Corporation
385 Terry Fox Dr.
Ottawa, Ontario, Canada, K2K OL1
613-670-2096
jshapiro@ciena.com

Introduction

The Cable MSO industry is at the beginning of the next major chapter in the evolution of their access network story. The deployment of fiber deeper into the access, with N+0 targets, and the introduction of digital optics to support new Remote PHY Nodes, MAC/PHY nodes, PON and/or wireless access points, the Cable MSOs will position themselves to deliver high capacity 10G+ bandwidth services to their customer. To complement these next generation broadband access networks, the Cable MSOs will also need to invest in capacity between aggregation sites (e.g. Secondary hubs), Converged Cable Access Platform (CCAP) sites (Primary Hubs) and Headends. This paper will explore how fiber deep solutions that integrate photonics, coherent modems and open intelligent software & applications will enable optimal scalability and deliver maximum total capacity.

Content

1. Integrated Photonic Solution

An integrated photonic solution is an architecture that makes use of scalable, programmable, and instrumented hardware with sophisticated software applications to change the way fiber deep networks are engineered, deployed, maintained, and operated. This in turn enables the optical layer to fully participate in multi-layer network optimization which adapts and optimizes network behavior and resources for the specific service metrics required at a specific point in time. Key components of an integrated photonic solution are outlined below.

- a. Integrated coherent optics including variable-bit-rate capacity modems that allow operators to optimize capacity for any stage of the network's life
- b. Flexible Reconfigurable Optical Add-Drop Multiplexers (ROADM) for the ability to increase service availability and optically switch traffic across any path in the network
- c. Flexible grid photonic lines, for the ability to maximize fiber capacity using next generation coherent technologies
- d. Highly instrumented physical layer to ensure real-time monitoring and predictability of the programmable infrastructure
- e. Application Programming Interfaces (APIs) and SDN applications to drive the infrastructure

2. Integrated Coherent Optics

Integrated photonics refers to the inclusion of high performance optics directly into the aggregation and distribution components of the fiber deep solution.

A fiber deep network must be designed to physically reach the intended number of subscribers and adapt to the changing capacity needs of the subscribers which are expected to grow significantly throughout the life of the network. To meet these requirements, it must include components that support many physical interfaces, and support differing levels of throughput and transport capacity. One such device is the Remote PHY (RPHY) aggregation platform located in the secondary hub or primary hub. This device will terminate the 10G signals from the fiber nodes and efficiently aggregate the traffic for local hand-off to the CCAP or to a packet optical platform for transport upstream towards the hubs which house CCAP's.



Figure 1 - Hub to Fiber Node

In addition to being low cost, low power, and low footprint to optimize the capacity and scalability of the solution, the aggregation solution should be able to scale to support full capacity on the subscriber interfaces along with the associated line side transport bandwidth. Coherent line side optics deliver the highest capacity and performance and enable these benefits with minimum cost and cabling.

The spectral efficiency and operational simplicity advantages associated with coherent technology is the reason 100G coherent systems form the foundation of backbone networks today, and why this technology is the optimal solution servicing headend locations in fiber deep architectures. Coherent optics use advanced Digital Signal Processing (DSP), with the ability to detect amplitude, phase, and polarization of a signal, to impart significantly more information across fiber optic cables than traditional intensity-modulated direct-detect 10G systems. Moreover, coherent systems also enable a simpler network, as they electronically compensate for signal distorting effects such as Chromatic Dispersion (CD). By deploying coherent systems, operators can completely eliminate optical compensators, and the associated Capital Expenditure (CAPEX), latency, and manual planning associated with these from the network.

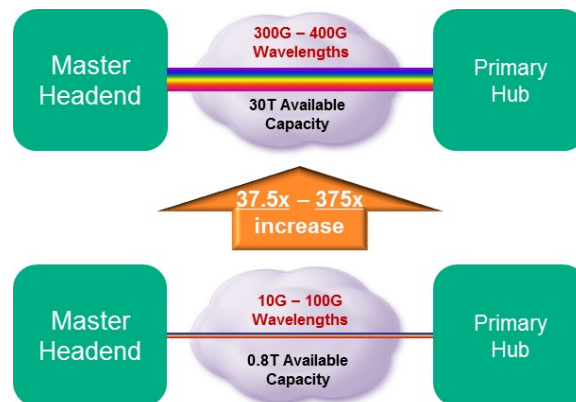


Figure 2 - Headend to Hub

Using coherent technology, numerous 10Gbps signals arriving from the fiber node can be efficiently mapped into a 100G coherent Dense Wave Division Multiplexing (DWDM) system and support up to 10T of capacity over a single fiber pair. Moreover, new imminently available coherent innovations are making it possible to radically drive down transport costs and improve scale of fiber networks. By leveraging the latest advancements in coherent optics and photonic capabilities, operators can now architect networks that can support up to three times that capacity between the headend and primary hub, as well as between the secondary to the primary hub segments, significantly improving scale and spectral efficiency of their fiber distribution networks.

Coherent systems today operate at symbol rates of 30-35Gbaud, and support up to 200Gbps of capacity per wavelength. Next generation coherent solutions, coming to market in 2017 and 2018, can operate at higher symbol rates and can process more information while still using a single set of electro-optics, dramatically reducing transport costs. Next generation coherent solutions can transport higher data throughput, such as 400Gbps of capacity, over a single wavelength. Resulting benefits include:

- Reduced transport costs, reduced footprint, and reduced power consumption, all resulting from the deployment of fewer coherent transponders
- Ability to maximize spectral efficiency and scale to higher capacity per fiber pair
- Simpler operations through the management of fewer wavelengths

These next generation coherent solutions also leverage more dense constellations and advanced DSP to offer a wide range of tunable capacity rates. An example of a programmable coherent modem, that can be tuned from 100G to 400G capacities in 50G increments, is shown in the figure below. Programmability in 50G increments allows cable operators to better match system capacity to available system margin. This translates into more bits carried at longer distances without requiring expensive Optical-Electrical-Optical (OEO) signal regeneration. The Forward Error Correction (FEC) implementation and DSP algorithms of the coherent solution will dictate its resulting system performance (how much capacity and unregenerated reach per channel), and how much ultimate capacity is achievable over fiber assets.




Figure 3 - Programmable coherent optics

In the future, coherent optics will also extend in to the access plant to support aggregation of 10G traffic from the fiber nodes to Outside Plant Aggregation Solutions. These optics will be optimized for reaches of 40 km to 100 km and drive capacities of 100G to 200G per wavelength.

3. Photonic Line System

The photonic line system is an important consideration in a fiber deep architecture. Providing the physical layer connectivity with as flexible an architecture as possible will enable maximum capacity but also enable new protection schemes increasing the availability of service without reserving bandwidth. Components of an integrated photonic layer include ROADM with directionless, colorless and contentionless capability supporting flexible grid applications. The photonic line system deployed for fiber deep can also be leveraged to provide connectivity for other services in the area thus increasing value and scalability of the network.

Flexible grid ROADMs provide several important functions including wavelength add / drop, switching, restoration and equalization required to optimize performance and allow the line side interfaces to operate at their maximum capacity. There are multiple configurations that can be deployed each providing

different levels of programmability: colorless (supporting any wavelength on any port), colorless and directionless (adding the ability to change the direction remotely) and contentionless (supporting multiple instances of the same wavelength on the add / drop device). They are programmable components critical to optimizing the benefit on an integrated photonic solution.

There are two primary benefits of flexible grid systems. First, a flexible grid photonic layer future proofs the network so it can carry any next generation higher bandwidth signal associated with higher baud rate coherent technologies. To achieve the desired system performance while operating at a faster symbol rate, wavelengths from these next generation modems will require larger spectrum than the traditional 50GHz seen in fixed grid systems today. MSOs will need to migrate to flexible grid systems to realize the full extent of economic savings associated with the new coherent technology.




Figure 4 - Example of optical signals requiring greater than 50GHz spectrum

The second benefit of a flexible grid photonic layer is that it allows the operator to squeeze carriers closer together, as shown below, and carry traffic across the least amount of spectrum. The operator manages the resulting media channel, also known as “superchannel”, as a single entity across the network, even though it may in fact consist of many optical signals. This use case is suitable for point-point applications, and is relevant for connections between secondary hub, the primary hub, and the headend node in MSO networks.




Figure 5 - Squeezing of Multiple Optical Signals Into One Media Channel

4. Instrumentation and intelligence

Optical Networks are largely perceived as rigid, static networks, which must be engineered for worst case scenarios. Worst case predicted A-Z capacity demands, worst case Service Level Agreements (SLAs), worst case margin allocation for worst case propagation conditions, etc. Essentially a one size fits all scenario, where the “one size” is engineered against best-guess predictions of worst case conditions. If the prediction turns out to be incorrect (e.g. A-Z demand is either larger or smaller than predicted), either more equipment needs to be ordered and installed with long lead time (demand larger than predicted), or deployed equipment remains unutilized/stranded (demand smaller than predicted). An intelligent integrated solution promises to radically change the way optical networks are engineered, deployed, maintained, and operated, by taking full advantage of agile, reconfigurable, and instrumented hardware with sophisticated SDN applications. This in turn enables the optical layer to fully participate in multi-layer network optimization which adapts and optimizes network behavior and resources for the specific service metrics required at a specific point in time. The implementation of real-time monitoring and intelligence capabilities in the photonic layer are becoming increasingly strategic capabilities of the network, as they allow cable operators to automate such tasks as configuration, provisioning, and troubleshooting. Full value of these capabilities is realized when the coherent optics and optical line system are part of an integrated solution, as full communication and visibility is possible between optics and line system, resulting in accelerated wavelength provisioning cycles and optimal system performance.

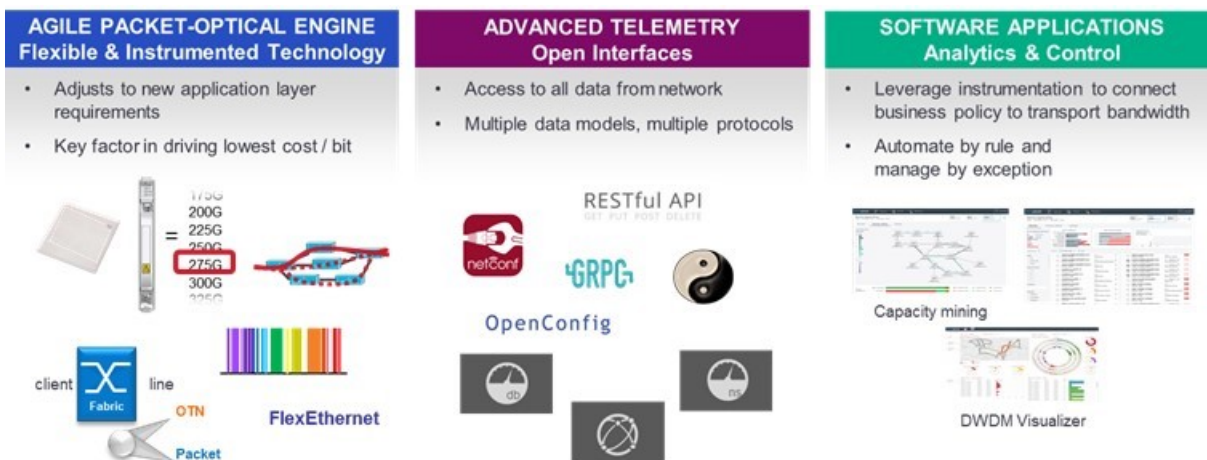


Figure 6 - Programability, Telemetry, and Analytics

An application that is enabled by the real-time monitoring and intelligence in the integrated solution is a Layer 0 (L0) control plane. A L0 control plane is an important component in simplifying optical network processes. It automates numerous network functions, using real-time network information to provide automated topology discovery and accelerated single-step service provisioning for faster turn-up of wavelengths, increased automation for efficient planning and operations, and photonic restoration for increased availability. Another important benefit of a Layer 0 control plane is the alternative protection architectures that it enables by facilitating wavelength re-grooming and switching without reserving bandwidth enabling operators to support higher overall capacity perform proactive network maintenance in a condensed maintenance window with fewer truck rolls.




Figure 7 - L0 Restoration

Another example of monitoring capability that can be used by the system to optimize performance and availability is the integrated Optical Time Domain Reflectometer (OTDR). Being able to find the location of fiber cuts and other impairments for the fiber deep network is critical to its operation. OTDRs can be expensive and require professional resources to run. The OTDR functionality can be integrated within the photonics platform and with the right software tools pin-point the fiber cut down to the street level. This capability eliminates the traditional lengthy troubleshooting step of sending technicians with test sets to either end of the failed span to localize the failure. Instead, the technician is dispatched to the precise fault location to promptly execute the repair. This quick turnaround results in increased network availability and reduced outage times. Integrated OTDR traces that can be run in-service provide additional benefits. Now, cable operators can proactively check for fiber degrades or bad repairs and ensure their network is operating with optimal performance.

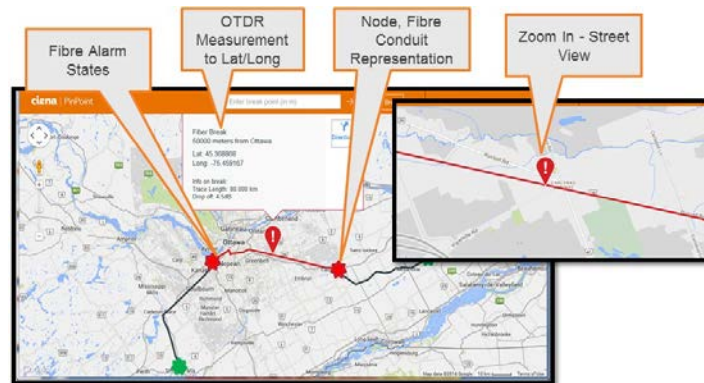


Figure 8 - Integrated OTDR Functionality

5. APIs and Advanced applications

As new platforms are built to support the hardware flexibility and programmability described above, they must also incorporate industry-standard, open APIs and common management interfaces. Modern, normalized data models and APIs are required to access the instrumented fiber deep network, and use high-performance telemetry to measure and predict, at any time and for various scenarios, system performance margin and resulting efficiency of the network.

New platforms built to support fiber deep applications must be designed to offer a simple, server-like deployment and operational model. Designed for intuitive installation and ease of operation, these fiber

deep platforms can offer rapid scalability and rack-and-stack simplicity to facilitate massive rollouts with minimal engineering effort.

By leveraging open APIs, such as REST, NETCONF, and gRPC, unique operational tools and scripts can be designed to specific MSO requirements. Examples include applications for network visualization, fault management, capacity management, and performance monitoring. Alternatively, the fiber deep platforms can be integrated into an existing operational paradigm via the open APIs.

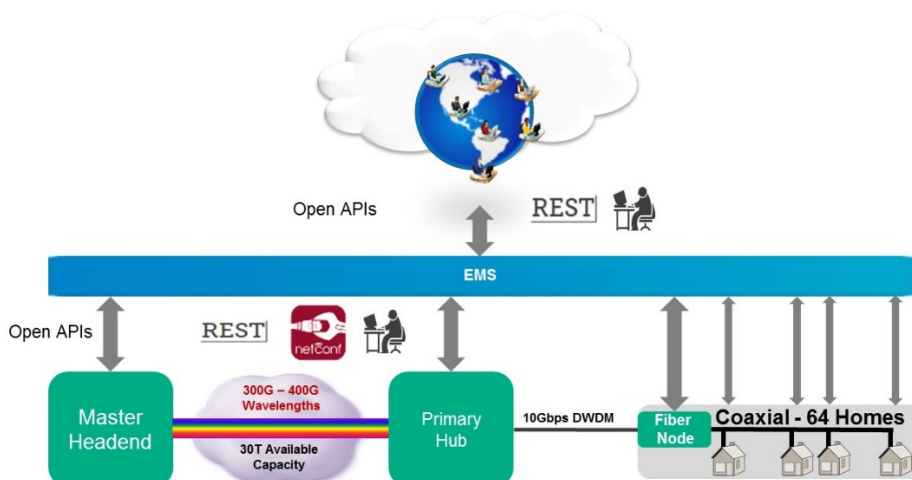


Figure 9 - System Automation

Platforms can also be managed through more traditional means with typical management interfaces, including CLI (Command Line Interface) and SNMP (Simple Network Management Protocol), or full lifecycle management software which enables out-of-the box planning, commissioning, and operation of the platforms without requiring any code development. This flexibility gives MSOs the choice to manage the platform through rich management software or through customized applications and scripts or back-office integration. Either way, open APIs, along with ease of use, let MSOs focus on growing their fiber deep networks without wasting effort on complex operations and integration.

6. SDN Applications

Advanced software applications abstract the complexity associated with advanced flexible technologies, enabling cable MSOs to fully operationalize and realize benefits associated with a fiber deep network. Applications should run 'off-box' in the cloud to take advantage of typical cloud computing and scale properties. Some examples of SDN applications include Bandwidth Optimization and Dynamic Restoration.

A Bandwidth Optimization application can be used to scale operations of the network, as well as convert the traditional static network to a dynamic, more strategic asset that can be used to generate new revenue streams. By accessing photonic instrumentation and exploiting properties of variable bit-rate coherent optics, the app can dictate the optimized capacity configuration and channel placement based on customer-definable margin policies so that cable MSOs can drive more efficiency from deployed assets. Taking it one step further, cable MSOs can leverage this intelligence and automation to convert excess system margin to additional capacity that can be used to drive new revenue streams such as Bandwidth-on-Demand services, without the need to deploy new hardware.

Another example of an advanced application that can be enabled by a solution with integrated photonics is Dynamic Restoration. Today's optical restoration is limited in that the restoration path must be pre-engineered so that the per-lambda capacity which is viable in the working path is also viable in the restoration path. This implies that either the length/propagation impairment of the restoration path is similar to (or shorter than) the working path, or regens must be pre-deployed. It also implies that the exact amount of spectrum that the restorable traffic occupies in the working path is available on the restoration path (and if it is not available, lower priority traffic that may be present on the restoration path needs to be pre-empted/dropped).

With Dynamic Restoration, a wider range of protection options are available, helping MSOs increase service availability and provide a higher quality experience to their end users:

- a. "Partial" restoration: If the restoration path is more challenging than the working path (typical situation), one has the option to "downshift" the capacity of the modems which are being re-routed to better match the more challenging propagation impairment of the restoration path. In this case, instead of all traffic being dropped, some of the traffic is restored along the more challenging path. The user can pre-determine which service to drop when the partial restoration option is used
- b. "Temporary full" restoration: Even when the restoration path is longer than the working path, it is possible in some scenarios to borrow sufficient dBs of margin to temporarily restore the full capacity and maintain services for end users

Advanced SDN applications, such as those mentioned above, will bring more capacity, scalability, and flexibility to the fiber deep network.

7. Small Cell MBH Application

With increased capacity and scalability, brings an abundance of new business opportunities. One such opportunity, being the growth of small cell Mobile Back Haul (MBH). Continued popularity of accessing applications and content is forcing Mobile Network Operators (MNOs) to continually expand their mobile network. As MNO's expand to Long Term Evolution (LTE) and LTE-Advanced (LTE-A) to accommodate packet-based mobile data services, whether man or machine.

Wireless networks are increasing becoming more between application end-users and their associated content, making the network a dominate factor in customer Quality-of-Experience (QoE). Small cells bring end-users and their mobile devices closer to the mobile network radio, vastly improving access performance. There are many emerging small cell technologies, including femto cells, pico cells, micro cells, WiFi cells, and small cells, the latter often assumed to include some or all terms, but all benefiting from integrated photonics. Small cells are deployed in two manners, as shown below one aggregating small and macro cell traffic to the Mobile Telephone Switching Office (MSTO), and secondly, they can be homed directly to the MSTO. Both applications are in use today, and deployed based on specific network requirements, deployment constraints (indoor or outdoor), and optical fiber availability.

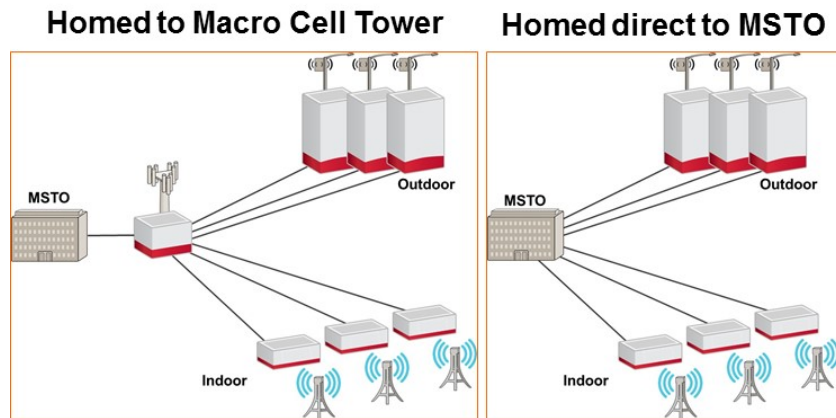


Figure 10 - Small Cell Back Haul Deployment

Current mobile radio technology (e.g. 2.5G 2.75G, 3G, pre-4G and 4G) are all deemed as ‘best effort’, in that most rarely achieve their maximum download or upload speeds. There are many factors why this happens, including large distances from mobile devices to macro cell towers, line-of-sight obstructions, indoor usage, transmission signal interference, and mobile device performance limitations. Removal of excessive equipment and optical transmission paths, as shown below can improve both active and protection distances.




Figure 11 - Optimized Optical Path

Small cells allow MNOs to better utilize wireless spectrum by offloading macro cell traffic, cable operators can ensure rigid Service Level Agreements (SLAs) using simple to own and operate packet based Operations, Administration, and Maintenance (OAM) tools and optimized optical path resiliency.

Conclusion

The advent of fiber deep is to push the optical-to-electrical conversion closer to the subscriber, resulting in a lower service group size, enabling higher throughput and capacity. “Fiber deep” means there will be more fiber nodes, fiber, and optical transceivers. Integrating optical connectivity and functionality can be seen as a viable and valuable approach to cost effectively improve network utilization.

Integrating DWDM optics and advanced coherent optical technology into switching and aggregation platforms can groom traffic from fiber nodes, enabling maximum capacity, optimal spectral efficiency, and ultimate economic savings - reducing space and power, while improving operational efficiency and performance. Moreover, advanced applications can be used to simplify the automation of services, enhance network reliability and enable new revenue opportunities using software enabled API's.

Abbreviations

API	Application Programming Interface
BSS	Backend Support Systems
CAPEX	Capital Expenditure
CCAP	Converged Cable Access Platform
CD	Chromatic Dispersion
CLI	Command Line Interface
CMTS	Cable Modem Termination Systems
DOCSIS	Data Over Cable Interface Specification
DWDM	Dense Wave Division Multiplexing
DSP	Digital Signal Processor
FEC	Forward Error Correction
IEEE	Institute of Electrical and Electronic Engineers
HFC	Hybrid Fiber Coax
ITU	International Telecommunication Union
L0	Layer zero
L2TP	Layer 2 Tunneling Protocol
L3VPN	Layer 3 VPN
LTE	Long Term Evolution
LTE-A	LTE-Advanced
MBH	Mobile Back Haul
MNO	Mobile Network Operator
MSO	Multiple System Operator
MTSO	Mobile Telephone Switching Office
NMS	Network Management System
OAM	Operations, Administrations and Maintenance
OEO	Optic-Electrical-Optical
OSI	Open Systems Interconnection
OTDR	Optical Time Domain Reflectometer
OTT	Over-The-Top
PMO	Present Mode of Operation
QoE	Quality of Experience
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RPHY	Remote PHY
SDN	Software Defined Networking
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNR	Signal-to-Noise Ratio
VLAN	Virtual Local Area Network
VPN	Virtual Private LAN
Web UI	Web User Interface
WDM	Wave Division Multiplexing

Bibliography & References

IEEE 802.3-2008, “CSMA-CD access method and physical layer specifications”

IEEE 802.3ba-2010, “CSMA/CD Access Method and Physical Layer Specification Amendment 4: Media Access Control Parameters, Physical Layers, and Management Parameters for 40Gb/s and 100 Gb/s Operation”

ITU-T G.652 (11/2016) Series G: Transmission Systems and Media, Digital Systems and Networks, Transmission media and optical system characteristics – Optical fibre cables

ITU-T G.653 (07/2010) Series G: Transmission Systems and Media, Digital Systems and Networks, Transmission media and optical system characteristics – Characteristics of a dispersion-shifted, single-mode optical fibre and cable

ITU-T G.694.1 (02/2012) Series G: Transmission Systems and Media, Digital Systems and Networks, Transmission media and optical systems characteristics – Characteristics of optical systems, Spectral grids for WDM applications: DWDM frequency grid

ITU-T G.709/Y.1331 Series G: Transmission Systems and Media, Digital Systems and Networks, Digital terminal equipments – General; Series Y: Global Information Infrastructure, Internet Protocol Aspects and Next-Generation Networks, Internet of Things and Smart Cities, Internet protocol aspects – Transports; Interfaces for the optical transport network

SFF-8083 SFF Committee, Specification for SFP+ 1x 10 Gb/s Pluggable Transceiver Solution (SFP10), Rev 3.1 September 13, 2014

SFF-8665 SFF Committee, Specification for QSFP+ 28 Gb/s 4x Pluggable Transceiver Solution (QSFP28) Rev 1.9 June 29, 2015

Insight-Driven Network Performance Management and Protection in the Cloud/IoT Era

An Operational Practice prepared for SCTE•ISBE by

Tony Kourlas
Director Product Marketing, ION
Nokia
600 March Road,
Kanata, Ontario
Canada K2K 2E6

Introduction

As streaming video, intelligent cloud and the Internet of Things (IoT) applications begin to dominate today's networks, they bring with them new challenges for service providers (SPs) to address. Subscriber demands for a perfect streaming and cloud experience is creating explosive growth in network bandwidth and complexity. A new generation of distributed denial of service (DDoS) attacks originating from cloud and IoT sources is bringing down critical parts of public network infrastructure. Any upsets in service increase customer dissatisfaction and churn.

At the root of the problem is lack of the visibility and control necessary to identify and resolve cloud/IoT network issues quickly and cost-effectively. Operators have petabytes of data at their fingertips. However, this data is collected in silos across multiple systems, requiring them to manually combine and correlate billions of data points amassed, and then organize them into a coherent report so they can try and sniff out issues. This process has been wrought with human error, and has not provided enough data to see exactly where problems lie. The primary tools of this process, deep packet inspection (DPI)-based appliances, were simply not designed to deal with cloud/IoT scale and complexity. In the end, problem resolution has become a costly guessing game that rarely resolves problems quickly enough to keep customers happy. That has led to ballooning capital and operational costs.

The industry has responded by evolving IP network analytics. A new generation of software-only solutions can ingest, combine, and correlate petabytes of siloed data from network, enterprise and cloud sources to provide a holistic view of the entire network and the applications that flow through it – in real time. SPs can, without hardware probes, track applications and services – not just at certain points in the network, but end-to-end, across 100 dimensions at the same time.

Network intelligence has also become actionable in real time. These solutions also provide SPs with the tools they need to quickly act on this data by creating baselines and triggers that alert on anomalies. Data from the network, data center and wide area network (WAN) can be used to trigger processes or real-time policies that increase customer quality of experience (QoE) and network security, while decreasing overhead and customer churn.

1. SP challenges and Needs in the Cloud/IoT Era

1.1. Ensuring Customer QoE

Cloud applications and services – including Netflix, Hulu, Twitch, YouTube and Facebook – make up more than 60 percent of network traffic today, and are expected to rise will rise to 80 percent by 2020, according to Nokia Bell Labs.¹ Yet in this environment, providers have very limited insight into which applications are running on their networks, and what impact this application traffic is having on performance and subscriber satisfaction. If their subscribers do not receive high-quality streaming, they will complain and eventually switch ISPs in search of a better experience. For them, “slow” is the new “down” when it comes to streaming speed and quality of over-the-top (OTT) content traffic. The *Nokia*

¹ *Bell Labs Consulting Inaugural Mobility Report* <https://pages.nokia.com/1503.bell-labs-mobility-report.html>

Acquisition and Retention 2016 Study found that internet quality is the most important driver for retention, and that those consumers most unsatisfied with that quality are more likely to churn.²

At the root of the quality problems is network congestion. To fix this, internet service providers (ISPs) have typically thrown more bandwidth and caches into the network, but they've done so blindly, making this an expensive and ineffective proposition.

Conventional methods for assessing network performance, including DPI, can be extremely expensive to deploy and scale, with limited visibility. DPI hardware cannot keep up with accelerating speeds and feeds, and is too expensive to deploy network-wide, which means it does not see much of the traffic flowing to and through a network. Because the classic DPI approach dissects every single packet in its path to see what's inside, it is completely blind to over 50 percent of network traffic. The end result is that ISPs lack the data they need to pinpoint underperforming areas of the network. They know there is a problem because subscribers have complained, but they have no idea how much capacity is needed, or where to place it. Time and money is typically wasted on a trial and error process—deploying bandwidth and caches effectively in the blind—with the hope of solving the problem. The end result is that providers typically fail to resolve service issues in a timely manner, and they fight a losing battle to keep customers happy.

Getting the visibility necessary to ensure optimal QoE lies not in tens of thousands of expensive hardware probes, but in software-based solutions that can scale to the largest networks and provide real-time, multi-dimensional (cloud and network) data. Such a solution must monitor tens of thousands of popular cloud applications and services in real time, and run analytics that track how this traffic flows to and through networks to reach subscribers – without the need for expensive probes, taps and monitors. It must then combine this multi-dimensional visibility and analytics with the ability to create alarms and trigger policies that ensure superior performance management and customer QoE.

1.2. Securing Networks from DDoS Attacks

Driven by IoT security holes and 10G cloud server uplinks, DDoS attacks are growing in frequency and intensity. In these attacks, hackers set their sights on the source of the connectivity to disable as many end users as possible, as quickly as possible. They hijack thousands of unprotected IoT devices and cloud servers, and leverage the combined flows to spin up terabyte-level attacks that can disable entire data centers in moments

Verisign's *Q1 2017 DDoS Trends Report* identified a 23 percent decrease in the number of attacks in Q1 2017,³ with average peak attack size increasing 26 percent compared to the previous quarter. Peak sizes were over 10 Gbp/s, while multi-vector attacks peaked at over 120 Gbp/s and around 90 million packets per second. DDoS attacks overall are expected to reach 100 million by 2019, up from 50 million DDoS attacks in 2016, according to Cybersecurity Ventures.⁴ These attacks create a firestorm of outages, a deluge of angry customers flooding call centers and a serious impact on profitability.

² *Nokia Acquisition and Retention 2016 Study*

<http://www.mediatelecom.com.mx/~mediacom/media/pdf/adquisition-retention-nokia-2016.pdf>

³ *Verisign Q1 2017 DDoS Trends Report*

http://forms.verisign.com/Q12017DDoS TrendsReport?utm_medium=Blog&utm_term=internal

⁴ *DDoS Attack Report*, Cybersecurity Ventures <http://bit.ly/2uA0b8A>

As networks are threatened with the skyrocketing strength and intensity of such DDoS attacks, the immediate reaction for many operators is to disregard their previous investments and accumulate large collections of costly mitigation-specific hardware. The problem with that approach is that these appliances operate by sending all traffic through scrubbing centers, yet some vendors charge based on tonnage inspected—not traffic mitigated. Without proper detection, scrubbing will grow with the network while the expense of traffic monitoring skyrockets.

The same big-data visibility and context that now allows SPs to proactively ensure reliable network performance can also be utilized to surgically mitigate networks from DDoS attacks. A software-based multidimensional analytics approach to network security is the best defense, as it automatically combines many sources of real-time streaming datasets from the network to detect and mitigate DDoS attacks in seconds. This broad visibility into network and cloud behavior also serves to minimize false positives and false negatives. It has the cloud intelligence required to recognize when surge in internet traffic is an attack, and when it is just a normal behavior from cloud sources such as Facebook or Amazon. Because SPs understand how cloud applications and services flow to and through networks, they quickly can identify the presence or potential for DDoS attacks, along with the cloud and IoT sources that are responsible for them. Software alerting drives action – routers can be called upon to drop traffic at the edge, or traffic can be sent to mitigation devices for more stateful analysis. (Figure 1.)

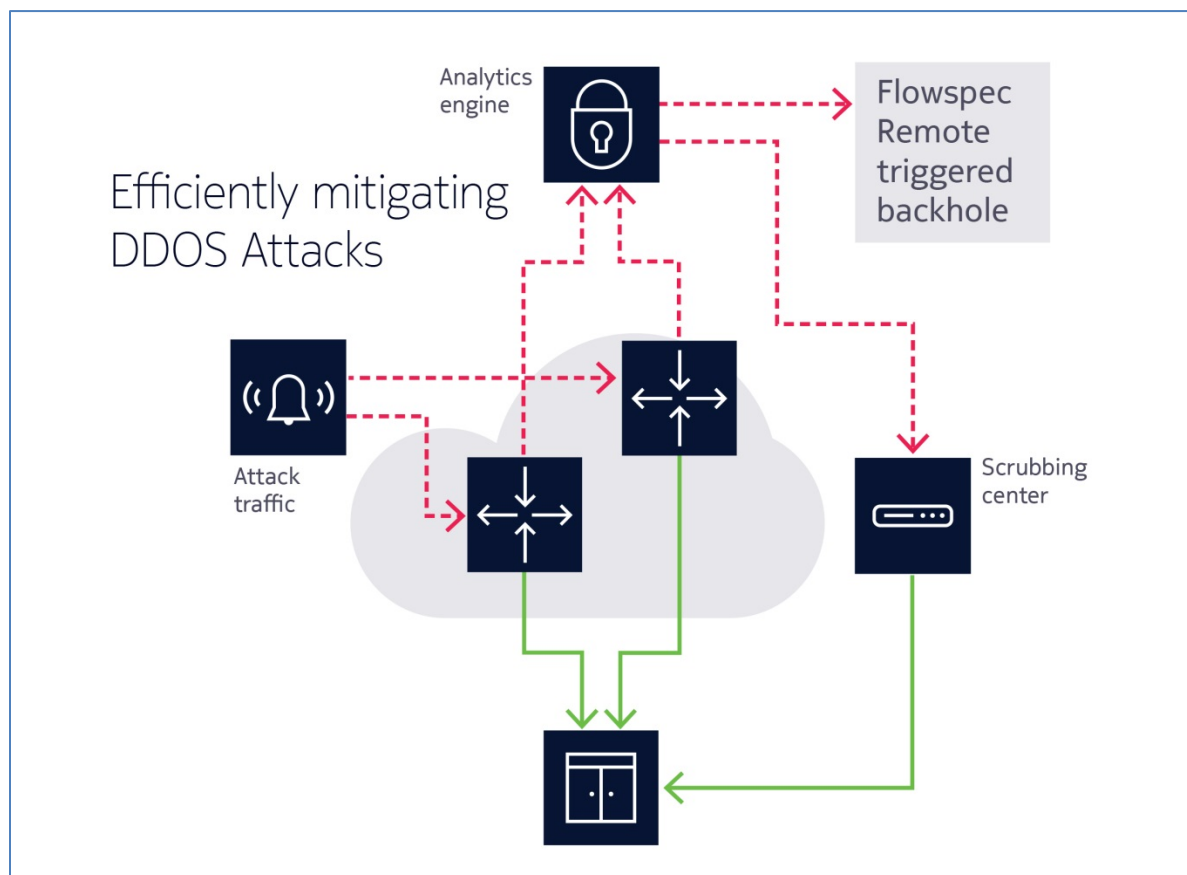


Figure 1 – Real-time insights allow ISPs to trigger on specific events and selectively mitigate DDoS attack sources.

1.3. Harnessing Big-Data Analytics to Drive SDN/NFV

Solving the network and service visibility problem requires coupling big data analytics with the dynamic control capabilities of software-defined networking / network function virtualization (SDN/NFV) orchestration platforms. Together, these elements become the cognitive “brain” that is capable of making real-time, automated corrections to critical networks so they can quickly adapt to changes in application demand, flow and traffic patterns, as well as automate actions that ensure ongoing service health and customer satisfaction. This all allows SPs to drive greater network efficiency, help assure quality and enhance security – all without manual intervention, and in real time.

Central to this process is automated monitoring, flagging and mitigation – the capability to set alerts on performance drops by any number of multidimensional variables, including customer segment, time of day when average bit rate (ABR) drops for a group of customers, and other key performance indicators. This requires rich analytics to drive the dynamic network configuration of software-defined SDN/NFV controllers.

2. The New Dimensions of Network Intelligence

Operations teams today typically have a lot of toolsets that primarily are focused on low-level sensing pieces, for example, “Is this interface up or down? How many tetrabytes are on the network currently? How is the CPU utilization?” and other variables. These are all useful metrics, but they don’t divulge any business-level data such as how Netflix is streaming in Peoria, or whether Chicago is up or down, or whether you are seeing something different on this peer from what you saw yesterday. In isolation, these low-level metrics don’t answer anything about customers’ services, as they aren’t mapped to today’s business questions.

Effectiveness in today’s business and technical landscape requires customizable network intelligence that offers several dimensions of capability:

- **Network Intelligence:** Visibility into tens of thousands of applications, without probes, in order to fully enable network optimization and informed discussions with content/network partners.
- **Service Intelligence:** Service assurance for cloud (OTT) applications (“How is Netflix or YouTube doing?”), while improving customer QoE and reducing churn, thereby lowering troubleshooting and support costs.
- **Subscriber Intelligence:** Visibility into granular usage dimensions such as daily traffic by subscriber (tonnage, category, site), which subscribers go over their plan tier and by how much, and use patterns – providing essential insight for how to give better service, bill more accurately, retain customers and predict cord cutters.
- **Security Intelligence:** Deeper, cloud-aware analytics, better attack detection, and more precise attack mitigation, minimizing the need for scrubbers, and eliminating the need for detection hardware.
- **Operational Intelligence:** Actively leveraging the same visibility and insight provided by all of these tactics to create custom alerts on business-level events.

All of these multidimensional analytics drive SDN/NFV operations, enabling dynamic network optimization and creating new service opportunities that will allow ISPs to better align network performance with customer expectations.

This new paradigm allows providers to bring topological and logical attributes to that lower-level data. Script-based application programming interfaces (APIs) allow them to answer questions about customers and their service, or above that, business-level objects, such as the performance of the OTT service as just a portion of the traffic – for example, Hulu on a certain node, or the YouTube peer. Simply being able to answer questions like that, and alert on abnormalities, is fundamentally a game-changing way of looking at the data. (Figure 2.)

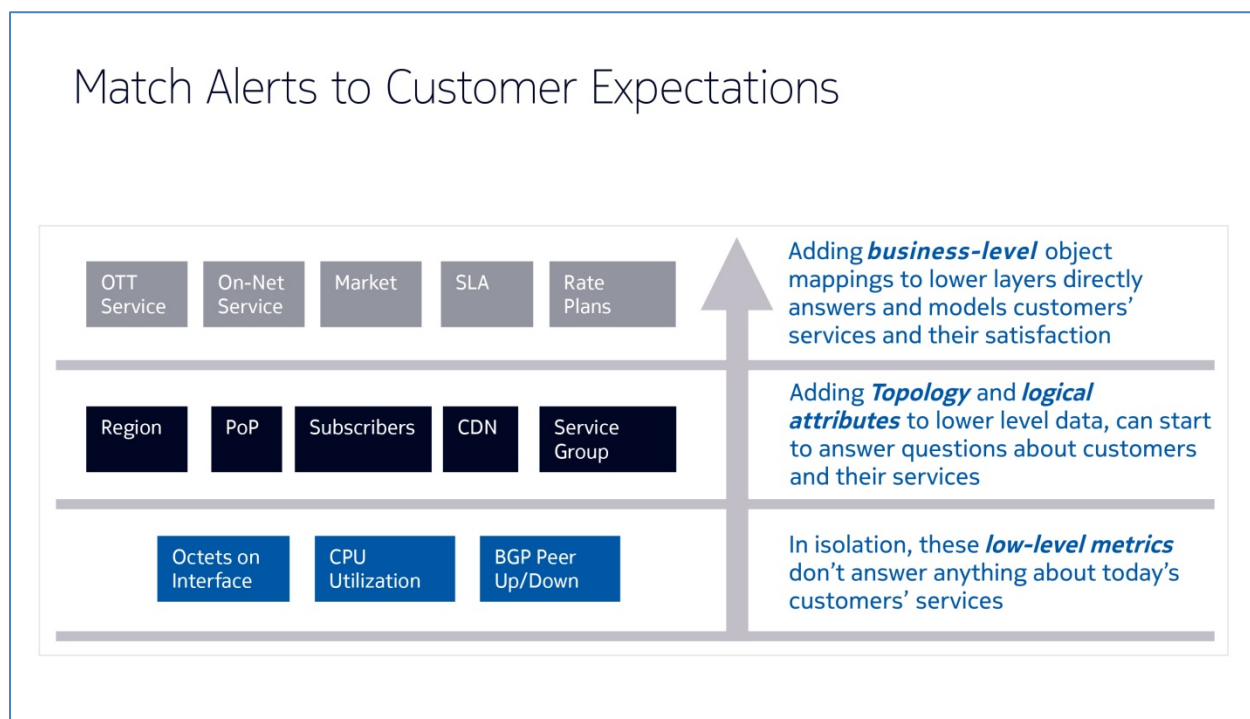


Figure 2 – Networks can be dynamically configured to better match alerts to customer expectations.

By employing an automated, dynamically configured network using a highly scalable software solution to leverage big data allows ISPs to build real-time alerts based on any number of those variables so that they can be proactive rather than reactive. Ultimately, they will also be able to create automated responses as well. The end result is a 360-degree solution that better understands the network to see what changes occur, and then repositions or reconfigures the network to make it more optimal for users.

2.1. A Modular Approach

One effective way to achieve this level of network intelligence is to combine several modules that work seamlessly together. An effective architecture includes a big-data engine/software platform; a massive map of the global service supply chain that adds visibility to all applications built onto the core platform; an analytics application that provides end-to-end network visibility and context-aware content

engineering; an analytics application that monitors customer QoE in real time; and a security module to enable real-time DDoS detection and mitigation.

A petabyte-scale big-data analytics engine can provide visibility into vast numbers of cloud applications and services, along with billions of IP addresses, tracking how traffic runs to and through networks to reach subscribers, in real time, and without the need for expensive probes, taps and monitors. Utilizing an ultra- high-capacity processor, it can analyze data against dynamic baselines based on network traffic, then provide alarms and, using existing router infrastructure to mitigate a bulk of traffic via flowspec or remote-triggered black hole, then take necessary actions when normal traffic flow and usage is disrupted.

This level of real-time, analytics-driven network and service automation can provide ISPs with greater network and application insight, control and DDoS protection.

2.2. Improving Network Performance

This approach takes advantage of information that is already available in internet infrastructure. By studying cloud applications and services, providers can unravel their supply chains to see what the related IP addresses are, where they're located and how they interact. When an IP flow reaches their network, those providers won't need DPI to tell them what application or service it is, how it landed on their peering router or how it traverses their network.

Armed with this data, operators can overlay this information onto their own topology to understand what traffic is on the network, where it is and its impact, even if it is encrypted—a level of visibility that's imperative for accurate performance management and quick identification of configuration issues. This approach also enables network issues to be resolved before customers complain about poorly streamed content.

2.3. Leveraging Distribution

This new level of dynamic network management employs multiple processors on a single multi-core unit – more on a single server than used to be available across an entire country. They operate as horizontally scalable clusters, providing resiliency that extends beyond what happens on any single server. With horizontal scaling, the cluster can adapt the resources and the compute as ISPs scale their services. This distributed file system provides the advantageous properties of sharing and replicating the data, fault tolerance and redundancy. Under this architecture, data can be *normalized*, meaning that domain name server (DNS), performance and topology data is converted into vectors, which then can be distributed across the streaming database, allowing all of the different parts of the cluster to operate on that data in parallel, both on the ingest and on the query. With these many dimensions at their fingertips, operators have the context needed quickly dive in deep and decipher service issues in an extremely flexible fashion. (Figure 3.)

Ultimately, this approach accomplishes three things:

1. Flexible pane-of-glass views across multiple data sources, enable extremely intelligent and deep analytics for an entire infrastructure.
2. Real-time correlation and an intelligent view into the data supports dynamic configuration.
3. Real-time alerts based on triggers and dynamic baselining allow operators to act on network drops or DDoS attacks in seconds.

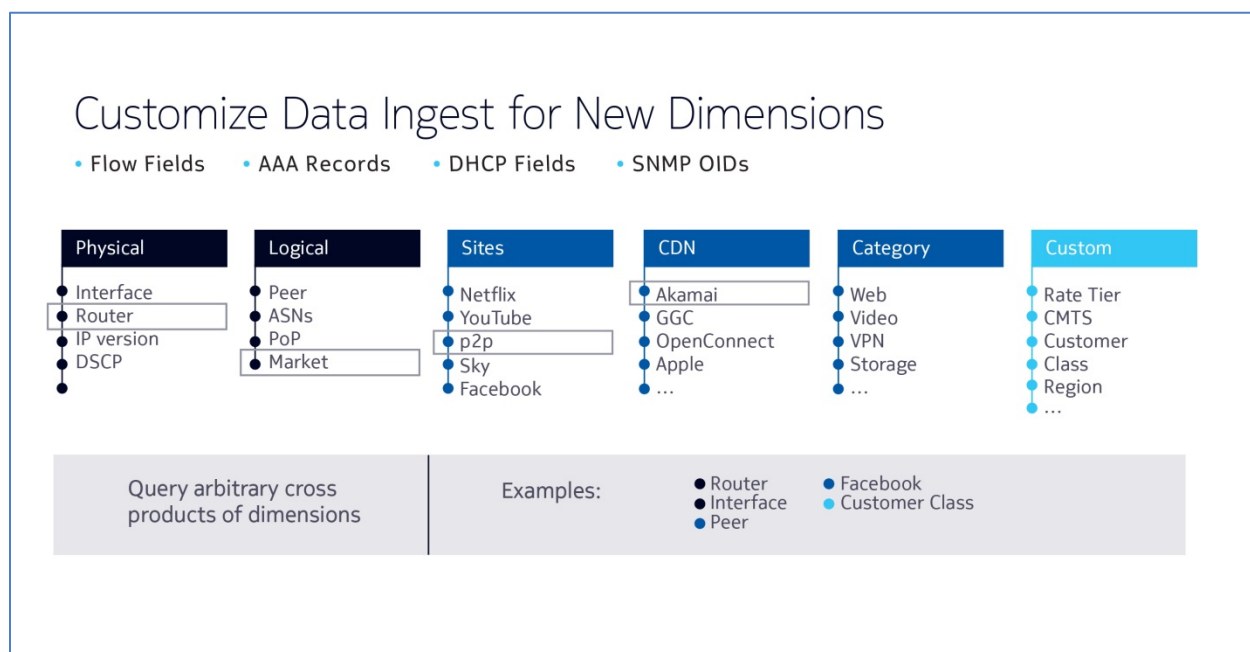


Figure 3 – Providers can use APIs to customize data queries across multiple dimensions.

2.4. Becoming Proactive Rather than Reactive

Typically, customer support teams actively monitor the network, as they have for many years, and when they determine that it is “green,” they return to other daily tasks. In the current era of cloud and IoT, however, it is not always enough to simply see that routers are behaving correctly – there could be higher-level issues afoot that are not able to be tracked with old monitoring tools. One example is a network operator that went through the complete checklist of tasks needed to ensure the network was in order, yet turned to Twitter to learn that an entire market had lost their access to a popular OTT service.

This caused an uncomfortable scramble to solve an issue far after it had negatively affected a multitude of customers. This is a scattershot, reactive approach that does no favors for the customers or the business. Using the new model of actionable intelligence, operators receive alerts through external systems such as Webhook, email, syslog, and simple network management protocol (SNMP) traps. They also can utilize flexible alert settings using baselines, trending and thresholds, creating a proactive rather than reactive service environment.

As an example, consider the situation where there is a set of data on a router, providing a baseline for operation and performance. Traditionally, the provider has had to pull that back every 5 to 10 minutes. Now the industry is moving into telemetry streaming, where the provider will set some kind of persistent query that can put into the network, setting alerts for anomalies. One typical set of parameters would essentially translate to “Look for all the video data between Houston and Washington, D.C. and flag any kind of anomaly in that particular part of the network.” The ISP will have an instant alert for that. When troubleshooting the root cause, operators can set alerts for Webhook, email, syslog, SNMP trap and other variables – whatever works in the system from a forensics perspective. The idea is that the

operator can build up different kind of queries to answer different questions about what has changed in the network, or what will trigger an alert and some kind of operations response. (Figure 4.)




Figure 4 – Alerts based on a vast number of variables can monitor network health and trigger focused, proactive resolutions.

3. Use Cases

Many SPs are already benefiting from enhanced real-time network analytics to increase subscriber satisfaction, reduce churn and secure their networks (Figure 5.) A few examples follow.

3.1. Enhance Video Performance Drops by Market

As video content consumes increasing amounts of available bandwidth, ISPs must build out their networks to keep subscribers happy. However, they have a severe lack of visibility that prevents them from solving or proactively avoiding problems. ISPs need a way to understand how applications perform across all parts of the network so they can quickly identify business level events such as “Netflix is down in Chicago” for quick remediation.

Unfortunately, IP network analytics were not built with web-scale services in mind. ISPs have historically collected both application and network data. But they have stored this data in silos, and have not had sufficient cross-correlation to identify the specific OTT streaming issues caused by congestion on a particular link. This lack of insight has made troubleshooting a very costly and inefficient process that has done little to improve overall service quality.

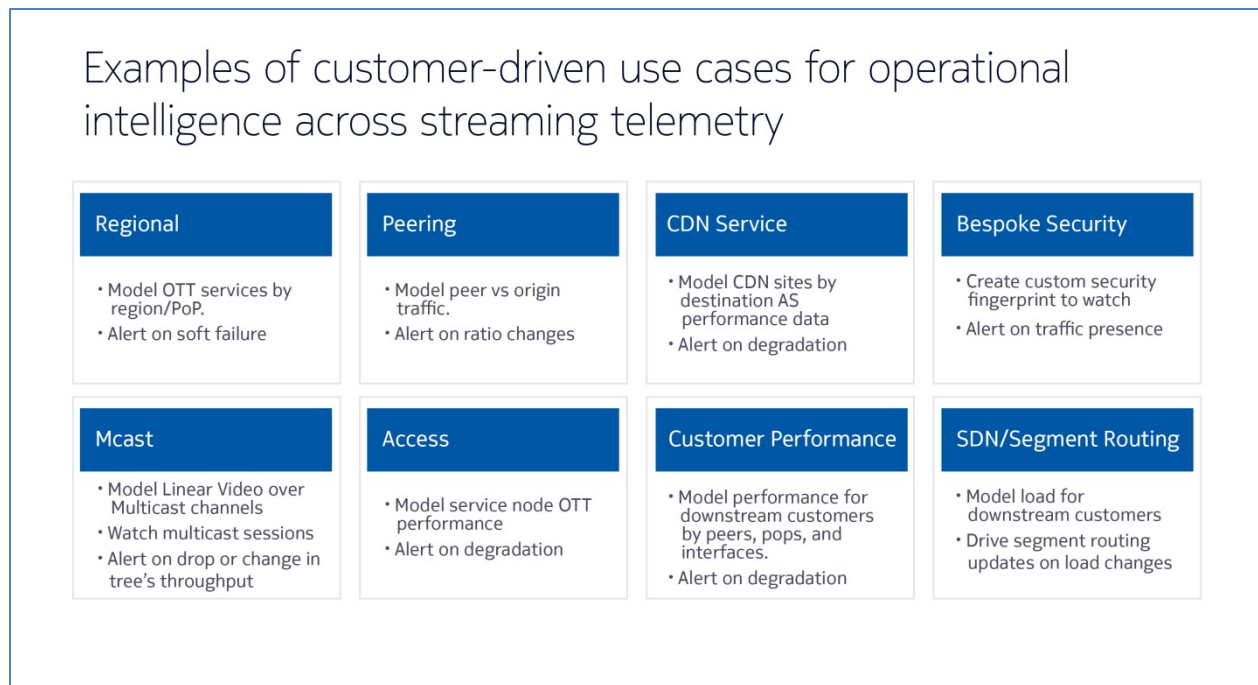


Figure 5 – Leading ISP and OTT providers are already using multidimensional intelligence for understanding and dynamically managing their networks.

One top ISP attempted to solve the visibility problem with DPI appliances. However, it was far too expensive across its entire network, and large swaths remained unmonitored. Since DPI is blinded by encryption, it provided no visibility into well over half of all video traffic flowing through the network. The ISP recognized that it needed a new software-based solution that could provide full multidimensional visibility and be far more cost-effective than customized hardware.

This ISP addressed the challenge with a massively scalable software-based solution, which identified applications and rapidly correlated those findings with all network data to immediately identify problems – all without looking at a single packet. This multidimensional analytics approach enabled the ISP to instantly monitor high-level events and visualize the impact of video traffic across any part of its network, alerting on any drops in any streaming degradation, and allowing speedy resolution of issues before customers had the chance to complain.

Using this solution, the ISP can now categorize every single flow to gauge and dramatically improve network performance by surgically adding the exact amount of bandwidth needed in just the right places. This ensures that all subscribers are receiving their desired content with the best possible quality. (Figure 6.)

Cloud intelligence: Network visibility to resolve delivery problems in moments

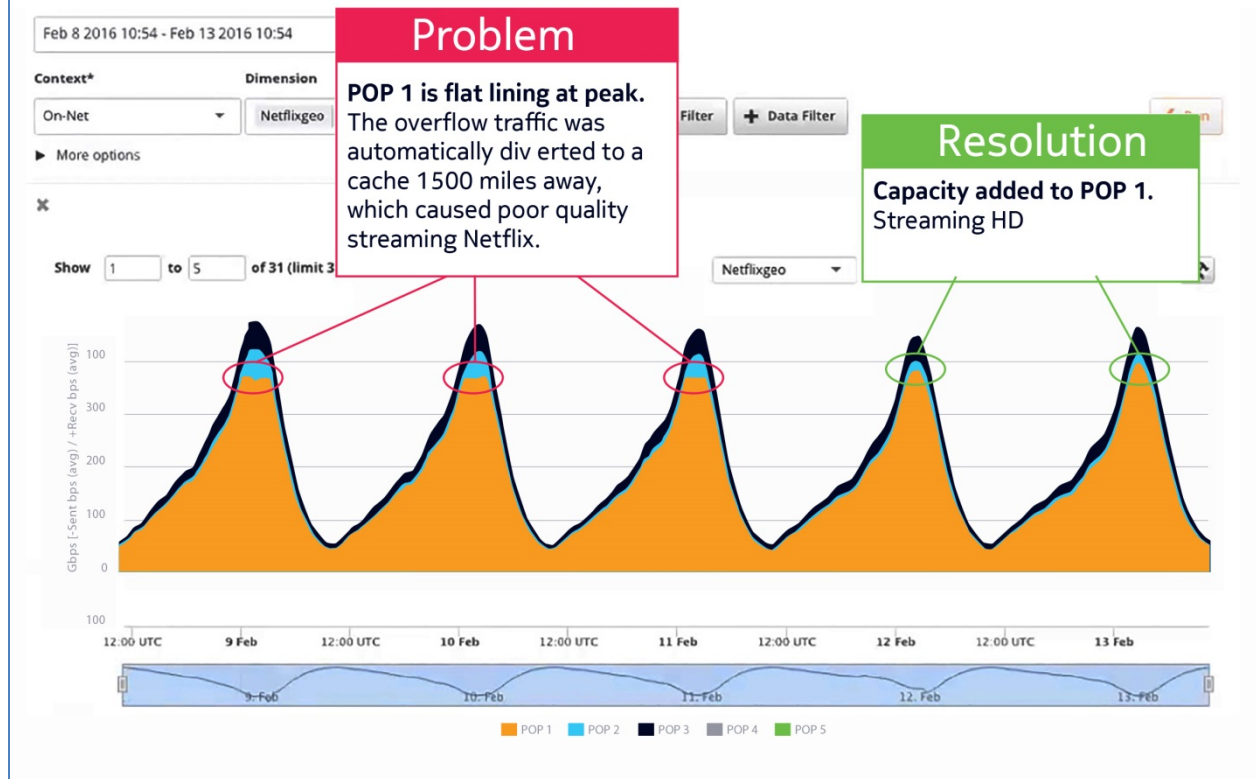


Figure 6 – SDN-based operational intelligence improves visibility to instantly identify and solve problems.

3.2. Take Immediate Action on Peering Misconfigurations

Historically, ISPs peered directly with one another to exchange traffic, and their contracts were based on tonnage exchanged. As the internet has become more complex, so have peering agreements. Ultimately, everyone involved is working together to deliver content to the end user when and where they want it, at the best quality. If peering relationships do not follow their contractual agreements, it could easily lead to clogged interfaces and poor streaming quality for the end user. Because of this, contracts detail what traffic can traverse which interface and when. It is imperative for an ISP to have a comprehensive view of its network that understands what traffic is flowing where if it is to actively monitor these relationships.

One ISP deployed a solution that alerted on configuration changes, and has found it to be highly beneficial. In one instance, Dropbox was just emerging, and at first used Amazon Simple Cloud Storage Service (S3) to distribute its content. However, as it grew, Dropbox decided build its own content delivery network (CDN), and began distributing traffic from both. This would look like suspicious brand-new traffic to a network that could not map and understand the change in IP flows. However, in this case an

operator could immediately see that it was the same traffic under a different name, and marked it as a valid traffic shift.

The same approach that identified all traffic and where it was traversing the network immediately alerted network operators to a change in peering traffic. There was a misconfiguration that caused peering traffic to enter the ISPs network on the incorrect interface, leading to a sharp drop in traffic entering at one port while flooding another. The real-time alert allowed for immediate intervention.

This same capability also can monitor the overall health of peers, alerting ISPs when performance thresholds are not met. (Figure 7.)




Figure 7 – ISPs can monitor peer health by setting a simple low watermark threshold.

3.3. Leverage Insight-Driven DDoS Mitigation

DDoS detection appliances are struggling to keep up with the increasing frequency, size and complexity of attacks. Only a software solution that leverages the same visibility to alert on performance drops can provide the necessary insight to protect even the largest networks while lowering OPEX.

By combining multiple sources of data to contextualize and identify the exact origin of all traffic traversing a network, software can most accurately flag and categorize that which is associated with a DDoS attack. Then, it can whitelist traffic that has been known to create false positives (such as Facebook traffic from a geo-location at a given time). This, in turn, drastically reduces the amount of traffic that must be monitored, and the amount that would be mistakenly backhauled to scrubbing centers. This further reduces the need for costly, specialized devices.

The benefit of accurate detection is easy to see in an example such as this. A digital video recorder (DVR) might sporadically send one or two DNS requests, and a system with this aptitude understands the context in which a DVR must operate. Therefore, if a DVR suddenly begins flooding a DNS server with requests, the software will create a real-time alert for security operators to investigate. The operator then has the choice of which type of mitigation should be used. It can be dropped at the edge with router mitigation, or a portion of the traffic can be diverted to a partner scrubbing appliance to be cleaned and reinjected into the network.

This new and innovative concept is able to immediately identify attackers based on intuitively derived traffic fingerprints so the attack traffic can be cut off at the edge of the network – before it affects a single target.

4. Conclusions

The internet is constantly evolving. We have gone from T1 to the era of cloud computing. We are seeing explosive growth, including that of IoT and new security challenges. This is the most interesting time there has ever been in terms of how quickly the nature of traffic is changing, and how dynamic network management has become.

We live in the world where disk, memory, CPU and computation are essentially infinite, and where market priorities have radically changed. Traffic management now is not just looking at the peer, but at dynamic data flows coming from across the world. A problem may be generated though a vast number of internal or external sources, whether from OpenConnect, traffic flowing through metro-scale datacenters, or even with a DDoS attack. The bottom line: data traffic now is far more dynamic and requires new sets of technologies to understand it.

All of this means that ISPs need visibility and operational intelligence throughout the network – not just at a single point, but end-to-end, with the capability to radically scale capacity in the most cost-efficient manner.

Networks now are competing on the management capabilities of those running them, the cost of operations and the ability to drive superior customer QoE. Effective management and protection is not just about managing data traffic in one or two dimensions, but about how ISPs build visibility into the edge data centers, the level of inner-connection in the edge, and their focus on customers.

With a new approach that considers every piece of the network puzzle, service providers can more accurately, cost-effectively and efficiently manage their networks. That, in turn, will save them a substantial amount of money by doing away with unnecessary buildouts, protecting critical assets and reducing customer churn.

Abbreviations

ABR	average bit rate
AP	access point
API	application programming interface
BGP	border gateway protocol
CDN	content delivery network
DDoS	distributed denial of service
DNS	domain name server
DPI	deep packet inspection
DVR	digital video recorder
HD	high definition
IoT	Internet of Things
ISP	internet service provider
NFV	network functions virtualization
OI	operational intelligence
OTT	over the top
POP	point of presence
QoE	quality of experience
RTBH	remotely triggered black hole
SDN	software defined networking
SNMP	simple network management protocol
S3	Amazon Simple Cloud Storage Service
SP	service provider
VSP	virtualized services platform
WAN	wide area network

Bibliography & References

Bell Labs Consulting's Inaugural Mobility Report <https://pages.nokia.com/1503.bell-labs-mobility-report.html>

DDoS Attack Report, Cybersecurity Ventures <http://bit.ly/2uA0b8A>

Nokia Acquisition and Retention 2016 Study
<http://www.mediatelecom.com.mx/~mediacom/media/pdf/adquisition-retention-nokia-2016.pdf>

Nokia FP4 Routing Chipset, New Routers and New Operations Methods; Appledore Research Group
<https://resources.ext.nokia.com/asset/201327>

¹*Verisign Q1 2017 DDoS Trends Report*
http://forms.verisign.com/Q12017DDoS TrendsReport?utm_medium=Blog&utm_term=internal

Sustained Throughput Requirements for Future Residential Broadband Service

Traffic Model for Bandwidth Estimates

A Technical Paper prepared for SCTE•ISBE by

Jeroen Wellen

Senior Member of Technical Staff
Bell Labs Consulting/Nokia
Antareslaan 1, Hoofddorp
The Netherlands
Jeroen.Wellen@bell-labs-consulting.com

Prudence Kapauan

Distinguished Member of Technical Staff
Bell Labs Consulting/Nokia
1960 Lucent Lane 9D-117
Naperville, IL 60563
Prudence.Kapauan@bell-labs-consulting.com

Amit Mukhopadhyay

Partner
Bell Labs Consulting/Nokia
600 Mountain Avenue
Murray Hill, NJ 07974
Amit.Mukhopadhyay@bell-labs-consulting.com

Introduction

The access network is the costliest investment for any operator and it is also the hardest to evolve to next generation technologies. Incumbent operators try to prolong their investment with minimal alterations to the “last mile” connection while new operators try to make their investments future proof as far into the time horizon as practical. While Fiber to the Home (FTTH) is the unquestioned leader from a performance perspective, the economics of the solution works out only in select morphologies and market conditions. Access technologies like Hybrid Fiber Coax (HFC) and Digital Subscriber Line (DSL) comprise the majority of the residential broadband connections today and Fixed Wireless Access (FWA) is also gaining momentum in certain markets. The challenge for DSL is that it can match FTTH performance for only limited loop lengths. HFC is largely constrained by the fact that it operates in a shared medium environment. FWA shares both the limitations – distance as well as shared medium.

The key is to balance the market-specific technology needs and the business justification of the capital and operational expenses. In other words, a technology that may be suitable for delivering the services to a particular geographic and demographic morphology, may be completely unsuitable for delivering a similar service to another morphology. Access networks are considered long term investments and operators make decisions based on revenue potentials and thus Return on Investment (RoI) over a relatively longer period of time is generally acceptable.

Choosing the right Access Network technology at the right moment is one of the most challenging decisions faced by Network Operators. How long the installed base can support the ever-growing demand of applications, especially video, and what technology to install or upgrade to, and when, are the key questions to be answered. HFC and FWA are particularly attractive technologies since they connect multiple homes to a single Access Node (a fiber node in HFC and a cell site for FWA) through a shared medium, allowing to benefit from stochastic multiplexing while minimizing drop costs. But the advantages of these technologies depend on how many homes can be connected per node. And, to assess the techno-economics of the solution, proper estimation of the expected traffic is needed – at present and during the foreseen lifespan of the investment.

This paper presents a model to estimate the 10 years maximum sustained throughput requirements for residential broadband services. The model is used to estimate the maximum throughputs during busy hour and to assess how upper throughput percentiles compare to the mean expected value. In addition, the impact of other effects such as disruptive video applications and loads during special events are investigated. The main purpose of the paper is to provide guidance to operators investing in either upgrading their current network or investing in new network technology. While multi-gigabit access connections often catch the headlines, that is not the sole criteria around which access networks need to be designed. It is imperative that shared medium technologies provide such peak speed connections to support demanding applications. In many circumstances, it is equally important to provide high quality sustained throughput connections at considerably lower speeds. The service targets and the choice of technologies are driven by market economics.

Study

1. Traffic Model

Although traffic demand is basically no more than just the downstream and upstream information rate, it can be viewed at various time scales for various objectives. For the purpose of network design, peak or line rates will determine the maximum distance or range of a technology and thus, by observing the household density, the required number of Access Nodes to cover a geographic area. But to assess equipment requirements per distribution point, or how many radio channels, cable modem terminations or uplink ports are needed at each Access Node, the Maximum Sustained Throughput must be estimated for heavy load conditions. To establish this value, estimations of mean values is not sufficient in that it provides no insight in the variance that is likely to occur during busy hour. A more advanced statistical method is required to cover excess situations that may occur with a specific probability. This was acknowledged in [1], where Monte Carlo methods were applied to forecast aggregate subscriber demands. To target specific Access Network scenarios, the model described here directly relates the number of simultaneous streams to local demographics and behavioral aspects such as sharing of devices and secondary device use. Moreover, here the impact of immersive Virtual Reality (VR) is targeted as an application that may disrupt bandwidth consumption in the coming decade.

What exactly is targeted by the Maximum Sustained Throughput is illustrated in Figure 1, where the hourly average traffic levels (grey on the left) is considered at the busy or busiest hour of the day (blue). Within that hour, we look for the minimum traffic level at which the users can make use of all network services for a specified quality of service, i.e. either with sufficient speed to ensure glitch-free video streams, or a web page refresh within a specified response time (green). Note that this Sustained Throughput level is resolved by an instant packet stream at a lower level, typically on/off switched at the peak or line rate of the access network (red on the right) or, most commonly, the subscribed rate. The Maximum Sustained Throughput is obtained from a percentile, in this study the 99th, of Sustained Throughput levels during busy hour.




Figure 1 - Definition of the Sustained Throughput and its relation to daily average throughput (left), hourly average and instant rate (right).

Forecasting the load per home connection is difficult, as it is based on a combination of various devices and user applications that evolve constantly over time, and user generations, old and new, that change their habits and preferences. Although national or global traffic forecasts can provide a high-level glimpse of the future by extrapolating historical growth rates [2], they are insufficient for access network designs. Access networks typically depend on features of (future) technologies that, in addition, are commonly deployed selectively to meet specific local markets and geographic conditions. Network usage by users has many dependencies whose effect is hard to determine today, let alone 10 years from now. Incorporating multiple parameters into the model adds complexity. As these values are difficult to obtain, such a model may not add any more accuracy. A more deterministic model that is tractable could prove more accurate. The model for the Sustained Throughput per household B_{HH} distinguishes between two types of traffic corresponding to the way data users consume data, either through streamed media or instant downloads:

$$B_{HH} = B_S + B_D$$

$$\text{with } B_{S/D} = \sum_{N_{S/D}} R_{S/d} \quad (1)$$

- Traffic from media streams, B_S , is directly consumed by users and originates from multiple active traffic streams N_S at a rate R_S . Typically, this involves video and audio either on dedicated devices or as part of (web) applications or games. N_S is determined by the number of users and their usage probability which is depends on the duration of use.
- Traffic from N_D consists of instant downloads of content amounting to B_D , and is initiated either by explicit user requests ranging from chat and email to file or web page downloads, or by automated processes such as background storage backup and file sharing. Usage probability here is determined by the duration of bursts rather than the actual use duration.

The Sustained Throughput is obtained by calculating the number of concurrent streams and downloads, multiplied by their corresponding rates. The Maximum Sustained Throughput is calculated for high-load conditions, i.e. during Busy Hour. Since the number of streams and downloads depends on the usage probabilities, not only the expected mean can be quantified, but also median value and higher percentiles.

Note that most services and applications can involve a combination of both traffic types previously described. Game applications can make use of media streams, content downloads as well as periodic updates of scene data, while email has both user initiated downloads and background updates. In any case, the number of active generators N_S and N_D is determined by the number of persons per household and their use of various applications that generate them.

1.1. Population

Household size projection is based on extrapolation of U.S. Census demographic data [3], as shown in Figure 2. Here, the 2016 household size is extrapolated by 10 years using the 2014 to 2016 average growth rates for each size while applying U.S. population projections [4]. This shows a drop from 2.46 persons per household in 2017 to 2.36 in 2027. Note that household sizes may differ significantly from the U.S. average for specific cases involving other countries, states, and various residential areas.




Figure 2 - U.S. Household Size Distribution (left) and Average Size (right) between 2017 and 2027, from [3,4].

1.2. Application Usage

Application usage is applied as a Usage probability p_u per resident and can generally be obtained from published reports from the number of users N_{users} of an application or service divided by the total or studied population N_{pop} , and multiplied by the minutes of use t_{use} per time frame Δt , e.g. minutes per day:

$$p_u = \frac{N_{users} t_{use}}{N_{pop} \Delta t} \quad (2)$$

For the Maximum Sustained Throughput, however, the usage probability during Busy Hour (typ. 8-9pm) $p_u(t_{BH})$ is especially of interest, and the average usage probability per day is modulated by the usage probability per hour. As reported in [5], in 2016, U.S. adults age 18+ spent 6h per day consuming media (average between September 26 and December 25 2016), of which about 5.5h is on TVs and connected devices, 21 minutes on PCs and Tablets and only 2 minutes on Mobile Phones. Neglecting concurrent use of multiple devices for the moment, modulating the average video usage probability of 22% with the hourly profile from [6], we get Figure 3, showing the Busy Hour peak at 9PM of around 52% for TVs, PCs and tablets. Here the data for adults 18+ are scaled to estimate the usage including teenagers. Although daytime use patterns may differ for that age group, it is assumed that this is less the case during the targeted busy hour. The same applies for the actual network load: the referred report does not distinguish media use at home from elsewhere and what mobile traffic is off-loaded to the home network. Given the dominance of TV streaming and the busy hour at 9PM however, it is presumed that the traffic primarily flows through the internet connection of households.




Figure 3 - Media usage (left) and device usage (right) probability per hour p_u per U.S. resident, based on [5,6].

1.3. Streamed traffic

The stream traffic per home connection is modeled as the sum of N_S active traffic streams, each contributing an amount R_s of traffic:

$$B_S = \sum_{i \leq N_S} R_{s,i} \quad (3)$$

Instead of associating traffic generators with devices or users, they are directly associated with active media streams per household. The model does require proper estimation of its two key components:

- The number of active streams N_S , and how it relates to
 - household size N_{HH} ,
 - the usage probability during Busy Hour p_u , but also effects such as
 - sharing of streams, such as radios and TV sets, by multiple users and
 - use of secondary streams on other devices
- The stream rates R_s , for
 - all possible device classes
 - today and in the future.

Streams can be associated with any video, audio or other media application, including gaming content and in-game communications, voice and video communications, and obviously, video. Although any media stream of any application would apply, for the maximum throughput mainly video streams (containing audio) will play a role. Typically, a distinction should be made between Access Network services that provide IPTV or other managed linear and non-linear TV (Triple Play) and those that do not. For this paper, the model assumes full TV services, either managed by the access network provider or from an Over-the-Top (OTT) subscription.

1.3.1. Stream usage

The number of concurrently active streams depends on household size and usage probability. From the data described in section 1.2 above, the video and audio usage probability can be derived as shown in Figure 4, which indicates that at prime-time almost half the population is watching video of some kind.




Figure 4 - Hourly video (left) and audio (right) usage probability $p_u(t)$ for U.S. residents, based on [5,6].

The number of active users N_U is modeled by applying a Binomial distribution with the household size N_{HH} and Busy Hour take rate, or usage probability $p_u = p_u(t = t_{BH})$ according to:

$$P\{N_U = k\} = \binom{N_{HH}}{k} p_u^k (1 - p_u)^{N_{HH} - k} \quad (4)$$

This model can be applied for individual households and, when their N_U can be assumed to be independent from each other, for Access Nodes that connect multiple households. As will be discussed later, this independence may not hold for special popular events.

To translate the number of active persons per household into the number of active streams, the number of shared views of a single stream should be observed as well as the concurrent use of multiple devices by a user (e.g. both TV and mobile phone).

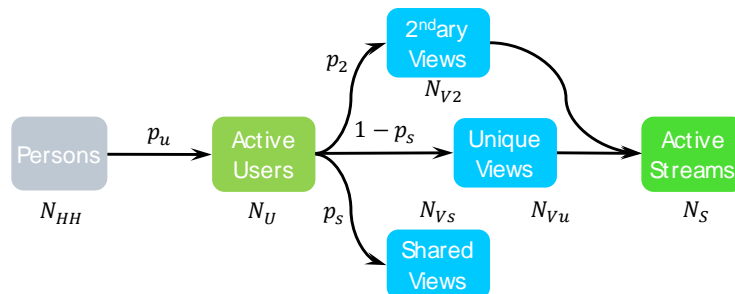


Figure 5 - Probability model to translate persons per household to active streams.

This is modelled as illustrated in Figure 5, where the number of shared streams N_{Vs} can again be modeled with sharing probability p_s for all active users that share a view. Conversely, to get the number of unique

views, N_{Vu} , a Binomial distribution is used with $1 - p_s$, insofar N_U exceeds one (otherwise there is no one else to share with):

$$P\{N_{Vu} = k\} = \binom{N_U - 1}{k} (1 - p_s)^k p_s^{N_U - 1 - k}, \quad N_U > 1$$

$$N_{Vu} = N_U, \quad N_U \leq 1$$
(5)

Although video is still viewed predominantly on TV sets, the number of televisions per U.S. household in fact exceeds the number of persons [7]. Therefore, sharing is assumed to involve no more than 50% of all views. By applying the probability that an active user uses a secondary stream, p_2 , the number of secondary views N_{V2} can be obtained:

$$P\{N_{V2} = k\} = \binom{N_U}{k} p_2^k (1 - p_2)^{N_U - k}$$
(6)

It is assumed that this probability is low, i.e. 5-10%, and that concurrent use of more than 2 streams per user is nihil. The total number of active streams is now simply the sum of unique and secondary views:

$$N_S = N_{Vu} + N_{V2}$$
(7)

Unlike usage probability, the estimates for both secondary usage, p_2 , and sharing, p_s , are rather speculative for now and needs proper evidence to determine to what extent they depend on user behavior, especially that of teenagers. Additionally, projections will depend strongly on new types of devices and future applications.

As indicated by Figure 6, although the average household size is expected to decrease from 2.46 persons today to 2.36 in 2027, a decline in shared usage will keep the number of active streams per household close to 1.3 between 2017 and 2027.




Figure 6 - Projection of persons on streams (mean values, left), accumulating into the active streams (distribution, right) per average U.S. household.

1.3.2. Stream rates

More difficult to predict is the actual traffic per stream. Shared devices, televisions and beamers, typically consume higher rates than mobile devices. Managed linear and non-linear video services often exhibit higher video rates than OTT video services, where adaptive stream rates are commonly used to limit data

consumption and server loads. Tablets and phones on the other hand show higher replacement rates than TV sets, while it is hard to predict the future adoption of virtual reality (VR) devices [8]. For this reason, again a probabilistic approach is chosen to estimate stream rates based on common resolutions and their adoption in the coming years as indicated in Figure 7. For 2017, this distribution is based on 37% of Standard Definition (SD) streams below 1k, and 47% around 2k and 16% 4k [9]. The model does not make a distinction between TV sets and mobile devices, as their resolution is presumed to evolve in parallel. 8k Video is assumed to be adopted as of 2020, while 16k starts in 2025. Note that these higher resolutions would represent Ultra High Definition (UHD) streams to both big-screen TVs, projectors and, for non-interactive streams, to VR goggles. Interactive immersive VR streams are addressed separately in section 1.3.3 below.




Figure 7 - Stream resolution distributions combined for current and emerging TV and mobile devices per household (left) and associated mean video rates and coding (right).

Sustained video rates for various stream resolutions will still be suppressed by using lossy compressions and codecs which in turn will evolve over time. The projections shown here are based on the analysis from [10], using the AVC and HEVC rates for 2k HD and 4k UHD video for a (subjective) quality score of 8 out of 10, and adding a transport overhead of 35%. This will apply to most managed TV and VoD services, albeit higher than many OTT streams. To account for quality and (adaptive) rate variations, the values shown in Figure 7 are varied by using a Gamma Distribution with σ/μ of 10%. Figure 8 shows an example for the 4k video sustained rate distribution in 2017. For the forecast, a compression improvement of 10% (MPEG2) up to 50% (H.266) per decade, or 1 to 5%/year is assumed, reflecting the migration to higher compression standards per year, due to device replacements and upgrades. 8k and 16k rate estimations are mere extrapolations of the lower rates, presuming that devices will have access to the processing power needed for future, i.e. H.266 or VP10, compression standards and codecs. The resulting average stream rate per household increases from 7 Mbps in 2017 to 12 Mbps in 2027 for an average resolution growing from 2k to 4.5k.




Figure 8 - 4k video stream rate variation in 2017 around the mean value to account for quality variation and rate adaptations.

1.3.3. Interactive immersive streams

Although the projected stream rates would include 16k video resolutions that can be expected from VR headsets, one aspect that is not covered is interactivity. To enjoy immersive experiences without nausea and motion sickness, VR stream rates will likely need to support 100 frames per second and millisecond response times in rendering new frames corresponding to a movement. Although technologies are yet to be developed for real-time network streaming that allows massive adoption, a scenario with interactive VR is included to assess the throughput sensitivity to new disruptive applications.




Figure 9 - Visual scope, 16:9 equivalent resolution and corresponding latency requirements for immersive VR.

To estimate potential video rates for immersive virtual reality applications, the basic properties of human vision are illustrated in Figure 9. Presuming perfect eye quality, the cones in the central foveal area are taken as a reference [11]. Translating 4.5 million cones per eye to a standard 16:9 ratio screen resolution equivalent ($4/3\sqrt{N_{\text{pixel}}}$) yields a modest 2.8k video stream. For a stereoscopic image, assuming 50% overlap and compression efficiency, this would result in a 4k stream. However, expanding a perfect retinal acuity of 20 to 30 arc seconds per cone to the full field of view (FoV, ca. 145° wide x 100° high) yields an equivalent resolution of 18k per eye. This means that to anticipate any possible eye movement you need to transfer 8 times as much information as a user actually can perceive within a FoV (similar for screen displays). By using eye tracking, the adaptive foveal stream of 4k would suffice, but then response

times far below 10ms are needed, including latency of viewing gear, server and transmission. Since directional head and body movements are slower than eye movement, response times can be more forgiving but still within 10 milliseconds to prevent motion sickness. A full spherical view ($360^{\circ 2}/\pi$ or about one billion hexagonally arranged cones) can instantly accommodate any change of view direction, but then an equivalent stereoscopic resolution of about 60k is needed. Obviously, when the viewer also moves around, physically or virtually, latency requirements will also apply here. When full sphere image streams can be compressed as sufficiently as HEVC projections promise then, for frame rates in the order of 100 fps, data rates around 200 Mbps can be expected. Since compression of spherical images will ultimately depend strongly on scenery, compression algorithms and certainly on device processing capabilities, both 200 and 400 Mbps is investigated. A second important parameter is the expected use probability. Although daily use patterns will probably resemble other video applications, the adoption of new devices is very uncertain. TV headset penetration rates up to 45% are expected [12] in the coming 5 years, but this includes smartphone-based devices. For stand-alone devices with capabilities projected here (18k per eye), scenarios are included for adoption as of 2020, growing to either 7% or 15% in 2027.

1.4. Download traffic

To translate the number of persons per household into the number of active downloads, the concurrent multiple use of devices (e.g. TV, PC, tablet and mobile phone) should be observed as well. For example, a member of the household may be using their PC to check email, browse the web or go on social media and then use their tablet or smart phone to do the same activity within the same busy hour time interval. The traffic generated in the household by instant content downloads can similarly be modeled as a sum of generators:

$$B_D = \sum_{i \leq N_D} R_{d,i} \quad (8)$$

The number of concurrent downloads N_D can be expressed again as:

$$P\{N_D = k\} = \binom{N_{HH}}{k} p_u^k (1 - p_u)^{N_{HH}-k}$$

According to (2), the download probability p_u is determined by the use duration. But unlike streaming content, this is the download or burst duration rather than the application use duration which, for a daily download size D per user, amounts to:

$$t_{use} = \frac{D}{R_d} \quad (9)$$

User initiated downloads do however depend on device use, so the hourly download probability of applications $p_u(t)$ can be derived by distributing the data consumption per device by the device use per hour. From the data described in section 1.2 above, Figure 10 results for various devices. For non-interactive data applications that do not rely on a specific response time, i.e. from autonomous devices and processes such as storage and backup, a constant uniform distribution profile is used.




Figure 10 - Application usage probability p_u per U.S. resident (left) and hourly distribution per device, based on [5,6] (right).

1.4.1. Download Rate

The burst rate R_d in (8) of application downloads is, apart from the response time t_r given by the burst data size d_b :

$$R_d = \frac{d_b}{t_r} \leq R_{max} \quad (10)$$

For unconstrained use one could argue this rate should be as high as possible, i.e. the maximum rate supported by both home network and access network R_{max} . But more realistically, also at the server side restrictions will apply, so that two principles should be observed. First, the amount of data that users can consume instantly is limited by perception restrictions. Most often, visual information in the form of shapes, images and photographs will be the most demanding (non-streamed) content, outnumbering plain text data by a factor of 1,000 to 10,000. Secondly, interactive view and browser applications apply progressive download techniques to present content before full download has completed. This applies to images, web pages and documents but also to interactive table views to online data sources.

In this study, the burst size is assumed to amount to 2.5 MB for an average web page [13], increasing 6% per year, parallel to the average stream rate (see Figure 7) while the response time is assumed 4 seconds (target web page download time [14]) decreasing to 0.5 seconds in 2027. Other application burst sizes and durations are merely estimates for typical applications and content types. But similar to stream rates, burst sizes will vary strongly and are therefore modeled using a Gamma distribution with a σ/μ of 50%.

With the burst rate, also the download probability can be established from (2), (9) and (10):

$$p_u = \frac{N_{users}}{N_{pop}} \frac{D}{\Delta t R_d} = \frac{N_{users}}{N_{pop}} \frac{D}{\Delta t} \frac{t_r}{d_b} \quad (11)$$

With $D/\Delta t$ the daily download size per user. As indicated, the burst probability is the product of user probability N_{users}/N_{pop} , the number of burst per day D/d_b and the time fraction of a burst $t_r/\Delta t$.

2. Calculations

2.1. Household traffic model

Although the model described allows for direct derivation of mean throughput values, the aim here is to assess the variation resulting from all assumed distributions. Especially for smaller Access Distribution Areas (DAs), e.g. for FTTdp and FWA, the network is commonly designed for a certain percentile well above the expected average, e.g. the 95th percentile. The cascade of probability distributions for users, usage and video rates prevents expressing the n-th percentile of the throughput directly. Although it is possible to scan through all permutations numerically to collect percentile values, here we revert to Monte Carlo simulation. For this purpose, thousands of households are calculated ($N = 10,000$ for each year) by using random samples for the various distributions. To suppress numeric noise when calculating low-probability conditions of high-percentiles, Latin hypercube (LHC) pseudo-random sampling is used to cover most of the variation ranges of population, device and stream distributions. A summary of the main model parameters is listed in Table 1 and Table 2.

Table 1 - Media stream model parameters

Parameter	Value
Busy Hour Usage probability p_u	$\mu = 45\%$, $\sigma = 23\%$
Secondary Streams per Active User p_2	$10\% \pm 10\%$, CAGR +2%/year
Shared Stream fraction p_s	$50\% \pm 50\%$, CAGR -2%/year
Video rates per stream (@60 fps, 35% overhead) R_i :	Gamma distributed, $\sigma/\mu = 10\%$
• $\leq 1k$	$\mu = 2$ Mbps, CAGR -1%/year
• 2k	$\mu = 8$ Mbps, CAGR -2%/year
• 4k	$\mu = 16$ Mbps, CAGR -3%/year
• 8k (as of 2020)	$\mu = 27$ Mbps, CAGR -4%/year
• 16k (as of 2025)	$\mu = 43$ Mbps, CAGR -5%/year
• Interactive VR (if included, as of 2025)	$\mu = 200/400$ Mbps
Audio BH Usage probability p_u (Radio & VoIP)	$\mu = 8.8\%$, $\sigma = 4.4\%$
Audio rate per stream (Radio & VoIP)	0.128 Mbps

Table 2 - Data download model parameters

Application	Users per resident	Data Use D MB/day	Burst Size d_b [MB]	Burst Time t_r [s]
Browsing	76%, CAGR +2%/y	89.1 +8%/y	2.5 +6%/y	4 -21%/y
Content	65%, CAGR +3%/y	46.5 +15%/y	5 +6%/y	4 -21%/y
Productivity	22%, CAGR +5%/y	97.2 +14%/y	10 +6%/y	4 -21%/y
Storage/backup	24%, CAGR +5%/y	84.7 +8%/y	50 +6%/y	300 -21%/y
Gaming	50%, CAGR +1%/y	10.6 +5%/y	0.01 +6%/y	0.1 -21%/y
File sharing	3%, CAGR -3%/y	2,700 +2%/y	200 +6%/y	600 -21%/y
E-mail	99%, CAGR -2%/y	3.1 +3%/y	5 +6%/y	10 -21%/y

2.2. Distribution Area traffic model

For Distribution Areas, the aggregate throughput of a collection of connected households is calculated by resampling the data obtained per household. This is done to limit computation time without having to rely on less accurate fitting to, e.g., Gamma distributions.

2.2.1. Multicasting

The total aggregate throughput at the line side of the Access Node is calculated as the sum of all connected home throughputs, where it is assumed that no multicasting mechanisms are available at the link layer. For throughput at the uplink of an Access Node, but also access technologies on shared media that do support link multicasting, e.g., HFC, PON and FWA, throughput calculations of linear TV streams must account for savings from broadcast streams. The same would also apply for IP multicast on Access Node uplinks, depending on group numbers and stream popularity distributions. Increased use of Unicast streams, i.e. for non-linear and time-shifted viewing, will reduce multicast savings. When Access Nodes provide caching however, uplink transport savings can be obtained as indicated by the Hit Ratio. Figure 11 shows content popularity, commonly modeled as Zipf-Mandelbrot distributions and the resulting Multicast Gain that can be obtained. For typical multicast and VoD shape parameters, 80% of multicast traffic can be saved for 1000 views assuming 250 channels. For VoD and time-shifting, usually more content is available so that, for 2000 titles, 30% of unicast traffic could be saved for 1000 viewers. If unmanaged OTT streams are cached at all, this mostly takes place deeper in the network, without impact on the load of Access Nodes or their uplinks. Although Figure 11 suggests that significant load reductions can be achieved especially for multicast video in larger Distribution Areas, this effect has not been included in the remainder of this analysis.




Figure 11 - Managed content popularity distributions (left) and corresponding Multicast Gain and Cache Hit Ratio (right).

2.2.2. Usage dependence

Although resampling provides an efficient way of calculating aggregate throughputs based on the statistics per household, it disregards one aspect that may play an important role. It implicitly assumes independence of the usage probability among different households which, especially for TV video streams may not apply in case of popular sports or newsworthy events. To see the impact of such events, an alternative scenario with a view probability of 90% instead of 45% is calculated which corresponds to the total population of 10 years and older watching TV at 9pm.

3. Results

3.1. Sustained Throughput per U.S. household

3.1.1. Mean Sustained Throughput

Using the distributions and assumptions discussed in sections 1.3, 1.4 and 2, the Sustained Throughput is calculated for each year. Figure 12 shows contributions of streams and downloads to the mean throughput values during Busy Hour between 2017 and 2027, which are summarized in Table 3. The slight curving at 2020 and 2025 correspond to the uptake of 8k and 16k video respectively. The average throughput in 2017 is about half the BH rate and would amount to 4 Mbps. As a reality check, this corresponds to 1.3 TByte of video data per month, increasing to 2.4 TByte/month in 2027. Again, this would represent a situation where linear TV services are taken either from the Access Network provider or from another IPTV or OTT TV subscription.

Table 3 - Mean Sustained Throughput per U.S. Household

Mean Sustained Throughput per HH	2017	2027	CAGR
BH Media Streams	7.9 Mbps	14.2 Mbps	6.7%
BH Data Download	0.2 Mbps	0.7 Mbps	12.9%
BH Total Throughput	8.1 Mbps	14.8 Mbps	6.9%
Daily Average Throughput	4.0 Mbps	7.4 Mbps	7.0%
Data per day	43.2 GB	79.5 GB	7.0%
Data per month	1.3 TB	2.4 TB	7.0%

These numbers indicate almost a doubling of growth from 8 to 15 Mbps per household. This results from the increased number of streams per household (Figure 6) and especially the mean rate growth (Figure 7). The impact of streaming video is clear, as it contributes to 95% of the total volume per household.




Figure 12 - Mean Sustained Throughput per U.S. average household contributions from Streamed Media and Data downloads during busy hour.

3.1.2. Throughput Distributions

Using the distributions and assumptions discussed in sections 1.3, 1.4 and 2, the Sustained Throughput is calculated for each year. Figure 13 shows the results for 2017 and 2027, indicating a growth of the mean Throughput value from 8.8 to 15.6 Mbps per household. This results from the increased number of streams per household (Figure 6) and especially the mean stream rate growth (Figure 7). The Probability Density Functions (PDF) show the long tail shape as they propagate from the Exponential-like shapes of both the stream number and size distributions.

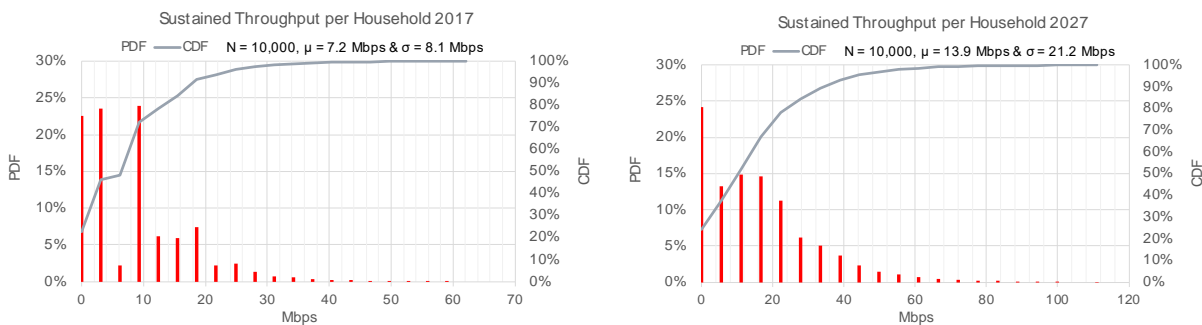


Figure 13 - Probability Density Function (PDF) and Cumulative Density Function (CDF) for the Maximum Sustained Throughput per U.S. household in 2017 (left) and 2027 (right).

3.1.3. Percentiles vs Mean value

A more significant value for network design is the upper range of the expected throughput values. Figure 14 and the summary in Table 4 show the various percentiles for the Throughput values between 2017 and 2027, indicating that to design for e.g. 95% of all households in 2027, a throughput of up to 43 Mbps per household would need to be considered rather than the mean value of 14.8 Mbps. For 99% this value reached 71 Mbps. For an Access Node, the difference between mean and high variations depends on how many households are actually connected per Access Node, and thus the Distribution Area size as analyzed next.




Figure 14 - Percentile range for the Sustained Throughput per U.S. household between 2017 and 2027.

Table 4 - Sustained Throughput percentiles per U.S. Household

Sustained Throughput per HH	2017	2027	CAGR
Mean	8.1 Mbps	14.8 Mbps	6.9%
50% (Median)	2.5 Mbps	10.3 Mbps	17.3%
95%	21.8 Mbps	46.5 Mbps	10.0%
99%	35.1 Mbps	71.4 Mbps	11.6%

It may be noted that networks in areas with higher revenue potentials will have to be designed to cover 95th or even 99th percentile throughputs as residents will expect service quality comparable to FTTH. Additionally, sufficient headroom is to be kept at distribution points as margin on top the sum of sustained throughput of all users, in order to cater for bursty applications or “speed tests” that users may perform from time to time to check their experience versus advertised headline speeds.

3.1.4. Stream use Sensitivities

To assess the model sensitivities to the major parameters, results from the simulations are shown in the scatter plots of Figure 15 and Figure 16. As expected, the Busy Hour usage (Figure 15, left) linearly drives the number of streams and thus the throughput per household. A more detailed analysis of high traffic loads during special occasions is analyzed in section 3.2.1 below.

The use of secondary simultaneous streams (Figure 15, right) has a much lower impact as its range is a factor 4 smaller than that of p_u and it impacts the number of active streams only relatively.

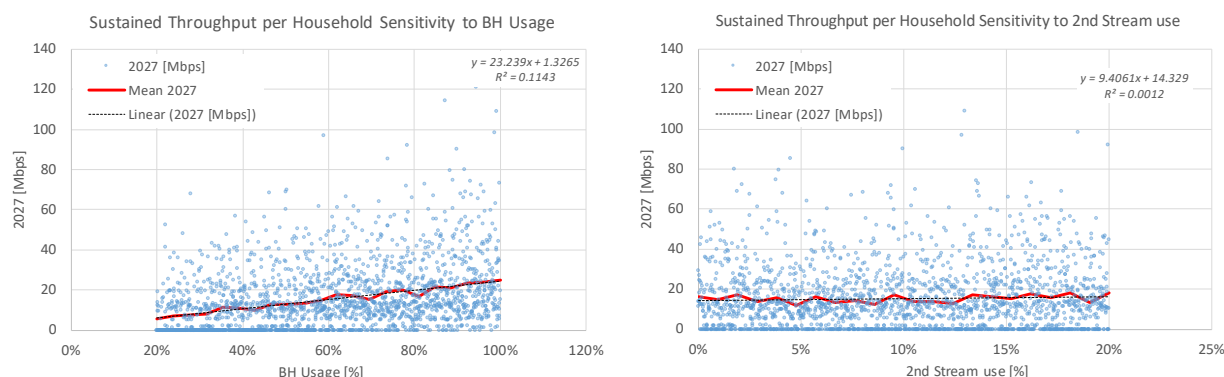


Figure 15 - Sensitivities of the Sustained Throughput per U.S. household in 2027 to variations in Busy Hour Usage p_u (left) and secondary stream use p_2 (right).

As indicated by Figure 16, the effect of Stream Sharing is more visible, but smaller than (BH) Usage: it only dampens active streams when more than one active user is present. Since an average household has only 1.3 active viewers during busy hour (45% of 2.45 persons), sharing only applies to 0.3 views per household or 25%.




Figure 16 - Sensitivity of the Sustained Throughput per U.S. household in 2027 to variations in Stream Sharing p_s .

3.1.5. Disruptive applications

To determine the impact of new disruptive applications, a scenario has been calculated where 200 Mbps streams associated with immersive, interactive VR were added to the portfolio increasing from 0.1% of the video streams in 2020 to 7% in 2027. The results are shown in Figure 17 and summarized in Table 5. The Mean Sustained Throughput doubles to 15.6 Mbps, the 99th percentile Maximum Sustained Throughput jumps from 70 to 241 Mbps. This suggests that even for a low uptake of applications such as Immersive VR, the Maximum Sustained Throughput will be impacted. Higher take rates of these services can dramatically increase throughput requirements.




Figure 17 - Sustained Throughput per U.S. household during busy hour including interactive VR Streams for take rate up to 7% in 2027 and 200 Mbps per VR stream.

Table 5 - Impact of Immersive VR streams on Sustained Throughput per U.S. Household (200 Mbps VR rate)

2027 Sustained Throughput per HH	without VR	with 7% VR	15% VR
Mean Average Throughput	7.4 Mbps	15.6 Mbps	23.7 Mbps
Mean BH Throughput	14.8 Mbps	31.9 Mbps	48.3 Mbps
95%	46.1 Mbps	207 Mbps	223 Mbps
99%	70.4 Mbps	242 Mbps	419 Mbps

Figure 18 shows the situation when take rates are doubled to 15% in 2027. Although a higher take rate increases the throughput probabilities, it is not until 2027 when the probability of two simultaneous VR streams pushes the 99th percentile Maximum Sustained Throughput beyond 400 Mbps.




Figure 18 - Sustained Throughput per U.S. household during busy hour including interactive VR Streams for a 15% take rate at 200 Mbps per VR stream.

That these maximum rates are directly driven by the VR rate is shown in Figure 19, where a 400 Mbps VR rate is assumed, which basically doubles the Maximum Sustained Throughput. As such, close monitoring of how VR gear capabilities will translate into stream rates is essential in the coming years.




Figure 19 - Sustained Throughput per U.S. household during busy hour including interactive VR Streams for a 7% take rate at 400 Mbps per VR stream.

3.2. Sustained Throughput for Distribution Areas

Applying the per household throughput distribution to simulation of a collection of connected households gives the results shown in Figure 20. As the size of DAs increase, averaging between small and large households and various devices results in less deviation and, as the central limit theorem (CLT) predicts, a throughput distribution that resembles a Normal distribution.

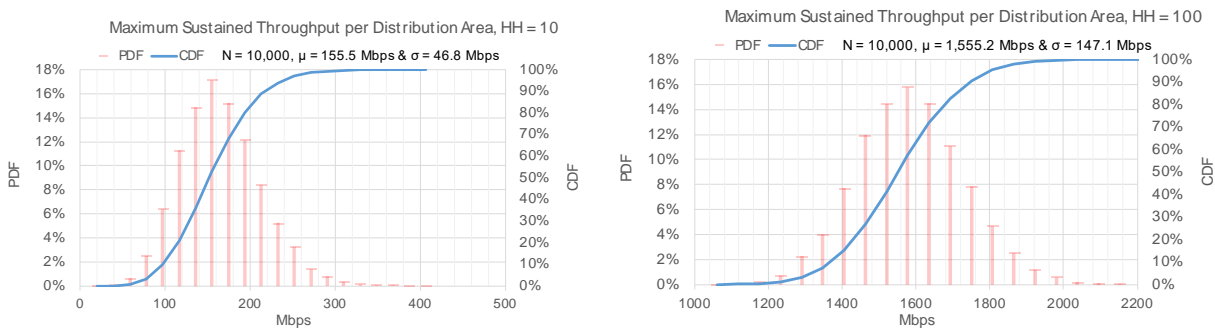


Figure 20 - 2027 Throughput for a Distribution Area with 10 (left) and 100 households (right).

This presumes some independence of the connected households which in practice, for specific geographic areas characterized by homogeneous income patterns, may actually-not be entirely satisfied and should be accounted for by selecting the proper household size distribution and possibly adjusting the stream distribution. For this study, relying on U.S. averages, the impact of DA size on the total Sustained Throughput is depicted in Figure 21. It shows that for smaller areas, i.e. < 100 households, the relative variance σ / μ is too high to ignore.




Figure 21 - 2027 Throughput for varying Distribution Area size per percentile

For $N=10$, the 95th percentile value is 80% higher than the mean value where for $N=100$, this is only 23%. For $N \gg 100$, the variance drops further and designs can assume mean throughput values similar to core transport networks. This is illustrated in Figure 22, where the DA capacity per connected household drops from 70 Mbps, the Maximum Sustained Throughput, for 1 household to 15 Mbps, almost the Mean Sustained Throughput per household.




Figure 22 - 2027 Capacity per household required to support the 99th percentile of the total Sustained Throughput in Distribution Areas

3.2.1. High-load conditions

To see the impact of massive views during special sports or other newsworthy events, a scenario is calculated with a double view probability of 90%, i.e. 9 out of 10 residents watching a video stream concurrently. As indicated in Table 6, a high view probability increases the mean throughput by 75% both for the daily average and Busy Hour. The Maximum Sustainable Throughput increases only by 27%, because these include households with more viewers, and the sharing of streams exceeding 1 is more likely in big households than in average sized households. Since it is unlikely that 90% of the population is watching video concurrently even during rare events, an additional 30% can be regarded as a solid margin to design for these circumstances.

Table 6 - 2027 Sustained Throughput per U.S. Household during high load events

2027 Sustained Throughput per HH	45% Viewers	90% Viewers	increase
Mean Average Throughput	7.3 Mbps	12.7 Mbps	74%
Mean BH Throughput	14.8 Mbps	25.9 Mbps	75%
95%	46.5 Mbps	60.5 Mbps	30%
99%	71.4 Mbps	90.6 Mbps	27%




Figure 23 - 2027 BH Throughput variation per Distribution Area size during high load events.

Conclusion

Video will continue to drive residential access network loads in the coming 10 years. While peak throughputs per household may cross 1 Gbps, the mean expected sustained capacity per household by 2027 will not likely exceed 7.5 Mbps on average and 14 Mbps during Busy Hour, but that will not support the Maximum Sustained Throughput of households.

The 99th percentile Maximum Sustained Throughput per household will be 70 Mbps in 2027 during Busy Hour but during high-load conditions, when most of the population is watching video concurrently, this will increase up to around 90 Mbps. The margin to consider for network design is therefore 30% higher than normal viewing conditions.

When disruptive interactive applications such as Immersive VR will represent only 7% of video streams, the Maximum Sustained Throughput may increase to 240 Mbps or, for 15% take rates, to 420 Mbps in 2027. Most of this rate will depend on new interactive and immersive video services and eventual device capabilities. Higher adoption of high end applications can dramatically increase the service requirements.

Table 7 - 2027 Sustained Throughput per U.S. Household for engineering considerations

Maximum Sustained Throughput (99th percentile Busy Hour)	2017	2027
U.S. household	35 Mbps	71 Mbps
U.S. household - High Load	45 Mbps	91 Mbps
U.S. household – 7% Disruptive VR	-	242 Mbps
U.S. household – 15% Disruptive VR	-	419 Mbps

The throughput variance per Access Node rapidly drops with higher number of homes connected, allowing taking advantage of statistical multiplexing. The standard deviation of the Throughput for 100 connections is only a third compared to 10 connections, while 1000-household DAs can design for only a few percent above mean average Busy Hour values. However, it is important to engineer Access Nodes with sufficient headroom to accommodate at least one “Speed Test” on top of sustained throughput, in order to meet customer expectations as well as any regulatory requirements in the foreseeable future. Many regulatory authorities in Europe have discussions on “truth in advertisement” in the public domain.

Scale is therefore key as fiber pushes further towards the home and Distribution Areas shrink with growing speed. For Access Networks, either HFC, FWA, FTTH or xDSL, proper traffic forecasting is the starting step for a proper design. With the appropriate demographics indicators and (video) service projections, the presented model is well suited to provide key insights into the Maximum Sustained Throughput and other traffic characteristics. As the projections for future video usage, application bandwidth requirements as well compression techniques emerge, the current projections will evolve over time.

Abbreviations

Acronym	Meaning
5G	5 th Generation (Wireless Systems)
AR	Augmented Reality
AVC	Advanced Video Coding
BH	Busy-Hour
CAGR	Compound Annual Growth Rate
CDF	Cumulative Distribution Function
CPE	Customer Premises Equipment
DA	(Access) Distribution Area
FoV	Field of View
FTTdp	Fiber to the Distribution Point
FWA	Fixed Wireless Access
HD	High Definition
HEVC	High Efficiency Video Coding
LHC	Latin hypercube
Mbps	Megabits per second
μ	Mean Average Value
OTT	Over-The-Top (unmanaged services)
PDF	Probability Density Function (also abused as Probability Mass Function)
QoE	Quality of Experience
σ	Standard Deviation
UHD	Ultra-High Definition
VoD	Video on Demand
VR	Virtual Reality

Bibliography & References

¹ E. Harstead and R. Sharpe: Forecasting of Access Network Bandwidth Demands for Aggregated Subscribers Using Monte Carlo Methods, IEEE Communications Magazine • March 2015, <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7060505&tp=?ALU=LU1005515>

² BLC Mobility Report, <https://pages.nokia.com/1503.bell-labs-mobility-report.html>

³ U.S. Census Bureau, <https://www.census.gov/data/tables/time-series/demo/families/households.html>

⁴ U.S. Census 2015-2060 National Population Projections, <https://www.census.gov/data/tables/2014/demo/popproj/2014-summary-tables.html>

⁵ Nielsen Q4 2016 Comparable Metrics report, <http://www.nielsen.com/us/en/insights/reports/2017/the-comparable-metrics-report-q4-2016.html>

⁶ Nielsen 2015 Audience report, <http://www.nielsen.com/us/en/insights/reports/2015/the-total-audience-report-q2-2015.html>

⁷ <http://www.nielsen.com/us/en/insights/news/2010/u-s-homes-add-even-more-tv-sets-in-2010.html>

⁸ Massive VR or AR adoption will likely not occur until power consumption and battery storage allows for affordable, comfortable goggle form factors.

⁹ Consumer Technology Association (CTA), 19th Annual Consumer Technology Ownership and Market Potential Study, via <https://www.cta.tech/News/Press-Releases/2017/May/A-Smartphone-Surprise-U-S-Ownership-Hits-Record.aspx>

¹⁰ Tan et al.: Video Quality Evaluation Methodology and Verification Testing of HEVC Compression Performance, IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, VOL. 26, NO. 1, JANUARY 2016, pp. 76-90,
<http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7254155>

¹¹ Purves D, Augustine GJ, Fitzpatrick D, et al., editors., Anatomical Distribution of Rods and Cones, Neuroscience. 2nd edition., National Center for Biotechnology Information, U.S. National Library of Medicine, <https://www.ncbi.nlm.nih.gov/books/NBK10848/>

¹² How Virtual Reality Hype Exceeds Demand—For Now, Fortune,
<http://fortune.com/2016/10/13/virtual-reality-headsets/>

¹³ The web is doom, <https://mobiforge.com/research-analysis/the-web-is-doom>

¹⁴ How Fast Should A Website Load in 2017?, <https://www.hobo-web.co.uk/your-website-design-should-load-in-4-seconds>

MSO's Health Over Cable

The Ways We Can Add Value

A Technical Paper prepared for SCTE•ISBE by

Mark Bugajski
SVP Advanced Technology
ARRIS
3871 Lakefield Dr.
Suwanee, GA 30024
mark.bugajski@arris.com

Paul Moroney
SVP Advanced Technology and GM Security Solutions
ARRIS
6450 Sequence Drive
San Diego, CA 92121
paul.moroney@arris.com

Introduction

Over the next 35 years, there will be a massive increase in the aging population that will require health care, despite scarcer human resources to provide that care. The Internet of Things (IoT) aims to connect medical devices to service providers and create The Internet of Health Things (IoHT). By transforming raw data into actionable information and communicating that information to everyday objects, machines, and people, the IoHT is becoming a vital tool to the healthcare industry. Although the IoHT promises to significantly lower the costs of healthcare to the aging populations of developed countries, the current Internet of Health Things are Over-the-Top (OTT) based, highly fragmented, and challenging for the average patient or consumer to use. Cable operators are in an advantageous position to partner with care-giving providers to create managed networks, use set-top boxes (STBs) as service portals, use STBs as medical devices with a built-in Bluetooth Low Energy (BLE) interface and utilize condition-specific Video on Demand (VoD). Most importantly, cable operators are able to provide a secure cable network, ensuring that a patient maintains their privacy.

Bringing Health Monitoring Services Home via Cable Networks

As is, modern nations will not be able to sustain the current level of care needed for baby-boomers and generations beyond them. In 2015, just 2.7 % of the world's population was over the age of 70, roughly 190 million people (as seen in Figure 1). By 2050, the percent of people over the age of 70 is expected to rise to 6% of the world's population (as seen in Figure 1).

Data source for Figure 1 and Figure 2: (www.populationpyramid.net)



Figure 1 - Global Aging Trends

In the United States alone, 4.5% of the population, or 14 million people, were above the age of 70 in 2015 (Figure 2). This number is expected to rise to 32 million or 8.2% of the population by 2050 (Figure 2).




Figure 2 - USA Aging Trends

To address the care to the aging population problem, there has been a great deal of technological innovation aimed at producing IoHT devices that will offset the associated costs.

Currently, there is much work on miniaturization, portability, and significant cost reduction for IoHT devices that can do various things like monitor human vital signs and track the number of steps taken during the day using OTT mobiles and various wrist bands. Many health and wellness monitoring devices are already widely available in retail as seen in Figure 3.



Figure 3 - Connected Medical Devices and Associated Challenges

- BLE Connected Thermometers use wireless Bluetooth to send data from temperature sensors to smart phones and tablets, making the data easier to understand and send to relevant healthcare providers

- ZigBee Alarm Buttons allow users who have fallen, are having an acute medical problem, or are under threat of fire or burglary to activate wireless alarms and panic buttons to alert monitoring personnel
- ZigBee Motion Detectors trigger lights to illuminate darkened hallways or doorways and to turn on children's nightlights. They also send text and e-mail alerts when motion is detected in empty houses, ensuring safety and security
- BLE Connected Weight Scales allow users to send their weight and BMI data into relevant apps that track health and dieting, while BLE Pulse and Oxygen Level Meters measure and wirelessly store blood oxygen levels, heart rates, and perfusion indexes
- BLE Connected Breathalyzers send blood alcohol levels to smartphones, and BLE Glucose Meters monitor glucose activity in real time, sending high and low alerts and alarms, and sharing data with selected followers
- Most popular, perhaps, are BT Fitness trackers which track, calories burned, steps taken, stairs climbed, and more

Though these devices are available and widely utilized, there are certain challenges in the existing technology as illustrated on Fig. 4, including the app diversity, connection to the caregiver, manageable and easy to understand regimes, compliance enforcement, data analysis, and follow-up with patients.

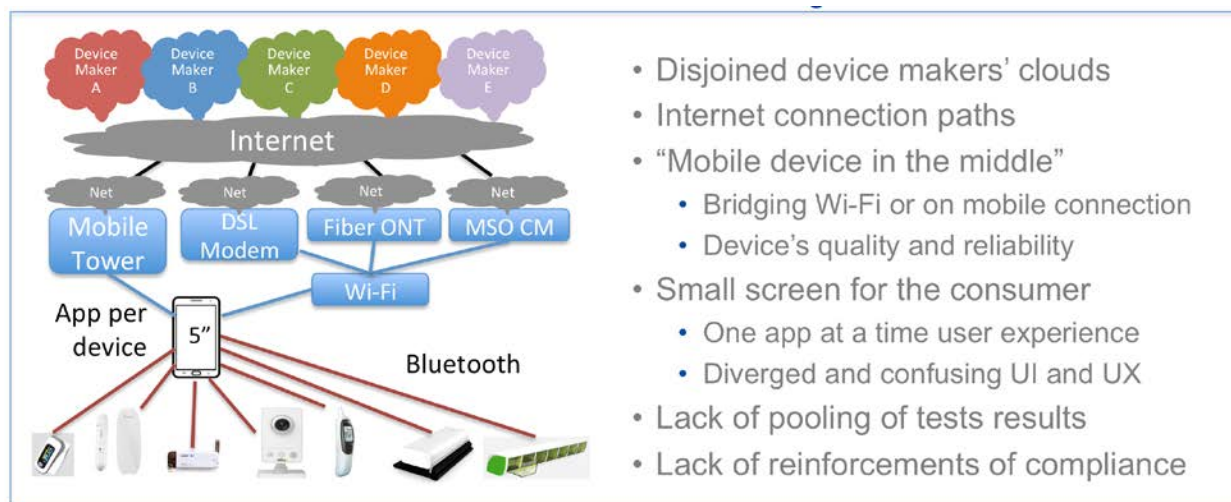


Figure 4 - Health over Internet Challenges

Furthermore, there are additional challenges in regards to the OTT connections. These include disjoined device makers' clouds, different Internet connection paths to consumers, reliance on home Wi-Fi or on mobile connections, and dependence on mobile device's quality and reliability. A too-small screen for the consumer, one app at a time user experience, diverged and confusing UI and UX, lack of pooling of tests and their results, and lack of reinforcements of regime compliance provide addition challenges to connections.

The current Health over Internet (OTT) method of communication with monitoring devices includes connection to the home gateway plus Wi-Fi connection to mobile devices. For example, smart phones and

tablets use Bluetooth to link to the point-of-use device like a thermometer, blood pressure cuff, pulse and oxygen level meters, scales, or other apparatus. Currently, the Wi-Fi and Bluetooth connections are not monitored by the service providers for quality and reliability. That quality and reliability of the measurement procedure depends upon several factors many of which are often outside the monitored person's or the patient's control. The required steps are shown in Figure 5.



Figure 5 - Health over Internet Required Measurement Steps

The lean-forward consumer experience requires the patient or care receiver to go through required steps. Initially, the monitored person must remember to take their measurements. For the purpose of this technical paper, “measurements” can be any information required for the health care application including: temperature, blood pressure, blood sugar reading, whether a scheduled medicine was taken, etc. The mobile device must not only be turned on, it also needs to be reliably connected to a Wi-Fi network. The battery inside the mobile device must have enough of a charge or have a connection to a power source. The appropriate app must then be found and launched by the patient. Usually, several on-screen navigational steps need to be undertaken to get to the monitoring device control. The monitoring device needs to be connected to the mobile device's Bluetooth. Only then can the test be initiated. After the test is complete, the results can be uploaded to the service provider's cloud.

Given all these steps, the current system of operation creates many points for error and failure as illustrated on Figure 6. Some of these steps are often difficult for the monitored person or patient to remember and correctly sequence as they interacts with the system. More importantly, the patient may not even remember to take the measurements in the first place. If the mobile device is misplaced or has a poor or intermittent Wi-Fi connectivity, then the measurement may fail to record or cut off early. The battery may drain, causing difficulty, or the device may go to sleep too early and interfere with the measurement taking. The Bluetooth radio may be disabled to save battery. The device may also have intermittent Bluetooth connection or drop the Bluetooth link altogether.

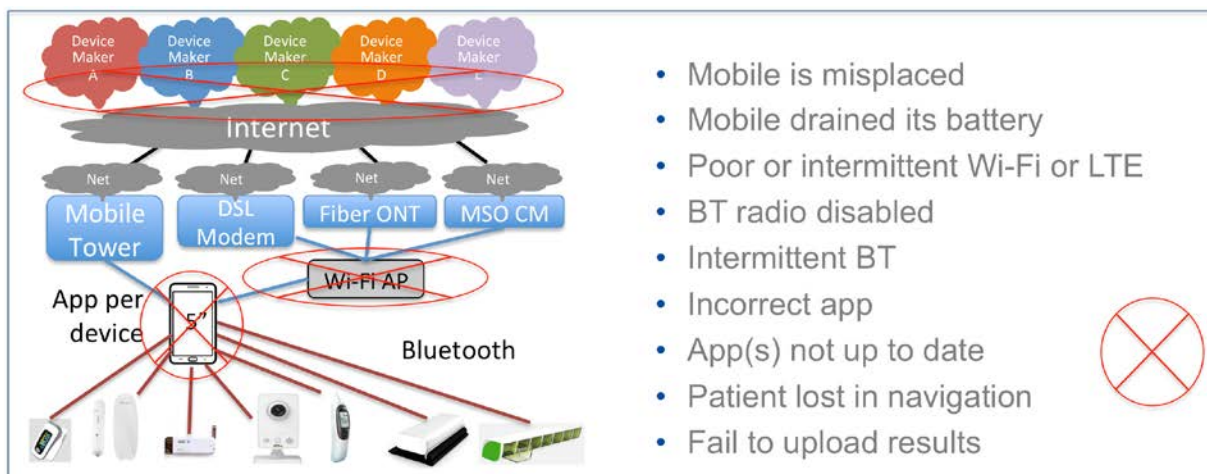


Figure 6 - Health over Internet Associated Risks

The incorrect app may start or the patient may start the incorrect app, and the patient may be very confused with the navigational steps. Finally, due to a variety of reasons, the device may fail to upload the test results.

In contrast, when a smart MSO set-top device with a built-in BLE interface is used to connect to the monitoring devices over a private MSO network, a much simpler and intuitively more reliable architecture shown on Figure 7 is being created. We can call it a **Health over Cable (HoC) System**. It consists of a cable operator owned set-top box (STB) to connect to all home medical devices over a totally private Intranet. It is assumed that the MSO partners with a healthcare or monitoring care services provider will connect in the MSO cloud.

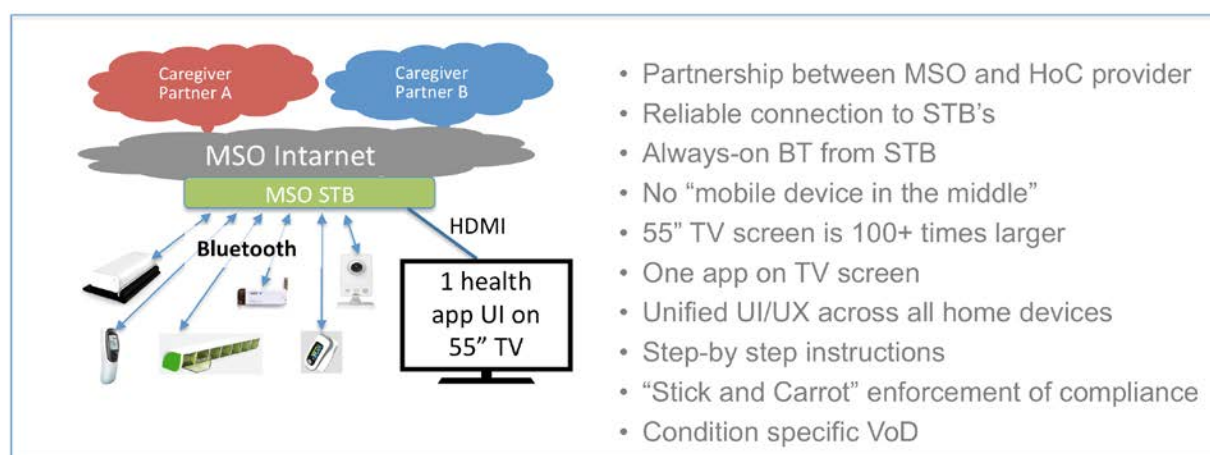


Figure 7 - Health over Cable Advantages

Cable operators are in a unique position to provide improvements in these areas. In order to offload hospital beds and provide real care to non-critical patients, cable operators could make the developments and take steps needed to provide beneficial solutions to their subscribers. The monitoring and care-giving portal in these healthcare devices must be extremely reliable and easy to interact with. Even people with a

limited understanding of technology, such as the elderly or mentally impaired, need to be able to understand and operate the simple interactive interface.

The same BLE fitted STB drives a large TV screen to become a portal to the patient. The patient interacts with the system using voice enabled remote or several buttons on such. The connection to the STB is wired and as such is very reliable. The BLE radio in the STB is “always-on” powered by the same power supply as the STB. There no additional wireless or mobile Wi-Fi or LTE connection required, because the device in the middle as it has been eliminated from the HoC architecture.

A large TV screen, 100 times larger than mobile device’s screen is used to interact with the patient. There is no need to touch any screens. One application replaces individual, per-device mobile apps. It features very simple User Interface and common experience across all medical devices. Step-by-step instructions are available with a single click of a remote’s button. The MSO can participate in the compliance enforcing features of the HoC system by motivating by manipulating video playout and awarding with VoD assets. On the education side, the MSO can make condition and life style specific content playable from the same app.

This architecture creates a comfortable lean-back experience for the patient. Instead of relying on an own individual’s memory, on-screen notifications remind and/or alert the patient to take the measurements. The HoC user experience is illustrated on Figure 8 with a number of simple steps needed to interact with the system.




Figure 8 - Health over Cable User Experience

The on-screen notifications are being forced by an always-connected cloud. The patient does not have to worry or be confused about the order of steps because on-screen instructions direct and instruct the patient to correctly start and use the monitoring device. Prior to taking the measurements, the Bluetooth connection is automatically enabled. This allows for a quick connection to the device. Visual and audio prompts then help the patient take the measurements correctly.

The possible problems with HoC user experience are illustrated on Figure 9. The removal of a mobile phone from the connection to the patient makes this solution inherently more robust and easier to use. The same set of devices can be operated from a single TV screen using a remote or own voice to navigate.



Figure 9 - Health over Cable and itUser Experience

Another HoC user experience (UX) option stops TV programming altogether and notifies the patient to take their measurements. In order not to miss any of the programming, the TV program is then recorded to a local or cloud-based DVR (as shown on the screenshot of the TV screen on Figure 10). After the alert, the patient is taken to the device control screen and STB will wait for their input (voice control or a click on the remote) to start the measurement.

- Care recipients and their family can be alerted by the on-screen and audio notifications from the HoC providers
- TV stops playing the content if measurements are not taken
- The TV screen will assist with measurements

Figure 10 - Health over Cable User Experience

The measurement will be shown and immediately uploaded to the cloud. Then, if needed, TV programming will resume and/or a reward can be awarded to the patient for complying with the program. The reward could be in a form of free VoD movie or a temporary access to premium channels.

The TV screen, which is powered by the always-on STB and voice-enabled remote, are in exactly the right place to deliver the condition-specific VoD content, educate the patient, connect with their nurse and doctor, and bridge the healthcare community.

There are a significant number of advantages to using the patient's TV screen as a portal that brings connected health and wellness to the care-receiving household:

- A TV with an average screen size of 55 inches is more than 100 times larger than a mobile device's screen, which has an average screen size of 5.5 inches
- Due to the TV's larger size screen, there is real estate for much larger user interface, which will help enable the patient or care receiver to clearly see the information on a big screen as illustrated on Figure 11
- The patient's or care receiver's family members or household members can be alerted by the on-screen and audio notifications. Even from a distance, the large screen allows family members and visiting caregivers to view the information and be attuned to the care being administered as illustrated on Figure 12

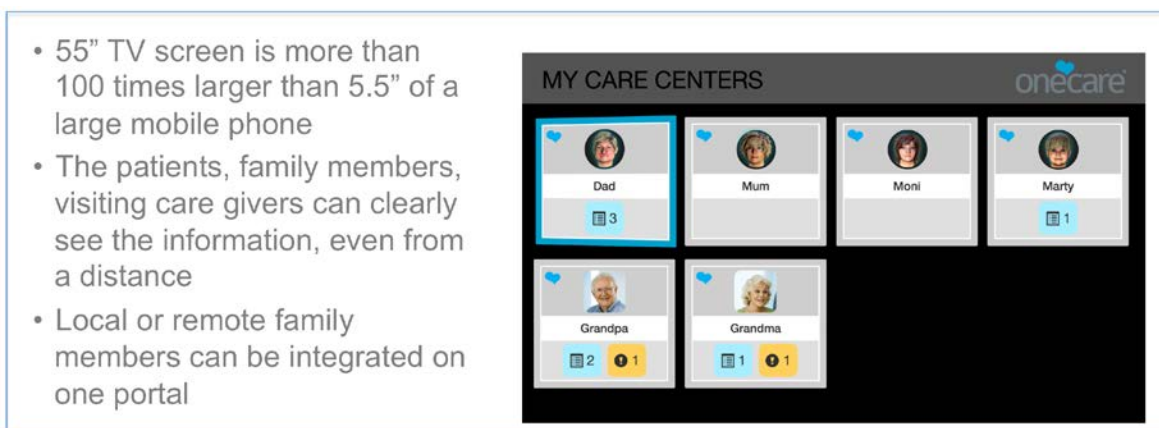


Figure 11 - Health over Cable User Family Portal

These notifications from the service providers help family or household members assist the monitored person in measurement taking activities. Furthermore, with the care recipient's consent, the care-giving service provider can deliver condition-specific on-demand video content to the TV screen.

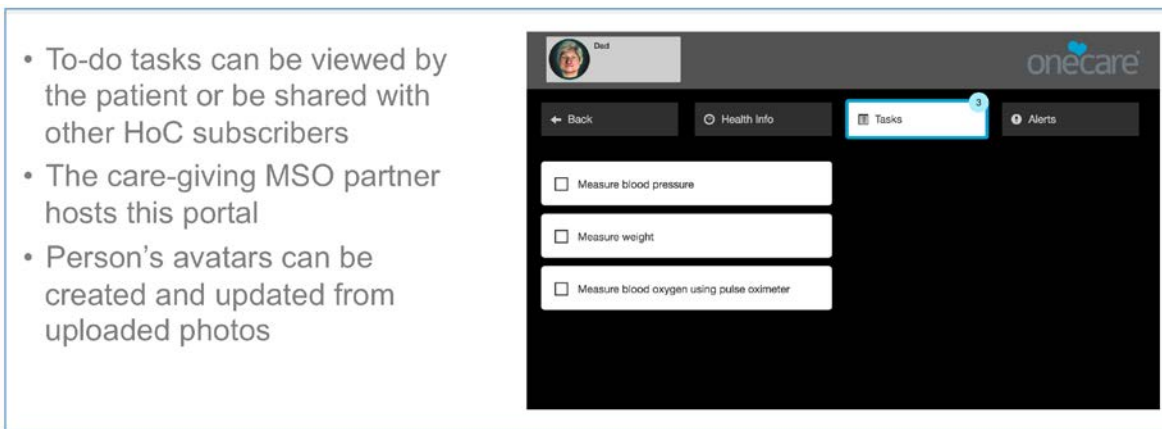


Figure 12 - Health over Cable Individual User Portal

The condition-specific content can be of a linear, lean-back nature. It can also be interactive to ensure that the patient pays attention to it and help them to remember the do's and don'ts associated with the condition they have and the measurement taking activities. The condition-specific content could also be expanded to recommend diets, exercises, and changes to one's lifestyle and behavior. Individual monitoring tasks can be viewed by the patient only or can be shared with other HoC service providers.

Also, the service provide could set up a content exchange portal for patients with similar conditions. Patients with the same or similar conditions can browse through the portal, select it for viewing, and interact with it in many ways as illustrated on Figure 13. Other patients can share links to any relevant content that may be of interest to the patient. The care-giving MSO partner hosts the portal, and the patient's avatar can be created and updated from uploaded photos.

On-Screen HoC VoD Experience

- The care-giving service provider can deliver condition-specific, on-demand video content to the TV screen
- The content can be of a linear, lean-back nature but it can also be interactive to ensure that the patient pays attention to it .
- It can be expanded into recommending diets, exercises and changes to one's lifestyles and behaviours
- The service provider can setup a content exchange portal for the patients with similar conditions

Cancer Care
 Share your story about how using health information technology helped you meet a goal related to cancer care.
 \$8,350 IN PRIZES
[LEARN MORE](#) [Health.gov](#)

Family Caregivers
 Create a video that shares how you use information technology to help manage healthcare for a loved one.
 \$8,350 IN PRIZES
[LEARN MORE](#) [Health.gov](#)

Figure 13 - Health over Cable VoD Experience

Cable service providers and application developers will find that there are many significant advantages to using the STB as a service portal. STBs are always connected. Connection to the Cloud is under total control of the service providers. This kind of connection to the Cloud is more reliable than a connection to any other device, and it is monitored and managed 24/7 each and every day. The STB is less likely to be unplugged from the wall, and unlike mobile devices, STBs are always connected to the Internet, even when they are not used. STBs do not need local health-care application storage. Rather, they only need an engine to render a cloud-based app. STB variety is limited, and the number of apps used in manageable. With STB, there is no need to support generations of mobile devices, and there is no need to support generations of OSs.

Cable providers are also able to offer improvements by ensuring reliable connections between the care-giving cloud and the patients. The cable modem termination system (CMTS), the home gateway, and all the connections between it and the medical devices need to be rock solid. Cable service providers have what it takes to create and maintain strong connectivity, because they have already invested heavily in very reliable networks. These networks are fiber-deep, with diverse routing and powering infrastructure. This infrastructure is able to be baked-up by alternative sources of energy. The central office DOCSIS systems have been designed and built by companies with deep roots in telecommunications. Therefore, they feature very high uptime specifications.

Modern STBs have built-in DOCSIS Cable Modems that are separate from the Internet Cable Modem. STBs are being installed by the MSO technicians what guarantees their reliable connection. Cable companies have taken major chunks of broadband services away from Telcos. This did not happen by chance. It happened because of superior connectivity and speeds offered by cable companies. Cable Home Gateways undergo tough and comprehensive compatibility and reliability tests performed by independent parties prior to their deployment in the field. Some provide lifeline telephony services at par with legacy telco equipment. Some modern STBs already feature low-energy Bluetooth interfaces to connect the remote control. These can be compatible with the medical devices. New generations of STBs are being specified by the service providers and OTT companies, including Android TV require BLE interface.

Cable providers can consolidate the fragmented OTT apps from one-per-device and integrate them into an intuitive, easy to interact with web app that would be under the control of the service provider. The new app could use the entertainment content as a reward for good behavior, compliance with the monitoring regime, and the taking of medication.

Today's health monitoring system landscape is highly fragmented. Every manufacturing company wants consumers to use their individual apps. They argue their app layer security and look for future recurring subscription revenue. However, these apps do not allow for the creation of very useful if-this-then-that (IFTTT) sequences of events across multiple devices. They do not allow for the pooling of measurements into sessions that could be reminded of via notifications and automatically invoked for the consumer. Such a pooling would be very useful for patients, particularly for those patients who are being asked to do multiple tests and measurements in a continuous session.

STB portals, on the other hand, have ample TV screen space. This screen space could show, at the same time, multiple devices that need to be operated. It could guide the patient through any procedures in a step-by-step manner. The procedures might include anything from taking the patient's pills from the connected dispenser, to taking the patient's temperature, and to performing blood pressure measurement. The application could urge patients to get on a weight scale and instruct them to test pulse and oxygen levels.

A key advantage of having the connected health application on the cable service provider's STB is the ability to link the cable service provider's or patient's VoD system to the application. An innovative rewards system can be integrated with the app to offer rewards to those who comply with the regime they agree to follow. For example, a free movie could be granted to the patient or care receiver when they follow the schedule and report the measurements.

Health, the IoT and Security

There is little doubt that the emerging world of IoT devices specific to the HealthCare industry can provide extreme value in terms of controlling costs, access to health care, more accurate and current monitoring of treatment programs, and remote diagnostics. However, it is also critical to examine the risks that exist with these new devices.

At a minimum, the consumer data accessible through such IoT devices should be considered extremely private, and fully protected under the federal HIPAA rules. The relevant data must be accessible only to the intended parties, and only for the intended uses. For example, blood pressure data must be accessible to the family physician or specialist, and only for the purposes of medical diagnostics. Even the proper doctor should not employ the data for any other purpose. Thus not only the device must understand the

proper connection for data transfer, but the processing software at the doctor's office must be designed for diagnostic uses alone, to the extent possible.

Secondly, malicious intent must be blocked to the maximum extent possible. Hacking to steal data is certainly not desired, but hacking that plants a virus and defeats the intent of the IoHT device could cause improper readings. Not only can normal data be changed to show abnormal readings, but abnormal readings might be blocked, disguising or delaying real medical needs and treatment. Health devices can be held hostage, in the sense that private data can be threatened with exposure unless some ransom is paid. For health devices intended to accept commands for remote treatment, the impact can be even more devastating, including improper dosages and even death. The need for security in all phases of the device design and usage cannot be understated.

Further, hacking should not be viewed as strictly a remote access concern, although that is the primary attack mechanism. It should also be extremely difficult to alter such a device physically, when an attacker has the opportunity for direct access to the device. This access would include factory access during the manufacturing process, warehousing access, shipping access including interception and replacement, and even access once a device has been installed in the home. Where a consumer's health is concerned, everything must be considered.

In addition to IoHT device security design having regulatory aspects (HIPAA) as well as a safety concerns (hacking), it is important to look at this as more than just "meeting requirements" and "preventing problems." Secure design of such devices can be a competitive advantage. This type of device and marketplace carries the very real opportunity for increased value for increased security. Most consumers can accept that health diagnostic and treatment is not an area to choose the lowest price; reputation based upon better security in design and in process can translate to better value and thus greater market share.

Secure Device Design

What are the general guidelines for designing health targeted IoT devices in a secure manner?

1. Application layer communication protocols and security algorithms defined by a reputable industry consortium, such as OCF [Open Connectivity Foundation]. Health diagnostics and treatment is no place for creativity in this domain. Protocols and algorithms must have survived the test of time, and been subject to broad scrutiny and analysis. A unique algorithm invented during the design process may seem to present advantages of various types, yet time and again such an approach has proven disastrous in the security field. The protocols defined for health environments will ensure that all private data never appears in transit in unencrypted (clear) form.
2. HIPAA rules also require that private data stored in the device and in the servers in the infrastructure medical systems must be encrypted.
3. The software that runs on such an IoHT device must be updatable in the field. The manufacturer has an on-going obligation to keep these devices up to date in terms of flaws discovered after shipment. Any such update scheme must be secure; as such schemes are not necessarily standardized, they must use industry accepted techniques, and be scrutinized for flaws during the design process. Further, solutions cannot be burdensome to the consumer. As an example, mailing a security update notice to the consumer, requesting that they take some obscure action involving pressing hidden reset buttons, and even worse, loading code manually, are not going to be acceptable ways to assure then on-going safety and protections assumed in these devices.

From a long term view, the responsibility for selling secure health industry IoT devices is not just about shipping working devices and updating them as needed. There is an entire lifecycle to consider. Any company that choose to exit the business, or drop a product line, still must commit to continuous update, or arrange for another company to “take over” that process. Further, when devices reach end-of-life, there needs to be some procedure for disposal, some proper handling recommendations, as private data may still reside in the device.

4. A unique identity is critical for each device. Most likely, it is a specific requirement of the industry protocols described in the first item above, and often in the networking layers below the application layer. Diagnosis and treatment are always specific to an individual, so the system approach must uniquely identify the device or devices in use by any specific individual. To achieve this, devices likely have factory identities installed by the manufacturer, which are used to bootstrap a process for installing application level identities relevant to the application ecosystem in use (again, for example, OCF).
5. Secure software design practices should be followed. Many scanning tools exist to discover vulnerabilities in software, and the use of one or more of these should be standard practice.
6. To minimize hacking, devices should include physical robustness protections. Devices should be difficult to disassemble, and once disassembled, critical data should not be exposed on easily accessible busses or circuit connections. Test modes should not exist that defeat any protections. Examples for some of these practices are described in NIST publication [FIPS levels] and in the example Robustness Rules of Appendix C of the DTLA License document. This last document comes from the content protection industry, which shares some of the same needs as the health industry.
7. In general, all aspects of the secure processing within the infrastructure and the device should be documented and reviewed. Security process checklists often assist, with entries for all aspects of what can be done to protect the integrity and privacy of health data and device secrets. Any audits and reviews of this type should be recorded, and made available as the need arises. If any industry validation group is created for assurance, the manufacturer should submit designs to that group, and market compliance to its requirements.

Some Security Specifics:

1. All software resident on the consumer device must be signed cryptographically, to ensure that a trusted party wrote the code that it executes. Most often, the cryptographic algorithm is RSA, although code signing with elliptic curve cryptography can also be done. Note that the trusted party creating and signing the code is likely to be the device manufacturer, and there are obligations to perform such signing operations very securely, never exposing the signing key.
2. Software must carry a version number, also signed or otherwise secured, so that any device can determine if it has up-to-date code. Any device that has software that is out of date would initiate a download of the update. Devices must be able to check to see if their software is current; if such check is unable to be performed, there may be an attack underway, and a proper response is required.
3. Signature checking in the device must be performed in such a way that it can be traced to a “hardware root of trust,” making it extremely difficult to circumvent the signature check process. Generally, this requires the device to have a single secure silicon device implementing a secure boot process at power up, which builds trust a layer at a time.
4. Each device should have a unique manufactured cryptographic identity, typically composed of an elliptic curve cryptographic key pair, and a certificate attesting to the authenticity of the public key of the pair. The certificate should originate from a trusted Certificate Authority, and the

private key of the pair should be protected within the IoT device so that its use is not exposed outside of the secure silicon device. This secure identity can anchor the download of application level secure identity information.

Once these requirements are fully appreciated and understood, it is easy to see why Cable Home Gateway devices are ideal IoHT “concentrators” or partners in the home solution for health applications. Cable gateways implement all of the secure design principles listed above, as they often include content protection as well as privacy features for their normal operation support. Their connection to the Internet and thus any health systems in the infrastructure is highly reliable and secured, and they often include home networking radios within the design, making them ideal for in home connectivity to various IoHT devices.

Further, such devices typically operate in a closed system, with signed code, secure boot, and well controlled access. Compared to today’s mobile devices, they represent a far more secure choice for IoHT functionality. It is far too easy to install applications on today’s mobile phones with their more open ecosystem that either steal data directly or allow hackers to gain access.

Conclusion

Soon, we will experience a wave of baby-boomers entering the phase of life in which they will need health care, and consequently, there will be an increase in the pressure on the global health systems. The Internet of Health Things (IoHT) is promising to significantly lower costs of moving health monitoring to the aging population’s homes by developing standards based devices that can be easily connected to the cloud.

Cable operators are in an advantageous position to partner with care-giving service providers to enable Health-over-Cable (HoC) offerings to the patients at their homes. Health/wellness monitoring devices can be connected via a managed network. Set-top boxes (STBs) can connect to the medical devices using built-in BLE interfaces. TV screens connected to STBs can act as HoC services portals. On-screen notification and reminders can force the patient to comply with the regime. Patients can be rewarded by MSO’s video offerings. Condition-specific V-D can be streamed to the patient’s homes. Most notably, this can all be accomplished in a secure and private cable network.

There are some important things to remember when offering improvements to the Internet of Health Thing. Security for all is paramount. First of all, the normal HIPAA level of privacy and protocol security would be designed into the system. Then, patients would have to be assured that malicious hacking of all types would be deterred. They would also have to be assured that a response to any detected incident would be immediate and swift. Their health security and privacy must be the most important factor in the creation of the system.

Cable operator equipment is the ideal place for true physical security and threat mitigation. Therefore, cable operators are ideally positioned to monitor and control these devices. After all, they have a long history and a lot of experience with video content protection, secure data delivery, and secure software updates.

Abbreviations

BLE	Bluetooth Low Energy
CMTS	Cable Modem Termination System
DOCSIS	Data Over Cable Service Interface Specification
HoC	Health over Cable
HoOTT	Health over OTT
Hz	hertz
IFTTT	if-this-than-that
IoHT	Internet of Health Things
IoT	Internet of Things
ISBE	International Society of Broadband Experts
MSO	multiple system operator
OS	Operating System
OTT	Over-the-Top
SCTE	Society of Cable Telecommunications Engineers
STB	Set-top Box
UX	user experience
VoD	Video on Demand

Bibliography & References

Population Pyramids of the World from 1950 to 2100, aging charts and stats:

<http://www.populationpyramid.net>

Digital Transmission Licensing Administrator (DTLA): <http://www.dtcp.com/agreements.aspx>

National Institute of Standards and Technology (NIST) publication [FIPS levels]:

<http://csrc.nist.gov/groups/STM/cmvp/standards.html>

Open Connectivity Foundation: <https://openconnectivity.org/>

The Office of the National Coordinator for Health Information Technology, a division of the U.S. Department of Health and Human Services: HealthIT.gov

Powerful LPWAN Solutions for IoT

How Low Power Wide Area Networks will Accelerate Smart City and Connected Business Initiatives

A Technical Paper prepared for SCTE•ISBE by

Chris Kocks

Director, IoT Practice
Pure Integration, LLC.
13454 Sunrise Valley Drive, Suite 500
Herndon, VA 20171
678-467-7458
Chris.Kocks@pureintegration.com

Technical testing contributions by:
Igor Sabaldash, pureIntegration
Michael Grudsky, Semtech
Brian Dunn, Semtech

Introduction

1. Abstract

The rapid growth of IoT has dramatically expanded the number of wireless devices in the home and around us. This expansion has been remarkable considering large dependence upon short range, local area networks, and antiquated cellular networks for longer range. One of the key changes enabling an explosion of IoT solutions is the deployment of LPWAN (low power wide area network) with accessibility measured in kilometers rather than meters. Greater ranges open the flood gates to more comprehensive hybrid solutions that are driving smart cities, business, automotive, and industrial connected initiatives. What are the use cases and value that will be enabled with LPWAN? What are the advantages over similar 5G cellular solutions being developed? Which platforms provide the most promise? How will operators position as key players in this new ecosystem? Get ready see real examples of LPWAN use cases and how operators will drive the next wave of evolution in IoT. This paper will share some approaches and examples of how LPWAN will drive value. Readers will also learn about leading solutions in this space and their unique advantages.

Audience:

The audience for this paper is anyone planning or designing an IoT solution to serve home, business, municipality, and industrial purpose. Product executives, solution architects, and operations professionals will be interested to understand the strategic opportunities and impacts in addition to specific use cases and value propositions for leveraging LPWAN solutions. Readers will get an overview of the technology, its applications, and some of leading wireless solutions in this space with their unique advantages.

2. Preview

Wireless technologies have played an important role in technology advancements for decades. In fact, we are amid perhaps the greatest technology transformations ever. Many experts have tagged this transformation Industry 4.0, as in the fourth massive shift in technology since the advent of commercial implementation of mechanization and steam/water power. What makes this shift greater than all others is the speed, reach, and predictive intelligence of new wireless technologies and accompanying applications.

Connectivity, home automation, smart buildings, energy conservation, security, health monitoring, business applications, transportation, agricultural and industrial applications are all driving factors for wireless communications. The number of wireless protocols and technologies have grown to meet various needs across industries. These multiple wireless technologies have many different requirements and benefits in terms of bandwidth, cost, privacy, installation, and operation. The development of Internet-connected technologies particularly requires implementing solutions that harness value, savings, and improve quality of service while remaining safe from increasing security threats.




Figure 1 - IoT Market Forecast from IHS

In past decades, wireless technologies have focused heavily on short-range Wi-Fi, long-range cellular, and satellite for global reach. In recent years, the acceleration of connected devices and sensors in every industry and aspect of our lives has created great demand for short and mid-range reach. Millions (soon to be Billions) of connected devices communicate simple status infrequently requiring very low bandwidths and small batteries that last for years. Current cellular protocols require significant bandwidth, large batteries, at higher cost compared to newer LPWAN protocols specifically designed for large scale, efficient communications.




Figure 2 - IoT Economic Impact Forecast from McKinsey

Several standard-based LPWAN networking protocols allow fast growth and implementation of self-healing mesh networks, which are more reliable network arrangements. Some of these networking protocols, including LoRaWAN, SigFox, Senet, The Things Network, and RPMA are based on protocols, operating in the unlicensed 915Mhz range. They enable cost-effective communication between devices with low bandwidth, low power, and low installation costs.

This paper describes several commercial flavors of LPWAN solutions that have been developed to meet recent monumental changes in the industry. We will illustrate frameworks based on existing solutions for practical, readily usable and hardware independent low power solutions. We will demonstrate compelling evidence for the real power offered by LPWAN solutions.

This paper highlights the importance of low-power networks and cites some of the important use cases that can be fulfilled. The focus is on long-range scenarios for smart cities, transportation, agriculture, and industrial. It also includes several localized short-range home & business scenarios for comparison.

Additionally, this paper includes test results and examples for professionals to consider to help them be most effective in deploying LPWAN wireless platforms.

Wireless Technologies

1. Historic Overview

Let's begin with a short history of wireless technologies for context. What do you think was the first invention of wireless voice communication? Many would guess the telegraph, or the telephone, or more recently Wi-Fi. The key to the answer is “wireless voice” and very few would know the photophone was the first *successful* wireless voice communication invention in 1880. In 2016 I attended a wireless association event in D.C. and walked past the Franklin School with a colleague after dinner on L street near the White House. The historic building caught our attention and we were surprised to stumble upon a plaque fixed to the wall of the building highlighting the achievement by Alexander Graham Bell and his associate Charles Sumner Tainter. The historic marker explained the photophone invention transmitted voice via a beam of light and was considered by Graham as his greatest achievement. The photophone paved the way for fiber optics which achieved wide-spread implementation 100 years later in the 1980s.




Figure 3 - Historic Wireless Invention

2. Modern-day Wireless

Fast-forward to modern day wireless technologies. Wi-Fi has been the most prevalent standard for local area connectivity. Cellular has been most prevalent for wide area communications. Satellite has been most prevalent for long-distance global communications. Each of these wireless technologies began to serve very specific uses with unique requirements. With the rapid proliferation of connected devices, cloud-based services, and global mobile device usage these wireless technologies are beginning to serve in a complimentary hybrid wireless network.

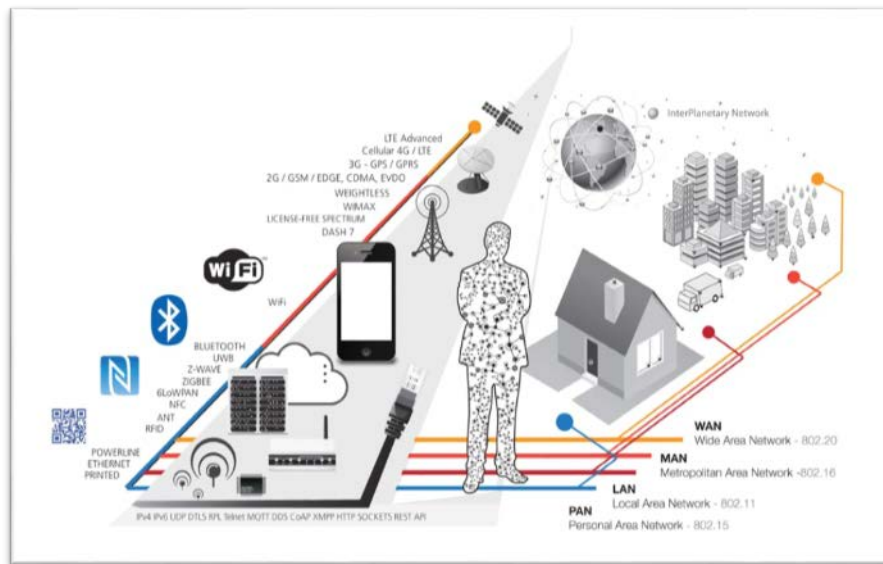


Figure 4 - IoT Wireless Technologies Spectrum by Postscapes.com

3. Short-Range Wireless

For the purpose of this paper we will characterize short-range wireless as less than 300 meters range, usually in a single premise or within shouting range. The Wi-Fi protocol is the most popular short-range consumer wireless solution due in large part to the standard being adopted many years ago by the laptop and portable device industry. Connectivity became extremely important for traveling professionals, students, and growing numbers of startup companies. Rabid desire for access to email, internet sites, and video content has driven the growth of indoor and outdoor public Wi-Fi deployment to meet consumer and business demand for access.

Bluetooth technology usage exploded in commercial usage in the early 2000's with wireless headsets, portable speakers, audio equipment, and more recently for rapid growth in wearables and sensors. Advertisers are also employing Bluetooth beacons for proximity-based ads. Most of these devices run on small batteries with wearables especially enabled with rechargeable batteries. Bluetooth has evolved through several generations of specifications, adding more advanced capabilities for bi-directional communications, greater security, and battery usage.

ZigBee and Z-wave protocols provide the basis for many home security and home automation systems due to their greater focus on securing device access and network traffic. Large cable and telco operator home security platforms have seen much success and high QOS with these solutions.

Wi-Fi, Bluetooth, ZigBee, and Z-Wave protocols are all based on IEEE 802.x specifications operating in the 2.4Ghz and 5Ghz ranges. Their standardization has enabled wide spread industry adoption and commercial success for short-range wireless communication.

In industrial, retail and other industries we have seen healthy adoption of other short-range technologies like RFID (radio frequency identification) and NFC (near field communication) technologies. RFID applications include asset tagging, electronic toll collection, and building access credentials. Multiple low and UHF frequencies are used by various vendors. Standards have also helped the adoption of RFID through a mix of several bodies including ISO, EPC, IEC, Dash7, ASTM. NFC is often used for asset tracking and payment transactions. Standards have been developed by ISO, and GSMA continues to develop standards for mobile device support of NFC.

Short-range wireless solutions are very capable of managing large numbers of sensors and large volumes of transactions as long as sensors and devices are in close proximity with their gateway or source of network connectivity. Range becomes an issue as use cases include highly mobile, remote locations, or broad areas that need to be covered. Typically, the longer the range, the more power is required to transmit and receive a signal. Sensors with significant physical obstructions (concrete walls, buildings, hills) or sensors at the extremities of the network range will use significantly more power to maintain adequate signal strength. That may be ok for powered devices; however, it becomes a major problem as batteries drain much faster, becoming an operational nightmare for companies needing to deploy thousands of sensors. As more sensors are being deployed for smart city, agricultural, and industrial applications we have a growing need to go beyond short-range use cases.

4. Long-Range Wireless

We will characterize long-range wireless as greater than 300 meters, typically serving a several kilometer range. Mobile phones are probably the most recognized commercial device utilizing long-range wireless. Common cellular protocols have evolved over generations from 1G Analog, 2G GSM & IS-95, 3G UMTS & IS-856, 4G LTE & WiMAX, and soon to be 5G protocols. Do you remember your first mobile phone? Mine was a bag phone that we used in the car for frequent long highway trips back in the early 90s. Even though reception was spotty, we still managed to benefit from analog voice coverage along much of our drive. In current day, we benefit from very good coverage in most areas with high quality digital data, large video bandwidth, and voice support; although, just shy of the high QOS expectations and performance of land lines.

Telcos and industry trade bodies GSMA, 3GPP, etc. have been working hard to advance the specifications for 5G solutions, conducting significant trials throughout 2017. Telcos have been working on variants of low-latency and low-power alternatives to support the “connected” world. Low-latency applications like assisted or self-guided vehicles manage constant communication, requiring significant bandwidth and power. 5G technology is a good fit for these use cases, at a reasonable cost.

Conversely, low-data-rate applications like connected meters for water, power, sewage, and lighting use infrequent bursts of small data packets to report status. The compact nature of these applications has spurred rapid development of “low-power” technologies that enable sensors to run on small batteries for years. With aggressive plans for low-data-rate device deployments, the industry is focusing on solutions

with a better cost equation. The following examples of LPWAN are intended to be representative of the technology sub-groups, rather than an exhaustive list. The industry is moving so fast that any roundup of solutions will change dramatically a year later.

LPWAN (Low Power Wide Area Network) technologies have been developed to manage large-scale deployments with low-bandwidth requirements, and operate at a very low cost compared to modern high bandwidth alternatives. There are several competing flavors of LPWAN deployed around the world, including those based on the proprietary LoRa radio CSS (Chirp Spread Spectrum), UNB (Ultra Narrow Band), RPMA (Random Phase Multiple Access), Narrow Band IoT, LTE-MTC, and more.

The LoRa Alliance with their 300+ members has made a large impact with world-wide deployment of their LoRaWAN protocol over the past couple of years. This group of solutions includes open-source solutions from The Things Network in Europe and Carnegie Mellon's OpenChirp management layer. Commercial platform solutions are provided by Senet, Everynet, ThingPark by Actility, and most recently Comcast machineQ. These solutions operate in the unlicensed 868MHz range in Europe and the 915MHz range in North America. The great number of members and solutions has strengthened and matured LoRaWAN tremendously in just a few years-time. This space is crowded with platforms and success for each platform will be influenced by serving large industries and building strong partner ecosystems.

Ultra Narrow Band solutions were first developed by Telensa to utilize very small bandwidth packets measured in hundreds of Hz (instead of KHz or MHz) operating in unlicensed sub gigahertz range for greater penetration of walls and ground. SigFox is a leading UNB-based commercial solution using the ISM radio band. The solution excels at long ranges and ground penetration; however, the maximum payload/message size is roughly one fifth that of LoRa's. The company is based in France that has had success in Europe and North America. There are additional variants promoted by the Weightless SIG and adopters.

Cellular-based variants of LPWAN include LTE-MTC, an evolution of LTE for machine type communications and NB-IoT by 3GPP. These solutions leverage existing "in-band" spectrum or "standalone" unused frequency. Benefits of NB-IoT solutions include no gateway cost as all devices connect to base station towers, which has some economies of scale in large deployments. NB-IoT often has better signal performance in urban or dense environments in addition greater bandwidth capacity.

Another LPWAN solution by Ingenu is RPMA (random phase multiple access) was designed for broad coverage beyond 300 square miles and over 600 kbp/s uplink. This solution operates in the 2.4GHz range which can be impacted by overall traffic, and has lower capability to penetrate walls and floors.

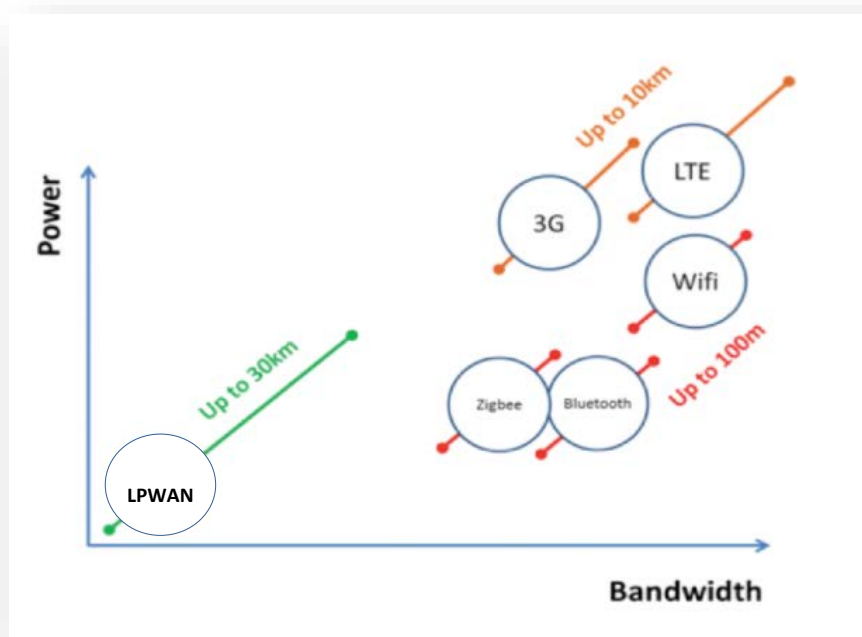


Figure 5 - Wireless Power vs. Bandwidth Niche

The following is a comparison of common wireless technologies.

Table 1 - Wireless Feature Comparison

LPWAN	Short-Range	Cellular 4G LTE
Range is 6km-10km	Range is 100m-300m	Range is 10km-20km
Unlicensed spectrum 915Mhz	Licensed Spectrum 2.4Ghz	Licensed Spectrum
Long battery life (8-10yrs)	Short battery life (1-2yrs)	Short battery life (hrs or wks)
High latency	Low-Medium latency	Low latency
Low data rates (10s of Kbps)	Medium data rates (100s of Kbps)	High data rates (Gbps)
Low cost sensors	Low cost sensors	High cost sensors
Medium cost gateways	Medium cost gateways	High cost base station
100s of devices/gateway	100s of devices/gateway	1000s of devices/base station
High penetration of walls/floors	Medium penetration of walls/floors	Medium penetration of walls/floors
Better in rural and indoor	Better indoor	Better in urban and indoor

LoRaWAN Range Testing

1. Objectives and Scope

The overall objective of the LoRaWAN range testing was to demonstrate the long distance and strong signal capabilities of the technology compared to other wireless technologies.

The scope of testing included the use of:

1. Stationary Semtech Pico Cell gateways
2. Roving Microchip sensors
3. Suburban and rural environments
4. Ground floor, subway level, and above 150 meters high

2. Test Devices and Environment

For our initial testing, we assembled a new prototype Semtech pico cell LoRa gateway on a Raspberry Pi running RDK-B software along with a GPS sensor, all connected to the Semtech test environment. Thanks to Semtech, the founder of LoRaWAN and the LoRa Alliance, for graciously working with us to conduct testing.

Our first round of testing the prototype was intended to shake down the environment and prove the concept of using a USB form-factor pico cell on the RDK-B platform. After making several changes to the packet forwarder and RDK environment we were on our way to building out a robust test environment with standardized test equipment.



Figure 6 - LoRaWAN Test Device Prototype

In the next round of testing we modified the environment to enable a more standardized and repeatable setup that works in the lab, home, or field. In this implementation, we added a 7" touchscreen along with a slim battery power source to enable portability and ease of controlling the Linux environment without having to remote into a headless device. Not only does this setup enable us to test mobility of sensors, but also mobile gateway scenarios.

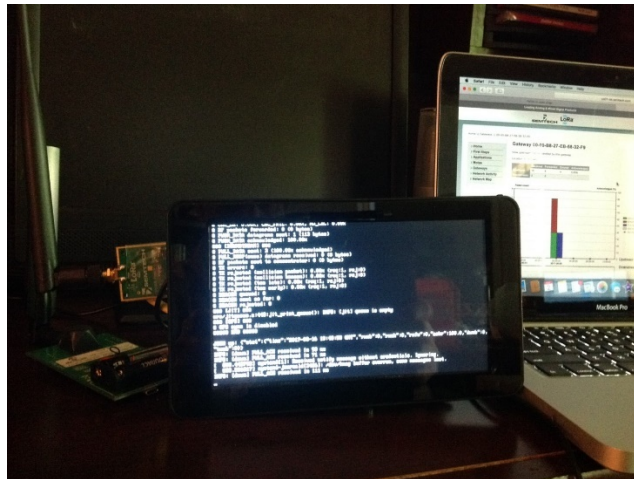


Figure 7- LoRaWAN Mobile Gateway Device

The current test environment has evolved to include cable operator all-in-one gateway devices hosting an integrated LoRa radio operating on multiple leading network servers. As a result, we were able to conduct tests on stationary commercial-grade gateways in addition to mobile gateways to cover dynamic scenarios.



Figure 8 - LoRaWAN Commercial-grade Test Devices

3. Test Scenarios

For the purpose of this paper we focused primarily on a stationary gateway, taking measurements from sensors at different distances (ranges) under multiple conditions. Our range testing was conducted using several common scenarios and use cases; however, we present data from a primary basic scenario to illustrate the power of LPWAN. The intent is to also illustrate performance in examples that were as close to real-life as possible.

1. Range of multiple distances from a single gateway in a suburban environment.

2. Handoff from one gateway to another in a suburban environment.
3. Range of multiple distances from multiple gateways in an urban environment.
 - a. Ground level
 - b. 150 Meters above ground (future)
 - c. Subway level (future)
4. Range of multiple antennas with different dbi levels.

4. LoRaWAN Test Results

The following test results were conducted by Igor Sabaldash of pureIntegration. Baseline Test: with a 0 dbi antenna on both the Mote and Gateway while roaming, beginning at 1m distance from the gateway reference implementation.

Table 2 - Baseline Test Data

Time	Mote	Seq	Freq (MHz)	Mod	BW (Hz)	SF	Coding Rate	ADR	Gateway	Chan	RSSI (dBm)	SNR (dB)
7/9/2017 19:17	00-00-00-00-15-22-46-89	52	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-71	7.5
7/9/2017 19:17	00-00-00-00-15-22-46-89	51	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-79	-6.2
7/9/2017 19:16	00-00-00-00-15-22-46-89	50	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-84	-0.2
7/9/2017 19:16	00-00-00-00-15-22-46-89	49	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-83	1.8
7/9/2017 19:16	00-00-00-00-15-22-46-89	48	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-84	N/A
7/9/2017 19:16	00-00-00-00-15-22-46-89	47	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-84	-5
7/9/2017 19:16	00-00-00-00-15-22-46-89	45	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-89	-6.2
7/9/2017 19:16	00-00-00-00-15-22-46-89	44	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-87	-3.8
7/9/2017 19:15	00-00-00-00-15-22-46-89	43	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-88	-8
7/9/2017 19:14	00-00-00-00-15-22-46-89	36	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-84	-12.5
7/9/2017 19:12	00-00-00-00-15-22-46-89	20	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-88	-11.8
7/9/2017 19:12	00-00-00-00-15-22-46-89	14	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-84	-7
7/9/2017 19:11	00-00-00-00-15-22-46-89	13	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-87	-4.5
7/9/2017 19:11	00-00-00-00-15-22-46-89	12	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-87	-3
7/9/2017 19:11	00-00-00-00-15-22-46-89	11	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-83	-3.8
7/9/2017 19:11	00-00-00-00-15-22-46-89	9	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-83	-2.8
7/9/2017 19:11	00-00-00-00-15-22-46-89	8	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-83	-7.2
7/9/2017 19:11	00-00-00-00-15-22-46-89	7	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-81	5.5
7/9/2017 19:10	00-00-00-00-15-22-46-89	6	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-73	3.5
7/9/2017 19:10	00-00-00-00-15-22-46-89	5	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-76	3.5
7/9/2017 19:10	00-00-00-00-15-22-46-89	4	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-67	9
7/9/2017 19:10	00-00-00-00-15-22-46-89	3	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-66	10
7/9/2017 19:10	00-00-00-00-15-22-46-89	2	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-75	5.8
7/9/2017 19:10	00-00-00-00-15-22-46-89	1	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-66	9.2




Figure 9- Baseline Test GPS Map




Figure 10 - Baseline Test RSSI Range

Test # 2: With 2 dBi antenna on a mote, and 0 dbi antenna on the gateway.

Table 3- Test Cycle #2 Test Data

Time	Mote	Seq	Freq (MHz)	Mod	BW (Hz)	SF	Coding Rate	ADR	Gateway	Chan	RSSI (dBm)	SNR (dB)
7/9/2017 19:33	00-00-00-00-15-22-46-89	57	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-72	9
7/9/2017 19:33	00-00-00-00-15-22-46-89	56	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-55	11.5
7/9/2017 19:33	00-00-00-00-15-22-46-89	55	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-53	9
7/9/2017 19:33	00-00-00-00-15-22-46-89	54	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-66	10.5
7/9/2017 19:33	00-00-00-00-15-22-46-89	53	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-59	10.5
7/9/2017 19:33	00-00-00-00-15-22-46-89	52	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-71	8.5
7/9/2017 19:32	00-00-00-00-15-22-46-89	51	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-63	11
7/9/2017 19:32	00-00-00-00-15-22-46-89	50	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-73	7.8
7/9/2017 19:32	00-00-00-00-15-22-46-89	49	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-82	4.8
7/9/2017 19:32	00-00-00-00-15-22-46-89	47	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-84	-5.8
7/9/2017 19:32	00-00-00-00-15-22-46-89	46	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-84	0.2
7/9/2017 19:32	00-00-00-00-15-22-46-89	45	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-78	3.5
7/9/2017 19:31	00-00-00-00-15-22-46-89	44	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-80	1.5
7/9/2017 19:31	00-00-00-00-15-22-46-89	41	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-84	-3.5
7/9/2017 19:31	00-00-00-00-15-22-46-89	40	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-82	-3.8
7/9/2017 19:31	00-00-00-00-15-22-46-89	38	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-84	0.2
7/9/2017 19:31	00-00-00-00-15-22-46-89	37	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-83	-3.2
7/9/2017 19:30	00-00-00-00-15-22-46-89	36	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-85	-4.2
7/9/2017 19:30	00-00-00-00-15-22-46-89	34	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-84	-7
7/9/2017 19:30	00-00-00-00-15-22-46-89	33	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-84	-11
7/9/2017 19:29	00-00-00-00-15-22-46-89	26	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-85	-9.5
7/9/2017 19:29	00-00-00-00-15-22-46-89	25	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-84	-13
7/9/2017 19:29	00-00-00-00-15-22-46-89	24	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-84	-4.5
7/9/2017 19:29	00-00-00-00-15-22-46-89	23	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-85	-8.2
7/9/2017 19:29	00-00-00-00-15-22-46-89	22	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-85	-7
7/9/2017 19:28	00-00-00-00-15-22-46-89	20	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-84	-2.8
7/9/2017 19:28	00-00-00-00-15-22-46-89	19	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-85	-6.5
7/9/2017 19:28	00-00-00-00-15-22-46-89	18	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-83	-5.2
7/9/2017 19:28	00-00-00-00-15-22-46-89	17	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-82	-4
7/9/2017 19:28	00-00-00-00-15-22-46-89	16	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-83	-2.8
7/9/2017 19:28	00-00-00-00-15-22-46-89	15	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-84	-6.5
7/9/2017 19:27	00-00-00-00-15-22-46-89	14	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-83	4.5
7/9/2017 19:27	00-00-00-00-15-22-46-89	13	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-85	3.5
7/9/2017 19:27	00-00-00-00-15-22-46-89	12	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-83	-4.2
7/9/2017 19:27	00-00-00-00-15-22-46-89	11	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-78	2
7/9/2017 19:27	00-00-00-00-15-22-46-89	10	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-87	1.2
7/9/2017 19:27	00-00-00-00-15-22-46-89	9	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-70	5.5
7/9/2017 19:27	00-00-00-00-15-22-46-89	8	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-75	6.5
7/9/2017 19:26	00-00-00-00-15-22-46-89	7	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-68	1.2
7/9/2017 19:26	00-00-00-00-15-22-46-89	6	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-79	1.2
7/9/2017 19:26	00-00-00-00-15-22-46-89	5	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-58	9.2
7/9/2017 19:26	00-00-00-00-15-22-46-89	4	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-55	8
7/9/2017 19:26	00-00-00-00-15-22-46-89	3	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-53	11.2
7/9/2017 19:26	00-00-00-00-15-22-46-89	2	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-61	4.5
7/9/2017 19:26	00-00-00-00-15-22-46-89	1	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-53	11




Figure 11 - Test Cycle #2 GPS Map



Figure 12 - Test Cycle #2 RSSI Range

Test # 3: With 6 dBi antenna on a mote, and 0 dBi antenna on the gateway.

Table 4- Test Cycle #3 Test Data

Time	Mote	Seq	Freq (MHz)	Mod	BW (Hz)	SF	Coding Rate	ADR	Gateway	Chan	RSSI (dBm)	SNR (dB)
7/9/2017 19:51	00-00-00-00-15-22-46-89	51	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-71	10.5
7/9/2017 19:51	00-00-00-00-15-22-46-89	50	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-61	11.5
7/9/2017 19:50	00-00-00-00-15-22-46-89	49	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-65	11.2
7/9/2017 19:50	00-00-00-00-15-22-46-89	48	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-75	11.2
7/9/2017 19:50	00-00-00-00-15-22-46-89	47	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-91	9.2
7/9/2017 19:50	00-00-00-00-15-22-46-89	46	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-87	2.5
7/9/2017 19:50	00-00-00-00-15-22-46-89	45	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-86	-3
7/9/2017 19:50	00-00-00-00-15-22-46-89	44	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-90	4.2
7/9/2017 19:50	00-00-00-00-15-22-46-89	43	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-91	-2.5
7/9/2017 19:50	00-00-00-00-15-22-46-89	42	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-90	-2
7/9/2017 19:49	00-00-00-00-15-22-46-89	41	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-82	1.5
7/9/2017 19:49	00-00-00-00-15-22-46-89	40	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-88	5.2
7/9/2017 19:49	00-00-00-00-15-22-46-89	39	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-91	-3.2
7/9/2017 19:49	00-00-00-00-15-22-46-89	38	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-89	-5.5
7/9/2017 19:49	00-00-00-00-15-22-46-89	37	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-93	-7
7/9/2017 19:49	00-00-00-00-15-22-46-89	36	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-90	0.2
7/9/2017 19:49	00-00-00-00-15-22-46-89	35	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-89	-0.5
7/9/2017 19:49	00-00-00-00-15-22-46-89	34	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-93	-6.5
7/9/2017 19:48	00-00-00-00-15-22-46-89	33	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-91	-9.2
7/9/2017 19:46	00-00-00-00-15-22-46-89	18	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-91	-10.8
7/9/2017 19:46	00-00-00-00-15-22-46-89	16	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-91	-5.2
7/9/2017 19:46	00-00-00-00-15-22-46-89	14	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-90	-6
7/9/2017 19:46	00-00-00-00-15-22-46-89	13	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-90	-4.8
7/9/2017 19:46	00-00-00-00-15-22-46-89	12	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-93	-5
7/9/2017 19:45	00-00-00-00-15-22-46-89	11	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-91	-1
7/9/2017 19:45	00-00-00-00-15-22-46-89	10	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-89	3.5
7/9/2017 19:45	00-00-00-00-15-22-46-89	9	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-81	3.8
7/9/2017 19:45	00-00-00-00-15-22-46-89	8	903.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	7	-89	5.8
7/9/2017 19:45	00-00-00-00-15-22-46-89	7	902.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	1	-84	3.2
7/9/2017 19:45	00-00-00-00-15-22-46-89	6	902.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	0	-79	1.2
7/9/2017 19:45	00-00-00-00-15-22-46-89	5	903.1	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	4	-77	10.2
7/9/2017 19:45	00-00-00-00-15-22-46-89	4	903.3	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	5	-76	9.8
7/9/2017 19:44	00-00-00-00-15-22-46-89	3	902.9	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	3	-60	9.5
7/9/2017 19:44	00-00-00-00-15-22-46-89	2	902.7	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	2	-51	8.8
7/9/2017 19:44	00-00-00-00-15-22-46-89	1	903.5	LoRa	125000	SF10	4/5	off	00-00-B8-27-EB-0C-F9-F3	6	-70	9



Figure 13 - Test Cycle #3 GPS Map




Figure 14 - Test Cycle #3 RSSI Range

The following testing was conducted by Semtech's Michael Grudsky and team in Iowa to demonstrate performance in the unobstructed open range. The tests included a gateway positioned on an 80-foot-high pole and a mobile GPS sensor communicating with the gateway. The image below illustrates the coverage obtained in the open range was the size of the Chicago metropolitan area. Test results courtesy of Michael Grudsky from Semtech.

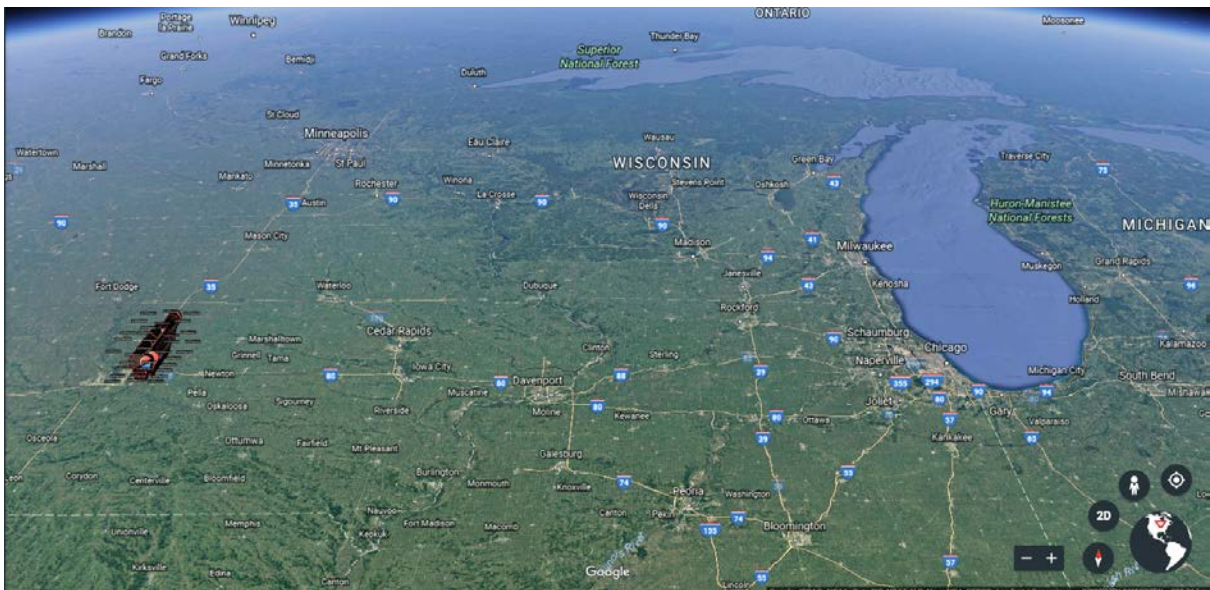


Figure 15 - Open Range Testing

A closeup view of the range testing reveals successful LoRaWAN packet transfers up to 50 km from the gateway.




Figure 16 - Open Range Testing2

The signal strength ranged from -68dBm to -125dBm at the extreme north end of the grid.

Conclusion

Overall, the test results illustrate that LPWAN wireless technology is very powerful for use cases requiring low-data-rates. This is supported by industry and engineering research that has been conducted along with actual deployment results around the world. LPWAN was created for a specific purpose, to communicate with and manage large numbers of sensors/devices in the field that report status infrequently, at low data volumes. This wireless technology was not created for low-latency applications that require constant communication with large amounts of data, like vehicle-to-vehicle applications.

Testing and real-life deployments indicate advantages in better penetration of walls and ground than cellular & protocols that use 2.4 Ghz ranges due to lower frequencies. LPWAN has a lower noise floor and link budgets than higher-power solutions due to physics of low power transmission versus noise levels with higher power solutions.

Because LPWAN is focused on low-data-rates, it has been optimized to enable sensors/devices to run on small batteries that may last for over 10 years. Battery lives measured in tens of years helps reduce operating cost related to battery replacement and equipment maintenance.

LPWAN has arrived at a time when radio, chip, and component costs have dropped to a level where BOM incremental cost is under \$10; while continuing to decrease. When quantities of sensors/devices needed for deployment are measured in thousands or millions, the economies of scale quickly make for compelling business cases in many industries.

Add to the equation that unlicensed spectrum enables solution providers to create more cost-effective products, while network operators have an economic advantage over operators of licensed spectrum (especially considering the cost to purchase spectrum in addition to operating cost).

The ecosystem for LPWAN is also very strong. SigFox and Ingenu have an impressive installed base around the world. The Things Network and other open solutions have grown via grassroots movements very rapidly over the past 2 years. Semtech as the creator of LoRaWAN and the LoRa Alliance has gone from 15 members 2 years ago to over 300 members which include IBM, Cisco, Orange, Comcast, and many other giants across multiple industries. The standards are in place and the industry adoption is arguably one to beat.

It is a compelling business model that takes advantage of many strategic advantages. LPWAN includes advantages in regulatory (unlicensed), cost (startup & operations), technology (range & penetration), customers (right use cases), and competition (strong ecosystem). While not perfect, and recognizing some of the advantages of other competing technologies; LPWAN has established a solid niche across many industries and will likely continue to be successful for years to come.

Abbreviations

AP	access point
bps	bits per second
Kbps	kilobits per second
Gbps	gigabits per second
FEC	forward error correction
GSM	global system for mobile communications
GSMA	GSM Association
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
GHz	gigahertz
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers
LPWAN	Low power wide area network
LoRa	Long range
IEEE	Institute of Electrical and Electronic Engineers
QOS	quality of service
RPMA	random phase multiple access

Bibliography & References

<http://www.cbronline.com/news/internet-of-things/smart-technology/industry-4-0-smart-investments-smart-factories>

http://www.semtech.com/wireless-rf/internet-of-things/what-is-lora/?utm_content=56831965&utm_medium=social&utm_source=linkedin

<http://www.techiexpert.com/2017/industry-4-0/>

<http://www.plattform-i40.de/>

<http://www.automation.com/automation-news/article/industry-40-only-one-tenth-of-germanys-high-tech-strategy>

<https://www.forbes.com/sites/louiscolumbus/2016/11/27/roundup-of-internet-of-things-forecasts-and-market-estimates-2016/#ff9fc4e292d5>

IoT for Peace of Mind

A Technical Paper Prepared for SCTE•ISBE by

Arun Ravisankar

Senior Engineer, Comcast Innovation Labs

Comcast Corporation

1701 JFK Blvd.

Philadelphia, Pa 19103

215-286-7558

Email: Arun_Ravisankar@comcast.com

Introduction

The telecommunications industry is hard at work to innovate, design and develop products that help consumers. What does a consumer look for when he/she plans to introduce a new product into his/her lifestyle? For example, the advent of smartphones changed the way mobile phones were used and perceived. What did these smartphones offer that influenced a paradigm shift in the industry? Mainly they led to a new era of computing. The influx of many applications helped solve some of the important, yet not necessarily critical, needs of consumers. The “app store” became a crucial platform where application developers and consumers could interact and share views. The main contributions of these technologies in consumers’ lives were ease of use, and, importantly, exposure to various tools that made life easier. One thing that the consumer looks for in any technology or product is “peace of mind.” Although the term “peace of mind” is decidedly subjective, it will be one of the major driving factors in product development, and is enormously applicable for the Internet of Things (IoT.) For example:

- * It would be immensely helpful for a consumer to know that his/her home is secure, while the family is out on a vacation. Design goal: That family members could remotely monitor the house and thus be less worried about security.

- * Beyond “just” security, the IoT can help making lives safe and efficient, with reduced anxiety. A busy mother would be relieved to know that she could check or turn off the stove remotely and need not worry about having left it on. Likewise, for that moment of fear about the garage door: Did I remember to shut it? And, from an efficiency point of view: Managing energy costs with smart bulbs and thermostats.

When IoT is applied to health and wellness applications, the peace of mind impact that the technology brings to the consumer is enormous. According to [Forbes](#), “it costs families more to care for a frail older adult than to raise a child for the first 17 years of life.” This is a growing concern. An [AARP publication](#) reported in late 2015 that the population of adults 85 and older in the U.S. will roughly triple between 2015 and 2060 – making it the fastest-growing age group over this time.

Apart from Eldercare, the IoT will help individuals to track personal health and well-being, with the use of various fitness-based devices. In examining the most prevalent home security and tele-health providers, one thing is paramount: Providing “peace of mind.” Figure 1 illustrates how this premise surfaces in brand marketing.

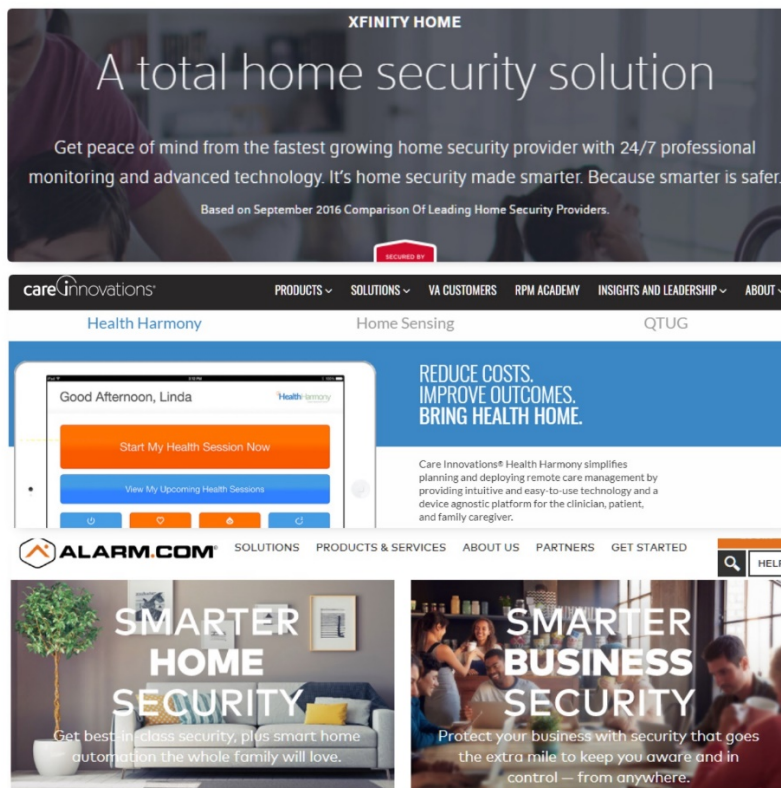


Figure 1 - How most IoT systems market their solutions

In this paper, we examine how IoT technologies can offer solutions to consumers to enhance lifestyles. This includes connectivity technologies, use cases and the perspective of a cable Multiple System Operators (MSO.)

IoT for Peace of Mind

1. Internet of Things

This term needs little in terms of elaborate definitions, at this point in time, a simple search shows that it can be defined as “interconnection, via the Internet of computing devices embedded in everyday objects, that enables them to send and receive data”. IoT can be applied to wide variety of applications, as characterized briefly here.

The number of connected devices is growing at a seemingly exponential pace -- and even modest estimates point to tens of billions of connected devices within the next 2 years. Any stroll across any home electronics store indicates that most of the new arrivals are connected, in some form, and have some amount of computing power inside --, be it refrigerators, washing machines, or vacuum cleaners, to name a scant few. The number of connected cars is increasing; autonomous cars are no longer science fiction.

The roster of IoT-related applications that relate to MSOs includes:

Home and Business Security Applications: To secure the premise, providing alerts based on rules established by consumers. That includes sensors that detect motion, for, doors, windows, and other ingress/egress locations, via cameras. Rules can be easily set and modified to generate multi-dimensional (multi-screen) alerts.

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 parallel 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.

Home Automation: To enable smart home applications, connect and control third-party smart devices, and incorporate the sensing equipment used for monitoring and automation.

Automotive: From connected cars to the autonomous vehicular future, the existence of voice control, heads-up displays, and self-parking systems are proliferating. Car-to-cloud connectivity is the enabler of vehicles that are safer -- which makes roadways safer, less congested, and more accessible.

Smart Energy: Through the "things" of the IoT, the power grid can share information in real time to distribute and better manage energy more efficiently. The IoT is already active in terms of Smart Cities, smart/connected communities, and long-range IoT networks.

Smart Manufacturing: IoT technologies can enable factories to unlock operational efficiencies, optimize production, and increase worker safety.

Retail: For retailers, the IoT offers substantial opportunities to increase supply chain efficiencies, develop new services, and reshape the customer experience.

2. Applications and Use Cases

The implications and intersections between the Internet of Things and peoples' lives is considerable, if not infinite. The scope ranges far wider than this paper can address. For that reason, this paper focuses on Home Automation/Security and Healthcare applications that have successfully used IoT in ways that provide peace of mind.

In the following sections, a brief case study is presented that combines the IoT with analytics and machine learning to solve for enhanced product/service performance.

3. Health and Wellness

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 the IoT could play a major role is in eldercare. Most studies show that there is significant growth in the Senior population in the coming years, as the Baby Boomer generation enters its twilight years. The IoT is poised to build applications to facilitate aging-in-place -- because personal independence is a high priority for Senior citizens.

The following images illustrate the trend in population growth, as well as technology adaptation and penetration by among age group.

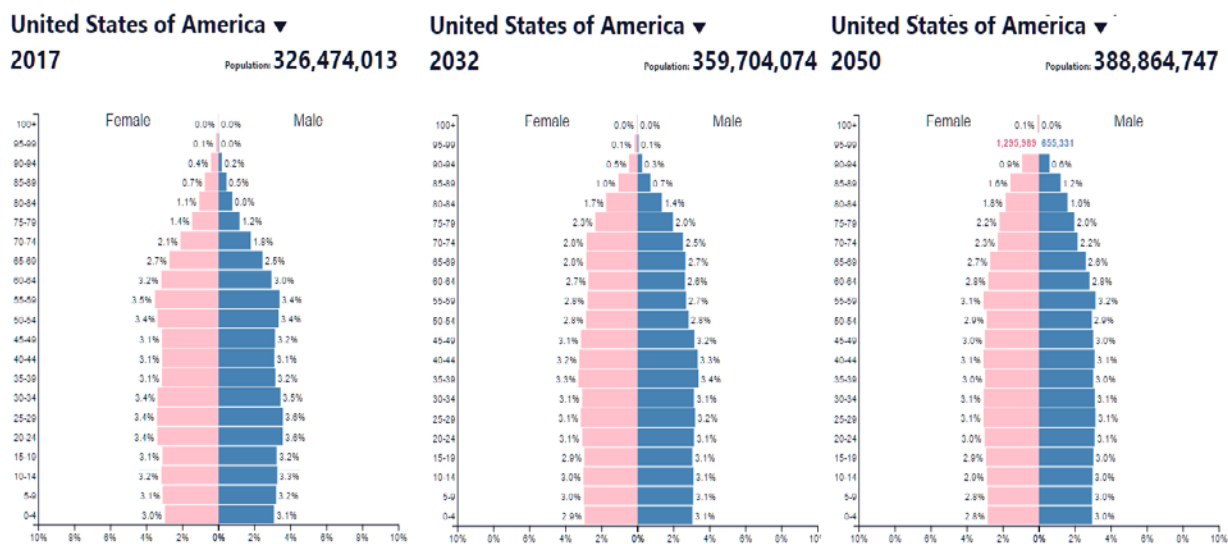


Figure 2 - Population Pyramid of USA in 2017, 2030 and 2050

Figure 2 shows overall trends in population growth. We can see that there is a constant increase in the Senior population, which is an indication that IoT and related technologies could play a major role.

Also, the technological advancements surfaced by industrial revolution have enabled most of us with Internet connectivity. It follows that technology adoption is rising among the Senior community. Soon enough (2050), “digital natives” (loosely defined as those born after 1985) will be in their twilight years.

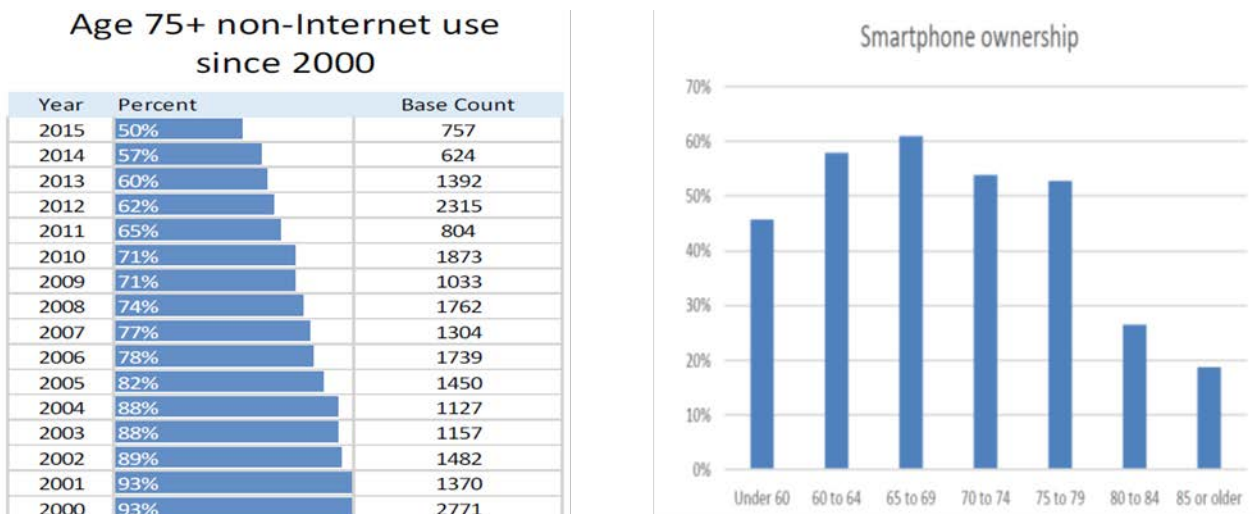


Figure 3 - Technology Adaptation among Seniors

Figure 3 shows a gradual reduction of non-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.

3.1. Why IoT for Healthcare/Eldercare?

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.
- 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 with chronic conditions.

3.2. Connected Health Applications

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 invaluable source of 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 incapacitation; fall detection or incapacitation could be used to trigger a PERS ([Personal Emergency Response Systems](#)) event

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.

Patient Adherence: Ensuring that patients adhere to doctor prescriptions; providing appropriate reminders to patients and family members, like reminders 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.

4. Home Monitoring and Security

IoT-based applications for home monitoring can be grouped into two broad buckets: Home Automation and Home Security.

4.1. Home Automation

Home automation or monitoring is loosely defined as local and remote access to in-home IoT devices that provide information about the home. Examples include:

- Motion detection (sensors or cameras) and monitoring equipment (possibly linked to local or remote recording)
- Sensors for changes in windows and doors
- Catastrophe sensors, for fires, floods, and weather-related danger,
- Sensors to indicate events that are of concern (out of normal ranges)
- Alerting end users .

With the known and predicted increase in the overall number of connected devices, there is a need to provide applications and services that utilize these resources and improve quality of life. These services also help manage energy efficiently and help proactively manage the safety of the premises.

Most operators provide smart home applications packaged with sensors and detectors, like motion sensors, door window sensors, door locks, fire sensors, flood sensors and others.

Most applications also provide hooks to add third party devices like door locks, thermostats. A typical smart home would look as shown below:



Figure 4 - A Typical Smart Home with Connected Devices

Most research indicates an increase in the volume of smart devices. The graphic, below, shows the ranking of devices that are most desired in a home monitoring application:

Most Desired Capabilities for Home Monitoring and Management System (Q2/12)

"Q7005. Which of the following capabilities would you most desire in a home monitoring and management system?"

The ability to automate, control and monitor...

(Among All BB HHs, n=2,517, ±1.95%)




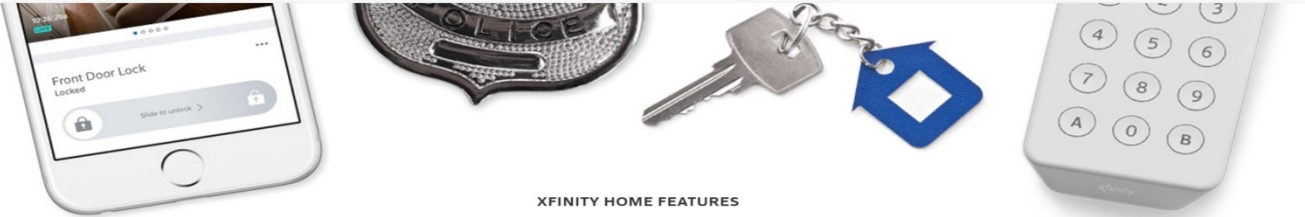
Figure 5 - Most desired capabilities in a smart home

4.2. Home Security

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 forgot to close doors or windows. A typical home security system, when set in “armed” mode, will detect any of such anomalies and alert the user. This is a huge relief for the home owner.

From a home security perspective, it is more important to know presence of a potential intruder, if there is a window breakage. Also important is redundancy in connectivity, like a backup connectivity option in case the main connectivity (usually cable Internet) is down.



XFINITY HOME FEATURES


Connected and protected

Protect your family and home and look after them from anywhere with an Internet connection using your smartphone, tablet or computer, with XFINITY Home.

SMART HOME CONTROL

Stay connected to your home from anywhere

In addition to remotely arming and disarming your system with the XFINITY Home app, you can also look after and control your home from anywhere. With additional equipment, get live video monitoring so you can see that your kids got home safely, control your thermostat and turn on lights so you never come home to a dark house again, and much more.



VIDEO MONITORING

Be in the know, even when you're away

See what's happening in and around your home in real time with live video monitoring, included in your monthly XFINITY Home service; you'll just need to add a camera to your system. For an additional \$9.95/mo. per camera*, rewind, review and share up to 10 days of footage in the cloud with 24/7 Video Recording.

*Taxes and fees extra. Pricing subject to change.

[Learn About Video Monitoring](#)

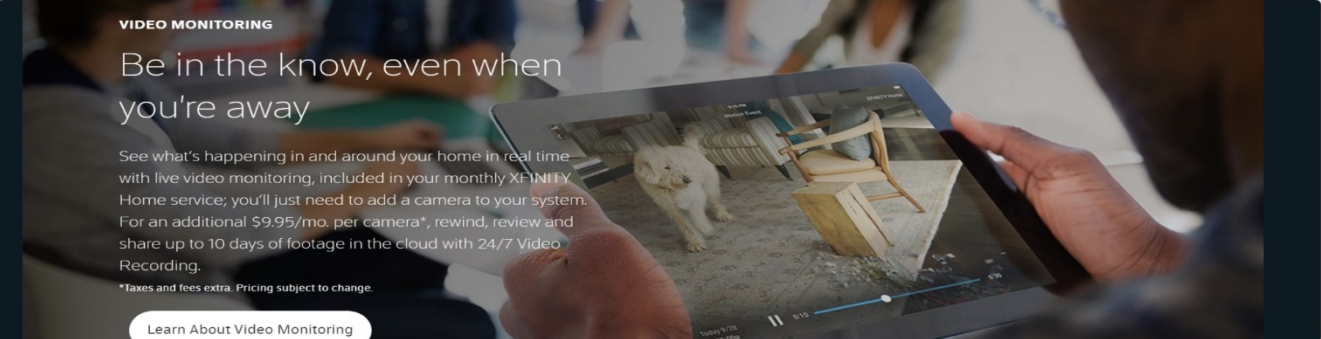


Figure 6 - Xfinity Home offerings for Home Security

5. Technology Overview

5.1. Network Topology

Most IoT solutions and applications discussed in the previous sections rely on a similar technological base, in terms of devices and network architecture. What differentiates the solutions are the end user applications and rules of operation. In this section, we describe a brief overview of the technology and network architecture. We will also look at security implications.

Like any other application stack, an IoT application also involves multiple layers implemented in hardware and software. The complexity and internal architecture of the hardware and software

components depends on the protocol being implemented by a device or gateway. In order to understand this better, we need a network topography of a typical IoT network:



Figure 7 - A typical IoT Network Showing Multiple Network Protocols

In any IoT network, multiple network protocols are involved, hence the need for a central gateway that can translate between protocols and forward packets for analysis. In the example above, the home gateway would be the IoT gateway, as it would have the necessary hardware and software to interact with devices which may use a different protocol.

For example, there may be a rule set to turn on a light bulb upon detecting motion in the hallway. In this case, the motion detector could be a [ZigBee-based device](#), and the light bulb could be based on [BLE \(Bluetooth Low Energy\)](#). In this case, the gateway would bridge the connection between the lightbulb and the motion sensor. The illustration, below, shows a sample sequence diagram of an IoT application:

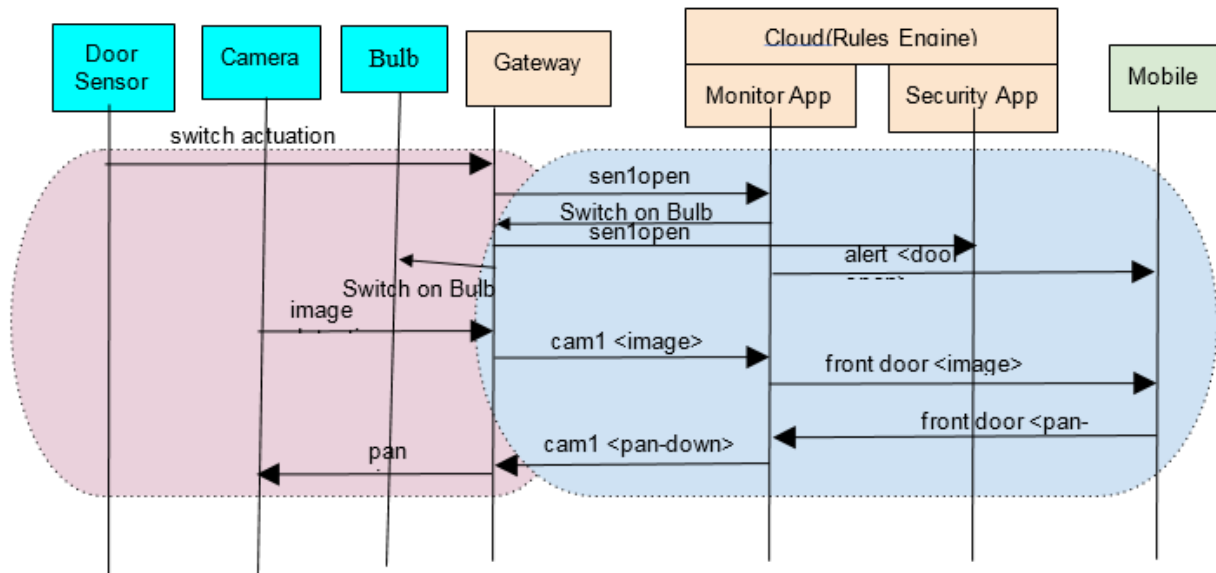


Figure 8 - Sample Sequence Diagram of IoT Use Case

As described above, the gateway needs to handle multiple network protocols. Table 1 shows the list of protocols and messages the gateway handles in the above sequence diagram:

Table 1 - Table Showing Network Protocols Handled by an IoT gateway

Gateway	
Camera	Wifi(TCP/UDP)
Bulb	BLE
Door Sensor	Zigbee

As seen in Table 1, the gateway handles messages from multiple devices, implemented using different networking protocols. The gateway implements these protocols, and the devices are paired with the gateway. Once paired, the devices (sensors) would function and exchange messages with gateway. The gateway would decode these messages and pass them to the cloud or rules engine for processing.

5.2. Technologies of great promise

In the previous section, we discussed about how a typical IoT network can interact to implement the desired use cases. In this section, we highlight technologies that show promise in enhancing the connectivity and performance of IoT applications.

5.3. RF sensing for IoT applications

Most IoT networks employ passive infrared/PIR-based motion sensors and cameras for surveillance applications. While these sensors work well, and provide the desired results, there are some challenges. These are line-of-sight sensors, for starters, and thus cannot detect motion through a wall. Cameras are used mostly for outdoor surveillance, and are often perceived as being intrusive for indoor tracking (activity monitoring).

These challenges could be easily overcome by using RADAR (Radio Detection and Ranging)-based sensing. Radar has been historically used in aviation and military applications, but holds a lot of promise for residential and commercial applications. Radars can detect presence and motion through walls, while providing good coverage, without cameras.

Microwave devices that are compact, accurate, reliable, and inexpensive are currently commercially available. Over the past few years, attempts to apply such devices to biomedical measurements has increased. Although some studies applied these devices to medicine and health care, such research is still in its infancy. Nonetheless, radio-frequency sensing techniques -- originally developed for military applications, and later applied to search and rescue operations locating earthquake survivors buried under rubble -- all carry plausible applicability for health care and home monitoring use cases, like aging in place and smart homes.

Doppler Radars could be used to implement motion classification, which could be used for applications like activity monitoring, fall detection and personal emergency response systems/PERS.

Traditionally, fall detection applications employ a wearable or push button device that needs to be activated after a fall occurs. The intent is to help the patient to trigger emergency assistance. Using radars for fall detection is non-intrusive and does not need require manual intervention (pushing the button) to trigger an alarm.

This matters because in many cases, people lose consciousness after falling, which obviates the applicability of such “push to enable” help calls. A radar, by contrast, monitors activities and can both auto-detect a fall, and raise an alarm. With machine learning capabilities, radar-based devices can also learn over time, to then more accurately detect falls.

Another form of activity monitoring enabled by radar is biometric, or the application of statistical analysis to biological data. Specifically, radar can provide effective, non-invasive and non-restrictive sensing techniques to acquire vital signs. For instance, radar could be used to monitor characteristics including body surface vibrations, mental state, and sleep apnea.

For instance, radars can detect minute vibrations on the body surface, such as those induced by heartbeat and respiration. Simple equipment can be used to self-monitor certain medical parameters or conditions, as well as to acquire related data required for Senior living homes as well as medical facilities.

The diagram below shows a sample application of radar installed in a room. The radar detects the presence, position and posture of people. This would be helpful to build an activity monitoring application



Figure 9 - Sample of a Radar based application used for activity monitoring

Radars have already been used in cars to provide features like collision warning and prevention systems, parking assistance and blind spot warnings. Extensive research has been done to ensure human safety while using radar systems, and have been proven to be safe.

Radars are thus a promising sensing technology that IoT systems could use to enhance range and performance.

5.4. Sensors to Insights

Most, if not all, sensors generate data of some form. In order to build an application that can learn using the data from the sensor, we need to add data analytics and deep learning to the IoT system. With this the system will learn over time, and will help improve performance and reactions to events. By incorporating machine learning and data analytics skills, the system could know the local environment better, which helps in building a sophisticated and customized application.

Consider, as an example, the use of a sensor-based IoT application, deployed to monitor activity in a Senior care facility. If we implement deep learning, using the data provided by the sensors, then the system would better learn more about the individual being monitored, which in turn would help build a better ADL monitoring application, customized to the individual. The system can learn that the patient needs to take glucometer test, or a medication, at 8 AM, every day, and can issue alerts if data is not available -- indicating that the test or medication wasn't taken. Apart from that, the system can also check for any drastic change in readings and could initiate appropriate actions.

There three important steps in deep learning that should be part of any deep learning system, as shown in the diagram below:




Figure 10 - Important steps in deep learning

Data analytics of critical health data comes with challenges, particularly related to how data security is handled. Most of these requirements are regulated by [HIPAA](#) compliance. In particular, secure data storage and analytics engines need to be implemented.

5.5. Using Video and Audio Analytics

Machine learning algorithms and AV analytics also hold promise in implementing monitoring solutions. With the advent of far-field microphone devices, there are methods available to train a system to look for specific sounds and trigger alerts. This could be used in a variety of applications, like home monitoring and smart cities.

Video analytics, on the other hand, uses a deep learning system that learns objects and people as seen from a video camera. Although using video is sometimes perceived as being intrusive, it could be used in certain applications to give a better description of a scene, and to help construct an IoT-based application.

5.6. Long Range IoT Networks

Most IoT services and applications currently rely on a broadband network (cable or wireless) to backhaul data. This setup would need a gateway to pass the data along to the cloud for processing (See example in [Section 5.1](#))

However, long range IoT networks like [LTE-M](#), [LoRa](#) or LoRaWAN, provide IoT environments in which the devices can directly connect to the Wide Area Network, without a gateway. Devices could be battery operated and could connect to a network that could span a wide geographical area. The following are some of the Long Range IoT networks that help enhance connectivity and improve IoT services:

5.6.1. LoRa

[LoRaWAN](#)[™] is a Low Power Wide Area Network (LPWAN) specification intended for wireless battery operated things in a regional, national or global network. Communication between end-devices and gateways is spread out on different frequency channels and at different data rates. The selection of the

data rate is a trade-off between communication range and message duration. The license-free, sub-GHz frequencies used in North America (915 MHz band) are as follows:

- Upstream: 64 channels numbered 0 to 63, DR0 to DR3
- Upstream: 8 channels numbered 64 to 71, DR4
- Downstream: 8 channels numbered 0 to 7, DR8 to DR13

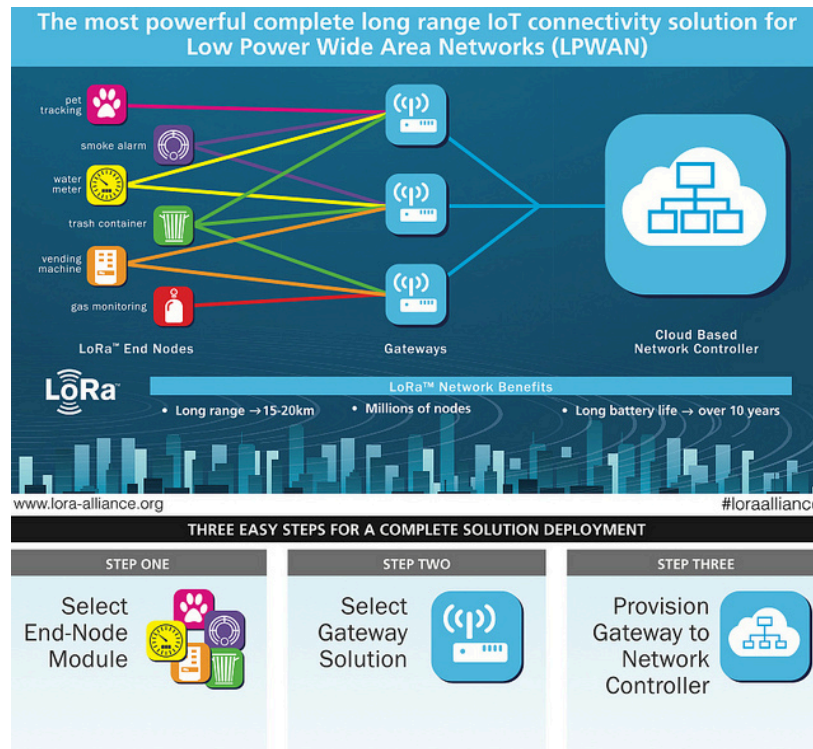


Figure 11 - LoRa Network diagram

5.6.2. LTE-M

LTE-M is the abbreviation for LTE Cat-M1, or Long Term Evolution (4G), category M1. This technology was designed for battery-powered Internet of Things devices to connect directly to a mobile 4G network, without a gateway.

LTE-M is a low power wide area technology which supports IoT with lower device complexity and extended coverage, while allowing the reuse of the LTE-installed base. This allows battery lifetime as long as 10 years or more for a wide range of use cases, with the modem costs reduced to 20-25% of the current EGPRS (Enhanced Global Packet Radio Service) modems.

Some of the advantages that long range networks provide include:

- **Cost efficiency:** Devices can connect to networks with chips that are less expensive to make, because they are half-duplex and have a narrower bandwidth.
- **Long Battery Life:** Devices can enter a "deep sleep" mode called Power Savings Mode (PSM) or wake up only periodically while connected.

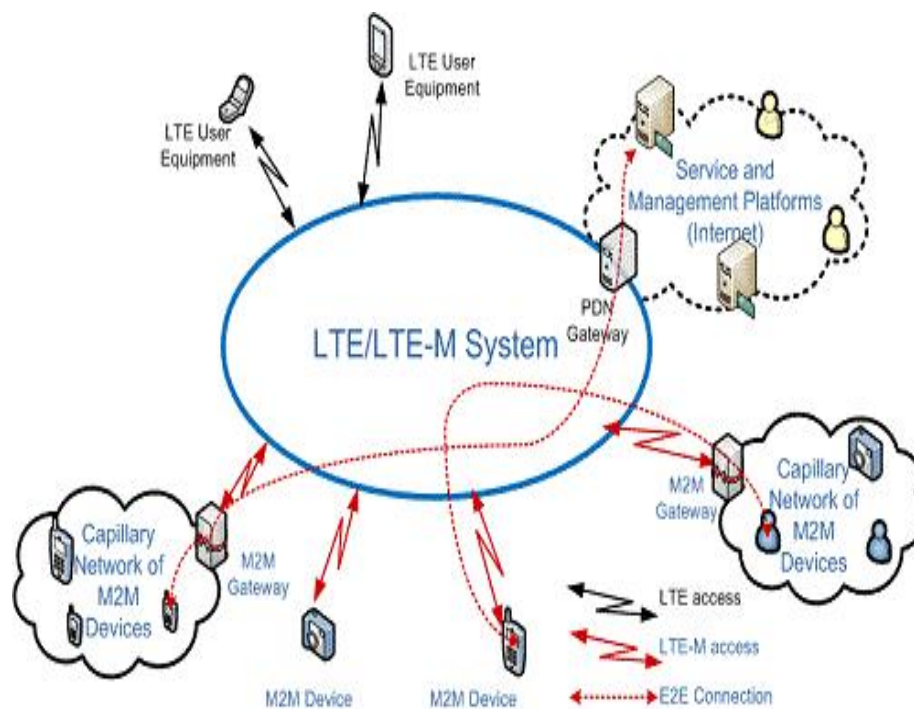


Figure 12 - LTE-M System diagram

Conclusion

Multiple potential use cases and technological viability indicate that IoT technologies can definitively provide *peace of mind* to consumers. There can be innumerable examples that can address peace of mind factor in a consumer's life -- whether it involves tracking a child, to ensure that he/she left school and reached home, or being able to check on elderly parents to ensure that they're doing ok, and, most importantly, to be alerted when they need help.

We examined the technical overview IoT networks, and how the IoT is solving current problems, including how operators could incorporate other newer technologies to improve the system.

We explained how devices using different protocols can be implemented in a way that facilitates suitable communication, and reviewed how machine learning and data analytics could help to more accurately predict any anomalies in the system, identify corrective actions.

The role of machine learning and data analytics is key to the longer term enhancement of IoT services, as is the continued work to adapt known technologies, like radar, to add value to IoT based services.

Acronyms and Abbreviations

IoT	Internet of Things
HIPAA	Health Insurance Portability and Accountability Act
LoRaWAN	Long Range Wide Area Network
LTE-M	Long Term Evolution (4G), category M1
RADAR	Radio Detection and Ranging
IP	Internet Protocol
STB	Set Top Box
6LowPAN	Internet Protocol (IPv6) and Low-power Wireless Personal Area Networks (LoWPAN)

Deploying and Optimizing the Next Generation Wireless Home

A Technical Paper prepared for SCTE•ISBE by

Steven R. Harris

Senior Director Technical Field and Engineering Education

SCTE / ISBE

140 Philips Road

Exton, PA 19341-1318

610-594-7324

sharris@scte.org

Introduction

Delivering a seamless managed wireless and Wi-Fi experience for a residence or commercial subscriber is quickly becoming the de-facto standard when judging an MSO's services. The MSO's are shifting their attention to the quality of experience (QoE), proactive (e.g. PNM for Wi-Fi) wireless ecosystems and carrier grade versions of Wi-Fi for multimedia services like ultra-high definition (UHD), high definition (HD), audio and voice over Internet protocol (VoIP). The next generation wireless and Wi-Fi are being used to provide whole home coverage for cable subscribers. The wireless networks must be seamless, reliable and optimally designed to support multimedia services, faster data services (e.g., gigabit) and whole home connectivity.

When it comes to in-home wireless readiness, making sure that the customer's data, video, voice services along with low power long-range wireless Internet of things (IoT) (e.g., 802.15.4, Sigfox, LoRa wide area network, etc.) will run smoothly long after the technician has left the premises. Operators must also leverage the latest Wi-Fi amendments of 802.11x, address spectrum concerns and solve challenges with in-home wireless interoperability. In addition, the development of operational practices to not only identify the QoE throughout the home, but also characterize ways through analysis to improve the user experience when expectations are not met.




Figure 1 - Wireless at the Premises

A big part of the next generation home is wireless, in particular Wi-Fi. IEEE 802.11ac second wave access points or gateway routers (GWRs) are expected to dominate the global wireless local area network (LAN) market by 2018, with more than 80%¹ units shipped.




Figure 2 - 802.11ac GWRs Expected to Dominate the Global WLAN Market

Subscribers want a consistent wireless experience while operators are shifting to a managed carrier grade wireless experience. Carrier grade is not defined by the maximum throughput of a technology but by the minimum. Carrier grade key metric: supports up to four (4) simultaneous HD video programs, just as good as a quadrature amplitude modulation (QAM) channel over a hybrid fiber coax (HFC). This requires 40 Mbps minimum at a low packet error rate (PER) of 10^{-6} or better, 90% coverage or better of the living space on both 2.4 GHz and 5 GHz². Multiple-input multiple-output (MIMO) and wider channels (40 MHz for 802.11n and 80 MHz for 802.11ac) allows a GWR to achieve higher rates for carrier grade metrics. Other technology that contributes to a delivery of a carrier grade experience is the network discovery and selection process; authentication and encryption methods used; and roaming features supported by a GWR. The Wireless Broadband Alliance (WBA) developed solutions and guidelines to improve the QoE for carrier grade Wi-Fi networks³.

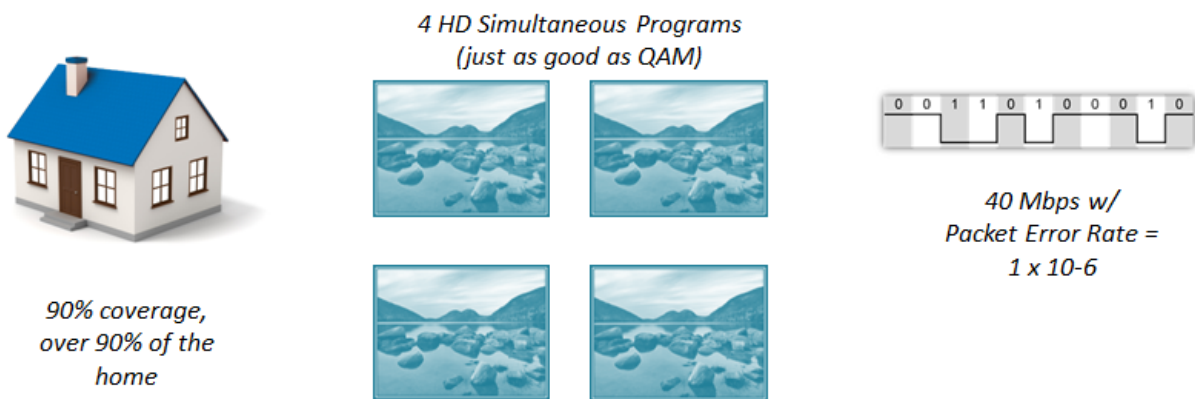


Figure 3 - Carrier Grade Wi-Fi

Managing the carrier grade Wi-Fi experience is equally as important in wireless communications. Cable operators will and must truly own the wireless ecosystem using managed devices (e.g., WGR) within the home. Coupling this with proactive network maintenance (PNM) for Wi-Fi (a working group within SCTE-ISBE), allows the operator to mitigate many of the problems associated with management of a wireless network. In addition, vendor incompatibles will degrade the QoE, driving the need for technologies like reference design kit for broadband (RDK-B). RDK-B is an open source software that

allows MSOs to accelerate innovation in video and broadband networks⁴. The next generation subscriber will require the management of not just Wi-Fi but multiple low powered long-range wireless standards (e.g., ZigBee, Z-Wave, 802.11.15.4, Bluetooth Low Energy, LoRaWAN) requiring common access hubs or WGRs. Improved data collection is part of the service allowing the MSO to collect information via technical report (TR) 69, simple network management protocol (SNMP), Internet protocol detail records (IPDR), etc. Finally, the smartphone will be used as the central control point and gateway for these wireless devices.

802.11 Features

For an operator to provide the next generation managed carrier grade QoE, the features of the IEEE 802.11 amendments must be leveraged and deployment to the subscriber. The newer 802.11 standards and amendments improve RF multiplexing and modulation; support wider channels, more spectrum access, efficient media access control (MAC) layer communication and more spatial streams (multipath).

The current trend for the subscriber's home is to offer a gigabit wireless service to support a gigabit Internet modem service. Technologies like high throughput (HT) 802.11n will need to be replaced with newer WGR versions using 802.11ac, as most 802.11n devices offered 100 Mbps at best and few 802.11n devices offered the 600 Mbps as the maximum achievable of the amendment. 802.11ac offers a very high throughput (VHT) service, supports a roadmap to gigabit speeds of 7 Gbps and builds on 802.11n/g MIMO technology. However, most 802.11ac wave 1 WGRs come close to 1 gigabit while wave 2 WGRs will perform in the 2 Gbps range⁵. Operators are under pressure today to offer the newer waves of the 802.11ac amendment supporting 4x4 multi-user MIMO (MU-MIMO) and future 8x8 MU-MIMO.

802.11ac will allow dramatically faster broadband speeds, because of the physical changes in spectrum, multiplexing, modulation, channel sizes, spatial streams and efficient MAC. Additional spectrum may be leverage in 802.11ac over 802.11n, further opening the 5 GHz band, up to eight times the spectrum of the 2.4 GHz band⁶. Even though the 802.11ac amendment supports larger 160 MHz channels, it may be difficult to bond channels to support these size channels due to dynamic frequency selection (DFS) for radar usage. DFS limits a 5 GHz channel reuse, up to 16 channels may be used by radar systems. Other technologies like adaptive antenna technology, MU-MIMO, polarization features and channel optimization increase the potential for 802.11ac.

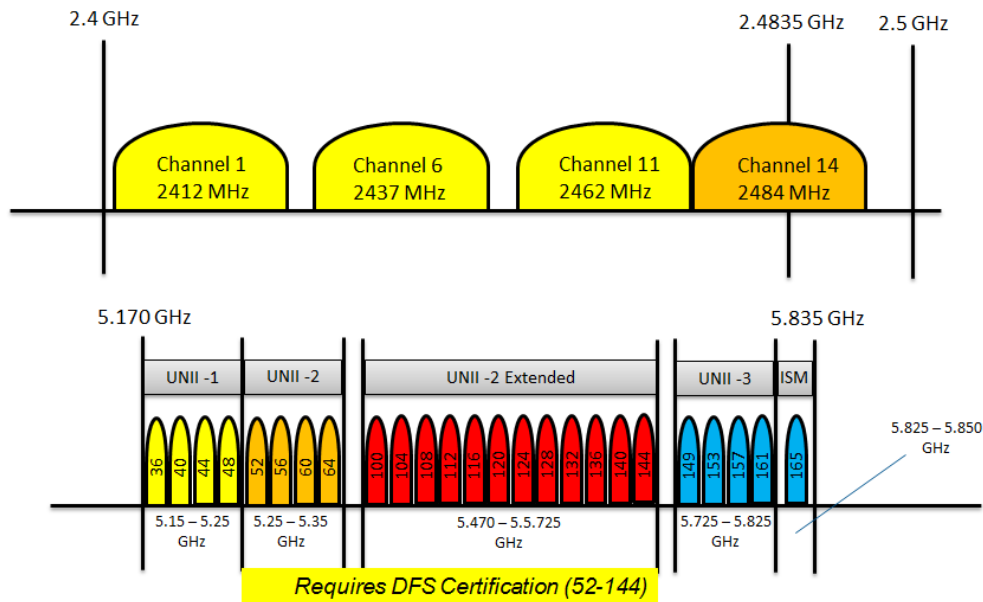


Figure 4 - 2.4 GHz and 5 GHz Channels (USA Only)

1. Multiple-In Multiple-Out (MIMO)

Multiple-in multiple-out (MIMO) radios were improved with 802.11n to provide better communication than earlier single-input single-output (SISO) communication. MIMO uses multipath RF propagation to increase the reliability and data rate of a Wi-Fi network. MIMO reliably allows multiple streams of data to be delivered to a subscriber's 802.11n customer premises equipment (CPE) using adaptive antenna diversity and multipath. In this example, the CPE must support more than a single antenna, a common issue for CPE. In addition, multipath might also create reflected streams, causing fading of a wireless signal where MIMO capability does not exist. The MIMO data rate function uses a technique known as spatial division multiplexing (SDM). In SDM, the data at the WGR is split into a number of spatial streams and transmitted through separate antennas to corresponding antennas at the receiver. The key here, is "corresponding antennas", most cable subscriber's 802.11n or earlier CPE did not support MIMO SDM.

In 802.11n MIMO, having three (3) spatial streams supports a data rate of 450 Mbps, enabling a far greater utilization of the available bandwidth if the CPE supports the feature. Since most operators want a balance of reliability with throughput achieving (3x3) or 450 Mbps is not possible using 802.11n as most customer owned devices offer around a 100 Mbps service. The current 802.11n standard allows for up to four spatial streams (4x4) or 600 Mbps, while 802.11ac allows eight spatial streams (8x) or 7000 Mbps.




Figure 5 - 802.11n or 802.11ac 4x4 MIMO Example

802.11n's MIMO capabilities were designed for a single user, or SU-MIMO. While 802.11ac's MIMO is designed for multi-user or MU-MIMO to alleviate congestion delays. MU-MIMO allows more simultaneous bandwidth, instead of broadcasting the signal equally in all directions, the base station is only sending RF in the client's direction, the downstream direction only. However, just like in 802.11n clients and WGRs they need to support MU-MIMO. Initial waves of the 802.11ac amendment will support a 4x4 MU-MIMO, while 8x8 MU-MIMO technology may have a prohibitive cost model associated, making 4x4 more attractive.

Leverage technologies like MU-MIMO may add improved data rates for the subscriber. The modulation and coding scheme (MCS) index of a WGR is also an important factor for overall data throughput. The higher the MCS index the higher the potential data rate as shown in the table. A single MCS index applies to multiple guard intervals (GI) of 400 ns and 800 ns; and channel sizes as shown in the table. Just like the MU-MIMO feature both the client and WGR must support the same level of MCS, channel size and GI.

Table 1 - Theoretical Mbps Data Rate Sample of 802.11n and 802.11ac

IEEE	MCS Index	Spatial Streams	Modulation	Coding Rate	20 MHz 800 ns GI	20 MHz 400 ns GI	40 MHz 800 ns GI	40 MHz 400 ns GI	80 MHz 800 ns GI	80 MHz 400 ns GI	160 MHz 800 ns GI	160 MHz 400 ns GI
802.11n	7	1	64 QAM	5/6	65	72.2	135	150	n/a	n/a	n/a	n/a
802.11ac	7	1	64 QAM	5/6	65	72.2	135	150	292.5	325	585	600
802.11n	15	2	64 QAM	5/6	130	144.4	270	300	n/a	n/a	n/a	n/a
802.11ac	15	2	64 QAM	5/6	130	144.4	270	300	585	650	1170	1300
802.11n	23	3	64 QAM	5/6	195	216.7	405	450	n/a	n/a	n/a	n/a
802.11ac	23	3	64 QAM	5/6	195	216.7	405	450	877.5	975	1755	1950
802.11n	31	4	64 QAM	5/6	260	288.9	540	600	n/a	n/a	n/a	n/a
802.11ac	31	4	64 QAM	5/6	260	288.9	540	600	1170	1300	2340	2600

2. Antenna Technology

Antenna technology provides three things to a WGR radio: gain (dBi, dBd), pattern and polarization.

When working with wireless equipment it is important to understand gain. Relative antenna gain is determined by two different scales. The first scale, dBi, describes the gain of an antenna relative to a perfect antenna that radiates in a 360° pattern called the theoretical isotropic radiator. The second scale, dBd, is antenna gain relative to a dipole antenna. Most antenna gain is measured in dBi, on occasion a dBd antenna gain may appear, just add +2.14 dB to the dBd value to determine the corresponding dBi value.

Antenna energy patterns may vary between WGRs. All antennas are referenced to an isotropic radiator or sphere, antennas are used to shape the RF energy pattern. There is no overall best or worst antenna. Each antenna type is designed to produce a different shape or energy pattern. There are three main types: omnidirectional, semidirectional and highly directional. Omnidirectional are the most common at a residential or small commercial location, operating similar to how light is radiated from a lamp. While the semidirectional is like a wall sconce light and the highly directional is like a spot light. The latter two types of antenna are more common in a CableWiFi™ or access network deployment.




Figure 6 - Isotropic Radiator with Directional and Omnidirectional Patterns

Adaptive antenna technology or beamforming allows wireless RF energy to be shaped and directed in different ways for clients within a home or business to gain better capacity and data throughput. Adaptive antenna technology controls multipath effects (diversity) of wireless and is well suited for high-density venues worldwide. Adaptive antenna technology is also effective on strand mounted access points (APs) for CableWiFi solutions. Adaptive antenna technology may increase the signal gain at the user's device by as much as +6 dB, while improving S/N with reduced noise. The signal gains depend on the number of antennas and the algorithm being used by the WGR. Operators who deploy this technology may be able to improve the wireless QoE for the subscriber.




Figure 7 - Adaptive Antenna Technology or Beamforming

Another consideration when it comes to antenna technology is polarization. Antennas support horizontal and vertical polarization or orientation of the RF waves which is extremely important to successful wireless communications. Antennas are to be aligned with the same polarization, it is irrelevant if horizontal and vertical polarization is used⁷. It must be noted that some clients may not support both types of polarization.




Figure 8 - Horizontal and Vertical Polarization

3. 802.11x

802.11, the Wi-Fi standard, was developed by the Institute of Electrical and Electronics Engineers (IEEE). The 802.11 designation is a set of wireless standards developed for the wireless local area network (WLAN). Many of the popular 802.11x standards are utilized at the customer premises (e.g., 802.11g, 802.11n and 802.11ac) while new emerging standards (e.g., 802.11ad, 802.11af, 802.11ah and 802.11ax) will start making their way into the home and business.

The HT 802.11n amendment was released in 2009 allowing wireless communication over the ISM 2.4 GHz and U-NII 5 GHz bands. The amendment increased data rates up to 600 Mbps using high throughput OFDM (HT-OFDM), MIMO spatial streams and channel bonding up to 40 MHz channels. The amendment is backward compatible with all 802.11 standards, but usually not 802.11 prime due to the speed limits of the amendment. 802.11n, as many of the previous versions of 802.11x Wi-Fi are no longer used or should be used to provide wireless services for gigabit service rates.

The VHT 802.11ac amendment was released in 2013 increasing data rates to 1,300 Mbps or 1.3 Gbps (wave 1 release), 3.5 Gbps (wave 2 release) and up to 7,000 Mbps or 7 Gbps (in future wave releases) – using strictly the U-NII 5 GHz bands. However, there are dual band options available supporting the legacy 2.4 GHz spectrum for subscribers. As mentioned to increase the data rate MU-MIMO was introduced along with wider bonded channels. Increased channel sizes of 80 MHz and 160 MHz allow for an increase in the data rates, however 160 MHz channels will be a challenge to implement.




Figure 9 - 802.11 Amendments

In both 802.11n and 802.11ac, channel bonding is a possibility. Wider channels provide more data throughput in a wireless network. 802.11n supports 20 MHz and 40 MHz channels while 802.11ac supports 20 MHz, 40 MHz, 80 MHz and 160 MHz channels. Note, achieving 160 MHz wide channels over an 802.11ac network is not likely possible due to the limited amount of spectrum available to the WGR in the 5 GHz band.

The 802.11ad amendment was released in 2012, increasing data rates to 7 Gbps over the 60 GHz band, however over a short range. The amendment also supports legacy 2.4 GHz and 5 GHz bands using a fast session transfer feature. One idea for a wireless home is to use 802.11ad in a room (e.g., living room) while 802.11ac is the backbone between rooms.

A newer amendment 802.11ax or high efficiency WLAN (HEW) will be released in 2019 offering four (4) times the throughput of 802.11ac (up to 28 Gbps in the specification), using new multiplexing and modulation. The multiplexing is similar DOCSIS 3.1's upstream, called OFDMA. While the modulation moves from 256 QAM to 1024 QAM in 802.11ax.

The 2014 amendment, 802.11af, White-Fi or Super Wi-Fi allows wireless in the newly opened TV whitespace frequencies of 470 MHz to 710 MHz (Europe) and 54 MHz to 698 MHz (USA). Operators need to keep in mind, low bandwidth frequencies will produce lower data rates of 26.7 Mbps or 35.6 Mbps⁷ depending on the width of the channel. To achieve higher data rates channel bonding will be a requirement for this amendment.

Similar to 802.11af where frequencies below 1 GHz are utilized, 802.11ah or HaLow will offer low power wireless for IoT or M2M. The new physical (PHY) layer that offers the low power will mean lower data rates, however much longer distances for sensor style networks.

To improve the above major amendments additional amendments were created to enhance wireless communications. To improve RF data collection and sharing (AP discovery), 802.11k provides radio resource management (RRM) to improve network performance (e.g., RSSI, utilization). 802.11h provides added value for 802.11k called transmit power control, or TPC, for the frequency bands. While 802.11r provides air interface fast basic service set (BSS) transition methodology that enables a device to have fast secure roaming. To allow improved secured roaming of a device the 802.11u amendment provides interworking with external networks, the basis for HotSpot 2.0. 802.11v provides wireless network management enhancements for network management (client config). Lastly 802.11e or Wi-Fi multimedia (WMM) defines procedures for packet classification and prioritization.

Table 2 - 802.11 Amendment Summary

Amendment	Feature
802.11n	Current version of Wi-Fi supporting up to 600 Mbps
802.11ac	Current version of Wi-Fi supporting up to 7 Gbps
802.11ax	Four (4) times the throughput of 802.11ac (up to 28 Gbps in the specification)
802.11ah	Low power wireless for IoT or M2M
802.11af	Wi-Fi over frequencies of 470 MHz to 710 MHz (Europe) and 54 MHz to 698 MHz (USA)
802.11k	RRM to improve network performance
802.11h	TPC for the frequency bands
802.11r	BSS transition, enabling fast secure roaming
802.11u	Interworking with external networks
802.11e	Packet classification and prioritization
802.11v	Network management and client configuration

Challenges with In-home Wireless

1. Installation Concerns

A few of the top issues for MSOs with in-home wireless are range (requiring additional devices in the home), throughput/speed (packets/sec), interference, congestion (too many users/connections), lost credentials, configuration issues and wireless device age/compatibility. In addition, a workforce that fully understands the technology to deploy it right the first time is a major barrier as few field operation folks are certified in wireless⁸.

Installation focus areas for MSOs are field operation staff education, customer education, pre-work, post-work, optimal RF propagation, using benchmarked devices, using operator managed devices (PNM for Wi-Fi and RDK-B⁴), following best practices for installation, easy customer installation and proper security configuration.




Figure 10 - No DOCSIS, NO Wi-Fi

Standardizing the installation focuses on several key areas that are important to a best in class subscriber QoE with a wireless network. Field operation staff must be prepared and trained to support wireless services at the premises. In addition to field operation training, customer education and the expectations on how the wireless network will operate as compare to the needs of the subscriber must be addressed.

As with any cable service installation, pre-work steps must be performed. Pre-work consists of qualifying the drop connection since wireless is a service that utilizes the HFC or optical distribution network (ODN) to send and receive data from the Internet or cable network. In addition to the drop, verification of the modem service is operating correctly as the wireless network will connect to the modem. Performing a wireless site survey to understand the operating environment and talking with the subscriber about the wireless coverage in the BSS area is a key step to providing a good wireless service. Identify items that may reduce RF coverage, range or performance for the wireless network and discuss with the subscriber. Site surveys are excellent troubleshooting tools, as they provide visibility into the network. Many survey tools are available for smart phones, relatively easy to deploy. The site survey results are used to determine optimal placement of the wireless WGR. Finally, make sure the CPE is matched to the wireless equipment being installed by the operator. Customer education is required in situations where older or non-compatible CPE exists at the premises.



Figure 11 - Survey Maps of 5 GHz

Operators should also perform post-work, such as an exit survey to verify RF coverage, interference and wireless channels are utilized correctly. All wireless CPE must have its connectivity verified with the customer before closing the job. Each of the installed WGRs will require continuous measurement and management by MSOs, on the same lines as how the industry monitors modem health. Tools like PNM for Wi-Fi and/or RDK-B are vital to the wireless QoE.

As previously mentioned, one of the device installation challenges include lack of device knowledge; how the CPE operates and what is the best way to troubleshoot CPE. Having the knowledge to perform manual input on wireless devices and the use of autoconfiguration scripts are equally important. These knowledge gaps may be due to inconsistent lack or poor tracking of field operation training - implementing formal standards and tracking for refresher training as new technologies are presented is recommended. For example, using online training programs such as SCTE-ISBE VirtuLearn, SCTE-ISBE LiveLearning and SCTE-ISBE Chapter Webinars to deliver training updates are excellent resources to keep current⁸.

In addition to training, many operators are moving to deploying benchmarked devices, these are CPE that are tested and verified before they are deployed to the field. Understanding the features and limits of wireless devices will help ensure QoE. Benchmarking allows an MSO to understand when to use one wireless device over another device at the premises. Along with benchmarking, using consistent tools to verify wireless health is important to a consistent experience.

Define a customer-owned device support strategy, target devices with highest penetration. Educate field operation staff to support, troubleshoot and statically configure the highest penetrated devices. Always verify all devices that need to be connected to the wireless network.

Security configuration includes configuring Wi-Fi protected access (WPA)/WPA2, service set identifier (SSID), enabling Radius and disabling MAC-address filters as customer will want to change devices on the wireless network from time to time. WPA2/AES is the greatest security and preferred by operators. It is important for an operator to have an administration login to perform advanced level configuration and troubleshooting, even remote access by CSRs. Security may also create issues with lost credentials, a portal or application may alleviate subscriber frustrations here by providing methods to easily apply or reset security options.

2. Range and Coverage

Strive to place the WGR centrally at the premises, both vertically and laterally, and ensure that the antennas are directed so that there are no nulls (end of antennae) pointing at key coverage areas. Another consideration is switching the WGR to one with additional antennas to support new features like MU-MIMO and adaptive antenna technology. A WGR with MU-MIMO capability will improve the wireless range and coverage for a subscriber. Operators may add additional WGRs (e.g., MoCA to Wi-Fi devices) to cover remote parts of the premises, e.g. basements, outdoor suites/garages, etc. Finally ensure that there are no key blockages near the WGRs like chimneys, refrigerators, metal objects, or similar.

An important aspect to coverage is understanding wireless device's capability for range to provide the appropriate CPE for the MSO's data speed tier. In other words, will the device work at the location, size of home, number of devices, type of devices or offer features expected by the customer? One way to mitigate a poor installation in this scenario would be to have each wireless device brand benchmarked. Real-world wireless performance of MSO CPE must be tested and understood. The site survey will aid in determining the best way to provide wireless signals, boost wireless performance or selecting an alternative device such as an enterprise solution.

The WGR needs the capability to perform continuous measurement and mitigate immediate issues to optimize range and coverage, ongoing after the install is complete. Many of the newer WGRs include dynamic channel allocation (DCA), a feature that optimizes the capacity in entire network based on effect of interference. TPC changes operating channel if the RSSI can be improved by a sensitivity value (e.g., 5 dB). While other WGRs (with 802.11k/v) deal with the behavior of sticky clients that will not roam, clients that associate with a GWR and hang on rather than moving to a nearby WGR.

In addition, the latest WGRs have features to reduce co-channel interference (CCI), dealing with issues where one channel change can cause other WGRs in an MDU scenario or neighborhood to also change channel (RF channel ripple effect). Non-Wi-Fi interference mitigation detection systems should classify non-Wi-Fi interference as either persistent or spontaneous, blocking persistent channels. Note that many wireless clients typically do not choose the best channel or serving WGR without some help by the MSO or software algorithms.

3. Throughput / Speed

Especially for older wireless devices at the premises, data throughput and speed are a major concern of our subscribers. Many of the older devices have a varying level of support with latest 802.11x amendments. As mentioned, range was also a factor when it comes to customer throughput and speed, ensuring proper location of the WGR and understanding building construction will improve range/coverage related speed issues.

CPE such as a wired router between the AP and modem must also be considered when looking to improve performance for a subscriber. Wired routers must support the desired packet per second (PPS) rate of the operator's speed tier selected by the subscriber. Wired router configuration issues may rob speed (e.g., logging or firewall settings) too. To ensure a level of QoE and to improve the wireless capacity, install wired twisted pair connections to stationary devices like television displays, desktop computers, IP security cameras, etc.

Another area that affects throughput and speed is the MCS index supported by both the WGR and client. There are MCS index values for 20 MHz and 40 MHz wide communication over 802.11n while 802.11ac extends these index values for 80 MHz and 160 MHz channels. The wider and higher MCS index values support greater throughput and speeds as higher channels will use better orders of QAM (e.g., 256).

Also, collisions and/or packet back-offs due to non-persistent CSMA cause large throughput loss in wireless networks. Finally, congestion may occur where too many users or devices exist in the wireless network.

4. Interference

Understanding what affects wireless signals is important, and many lack a full understanding of signal characteristics that affect RF attenuation. There are many possible symptoms when interference is encountered by a technician in the home or business. A common concern is low signal power levels coming from the WGR. These lower powered signals will not be able to propagate when interferers exist in proximity to the network. Interferers will cause wireless signal dropping, even when close to a WGR, creating speeds much slower than Ethernet or MSO modem tiers. Subscribers consuming multimedia content will experience freeze frame viewing of video or stop/start of audio.

To install an unlicensed wireless network, it is important to determine the intentional and unintentional interferers at the customer premises by performing a site survey of the RF spectrum at the location. Since wireless, in particular Wi-Fi, is an unregulated spectrum with many devices now sharing this RF space it is a recommend step for installation. Intentional interferers are devices designed to send and receive RF signals such as laptops, Bluetooth devices, game controllers, untethered cameras, cordless phones, baby monitors, etc. In addition, some subscribers may have additional Wi-Fi devices on overlapping channels within the BSS. Unintentional interferers are devices that are not designed to send and receive RF signals such as microwave ovens, florescent lamps, etc. Microwave ovens may reduce throughput and capacity of wireless access by as much as 50 percent⁵.



Figure 12 - Common Interferers

Spectrum

1. RF Propagation

Radio frequency, or RF, is a term that refers to AC having characteristics such that, if a current is applied to an antenna, an electromagnetic (EM) field/wave is generated suitable for wireless communications. However, optimal RF propagation is determined by many factors.

Light waves like infrared have different characteristics than wireless RF waves, they require VLoS. Visual line of sight (VLoS) is where a field operation technician can see from one point to another point, a client to a WGR. RF energy may penetrate many types of substances in the home, and therefore visual LoS is not always required for an RF link such as Wi-Fi. Radio frequency line of sight (RFLoS) is when two RF transceivers (transmitter/receiver) can 'hear' one another. A good example of RFLoS is a WGR and smartphone client connection communicating in a home through wood/plaster walls.

One goal in wireless deployment is clear LoS between the WGR and its client when possible however, most of the time when a wireless RF signal travels over the air (OTA) it is weakened due to many effects such as attenuation (e.g., free space path loss or objects in the LoS), absorption, reflection, scattering, refraction, diffraction, Rician fading / multipath, Rayleigh fading, antenna polarization and low antenna gain. The idea for a quality wireless installation is to try to minimize these propagation impairments when selecting access point or WGR placement.

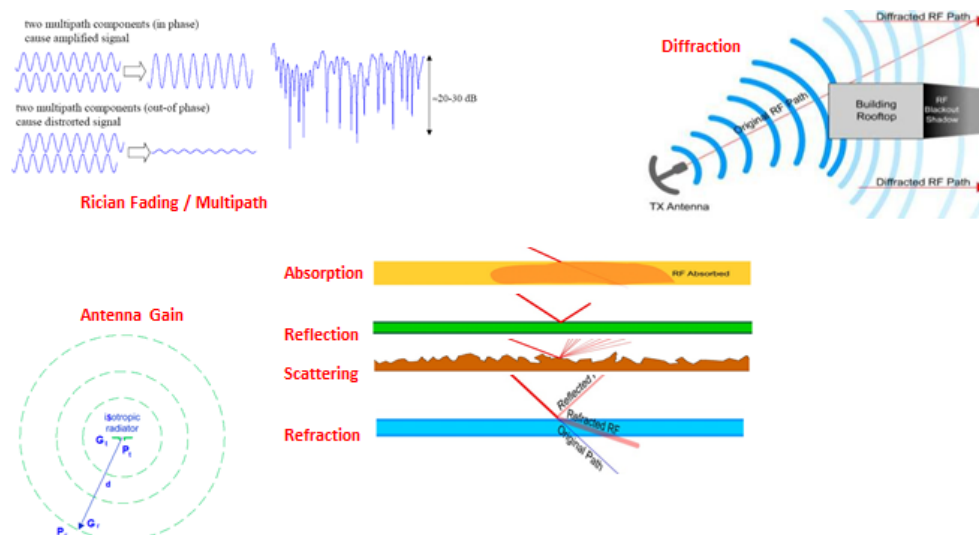
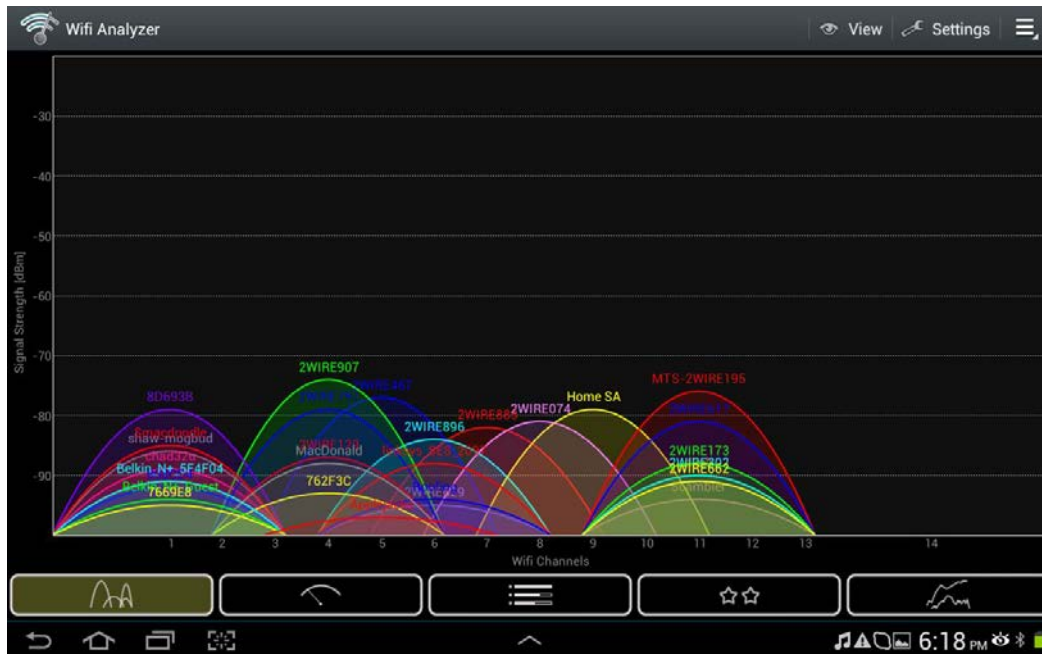


Figure 13 - RF Propagation Effects that Attenuate Signals

The spectrum of 2.4 GHz is only capable of handling 3 channels (shown below) of Wi-Fi (e.g., 1, 6 and 11). In addition, MSOs are using the RF space between these channels for low power 802.15.4 wireless communication. Due to the limitations of 2.4 GHz, the 5 GHz is a preferred band for Wi-Fi communications, offering 5x more spectrum than 2.4 GHz⁶.



There is a myth that there is no interference in 5 GHz band of Wi-Fi because there are less devices in the spectrum. While it may be generally true that fewer devices currently operate in the 5 GHz band that are causing interference as compared to 2.4 GHz devices, however this will evolve over time. Just as subscribers moved from 900 MHz to 2.4 GHz to avoid interference, the "band jumping" effect will catch up with the 5 GHz spectrum. Some signals already exist at 5 GHz, they include cordless phones, radar, perimeter sensors, and digital satellites. DFS is a process to detect radar signals and select alternate frequencies for communication over 5 GHz. DFS may limit an 802.11ac WGR's ability to achieve 160 MHz channels.

In addition to band jumping, higher frequencies like 5 GHz attenuate more than lower frequencies like 2.4 GHz. The reason cable operators introduce tilt into cable low and high RF transmission over a coax access network like the HFC. The higher attenuation of 5 GHz must be taken into consideration as it will reduce RF propagation, range and coverage.



Figure 15 - 2.4 GHz vs 5 GHz Coverage

One of the big items cited by the MSOs is the lack of a site survey performed at the subscriber premises, just dropping the WGR where the modem is located is not a best practice. This practice may create a sub-optimal location for RF to propagate throughout the premises. Proper attention to the location of WGR and modem are important to the QoE. Following company defined pre-installation steps, such as a site survey reveals information important to determining the location of the WGR.

Also, to note that wireless clients have a lower transmit power than a WGR and or obstacles create issues such as hidden node. A hidden node occurs when two clients cannot “hear” or “see” each other (e.g., obstacles, distance or technology) and their traffic causes collisions at the WGR.




Figure 16 - Wireless Spectrum Used by Operators

2. Wireless 802.15.4

Wireless 802.15.4 is positioned to take over as the de-facto wireless standard in the connected home. Technologies like ZigBee, WirelessHART, ISA100.11a, 6LowPAN and Bluetooth and others (LoRaWAN, IC3, Nivis, etc). These devices provide network sensing communication-ability.

ZigBee is a full wireless protocol suit that many MSOs have adopted for IoT applications such as home security/lifestyle. ZigBee was built on top of the 802.15.4 standard and is defined by the ZigBee Alliance. IEEE Wi-Fi is too power intensive for most low power IoT applications, allowing ZigBee to offer a lower cost per chip. Low power wireless technology such as ZigBee offers a mesh topology, common for sensor networks utilized in the connected home that requires communication with a WGR.

ZigBee requires no existing or additional wiring for wireless clients. Additional connections may be made, beyond the four (4) port switch provided by most CE gateways.

802.15.4 technology suffers from many of the same Wi-Fi interferers like; baby monitors, cordless phones, microwaves, co-channel, metal, etc. Intermittent connections due to distance from WGR still exist in low powered networks. As with Wi-Fi, 802.15.4 requires proper training of wireless RF propagation theory to be properly installed. Finally, precautions need to be made for ensuring the security of these networks.

Wireless highway addressable remote transducer (WirelessHART) is an IoT standard also built on 802.15.4 for industrial automation. It is a self-organizing, self-healing and time-synchronized network.

ISA100.11a also is an IoT standard built on 802.15.4, supporting frequency hopping and mesh routing at the data link layer. ISA100.11a supports IPv6 along with reliable transport using TCP. While the 6LowPan network allows IP to be applied even on the smallest sensor devices.

To build IoT data collection devices and networks for smart cities and municipal organizations a protocol known as LoRaWAN is being trialed by some operators around the globe. Some of the other organizations include public utilities, automotive and healthcare.

User Experience

1. Video over Wi-Fi

Managed video cannot afford the issues of Wi-Fi, video is intensive and in many cases still using MPEG technology today. However more MSO are deploying IPTV and video over Wi-Fi. Operational savings can be realized if MSOs can rely on Wi-Fi for managed video distribution. If Wi-Fi is problematic, it will have catastrophic support costs – already seen in many initial Wi-Fi deployments for HSD throughout the world. Problematic Wi-Fi will create huge churn on video subscribers at a time that the MSO is trying to win them back and compete with consumer electronic (CE) and OTT user experiences. Video over Wi-Fi needs to be as reliable as QAM based video!

Video services over untethered devices is how many our customers will consume video content in the future. The speed tier must be checked for a subscriber account to verify the appropriate wireless device is chosen to support the video installation. In cases where devices do not meet expectations, appropriately

set subscriber expectations through education about speed, range, security, device obsolescence, hardware (e.g., WGR), etc.

2. MoCA

MoCA is the foundation for distributing content around the home and jumps splitters to form a full mesh and peer to peer wired network in the home using existing coax cable. Gigabit Ethernet and MoCA will both consistently deliver DOCSIS 3.0 modem speeds of 400+ Mbps MAC throughput (MoCA 2.0) download and that 802.11n is also capable of delivering DOCSIS 3.0 cable modem speeds of 100+ Mbps download (802.11ac supports 500+ Mbps) under some circumstances while getting very close using a variety of client devices and in many locations throughout the home. Enhanced MoCA mode supports 800+ Mbps MAC throughput. MoCA to Wi-Fi extenders are used to improve the Wi-Fi QoE, designed for MSOs to boost/extend Wi-Fi coverage in home network for video and data installations.

3. Better Tools

The first step in improving the wireless environment is having eyes to see the Wi-Fi RF, such as channel IDs used, signal strength, 802.11 standard, S/N, co-channel interference (CCI), adjacent channel interference (ACI), overlapping channel, etc.




Figure 17 - Wi-Fi RF Spectrum Analysis

Take signal strength readings during the site survey using an analyzer before installation. Standardize wireless signal spectrum analyzers installed on laptops/smart phone to report consistent metrics across the MSOs footprint. For example, uses the standard RF meter as some vendors support wireless adapters kits for the meters. These tools must support the 5 GHz band, as more and more devices are moving into this spectrum. Tools may be used to verify whole home connectivity during the post installation steps. As with any new tool, it is important for field operation installers to understand the tools, how to properly use them to locate and identify information quickly.

Wireless signal spectrum analyzers need to identify co-channel interference (CCI) or in band noise as other signals maybe located in the same frequency band. In addition, wireless ACI must be characterized,

as it produces interfering signals located in adjacent bands. Tools should account for inter-symbol interference (ISI) or zero ISI due to multipath of the wireless signal. To characterized overlapping channels, multiple system interference (MSI) must be accounted for due to interfering signals exploiting the same frequency band (e.g., Wi-Fi, 802.15.4)

At the packet level, collision must be minimized when two packets are transmitted simultaneously over a wireless network. Interference reduction will minimize collisions at the PHY layer (mitigation) and can be managed at the MAC layer (avoidance).

A measure of RSSI using power in dBm and typically shown as a negative (-) value provides indicators of interference, high noise floor or low S/N. One caveat about RSSI, vendors may implement this measurement differently (scale of 0-60, scale of 0-255) as it is not a standard. Most vendors recommend a -65 dBm value or better for RSSI⁶ in the 2.4 GHz and 5 GHz bands to be optimized for untethered devices. Be sure to benchmark the devices and set a best practice for RSSI measurements. A higher S/N will produce a higher data rate, the same as a wired coax network. When using MU-MIMO and modulation orders beyond 802.11n, 802.11ac must have a significantly higher S/N, requiring cleaner spectrum and better WGR to client distances. To estimate a 30 dB S/N, take the noise floor (e.g., -95 dBm) minus the RSSI value (-65 dBm).

Table 3 - Sample S/N Values for 802.11n and 802.11ac

Protocol	MCS Index	Modulation	Channel Width (MHz)	Minimum S/N (dB)	RSSI (dBm) <small>*NF = -95 dBm</small>
802.11n/ac	5	16 QAM	20	18	-77
802.11n/ac	5	16 QAM	40	21	-74
802.11n/ac	6	64 QAM	20	20	-75
802.11n/ac	6	64 QAM	40	23	-72
802.11n/ac	7	64 QAM	20	25	-70
802.11n/ac	7	64 QAM	40	28	-67
802.11ac	8	256 QAM	20	29	-66
802.11ac	8	256 QAM	40	32	-63

4. Best Practice Summary

DOCSIS Layer – As like other services within the premises that operate on top of a modem service, verify the DOCSIS service is functioning correctly.

RF Surveying - Locate sources of interference using wireless analysis tools. Determine cause of weak RF signals; large premises, premises with several floors and long distances to WGR. Locate causes of signal loss; brick, concrete, block, metal duct work, etc. Metal studs or heating duct may have a loss of 26 dB! Location of WGR, LoS is the best, open area without obstructions and elevated to laptop working height. 5 GHz provides less coverage than 2.4 GHz. Conduct walkthrough of premises, measure signal strength and identify open channels. Place WGR in an ideal location(s).

Security – no WEP, no open networks; Some MSOs feel AES is best; Encryption compatibility is important with older clients; Understand security options available on devices (IPSec passthru, parental controls, firewall, etc.) Standard passphrase; combination of modem MAC address and phone number is an example. To improve security at the physical layer consider reducing the amount of RF power

transmitted by WGR. MSO portal or app can be developed to help subscriber configure their own security features.

WGR/Client Setup - Large premises, premises with several floors and long distances to WGR require multiple WGRs and both 2.4 GHz and 5 GHz bands; Use different RF channels on each of the WGRs. Operate in 5 GHz; Multiple devices and HD video applications work best with 5 GHz and good RF coverage; Speed tests at 24, 50 and 100 Mbps work best with 5 GHz and excellent RF coverage. eRouter specification from CableLabs offers MSOs control over the wireless eco-system, also allows PnP gateway (not installation). As part of post work test the installation with client devices and subscriber. Allow auto channel selection feature where possible, network is adaptable. Where possible directional antennas offer higher signal gain, omnidirectional may waste RF energy. Many mobile devices may have throughput limits, verify all client devices.

Training - Training on the new wireless and Wi-Fi aspects; Everyone needs to be trained and certified to deploy, support and troubleshoot wireless networks; Schedule refresher training early and often. Utilize SCTE-ISBE working group, chapter and educational resources.

Customer Education - Interview subscriber about the types of devices, number of devices, habits of usage, etc. Review all aspects of installation with the subscriber.

Repeaters - Repeaters are often used to improve the user experience with wireless networks. The repeating devices act like an RF router: receiving a weakened RF signal, regenerating the RF power levels, and retransmit the RF signal. An important field tip, RF signals arriving at the repeater's receiver must have a good S/N. Using repeaters properly many effectively double the coverage and extend the range.

Conclusions and Recommendations

Wireless communications in the home, business and access network are still evolving. Many challenges will still need to be solved, however it forces our industry to continually improve the experience. QoE and a carrier grade version of wireless is becoming more important than previous versions of best effort communication. The wireless experience will determine the MSO's image, operator will need to own and manage the ecosystem.

Leveraging newer amendments like multiple radios, MU-MIMO, beamforming, interference mitigation, etc. will be critical to the QoE. Following best practices for deploying benchmark devices will be vital to customer satisfaction. Wireless deployments require better tools that provide key metrics and data (e.g., S/N, MCS, ACI, CCI, etc.) for both Wi-Fi and low powered wireless 802.15.4 radios. The 802.11ac amendment sets the foundation for carrier grade deployments while the 5 GHz band will alleviate some of the congestion issues with 2.4 GHz to improve the QoE.

Educating our field operations workforce in wireless is highly recommended and most important, due to large call volumes and truck rolls related to wireless technology in our industry. SCTE-ISBE has a working group that has been developing wireless best practices for our industry, free to join. SCTE-ISBE has a training and certification program for broadband professionals, learn more at SCTE.org/BWS. In addition, the SCTE-ISBE has partnered with the certified wireless network professional (CWNP.com) to provide a cable version of the CWNP wireless career path.

Abbreviations

AC	alternating current
ACI	adjacent channel interference
AES	advanced encryption standard
AP	access point
BSS	basic service set
CCI	co-channel interference
CPE	customer premises equipment
CSMA	carrier-sense multiple access with collision avoidance
dBd	decibel relative to a dipole antenna
dBi	decibel relative to an isotropic radiator
DCA	dynamic channel allocation
FSPL	free space path loss
Gbps	gigabit per second
GI	guard interval
HEW	high efficiency WLAN
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of things
ISM	Industrial, Scientific and Medical
M2M	machine to machine
MAC	media access control
MCS	modulation and coding scheme
MIMO	multiple-input multiple-output
MoCA	Multimedia over Coax Alliance
MU-MIMO	multi-user multiple-input multiple-output
ns	Nanosecond
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OTT	over the top
PER	packet error rate
PNM	proactive network maintenance
QAM	quadrature amplitude modulation
QoE	quality of experience
RDK-B	reference design kit for broadband
RFLoS	radio frequency line of sight
RRM	radio resource management
RSSI	receive signal strength indicator
S/N	signal to noise ratio
SDM	spatial division multiplexing
SISO	single-input single-output
SSID	service set identifier
TPC	transmit power control
UNII	Unlicensed National Information Infrastructure
VLoS	visual line of sight
WGR	wireless gateway router

WLAN	wireless local area network
WMM	Wi-Fi multimedia
WPA	Wi-Fi protected access

Bibliography & References

- [1] Infonetics Research; Global Unit Shipments
- [2] Quality of Service on Carrier Grade Wi-Fi; WBA; CableLabs
- [3] WBA; <https://www.wballiance.com/carrier-wi-fi-services/>
- [4] Reference Design Kit for Broadband; <http://rdkcentral.com/>
- [5] Cisco; <https://www.cisco.com/c/en/us/solutions/collateral/enterprise-networks/802-11ac-solution/q-and-a-c67-734152.html>
- [6] Wi-Fi Design Evolution; Xirrus, Inc.
- [7] The Certified Wireless Network Administrator, 4th Ed.; Official Study Guide
- [8] SCTE-ISBE; www.scte.org

Optimizing and Protecting the Value of Unlicensed Spectrum

A Technical Paper prepared for SCTE•ISBE

Narayan Menon, CTO/EVP of Engineering, XCellAir

Angelo Cuffaro, Senior Director R&D, XCellAir

Jean-Louis Gauvreau, Director R&D, XCellAir

Todd Mersch, EVP of Sales & Marketing, XCellAir

Introduction

The primary wireless access technology at the disposal of cable operators today is Wi-Fi. The ubiquitous proliferation, low device cost and lack of spectrum license has made Wi-Fi the clear choice. However, as cable wireless services transition from best effort to managed broadband, the ability to deliver reliable high-performance access using unlicensed spectrum is critical. This paper describes the challenges, techniques and results obtained from intelligently optimizing unlicensed spectrum. It also discusses ways to enable harmonious coexistence of Wi-Fi with new technologies such as LTE License Assisted Access (LTE-LAA), LTE – Wi-Fi Aggregation (LWA) and MuLTEfire (standalone LTE), exploiting techniques being leveraged in today's Wi-Fi optimization landscape.

Unlicensed Spectrum – the Congestion Problem

Cable operators have been aggressively expanding their Wi-Fi networks over the last couple of years. This has been driven by goals to expand the customer base, reduce churn and deliver a variety of value-adding services with a high-quality user experience. These could include managed voice, video, gaming and IoT applications.

As Wi-Fi networks expand and densify, radio resource optimization becomes critical to delivering a high quality of experience. Wi-Fi works on a Listen Before Talk (LBT) basis, where devices contend politely for access to the medium (channels). If a channel is in use by a device, other devices wait for the channel to become free before accessing it.

In a dense Wi-Fi network, the large number of contending devices and access points, and a high level of user activity, can combine to cause heavy contention, resulting in congestion. This is true even in a really “well-behaved” Wi-Fi cluster, where devices and access points can all see one another, and defer to each other gracefully when contentions occur. Excessive medium contention results in long wait times for packet transmission opportunities, high latencies and low throughputs.

This issue is likely to be exacerbated when other technologies, such as LTE, start using unlicensed spectrum. LTE-LAA devices will increase contention levels on Wi-Fi channels. Wi-Fi devices will have to also compete with LTE-LAA terminals and access points, and LTE-LAA devices will be contending with Wi-Fi and other LTE-LAA endpoints.

Addressing congestion and interference issues requires the use of radio resource management techniques. Techniques for intelligent allocation of Wi-Fi channels, steering of client devices to a less congested Wi-Fi band, power control and device mobility between access points in a multi-AP environment become critical in making Wi-Fi work well despite the increased contention from Wi-Fi and other technologies using unlicensed spectrum.

Automated management and optimization solutions must include powerful radio resource management (RRM) tools to dynamically provision and tune radio resources. Running RRM algorithms on a cloud server, and interacting with the AP to set and change channels, bands, power levels and other parameters provides scalability, maintainability and a centralized view to deliver a network level solution. The RRM schemes must work proactively to avoid congestion and interference, and steer client devices between

APs to improve coverage. The end results are the optimal usage of available Wi-Fi capacity, dramatically improved latency/ jitter/ throughput performance, and a significantly enhanced quality of experience.

Environmental Observations

XCellAir carried out a study to characterize channel usage in a real-life Wi-Fi environment. The observations were done in downtown Montreal, a busy urban environment with multi-storied office buildings and dense Wi-Fi deployment. At any given time, anywhere from 100 – 250 access points (APs) are visible in the immediate vicinity of a Wi-Fi network.

The objective of the study was to observe Wi-Fi channel availability, i.e. to see how much bandwidth headroom was available at a given time in a typically busy deployment scenario. The study gathered per-channel utilization levels over several periods of time. This was done for both 2.4 and 5GHz bands.




Figure 1 - Per-Channel Utilization Levels in 2.4GHz Band

Figure 1 shows the channel utilization pattern observed for the 2.4 GHz channel over a period of several hours. On the x-axis are the 2.4 GHz channels – 1 to 11. Channel utilization levels (0 – 0.9 or 90%) are shown on the y-axis. Channel utilization reflects the degree (in percentage terms) to which a channel has been occupied by Wi-Fi devices over a measurement period. Each vertical spike on the graph reflects utilization data for the given channel over a 15-minute time period. The horizontal trend lines running through the graph indicate the minimum (red line), average (blue) and maximum (yellow) utilization levels per channel.

Some interesting conclusions can be made from Figure 1:

- Spikes in channel occupancy are visible on several channels at different points in time. These reflect inflection points at which contention and congestion levels increase, and service quality starts to degrade, e.g. high packet error rates, high latencies and jitter, low throughputs etc.
- However, based on average utilization levels, there is bandwidth headroom available at any given time – more on some channels than others. Not all channels typically experience high occupancy at the same time.
- Considering an allocation pool of even five of the channels (e.g. channels 1, 4, 6, 8 and 11), up to two channels' worth of aggregated bandwidth is available on average (considering a maximum channel occupancy of 80 – 85%).
- A dynamic channel allocation algorithm can unlock this bandwidth by moving access points from heavily loaded channels to less occupied ones.




Figure 2 - Per-Channel Occupancy in 5GHz Band

Figure 2 shows a similar channel occupancy view observed for the 5 GHz band. In general, this band is less loaded today than the 2.4 GHz band. A band steering feature (i.e. the ability to move dual-band capable devices on a tightly loaded AP from the 2.4 GHz channel to a less loaded 5GHz channel) takes advantage of this loading imbalance by moving active devices from 2.4 to 5 GHz to leverage the clearer channels and higher bandwidth in the 5 GHz band.

A separate study conducted by XCellAir looked at access point products deployed in the Montreal downtown neighborhood, and observed channel change behavior exhibited by the APs. The conclusion was interesting – less than 8% of the surveyed APs changed channels, once they powered up. Without RRM capabilities, APs stay rooted to a channel regardless of how high the congestion is, and are unable to leverage headroom available on other channels.

RRM Benchmarking & Impact on Quality of Experience

This section discusses a representative sample of test results highlighting impacts of congestion / interference on service quality, and the performance-enhancing influence of two categories of optimization tools – dynamic channel management and band steering. While these results are for Wi-Fi, they are relevant to multi-technology coexistence scenarios in the unlicensed band, e.g. LTE-LAA coexisting with Wi-Fi.

1. Dynamic Channel Management

Dynamic channel management algorithms can mitigate Wi-Fi channel congestion and interference. They detect developing congestion; at a configurable trigger point, if the algorithms decide to change the AP's operating channel, they select a cleaner (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 KPIs to degrade to poor levels as a result of congestion and interference.

Test Setup & Methodology

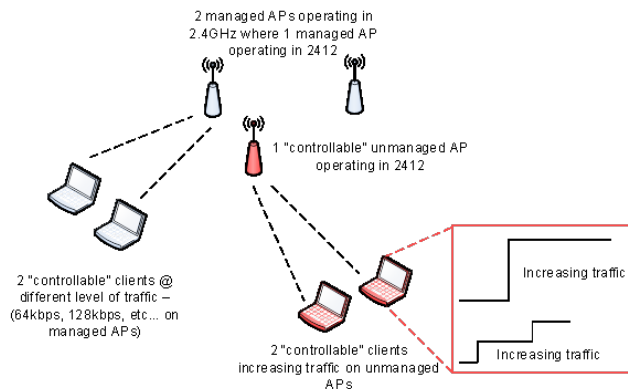


Figure 3 - Example Lab Test Setup

The test setup outlined in Figure 3 was used to create channel congestion and observe impacts on key operational metrics, e.g. latency, jitter, throughput, etc. A “target” AP was used as an observed target, with multiple client devices connected to it. Clients were distributed at distances of 3 – 10 meters from the access point. A mix of Voice over Wi-Fi (VoWiFi) and video traffic was run through the target AP. End-to-end metrics such as jitter, latency and throughput for the Voice over Wi-Fi (VoWiFi) and video traffic were measured using the IxChariot tool.

Separate “aggressor” APs were set up on the same channel and loaded with multiple clients sending iPerf and Youtube data. Identical tests were run with RRM disabled and enabled. The objective was to load up the channel to a congested level and: (a) observe the deteriorative impacts of rising congestion on service quality on the target AP (without RRM); and (b) observe the congestion mitigation and resultant quality improvement when RRM was enabled.

Test Results

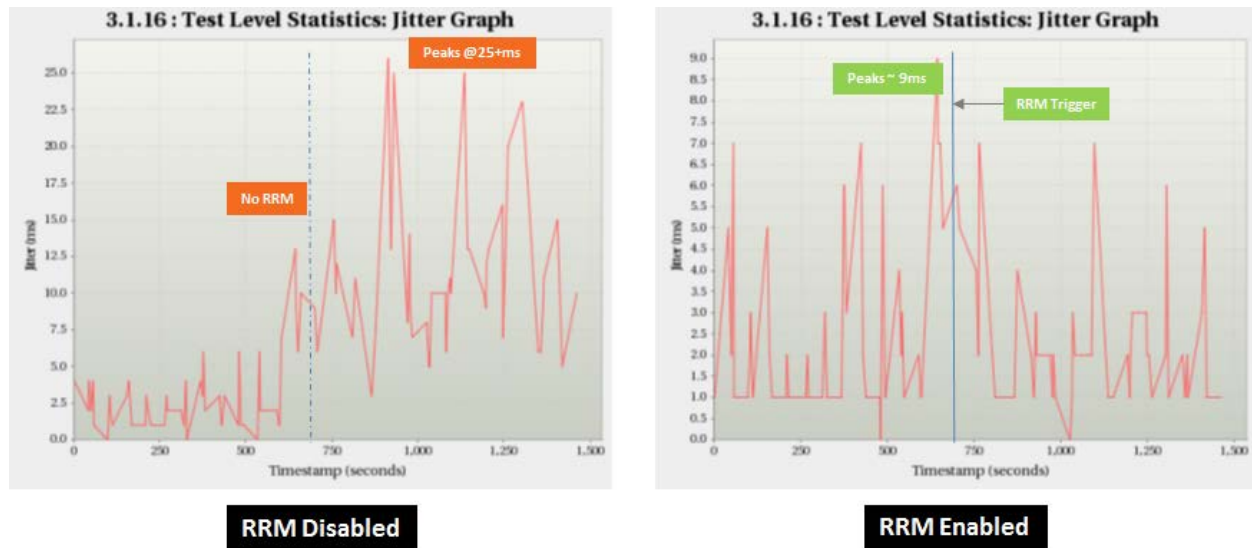


Figure 4 - Impacts on Jitter

Figure 4 depicts the impacts on jitter for VoIP, with RRM turned off and on respectively. With RRM disabled, jitter reaches a peak of 25 ms, and stays in a high range subsequently, hitting high peaks frequently. With RRM enabled, jitter is allowed to reach a peak of only 9 ms before RRM kicks in and brings congestion under control. Jitter subsequently settles into a 1 – 7 ms range, compared to the 5 – 25 ms range demonstrated without RRM.




Figure 5 - Impact on VoWifi Latency

Figure 5 illustrates the impacts of channel congestion and RRM on one-way latency. Latency increases by an order of magnitude as congestion builds up, and without RRM, peaks at around 580 ms and persists within the 0 – 350 ms range. Latency levels above 100 ms result in noticeable interactivity and echo issues with voice calls. With RRM enabled, the congestion situation is corrected before latency escalates; in this case, latency settles down in an excellent 0 – 5 ms range, with a much lower and far more acceptable peak value.




Figure 6 - Impact on Throughput

Figure 6 and Figure 7 respectively show the impacts on throughput and Mean Opinion Score (MOS) for VoWiFi. RRM keeps throughput much steadier and minimizes dips in throughput levels. The MOS drops to unacceptable levels without RRM, but hovers between excellent and acceptable levels with RRM enabled.




Figure 7 - Impact on MOS score for VoWiFi Call

2. Band Steering

RRM can also support band steering for associated clients, which can be steered to a different radio/frequency band to improve QoE. When an AP or specific radio on the AP is overloaded, and a channel change does not mitigate the situation, band steering functionality moves clients from the 2.4GHz band to the 5GHz band, or vice versa, to mitigate radio overload scenarios and improve QoE.

Test Setup & Methodology

To assess the performance impacts of band steering, a lab test setup was used to create radio overload and observe impacts on key operational metrics, e.g. latency, jitter, throughput, etc. On the target AP, multiple observable traffic flows (VoWiFi, video and data) were run between connected clients on the 2.4 GHz radio. End-to-end metrics such as jitter, latency and throughput for the VoWiFi and video traffic (“target traffic”) were measured using the IxChariot tool.

Additional clients were then connected to the same (2.4 GHz) radio to introduce “aggressor” traffic on the radio. When the radio became overloaded, band steering kicked in and moved the “target traffic” flows from the 2.4 GHz radio to the 5 GHz band. The band steering algorithm factors in path loss adjustments for the move from 2.4 GHz to 5 GHz, i.e. given that range reduces when a client moves to 5 GHz, the algorithm only moves clients that have strong enough signal on 2.4 GHz to begin with; a calculation is done to ensure there is sufficient “coverage buffer” to withstand the path loss on moving to 5 GHz. The Chariot tool measured performance KPIs before and after the steering action.

Test Results




Figure 8 - Band Steering – Impact on MOS Score of VoWiFi

Figure 8 reflects the impacts of developing congestion and the band steering action on the quality of the VoWiFi call in progress (reflected by the MOS score). When the aggressor traffic raised congestion on the 2.4 GHz to high enough levels, the measured MOS deteriorated to unacceptable levels (staying between 1 and 2.25). When band steering kicked in and moved these clients to the much cleaner 5 GHz band, voice quality improved immediately and moved back up to acceptable levels.




Figure 9 - Band Steering – Impact on Delay Factor

For the same test, Figure 9 shows the impacts of overload and band steering correction on delay factor for the video traffic flow. The delay factor (DF) KPI is measured in milliseconds, and is a time value reflecting the data buffer that has to be maintained to eliminate time distortions (jitter). DF levels reach around 1,300 milliseconds when the radio becomes congested, and settle back down to acceptable levels after the traffic flows are steered to the 5 GHz band. DF values in the 0 – 100 ms range are generally deemed acceptable.

3. Conclusions from Lab Tests

In general, the lab tests demonstrated the following trends:

- RRM brings about order-of-magnitude improvements in latency, jitter, throughput and other key service quality metrics
- With RRM enabled, we see significantly reduced metric spikes / dips. In other words, the KPI values deteriorate far less. Also, with RRM kicking in, we see much better settled metric levels. RRM ensures that congestion is quickly reduced, and that the KPIs that control service quality do not get out of hand.
- Without RRM, we see that the issues persist and escalate, resulting in much greater deterioration, drastic impacts to service quality and possibly eventual loss of service.
- These tests measured the impacts of congestion and RRM on real traffic flows, and it is worth noting that the results from the tests are fully applicable to similar traffic running in real world environment.

Enabling Wi-Fi Coexistence with New Technologies in Unlicensed Spectrum

While this paper has discussed RRM techniques and performance results in the context of Wi-Fi, they are equally applicable to cross-technology coexistence in the unlicensed band. Spectrum sharing and IoT are two areas that will introduce additional technology types into the unlicensed band. Flavors of LTE are poised to share spectrum with Wi-Fi; LTE License Assisted Access (LTE-LAA) and MulteFire will seek to share the 5 GHz band with Wi-Fi. IoT will stimulate increased use of technologies like ZigBee and Bluetooth, which also operate in the 2.4 GHz band.

Congestion and interference will multiply in severity when these new technologies operate in the band and contend with Wi-Fi and each other for the same set of resources. This will be driven by the many more devices and access points now “queueing up” to access the same set of unlicensed band channels. The 5 GHz band is relatively clear today, but can get significantly congested with the arrival of LTE flavors. The already crowded 2.4 GHz band will gain more popularity with IoT operating in it.

Some of the LTE – Wi-Fi spectrum sharing issues will be addressed by the lower layers of the LTE system. The new LTE-based unlicensed band technologies will likely support Wi-Fi –like schemes like Listen Before Talk (LBT), Dynamic Frequency Selection (DFS), Discontinuous Transmission and others, in an effort to minimize interference with Wi-Fi. In effect, these technologies will exhibit Wi-Fi –like behavior, at least with respect to channel access etiquette. These aspects are being standardized within 3GPP.

However, radio resource management schemes will be as important for interference and congestion avoidance, and to enable optimal sharing of unlicensed band resources between devices belonging to different technologies. RRM is a necessary complement to schemes like LBT, and can help reduce channel contention and the amount of time devices end up in “wait mode”.

What is also interesting is that centralized optimization schemes (like XCellAir’s) that rely on broader system performance metrics such as channel quality, device QoS parameters etc. can assess and mitigate congestion / interference impacts in a multi-technology environment without necessitating any cross-technology interactions. RRM can be used by a Wi-Fi system, for example, to protect itself from other Wi-Fi or LTE systems, without the system having to know of or communicate with the other systems present.

1. Channel Management

Dynamic channel management can be used by any system to protect itself from interference from other devices in the band. For example, a Wi-Fi system can proactively be allocated the best channels to mitigate interference:

- Between Wi-Fi devices in its own network (same operator)
- To / from devices in other Wi-Fi networks (other operators)
- To / from devices in LTE systems operating in the same band (same or other operators)
- To / from devices in Bluetooth or ZigBee systems operating in the same band.

The same schemes can be applied to an LTE-LAA / MulteFire system, or to an IoT system, to coordinate within its own network or protect it from other LTE or Wi-Fi systems. In a scenario where the operator is deploying both Wi-Fi and LTE devices, RRM can facilitate coordinated resource allocation across Wi-Fi and LTE. The performance improvements discussed in this paper are fully obtainable in such a multi-technology scenario.

2. Band Steering

Wi-Fi systems can continue to leverage band steering to keep devices in the best possible band. The 5 GHz band is a lot less congested than 2.4 GHz today, but that is likely to change with LTE devices entering the unlicensed band. Wi-Fi systems can use band steering to use the two bands optimally – move devices out of 5 GHz if high congestion is being caused by LTE, for example. Similarly, devices can be moved out of the 2.4 GHz band if IoT systems cause congestion there.

LTE-LAA and other systems can also use steering techniques to move devices and traffic flows between the licensed and unlicensed bands.

3. Power Control

In high-density deployments, Wi-Fi systems can use power control to minimize interference to other devices (Wi-Fi, LTE, others). Conversely, power can be increased to fill coverage gaps in pockets where deployment is less dense.

Conclusion

Clearly, radio resource management (RRM) is a critical necessity for optimal operation of networks in the unlicensed band. These techniques are readily applicable to Wi-Fi today, especially as cable operators and Internet providers roll out large and dense Wi-Fi networks. As use of the unlicensed band grows to include LTE and IoT devices, congestion and interference issues will worsen significantly, and the use of RRM schemes will become even more critical.

Resource contention, congestion and interference can cause harmful degradations in service quality, as tests described in this paper illustrate. Equally importantly, optimization schemes such as dynamic channel management, band steering and power control can help systems operating in unlicensed bands mitigate congestion and prevent severe service degradation - as quantified in this paper. The end results are the optimal usage of available unlicensed band capacity, dramatically improved KPIs and service performance, and a significantly enhanced quality of experience. These benefits are available to any type of system operating in this band – Wi-Fi, LTE and others.

It is also clear that even in dense network deployments, there is generally some spare headroom that can be leveraged by a smart RRM algorithm. XCellAir's studies have shown that there is, in general, spare channel bandwidth in the system – largely unutilized because most existing systems are unable to dynamically reallocate resources. An intelligent RRM scheme can unlock this free bandwidth for use by access points experiencing congestion or interference.

Can eMTC IoT Be Supported over the HFC Network

A Constructional Proposal of MOVING IOT into HFC

A Technical Paper Prepared for SCTE•ISBE by

Shengbo Ge

Shanghai, China

Cable HFC/CRDC CHG/ Cisco Systems (China) Research and Development Co.,

Ltd 16 Floor, Block C of Keji Building, 900 Yishan Road

86-21-24014382

shge@cisco.com

Introduction

This proposed solution combines eMTC wireless technology with an HFC switch node in order to leverage IoT services over the HFC network. The new designed enhanced machine type communications cellular device (eMTCcell) combines with HFC switch node to transmit IoT data over The HFC network. The new designed HFC optic switch node can work as an outdoor Ethernet switch, its digital uplink fiber acts as the role of IoT mass data backhaul.

Hybrid fiber-coaxial (HFC) Network

Hybrid fiber-coaxial network is a traditional solution for broadcast video RF signal transmission that combines optical fiber and coaxial cable. It has been commonly employed globally by cable television operators since the early 1990s.

In recent years, The HFC network has gradually evolved to support high definition video on demand (VOD) services and internet broadband access services. As both digital and analog fibers are deployed between the head-end and fiber optic node in the field, the HFC network has the capability to act as the backhaul role of aggregated data from access network.

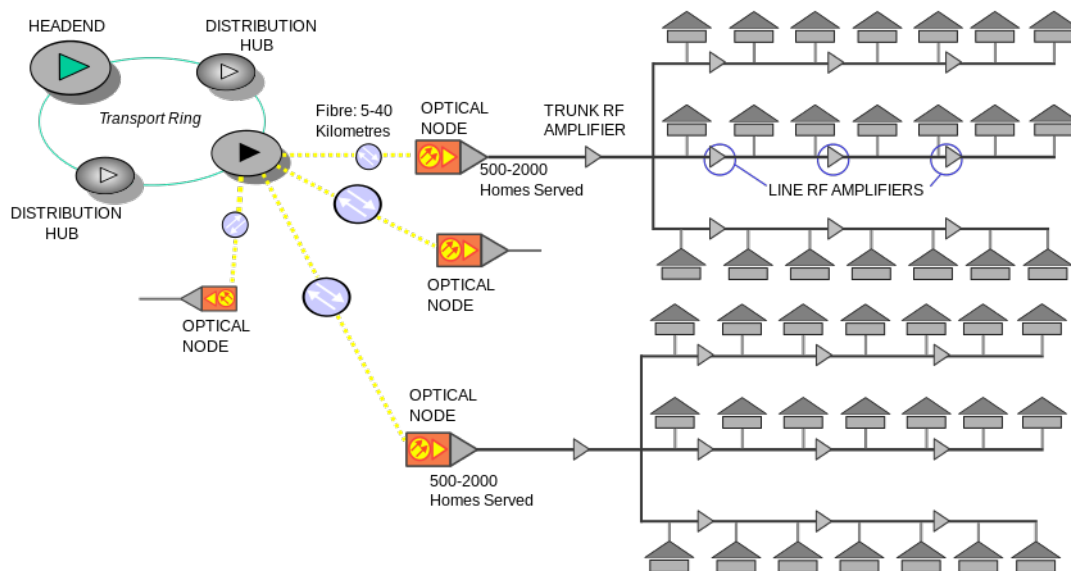


Figure 1 - A Common HFC Architecture (Source: WIKIPEDIA)

HFC Optic Node Evolves to Digital Device

A traditional linear fiber optic node has a RF optical receiver, which converts the downstream optically modulated signal coming from the head-end or hub to an electrical signal going to the homes. Today, the digital fiber optic node can work as an Ethernet switch providing broadband and narrowband access.

The optical portion of The HFC network provides a large amount of flexibility. If there are not many fiber-optic cables to the node, wavelength division multiplexing can be used to combine multiple optical

signals onto the same fiber. Optical filters are used to combine and split optical wavelengths onto the single fiber.

From an architecture perspective, the HFC network provides the foundation for mass data transmission of IoT services.

Why eMTC?

NarrowBand IoT (NB-IoT) is a Low Power Wide Area Network (LPWAN) radio technology standard that has been developed to enable a wide range of devices and services to be connected using cellular telecommunications bands. {Wikipedia}

eMTC (Enhanced Machine Type Communication, often referred to as LTE-M), is part of 3GPP Release 13 and includes additional cost reduction measures, 1Mbps data rates in the uplink and downlink, and reduced transmit power.

Both NB-IoT and eMTC are typical solutions of Internet of Things mass data wireless transmission based in LTE technology. This paper focuses on eMTC as a detailed example.

	LTE Cat 1	LTE Cat 0	LTE Cat M1 (eMTC)	LTE Cat NB1 (NB-IoT)	EC-GSM-IoT
3GPP Release	Release 8	Release 12	Release 13	Release 13	Release 13
Downlink Peak Rate	10 Mbps	1 Mbps	1 Mbps	250 kbps	474 kbps (EDGE) 2 Mbps (EGPRS2B)
Uplink Peak Rate	5 Mbps	1 Mbps	1 Mbps	250 kbps (multi-tone) 20 kbps (single-tone)	474 kbps (EDGE) 2 Mbps (EGPRS2B)
Latency	50-100ms	not deployed	10ms-15ms	1.6s-10s	700ms-2s
Number of Antennas	2	1	1	1	1-2
Duplex Mode	Full Duplex	Full or Half Duplex	Full or Half Duplex	Half Duplex	Half Duplex
Device Receive Bandwidth	1.08 - 18 MHz	1.08 - 18 MHz	1.08 MHz	180 kHz	200 kHz
Receiver Chains	2 (MIMO)	1 (SISO)	1 (SISO)	1 (SISO)	1-2
Device Transmit Power	23 dBm	23 dBm	20 / 23 dBm	20 / 23 dBm	23 / 33 dBm

Figure 2 - eMTC Preliminary Specification (Source: 3GPP)

As one of the Low Power Wide Area (LPWAN) technologies, eMTC supports geolocation and mobility applications with licensed frequency spectrum, eMTC's downlink/uplink 1Mbps peak rate is higher than traditional GPRS and Zigbee.

eMTC IoT over HFC Architecture

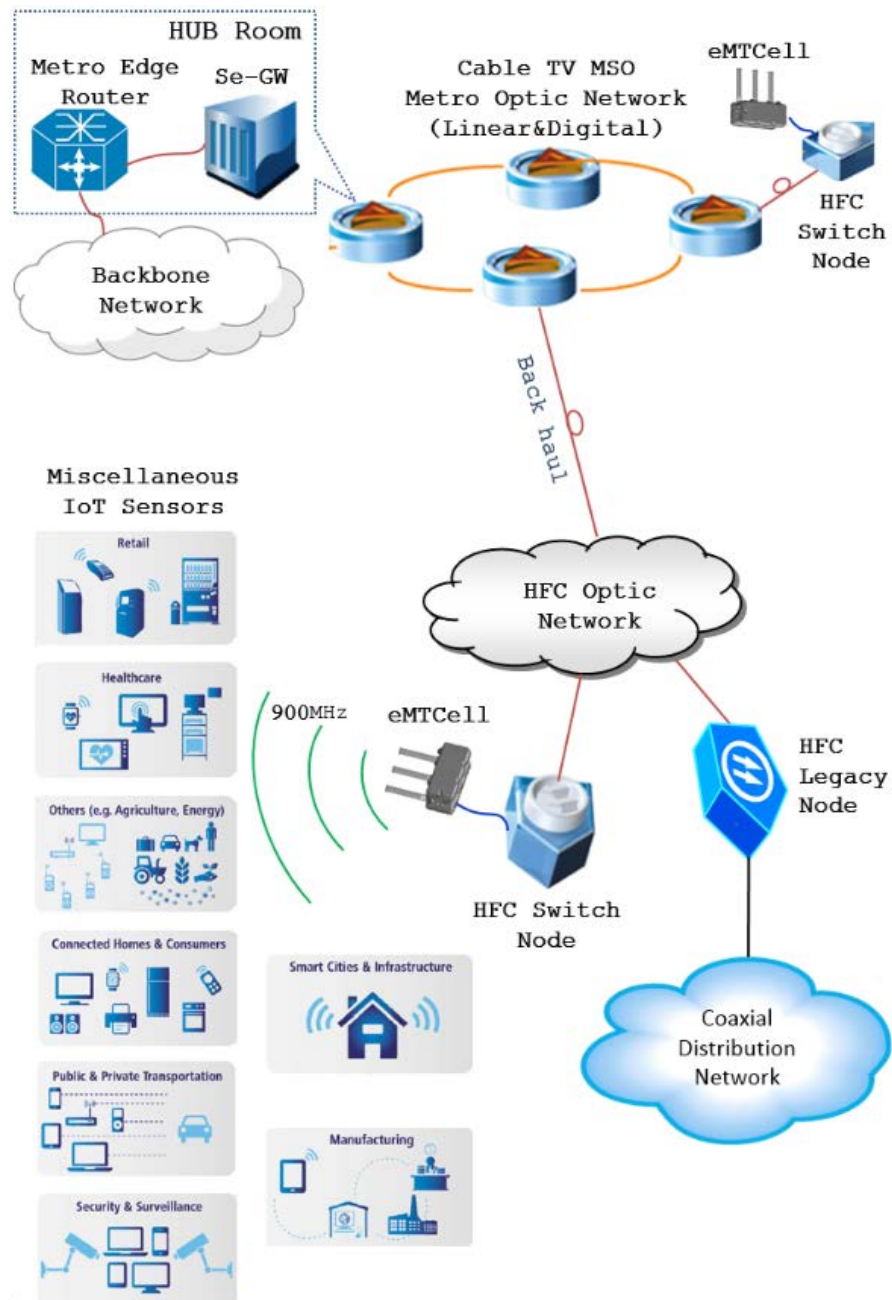


Figure 3 - eMTC over HFC Architecture

This architecture demonstrates how eMTC and IoT combine with the existing the HFC network. In many HFC networks, the fiber optic node is located within the last-mile distance to the subscribers, and eMTC technology typically covers a one mile semi-diameter area. eMTCCell is considered to be designed as a standalone device nearby to the HFC fiber optic node or as a built-in module to the HFC node. eMTCCell

is a key device which connects IoT sensors and the HFC network, and functions like an outdoor Femtocell.

eMTCCell Conception and Preliminary Specifications

To realize eMTC technology over the HFC network, a new eMTCCell is to be designed, whose deployment location is outdoor and provides a 1-mile diameter coverage capability.

The proposed standalone eMTCCell includes these key features:

- Compact design, easy for field installation
- Outdoor waterproof housing
- eMTC (Rel.13) standard support
- IP protocol support
- Fast Ethernet RJ45 interface
- (power over Ethernet) PoE support
- Low power consumption
- Low cost
- Remote EMS support

eMTCCell is defined as: a small-sized, low power cellular wireless access device. It must connect to the nearby HFC optical fiber node's Ethernet switching module to communicate with the core broadband network. RF coverage is provided via single-input-single-output (SISO) antenna, supports access capability of up to 10,000 IoT sensors, and antenna port average transmitting power is expected to less than 5W.

The eMTC chipset or eMTC /NB IoT dual mode chipset vendors include Qualcomm, Marvell, Intel, ZTE, Sequans, MTK, Nordic and Altair. The industry predicts some chipsets being released before 3Q2018.

Table 1 – eMTCCell Proposed Specifications

Housing Dimensions	
Dimensions (mm)	(compact design)
Operating Temperature	–20°C to +55°C or –4°F to +131°F
Relative Humidity	10% ~ 90%
RF Parameter	
Antenna Port Power (average)	<5W
Operating Frequency	900MHz
Receive Bandwidth	1.08MHz (carrier bandwidth 1.4MHz)
Electrical Parameter	
Power Supply	PoE , 36-90 VAC over Coax
Max. Power Consumption	TBD
Surge Voltage	2 KV (composition surge) 6 KV (ring surge)

Note- CableLabs specification controlling is recommended




Figure 4 – eMTCCell Ethernet Connection and Power Supply

HFC Switch Node Key Features

HFC switch node is the new type of HFC digital node whose typical features should include:

- Outdoor waterproof housing
- Build-in switch module which provides 10Gbps uplink capability and multiple Fast Ethernet and Gigabit Ethernet access ports
- Support RJ-45, SFP, SFP+ interfaces
- Support port-based VLAN
- I-temp switch chipset and I-temp optical transceivers
- Support PoE
- Support PoC
- Operating Temperature Range -40° F to 140° F (-40° C to 60° C)

System Functional Overview

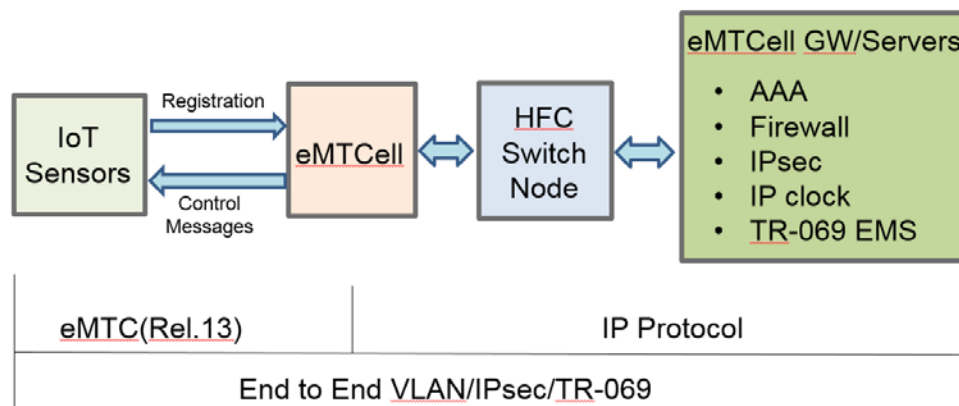


Figure 5 – System Functional Overview

The end-to-end communication mechanism includes two key protocols: eMTC protocol and IP protocol. The eMTC(Rel.13) protocol is responsible for wireless communication between IoT sensors and eMTCCell. The IP protocol is responsible for communication between the eMTCCell and the eMTCCell gateway and relative servers.

The Security Gateway (Se-GW) in the head-end provides IoT data security. Each eMTCCell builds an encrypted tunnel to the Se-GW using IPsec technology. The Se-GW supports numerous encrypted tunnels, switching users in and out of different tunnels.

Other servers are required to be responsible for authentication, authorization and accounting (AAA), IP clock, Element Management System (EMS), etc.

eMTCCell over HFC Application Example

1. eMTCCell Deployment

Cable industry MSOs have deployed large numbers of fiber optic nodes covering geographic regions of cities and country-sides, eMTCCell devices can be paired with nearby HFC nodes (<100m) and installed on: poles, aerial strands, or on top of or on the side of buildings. The compact housing design allows easy field installation. CAT-5 shielded cable is used to connect between the eMTCCell and HFC fiber optic node. eMTCCell obtains power via the power over Ethernet (PoE) or power over coaxial (PoC) cables.

2. eMTCCell Initiation

When the eMTCCell is powered on, it detects the connection to the node and Se-GW, and obtains an IP address from the DHCP server, downloads the latest firmware image, senses the wireless surroundings, and auto-adapts the relative parameters. All of these actions are self-initiated.

If the eMTCCell detects a nearby eMTCCell's signal in the same frequency, it automatically reduces the transmission power and coverage area to avoid interference.

3. IoT User Authentication Mechanism

In mobile IoT applications (e.g., intelligent watch), if an IoT sensor moves from location A to location B, the behavior triggers a re-authentication:

- In a typical application, a group ID is assigned to the eMTCCell devices in the same management group, and the IoT sensor's user ID is registered under the group ID. If the sensor moves among eMTCCell devices belonging to the same group ID, the management system regards the user as legal, and re-authentication is not required.
- The eMTCCell devices in the same group are assigned to the same Location Area Code and Route Area

Code. As seen from the management system, the movement among these eMTCCell devices belongs to the same area.

- The eMTCCell gateway monitors the mobile user's relocation behavior. If the source ID and destination ID belong to the same group, the eMTCCell gateway processes the operation and sends request messages to the destination eMTCCell devices. This mechanism avoids multiple re-authentication requests and terminates the inter-eMTCCell requests when the user relocation happens within the same group.

Conclusion

The evolving HFC network has made it possible to support IoT services, even though the end-to-end solution will have scale and security challenges ahead.

Abbreviations

AAA	<i>authentication, authorization, accounting</i>
AP	access point
bps	bits per second
DHCP	dynamic host configuration protocol
DOCSIS	data over cable service interface specifications
EMS	element management system
eMTC	enhanced machine type communications
eMTCCell	enhanced machine type communications cellular device
FEC	forward error correction
GPRS	general packet radio service
GW	gateway
HFC	hybrid fiber-coax
Hz	hertz
IOT	internet of things
ISBE	International Society of Broadband Experts
LPWAN	low power wide area network
LTE	long term evolution
MSO	multi-service operator
NB-IOT	narrow band internet of things
POC	power over coaxial
POE	power over Ethernet
RF	radio frequency
SCTE	Society of Cable Telecommunications Engineers
SFP	small form-factor pluggable
VLAN	virtual local area network
VOD	Video on demand

Bibliography & References

Hybrid Fibre Coaxial Cable; Services.eng.uts.edu.au

3GPP Low Power Wide Area Technologies; GSMA White Paper

LTE-M & 2 Other 3GPP IoT Technologies To Get Familiar With, Brian Ray; Link Labs

Principles for Interoperability in the Internet of Things

A Technical Paper prepared for SCTE•ISBE by

J. Clarke Stevens

Principal Architect, Emerging Technologies
Shaw Communications
2420 17th Street
Denver, CO 80202
587-393-0605
clarke.stevens@sjrb.ca

Introduction

Perhaps the biggest problem with the emerging Internet of Things (IoT) is that there are so many standards and proprietary systems. Apple, Google and Amazon each have proprietary approaches. Standards on relevant IoT topics are available from IEEE, IETF, W3C, ISO/IEC and virtually every other standards body. Even organizations trying to comply with standards are forced to choose between incompatible options. The integrated Smart Home and other IoT applications and systems cannot achieve their promise if the billions of connected things can only connect with a limited subset of the other devices. This paper will outline a number of principles that can enable interoperability, scalability and security while including all of the devices that need to be connected. The Open Connectivity Foundation (OCF), as one of the premier standards organizations trying to solve the interoperability problem, will be measured against these principles. The SCTE IoT working group is investigating OCF and several other IoT topics to determine if there is work for SCTE to do to provide guidance for its members.

Basic Architecture

The prospect of the Internet of Things is one of the most exciting prospects to come along in computing for some time. Science fiction has long anticipated this development, but it has always been too demanding as an application, too expensive in terms of resources, or too impractical in terms of size or speed. All of those barriers are now coming down. Advances in hardware have reduced size and cost. The size reduction has in turn increased compute capability and improved speed to the point that a sophisticated computing device can be about the size of a piece of confetti and process data in essentially real time.

What remains to do is build the infrastructure that can combine this technology with cooperation between organizations so that the same scale and benefits of the Internet can be realized in the Internet of Things. The true benefit of the Internet of Things comes from the combination of computing and communication. The communication is only useful if the various devices are speaking the same language. This is the role of standards.

1. Data Modelling

In defining how the Internet of Things should communicate, it is important to firstly understand what things will be communicating. For the most part, these “things” are real-world devices like lights, refrigerators, lathes and carburetors. Standards are often tripped up by disagreeing on how to define these real-world objects. However, the very fact that the objects exist in the real world implies a certain compatibility between any of the various ways to describe them. This is the fundamental principle behind common data modelling. If modelling method “A” describes an object and modelling method “B” describes the same object, there must be a mapping between “A” and “B.” Similarly, if model “C” describes the same object, there must also be a mapping between “A” and “C,” and by the transitive property between “B” and “C.”

Of course, another problem between standards is what they choose to model. Some standards may only model home appliances while others may concentrate on healthcare devices. It’s hard to map a digital blood pressure monitor to a stereo system.

One way to overcome this problem is to map things at a more atomic level. Both a blood pressure monitor and a stereo system have a switch and a display. If you take all of the atomic components of each system and map them to a common model, you can get some basic interoperability even between items as different as a blood pressure monitor and stereo system.

Establishing a common data model that is a superset of all the atomic resource that are used to compose the devices of interest is a logical way to create general compatibility. Each complex device is just a collection of all the atomic resources it contains. Similar items (like refrigerators) from different manufacturers may have different features, but will still share a majority of common atomic resources like lights, thermostats, switches, etc.

By defining mappings between device representations in any standard and the common data model, a very complete interoperability model can be established.

2. RESTful Architecture

Once a common data model has been set up, the interactions with the data model must be defined. This is where RESTful architecture comes into play. An extended RESTful model can be defined by the following actions: Create, Read, Update, Delete and Notify (CRUDN). In other words, device models can be created and deleted. Their state can be read and set (updated) and devices can send notification if an anticipated event occurs. A very large number of real-world devices can be defined, observed and operated using only these principles.

Table 1 - CRUDN view of a RESTful interface

CREATE	A resource must be created before it can be used. PUT is normally used.
READ	Read allows the current values of the resource to be determined. Get is normally used.
UPDATE	Update allows the current values of the resource to be set. Generally, POST is used so the elements of the structure can be written selectively.
DELETE	Delete is used if the resource is no longer needed.
NOTIFY	Notify is used to get information initiated by the server. It is usually communicated using an OBSERVE function or a Publish/Subscribe method.

3. Security

One of the biggest fears generated by the Internet of Things is security. If hackers can already sabotage your hard drive or steal your passwords, imagine what damage they could do if they had access to the door locks on your house or could control your oven. This is why a security strategy must be built-in to any IoT ecosystem from the start. The security system must not only employ the latest security technologies, but must also anticipate that it will be hacked and have a plan for how to survive the hack and update its own systems. This sort of security can not succeed if it is “added” after the system already works. It must be designed in from the start.

4. Bridging

Standards succeed when they are agreed upon by the key implementers in any ecosystem. This is why cellular phones can call other cellular phones (it wasn’t always that way). One way to make sure

everything works together is to get all the key players to choose the same solution. While attractive, this rarely works quickly or without several standards failing the test of time.

Another way to do this is to define bridges between standards. By using the common data model described in section 1, a mapping can be described that maps another standard into the common data model and out of the common data model. If several standards define this two-way mapping, they can all communicate by transitioning through the common data model. A bridge is a device that implements these two-way mappings. By using a bridge, a standard can take advantage of the common data model without the need to adopt it or convert millions of deployed devices to use it.

Organizational Pillars

The architecture described in the sections above provides the theoretical construct for interoperability, but a standard is always open to interpretation. Even a well written standard can be understood differently by different brilliant engineers. In order to ensure interoperability, more is needed. Interoperability can be better enabled by building an ecosystem with three pillars.

5. Open Standard

The first pillar is an open standard. A standard defines as unambiguously as possible all of the details for implementing the architecture. Standards use very precise language and lots of coding examples. An open standard is developed in view of many critics and reviewers. Eventually, it is made public so anyone can access it for review, implementation and improvement. An open standard also generally has very generous usage terms including minimally restricted use, fair and often free licensing terms and a regular process for finding and correcting errors. An open standard is a key element of a maximally interoperable system.

6. Open Source Implementation

Having an open standard is a good start, but relatively few people have the patience or motivation to read a standard and write compliant code for it from scratch. That's one reason it's important to have an open source implementation. With open source, the code that implements the standard is created, tested and shared by a community. By starting with a solid code base, individual developers can begin to implement a new product with basic code that already works. They just need to add the features that are required by their new product. With many people reviewing and using the open source code on a regular basis, errors are more easily discovered and corrected. The community works together to maintain and improve the code base to the benefit of all.

7. Certification Test Tool

The final organizational pillar is the certification test tool (CTT). The CTT tests the code to make sure it is compliant with all the requirements laid out in the open specification. By automatically testing each specification requirement, any implementation (private or open source) can be verified to be a valid implementation of the specification. If a device passes the tests in the test tool, it should logically work with other devices that pass the test tool.

By using each of the three organizational pillars, a standard that enables interoperability theoretically, can test a “canonical” open source implementation and verify it is compliant with the specification. The CTT

can also test other private implementations. The combination of open source, open standard and a CTT provides multiple checks. When errors are discovered, they can be tracked to errors in the implementation, errors in the CTT, or even errors in the specification.

Additionally, “plug fests” can be held where products that are expected to be interoperable by virtue of passing the CTT can actually be tested with other products that also pass the CTT. This is a nice final test to verify interoperability.

Tools and Support

The features described to this point help to ensure that interoperability and security can be developed across multiple products and ecosystems. However, if it is too difficult to build these interoperable products, nobody will build them. This is why it is important to develop tools and a support system to help product developers.

8. Crowd-sourcing Data Models

One of the best ways to help developers is to provide a complete resource repository of atomic data models. While this can be done in a number of ways, one of the most efficient is to have a common place where anyone interested can create a model and submit it. Of course, this would quickly become unwieldy without a bit of control.

One way to manage this would be to have an online tool where users could log in and make submissions. If these submissions could then be reviewed by a team of experts before being accepted, consistency could be assured and duplication could be avoided.

If this tool also allowed for submission of mappings between different ecosystems and the common data model, there would be a single authoritative source for data model resource interoperability. Furthermore, the accuracy and completeness of these models would be encouraged by support of the community that would use them.

9. Support for Multiple Platforms

One of the real benefits of the Internet of Things is the vast variety of devices that can be built on numerous ecosystem platforms. For some products, a minimal platform is required that can run on a button-cell battery for several years. For other applications, a more capable platform is required that can process video signals or respond instantaneously.

These platforms are not going to all support the same operating system or state machine. They won't support the same programming languages. This is why it is important to have an architecture and implementation that can support a wide variety of platforms. It is also important to maintain a number of these platforms to give programmers and product developers a head start.

10. Tutorials and Developer Support

Another critical tool for developers is an example. A working example is far more useful than several pages of specification text. Moreover, a tutorial that takes a developer step-by-step through the development process is much more valuable than a simple example.

With the right development tools and a good step-by-step example, a programmer can be productively programming in a couple of hours rather than a couple of weeks.

Open Connectivity Foundation

Now that we've laid out the basic requirements for a secure and scalable IoT system, let's look at the Open Connectivity Foundation (OCF) to see how well it stacks up to the principles for interoperability that have been described. The OCF is an open standards organization with over 300 members worldwide.

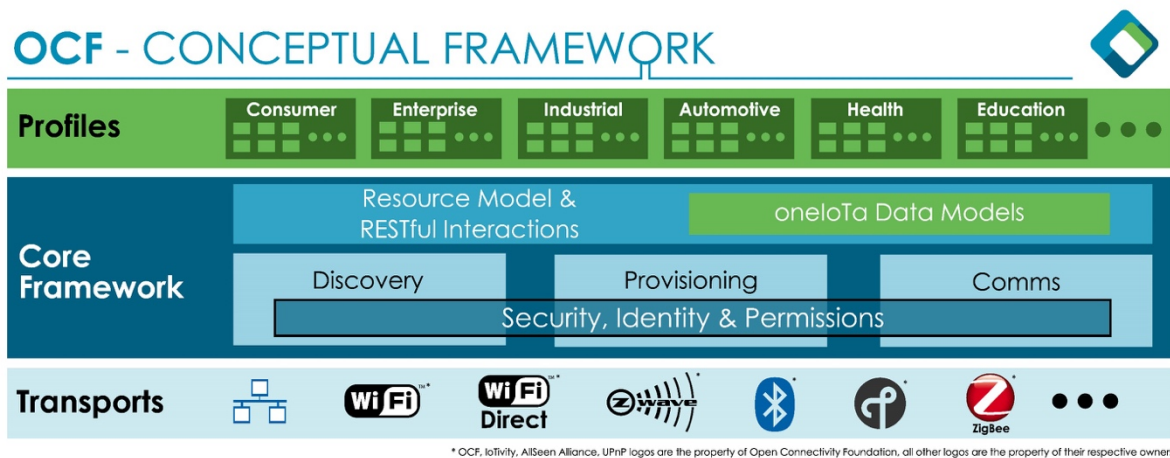


Figure 1 - OCF conceptual framework

11. Common Data Model & RESTful Architecture

OCF uses a common data model comprised of atomic resources (currently over 100) defined using JSON schema for payloads and Swagger files describing a RESTful interface. OCF uses a CRUDN set of actions on the atomic resources and constructs complete devices as collections of resources. The interface for a device is completely described by the device data model. A “client” or control point interface can operate a device by introspecting the device description and using the CRUDN interface on the device data model.

12. Security from the Start

Robust communication security is designed into OCF from the start. At its most secure, OCF uses a public-key infrastructure with credentials installed into the device at manufacture time. Additionally, OCF uses link-based security. Insecure ecosystems can interface with OCF, but capabilities of these insecure ecosystem will be limited.

13. oneloTa and Derived Models

oneIoTa (oneIoTa.org) is an online tool for crowd-sourcing resource data models. It is a basic Integrated Development Environment (IDE) that includes back-end processes for submitting and reviewing data

models with groups of experts. oneIoTa provides the definitive collection of OCF data models that comprise the common data model. It also contains resources from other organizations like AllJoyn and UPnP. Other organizations are encouraged to contribute their data models.

oneIoTa is also the repository and tool for “derived” data models that describe the mapping between other ecosystems and the OCF common data model.




Figure 2 - Three pillars of OCF

14. IoTivity

IoTivity is the official open source implementation of the OCF standard. At each release event, IoTivity will provide a complete implementation of OCF that passes the certification test tool. IoTivity is managed by the Linux Foundation and funded by OCF. In addition to a complete implementation of OCF, IoTivity includes a number of examples and implementations for several platforms.

15. Certification Test Tool and Authorized Test Facilities

The OCF certification test tool (CTT) is an automated test tool that implements a complete test of OCF at each release event. CTT is under continual development. IoTivity must pass the CTT at each release event. CTT is also the tool that is used by OCF authorized test facilities to certify real products.

16. Developer Community and Tools

OCF has a marketing group that provides resources and instruction events to encourage development and support for devices based on OCF. The Tools Task Group in OCF will develop and distribute tools to assist developers in the creation of OCF products. OCF also encourages independent groups with different interests to support OCF. Implementations for particular ecosystems, platforms and applications adds to the diversity of the development community.

17. Current Status

OCF has not suddenly solved IoT interoperability, but progress is being made. OCF currently has over 350 members and liaison relationships with about 20 other standards organizations. OCF version 1.0 will be publicly released this fall. Each OCF biannual release will include synchronization between the open standard, the open source implementation (IoTivity), and the certification test tool. There is a well-defined process for continual introduction of new vertical markets and associated use cases. Mainstream devices that are OCF native will start showing up this Christmas along with bridges to existing ecosystems. There is still a long way to go, but progress is being made.

Conclusion

In order to get the full benefit of the Internet of Things, a maximal set of things must work together regardless of development platform or underlying ecosystem. The principles described in this paper can help deliver on the promise of IoT by creating a scalable system that can describe virtually any product and interoperate with virtually any other ecosystem. Furthermore, the three pillars of an open standard, an open source implementation and a certification test tool ensure that the implementations of OCF meet the expectations of the theory. Finally, the support community around OCF aims to ensure that OCF-based products can be successfully implemented as easily as possible.

Creating an Internet of Things that parallels the opportunity of the Internet is no small task and the success of OCF is in no way guaranteed. However, by using the interoperability principles described in this paper the chances of success for OCF are improved substantially.

Abbreviations

CTT	Certification Test Tool
OCF	Open Connectivity Foundation
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IoT	Internet of Things
ISBE	International Society of Broadband Experts
ISO/IEC	International Standards Organization/International Electrotechnical Commission
SCTE	Society of Cable Telecommunications Engineers
W3C	World Wide Web Consortium

High Dynamic Range for HD and Adaptive Bitrate Streaming

A Technical Paper prepared for SCTE•ISBE by

Sean T. McCarthy, Ph.D.
Independent Consultant
Sean McCarthy, Ph.D. Consulting
236 West Portal Avenue, #293
San Francisco, CA 94127
415-518-5287
sean.mccarthy@comcast.net

Introduction

When High Dynamic Range (HDR) began to emerge a few years ago as a great new television experience, it was almost invariably thought of as only one part of Ultra HD/4k along with a wider range of colors and higher bit depths for video coding. Now though, many are wondering if HDR can be an independent contributor that makes for better TV experiences even without higher 4k resolutions. Indeed, several MVPDs, broadcasters and industry associations are moving towards enhanced programming and distribution that marries HDR with HD and sub-HD resolutions. The next-generation television standard, ATSC 3.0, for example, enables distribution of HDR and Wide Color Gamut (WCG) content at any resolution. Similarly, leading internet-based streaming services are leveraging multi-resolution adaptive bitrate (ABR) protocols to deliver HDR experiences.

The potential wrinkle is that HDR is still often tied to Ultra HD/4k/10-bit in the content creation studios and production houses. The HDR highlights that make HDR shine are often very small and localized. Thus, converting original Ultra HD HDR content to lower resolutions for HD and ABR streaming runs the risk of obliterating those visually potent HDR highlights. When the resulting video is compressed with High-Efficiency Video Coding (HEVC), HDR distortions could be made worse, particularly for aggressive streaming profiles.

In this paper, we use a recently developed method for measuring distortion in HDR video to quantify the impact of encoding original Ultra HD HDR content at HD and ABR resolutions. Our method is intrinsically independent of any particular HDR transfer function and may thus be applied to HDR content having Hybrid-Log Gamma (HLG), Perceptual Quantizer (PQ), or other transfer characteristic. Specifically, we show:

- 1) That each frame of HDR video can be represented as the sum of a Spatial Detail signal and a Foundation Image. The Spatial Detail signal contains the localized contrast variations associated with HDR features. The Foundation Image contains the more smoothly varying large-area contrast variations associated with textures and backgrounds.
- 2) The Spatial Detail signal is the main contributor to HDR distortions introduced by rescaling and compression. The Foundation Image contributes little to total HDR distortion.
- 3) Quantifying the correlation between the Spatial Detail signal of processed images and corresponding original images provides a systematic method to choose the best combinations of HEVC-encoded resolutions and bitrates to minimize overall HDR distortion.

Our key objectives in this presentation is to provide a practical method that can help ensure great HDR experiences at any resolution.

Content

1. Background

IP-based protocols and streaming are rapidly becoming powerful methods for distributing television to all screens from the big screens in living rooms to smaller-screen smartphones and tablets. At the same time, displays big and small are becoming much more capable of rendering the deep darks and bright highlights that make HDR^{1,2} special.

Ideally, content producers and distributors would like HD variants of Ultra HD/4k HDR to differ as little as possible from the original full-resolution content. We would also like to create a consistent HDR experience across all screens, big and small. A challenge is that small screens tend to have lower resolution than big screens. They also tend to be used more often at the edge of lower-bandwidth wireless and more congested networks.

To create high-quality consistent HDR experiences, we need a method to quantify the differences between studio-approved full-resolution HDR content and corresponding HD and sub-HD variants that might be distributed over DOCSIS and adaptive streaming protocols that switch between encoded resolutions as network conditions vary.

Service providers choose specific combinations of encoded resolution and bitrate to enable best-quality IP and adaptive streaming. The set of combinations is called an adaptation set, or sometimes called an encoding ladder. In adaptive streaming, client players choose the specific resolution-bitrate pair available in the adaptation set based on the clients' available bandwidth and other performance criteria.

The key questions for service providers are the following. Which combinations of resolution and bitrate should be included in HDR adaptation sets to minimize HDR distortions and promote consistency across screens? Which combinations minimize the visibility of switching between the rungs of the encoding ladder? This paper provides methods to help answer these questions.

2. Test Sequences & Preparation

In this study, we used the Meridian³ HDR test content as shown in Figure 1. Meridian is Ultra HD HDR 3840x2160 60 fps test content graded to 4000 cd/m² with transfer characteristics and color primaries specified by ITU-R BT.709⁴. (See the note below regarding BT.709 HDR transfer characteristics.) Meridian was professionally produced by Netflix and is publicly available for download⁵ in Interoperable Master Format (IMF) (see SMPTE OV 2067⁶) for which the video essence is an MXF⁷ file containing JPEG2000⁸-encoded 10-bit YCbCr 4:2:2 video data. To perform the calculations presented in this paper, we extracted data for individual frames using ffmpeg⁹.

(Note on ITU-R BT.709 transfer characteristics for Meridian. HDR content is typically encoded using either HLG, an Opto-Electronic Transfer Function (OETF) specified in ITU-R BT.2100¹⁰, or PQ, an Electro-Optical Transfer Function (EOTF) specified in SMPTE ST 2084¹¹. Thus, it is perhaps surprising that Meridian makes use of ITU-R BT.709, an OETF originally intended for High-Definition TV (HDTV) and which was standardized before the emergence of HDR. For the purposes of this paper, it is not an issue. Our methodology is independent of HDR transfer characteristics, though the results of our method could be expected to provide tighter correlation with subjective video quality when applied to PQ- or HLG-encoded video.

It may be that BT.709 was used because Meridian was produced before HDR-support in IMF was standardized. The availability of Meridian was announced³ in September 2016 and IMF did not include HDR support until shortly before that with the publication of SMPTE ST 2067-21¹² in July 2016.)



Figure 1 - Ultra HD HDR Test Sequence Used in this Study



Figure 2 - HD HDR Test Sequences Used in this Study

There is a dearth of contribution-quality 4k HDR content available for testing. Thus, we also used the HD-HDR WCG test sequences shown in Figure 2. These sequences were created by the “HdM-HDR-2014 Project”^{13,14} to provide professional quality cinematic wide gamut HDR video for the evaluation of tone mapping operators and HDR displays. All clips are 1920x1080p24 and colour graded for ITU-R BT.2020¹⁵ primaries & 0.005-4000 cd/m² luminance. We converted the original colour graded frames (RGB 48 bits per pixel TIFF files) to Perceptual Quantizer (PQ) YCbCr v210¹⁶ format (4:2:2 10 bit) using the equations defined in ITU-R BT.2020, ITU-R BT.1886¹⁷, and ITU-R BT.2100.

For each HdM-HDR-2014 test sequence, we created 50 variants having different encoded resolutions and bitrates. Of the 50 variants, 20 were raw uncompressed versions used to isolate the impact of different rescaling algorithms on HDR distortion. The remaining 30 variants for each test sequence were compressed using HEVC for each combination of encoded resolution (1920x1080, 1440x1080, 1280x720, 960x540, 720x540, and 640x360) and bitrate (10000, 3000, 1000, 300, and 100 kbps).

All rescaling was performed in Matlab¹⁸ using the *imresize* function. All compression was performed using command-line x265¹⁹ (main10 profile).

3. Decomposing an HDR Frame into Spatial Detail and Foundation Images

Our approach to measuring HDR distortion begins with decomposing each HDR frame (A) into a Spatial Detail signal (B) and a Foundation Image (C), as illustrated in Figure 3. The pixel-by-pixel sum of the Spatial Detail signal and Foundation Image is equal to the original image.




Figure 3 - Decomposition of an HDR frame into Spatial Detail and Foundation Image

The method of calculating the Spatial Detail signal is described in the Appendix and detailed more thoroughly in previous publications²⁰⁻²³. In brief, the Spatial Detail signal can be thought of as the condensed essence of the original image. It isolates the features and details that are unique to the original while minimizing statistically expectable characteristics that the original image shares with images as a statistical class.

Images of natural and other complex scenes have an interesting statistical property: They have spatial-frequency magnitude spectra that tend to fall off with increasing spatial frequency in proportion to the inverse of spatial frequency²⁴. The magnitude spectra of individual images can vary significantly; but, as an ensemble-average statistical expectation, it can be said that “the magnitude spectra of images of natural and other complex scenes fall off as one-over-spatial-frequency.”

The Spatial Detail signal is effectively the result of de-emphasizing the statistically expectable one-over-frequency characteristic. As such, the Spatial Detail signal emphasizes the unique unexpected details in an image.

The Foundation Image is obtained by simply subtracting the Spatial Detail image from the original HDR image. As such, the Foundation Image may be thought of as a special kind of low-pass filtered version of the original HDR image in which the unique spatial details are selectively attenuated. Although perhaps difficult to appreciate on the printed page, the Foundation Image gives the visual sensation of being out of focus.

One way to think of the Foundation Image is that it represents the overall contrast and luminance range of the HDR image, whereas the Spatial Detail signal can be thought of as representing localized contrast and luminance variations.

The concept of decomposing an image into coarse and fine components is not new. Indeed, it is the basis for scalable video encoding schemes such as Scalable High Efficiency Video Coding²⁵ (SHVC) that is included in ATSC 3.0²⁶. Wavelet-based video compression, such as JEG2000, also relies on decomposition.

We use decomposition for a different reason in this paper. SHVC, for example, provides for a base layer that that can be decoded and displayed. SHVC also provides enhancement layers that can be added to the base layer to improve quality and detail. The amount of enhancement that can be achieved depends on bandwidth availability and the capabilities of the decoder and display. SHVC is thus a scheme to optimize

the quality and detail of a displayed video for various network and user-device situations. JPEG2000 combines layers of differing detail to achieve target bitrates and visual quality.

We do not intend the Foundation Image to ever be displayed, nor should it be thought of as the lowest-resolution variant in an encoding ladder. Rather, our decomposition into a Spatial Detail signal and Foundation Image is strictly a mathematical way of concentrating most of the distortions introduced in processing HDR video into a more concise signal than the original HDR video. The resulting concentration of visual detail provides a useful way to quantify HDR distortion, as described below.

4. Spatial Detail Signal as a Guide to Features and Textures

The histogram of the Spatial Detail signal is shown in Figure 4. A special property of the Spatial Detail signal is that the shape of its histogram, which resembles a two-sided exponential distribution, is preserved across images. For example, the Spatial Detail signals for the different HdM-HDR-2014 test sequence all have Spatial Detail histograms that resemble two-sided exponential distributions. Like scale-invariance, the shape of the histogram appears to be an expectable statistic of natural and complex images.

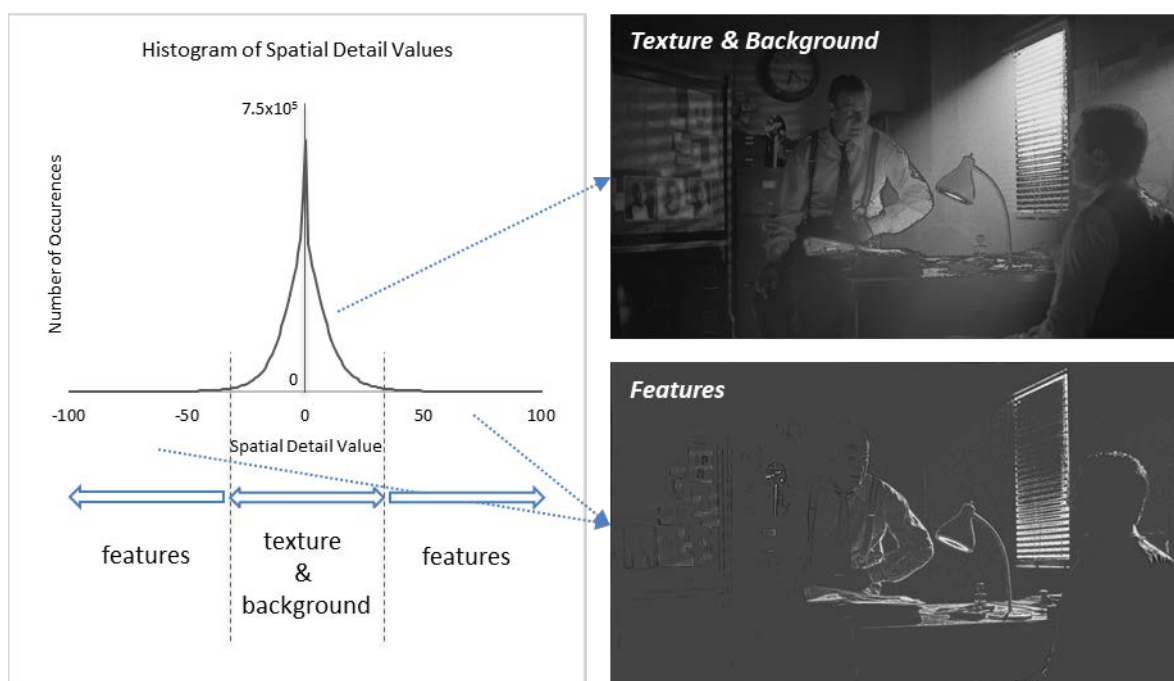


Figure 4 - Spatial Detail Signal as a Guide to Features, Textures, and Background

Another special property of the Spatial Detail signal is that its values provide a convenient guide to the parts of an image that would tend to be called features and the parts that would tend to be called textures and background. Large absolute values of the Spatial Detail signal tend to correlate with features. Small absolute values tend to correlate with textures and background. The images in Figure 4 were created by creating a masking image by applying a threshold to the absolute value of the Spatial Detail signal, and then applying the mask to the original luma values of the HDR frame data. The “Features” image are the original luma values for every pixel at which the absolute values of the Spatial Detail signal exceeded the applied threshold. The “Textures & Background” image is the complement of the “Features” image.

5. Spatial Detail Distortion is Most of the Total HDR Distortion (I)

An advantage of decomposing each HDR frame into a Spatial Detail signal and Foundation Image can be appreciated by examining the relative contribution of each to total HDR distortion.

Mean-Squared Error (MSE) and Peak-Signal-to-Noise Ratio (PSNR) are common and equivalent metrics of video distortion^{27,28}. (PSNR is proportional to the logarithm of MSE.) MSE is the average over all pixels of the squared difference between an original image and a corresponding encoded variant.




Figure 5 - Relative Contribution of the Spatial Detail Signal and Foundation Image to Total MSE

MSE values calculated for four encoded resolutions and four candidate rescaling algorithms are shown in Figure 5. The candidate rescaling algorithms are nearest-neighbour interpolation, bilinear interpolation, bicubic interpolation, and lanczos3 resampling²⁹. The MSE values shown in Figure 5 are for “Meridian”. Lower MSE values indicate that the rescaled variant is less distorted from the original in terms of squared-error (and thus PSNR). The data show that lanczos3 resampling provides the lowest MSE values for all resolutions and should thus be considered the best choice in constructing adaptive streaming variants. If other considerations such as processing demands are significant, bicubic interpolation can deliver nearly as good results.

Total MSE is the mathematical sum of the sum of the Spatial Detail MSE by itself, the Foundation Image MSE by itself, and a contribution from the covariance of Spatial Detail signal and the Foundation Image.

The data in Figure 5B show that the Foundation Image MSE is a small fraction of the total MSE. The Foundation Image contribution to total MSE increases with progressively more aggressive downscaling, which indicates that rescaling progressively distorts the underlying smoothly-varying luminance of the encoded video. Yet, even for the 4-fold downscaling from the original 3840x2160 to 960x540, the error associated with the Foundation Image is only 10-15% of the total error. Most of the total error is attributable to distortion of the Spatial Detail signal (approximately 75% for 1920x1080, ~65% for 1280x720, and ~55% for 960x540. For 640x360 – a 6-fold downscaling -- no single portion of MSE is

the majority though the Spatial Detail contribution remains largest. (The data shown in Figure 3B are for lanczos rescaling.)

6. Spatial Detail Distortion is Most of the Total HDR Distortion (II)

Figure 6 and Figure 7 show another way of visualizing distortions introduced by processing HDR content. The data in Figure 6A show the original luma code values (horizontal axis) plotted against average encoded luma code values (vertical axis). The data are for the 3840x2160 Meridian HDR test content rescaled at four sub-4k resolutions (1920x1080, 1280x720, 960x540, and 640x320 using lanczos3 resampling). The rescaled luma values display an almost perfect linear correlation with the original values: The encoded values fall along a straight line with a slope of 1. The exceptions, in this example, are for bright regions having code values above ~900 (dashed square) and dark regions having code values less than ~75 (dashed circle with arrow) within the valid 10-bit luma range of 64 to 940. (For some other test content, not shown, the elevation of dark luma values is more significant). The data in the bright range of Figure 6A is shown close-up in Figure 6B. For the Meridian Ultra HD HDR test content, rescaling results in a progressively decreasing luma values with increasingly aggressive downscaling.




Figure 6 - Correlation Between Original and Rescaled Luma Values

The distortions of very dark and very bright regions appear to be a result of spatial averaging during resizing. For a finite range of allowed luma values, values near the limits of the range would tend to be averaged with neighbour values nearer to the center of the range, particularly for small isolated bright and dark regions. Thus, extreme values would tend to be pulled away from the upper and lower limits towards more central values during the resizing process. The data shown in Figure 6 indicate that there is room to develop HDR-sensitive resizing algorithms that are better than even lanczos3 resampling for use in multi-resolution HDR video services.

(A note on calculating average code values – Each average code values in the rescaled and/or encoded frame was calculated by finding all the pixel locations in the original frame that have a specific code value, for example, 312 out the possible range of 64 to 940 for 10-bit encoding. The encoded code values

for the corresponding pixel locations tend to have a distribution of values because of scaling and/or compression. We use the average over the distribution. Thus, the data in Figure 6, and other similar figures in this paper, provide a view of the correlation between expected values (original value) and the corresponding average code value in the processed image.)

The data plotted in Figure 7A & B provide more details on the nature of the luma distortion evident in Figure 6. In Figure 7A, there is negligible apparent distortion of the Foundation Image component of the rescaled image compared to the Foundation Image component of the original. (Data for all rescaled resolutions are plotted in Figure 7A. The data lay on top of each other thus giving the appearance of a single line.) On the other hand, the Spatial Detail signal of the rescaled image (Figure 7B) is significantly different from the Spatial Detail signal of the original. If there were no Spatial Detail distortion, the data would lay along the dashed line having unity slope. Thus, the conclusion is that luma distortion evident in Figure 6 is mainly a result of distortion of the Spatial Detail signal with hardly any contribution from the Foundation Image component.

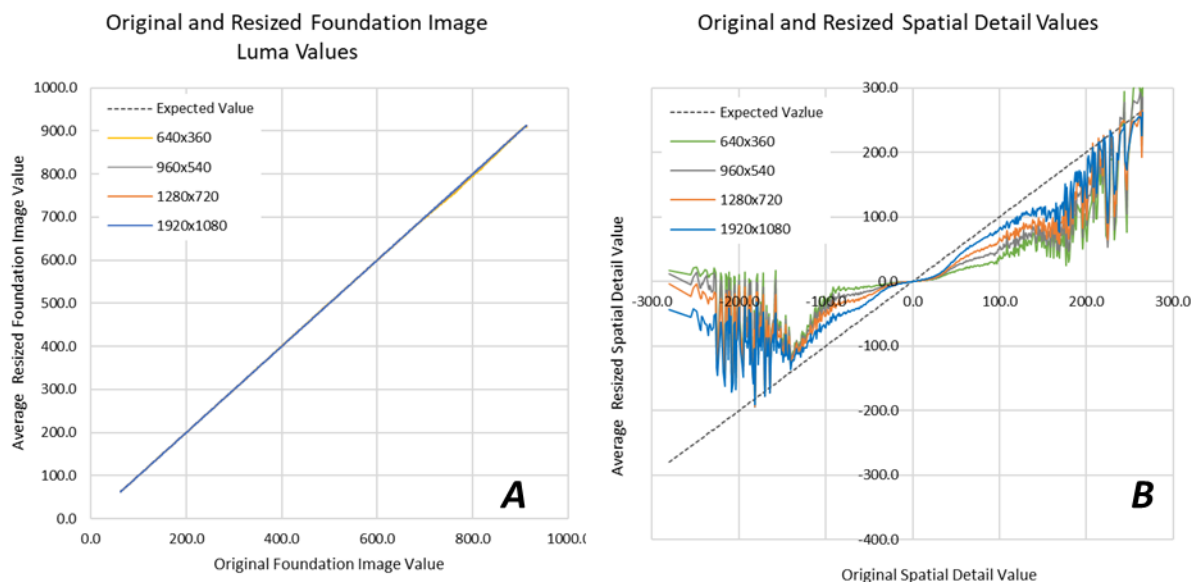


Figure 7 - Correlations for Rescaled Foundation Image and Spatial Detail Values

The distortion of the Spatial Detail component has several interesting features that provide insight into the effects of rescaling.

First, Spatial Detail signals of the rescaled images tend to follow a progressively shallower slope with increasingly aggressive downscaling. This indicates a progress systematic loss of local contrast in the HDR image. That is, the range of Spatial Detail values becomes narrower with rescaling. (The histogram of the Spatial Detail signal of rescaled images becomes narrower and more peaked.)

Second, the slope of the Spatial Detail signal is near zero near the origin of Figure 7B. This is the range in which the original Spatial Detail values have small absolute values. In other words, this is the range that tends to identify the regions in the original luma image that can be thought of as texture & background (see Figure 4). The conclusion is that rescaling more severely smooths low-amplitude textures and film grain and has a relatively lesser impact on larger-amplitude features.

Third, the average values of Spatial Detail signal of the rescaled image become progressively noisier and decorrelated with the value of the Spatial Detail signal of the original as the absolute value of the original Spatial Detail increases. We find that the decorrelation is content- and compression-dependent. Thus, quantifying the decorrelation can be a useful metric, as we show below.

7. Choosing Best Combinations of Resolution & Bitrate

HEVC compression introduces its own distortions beyond those introduced by rescaling. Total distortion depends on the interaction of rescaling and bitrate. Reducing encoded resolution at an encoded bitrate can actually reduce total distortion and increase video quality; but only to a point beyond which further reductions in encoded resolution increases total distortion. Minimizing distortion is a matter of finding the best balance between encoded resolution and bitrate.

In this study, we quantify HDR distortion by measuring the correlations between processed and original Spatial Detail signals for many bitrate-resolution combinations. Recall that the Spatial Detail signal represents most of the total distortion.

We measure correlation with the coefficient of determination, R^2 , the square of the Pearson correlation coefficient, R , (see ref. 30). R^2 quantifies the “goodness of fit” of data to a linear regression line. The expectation is that the average code value of the encoded frame would be the same as the corresponding code value of the original frame. Mismatches are a manifestation of a lack of correlation and result in a lower value of R^2 , which has a range of 0 to 1.




Figure 8 - Spatial DetailCorrelations for HEVC Compressed and Rescaled HDR Content

The data plotted in in Figure 8 illustrate Spatial Detail correlations for HEVC-compressed versions of the HdM-HDR-2014 test content. Each test sequence was encoded at 1920x1080 resolution and bitrates from

100 kbps to 10 Mbps. Though not shown, we also encoded the test clips at the same bitrates for each downsampled resolution 1440x1080, 1280x720, 960x540, 720x540, and 640x360. We used the HD-HDR test sequences because they comprise a richer more complete test set, and thus more reliable, than the single 4k-HDR example of the Meridian test sequence.

In Figure 8, there is a general overall linear relationship between encoded and original Spatial Detail, but there are three significant differences worth noting. First, the magnitude of the variation around the fitted straight-line (dashed) increases with decreasing bitrate. This indicates that Spatial Detail correlation decreases with decreasing bitrate. Second, the slope of the fitted straight-line decreases with decreasing bitrate. This indicates that localized contrasts of textures and HDR highlights are diminished with decreasing bitrate. Third, the slope of the correlation is flatter near the origin (small Spatial Detail values) than for large Spatial Detail values. This indicates that low contrast textures and details (such as those typically associated with faces and background textures) are systematically impacted more severely by HEVC compression than are high contrast HDR textures.

(Although not the main purpose of this paper, it is worth pointing out that Spatial Detail plots such as Figure 8 could be a useful engineering tool. Such plots provide a quantitative view of how compression and processing algorithms differentially affect textures, features, noise, and overall Spatial Detail contrast. Video encoder developers and video compressionists could use such information to optimize algorithms and parameter selection to improve visual quality and compression efficiency.)

We calculated an R^2 value for each of the 30 permutations of HEVC bitrate and encoded resolution for each of the 5 test sequences. We analysed the resulting R^2 values in the spirit of ATIS-0800061³¹, “Methodology for Subjective or Objective Video Quality Assessment in Multiple Bit Rate Adaptive Streaming.”

Table 1 summarizes the goodness-of-fit, R^2 values, for the bitrate-resolution combinations used in this study. The cells highlighted in green indicate which encoded resolution maximizes the goodness-of-fit for each bitrate. In other words, the green-highlighted cells correspond to the bitrate and resolution combinations that maximize the correlation between encoded Spatial Detail and original Spatial Detail. These are the bitrate-resolution combinations that best preserve the textures, highlights, and local contrast variations in HDR video. Higher resolutions at lower bitrates introduce compression distortion. Lower resolutions at higher bitrates introduce rescaling distortion. Our quantification of Spatial Detail correlation highlights the bitrate-resolution combinations that provide the best balance.

Table 1 – Choosing Bitrate & Resolution Combinations based on Spatial Detail Correlation

Resolution	Bitrate (kbps)				
	100	300	1000	3000	10000
1920x1080	0.900	0.933	0.946	0.966	0.983
1440x1080	0.893	0.939	0.949	0.968	0.982
1280x720	0.919	0.942	0.955	0.965	0.976
960x540	0.903	0.934	0.946	0.955	0.962
720x540	0.892	0.924	0.935	0.943	0.950
640x360	0.869	0.897	0.902	0.909	0.914

The data in Table 1 are the average over all HdM-HDR-2014 test sequences and thus represent a compromise trade-off for any particular test sequence. The compromise is arguably unfair at low bit

rates. As illustrated in Table 2, some test sequences benefit from lower-resolution encoding at low bit rates whereas others benefit from higher-resolution encoding.

Table 2 illustrates the use of Spatial Detail correlation to choose the best variants to include in HDR adaptation in a content-aware manner, also called per-title encoding. Note that the *carousel_fireworks* test sequence benefits from lower resolution variants than does the *bistro* test sequence. The *carousel_fireworks* test sequence is more challenging because it has a wider range of luminance and more motion. From a practical standpoint, Table 2 is a guide to setting encoder parameters to produce the best content-specific per-title encoding ladders.

Table 2 – Choosing Content-Dependent Bitrate & Resolution Combinations

Resolution	bistro					carousel_fireworks				
	100	300	1000	3000	10000	100	300	1000	3000	10000
1920x1080	0.958	0.987	0.995	0.997	0.999	0.841	0.848	0.878	0.942	0.985
1440x1080	0.956	0.990	0.995	0.997	0.998	0.848	0.879	0.898	0.956	0.989
1280x720	0.979	0.990	0.993	0.995	0.996	0.841	0.901	0.932	0.966	0.986
960x540	0.963	0.985	0.987	0.988	0.988	0.860	0.906	0.927	0.948	0.968
720x540	0.966	0.983	0.988	0.988	0.989	0.866	0.916	0.931	0.946	0.960
640x360	0.950	0.966	0.967	0.967	0.968	0.869	0.912	0.920	0.927	0.929

Conclusion

In this paper, we provided methods to quantify HDR distortions and take steps to mitigate those distortions by selecting the best combinations of bitrates and resolutions to include in HDR adaptation sets used in adaptive streaming services.

A key part of our approach is to decompose HDR video into two components for the purpose of mathematical analysis. A Foundation Image contains the overall luminance and contrast variations in HDR video. A Spatial Detail signal contains the localized luminance and contrast variations.

We showed that most of the distortion in encoded HDR content is a result of distortions of the Spatial Detail signal.

We proposed that the correlation between the encoded and original Spatial Detail is a useful metric by which to design encoding ladders; i.e., the best combinations of bitrate and resolution for HDR television services. This approach maximizes the similarity of the local contrasts and highlights between encoded and original HDR videos. These local contrasts and highlights add significant impact to HDR video and represent a large part of the content creator's original intent. Significantly, the use of Spatial Detail correlation can be used to design content-aware per-title encoding ladders as well as overall generic encoding ladders.


We also provided data that indicate that plots of Spatial Detail correlation could be a useful engineering tool. Such plots provide insight into the relative impact of compression and processing on different characteristics of HDR video, such as: noise, textures, and features.

In addition, we provided data that indicates that there is room to develop better HDR-sensitive resizing algorithms. The data show systematic distortions in very dark and very bright regions. New luminance-adaptive resizing algorithms might be worth investigating to mitigate such distortions.

Our next steps are to evaluate our proposed methods on a larger set of HDR data. Current publicly available HDR test material is limited and was produced before the latest HDR international standards were published. Our methodology is independent of any particular HDR transfer characteristics. Nonetheless, the application of our methodology to PQ and HLG HDR transfer characteristics might provide better or worse agreement with human opinions of video quality. This is the topic for a follow-on paper.

Appendix

The method of creating the Spatial Detail signal can perhaps best be understood by thinking of an image in terms of spatial frequency spectra as illustrated in Figure 9 (only the luma channel is shown). Any 2-dimensional array of pixel values can also be represented without loss of information as the product of a magnitude spectrum and a phase spectrum in 2-dimensional spatial frequency space. Spatial-frequency spectra can be obtained from an image pixel array by performing a 2-dimensional Fast Fourier Transform (FFT2). The pixel array can be recovered by performing a 2-dimensional Inverse Fast Fourier Transform (IFFT2). FFT2 and IFFT2 are well known signal processing operations that can be calculated quickly in modern processors.



$$FFT2(Y(x, y)) = \mathbb{Y}(k_x, k_y) = |\mathbb{Y}(k_x, k_y)| * \exp(i\theta_Y(k_x, k_y))$$

Figure 9 - Representation of a Video Frame in Terms of Spatial Frequency

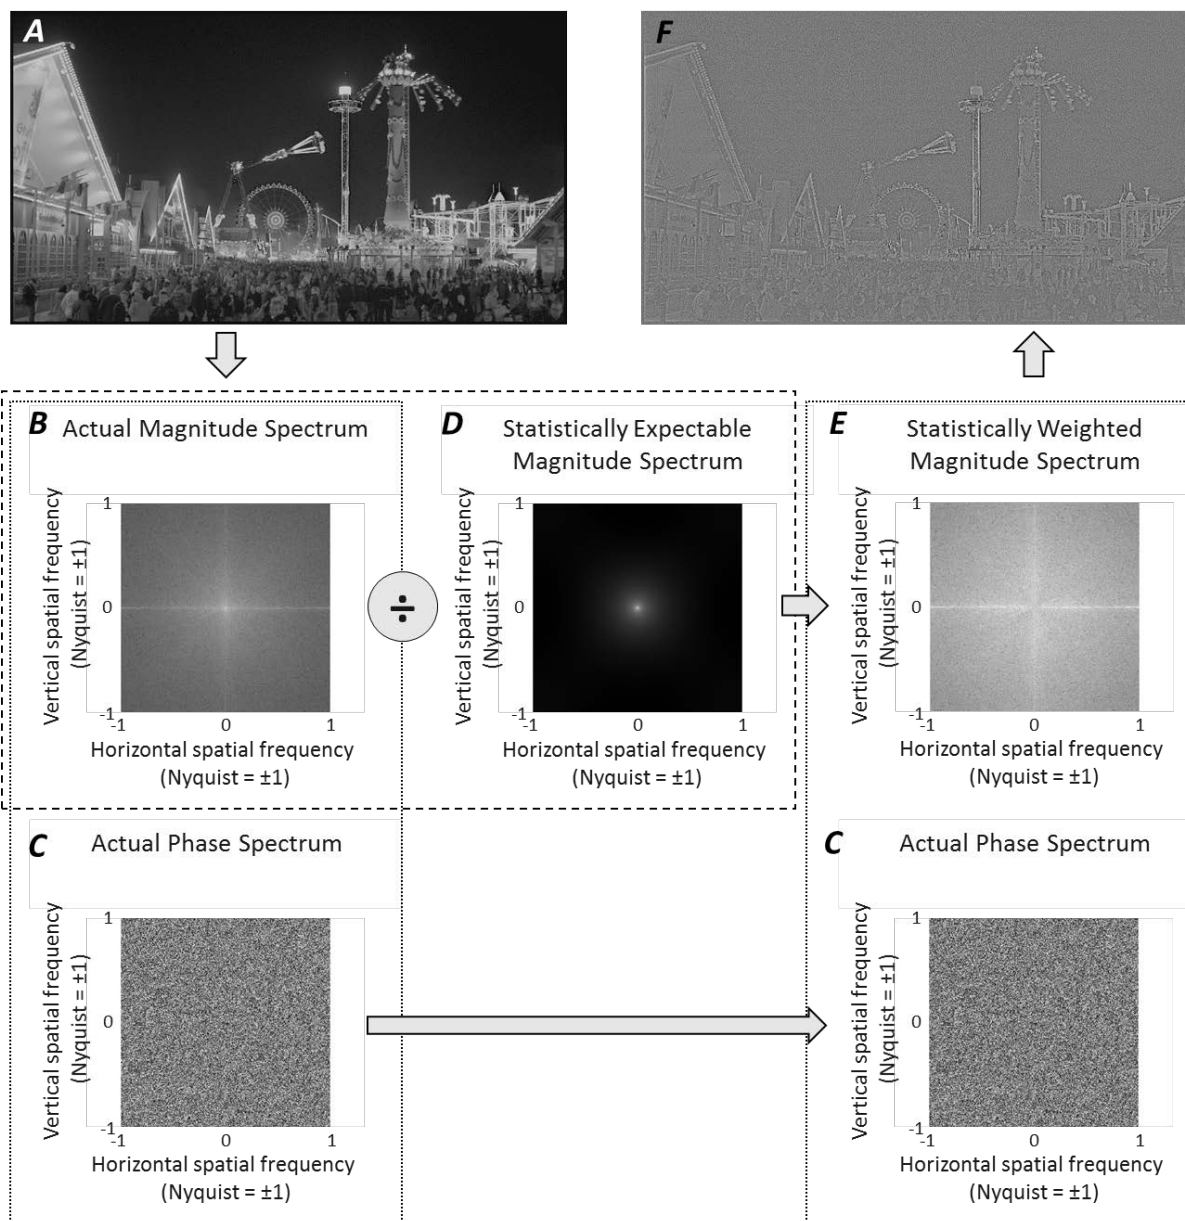


Figure 10 - Method of Calculating the Spatial Detail Signal

The Spatial Detail signal is calculated as illustrated in Figure 10. First, the magnitude (**B**) and phase spectra (**C**, shown twice) are calculated from the image pixel array (**A**). Next, a predetermined archetype of the statistically expectable one-over-frequency magnitude spectrum (**D**) is divided into the actual magnitude spectrum to produce a statistically weighted magnitude spectrum (**E**). Third, the statistically weighted magnitude spectrum is combined with the actual phase spectrum (**C**). Finally, a 2-dimensional Inverse Fast Fourier Transform is applied to produce a pixel array that we call the Spatial Detail signal (**F**).

Abbreviations

4k	3840x2160 resolution
ABR	Adaptive Bitrate
ATIS	Alliance for Telecommunications Industry Solutions
ATSC	Advanced Television Systems Committee
DOCSIS	Data Over Cable Interface Specifications
EOTF	Electro-Optical Transfer Function
HD	High Definition
HDTV	High-Definition Television
HD-HDR	High Definition-High Dynamic Range
HdM	Hochschule der Medien (Stuttgart Media University)
HDR	High Dynamic Range
HEVC	High Efficiency Video Coding
HLG	Hybrid-Log Gamma
IMF	Interoperable Master Format
IP	Internet Protocol
ITU-R	International Telecommunication Union - Radiocommunication Sector
JPEG2000	Joint Photographic Experts Group-2000
MXF	Material Exchange Format
MSE	Mean-Squared Error
MSO	Multiple-System Operator
MVPD	Multichannel Video Programming Distributor
OETF	Opto-Electronic Transfer Function
PQ	Perceptual Quantizer
PSNR	Peak Signal-to-Noise Ratio
RGB	Red Green Blue
SHVC	Scalable High Efficiency Video Encoding
SMPTE	Society of Motion Picture and Television Engineers
TIFF	Tagged Image File Format
Ultra HD	Ultra High Definition
v210	Quicktime Raw Component Video Picture Format
WCG	Wide Color Gamut
YCbCr	Luma and Chroma Video Picture Format

Bibliography & References

1. ITU-R Report BT.2246-2 (2017) “The present state of ultra-high definition television.”
2. ITU-R Report BT.2390-2 (2017) “High dynamic range television for production and international programme exchange.”
3. Roettgers, J. “The Story Behind ‘Meridian’: Why Netflix Is Helping Competitors With Content and Code.” Variety, Sep 15, 2016.
4. ITU-R BT.709-6 (2015) “Parameter values for the HDTV standards for production and international programme exchange.”
5. Xiph.org <https://media.xiph.org/video/derf/>
6. SMPTE OV 2067-0:2017 “SMPTE Overview Document – Interoperable Master Format – Overview for the SMPTE 2067 Document Suite.”
7. SMPTE ST 377-1:2011 “SMPTE Standard – Material Exchange Format (MXF) – File Format Specification.”
8. ISO/IEC 15444-1:2016 “Information Technology – JPEG 2000 image coding system: Core coding system”
9. ffmpeg.org <https://www.ffmpeg.org>
10. ITU-R Rec. BT.2100-0 (2016) “Image parameter values for high dynamic range television for use in production and international programme exchange.”
11. SMPTE ST 2084:2014 “SMPTE Standard – High Dynamic Range Electro-Optical Transfer Function of Mastering Reference Displays.”
12. SMPTE ST 2067-21:2016 “SMPTE Standard – Interoperable Master Format – Application #2E”
13. Froehlich, J., et al., “HdM-HDR-2014 Project,” <http://www.hdm-stuttgart.de/~froehlichj/hdm-hdr-2014>
14. Froehlich, J., Grandinetti, S., Eberhardt, B., Walter, S., Schillin, A., and Brendel, H. 2014. “Creating cinematic wide gamut HDR-video for the evaluation of tone mapping operators and HDR-displays,” Proc. SPIE 9023, Digital Photography X
15. ITU-R Rec. BT.2020-2 (2015) “Parameter values for ultra-high definition television systems for production and international programme exchange.”
16. “V210 Video Picture Encoding” at Library of Congress (Sustainability of Digital Formats). <https://www.loc.gov/preservation/digital/formats/fdd/fdd000353.shtml>
17. ITU-R Rec. BT.1886 (2011) “Reference electro-optical transfer function for flat panel displays used in HDTV studio production.”
18. MathWorks, MATLAB. <https://www.mathworks.com/>
19. x265 (x265-64bit-10bit-2017-05-01.exe) <https://builds.x265.eu/>
20. McCarthy, S.T. 2017. “A Biologically-Inspired Approach to Making HDR Video Quality Assessment Easier” SMPTE Motion Imaging Journal, vol. 124, no. 4, pp 47-58
21. McCarthy, S.T. 2014. “Theory and practice of perceptual video processing in broadcast encoders for cable, IPTV, satellite, and internet distribution,” Proc. SPIE 9014, Human Vision and Electronic Imaging XIX
22. McCarthy, S. 2012. “A Biological Framework for Perceptual Video Processing and Compression,” SMPTE Mot. Imag. J., 119(8):24-32, Nov/Dec.
23. McCarthy, S.T. and Owen, W.G. 2006. “Apparatus and Methods for Image and Signal Processing”. US Pat. 6014468 (2000). US Pat. 6360021 (2002), US Pat. 7046852 (2006)
24. Field, D.J. 1987. “Relationship between the statistics of natural images and the response properties of cortical cells,” J. Opt. Soc. Am. A. Vol. 4, No. 12

25. Boyce, J.M., Ye, Y., Chen, J., and Ramasubramonian, A.K. "Overview of SHVC: Scalable Extensions of the High Efficiency Video Coding Standard." IEEE Trans Circuits and Systems for Video Technology, Vol. 26, No. 1, 2015.
26. ATSC A/341:2017 "Video – HEVC"
27. Winkler, S. 2005 Digital Video Quality: Vision Models and Metrics, John Wiley & Sons
28. Wang, Z. and Bovik, A.C. 2009. "Mean squared error: love it or leave it? - A new look at signal fidelity measures," IEEE Signal Processing Magazine, vol. 26, no. 1, pp. 98-117
29. "Lanczos resampling" https://en.wikipedia.org/wiki/Lanczos_resampling
30. Bansal, G. 2017. "What is the difference between coefficient of determinations and coefficient of correlation?" <http://blog.uwgb.edu/bansalg/statistics-data-analytics/linear-regression/what-is-the-difference-between-coefficient-of-determination-and-coefficient-of-correlation/>
31. ATIS 0800061 (2013) "Methodology for subjective or objective video quality assessment in multiple bit rate adaptive streaming"

Single-Layer HDR Video Coding with SDR Backward Compatibility

A Technical Paper prepared for SCTE•ISBE by

David Touze

Research Engineer
Technicolor

975 avenue des Champs Blancs - ZAC des Champs Blancs - 35576 Cesson Sevigné, France
+33 2 99 27 31 08
david.touze@technicolor.com

Leon van de Kerkhof

Chief Architect
Philips

leon.van.de.kerkhof@philips.com

Introduction

The arrival of the High Efficiency Video Coding (HEVC) standard enables the deployment of new video services with enhanced viewing experience, such as Ultra HD broadcast services. In addition to an increased spatial resolution, Ultra HD will bring a wider color gamut (WCG) and a higher dynamic range (HDR) than the standard dynamic range (SDR) HD-TV currently deployed. Increasing of dynamic range, i.e. the luminance ratio of bright over dark pixels, has been shown to dramatically improve the user experience. Increasing gamut and dynamic range are two faces of the same coin as they basically augment the color volume to which pixels belong. Furthermore, luminance and colors are intrinsically linked in legacy workflows that are non-constant luminance: the signal non-linearity is not applied directly to the luminance, but instead the non-linear luminance is a combination of non-linear quantities (typically RGB).

Different solutions for representing and coding HDR/WCG video have been proposed [1][2] [3] [4]**Error! Reference source not found.**. As stated in [5][6][7][8]. SDR backward compatibility with decoding and rendering devices is an important feature in video distribution systems, such as broadcasting or multicasting systems. The coming American broadcast standard ATSC 3.0 is expected to emit both SDR BT.709/2020 and HDR BT.2020 streams. The European DVB standard has already introduced SDR UHDTV in the BT.2020 color space and will extend it to HDR BT.2020 soon. Peak brightness is expected to migrate from legacy 100 nits to about 1000 nits, but compression solutions should be flexible enough to handle future higher brightness as well as non-broadcast applications that may take advantage of more nits.

Dual-layer coding, for instance using the scalable extension of HEVC (a.k.a. SHVC) is one solution to support SDR backward compatibility. However, due to its multi-layer design, this solution is not adapted to all distribution workflows. An alternative is to transmit HDR content and to apply at the receiving device an HDR-to-SDR adaptation process (tone mapping). One issue in this scenario is that the tone mapped content may be out of control of the content provider or creator. Another issue is that a new HDR-capable receiving device is needed to apply this tone mapping for existing SDR displays. Alternatively, the Hybrid Log Gamma (HLG) transfer function [2] has been designed as a straightforward solution to address the SDR backward compatibility, that is, an HDR video graded on a display using the HLG transfer function can be in principle directly displayed on an SDR display (using the BT.1886 transfer function [2]) without any adaptation. However, this solution may result in color shifting when the HLG-graded video is displayed on an SDR rendering device, especially when dealing with content with high dynamic range and peak luminance [10][11][12]. Also, there is no way to optimize the brightness and contrast of the SDR image.

The proposed Single Layer SDR backward compatible HDR video distribution solution detailed in this paper, named SL-HDR1, and standardized in ETSI TS 103 433 specification [13]**Error! Reference source not found.**, aims at addressing these issues. SL-HDR1 leverages SDR distribution networks and services already in place. It enables both high quality HDR rendering on HDR-enabled CE devices, while also offering high quality SDR rendering on SDR CE devices.

The main features of the HDR distribution system are as follows:

- Single layer with metadata: SL-HDR1 is based on a single layer coding process, with side metadata that can be used at post-processing stage. The metadata payload corresponds to a few bytes per picture, GOP or scene.

- Codec agnostic: SL-HDR1 does not impact the core codec technology and is codec independent. SL-HDR1 is based on an encoding pre-processing applied to the HDR input, and on a corresponding decoding post-processing (functional inverse of the pre-processing) applied to the reconstructed video from decoding. Use of a 10-bit codec is recommended, since an 8-bit codec could introduce artefacts such as banding effects, due to having too few codewords available for the precision required for HDR content.
- Enable SDR backward compatibility: a decoded bitstream can be displayed as is on an SDR display. The color fidelity is preserved compared to the HDR version. An additional post-processing is applied to convert the decoded SDR version to HDR, thanks to the metadata, with preservation of the HDR artistic intent.
- Enable preserved quality of HDR content: there is no penalty due to the SDR backward compatibility feature; coding performance compared to HLG are improved, in particular in terms of color impairments.
- Enable adaptation of the HDR content to the HDR display capabilities: if the HDR content peak brightness is higher than the HDR display peak brightness, the post-processing adapts the HDR content to display peak brightness, preserving all details and HDR artistic intent.
- Limited additional complexity: the pre- and post-processing steps are of limited added complexity; in particular the involved operations are only sample-based, without inter-sample dependency.
- Independent from the input OETF: the pre- and post-processing operate in linear-light domain, and are therefore independent from the input OETF.

The document is organized as follows. The solution overview is presented in section 1. Section 2 describes the HDR-to-SDR decomposition and section 3 the HDR reconstruction process. Section 4 relates to the metadata signaling. Section 5 details the display adaptation feature. Section 6 presents tests results, assessing the SDR quality and the HDR compression performance of SL-HDR1 comparatively to distribution solutions based on PQ and HLG transfer functions. Conclusion section provides closing remarks.

Content

1. SL-HDR1 System Overview

Figure 1 shows an end-to-end workflow supporting content production and delivery to HDR and SDR rendering devices. The core of the HDR distribution solution SL-HDR1 corresponds to yellow and green boxes. SL-HDR1 involves a single-layer SDR/HDR encoding-decoding, with side dynamic metadata. At the distribution stage, an incoming HDR signal is decomposed in an SDR signal and content-dependent dynamic metadata. The SDR signal is encoded with any distribution codec (e.g. HEVC Main 10) and carried throughout the existing SDR distribution network with accompanying metadata conveyed on a specific channel or embedded in the SDR bitstream. The dynamic metadata are typically carried in an SEI message when used in conjunction with an HEVC codec. The post-processing stage is functionally the inverse of the pre-processing and performs the HDR reconstruction. It occurs just after SDR bitstream decoding. The post-processing takes as input an SDR video frame and associated dynamic metadata in order to reconstruct an HDR picture. Single-layer encoding/decoding requires only one encoder instance at HDR encoding side, and one decoder instance at player/display side. It supports the real-time workflow requirements of broadcast applications. The dynamic metadata are produced by the HDR decomposition process and remain internal to the distribution process. They do not need to be conveyed to the rendering device. Additional metadata, originated from the production/post-production, can optionally be distributed and conveyed to the rendering device.

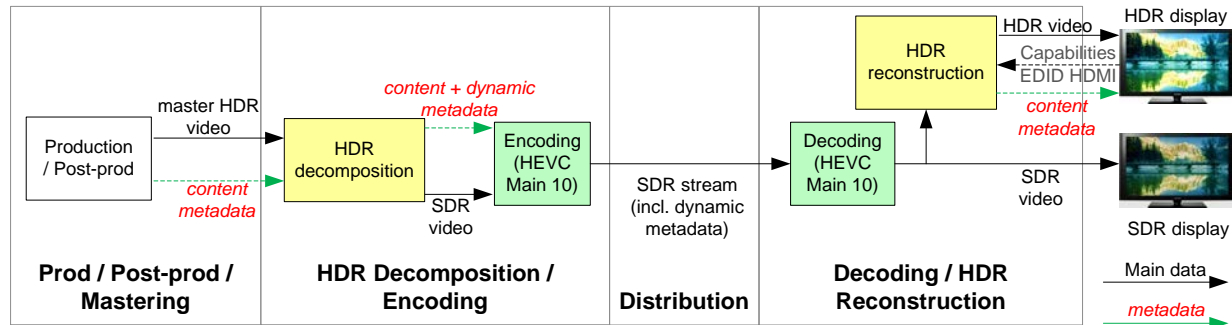


Figure 1 - Example of HDR end-to-end system.

The block diagram in Figure 2 depicts in more details the HDR decomposition and reconstruction processes. The center block included in red dashed box corresponds to the distribution encoding and decoding stages. The left and right grey-colored boxes respectively enable format adaptation to the input video signal of the HDR system and to the targeted system (e.g. a STB, a connected TV). The yellow boxes show the HDR specific processing. The first step of the HDR decomposition process linearizes the input HDR content, allowing the system to ingest every HDR production format such as PQ (display referred), HLG (scene referred) or any other production format. For the HLG case, as it is a relative format, the peak brightness of the HLG content needs to be provided to the system. The linearized content is then independent from the input format and allows the system to always work in the same consistent linear-light domain. The core component of the HDR decomposition stage is the HDR-to-SDR conversion that generates an SDR video from the linear-light HDR signal. Optionally, gamut mapping may be used when the input HDR and output SDR signals are represented in different color spaces. This

optional gamut mapping may be introduced either before or after the HDR-to-SDR conversion. The decoder side implements the inverse processes, in particular the SDR-to-HDR reconstruction step that inputs the SDR video provided by the decoder and that transforms it back to an HDR video at a peak luminance adapted to the HDR display capabilities.

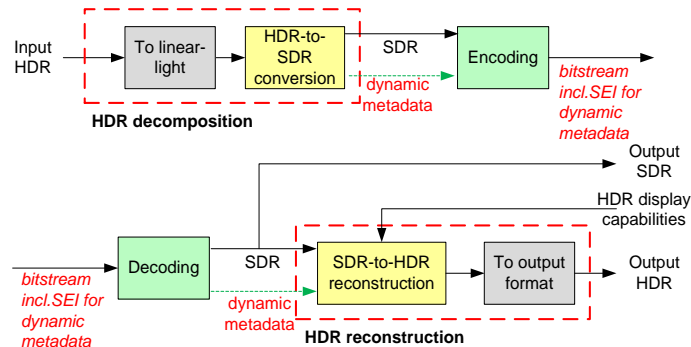


Figure 2 - HDR system architecture overview.

2. HDR-to-SDR Decomposition Process

The HDR-to-SDR decomposition process aims at converting the input linear-light 4:4:4 RGB HDR signal to an SDR Y'CbCr 4:2:0 compatible version. The process uses side information such as the color primaries and gamut of the container of the HDR and SDR pictures. The process operates without color gamut change: the HDR and SDR pictures are defined in the same color gamut. If needed, a gamut mapping processing may be applied either on the HDR pictures or on the SDR pictures. In the former case the HDR picture is converted from its native color gamut to the target color gamut before the HDR-to-SDR decomposition process. And in the latter case the SDR picture generated by the HDR-to-SDR decomposition process is converted from its native color space to the target SDR color gamut.

The HDR-to-SDR decomposition process is depicted in Figure 1 and Figure 3. It is primarily based on the HDR content analysis (picture per picture) in order to derive a set of mapping parameters that will be further used to convert the HDR signal into SDR (step 1). Once the mapping parameters have been derived, a luminance mapping function, noted TM , is obtained. In step 2, the luminance L , derived from the HDR linear-light RGB signal, is mapped to an SDR luma signal using the luminance mapping function TM (step 2). The chroma components are then derived (step 3). A final color correction is applied in order to match the SDR colors to the input HDR signal colors (step 4). Steps 2 to 4 are detailed in the following sub-sections.




Figure 3 - Synoptic of HDR-to-SDR Decomposition Process

2.1. Luminance mapping

The luminance mapping (step 2) aims at converting the input linear-light luminance signal, derived from the HDR linear-light RGB signal, into an SDR luma signal using a luminance mapping function TM . This is done according to the following equations:

$$L = A_1 \times \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (1)$$

$$Y_{tmp} = (LUT_{TM}(L))^{\frac{1}{2.4}} \quad (2)$$

where $A = [A_1 A_2 A_3]^T$ is the conventional 3x3 R'G'B'-to-YCbCr conversion matrix (e.g. BT.2020 or BT.709 depending on the color space), A_1, A_2, A_3 being 1x3 matrices.

The mapping function or look-up-table TM is built as follows. The mapping is based on a perceptual transfer function, and uses a limited set of control parameters, that have to be further conveyed to the post-processing in order to be able to invert the luminance mapping process. The input linear-light luminance signal L is first converted to the perceptually-uniform domain based on the mastering display peak luminance, using a perceptual transfer function illustrated in left picture of Figure 4. This process is controlled by the mastering display peak luminance parameter. To better control the black and white levels, a signal stretching between content-dependent black and white levels (parameters *blackLevelOffset* and *whiteLevelOffset*) is applied. Then the signal is tone mapped using a piece-wise curve constructed out of three parts, as illustrated in Figure 5. The lower and upper sections are linear, the steepness being determined by the *shadowGain* and *highlightGain* parameters. The mid-section is a parabola providing a smooth bridge between the two linear sections. The width of the cross-over is determined by the *midToneWidthAdjFactor* parameter. The curve can be further fine-tuned using a piece-wise linear corrective function. Then the signal is converted back to the linear light domain based on the targeted SDR display maximum luminance of 100 cd/m², as illustrated in the right picture of Figure 4. The resulting signal is the SDR luma.

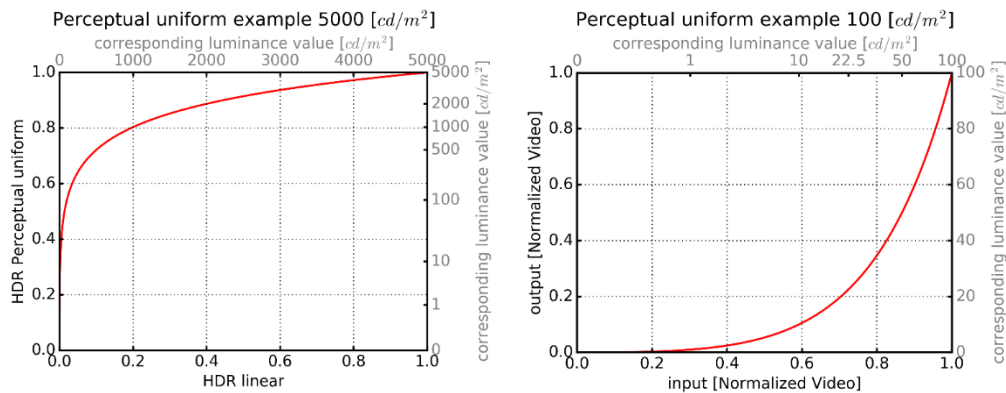


Figure 4 - Example conversion curves for converting from linear light to perceptual domain (left, with peak luminance 5000 cd/m²) and back to SDR linear light (right).




Figure 5 - Tone mapping curve shape example.

2.2. Chroma components derivation

The chroma components are derived as follows (step 3). First a square root is applied to input HDR linear-light R , G , B and L values to reproduce a transfer function close to the BT.709/BT.2020 OETF (the usage of a square root guarantees the reversibility of the process). Then the resulting squared-root R , G , B values are scaled by the squared-root L value, which results in a gamma-sized SDR version of the input R , G , B signals. The resulting R , G , B signal is converted to chroma components U_{tmp} , V_{tmp} :

$$\begin{bmatrix} U_{tmp} \\ V_{tmp} \end{bmatrix} = \frac{1}{\sqrt{L}} \times \begin{bmatrix} A_2 \\ A_3 \end{bmatrix} \times \begin{bmatrix} \sqrt{R} \\ \sqrt{G} \\ \sqrt{B} \end{bmatrix} \quad (3)$$

2.3. Color correction

A final color correction is applied in order to match the SDR colors to the input HDR signal colors (step 4). First the chroma components are adjusted by a scaling factor $1/\beta(Y_{tmp})$, where $\beta(Y_{tmp})$ is a function that enables to control the color saturation and hue and that is constructed by matching primaries and white points between the SDR and the HDR gamut.

$$\begin{bmatrix} U_{sdr} \\ V_{sdr} \end{bmatrix} = \frac{1}{\beta(Y_{tmp})} \times \begin{bmatrix} U_{tmp} \\ V_{tmp} \end{bmatrix} \quad (4)$$

Then the luma component is adjusted to further control the perceived saturation, as follows:

$$Y_{sdr} = Y_{tmp} - \text{Max}(0, a \times U_{sdr} + b \times V_{sdr}) \quad (5)$$

where a and b are two control parameters. This luma adjustment step helps in recovering the color perception difference that occurs when a specific color is rendered at different luminance level.

As demonstrated in [13], this color correction step is fundamental to control the SDR colors and to guarantee their matching to the HDR colors. This is in general not possible when using a fixed transfer function.

3. HDR Reconstruction

The HDR reconstruction process is depicted in Figure 6. This section describes the reversible process without taking into account the display adaptation feature that is detailed in section 5. From the input dynamic metadata (detailed in section 4) a luma-related look-up table, *lutMapY*, and a color correction look-up table, *lutCC*, are derived. The next step consists in applying the SDR-to-HDR reconstruction from the input SDR picture, the derived luma-related look-up table and color correction look-up table. This process produces an output linear-light HDR picture. An optional gamut mapping can be applied when the color spaces of the SDR picture and of the HDR picture are different (either before or after the SDR-to-HDR reconstruction).




Figure 6 - Overview of the HDR reconstruction process.

The SDR-to-HDR reconstruction process is the functional inverse of the decomposition process. However, for implementation complexity reasons, some operations are concatenated or applied in a different order. This is the case for the final R, G, B reconstruction step described below. The operation reordering and concatenation allows this step to be implemented as a single “output” LUT and this method optionally allows the addition of an output EOTF (such as PQ or HLG) in the same “output” LUT. The LUT *lutMapY* actually corresponds to the inverse of the square-root of the mapping LUT *TM*. The post-processing color correction LUT *lutCC* is actually linked to the pre-processing color correction LUT β and the tone mapping LUT *lutMapY* by the following equation:

$$\beta[Y] = 2^B \times lutMapY[Y] \times lutCC[Y] \quad (6)$$

where B is the bit-depth of the luma signal. It can be demonstrated that the inverse of *lutCC* is close to a linear function.

The HDR reconstruction process performs the following successive steps for each sample Y , U (Cb component), V (Cr component), of the SDR picture. First U and V are centered (by subtracting the chroma offset, e.g. 512 for a 10 bits signal). Then the variables Y_{post} , U_{post} and V_{post} are derived as:

$$Y_{post} = Clamp(0, 2^B - 1, Y + Max(0, a \times U + b \times V)) \quad (7)$$

$$\begin{bmatrix} U_{post} \\ V_{post} \end{bmatrix} = lutCC[Y_{post}] \times \begin{bmatrix} U \\ V \end{bmatrix} \quad (8)$$

The reconstruction of the HDR linear-light R , G , B values is made up of the following steps. A parameter T is first computed as:

$$T = k0 \times U_{post} \times V_{post} + k1 \times U_{post} \times U_{post} + k2 \times V_{post} \times V_{post} \quad (9)$$

where $k0$, $k1$, $k2$ are predefined parameters that depend on the coefficients of the R'G'B'-to-Y'CbCr conversion matrix A . The intermediate values R_{im} , G_{im} , B_{im} are derived as follows:

$$\begin{bmatrix} R_{im} \\ G_{im} \\ B_{im} \end{bmatrix} = A^{-1} \times \begin{bmatrix} \sqrt{1-T} \\ U_{post} \\ V_{post} \end{bmatrix} \quad (10)$$

A clamping is done to $0, \sqrt{L_{HDR}}$, where L_{HDR} is the HDR mastering display peak luminance.

Then, linear-light R , G , B values are obtained by the following equation:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = (lutMapY[Y_{post}])^2 \times \begin{bmatrix} R_{im}^2 \\ G_{im}^2 \\ B_{im}^2 \end{bmatrix} \quad (11)$$

It can be demonstrated that equations (9) to (11) invert the pre-processing operation of (1) to (3), that is, the conversion of the HDR version R , G , B into chroma components. When T is larger than 1 (which is in principle not possible, but may happen because of quantization and compression), U_{post} and V_{post} are scaled by $1/\sqrt{T}$, the resulting T becoming equal to 1. As this scaling applies simultaneously on the two chroma components, the resulting hue remains stable.

4. Metadata Description

The post-processing uses the LUTs *lutMapY* and *lutCC*, and the parameters a , b , $k0$, $k1$ and $k2$, as dynamic data. These data enable to finely control the texture and colors of the SDR version, and to ensure a good fit to the HDR intent. The LUTs *lutMapY* and *lutCC* are conveyed either using a limited set of parameters (parameter-based mode), or explicitly coded (table-based mode). In both cases, the metadata payload corresponds to a few bytes per video frame or scene. The parameter-based mode may be of interest for distribution workflows which primary goal is to provide direct SDR backward compatible services with very low additional payload or bandwidth usage for carrying the dynamic metadata. The table-based mode may be of interest for workflows equipped with low-end terminals or when a higher level of adaptation is required for representing properly both HDR and SDR streams.

Next to the dynamic metadata, the system uses information that define the properties of the mastering display used when grading the HDR content, as defined in Mastering Display Colour Volume (MDCV) message. This is static information (typically fixed per program) required by the post-processing. It comprises the color gamut of the SDR/HDR signal and the mastering display peak luminance.

In the parameter-based mode, the metadata for reconstructing *lutMapY* consist of the parameters mentioned in section 2.1. For reconstructing *lutCC*, a default pre-defined LUT is used at the post-processing side, and a piece-wise linear table made of at most 6 points is used as a scaling function to adjust the default table. These parameters are conveyed using the parameters defined in the SMPTE ST 2094-20 specification. Typical payload is about 70 bytes per scene, including Mastering Display Colour Volume (MDCV) message. In the table-based mode, *lutMapY* and *lutCC* are explicitly coded using the parameters defined in the SMPTE ST 2094-30 specifications. Typical payload is about 186 bytes per scene, including Mastering Display Colour Volume (MDCV) message. In both cases, the metadata are limited to the codec space. They do not come from the production side, and do not need to be conveyed outside the decoding platform. They are conveyed using standardized metadata containers.

The usage of dynamic metadata allows a fine control of the SDR texture (using the tone mapping LUT *lutMapY*) and of colors (using the color correction LUT *lutCC* and the parameters a , b , $k0$, $k1$ and $k2$). This guarantees the preservation of the HDR texture and intended colors in the SDR version, as illustrated in pictures in next section. High SDR and HDR video quality is obtained, without any strong limitation of the dynamic range and peak luminance (no limitation to peak luminance of around 1000-1500 nits). This also gives high flexibility which enables to easily adapt the system (for instance thanks to the easy control of the dynamic metadata payload) to the distribution workflow.

5. Display Adaptation

The display adaptation feature is only active with parameter-based metadata mode and is based on the Tone Mapping and Inverse Tone Mapping computation blocks of the system.

On the HDR decomposition side (see Figure 3), the luminance mapping block (step 2) computes a Tone Mapping curve based on the dynamic parameters described in section 2.1 and based on the mastering display peak luminance and the targeted SDR display maximum luminance of 100 cd/m².

On the HDR reconstruction side (see Figure 6), the *lutMapY* derivation block computes an Inverse Tone Mapping curve based on the same dynamic parameters, the same SDR display maximum luminance of 100 cd/m² and the same HDR mastering display peak luminance. The resulting Inverse Tone Mapping curve is the inverse of the Tone Mapping curve computed by the HDR decomposition block, as depicted in Figure 7.

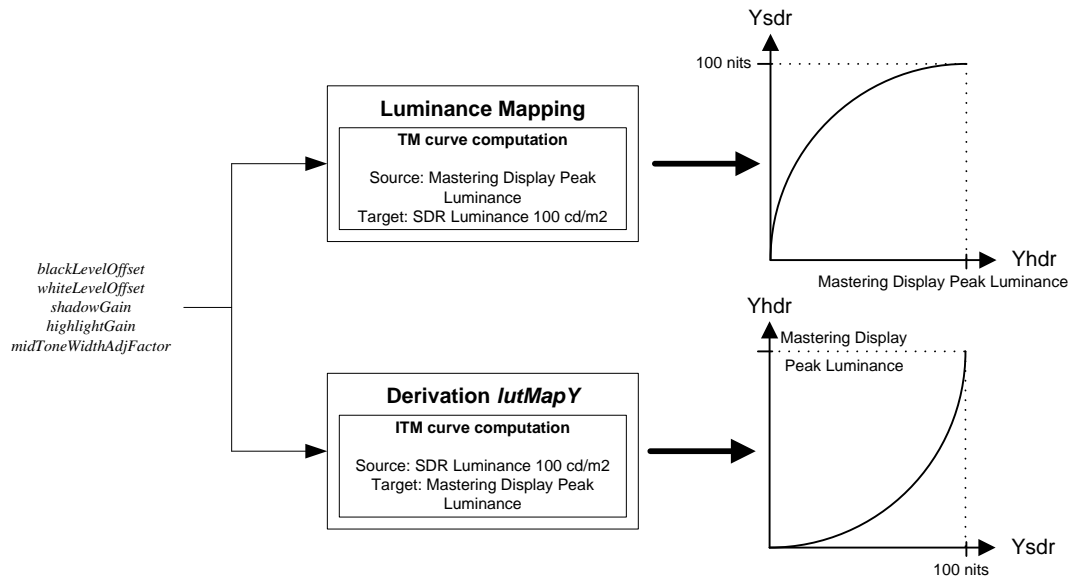


Figure 7 - Tone Mapping and Inverse Tone Mapping computation blocks

The display adaptation feature is an extension of the HDR reconstruction process and takes place in the *lutMapY* derivation and *lutCC* computation process, as shown in Figure 8.

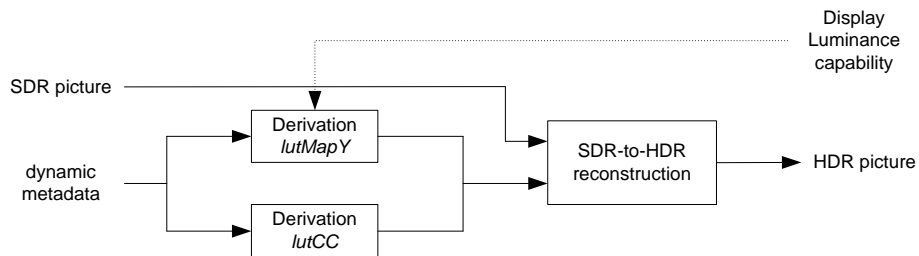


Figure 8 - Overview of the HDR reconstruction process with display adaptation

Given that the attached display provides its peak luminance capability (either through the EDID data of the HDMI connection to the display or inherently when the processing is integrated in the display device), the solution consists of a cascaded calculation of the previously described Inverse Tone Mapping curve and an adapted Tone Mapping curve. This added Tone Mapping block uses the same dynamic parameters and the same input mastering display peak luminance but now relies on the provided presentation display peak luminance capability as the target output peak luminance, as shown in Figure 9:

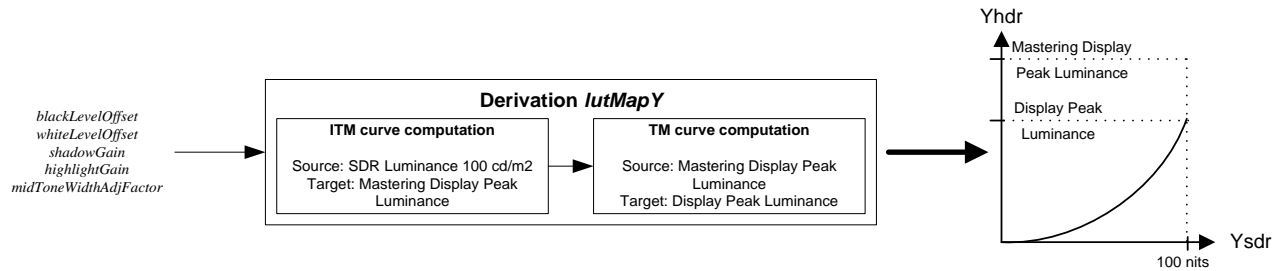


Figure 9 - Inverse Tone Mapping curve computation with display adaptation

The cascading of these two processes is easily implemented in the *lutMapY* LUT with no added complexity.

This process in essence only adapts the luminance of the signal and aims to preserve the artistic intent for a presentation display with a different peak luminance than the targeted SDR or the original HDR peak luminance by means of a reconstructed adapted signal.

6. Performance Evaluations

Two reference distribution solutions, PQ and HLG, can be considered for evaluating an HDR video distribution solution. PQ [14] is a non-SDR backward compatible transfer function, based on the human contrast sensitivity model developed by Barten [15] more adapted to HDR than the usual gamma. HLG [2] is a new transfer function aiming at offering some level of SDR backward compatibility. The two next sections report comparative results of SL-HDR1 vs. distribution solutions using these two transfer functions.

6.1. HDR compression comparison with PQ transfer function

This section reports HDR compression results of SL-HDR1, compared to the non-SDR backward compatible solution HDR10 based on the PQ transfer function. The considered test sequences, selected for the MPEG HDR and WCG Call for Evidence[5], are natively in RGB, 4:4:4, linear-light format, represented in BT.2020 color primaries container. The complete conversion and coding, decoding and back-conversion chain, for all tested solutions has been performed in BT.2020 color primaries container. All contents have been compressed using the MPEG reference HEVC encoder HM 16.2.

Several metrics to assess HDR image objective quality have been proposed at MPEG [5]. Unfortunately, none of them is very satisfactory in the sense that better metric values do not necessarily imply better subjective quality. For this reason, metric values have to be considered with care and are provided in Table 1 for both 4:2:0 subjectively optimized luminance mapping parameters (the automatic tuning is adjusted to get optimal visual SDR quality, left column) and 4:4:4 metric oriented parameter optimization (the automatic tuning is adjusted to get high objective metric DE100 values, right column). The gains are expressed in average bitrate savings compared to the HDR10 reference using the Bjøntegaard-Delta-rate measure [16] and do include the overhead due to the metadata (encoded in dedicated SEI message) associated to the proposed method. A negative value of x% indicates that SL-HDR1 requires x% less bitrate than the reference (PQ/HDR10) for a similar quality.

The metric DE100 is more color oriented and the metric L100 evaluates the luminance quality only. Both metrics are based on an extension of the well-known Lab 2000 color space [17]**Error! Reference source**

not found.. It is observed that high metric gains are possible, but subjectively optimized compressed videos show much less gains, proving once again the inconsistency of the objective metrics. In this test, it has been verified by subjective tests performed by non-naïve viewers that all visually optimized videos have a quality at least comparable to the best proponents of the MPEG HDR and WCG Call for Evidence [5]**Error! Reference source not found..** The artifacts observed on the reconstructed HDR content include typical compression artefacts already observed in SDR compression such as blocking artefacts and also inconsistent color patches generation especially for color areas reaching the color gamut boundaries.

Table 1 - HDR compression performance on MPEG metrics compared to HDR10.

Sequence	Resolution	Peak luminance	420 subjective optim.		444 DE100 optim.
			DE100	L100	DE100
Market3	HD	4000 nits	-30.4%	-25.1%	-71.8%
AutoWelding	HD	4000 nits	-14.8%	-2.8%	-38.8%
ShowGirl2Teaser	HD	4000 nits	-34.1%	-8.9%	-56.8%
StEM_WarmNight	HD	4000 nits	-17.7%	-2.8%	-41.1%
BalloonFestival	HD	5000 nits	-3.1%	-13.6%	-61.3%
Average			-20.0%	-10.6%	-53.9%

It should be noted that SDR backward compatibility imposes a natural balance, but not optimal for HDR compression, between chroma and luma. The balance may be compensated by an adequate adjustment of the chroma quantization parameter (controlled by syntax element QPChromaOffset in HEVC). It has been observed that the tuning of this parameter also improves the quality of HDR10 anchors and it may as well improve the proposed solution. However this raises the issue of having to develop, for HDR content, particular encoders significantly different from encoders optimized for SDR video. An extra-advantage of SDR backward compatible solutions is that existing SDR encoders can be re-used.

6.2. HDR Compression Comparison with HLG Transfer Function

This section reports SDR quality evaluation and HDR compression results of SL-HDR1, compared to the SDR backward compatible solution based on the HLG transfer function. For these tests, visual evaluations made by an independent lab, under supervision of V. Baroncini (MPEG tests chair), have been performed. Two different tests were performed. First tests aimed at evaluating the SDR backward compatibility feature, by checking the quality of the SDR video generated from the HDR content. Second tests aimed at evaluating the HDR compression performance. For these two tests, comparative evaluation was performed between three solutions: HLG, SL-HDR1 with first tuning, SL-HDR1 with second tuning. First tuning tends to generate brighter pictures, i.e. picture with a higher average luminance level, than second tuning. In both cases, the tuning is fully automatic, but is performed according to one of these two different modes.

In the experiments, the HLG implementation from HDRTools software, version 0.12 (accessible at link <https://gitlab.com/standards/HDRTools/tags/v0.12>), made by the HLG designers, has been used to generate the HLG results. As recommended by the HLG designers, a system gamma correction was applied to the input linear-light RGB HDR content prior to converting it with HLG. The value of the

system gamma γ in the HLG pre-processing depends on the peak luminance L_{peak} of the HDR mastering display, and is derived as follows:

$$\gamma = 1.2 + 0.42 \times \text{Log}_{10}(L_{peak} / 1000) \quad (12)$$

6.2.1. Test Sequences

In order to have a future-proof evaluation, and to anticipate the evolution of HDR displays capabilities, sequences with various peak luminance have been used. The content color gamut is either BT.709 or P3D65, but all sequences are represented in BT.2020 color primaries container. All sequences are natively represented in EXR RGB 4:4:4 linear-light half-float format. The complete conversion and coding, decoding and back-conversion chain, for all tested solutions has been performed in BT.2020 color primaries container.

The test sequences are listed in Table 2.

Table 2 - Test Sequences for Comparative Tests with HLG.

Name	Sequence	Peak luminance	Content gamut	Container gamut	fps	Size	duration
S0	Market3	4000	709	2020	50	HD	8s
S1	EBU_04_Hurdles	3000	709	2020	50	HD	10s
S2	EBU_04_Starting	3000	709	2020	50	HD	10s
S3	EBU_13_LongJump	3000	709	2020	50	HD	10s
S4	HdM ShowGirl	5000	P3D65	2020	24	HD	10s
S7	CableLabs Rope	5000	709	2020	24	HD	10s

6.2.2. Evaluation of SDR Quality

The goal of these tests is to verify that colors and texture of the SDR generated by HLG and SL-HDR1 conform to those of the HDR content. The methodology described in [18] is used. For each solution, the viewers have to assess the conformity of the SDR displayed on an SDR monitor to the HDR displayed on an HDR monitor (Sim2). In particular, they have to check the conformity of colors and the texture preservation. This test set-up uses two displays that, when driven with a “black” input signal, become not visible to the viewing subjects. Furthermore, an opaque non reflective curtain is placed between the displays, in a way that, even if the viewer can still watch both displays, any visible interference due to reflections and indirect illumination among the displays is avoided.

Test results are depicted in the graphs of Figure 10 and Figure 11. The vertical axis depicts the Mean Opinion Score (MOS). A score of 10 corresponds to quality of reproduction that is perfectly faithful to the original. A score of 0 denotes a quality of reproduction that has no similarity to the original. A worse quality cannot be imagined. 5 corresponds to a fair quality. The average score value, with related confidence interval, is depicted for each tested solution. In all cases but one, SL-HDR1, with any tuning, was judged much better than HLG. Only for sequence S7, SL-HDR1 tuning1 and HLG are equivalent, and better than SL-HDR1 tuning2. The average scores and confidence intervals for the 6 sequences, and for the three methods, are depicted in Table 3. SL-HDR1 with both tuning outperforms HLG by around 1.7 MOS points.

Illustrative SDR pictures resulting from the HDR conversion by HLG (left) and by SL-HDR1, with first tuning (right) are depicted in Figure 12. Color hue and contrast issues can be observed in HLG versions. In general, texture losses are observed in HLG, especially in bright areas (wall in Market, ground track in Hurdles). And saturated colors (such as red and purple colors) suffer from noticeable hue shifts.

Table 3 - Average MOS and Confidence Intervals

	Average MOS	Confidence interval
HLG	4.82	0.35
SL-HDR1 tuning1	6.56	0.26
SL-HDR1 tuning2	6.54	0.22




Figure 10 - Average MOS values per sequence, with corresponding confidence interval.




Figure 11 - Average MOS values per sequence, with corresponding confidence interval.




Figure 12 - SDR version from HLG (left) and SL-HDR1 (right) of Market (top) and EBU Hurdles (bottom) HDR first picture.

6.2.3. Evaluation of HDR compression performance

The goal of these tests is to evaluate the quality of the reconstructed HDR using SL-HDR1 compared to the reconstructed HDR using HLG (that is, HDR coded with HLG transfer function, compressed then decompressed with HEVC Main10). The test procedure for the formal subjective evaluation uses one of the test methods described in [19], specifically the Degradation Category Rating (DCR) method. Four bitrates were used, adapted to each test sequence, as listed in Table 4. The encoding was performed using a professional HEVC file encoder from Elemental, Main 10 profile, with same settings for the three tested solutions. The Intra period was set around 1s (24 for 24 fps content, 48 to 50 fps content). The hierarchical-B GOP structure was of size 8, with 4 reference pictures, and B-pictures used as reference.

The average bitrate savings of SL-HDR1 compared to HLG for each sequence have been computed from the MOS vs. bitrate data to quantify the achieved bitrate savings. Table 5 shows the MOS Bjøntegaard-Delta-rate (BD-rate) for each sequence. BD-rate measures as described in [16] were used with MOS scores taking the place of the peak signal-to-noise ratio (PSNR) values that have been typically used with BD-rate measurements, with negative numbers indicating percentage of rate reduction at the same MOS quality. The estimated rate saving by using SL-HDR1 in place of HLG is quite significant. Only in one case (sequence S7), the compression performance with SL-HDR1 tuning 1 is similar to HLG. The average MOS BD-rate gain (that gives an estimation of the bitrate saving) of SL-HDR1 compared to HLG is of 24.6% for tuning 1, and 26.7% for tuning 2.

Table 4 - Bitrates (kbps) per Sequence

Sequence	R1	R2	R3	R4
S0	8000	5000	3000	2000
S1	8000	5000	3000	2000
S2	8000	5000	3000	2000
S3	6000	4000	2000	1000
S4	6000	4000	3000	2000
S7	3000	1700	1000	0700

Table 5 - MOS BD-rate savings measurements.

	SL-HDR1 tuning1 vs HLG	SL-HDR1 tuning2 vs HLG
S0	-33.1%	-27.5%
S1	-13.9%	-9.6%
S2	-38.5%	-50.7%
S3	-37.1%	-17.9%
S4	-23.9%	-41.4%
S7	-1.2%	-13.1%
Avg	-24.6%	-26.7%

Conclusion

The co-existence of SDR and HDR on one hand, and BT.709 and BT.2020 contents on the other hand is likely to happen and last for at least a few years. Some applications, like broadcasting, will benefit from SDR backward compatible solutions that avoid simulcasting versions with various ranges and gamut. Backward compatibility would make the transition from SDR HDTV to HDR UHD TV smoother with increased interoperability.

The proposed solution, SL-HDR1, addresses the SDR/HDR backward compatibility by offering a new dynamic reducer with consistent SDR/HDR colors. This solution also adapts to the wide range of HDR display brightness capabilities by tuning the content to the display capabilities while preserving the artistic intent. It shows solid compression gain compared to the conservative non-backward compatible HDR10 approach. Tests results also show that SL-HDR1 outperforms HLG in a statistically significant way. This holds for both tests, i.e. both the visual quality of the SDR and that of the reconstructed HDR are better for SL-HDR1 compared to HLG. For SDR visual quality, SL-HDR1 outperforms HLG by 1.7 points in Mean Opinion Score (MOS). For HDR compression, a bitrate saving of around 25% is obtained by SL-HDR1 compared to HLG. SL-HDR1 is compliant with a 4:2:0 distribution workflow as well as with existing HDR color spaces (namely BT.2020 CL CbCr), non-linearity (PQ or HLG EOTF) and bit-depth (10 bits) used as input to rendering devices.

The solution has been designed with a particular focus on low complexity and high performance. The pre- and post-processing are of very low added complexity. The involved operations are pixel-based, without inter-sample or temporal dependency. The complexity increase is very reasonable (a few operations and

LUTs) relatively to the HDR coding gain and the new backward compatible feature that are provided. Associated metadata can be encapsulated in the compressed bit-stream and would not require their transmission from the production to the display. The solution has been standardized as ETSI TS 103 433.

Abbreviations

ATSC	Advanced Television Systems Committee
BD-rate	Bjøntegaard-delta-rate
CE	Consumer Electronics
DCR	degradation category rating
EBU	European Broadcasting Union
EDID	Extended Display Identification Data
ETSI	European Telecommunications Standards Institute
fps	frame per second
GOP	group of pictures
HD	high definition
HDR	high dynamic range
HEVC	high efficiency video coding
HLG	hybrid log gamma
ITM	inverse tone mapping
kbps	kilo bits per second
LUT	look-up table
MDCV	mastering display color volume
MPEG	Moving Picture Experts Group
MOS	mean opinion score
OETF	optical to electrical transfer function
PQ	perceptual quantization
PSNR	peak signal to noise ratio
SCTE	Society of Cable Telecommunications Engineers
SDR	standard dynamic range
SEI	supplemental enhancement information
SHVC	scalability extension of HEVC
STB	set-top box
TM	tone mapping
TS	technical specification
TV	television
UHD	ultra high definition
UHDTV	ultra high definition television
WCG	Wide color gamut

Bibliography & References

- [1] SMPTE, “High Dynamic Range Electro-Optical Transfer Function of Mastering Reference Displays”, SMPTE ST 2084 (2014).

- [2] Borer, T. and Cotton, A., “A display independent high dynamic range television system”, International Broadcasting Convention (IBC), <http://www.bbc.co.uk/rd/publications/whitepaper309> (September 2015).
- [3] EBU, “High-Performance Single Layer Directly Standard Dynamic Range (SDR) Compatible High Dynamic Range (HDR) System for use in Consumer Electronics devices (SL-HDR1)”, ETSI technical specification ETSI TS 103 433 (May 2016).
- [4] Diaz, R., Blinstein, S. and Qu, S., “Integrating HEVC Video Compression with a High Dynamic Range Video Pipeline”, SMPTE Motion Imaging Journal, Vol. 125, Issue 1., pp 14-21 (February 2016).
- [5] Luthra, A., François, E. and Husak, W., “Call for Evidence (CfE) for HDR and WCG Video Coding”, ISO/IEC JTC1/SC29/WG11 MPEG2015/N15083 (October 2015).
- [6] Luthra, A., François, E. and Husak, W., “Requirements and Use Cases for HDR and WCG Content Coding”, ISO/IEC JTC1/SC29/WG11 MPEG2015/N15084 (October 2015).
- [7] François, E., Fogg, C., He, Y., Li, X., Luthra, A. and Segall, A., “High Dynamic Range and Wide Color Gamut Video Coding in HEVC: Status and Potential Future Enhancements”, IEEE TCSVT, Vol. 26, Issue 1, pp 63-75 (July 2015).
- [8] François, E., “Standardizing the landscape, High Dynamic Range & Wide Color Gamut Video”, Proceedings of DVB scene, Issue 45, pp. 7 (March 2015).
- [9] ITU-R, “Reference electro-optical transfer function for flat panel displays used in HDTV studio production”, Recommendation ITU-R BT.1886 (March 2011).
- [10] Luthra, A., François, E. and van de Kerkhof, L. 2016, “Report of HDR Core Experiment 7: On the visual quality of HLG generated HDR and SDR video”, Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, JCTVC-W0027 (February 2016).
- [11] Pindoria, M., Naccari, M., Borer T. Cotton, A., “Some considerations on hue shifts observed in HLG backward compatible video”, Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, JCTVC-W0119 (February 2016).
- [12] Holm, J., “Information and Comments on Hybrid Log Gamma”, Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T SG16 WP3 and ISO/IEC JTC1/SC29/WG11, JCTVC-W0132 (February, 2016).
- [13] Lasserre, S., Le Léannec, F., Poirier, T. and Galpin, F. 2016, “Backward Compatible HDR Video Compression System”, In Proc. Data Compression Conference, pp 309-318 (March, 2016).
- [14] Miller, S., Nezamabadi, M., Daly, S., “Perceptual signal coding for more efficient usage of bit codes”, SMPTE Motion Imaging Journal, vol. 122, no. 4, pp 52-59, (May-June 2013).
- [15] Barten, P. G. J., Formula for the contrast sensitivity of the human eye, in Proc. SPIE – Image Quality and System Performance, vol. 5294, pp 231–238, San Jose, (December 2003).
- [16] Bjøntegaard, G., “Improvements of the BD-PSNR model”, ITU-T SG16/Q6, 35th VCEG meeting document VCEG-A111, Berlin, Germany (July 2008).
- [17] Sharma, G., Wu, W., Dalal, E. N., “The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations”, Color Research & Applications 30 (1), pp 21–30, (2005)
- [18] Baroncini, V., “Proposed test methodology of SDR quality evaluation with HDR reference”, MPEG document ISO/IEC JTC1/SC29/WG11 MPEG2015/M37407, Geneva, (October 2015).
- [19] International Telecommunication Union – Telecommunication Standardization Sector; “Recommendation ITU-T P.910 Subjective video quality assessment methods for multimedia applications” (September 1999).

Upstream Challenges With DOCSIS 3.1

White Paper

A Technical Paper prepared for SCTE•ISBE by

Jan Ariesen
Chief Technology Officer
Technetix Inc

Introduction

The introduction of DOCSIS® 3.1 means that higher modulation schemes, higher levels and higher bandwidth increase the load on the cable network.

The downstream increases from several QAM channels used in DOCSIS 3.0 in spectrum carrying up to 862 MHz with the addition of a couple of OFDM channels in spectrum expanding up to 1.2 GHz. All these changes greatly increase the load on the equipment used in the network.

However, this is nothing compared to what happens in the upstream. The big change is the bandwidth from the current low split of 42/54 MHz in the USA or 65/85 MHz in Europe to mid-split in the USA 85-105 MHz and high-split in Europe 204/258 MHz. In the USA, the bandwidth has grown by 216% and in Europe by 331%. In addition to the phenomenal bandwidth growth, increased loss in the cables is a reality and as a result, DOCSIS 3.1 increases the maximum level of the upstream carriers by 6dB to 57dBmV per 4 channels.

This growth in level and bandwidth creates four main challenges:

- Passive intermodulation (PIM) in in-home splitters
- 2nd order distortion in in-home amplifiers
- Overloading of the upstream amplifier in access amplifiers
- The downstream reducing the quality of the upstream in the access amplifiers

This White Paper summarizes the four challenges detailed and describes what to do to prevent possible problems.




Figure 1 - Cable network

1. Passive intermodulation (PIM) in in-home splitters

Intermodulation is a phenomenon which occurs in almost all components used in broadband networks, both active and passive devices. It is therefore important to ensure this is reduced for improved network performance. Intermodulation is caused by non-linear transfer characteristics of components.

1.1. Theory and background

In the Cable TV, non-linear components include ferrites and diodes in passives and transistors in active devices. In this section, the focus is on in-home splitters. A splitter is made with a ferrite core transmission line transformer.

This ferrite core transmission line transformer transfers energy from one port to the other via a magnetic field inside the ferrite core. See figure below:




Figure 2 - Splitter




Figure 3 - Non-linear transfer function due to core saturation




Figure 4 - Transfer characteristic of a ferrite

If due to surge, the input energy level (current) is too high the ferrite core becomes saturated. In cable TV applications, this energy is usually caused by a pulse, caused by lightning or a power surge, introducing a shift of the magnetic working point. Due to the saturation, the center point of the transfer characteristics of a ferrite moves up or down in the direction of the non-linear part of the graph. The result is that the transfer function is no longer linear and the signal becomes distorted.

In numbers:

The non-linearity changes the input-output characteristic of a splitter; the input output characteristic in the time domain of a linear system is given by:

$$y(t) = ax(t),$$

while this characteristic of a non-linear system can be approximated by a polynomial:

$$y(t) = a_1x(t) + a_2x^2(t) + a_3x^3(t) +$$

If a sinusoid with function $x(t)=A\cos(\omega t)$ is applied to a nonlinear system, multiples at the frequencies $2\omega, 3\omega, 4\omega, \dots$ arise at output $y(t)$ [1]. These multiples are called harmonics.

The products that arise when two or more different frequency signals are applied to a non-linear system are called intermodulation distortion (IMD) products. For example, if signal:

$$x(t)=A\cos(\omega_1 t)+B\cos(\omega_2 t)$$

is applied to a non-linear system, e.g. products arise at frequencies:

$$2\omega_1, 2\omega_2, \omega_1+\omega_2, \omega_1-\omega_2 \quad (= \text{second order IMD})$$

$$3\omega_1, 3\omega_2, 2\omega_1-\omega_2, 2\omega_2-\omega_1 \quad (= \text{third order IMD})$$

These are spurious unwanted products, which are located at frequencies of desired channels.

1.2. Measurement

The European Cenelec, "EN 60728-4:2008 Cable Networks for Television Signals, Sound Signals and Interactive Services

- Part 4: Passive Wideband Equipment for Coaxial Networks", describes the IMD measurements for 2nd and 3rd order products in passives. The measurement set up for IMD measurements is shown in Figure 1.

Two signals with different frequencies are applied to the output ports of a passive system, e.g. a splitter or a directional coupler. In the current specifications, two sine waves of 60dBmV at $F_1=60$ MHz and $F_2=65$ MHz are used. These are the two worst case scenario frequency channels of the upstream frequency band of 5 MHz – 65 MHz.

Using a spectrum analyzer, the 2nd order IMD product at 125MHz (F_1+F_2) is measured, as shown in Figure 2. In these measurements, the amplitude ratio between the fundamental product and second order product is measured.

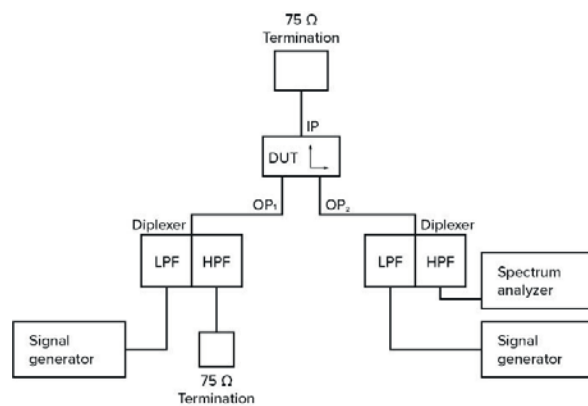


Figure 5 - Upstream IMD measurement setup




Figure 6 - Spectrum with IMD products

With a surge pulse to one of the ports, the ferrite is magnetized. In the test, the IMD is measured before and after the magnetization with a surge pulse:

- A demagnetized splitter can achieve approximately 110 dBc 2nd order intermodulation.
- A magnetized splitter will achieve less than 80 dBc 2nd order intermodulation.

If the frequencies of the carriers are at 34 MHz and 40 MHz, the second order is on 74 MHz, this is in the downstream band.

1.3. PIM results with DOCSIS 3.0

The upstream levels of DOCSIS 3.0 are 54dBmV with two carriers:

- The 2nd order product of these signals will be with a demagnetized splitter at -62dBmV
- The 2nd order product of these signals will be with a magnetized splitter at -32dBmV

With a downstream level of 0 dBmV this will create a C/N of 62 dB with a demagnetized splitter and 32 dB with a magnetized splitter. A magnetized splitter will already cause issues, as the CINR of the downstream signal will be critical for 256 QAM signals.

1.4. PIM results with DOCSIS 3.1

The upstream levels of DOCSIS 3.1 are 57 dBmV with four carriers:

- The 2nd order product of these signals will be with a demagnetized splitter at -50dBmV
- The 2nd order product of these signals will be with a magnetized splitter at -20dBmV

Table 1 - PIM by DOCSIS 3.x

	Upstream level 2ch	Intermodulation demagnetized	Intermodulation magnetized
DOCSIS 3.0	54 dBmV	-62 dBmV	-32 dBmV
DOCSIS 3.1	60 dBmV	-50 dBmV	-20 dBmV

With a downstream level of 0 dBmV this will create a C /N of 50 dB with a demagnetized splitter and 20 dB with a magnetized splitter. With a magnetized splitter communication will stop, as the CINR of the downstream signal will be too low.

The previous results have shown the PIM is a problem with DOCSIS 3.1 signals, this issue must be taken seriously. It is not just PIM levels that are an issue, there are many more carriers and therefore a greater PIM total. This combination with DOCSIS 3.1 increases the need for low PIM devices.

The increased frequency band to 1.2 GHz forces the splitter suppliers to go to smaller ferrites and these are even more sensitive to magnetization. To make a splitter that will meet the higher downstream frequency requirements and sustain a good PIM level is very difficult and costly. It is better to protect the ferrite from getting magnetized with a special circuit that will bypass the high RF signals but protect the low RF energy from surges and voltage pulses.

1.5. Conclusion: PIM in in-home splitters

- PIM has always been a critical specification for in-home splitters but the introduction of DOCSIS 3.1 dramatically increases the importance of this specification.
- PIM will reduce the quality of the downstream signals due to intermodulation from the upstream signals.
- There are good standardization methods to measure and qualify the in-home splitters.
- The new increased maximum frequency band of 1.2 GHz will force the splitter vendors to use smaller ferrites and these are more sensitive for PIM.
- It can be resolved with either a special protection circuit that protects the ferrite from being magnetized or using a more expensive high permeability ferrite.

2. Second Order Distortion in In-Home Amplifiers

With the addition of even more digital channels, cable network operators are particularly aware that in-home equipment should not degrade the performance of their network; both upstream and downstream.

2.1. Theory and background

The intermodulation and noise figure is important for the downstream amplifier stage but, it is the intermodulation that is key for upstream amplifiers.

In the upstream a poor intermodulation performance will result in:

- Generated beats in the downstream frequency spectrum but, these are normally filtered by the diplex filters
- Generated intermodulation products in the upstream spectrum




Figure 7 - Blockdiagram low split amplifier

The upstream (yellow) goes from the right side, via the 4-way splitter and the DPF (diplex filter) to the upstream amplifier. The upstream amplifier will also create 2nd order products adjacent to the upstream signal. This unwanted 2nd order product will be stopped in DPF 1 before it reduces the quality of the downstream (blue) signal.

As the levels of the downstream signals are low, it is important to use good duplex filters with enough isolation, otherwise these signals will reduce the quality of the downstream signal.

Signals below 15 MHz are not normally used due to the galactic noise creating too low a CINR noise. Therefore, signals are normally 15 MHz and higher.

Example 1, with a low split of 42/54 MHz,:

- $F1 = 20 \text{ MHz}$
- $F2 = 26 \text{ MHz}$
- 2nd order = $F1 + F2 = 46 \text{ MHz}$. This is outside the upstream band and is stopped by the duplex filter.

So, with a low split, the duplex filters are helping to prevent reduction in the quality of the signals.

If the split become mid-split 85/105 MHz or high-split 204/258 MHz, this becomes a different story.




Figure 8 - Blockdiagram mid split amplifier

With this split, the 2nd order products are in the same upstream band. For example 2, with a mid-split of 85/105MHz:

- $F1 = 20 \text{ MHz}$
- $F2 = 26 \text{ MHz}$
- 2nd order = $F1 + F2 = 46 \text{ MHz}$. This is in the middle of the upstream band.

This is also the case with the higher frequencies:

- $F1 = 76 \text{ MHz}$
- $F2 = 30 \text{ MHz}$
- 2nd order = $F1 - F2 = 46 \text{ MHz}$. Again, this is in the middle of the upstream band.

With a high split of 204/254 MHz, this only becomes worse. The 2nd order becomes much more important with a mid or high split than with a low split as IMD products are no longer out-of-band.

Note: this is a little different with a split of 65/85 MHz but also most of the 2nd order products are outside the band and will be filtered out by the diplex filter.

The best way forward is to design an upstream amplifier with reduced 2nd order products. There are several ways of building these amplifiers:

- Negative feedback amplifier
- Push-pull amplifier
- Balanced amplifier

All three have their own challenges, a summary of each follows.

2.2. Negative feedback

This is the most widely used circuit in cable TV amplifiers. It not only reduces intermodulation but also reduces gain and improves frequency response resulting in a wideband, flat frequency response amplifier. In fact, the improvement of the frequency response is usually the main reason for its use in cable TV applications.

Negative feedback also reduces noise and intermodulation. See figure 9 for the modelled amplifier.




Figure 9 - Negative feedback

The output voltage can be expressed as:

$$V_{of} = V_2 + V_{dn} \quad V_{of} = A_v \cdot V_1 + V_{dn}$$

$$V_{of} = A_v (V_i + V_f) + V_{dn}$$

$$V_{of} = A_v (V_i - K_f \cdot V_{of}) + V_{dn}$$

$$V_{of} = A_v \cdot V_i - A_v \cdot K_f \cdot V_{of} + V_{dn}$$

$$V_{of} (1 + A_v \cdot K_f) = A_v \cdot V_i + V_{dn}$$

$$V_{of} = \left[\frac{A_v}{1 + A_v \cdot K_f} \right] V_i + \left[\frac{1}{1 + A_v \cdot K_f} \right] V_{dn}$$

The conclusion is twofold:

- first part of the equation shows that the gain is reduced
- second part of the equation shows that the noise and distortion is reduced The equation shows that more feedback results in lower gain and lower distortion.

The most common form of negative feedback is the single transistor with resistive feedback circuit as shown in fig 10.

Analysis of the schematics show that feedback occurs through resistor (R_t) being in series with the input impedance. The output voltage (U_c) has a 180° phase shift with respect to the input voltage (U_b), hence the negative feedback of the circuit.




Figure 10 - Resistor as negative feedback

It is advisable for the collector current be made adjustable for best performance, since most transistors have a distinct point for lowest 2nd order distortion as a function of collector current. See figure 3 as an example of the BFG 135 specifications. The distinct minimum 2nd order point does not coincide with the lowest 3rd order intermodulation hence the need for an adjustable current.





Figure 11 - Intermodulation as a function of collector current (BFG-135)

2.3. Push-pull

This type of amplifier is widely used in all sorts of applications, including: cable TV power hybrids, GSM, linear RF power amplifiers, etc.

In cable TV applications this type is mainly used in hybrids; it is rare to find this type inside in-home amplifiers. There are exceptions where a push-pull architecture is used for the upstream amplifier or downstream amplifier.

Figure 12 shows a simplified circuit for analysis.

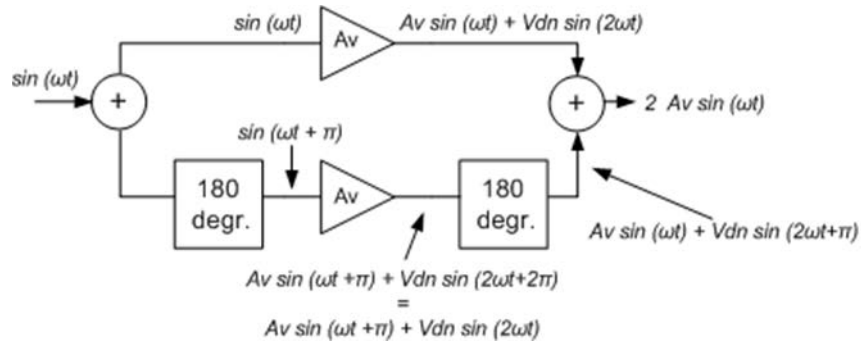


Figure 12 - simplified push-pull circuit

It is assumed the divider/combiner is lossless therefore the gain of the circuit equals $2 A_v$. In the real world, the divider loss is 3dB and therefore the gain of the circuit equals A_v .

From the figure above it can be seen that even (2nd, 4th, ...) order intermodulation products are cancelled. Odd (3rd, 5th,...) order intermodulation performance improvement is only marginal with two amplifiers in parallel.

In the practical world, cable TV type push-pull circuits use a balun as a combiner/divider. This is a relatively simple solution but, it has some drawbacks:

- Impedance of the amplifiers must be $75/2 = 37.5 \Omega$.
- Balance of the circuit (and therefore 2nd order cancellation) relies totally on the symmetry between the amplifiers used. A slight difference in amplifier impedance will result in imbalance because of the lack of isolation in the divider/combiner.
- Intermodulation performance is therefore influenced firstly input and output impedance and secondly the intermodulation performance of the individual amplifiers.




Figure 13 - BALUN as divider/combiner

2.4. Balanced amplifier

This type of amplifier is referenced only for completeness since the practical realization of the circuit has too narrow a bandwidth for cable TV applications.

The 90 degree delay lines used are frequency dependent, limiting the useable bandwidth of the amplifier. This type of amplifier is commonly used in microwave applications or in narrowband RF applications.

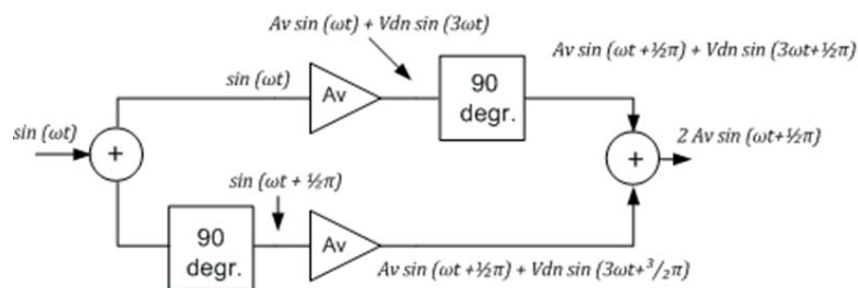


Figure 14 - Simplified Circuit of a Balanced Amplifier

As can be seen in figure 9, this type of amplifier has improved odd order intermodulation performance, while even order intermodulation is improved with 3dB w.r.t. a single amplifier. Some high-power RF amplifiers use this kind of circuit to combine two push-pull amplifiers achieving the best of both worlds (no even or odd order distortion products) but at the cost of 4 single amplifiers.

As previously stated, the cable TV band is wide and as it's not possible to create 90 degrees commercial delay lines, therefore this amplifier is unsuitable for cable TV implementations.

2.5. Conclusion 2nd order distortion in in-home amplifiers:

- Mid and high split drive the complexity of upstream gain stages in the in-home amplifiers.
- There are different styles and gain stages and all have their pros and cons.
- As the 2nd order by mid and high split becomes more important, the push-pull seems a good solution.
- The DPF in the in-home amplifier will prevent a lot of the intermodulation products but, this also increases the importance of a good DPF.

3. Overloading of the upstream amplifier in access amplifiers

3.1. Theory and background

An upstream amplifier in an access amplifier purpose is to overcome the loss in the subsequent cable.




Figure 15 - Signal Path Up and Downstream

Upstream signals (yellow) from the home arrive at the amplifier. The upstream signal at the second amplifier has the same input level as the signal on the first amplifier.

The signal on the output for the upstream of amplifier has a higher level to compensate the cable loss between the amplifiers.

3.2. DOCSIS 3.0 way of specifying the upstream performance

Historically, the performance of an upstream amplifier was provided via the so-called Noise Power Ratio method (NPR).

An upstream amplifier must meet a certain Carrier to Interference Noise Ratio (CINR) on a specified input level, over a specified input window, with a specified gain and over a specified frequency range. Such as:

- Input power $P=6 \text{ dBuV}/\sqrt{\text{Hz}}$
- Input window $\pm 12 \text{ dB}$
- Gain = 25 dB
- Frequency 5 – 65 MHz
- Min CINR, 50 dB

This results in performance of the upstream amplifier where, if the input level is between $-6 \text{ dBuV}/\sqrt{\text{Hz}}$ and $18 \text{ dBuV}/\sqrt{\text{Hz}}$, the output performance is always more than 50 dBc.




Figure 16 - CINR versus input power graph

3.3. DOCSIS 3.1 way of specifying the upstream performance

With the new DOCSIS 3.1 standard, the same performance is often requested but, only with an upstream band until 204 MHz and limited gain in dB's. This is because of the additional loss at higher frequencies, resulting in:

Input power $P=6 \text{ dBuV}/\sqrt{\text{Hz}}$

Input window $\pm 12 \text{ dB}$

Gain = 28 dB

Frequency 5 – 204 MHz

Min CINR, 50 dB

The difference between a channel of 6.4 MHz and $\sqrt{\text{Hz}} = 10\log 6,400,000 = 68 \text{ dB}$ This results a maximal power per channel of 6.4 MHz of:

$P_{\text{out}} = 6 \text{ dBuV}/\sqrt{\text{Hz}} + 12 \text{ dB (input window)} + 28 \text{ dB (gain)} + 68 \text{ dB (6.4 MHz)} = 114 \text{ dBuV} = 54 \text{ dBmV}$
204 MHz has 30 channels, this results in a total composite power of:

$54 \text{ dBmV} + 10\log 30 = 69 \text{ dBmV}$

This is too much power to be created in a mimic and a hybrid is required with an additional power consumption of 8 Watts.

In reality, this is never the case:

- 1) Generated beats in the downstream frequency spectrum but, these are normally filtered by the diplex filters
- 2) Generated intermodulation products in the upstream spectrum and a NF of 8 dB and NP of 2 dBuV the C/N towards the node of the 10 amplifiers together:

$C/N_{\text{upstream}} = \text{Input level} - (8 \text{ dB (NF)} + 2 \text{ dBuV (NP)}) - 10\log 10 (\text{Nb of amplifiers})$

$= \text{Input level} - 20 \text{ dBuV}$

With an input level of 10 dBmV, the C/N from the access network is still 50 dB. This means it is nil versus the noise from the home and this will lower the input but more importantly the output level of the upstream amplifier.

With the 2 points above, the output composite power on the output will drop a lot and this will make it feasible to use a normal mimic instead of a hybrid and this will improve the quality and reduce the power consumption.

3.4. Conclusion: overloading of the upstream amplifier in access amplifiers

- The current method of setting levels and specifying upstream amplifiers in access amplifiers will drive power consumption too high.
- The level currently used for upstream can be reduced because less amplifiers are combined behind an optical node, due to segmentation in the network.
- The new DOCSIS 3.1 channel drives a new way of alignment of the cable TV network for the upstream.

4. Downstream: reducing upstream quality in access amplifiers

4.1. Theory and background

A block diagram of an access amplifier is given below:




Figure 17 - Block Diagram Access Amplifier

As described in section 4, the input level of the upstream amplifier is approximately 10 dBmV.

The downstream signal of an access amplifier can be up to 50 dBmV in the real world.




Figure 18 - Block Diagram Access Amplifier with Levels

The isolation of a normal duplex filter is approximately 50 dB, this means that the downstream signal reaching the upstream amplifier is approximately 0 dBmV.




Figure 19 - Block Diagram Access Amplifier with Levels

The downstream signal will load the upstream amplifier and this will create harmonics in the upstream band.




Figure 20 - Block Diagram Access Amplifier with Levels

This will reduce the CINR in the upstream and, as a result, reduce the capacity.

4.2. Measurements

Extensive testing has been done to support the above issue:




Figure 21 - Test Setup

The upstream connected to the DUT via the DPF 204/258 has a gap at 119 MHz. In this channel, we measure the CINR. To be certain we don't overload the analyzer, we first measure the level of the upstream without filter and then the noise and interference in the gap is measured.

Results

Table 2 - Upstream CINR versus Downstream level

downstream signal (dBmV/6MHz)	CINR
-10	42.01
-8	42.01
-6	42.01
-4	42.01
-2	42.01
0	42.01
2	42.01
4	42.01
6	42.01
8	41.61
10	40.41
12	38.91
14	35.91
16	30.51
18	23.21
20	17.41
22	9.81

The results above show the upstream amplifier starts to produce intermodulation products caused by the downstream signals above 10dBmV in the upstream band. These intermodulation products are caused by the downstream and if multiple amplifiers are added on one node, the interference will add with 10 log the number of amplifiers.

New style amplifiers often have a better push-pull and more isolation than 40 dB in the duplex filters. This will result in a lower downstream signal coupled to the upstream amplifier and the push-pull can better handle the loading. Both improvements will reduce the impact of the downstream signal on the upstream quality.

4.3. Conclusion: downstream reducing the quality of the upstream in the access amplifiers

- Poor isolation in diplex filters will reduce the quality of the upstream due to loading of the upstream amplifier with downstream signals.
- The downstream intermodulation products in the upstream will look like CPD.
- This is solved by good DPF and amplifier stages.

Conclusion

Upstream needs with DOCSIS 3.1

There are several challenges to driving upstream through the existing cable TV network but, all of these challenges are resolvable with the latest technology.

PIM in splitters

- PIM are intermodulation products from the upstream signals that reduce the quality of the downstream signals.
- Standards drive the measurability.
- PIM can be reduced by better ferrites or special protection circuits.

2nd order in-home amplifiers

- Intermodulation products in in-home amplifiers are products created by upstream signals reducing the quality of upstream signals.
- The 2nd order problem in upstream amplifiers becomes critical with mid and high split.
- There are several types of in-home amplifiers and they all have their pro's and con's.

Overloading of the upstream amplifier in access amplifiers

- Overloading of the upstream amplifier in the access network can be solved by a new method of specifying and alignment of the upstream signals.
- The segmentations mean that levels on the input of the upstream amplifiers can be reduced.

Downstream: reducing the quality of the upstream in the access amplifiers

- The upstream amplifier can be overloaded by the downstream signals which will reduce the quality of the upstream signals.
- This problem can be solved by good DPF and the latest state-of-the-art amplifier technologies.

Abbreviations

CINR	Carrier to Interference Noise Ratio
CNR	Carrier to Noise Ratio
CPD	Common Path Distortion
DOCSIS	Data Over Cable Service Interface Specification
DPF	Diplex Filter
GHz	Giga Hertz = 1,000,000,000 Hz
IMD	Intermodulation
MHz	Mega Hertz = 1,000,000 Hz
NF	Noise Figure
NP	Noise Power = kTB (k = Boltzmann's constant, T = temperature in Kelvin, B = Bandwidth)
NPR	Noise Power Ratio
OFDM	Orthogonal Frequency Division Multiplexing
PIM	Passive intermodulation
QAM	Quadrature Amplitude Modulation
RAF	Radio Frequency

A Proposed End-to-End SDN Architecture for MSO

A Technical Paper prepared for SCTE•ISBE by

Mohcene Mezhoudi

Principal, Bell Labs CMTS
Bell Labs Consulting
600 Mountain Avenue
Murray Hill, NJ. 07479600
908-582-2776

Mohcene.Mezhoudi@bell-labs-consulting.com

Benjamin Y. Tang

Principal, Bell Labs DMTS
Bell Labs Consulting
401 Data Drive
Plano, TX. 75024
469-910-3698

Benjamin.Tang@bell-labs-consulting.com

Jean-Philippe Joseph

Principal, Fixed Access Domain Leader
Bell Labs Consulting
600 Mountain Avenue.
Murray Hill, NJ. 07479
908-679-5798

Jean-Philippe.Joseph@bell-labs-consulting.com

Enrique Hernandez-Valencia

Bell Labs Fellow, Partner
Bell Labs Consulting
600 Mountain Avenue.
Murray Hill, NJ. 07479
908-582-3144

Enrique.Hernandez-Valencia@bell-labs-consulting.com

Introduction

The advent of Software Defined Networking (SDN) is becoming strategically critical for telecom/datacom and information/cloud Service Providers worldwide. Discussion and writing on the application of SDN in MSO networks have been limited to specific parts of the network. SDN leverages a control plane model with open programmable interfaces to enable the agile and dynamic creation of new services and the orchestrated allocation of network resources. Through this, and by bridging the SDN architectures and solutions between broadband access, hub/headend sites, metro and backbone transport, MSOs will fully leverage SDN principles and their end-end networks. Further, SDN can address the following challenges faced by many MSOs:

- **Footprint** – Traditional regional boundaries limit ability to provide national or global services. Recent mergers and acquisitions to expand MSO footprint may impact Service Level Agreements (SLAs) and resiliency in the consolidated network. SDN can help mitigate this through a combination of telemetry, data analytics and proactive performance/resiliency management.
- **Scale/optimization**– Today’s rigid network architectures increase service delivery costs. SDN provides flexibility in dynamically scaling and optimizing resource utilization across IP/optical and throughout metro and backbone networks and in broadband access.
- **Growth** – Offering services primarily to residential customers limits growth. MSO expansion into new markets and offering competitive new service offerings will enable new growth opportunities. SDN enables the rapid creation of new innovative, revenue-generating services by using solutions such as dynamic and automated service provisioning, network slicing and service chaining end-to-end.
- **Capacity** – Bandwidth growth is straining broadband access networks. New architectures and technologies such as DOCSIS® 3.1 deep fiber, node splitting, and symmetrical Full Duplex (FDX) DOCSIS will significantly increase broadband access network throughput. SDN can help manage these new complex networks to reduce migration and operations cost.
- **Latency** – As new ultra-delay-sensitive applications emerge there will be a need to minimize end-end latency and maximize resiliency performance and user experience. This can be accomplished by moving latency sensitive functions and applications closer to the end-user and placing them under end-to-end SDN control.

We first present an MSO future network vision with key building blocks cemented by an end-to-end SDN architecture. We propose a multi-layer end-end SDN architecture that can be achieved by implementing an SDN controlled broadband access network and extending SDN to the metro and backbone network with integrated IP/optical transport. We argue for the importance of convergence in both the broadband access and transport network, the metro portion of the MSO network, and the criticality of an SDN-based IP/Optical transport integration. Convergence and integration, which are greatly aided by network virtualization through SDN/NFV, can provide MSOs with an agile, hyper-scalable and cost-effective network fabric to enable automated management, control and optimization of network resources, and the dynamic provisioning of all revenue-generating services. The roles and functions of SDN in broadband access and metro/backbone, as well as the cost benefit of IP/optical integration – in particular, multi-layer protection, will be presented in the remainder of this paper. We conclude by presenting end-to-end SDN based service orchestration that can greatly benefit MSOs in future service delivery.

Future MSO Network Vision

We envision that the future MSO SDN-enabled metro and broadband access network will be built based on four building blocks as depicted in Figure 1.

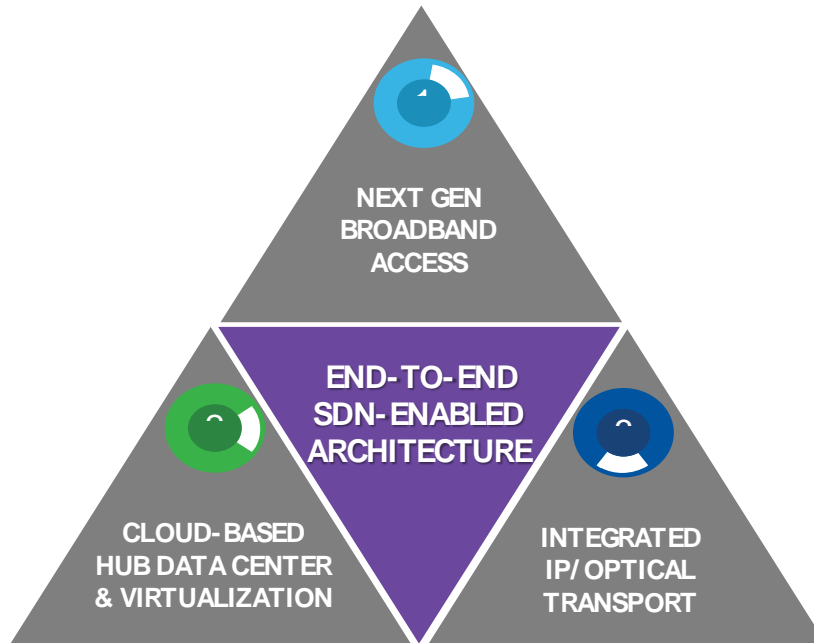


Figure 1 - Building Blocks of Future MSO Network

1. MSO Future Network Building Blocks

Next-Generation Broadband Access

There are multiple initiatives underway to build next generation broadband networks. A brief survey is provided here as background. The most fundamental initiative is the deployment of fiber deeper into the broadband access network and closer to end-users. This is leading to a reduction in the average number of coax amplifiers to N^{1+x} , where $x=6,5,4,3,2$, enabling multi-Gbps downstream service access rates (in conjunction with migration to DOCSIS 3.1). Deep-fiber architectures are referred to as Distributed Access Architectures (DAA).

DOCSIS 3.1 is in deployment with many MSOs and will enable higher Upstream (US) and Downstream (DS) bandwidths through improved modulation techniques, and move the US/DS split to provide additional US capacity. Full Duplex (FDX) DOCSIS, currently under specification development, will require a N+0 (i.e., fiber node deployed at the last amplifier location) and will theoretically provide 10Gbps symmetrical bandwidth (initially FDX will operate in a reduced spectrum and provide 4-5Gbps).

¹ N=the fiber node

Analog spectrum is being reclaimed for use in other services such as high-speed data. Many MSOs are expanding or considering expanding outside plant (OSP) spectrum to 1 or 1.2 GHz spectrum to create additional channels used for services such as high-speed data.

MSOs are investigating providing wireless services (licensed and unlicensed spectrum) beyond Wi-Fi. In providing such services, wireless broadband access points or small cells will need to be brought to within 100 to 200 meters of end users, with fiber backhaul of traffic into the network. Similarly, DAA nodes are also brought within about 100 meters of end users, creating an ideal location for a common platform that serves both wireless and wireline access technologies [1].

Cloud-Based Hub Data Centers and Virtualization

The future MSO hub/headend locations will also evolve to a Data Center-like architecture, or edge clouds, where Virtualized Network Functions (VNF) and Physical Network Functions (PNF) will be deployed to meet the performance requirements for future applications and to offload bandwidth in the transport network. For instance, the future MSO edge cloud must be located close enough to end-users to meet low latency requirements for critical applications such as Cloud RAN fronthaul, 5G enhanced broadband, and industrial IoT.

In general, the edge cloud site will host both the mobile and fixed network functions (e.g., vRAN, vS/PGW, vOLT, vCCAP) and common applications (vDVR, vCDN, etc.), which are virtualized on Commercial Off-The-Shelf (COTS) platforms, general-purpose servers, and specialized hardware accelerators. Within the edge cloud location, leaf and spine intra-hub networking architectures are leveraged to connect servers, compute platforms and storage platforms. The edge cloud and the broadband access network elements are under the control of a Software Defined Networking (SDN) controller. SDN with its separation of control and data planes, service/resource abstraction, open interfaces and programmability, provides MSOs with great flexibility and agility to deal with dynamic applications and services and drive service innovation in the most operations efficient and cost-effective fashion.

Integrated IP/Optical Transport

A smart IP/Optical transport network interconnecting edge and core cloud sites is a fundamental component of the future MSO network vision. The benefits, enablers, and challenges of IP and Optical transport integration have been well understood by network operators [2] for some time. IP and optical integration can be achieved in multiple dimensions: data plane integration (colored interfaces), multilayer control and resilience, multilayer planning, and management. Main motivations for adopting IP/optical integration are improved network efficiency across layers, reduced resources required for protection and hence TCO savings, and optimized network management and planning. One major use case of IP/optical integration is multi-layer protection taking full advantage of available protection mechanisms from IP and optical to minimize the cost of protection/restoration.

Key technology enablers for IP/optical integration include GMPLS/GMPLS-UNI allowing tight interaction between IP and optical and an automated and switchable photonic layer. The requirement of GMPLS-UNI, however, presents a risk of vendor lock-in as the physical integration of IP and optical layers might require a single vendor in most cases. With SDN-inspired protocols such as PCE, BGP-LS, and Segment Routing (now emerging as a more promising and future proof control plane technology enabler than plain GMPLS), IP/optical integration has received renewed and growing interest from operators. These new tools largely eliminate the challenge of vendor interoperability by reducing the

amount of service state maintained in the transport nodes. They also support open APIs to provide the basis for interoperability between network layers as well as between the network and orchestration layers in a multi-vendor environment. These SDN based protocols also add capabilities to rapidly instantiate multi-layer network element provisioning and turn up of bandwidths, and services, on demand. With SDN-enabled IP/optical integration, MSOs can create a scalable, reliable and cost effective tunable network fabric [3] with adaptive multi-layer resource optimization, automated network/service management, and dynamic service assurance.

2. Defining an End-to-End SDN Architecture for MSOs

SDN allows operators to effectively address multi-domain, and multi-vendor interoperability. While SDN implementation saw initial deployment in the enterprise data center space, it has quickly expanded to Next Generation Central Offices (NGCOs), the backbone and metro transport, and now the broadband access networks. Many of the benefits rendered by SDN regarding basic transport and connectivity services have been well documented. An industry survey [4] identified the top three SDN benefits sought by operators as automated service provisioning, service assurance through proactive monitoring and management, and IP/optical multi-layer optimization. Combined with segment routing, PCE, and traffic steering techniques, SDN allows for optimal allocation of network resources to meet dynamically changing service bandwidth demands and SLA requirements. For MSOs, an integrated SDN framework across broadband access, metro/backbone transport, and hub/headend (edge/core cloud) sites brings flexible end-to-end control and management of cloud-based applications and services, and hence, their agile, reliable and cost-effective interconnection.

For cable broadband access networks, the most critical benefits of SDN for an MSO – as identified by CableLabs [5] – are the extension of a common set of protocols and APIs to enable software programmability, automatic provisioning and management of network devices across broadband access and transport technologies. These new tools enable rapid creation of new services – based on the network element and service abstraction via UML and YANG data models and NETCONF procedures to manipulate the configuration of network devices, open interfaces, service chaining and intent-based networking analysis.

The proposed MSO future end-to-end network architecture, based on the building blocks described, is illustrated in Figure 2. This SDN-based architecture framework abstracts the network capabilities into “horizontal” domains such as access, metro and backbone transport, and hub/headend sites, as well as “vertical” domains – or operational layers - such as products/service, logical resources and physical resources (aka, the infrastructure).

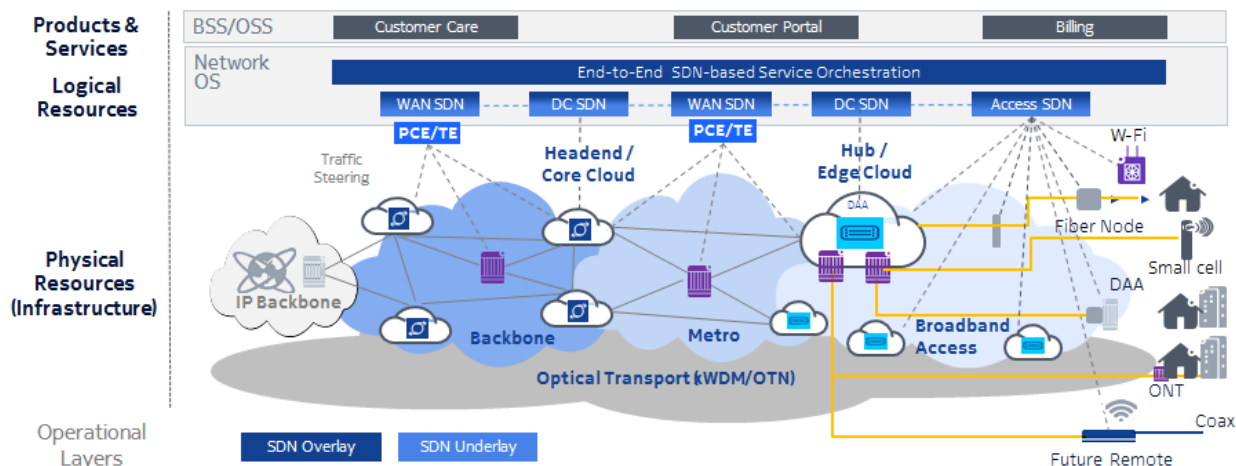


Figure 2 - MSO future SDN-based network architecture

From the network services viewpoint, a minimum of three “vertical” operational domains are envisioned:

- Physical resources: CPEs, network termination equipment (Cable modems, wireless access points), access aggregation nodes, IP edge functions, transport switches/routers and the servers/storage nodes supporting any of the network functions and applications.
- Logical resources: representing the logical abstraction of the network resources, which need also provide the ability to support “horizontal” decomposition into functional network domains such as broadband access, transport, and hub/headend sites. This functional decomposition is illustrated through the corresponding “Access”, “WAN” and “DC” SDN controllers.
- End-to-end SDN network service orchestration: including any connectivity required among communicating end-users, as well as the connections to any access any network functions – physical or virtual – delivering any contracted network services.

Note that the various SDN controllers illustrated here are intended to represent logical functionality. They may be implemented as separate logical entities, or as a single integrated controller, depending on performance, scaling and administrative needs.

3. Network Decomposition and Slicing Usage of Styles

Besides network connectivity, SDN also plays a pivotal role in abstracting network services and in managing and optimizing the consumption of the resources provided by the network infrastructure. Hence, consistently with SDN principles, it is necessary for the logical components of the network architecture to support the ability to further decompose any of these vertical layers into sublayers as may be required to better integrate specific network technologies or administrative needs.

In the case of the broadband access, metro and backbone transport infrastructure, it should be possible to further decompose the physical media (fibers, coax, etc.) from the optical/wavelength links and the IP flows that constitute the ultimate physical *underlay* of the transport infrastructure. This logical decomposition, on top of these physical infrastructures, will lead to many logical *overlays*, typically instantiated through L2/L3 packet tunneling technologies that provide the logical connectivity between end-users and network function end-points, and hence, instantiated the network slices that enable flexible,

dynamic and elastic delivery of services (e.g., SD-WAN, service chaining) as opposed to rigid and expensive separately built networks, as illustrated in Figure 3.

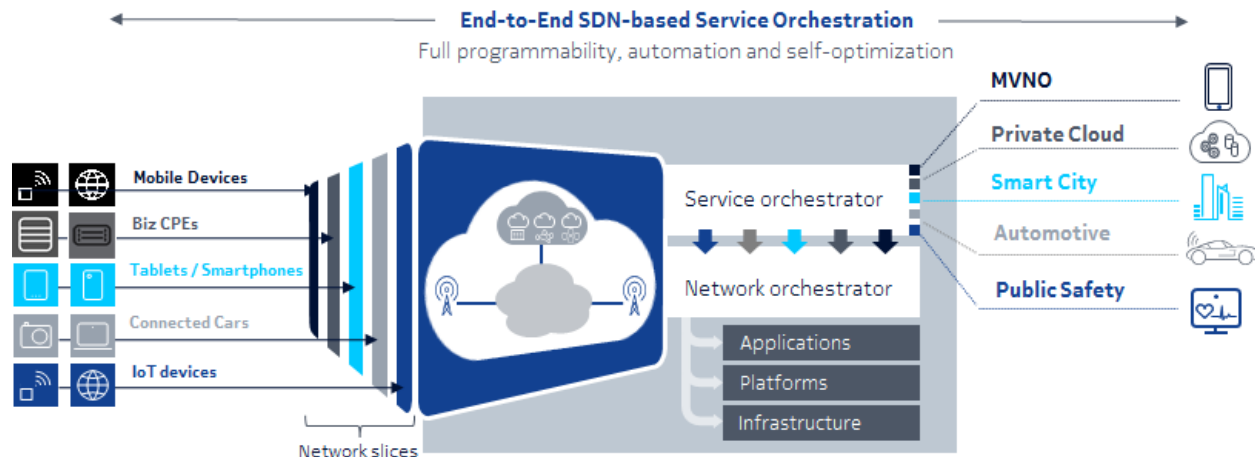


Figure 3 - Network slicing instantiated by SDN

SDN in Broadband Access

4. DSN Broadband Access Architectures

Today's technological trends such as the industrial Internet, the Internet of Things (IoT), smart home, streaming video content, bandwidth-on-demand and app-centric, mobile-friendly personalized communications are taking roots in today's environment - and the market competition continues to experience the entrance of service providers ready and willing to deliver these service packages to meet consumers demands.

MSOs are seriously evaluating architectural options to evolve their broadband access network to meet these trends and future demands for high bandwidth services, to improve performance, service agility, and achieve operational efficiency while reducing their overall cost structure. A key enabler of the future broadband access network is the introduction of SDN. Among the promises of SDN in the MSO's broadband access network include the following:

- **Open eco-system** – Today's broadband access network is characterized by an inflexible and closed environment with a monolithic access node, proprietary hardware and software. As a result, there is no direct access to functions, data, application and service layers, impeding flexibility and limiting scalability. The future broadband access is an SDN-enabled and open eco-system, which enables innovations and purpose-built applications in a multi-vendor and open-source environment. Further, the basic SDN architecture separates the management and control planes from the forwarding plane, thus allowing each part of the network to be independently optimized. The underlying network infrastructure is abstracted from the applications and network services, but the communication between layers is achieved with an open application layer interfaces (APIs) to allow direct programmability of each component of the network. Last, these

planes can be virtualized allowing various NVFs/PNF to be distributed to different locations in the broadband access network.

- Network slicing – As the future broadband access network is multi-service, medium agnostic and programmable, the ability to create virtual slices of the network infrastructure to support wholesale, differentiated services, to dynamically allocate the network resources and expose APIs for innovative service creation, becomes critical for MSOs to evolve their business models and tap into new revenue growth opportunities.
- Node management – The ability to deliver high-bandwidth services requires the MSO to re-engineer broadband access networks and adopt a fiber-deep strategy that places smaller remote nodes closer to the customers, sharing the available bandwidth among fewer users. As a result, the number of serving nodes will significantly increase and, perhaps, by one order of magnitude relative to the current Fiber Node (FN) to achieve an amplifier-free OSP with nodes placed at the last amplifier locations (n+0). Thus, a most robust node management system is required to minimize human interventions in the OSP and allow the MSOs to reduce related OpEx. SDN will enable automation and zero-touch provisioning, and automated activation, thus drastically reducing truck rolls and OpEx.

5. I-CCAP Evolution to Distributed Architectures

The current MSO hub architecture is based on CCAP, which combines the traditional CMTS with high-speed data functions and edge QAM functions used for video (thus Integrated-CCAP or I-CCAP). The natural evolution of HFC networks is driving fiber deeper due to SG size reduction and corresponding physical FN splits. DOCSIS 3.1 is designed to enable major gains in modulation, some that require shorter coax runs with fewer or even no amplifiers in the cascade. Legacy FN technology is adequate to leverage these realities. However, the increase in the number of SGs drives a need for additional space, power and cooling in the headend and hub locations where the CMTS+EQAM and/or I-CCAP equipment resides. There is also the issue of analog optics between the hub and FN being at capacity and operational expenditure (OpEx) concern as capacity requirements on feeder fiber increase along with SG growth in an FN Serving Area (FNSA).

Distributed Access Architecture (DAA) has arisen as the next architectural advance in cable broadband access networks [6]. It considers all reasons mentioned previously as well as the conclusion that digital feeder fiber provides a ubiquitous platform which can support all fiber deeper requirements of the broadband access network, including fiber to the premise cases. A common requirement in all DAA variants is digital fiber between the headend/hub and FN location. Where DAA variants differ is in the distribution of the functional elements of I-CCAP. One variant is R-PHY or RPD, which remotes the common PHY used by all HFC-based voice, data and digital video services into the next generation FN. R-PHY leaves the MAC at the existing head-end/hub CCAP location. The other main DAA alternative remotely locates both MAC and PHY in the FN to an R-MAC/PHY or (RMD) node. The RMD alternative enables the majority of remaining CCAP functionality, which is largely implemented in software, to be placed anywhere in the cable operator network.

DAA in its different flavors is currently maturing in the specification stage at CableLabs® with product expected in 2018.

6. SDN/NFV-based Broadband Access Solution Alternatives

MSOs are considering various options to introduce SDN into the broadband access network. SDN decouples hardware (e.g., PHY layer) and software (e.g., control plane). The software can be moved to the edge cloud or instantiated wherever needed in distributed servers, independently of hardware platforms. Likewise, virtualization enables network functions to migrate from dedicated hardware to virtual machines running on general-purpose hardware and specialized hardware accelerators. Thus, SDN provides the flexibility to centrally reconfigure networks in an automated fashion. A key SDN capability is to enable automated on-demand provisioning of services in the broadband access network.

Multiple SDN broadband access solutions have been recently proposed, inspired, and stimulated by the SDN paradigm successes in data centers. SDN carries the promise of enhanced and expanded services with a more manageable operational environment. Though many standards groups and open-source projects exist, deployments remain in their infancy.

In the following, we survey a sample of these proposals with an emphasis on broadband access architectures.

CableLabs introduced in [7] the basis for an overall SDN architecture model that applies to any underlying cable broadband access technology such as DOCSIS, EPoN etc. The proposal focuses on the data models and south-bound protocols, as the foundation for how the MSOs will deploy and manage their future programmable broadband access networks. It highlights the SDN controller and orchestrator functionalities to simplify the MSO network operations by abstracting out the complexity of the network devices and their configuration. The claim is that the gain from an SDN controller which enables programmable configuration, provisioning and dynamic control of various network devices and elements, will liberate operators to focus on developing new and revenue-generating applications and integrate them with the existing services seamlessly.

The authors in [8] emphasize that for the distinct characteristics of DOCSIS broadband networks (protocol, topology, and management capabilities), a new set of unique requirements is needed and modifications/adjustments must be applied to the existing SDN model in data centers and other networks. They present an overview of Comcast's SDN-driven implementation of programmable Service Flows through the integration of the OpenDaylight (ODL) PacketCable® Multimedia (PCMM) plugin. They acknowledge the progress demonstrated by this platform as a control plane for the creation of DOCSIS/PCMM service flow management. However, they stress that while further development in practical and scalable, deployment/implementation and field experience are still extremely essential, the path forward for SDN in the new programmable DOCSIS network is better defined and clear.

An SDN/NFV framework is presented in [9] to enable MSOs to deliver a set of new services (network access, infrastructure, consumer services and 3rd party applications) across multiple broadband access networks technologies (DOCSIS, DPoE, PON, Wi-Fi, Mobile, etc.). They elaborate on how operators would offer and manage them. They suggest a broadband access network architecture where the broadband access technology is abstracted from the services layer, allowing the technology vendors to optimize the ASIC based solutions effectively and only migrate the services and control/management plane to a COTS platform with compute, storage and switching capabilities. This will enable both operators and vendors to offer feature parity and reuse of certain components independent of the access

technology, thereby improving time to market for new services and features. In addition, they present views and assessment on what network functions can be virtualized to improve the agility of the broadband access network. And observe that while MSOs are experiencing multiple transitions (architectures, technologies, services, app-aware), they suggest that the SDN/NFV paradigm could be their means to realize synergies while addressing these transformations. They recognize SDN/NFV technology provides enormous potential benefits though it is still maturing and has achieved a lot of progress with vendors and operators actively proving its feasibility and performance capabilities and conducting real-network trials.

The integration of SDN concepts and the CableLabs DOCSIS, L2VPN, and DPoE definitions is investigated in [10]. The objective is to enable cable operators to realize fully automated end-to-end provisioning of commercial services and dynamic network reconfiguration through a single and simplified provisioning interface. The authors describe an overall broadband access architecture and its key elements. They define a set of requirements for the service provisioning API and the capabilities of DOCSIS-approach provisioning. They outline an automated DOCSIS-type provisioning framework with SDN. It is proposed that a network-layer service could be fully specified by the business systems by providing the appropriate attributes to all network elements. They assert that the model could be extended to other broadband access technologies beyond DOCSIS and DPoE technologies with further specifications of the requirements and the actions needed to be performed through the API. They supplement their approach with a description of a provisioning flow from the business service system (BSS) to the Cable Modem (CM). This involves multiple systems and human interventions. They proclaim that the proposed service provisioning framework allows the MSO to offer better-quality services at lower cost. This is achieved by reducing new services lead-time to market and eliminating manual and error-prone configuration of network equipment. Particularly for more involved business offerings where the service provisioning is reduced to a transaction rather than multiple steps with multiple complex human interaction and intervention.

The use of SDN Broadband access concepts to enable dynamic and flexible pairing of a separated MAC-PHY and R-PHY in a remote physical device (RPD) architecture is presented in [11]. The agile SDN infrastructure assigns various remote physical devices to various MAC resources (local or remote CCAP, or virtual CMTS) based on policies and network varying requirements. This results in new options for network operations savings as well as new load sharing and availability performance.

In this subsection, we focus on the introduction of SDN in the DAA-based broadband access network. Considerations pertaining to enabling SDN in an FTTH network, which MSO might deploy as an overlay network to serve their customers, are not addressed here. Among the DAA solution alternatives for the future HFC-based broadband access network, the following have gained more acceptance in the industry: non-virtualized Distributed CCAP Core (CCAP Core RPD) and virtualized DAA (vDAA). vDAA has two implementations: 1) Virtualized CMTS in a centralized location (vCMTS RPD); and, 2) Virtualized CMTS in the node (vCMTS RMD).

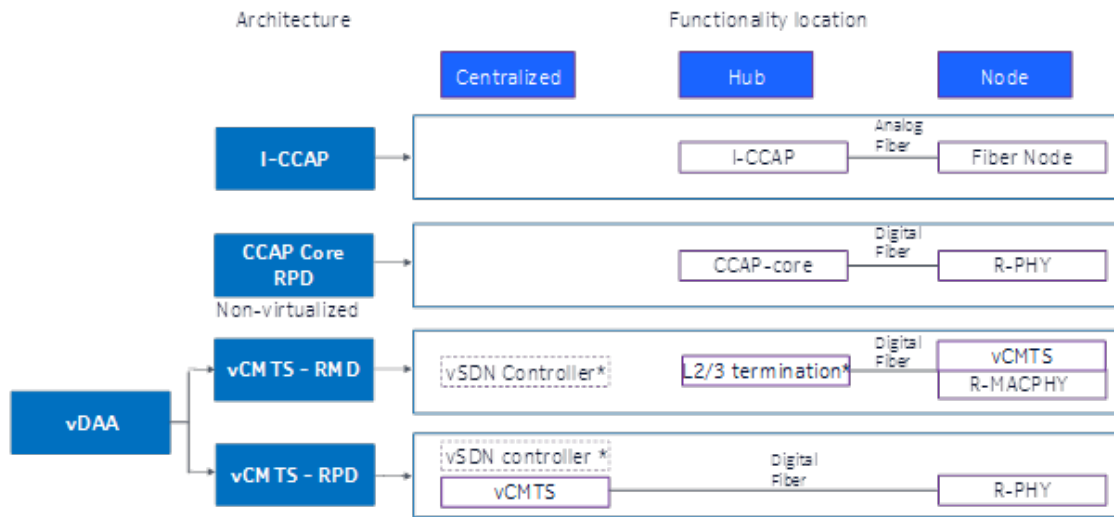


Figure 4 – DAA Options

As shown in Figure 4 below, the CCAP Core RPD solution leverages a CCAP core or DOCSIS MAC in the hub location and decouples the PHY function, which is deployed as RPD nodes in the OSP.

vCMTS RMD introduces a virtualized SDN controller (vSDN), which is deployed in a centralized location, hub or headend, and remoting the combined MAC and PHY layers to an RMD node deployed in the OSP. The most obvious benefit of this solution involves trade-offs among complexity, space, and energy consumption.

vCMTS RPD leverages a virtualized CMTS (vCMTS) function that can be deployed centrally in a hub/headend to support RPD nodes. This solution provides complete flexibility to the MSOs in the DAA space. A vSDN controller is also used and can be deployed in the hub or headend location.

SDN in Metro Networks

A Bell Labs study [13] shows that consumer and enterprise traffic in the metro network is expected to grow by 3.9x from 2015 to 2020, driven by video and cloud/data center traffic. In addition, cloud-based applications (e.g., network DVR, gaming, video streaming, and future AR/VR) and dynamic placement of virtualized network functions (e.g., vCDN, vRouter, vCCAP/vCMTS) deployed in distributed edge cloud DCs will lead to a substantial change in the traffic patterns of metro networks. This increase in network flexibility coupled with the trends towards real-time user-controlled service experience - such as self-portal for dynamic change of service attributes (e.g., bandwidth and performance/SLA), and the emergence of richer IoT and 5G applications based on network slicing, will create very dynamic and elastic traffic demands. These demands are not only North-Southbound (between servers and clients) but also have East-Westbound traffic (between data centers). Delivering such large and rapid changing traffic

demands will require a flexible transport network fabric, particularly across the critical metro networks. These trends will require MSOs to:

- Scale network capacity cost efficiently to sustain future traffic growth
- Allocate network resources dynamically to meet rapidly changing applications and services demands
- Automate network configuration and service provisioning in a multi-vendor, multi-layer environment
- Provide dynamic assurance of network/service performance and efficiency through use of KPIs and data analytics

IP/Optical transport integration via multi-layer protection, adaptive network resource optimization, and SDN based principles, protocols and APIs will provide MSOs a flexible and smart network fabric over which to implement an end-to-end SDN-based service delivery architecture to achieve sustainable network capacity and revenue growth. An SDN-based architecture framework further enables network programmability and automation to provide the flexibility and agility required to achieve dynamic and on-demand service provisioning and assurance.

7. Multi-layer Protection and Adaptive Resource Optimization

Several shortfalls exist today in most backbone and metro networks:

- Most operators are deploying IP and optical transport separately without taking full advantage of multi-layer protection mechanisms available – many of them are using IP protection and restoration mechanisms (MPLS, IP FRR or simply routing convergence) without leveraging optical protection capabilities. This approach risks higher cost of protection and missing critical SLA requirements (e.g., shorter failover time which can only be achieved by optical 1+1 protection).
- Furthermore, there is no differentiation of protection – traffic from all service classes is receiving the same level of protection in the IP/optical transport network, making it unnecessarily expensive for operators and results in a lost revenue opportunity for premium services with higher resiliency.
- IP and optical transport have a rigid network topology which does not adapt to dynamic traffic patterns and loads. This leads to inefficient routing and network resource utilization. For example, operators should be able to cost-effectively and dynamically establish IP shortcuts through the optical layer to address short-term and/or seasonal high-volume traffic demands more cost effectively.

Failure to address these shortfalls lead to an over-engineered network and high per-bit transport cost that is not sustainable to support future traffic growth. MSOs should develop a comprehensive strategy and solution toward multi-layer protection and optimization to achieve hyper-scalability, hyper-flexibility, and hyper-performance in its backbone and metro network. This includes:

1. A clear understanding of the business requirements for the quality of the services offered – for example, multiple classes of service (CoS), SLA and resiliency required for each CoS, as well as a business opportunity for generating additional revenue from creating additional and higher levels of CoS.
2. A well-defined differentiated, multi-layer protection solution using a mix of protection and restoration mechanisms available from IP and optical layers differently for multiple CoS options based on their SLA and resiliency requirements. An example of such a solution is given in Figure 5, where it is compared with an IP only single layer protection solution which is over-engineered

to meet the same SLA and resiliency requirements. For example, for the premium CoS, the multi-layer protection solution uses a mix of MPLS FRR (to meet < 50msec recovery time) and lower-cost optical guaranteed routing (GR) for resiliency against multiple failures, while the IP only protection requires a combination of MPLS path protection and FRR which is more expensive as the path protection would go over multiple intermediate IP nodes.

3. A set of agile IP and optical network resources made possible by emerging technologies such as segment routing (SR), path calculation engine (PCE), virtualization, CDC-F, Photonic switching, OTN switching, Variable Modulation OT's, Flexigrid and open ROADM.
4. Adaptive network resource optimization achieved by a single or hierarchy of SDN controllers capable of controlling and coordinating agile network resources across IP and optical and dynamically setting up the optimal traffic routing, network topology and resource allocation subject to the changing traffic demands so that network traffic can be optimally routed and protected at the least possible cost.

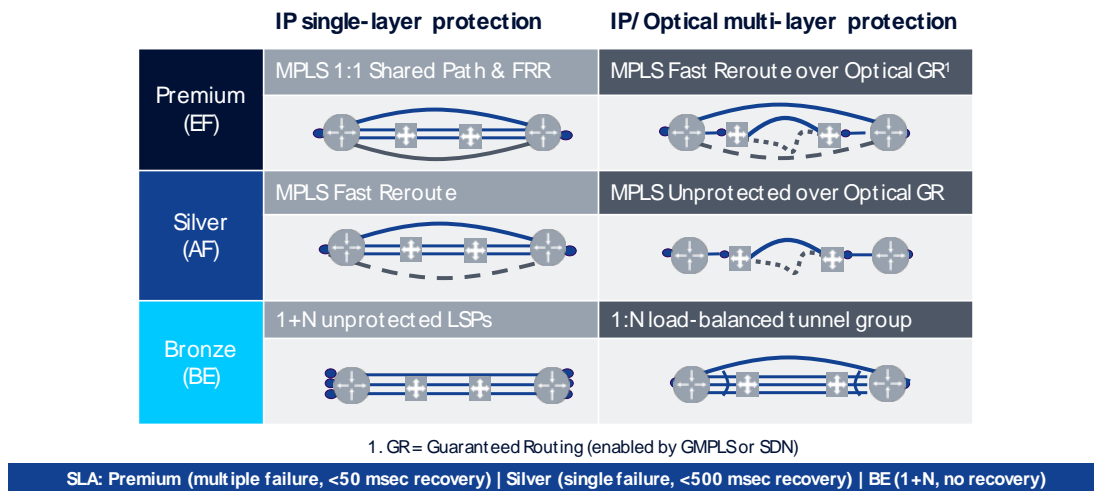


Figure 5 – An Example of Multi-layer Protection Solution

Adaptive network resource optimization can be executed at different timescales of the network optimization cycle as indicated by Figure 6. At real-time or near real-time, the SDN controller along with PCE, SR and intelligent traffic routing algorithms [14] provide a traffic steering framework allowing MSOs to dynamically regulate admission and routing of traffic subject to specific policies and optimization goals (e.g., maximized revenue, guaranteed latency and/or resiliency) – achieving all these based on a given fixed available network capacity. At the mid-term time scale (days/weeks/months), changes in link capacity and network topology can be triggered by the shift in traffic patterns and load to optimize cost and performance. Long-term, network capacity is optimized through the operator's semi-annual or annual (or longer) capacity planning process which will involve only automatic SDN control.

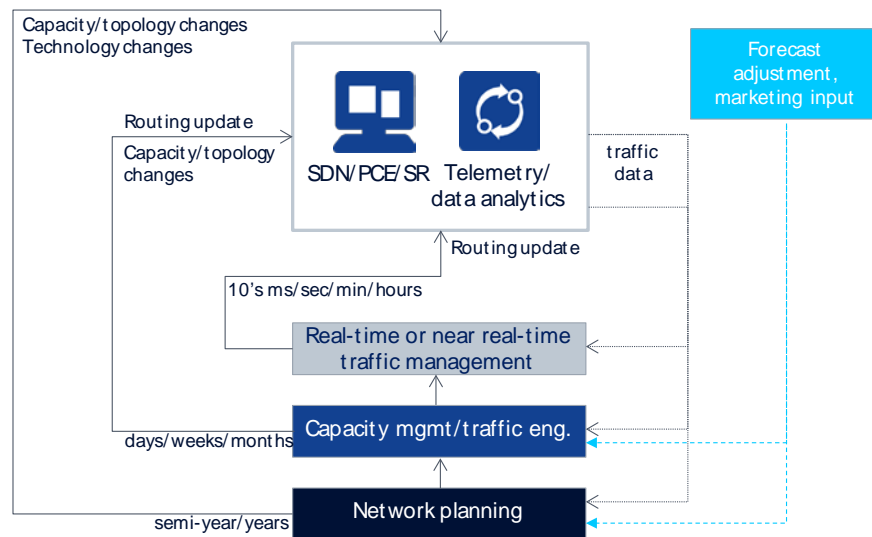


Figure 6 – Adaptive Network Resource Optimization at Various Time Scales

A Bell Labs study [15] examines SDN-enabled multi-layer protection in a typical medium-sized metro network with a mix of residential, mobile and enterprise services carried in three classes of services (Figure 7) and protected by a differentiated multi-layer protection strategy as defined in Figure 5. The study demonstrates a potential TCO savings of 30%-40% over 5 years compared to the scenario based on IP single layer protection (also described in Figure 5) in two cases where the traffic growth is driven by residential/mobile consumer services (traffic profile P1) and by enterprise services (traffic profile P2) separately. The TCO saving is due to not only better utilization of network resources for protection/restoration by using a mix of protection mechanisms from IP and optical, but also the adaptive resource optimization made possible by SDN where the IP network topology (“agile” logical topology) can be dynamically optimized to the changing traffic pattern that leads to more efficient traffic routing and further reduces the cost of protection. We depict in Figure 7 the 5-year TCO saving based on traffic profile P2 (enterprise driven) increases from 30% to 36% in the presence of a dynamically changing traffic pattern between enterprise data centers. Another metro SDN multi-layer protection case study for an MSO network operator [16] also shows a consistent 5-year TCO saving of 30%.




Figure 7 - A TCO Case Study of SDN-enabled Multi-Layer Protection in Metro

End-to-End SDN-based Service Orchestration

Network operators planning to implement carrier SDN expect the most important business benefits to achieve are the ability to generate more revenue from existing services and to create and deploy new services to market more rapidly. This priority is particularly driven from the perspective of a network operator's external facing product group as opposed to the network group who sees SDN primarily to improve network efficiency, cost per bit and interoperability. SDN achieves the dynamic and rapid creation of services by facilitating automated and programmable configuration and provisioning of connections and service components, which also reduces operations cost to maintain service profitability. Network automation goes beyond service orchestrating – it also addresses service assurance using monitoring and data analytics tools to detect and diagnose network problems, sometimes proactively, and respond to network demand ahead of time, delivering auto-healing and auto-scaling functionality.

These SDN benefits are of importance for MSOs as they are expanding footprint into new markets and business models where they need to quickly turn on existing and new competitive services with quality assurance at profitable operations cost.

Within the broadband access and metro/ backbone networks, MSOs are embracing the industry trend to re-architect existing hubs and data centers to centralized and edge clouds to host converged application/service platforms (e.g., SD-WAN, vDNS) and deploy a common NFV infrastructure for network functions such as vSTB, vCPE, vPVR, vCDN, vRouter and vCCAP/vCMTS. While these new capabilities promise to create more agility, cost efficiency, third-party innovation and new revenue opportunity, this cannot be accomplished without a multi-domain SDN control that orchestrates the allocation, activation, and monitoring of this virtualized services and network resources as well physical devices across the data centers and metro and broadband access network.

Multi-domain SDN orchestration in such case can be achieved through the use of underlay and overlay SDN controllers as illustrated in Figure 1. Underlay SDN controls the IP and Optical transport infrastructure layer and performs dynamic (re-)configuration of physical and virtualized forwarding network functions via network protocols such as OpenFlow, Nentconf, BGP-LS, and segment routing to create connections with required bandwidth and performance (latency, resiliency) through a programmable north-bound API to the clients requesting such connections. The utilization of network resources for establishing the connections can be optimized with the use of intelligent traffic steering algorithms and PCE by the underlay SDN. Underlay SDN is also responsible for proactively assuring quality, resiliency, and security of connections through constant data collection and analytics, and advanced support for cognitive, dynamic network control through the northbound API. Such intent-based API simplify the network knowledge required on the client side and allows the SDN control more flexibility to select and place network functions to provide the connections desired. All these underlay SDN capabilities can be implemented in a future network operating system as illustrated in Figure 8.




Figure 8 – Underlay SDN Implemented in a Network Operating System

8. Automated Service Provisioning and Assurance

Overlay SDN at the services layer enables flexible, dynamic and elastic allocation of required application/service functions through service chaining and policy enforcement and ties them together using the connections created by the underlay SDN. Overlay SDN is responsible for overall orchestration and delivery of services and can place requests to the underlay SDN for the elastic creation and changes of connections. This can start with customer service orders, generated by either manual tasks or customer-driven actions such as the ordering of a service through a service portal. The service ordering will then be fulfilled by the overlay SDN by identifying specific applications/services required, such as security gateway for encrypted tunnels, and determining the placement of such applications/services from a common pool of resources which may be located across multiple data centers and the connections required to stitch them together to deliver the customer's service order end-to-end. In large operator networks it is also desirable to provide the capability to segregate the provisioning and configuration of the connections among the network functions inside the data center (or hub/headend) – the intra-DC LANs - from the connections across the metro (and backbone) transport network elements. This segregation can be attained via separate DC SDN and WAN SDN controllers under the controller responsible for the End-to-End SDN-based service orchestrator. As illustrated in Figure 1, we refer to this collection of intelligent controllers and orchestrators the Network OS.

SDN orchestration with underlay and overlay SDN delivers automated service provisioning and assurance in the metro network for MSOs. The implementation of SDN orchestration includes SDN controllers and tools supporting Lifecycle Service Orchestration (LSO) which integrates orchestration, fulfillment, control, performance, assurance, usage, analytics, security, and policy of networking services.

Network operators are building up SDN controllers and tools for LSO based on open standards and open-source development such as OpenDaylight and Open Network Automation Platforms (ONAP). Contributed by both network operators and third-party developers, this offers open service creation and management platforms where concepts like microservices and service chaining are applied to create a highly scalable and extensible framework which can be used repeatedly for rolling out new services rapidly to the market in an open, multi-source environment.

Conclusion

SDN discussions and writings have focused on specific portions of MSO networks. This paper expands the application of the SDN paradigm to reconfigurable and programmable metro and access networks. This concept has been widely contextualized in the general scenario of SDN technology, with reference also to the related market place evolution. The reasons for an end-to-end SDN extension have been presented, considering the metro and access traffic evolution, the future network and business requirements, the trends of data center interconnect and edge cloud, the benefit in terms of end-user quality of experience, and the TCO savings. The building blocks for the future SDN networks have been reviewed.

While these building blocks are all critical to MSOs, an End-to-end SDN architecture allows MSOs to effectively address multi-domain, and multi-vendor interoperability. The SDN-based architecture framework proposed provides network abstraction capabilities into “horizontal” domains such as broadband access, metro and backbone transport, as well as “vertical” operational layers - such as products/service, logical resources, and physical resources. By deploying an end-to-end SDN architecture across metro and broadband access, MSO can automate and dynamically scale service provisioning, minimize end-to-end latency and maximize resiliency performance, optimize utilization of resources across IP/optical and throughout metro and broadband access, and create new innovative, revenue-generating services quickly using solutions such as network slicing and service chaining end-to-end. These SDN benefits are of importance for MSOs as they are expanding footprint into new markets and business models where they need to quickly turn on existing and new competitive services with quality assurance at profitable operations cost. This cannot be accomplished without a multi-domain SDN control that orchestrates the allocation, activation, and monitoring of this virtualized service and network resources as well physical devices across the data centers and metro and the broadband access network. Multi-domain SDN orchestration in such case can be achieved through underlay and overlay SDN to deliver automated service provisioning and assurance in the metro and the broadband access network for MSOs. The implementation of SDN orchestration includes SDN controllers and tools supporting Lifecycle Service Orchestration (LSO) which integrates orchestration, fulfillment, control, performance, assurance, usage, analytics, security, and policy of networking services.

Abbreviations

API	application programmable interface
BGP-LS	border gateway protocol - link state
CCAP	converged cable access platform
CDC-F	colorless directionless contentionless - flexible grid
CDN	virtual content delivery network
CMTS	cable modem termination system
DAA	distributed access architecture
FDX	full duplex
FFR	fast reroute
HFC	hybrid fiber-coax
ISBE	international society of broadband experts
MAC	media access control (layer)
MPLS	multi-protocol label switching
MSO	multi-system operator
NFV	network function virtualization
OpEx	operational expense
OT	optical transponder
OTN	optical transport network
PCE	path calculation engine
PHY	physical (layer)
PVR	personal video recorder
RMD	remote mac device
ROADM	reconfigurable optical add-drop multiplexer
RPD	remote phy device
SCTE	society of cable telecommunications engineers
SDN	software defined networks
STB	set atop box
API	application programmable interface
BGP-LS	border gateway protocol - link state
CapEx	capital expenditures
CCAP	converged cable access platform

Bibliography & References

Future Proofing Access Networks Through Wireless/Wireline Convergence, M.J. Glapa, et.al., SCTE Cable-Tech Expo 2017

IP and Optical Convergence: Use Cases and Technical Requirements, January 31, 2014

The Future X Network: A Bell Labs Perspective, Marcus Weldon, CRC Press, 2016

Delivering on the Promise of Multi-Layer Integration with SDN, Light Reading, July 2016

Virtualization and Network Evolution Open Networking – SDN Architecture for Cable Access Networks Technical Report, VNE-TR-SDN-ARCH-V01-150625, CableLabs, 2015

Optimizing Cable Access Network Evolution for Cost, Performance, and Competition, Jean-Philippe Joseph et.al. SCTE•ISBE Journal of Network Operations Volume1, Number2, 2016

SDNized Cable Access Networks, Karthik Sundaresan, 2015 Spring Technical Forum Proceedings, CableLabs, NCTA, SCTE

SDN Ground Truth: Implementing a Massive Scale Programmable DOCSIS Network, Sameer Patel et al., 2016 Spring Technical Forum Proceedings, CableLabs, NCTA, SCTE

Delivering Seamless Subscriber Aware Services over Heterogeneous Access Networks using SDN/NFV, Nagesh Nandiraju et al., 2015 Spring Technical Forum Proceedings, CableLabs, NCTA, SCTE

Software Defined Networking, DOCSIS Provisioning, and MSO Commercial Services, Kevin A. Noll et al., 2014 Spring Technical Forum Proceedings, CableLabs, NCTA, SCTE

SDN As A Matchmaker For Remote Phy Architecture, Alon Bernstein et al., 2015 Spring Technical Forum Proceedings, CableLabs, NCTA, SCTE

Virtualization and Network Evolution – Virtual Provisioning Interface Technical Report, VNE-TR-VPI-V01-170424, CableLabs

Metro Network Traffic Growth: An Architecture Impact Study, Bell Labs, 2015

Optimizing revenues with the PATH PCE – Quantifying the Benefits of Deploying a Centralized Path Computation Element*, Bell Labs, 2015

IP/Optical Integration with SDN, Ben Tang, SRExperts America, September 14, 2016

Metro Multi-Layer Protection Design, Mohcene Mezhoudi, Ben Tang, October 2016

Virtual Fiber – 100 Gbps over Coax

Coaxial Cable: The once and Future King

A Technical Paper Prepared for SCTE•ISBE by

Steven Krapp
Technical Marketing Director
MaxLinear, Inc.
skrapp@maxlinear.com

Special Thanks to:

Hans Wambach, Director of Access Architecture, Liberty Global International

Jan Ariesen, Chief Technical Officer, Technetix

Introduction

Much debate has been made about the medium used in the last mile of the access network. Fiber is attractive because of the low cost/bps of the equipment used to deploy it. However, fiber-to-the-home (FTTH) deployments have demonstrated that deploying fiber in the last hop is expensive. This is primarily due to factors other than the cost of capital equipment such as the cost labor and right-of-way. Fortunately, technologies such as data over cable system interface specification (DOCSIS®) 3.1, full duplex DOCSIS (FDX), and distributed access architecture (DAA) will allow operators to continue to leverage their installed coaxial cable to provide multi-gigabit symmetrical services.

In order to deploy DAAs, such as remote physical layer (R-PHY) and remote media access control (MAC) and PHY (R-MACPHY), operators are pushing the distribution network deeper into the plant to reach DAA nodes. The physical media to reach these nodes has been assumed to be strands of fiber. However, just like in the access network fiber can also be cost prohibitive in the distribution network.

This paper will discuss a technological solution called “virtual fiber.” Virtual fiber will allow operators to leverage the in-place coaxial cable that typically would be bypassed by a fiber link to connect DAA nodes. The first generation of virtual fiber will soon be available and will be able to achieve symmetrical speeds of 10 Gbps. Future generations of virtual fiber will be able to achieve speeds of 25 Gbps and higher providing operators with a deployment roadmap that can keep pace with demand.

Acknowledgements: Virtual fiber and virtual segmentation are technologies that I stumbled upon and to which I am lucky to have been exposed. Hans Wambach, Director of Access Engineering at Liberty Global and Jan Ariesen, Chief Technology Officer (CTO) at Technetix deserve much credit for incubating the ideas and products upon which this paper is based.

1. Access vs. Distribution and DAA

For the purposes of this paper the access network consists of the infrastructure and technologies used to attach subscribers to an operator’s network. It can be a Wi-Fi network, a cellular network, digital subscriber line, DOCSIS, etc. For data services over hybrid fiber-coaxial (HFC) networks that are built on centralized access architectures (CAA) the access network begins at the cable modem termination system (CMTS). The subscribers are attached to the access network which is located on the cable side interface (CSI) of the CMTS. The distribution network is attached the network side interface (NSI).

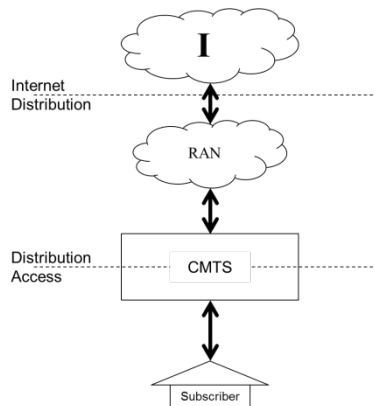


Figure 1 – CAA Networks

DAA comes in two primary flavors, R-PHY and R-MACPHY. In an R-PHY network the CMTS is split between MAC functions and PHY functions. In the simplest form of an R-PHY system, the MAC portion of the CMTS remains at the same location as the CMTS in a CAA network and the PHY portion is located closer to the subscriber. Between the MAC and the PHY is the converged interconnect network or the CIN. For the purposes of this paper the CIN in an R-PHY system will be considered as part of the distribution network.




Figure 2 – DAA Networks

It is expected that the link between the DOCSIS MAC and the R-PHY will consist of fiber with either dedicated or shared wave lengths. Likewise, an R-MACPHY solution is also assumed to have a fiber link between it and the last distribution router or switch in the operator's network.

The idea of virtual fiber is complementary to DAA, R-PHY and R-MACPHY. For an in-depth discussion of CAA and DAA architectures see [EMMENDORFER], for a great intro to R-PHY and its benefits see [CHAPMAN], and for an excellent discussion of why and how operators are deploying R-PHY see [SALINGER].

2. Introduction to Virutal Fiber

It has been noted that deploying fiber to the home can be expensive compared to reusing existing attachments such as coaxial cable, twisted pair, or even power lines [EMMENDORFER 2]. The same cost argument can be made for a portion of the backhaul links in an operator's network. For example, many urban communities require all cabling to be hidden or buried. In these areas labor costs can be prohibitive to deploying new network infrastructure. In these areas there is a large incentive to be able to reuse existing cabling, much of which is coaxial cable.

Virtual fiber is a technology that is being developed that can eliminate the need to dig up streets. Instead the existing coaxial cable is used to transmit 10 Gbps Ethernet signals in a point to point topology.

Consider the network in **Error! Reference source not found.** below:




Figure 3 – Example Coaxial Cable Plant

It consists of a node and several line extenders. 10 homes passed are drawn for simplicity. The node consists of single downstream service group (DS-SG) and a single upstream service group (US-SG). At some point

the operator determines that this node should be split. It would traditionally be done in the following manner shown in **Error! Reference source not found.** below:




Figure 4 – Example Coaxial Cable Plant after node Split

A second node has been added and new fiber has been pulled to that node. After the split half of the homes are connected to the old node and half of the homes are connected to the new node. This type of split is typical. Operators may perform these splits in an as-needed fashion or they may proactively deploy a number of nodes at once in a fiber deep fashion [HOWALD].

A virtual fiber adaptor (VFA) is a device that takes in 10 Gbps Ethernet on one port and emits a radio frequency (RF) signal on another port. Essentially it is a media converter or bridge connecting Ethernet and RF networks. With a pair of VFAs a point-to-point 10 Gbps coaxial connection could be created such as the following in **Error! Reference source not found.**:




Figure 5 – Point-Point Virtual Fiber System

Virtual fiber signals can coexist with traditional HFC signals. In a typical deployment the RF spectrum of system will look like the following:




Figure 6 – Virtual Fiber Spectrum Example

From the above you can see that the virtual fiber band is divided into two bands, Tx A and Tx B. This allows for simultaneous transmission and reception via frequency division duplex (FDD). Alternatively, a time division duplex scheme could be used. In the above example the first VFA, VFA A, transmits on the lower portion of the virtual fiber band, and the second VFA, VFA B transmits on the upper portion of the virtual fiber band. During link establishment the two VFAs can negotiate which portion of the spectrum they will use.

With virtual fiber the split network from **Error! Reference source not found.** can now look like the following:




Figure 7 – Example Coaxial Cable Plant After Virtual Fiber Node Split

The network in **Error! Reference source not found.** is logically identical to the network in **Error! Reference source not found.**. The same DS-SG exists in both. The difference is the RPD in the second node in **Error! Reference source not found.** is fed via fiber, whereas the same node in **Error! Reference source not found.** is fed its 10 Gbps link via the coaxial cable.

3. Overlay Virtual Fiber Example

In most if not all HFC deployments today there is a combination of broadcast or shared signals which are identical in frequency and content for multiple DS-SGs and there is a set of narrowcast (NC) RF signals which are unique in content for each DS-SG. The NC signals typically are at the same frequency in each DS-SG.

Prior to deploying virtual fiber and any node splits the existing cable plant is able to serve the needs of the attached subscribers for both broadcast and narrow cast services. Nodes splits are performed when the quantity of NC signals is no longer enough to meet demand. Increasing NC content is the driver for node splits. It is important to understand that the same broadcast signals that existed prior to the node split are usually adequate to meet demand for broadcast services after the node split.

Virtual fiber deployments can leverage the constant demand for broadcast services in a system by overlaying unique NC signals for each DS-SG, but simply repeating the broadcast signals. The block diagram of a virtual fiber node is shown in **Error! Reference source not found.**. It shows the logical combining of signals for the first node in a virtual fiber system:

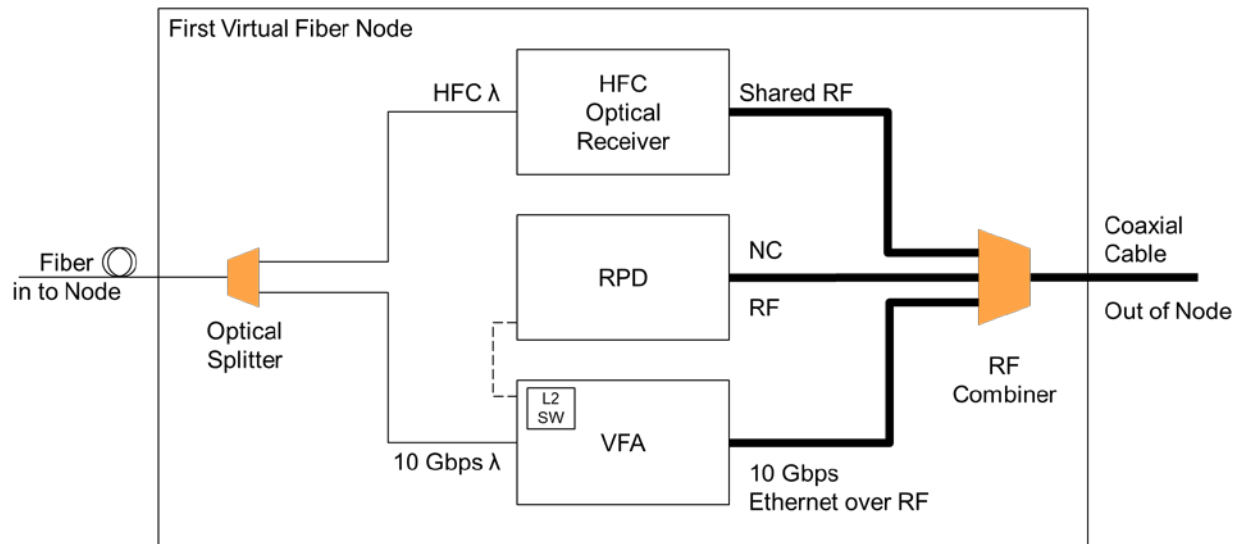


Figure 8 – First Virtual Fiber Node Block Diagram

In the above a layer two switch (L2 SW) has been added to the VFA. This allows packets to go to either the coaxial cable output of the VFA or to a second 10 Gbps port. In this example the second 10 Gbps port is the dashed line connected to the R-PHY device (RPD). This device is collocated in the same node housing. An HFC optical receiver is shown. It connects to a separate λ after the optical splitter. This λ is carrying traditional analog modulated signal used in HFC systems and provides broadcast services in this instance. The optical receiver is optional if all the broadcast signals are digital as the RPD can generate both broadcast and NC digital signals. However, if the broadcast signal contains analog TV signals then the HFC optical receiver would be required.

The second node in **Error! Reference source not found.** needs to receive the 10 Gbps signal sent by the first node and generate its own local RF signal. The following is the block diagram of this second node.

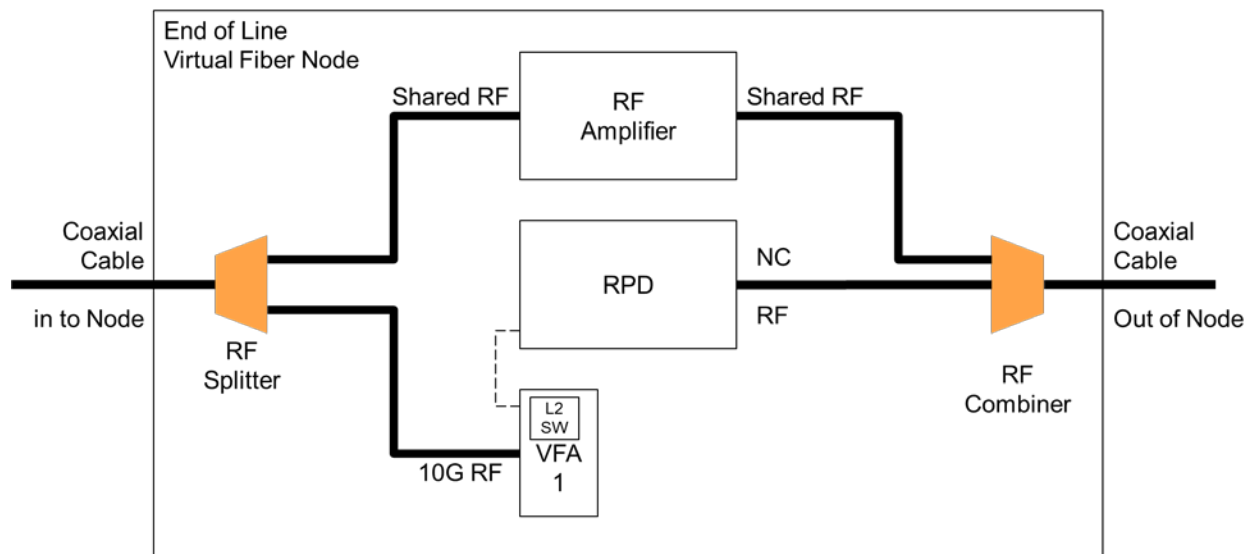


Figure 9 – End of Line Virtual Fiber Node Block Diagram

The block diagram above contains a VFA which receives the 10 Gbps signal and sends these signals to the colocated RPD. This RPD generates any NC for the local DS-SG. If the operator wishes to share any RF signals between the service groups this can be accomplished with the RF amplifier shown. As stated above an example would be broadcast video signals or analog video signals which cannot be generated by the RPD.

Note that **Error! Reference source not found.** only contains two nodes and two DS-SG. This is why the node in **Error! Reference source not found.** is titled “end of line.” Many deployments will require more than two nodes and more DS-SG. To accommodate multiple DS-SG from a single coaxial plant an in-line virtual fiber node is used. Its block diagram looks like the following:

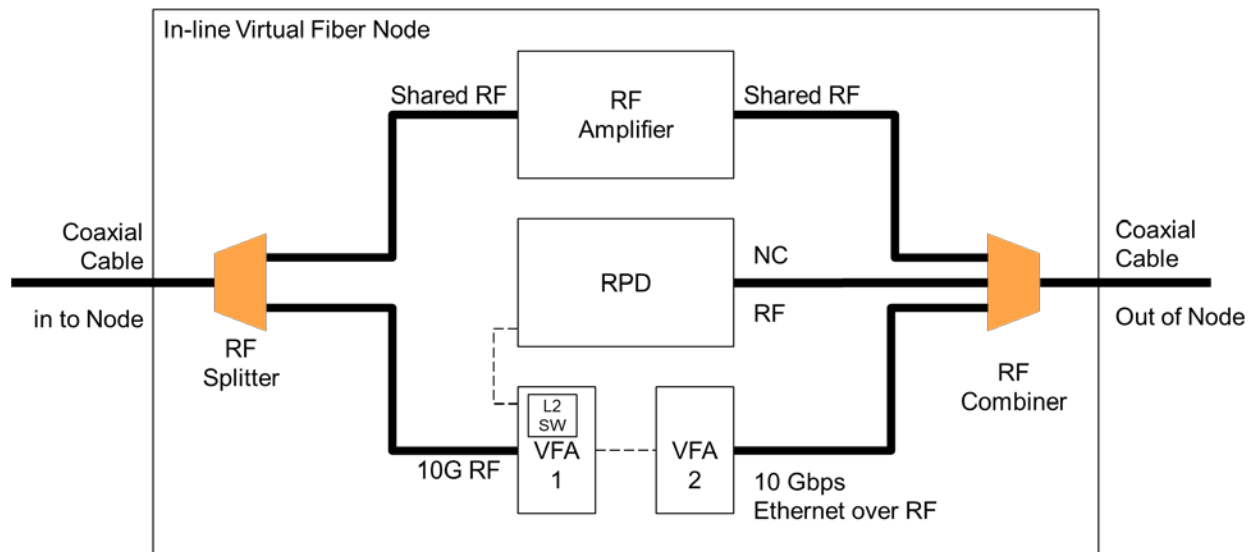


Figure 10 – In-line Virtual Fiber Node Block Diagram

The in-line node contains an RF amplifier for the same reason as the end-of-line node, but instead of a single VFA, a second VFA is used to propagate the 10 Gbps to the next node in the system.

The last component of a virtual fiber deployment that needs to be discussed is the virtual fiber line extender. **In Error! Reference source not found.** RF line extenders are shown in the network. These are two-way amplifiers and most likely they are limited to spectrum that does not include the RF spectrum associated with the virtual fiber band shown in **Error! Reference source not found.**. Thus, in order for the line extenders to propagate the virtual fiber signals each will need to be updated. The following figure shows the block diagram of the line extender after it has been updated to become a virtual fiber line extender:

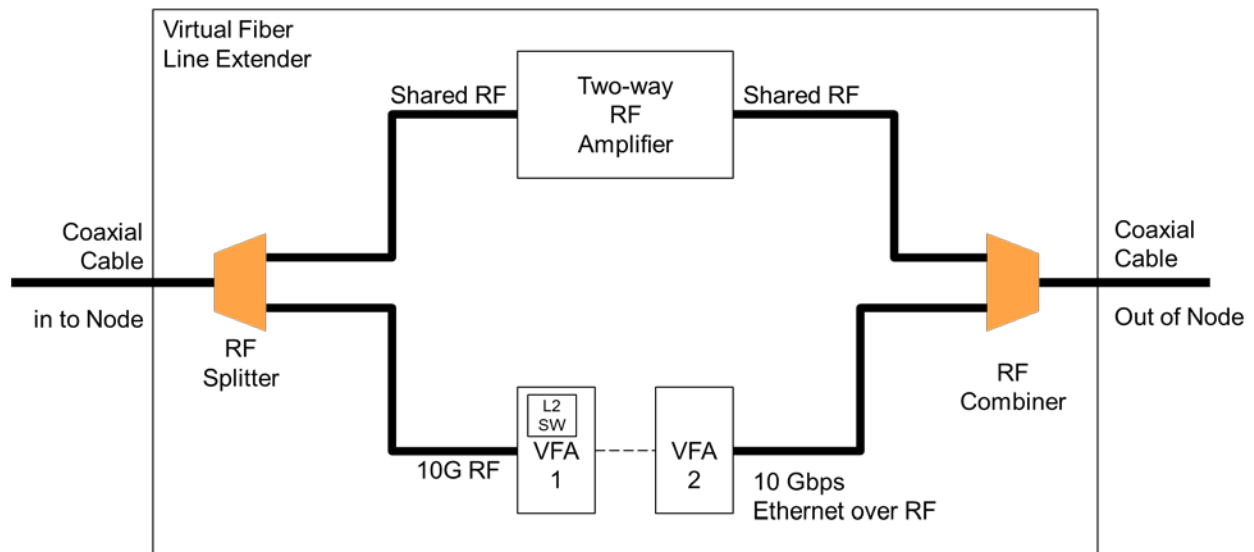


Figure 11 – Virtual Fiber Line Extender Block Diagram

There is no RPD in the virtual fiber line extender. The coaxial cable before the line extender and after the line extender is part of the same DS-SG and US-SG. The RF amplifier in this case is a two-way amplifier as both upstream and downstream signals will need amplification. This is unlike the virtual fiber nodes where it is assumed the RPD terminates any upstream signals. Like the in-line virtual fiber node, there are two VFAs used to propagate the 10 Gbps virtual fiber signal.

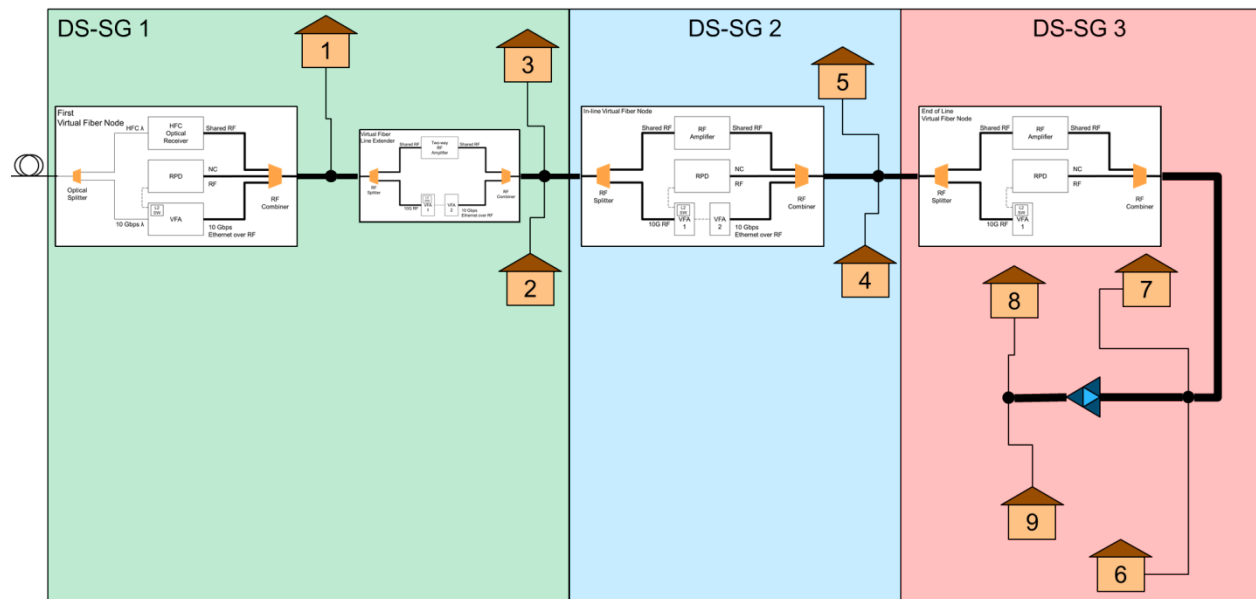


Figure 12 – Three DS-SG Virtual Fiber Network

Error! Reference source not found. represents a three DS-SG virtual fiber network that consists of all four elements discussed above. Home 1 is in DS-SG 1 and is serviced by a first virtual fiber node from **Error! Reference source not found.** Homes 2 and 3 are also in DS-SG 1 but they are farther away and a line extender is required to service them. The line extender shown is a virtual fiber line extender shown in **Error! Reference source not found.** After homes 2 and 3 there is an in-line virtual fiber node from **Error! Reference source not found.** This node is tapping off of the virtual fiber 10 Gbps link on the coax to feed its RPD and generate the local NC signals for DS-SG 2. It serves homes 4 and 5. The last virtual fiber element is an end of line virtual fiber node from **Error! Reference source not found.** It serves DS-SG 3 and homes 6, 7, 8 and 9. It is also tapping off of the virtual fiber 10 Gbps signal. This signal and its bandwidth are shared between DS-SG 2 and DS-SG 3. The last active element in **Error! Reference source not found.** is a standard HFC line extender. It's needed to boost signals to homes 8 and 9; however, it does not need to be a virtual fiber line extender because there are no more consumers of the virtual fiber 10 Gbps signal after it in the network. Thus, a standard line extender will suffice.

4. Virtual Fiber and Timing

RPDs require timing synchronization with the CCAP core. IEEE 1588 and the precision time protocol (PTP) specified by it is used on 10 Gbps fiber links to provide the needed synchronization. PTP and how it relates to DOCSIS has been covered in depth by [JIN]. The reader should take away that a virtual fiber system supporting RPDs will need to support PTP and that this can be accomplished with the VFA acting as a PTP transparent clock. At a very high level, a layer two switch that supports a PTP transparent clock passes a PTP sync message from one port to the other ports on the switch. Since the time spent in the switch can be non-deterministic the switch provides the time each sync message was in the switch. It can do this either by appending the duration to the sync messages, or by sending a follow up message that contains the duration of time that the sync packet was in the switch.

5. How Fast can Virtual Fiber Go

Virtual fiber like all communication systems is speed limited by two main factors: spectral efficiency of the channel (bps/Hz), and the channel width (Hz). With available technology it is possible to readily achieve 10 Gbps bidirectional communications with a signal that is located above traditional HFC signals. Thus, both overlay and non-overlay systems can be accommodated.

Achieving faster speeds up to around 25 Gbps can be accomplished simply via channel stacking. A brute force way to support channel stacking is by ganging multiple VFAs onto a single coaxial cable with each VFA's signal located at a different portion of the spectrum. However, advances in technology should allow a single VFA to support more MHz than the current generation and thus achieve 25 Gbps speeds in a single albeit wider channel.

While wider channels, either achieved via channel stacking or simply growing the MHz per channel, are interesting, there are diminishing returns. When using coaxial cable as a medium the system designer must deal with the fact that signal loss increases with the square root of the frequency of the signal. This can be overcome if the designer is willing to reduce QAM order and distance. [CLOONAN] shows that speeds of 100 Gbps are achievable, albeit for very short distances of about 100 feet.

Conclusion

Virtual Fiber allows for 10 Gbps symmetric links over coaxial cable in its initial form. With a proper VFA implementation virtual fiber will be transparent to the DAA devices which it serves including support for PTP. Thus, virtual fiber can support both R-PHY and R-MACPHY solutions. Faster data rates beyond 10 Gbps are possible with improved spectral efficiency and increases in total Hz used by the virtual fiber solution. This provides operators with a roadmap such that virtual fiber can be utilized to deploy cost effectively future proof R-PHY and R-MACPHY systems in areas where new fiber builds would be cost prohibitive.

While virtual fiber was initially conceived for high fiber cost areas, it can be used in typical HFC plants as an alternative to fiber. Virtual fiber is still in its infancy so it is not yet known if it will be effective or desired in all areas of the network. However, if virtual fiber is successful, there will be one less reason to replace coaxial cable with fiber and coax will continue to reign.

Abbreviations

bps	bits per second
CAA	Centralized Access Architecture
CIN	Converged Interconnect Network
CMTS	cable modem termination system
CSI	cable side interface
CTO	chief technical officer
DAA	distributed access architecture
DOCSIS	data over cable service interface specification
DS-SG	downstream service group
FDD	frequency division duplex
FDX	full duplex DOCSIS
FTTH	fiber-to-the-home
HFC	hybrid fiber coaxial
L2 SW	layer 2 switch
MAC	media access control layer
NC	narrowcast
NSI	network side interface
PHY	physical layer
PTP	precision time protocol
RF	radio frequency
R-MACPHY	remote mac-phy
RPD	remote phy device
R-PHY	remote phy
US-SG	upstream service group
VFA	virtual fiber adaptor

Bibliography & References

[EMMENDORFER] *ESTIMATING DOWNSTREAM PERFORMANCE AND DOCSIS 3.1 CAPACITY IN CAA AND DAA SYSTEMS*; Michael Emmendorfer, Brent Arnold, Zoran Maricevic, Frank O'Keeffe, and Venk Mutalik; SCTE 2015 Spring Technical Forum Proceedings

[CHAPMAN] *DOCSIS Remote PHY: Modular Headend Architecture (MHA_{v2})*; John T. Chapman; 2013 SCTE CableTEC Expo; available via:

https://www.cisco.com/c/dam/en/us/solutions/ns341/ns522/ns791/workshop_remote_phy_chapman_paper.pdf

[SALINGER] *Remote PHY: Why and How*; Jorge D. Salinger; SCTE 2014 Spring Technical Forum Proceedings

[EMMENDORFER 2] *AN ECONOMIC ANALYSIS OF BROWNFIELD MIGRATION CTTH VS. FTTH*; Michael Emmendorfer; SCTE 2016 Spring Technical Forum Proceedings

[CLOONAN] *LESSONS FROM TELCO & WIRELESS PROVIDERS: EXTENDING THE LIFE OF THE HFC PLANT WITH NEW TECHNOLOGIES*; Tom Cloonan, Ayham Al-Banna, Mike Emmendorfer, Zoran Maricevic, Frank O'Keeffe, John Ulm; SCTE 2015 Spring Technical Forum Proceedings

[HOWALD] *THE FIBER FRONTIER*; Dr. Robert L. Howald; SCTE 2016 Spring Technical Forum

[JIN] *TIME SCHEMES IN REMOTE PHY ARCHITECTURE*; Hang Jin and Yubin Chen, SCTE 2015 Spring Technical Forum

A Simple Overview of Blockchains

Why They Are Important to the Cable Industry

A Technical Paper prepared for SCTE•ISBE by

Steve Goeringer
Principal Architect
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
s.goeringer@cablelabs.com

Introduction

Blockchain technology has rapidly become one of the most discussed and visible emerging technologies. Gartner's 2016 Hype Cycle for Emerging Technologies shows blockchain near the peak of inflated expectations while technology mainstream adoption is still likely to be 5-10 years out. Other technologists and analysts have hyped blockchain even further, claiming it to be the most significant technological innovation since the internet. Recently, many researchers have started to consider whether blockchains can be applied to improving IoT Security or services. What is a blockchain? How is it transformational? This paper provides a quick primer into what blockchains are and why they have the potential to be uniquely valuable to cable network operators. The first part of the paper reviews the basics of how blockchains work. This is followed by a discussion of blockchain features and requirements that are relevant to network operators. The paper concludes by asking two key questions that will aid readers to find their own killer applications.

So... What is a Blockchain?

What is a blockchain? It's hard to find a simple definition that doesn't relate to a distributed database or contain a reference to Bitcoin. Perhaps a simplistic but concise definition is that a blockchain is an immutable, distributed ledger visible to the community implementing and using the blockchain. Immutable means that the information a blockchain contains cannot be changed. Distributed means that the information is replicated amongst many participants (in Bitcoin terms, nodes). Ledger implies that the blockchain records transactions. Visible to the community means that every transaction recorded in the ledger is visible to every participant – user or implementer – of the blockchain.

Another definition, closer to what a computer scientist might appreciate, is that blockchains are a method used to create securely linked lists of transactions. Secure in this context means cryptographically protected (authenticated and signed) and distributed amongst stakeholders. The list of transactions is linked by using a hash of a collection of transactions (a block) in the next collection of transactions. Fault tolerance against failure, including malicious or negligent actions of stake holders, is achieved using a consensus protocol (achieving a security goal referred to as Byzantine fault tolerance) [Lamport][Castro].

However, unless you already had a firm grasp of what a blockchain is, really, these definitions probably didn't really provide much insight. Let's come back to building that insight in a moment.

Some of the hype also talks about distributed ledgers. The phrases blockchain and distributed ledger are often used synonymously. In many ways, the term distributed ledger is more descriptive than blockchains. Perhaps it would be useful to think of blockchains as the technology and distributed ledgers the result of using blockchains (e.g., blockchains create distributed ledgers). However, in common usage the terms are used interchangeably.

Let's talk about why you might want to even care about blockchains or distributed ledgers. There is huge hype about why blockchains are important. Marco Iansiti and Karim Lahkhani describe blockchains as a foundational technology. They also think it will take quite a while for it to achieve its transformational process. [HRB] Alex Tapscott describes blockchain as the "next generation of the internet." He goes further, "The blockchain is the internet of value." [Miller]

The "big deal" is that we've never had a capability to create a distributed public history and make it available to the blockchain participants. Through the power of cryptography, we can create a secure

history of transactions that is much more expensive to change than it is to create, and in fact is practically impossible to change. And we can do so in a way that makes all those transactions visible to all the participants in a given network (or not – visibility is a choice). This makes blockchains uniquely valuable.

Why is an unchangeable, public, distributed ledger valuable? Some applications that have been discussed for blockchain include digital currencies, recording of real estate transactions, and registration of marriage licenses. Digital currencies can be used in countries experiencing hyper-inflation. Blockchains can be used to record public transactions and prevent corrupt officials from changing transactions illicitly after the fact. These are just a couple of examples of blockchain applications.

How Do Blockchains Work?

There are many blockchain implementations – Bitcoin, Ethereum, Hyperledger (a Linux Foundation Project), Multichain, BigChainDB. A comparison of several is provided by a companion paper [Hintzman]. They all differ in important details. However, the general notion is illustrated in Figure 1.

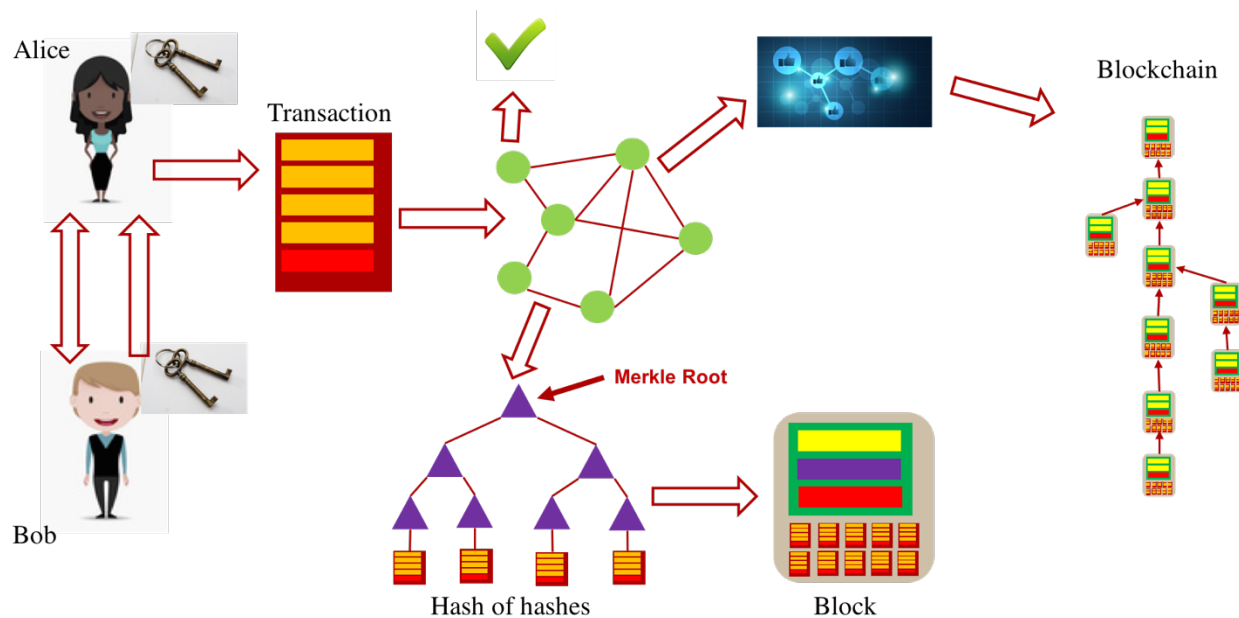


Figure 1 – How blockchains work

Alice and Bob want to record a transaction – perhaps Alice is buying something from Bob using a digital currency, or perhaps there is some event they are trying to permanently record. This explanation is use case neutral. Both Alice and Bob have created or been assigned public-private key pairs that will be used to attest this transaction between them. Generally, the transaction will be from Bob’s public key to Alice’s public key. (Readers are referred to Wikipedia’s entry on public key infrastructure for further details. [PKI]) After they negotiate the terms of their transaction, Bob provided his public key to Alice. Alice creates and sends a transaction to Bob using his public key and signs the transaction using her private key. (In reality, Alice uses an application to do all this. Digital currencies call this application a wallet.)

Alice submits that transaction to one or more nodes in a blockchain. The nodes comprise a network and Alice’s transaction may be submitted to one, many, or all nodes. For purpose of this discussion, the

elements that participate in the blockchain network will be referred to as nodes. Each node that receives Alice's transaction will validate the transaction according to some criteria (for example, authentication using public key infrastructure (PKI) or comparison of information in the transaction to a policy or list). Nodes will add valid transactions to stack or queue of transactions.

At some point, the collections of transactions in the queue get processed at each of the nodes. Usually, this is triggered by a time interval but other criteria are possible. First, the transactions are hashed (this might have been done when the transaction was generated by Alice). Hashing is a mathematical function (often referred to as a trap door) that computes a value that cannot be easily reversed [Diffie]. Trap door functions ensure it is hard to determine the original input given just the output of the hash function. These hashes are then aggregated and hashed again, producing a hash of hashes (such as a Merkle root, which use a tree structure as shown in Figure 1) [Merkle]. The transactions, the hash of hashes, a link to the immediately previously produced block (usually the hash of that previous block), and other information are encoded into a block. Proof-of-work may be performed on this block [Jakobsson]. (Proof of work requires application of computer resources to solve a problem, usually a mathematical computation, as an economic measure to discourage system misuse such as denial of service.) And, of course, this block is hashed.

The next step has the greatest variation amongst the different blockchain implementations. One or more blocks from all the nodes in the network need to be added to the blockchain (distributed ledger). All the blockchain network nodes that successfully create a block in time (systems that use proof of work have uncertainty) have a chance of having their block added. A consensus protocol and process is applied to select the block (or blocks) [Fischer]. This might be through voting, or by proof of work, or proof of stake, or some other scheme. Proof of stake consensus has block selection conducted using an a priori deterministic selection of blocks based on the stake (ownership or possession) a given submitter has in the blockchain. [POS] The goal of the consensus process is to make it hard for one of more nodes to compromise the overall, long term integrity of the blockchain. In some systems, the consensus process may allow more than one blockchain to exist at a time while consensus is still "debated". Eventually, however, the network should converge to a single chain.

And then the process starts over. On current blockchains, iteration is assumed to occur for eternity.

What Do Blockchains Achieve?

What is the result of the blockchain as described above? There are several achievements worth noting. Alice, Bob, and everybody else can see that a transaction occurred between their identities. If those identities are anonymous, then the transaction is visible but nobody knows the actual actors involved (see Figure 2). It's important to keep in mind that anonymity is a choice. Also, information contained in the blockchain can be encrypted or actually just be a hash of information that is stored elsewhere known to at least Alice and Bob.




Figure 2 – Blockchain visibility

We've also achieved layers of integrity as shown in Figure 3. The integrity of that transaction within the blockchain is verifiable by Alice's digital signature and the Merkle root in the block in which it saved. This is probably the real big deal about blockchains – we've created a verifiable history that is computationally challenging to change.

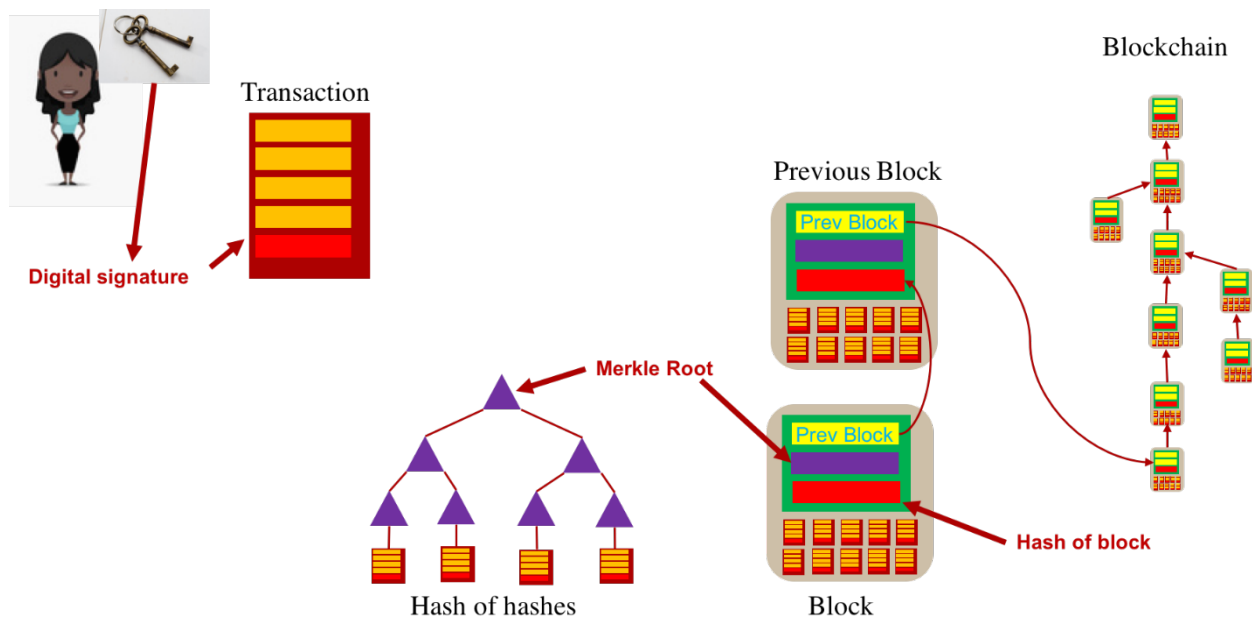


Figure 3 – Blockchains provide multiple layers of integrity

There are more features that further protect the integrity of the transaction history. Several nodes achieved consensus (at least 51%) on the validity of the blockchain after the new block was added before their proposed blockchain was accepted. This protects against malicious nodes and against malfunctioning nodes (in other words, Byzantine fault tolerance). As long as the consensus pool is large enough, the

nodes don't have to trust each other – they just assume more nodes are trustworthy than aren't. The resulting distributed ledger is replicated (usually in whole, but in part is possible) across many nodes assuring availability and also making it hard to change the distributed ledger. Finally, the integrity of the chain itself is verifiable by checking the chain of signatures from the genesis of the ledger all the way to the current block. See Figure 4.

The result is a linked list of transactions that are visible to blockchain participants, verifiable, and unchangeable. The transaction occurred and the ledger entry for that transaction on the blockchain cannot be changed. Consequently, there is no longer any need for participants in a blockchain to need to trust each other with regards to the nature of that transaction – the existence of, and the contents of the transaction can be treated as fact. A major misconception about blockchains is that they provide a basis of trust. A better perspective is that blockchains eliminate the need for trust.




Figure 4 – Blockchain consensus

Smart Contracts

One of the values not discussed above is that the transaction submitted by Alice may include executable code, or a script. This concept is loosely described as a “smart contract” [Stark]. This is an area still in its infancy – it still must be proven that small code snippets imbedded in blockchain transactions can really be secured against misuse. Assuming they are secure, smart contracts could be transformational by creating programmable currencies and transactions that execute automatically according to the conditions included in the contract. Conditions of the contract are programmed using a constrained and highly secure programming language. Only an authorized part (verified usually by possessing a private key) can execute the code.

So, in our blockchain example, the contract gets encoded by Alice. She includes it in her transaction (which is signed by her, so we know the contract is valid). See Figure 5. Then execution of the transaction can be conditional. Some examples:

- The transaction is not valid unless the recipient can respond correctly to a cryptographic challenge which is verified by the smart contract.
- The smart contract contains a counter that is decremented each time the contract is accessed and when the counter is zero, the transaction becomes unusable.




Figure 5 – Smart contracts using blockchain

How Can Cable Use Blockchains?

Now that we know what blockchains are, let's consider how they apply towards the cable industry. Blockchains can be used as platforms for orchestrating ecosystem-secure transactions. They convey transparency and visibility. They are immutable. And they are transaction-oriented. Given those strengths, it would seem that they might be widely applicable.

However, there are some factors that should be evident from the example above to consider first. For blockchains to be useful, they rely on relatively strong cryptography which is computationally intensive. In network engineering terms, this means transaction processing uses lots of power and may be relatively slow. Furthermore, they are distributed which means that information is stored very redundantly – perhaps at thousands of nodes. In fact, in most blockchain implementations, every node contains every transaction. And the size of that storage increases linearly over time per node, and geometrically across the network (because every node contains all the data). And finally, behind the consensus algorithm is the need for a community of actors to work together to implement the blockchain. This ultimately means governance of the code base and network participation terms. In summary, blockchains are resource intensive in terms of compute, storage, and networking and the stakeholders that implement the chain must be willing to work together.

Given that background, the transformational potential of blockchains can be hard to realize. There have been many use cases postulated that ultimately use blockchains as a secure database. Blockchains are very inefficient for data storage and data retrieval and all the public chains actually have databases used to present information from the blockchain (rather than running queries against the blockchain itself). Many of these use cases, however, appear to have worked very well. However, the benefits may have had more

to do with application programming interfaces that were optimized for the specific needs of these particular use cases.

How then, can we identify good use cases for the cable industry where blockchains may work well or use cases that might be transformation to our industry? The goal is to find opportunities that dramatically impact cable, and identify areas where we can reduce friction, speed time to market, and remove the need for trust. Considering the following two questions are helpful:

- *Can we use blockchains as platforms for digital transformation of the cable user experience?*
- *Can blockchains enable dynamically social user experiences for cable subscribers that mirror the sharing economy?*

Use cases that satisfy these questions are still under investigation. Three areas seem potentially very attractive: improving trust in content distribution, streamlining complex service delivery in the medical industry by leveraging cable, and providing improved secure digital content production and distribution. Evaluation of specific use cases should apply formal methods to determine blockchain applicability [Scriber].

A Blockchain for the Cable Industry

It may be beneficial to leverage public or other industry chains for some cable industry use cases. However, the cable industry might be well served by one or more industry-specific blockchains. Such blockchains can be designed to meet specific security, performance, and scalability requirements appropriate to regulated products that serve markets of millions of subscribers. Moreover, governance can be left wholly in the control of the cable industry stakeholders without compromise with the best interests of other sectors or parties. Finally, a cable industry blockchain may prove to be more economic in the longer term. Leveraging existing chains may involve transactions fees, integration and consulting costs, and features that provide little value to cable use cases while dramatically increasing processing and storage requirements.

This should be very practical because of the following: Given how many dozens of blockchains already exist, it is clear there is no specific technical or economic hurdle that prevents creation of an industry blockchain. The critical design decision issues appear to be how governance will be executed and maintained and, under that governance, who will implement the blockchains.

Can a single company benefit internally from a blockchain? Perhaps. If trust management between company elements is challenging, the visibility and immutability of a blockchain may prove useful. However, using a blockchain may appear easier in implementation than traditional database or transaction-logging mechanisms. As mentioned previously, this may largely be due to more modern, or more simply engineered application programming interfaces that are particularly easy to apply to the intended use cases. It must be remembered that maintenance of a blockchain requires significant processing and storage resources. Moreover, the rigid linked-list structure of a blockchain is not well suited for efficient searching or indexing. It seems that the best use cases must have the scope of an entire ecosystem for blockchains to serve their intended purpose in an efficient manner.

Conclusion

This paper has provided a quick, practical overview of the basic concepts of how blockchains work and some suggested use-cases that may be of interest to cable operators. Given that background, the paper has suggested that there may even be transformational use cases that are yet to be discovered, that are very applicable to the cable industry. However, the benefit of blockchains to those use cases must be carefully considered because blockchains can be inefficient and/or expensive. Where blockchains seem to fit well for cable services and applications, we may find that public or otherwise commercially-enabled chains may not be the best implementation path for cable operators. Rather, the cable industry should consider implementation of one or more dedicated blockchains.

Abbreviations

ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[HRB] The Truth About Blockchain. Marco Iansiti and Karim Lakhani. Harvard Business Review. January-February 2017. Online. Downloaded July 11, 2017. <https://hbr.org/2017/01/the-truth-about-blockchain>.

[Miller] Alex Tapscott on why the future of finance will be on blockchain. Podcast. Zack Miller. Tearsheet. July 22, 2016. Online. Downloaded July 11, 2017. <http://www.tearsheet.co/digital-currency/podcast-alex-tapscott-on-why-the-future-of-finance-will-be-on-blockchain>.

[Stark] Making Sense of Blockchain Smart Contracts. Coindesk. June 2016. Online. Downloaded July 12, 2017. <http://www.coindesk.com/making-sense-smart-contracts/>.

[Lamport] The Byzantine Generals Problem. Leslie Lamport, Robert Shostak, Marshal Pease. ACM Transactions on Programming Languages and Systems. July 1982. Online. Downloaded July 12, 2017. <http://dl.acm.org/citation.cfm?doid=357172.357176&CFID=784828639&CFTOKEN=84236530>.

[Castro] Practical Byzantine Fault Tolerance. Miguel Castro and Barbara Liskov. Proceedings of the Third Symposium on Operating Systems Design and Implementation. February 1999. Online. Downloaded July 12, 2017. <http://pmg.csail.mit.edu/papers/osdi99.pdf>.

[Hintzman] Comparing Blockchain Implementations. Zane Hintzman. Cable-Tec Expo 2017. October 2017.

[PKI] Public key infrastructure. Wikipedia. Online. Downloaded July 17, 2017. https://en.wikipedia.org/wiki/Public_key_infrastructure.

[Diffie] New Directions in Cryptography. Whitfield Diffie and Martin Hellman. IEEE Transactions in Information Theory. November 1976. Online. Downloaded July 12, 2017. <http://www-ee.stanford.edu/~hellman/publications/24.pdf>.

[Merkle] Comments in 2012 about the 1979 paper: A Certified Digital Signature. Ralph Merkle. Online. Downloaded July 12, 2017. <http://www.merkle.com/papers/Certified1979.pdf>.

[POS] Proof of Stake. Bitcoin wiki. Online. Downloaded July 17, 2017. https://en.bitcoin.it/wiki/Proof_of_Stake.

[Jakobsson] Proofs of Work and Bread Pudding Protocols. Markus Jakobsson and Ari Juels. Communications and Multimedia Security. Kluwer Academic Publishers. 1999. Extended Abstract Online. Downloaded July 12, 2017. <https://www.emc.com/emc-plus/rsa-labs/ps/breadpudding.ps>.

[Fischer] The Consensus Problem in Unreliable Distributed Systems (A Brief Survey). Michael Fischer. International Conference on Foundations of Computation Theory. August 1983. Online. Downloaded July 12, 2017. <http://zoo.cs.yale.edu/classes/cs426/2012/bib/fischer83consensus.pdf>.

[Scriber] A Framework for Determining Blockchain Applicability: Blockchain Architectural Fit. Brian Scriber. IEEE Software Magazine. 2017. Publication pending at time this article was written.

Comparing Blockchain Implementations

A Technical Paper prepared for SCTE•ISBE by

Zane Hintzman

Associate Engineer

Employee/CableLabs

11605 Destination Drive, Apt 5207

Broomfield, CO 80021

303-517-2664

z.hintzman@cablelabs.com

Introduction

Blockchain technology is now very popular because it provides a new tool to solve problems in a way people could not before. When people think of blockchains, they most likely think of Bitcoin, the most well-known implementation of a blockchain. There are many other blockchain implementations as well. Some of them are still in development, while others are currently running. Different implementations will vary in many ways such as their purpose, ease of participation, how governance is handled, and much more. To determine which blockchain implementation should be leveraged for a given application, it is important to be familiar with the differences between each implementation.

Possibly the most important consideration for a blockchain implementation is its purpose. This seems obvious, but is truly overlooked by many developers. Most existing blockchains are specialized for cryptocurrencies, and many of these blockchains use Bitcoin's codebase. These blockchains often provide ways to enable usage in non-monetary applications by storing application data within transactions. For example, some solutions store information off the blockchain, then create a hash of that information and submit it in a blockchain transaction. The hash is stored on the blockchain and can be referenced by anyone to validate the same information. However, these methods are often limited in use, so cryptocurrency blockchains might not be a good fit for other applications. This makes it very important to understand the purpose of the application utilizing a blockchain to ensure that it aligns with the intended usage of the implementation.

The next aspect to consider is the ease of participating in the blockchain. For example, connecting to a blockchain may require running a full node that stores the entire history of transactions and blocks. Some blockchains have lightweight clients that allow access to the blockchain's network without downloading the whole transaction history. In addition, it is necessary to check whether the blockchain ecosystem is open to anybody, or only a closed group. If the codebase is open-source, then it will typically be an open network. Developers can create a *software fork* from open-source code to make a similar but different blockchain implementation. But doing so will bring maintenance under one's own responsibility.

The governance of a blockchain can also impact how it operates and what it can do. Thus, one should consider who controls access to the ecosystem, and who enforces decisions to make changes to the blockchain. Depending on the implementation, it may or may not be possible to see who is making or enforcing these decisions. Most blockchains require consensus from all participants in order to agree to a change on the blockchain. Someone who has less control in the blockchain's network may find it difficult to utilize the blockchain for their specific application.

Another important aspect is how well the blockchain performs. This means measuring how fast transactions are accepted by the blockchain's network, how much bandwidth it uses, how much blockchain data needs to be stored, and in what way the data must be stored. More specific metrics that one should measure are how quickly blocks are added to the blockchain, block and transaction sizes, and transaction rates. It is also important to note how transactions are bundled within blocks, and what limitations that may cause.

Some people are interested in blockchain technology because of the enhanced security features they offer. Different blockchains often use varying algorithms for key parts of their systems. For example, there are many algorithms that can be used to create the proof-of-work or proof-of-stake necessary to add blocks to the blockchain. Blockchains may also use different scripting languages to run smart contracts. Some blockchains might implement a managed public key infrastructure (PKI) for strong validation of users.

Furthermore, they may have widely differing consensus models for obtaining agreement on the state of the blockchain.

This paper aims to compare many blockchain implementations based on these criteria. Bitcoin will be investigated first because many people are familiar with it. Then other implementations will be described and compared to each other, with Bitcoin as a starting point for discussion.

Comparison of Blockchain Implementations

1. Bitcoin

Bitcoin is a digital asset and payment system where users can perform transactions without any intermediary system. The original source code of Bitcoin includes an implementation of a blockchain. The Bitcoin blockchain stores data on transactions, which indicate the amount of currency that moved between two or more accounts. However, there are ways to encode additional data into transactions. Coinbase transactions, which are transactions that generate new currency, can encode 100 bytes of arbitrary data. But this is limited in use because it can only be used when a node successfully adds a new block. A more accessible option is to create a transaction that uses an OP_RETURN output. This type of transaction can encode 40 bytes of arbitrary data. Still, the primary focus of Bitcoin is digital currency transactions.

The main currency in Bitcoin is simply called *bitcoin* (BTC or B or ₿). The smallest unit of bitcoin is called a *satoshi*, named after Satoshi Nakamoto, the pseudonym for the person or persons who designed Bitcoin. One satoshi equals 1/100,000,000th of a bitcoin. Bitcoins are associated with Bitcoin addresses, and the Bitcoin network runs various scripts to validate the individual funds for each address. The scripts include both locking and unlocking scripts. Locking scripts specify the conditions for spending an output, whereas unlocking scripts contain information to satisfy the conditions of some locking script. Both kinds are written in a scripting language known simply as *script*. It is a stack-based language that is purposefully limited in scope to prevent bugs such as infinite loops.

The Bitcoin protocol is fully based on open-source software which can be downloaded from GitHub (<https://github.com/bitcoin/bitcoin>). The most common way that users participate on the blockchain is by exchanging bitcoin with other users. To get bitcoin, a user generates a private key from which they generate a public key. Then from the public key, they create a Bitcoin address to which others can send bitcoin. There are many third-party sources where bitcoin can be obtained at the current bitcoin exchange rate. A high-level description of this process is presented in a companion paper by Steve Goeringer, a Principal Architect at CableLabs (1).

Another way to interact with Bitcoin is to run a node, which can be run by anyone who has sufficiently powerful hardware. Users can run a full node, which stores the entire Bitcoin blockchain with the history of all bitcoin transactions. But running a full node requires significant hardware. It is recommended to have at least 145 GB of disk space and 2 GB of memory. Over time, more storage and memory will be required. A full node may also use a lot of CPU when it initially synchronizes with the Bitcoin network. An alternative is to run a lightweight client, also known as a simplified payment verification (SPV) node. This stores a user's wallet, but relies on third-party servers to view the blockchain and perform transactions. There are also web clients that allow you to see Bitcoin data on a web browser, but rely on third-party servers to interact with Bitcoin. Furthermore, the Bitcoin network has many application programming interface (API) calls that read from or write to the blockchain.

Bitcoin has two distinct blockchains: the main blockchain and the testnet blockchain. Both use the same underlying protocols, but only the main blockchain is considered a valid means of payment for real-world services. The purpose of the testnet is to allow people to test Bitcoin applications without potentially harming the main blockchain. The testnet is a good place to start testing blockchain applications to see if they can run efficiently with the Bitcoin protocol. Furthermore, Bitcoin also has a regression test mode, allowing users to create a private blockchain where they control when new blocks are created. This provides another environment for users to test Bitcoin blockchain applications. Refer to Bitcoin's website (<https://bitcoin.org>) for more information on how to connect to Bitcoin blockchains.

The Bitcoin blockchain is not owned by any single entity. All nodes in the blockchain's network have an impact on how the blockchain operates. Every transaction is transmitted to every full node, and all full nodes coordinate to obtain the same blockchain. The protocol can be updated across the network if the majority of full nodes agree to implement a particular change. All users, whether they run a full node or a lightweight node, are anonymous to the network. A user's Bitcoin address is used to distinguish different accounts in the network, and a user's wallet can contain more than one address. But there is no requirement to associate personal information with one's Bitcoin addresses. Therefore, it is not directly possible to find out who (meaning real people, not Bitcoin addresses) is participating in the Bitcoin network. Every node contributes to the network and aims to achieve consensus with other nodes.

Consensus is achieved through four general processes. First, each full node independently verifies each transaction based on several criteria. Second, full nodes work independently to aggregate transactions into blocks so they can earn a bitcoin reward. Third, all full nodes verify new blocks as they come in. Finally, full nodes independently select a blockchain with the most cumulative computation demonstrated through proof-of-work. The last step means nodes will choose the blockchain with the most blocks that have a valid proof-of-work.

New bitcoins are generated through a process called *mining*, and is performed by nodes called *miners*. First, a miner consolidates one or more transactions into a block. Next, they calculate a proof-of-work for that block so it can be approved by the network. This is done by creating a hash of the block. As long as the hash is not numerically less than the current difficulty target, the miner changes the *nonce* field in the block so the block's hash also changes. Once the hash is less than the target, the miner transmits the block across the Bitcoin network. Lastly, if the block is accepted by the entire network, the miner receives a reward in the form of newly generated bitcoin. When this happens, miners can no longer submit blocks at the same height as the accepted block. Thus, miners will want to complete these steps faster than other miners to get bitcoin rewards and avoid wasted effort. The most compute-intensive step is calculating the proof-of-work for the block. Miners that can compute a proof-of-work faster than other miners will have a better chance of getting bitcoin rewards.

The protocol is specifically designed so that it always takes around 10 minutes to generate a proof-of-work. This determines the minimum time to create a new block, sometimes called the block-release timing. After a specific number of blocks, the network will change the difficulty of creating a proof-of-work to ensure the process still takes about 10 minutes. The block-release timing is set to 10 minutes in order to decrease the probability of a *fork*, where there are multiple blockchains competing to be considered the main blockchain. The protocol is also limited in that the maximum size a block can have is 1 MB, and the minimum size of a transaction is around 200 bytes. Based on these constraints, the maximum theoretical transaction rate is 7 transactions per second. In contrast, Amazon can handle 1 million transactions per second on most days, and can process 600 transactions per second on days with

lots of internet traffic. Thus, Bitcoin may limit the number of transactions per second for large scale companies like Google and Amazon.

Bitcoin uses a few different algorithms for various aspects of the protocol. To generate a public key from a private key, it uses the Elliptic Curve Digital Signature Algorithm (ECDSA). According to the book *Mastering Bitcoin: Unlocking Digital Cryptocurrencies*, “Bitcoin uses a specific elliptic curve and set of mathematical constraints, as defined in a standard called secp256k1, established by the National Institute of Standards and Technology (NIST)” (2). To create a Bitcoin address, the algorithms SHA256 and RIPEMD160 are used sequentially to hash a public key into the Bitcoin address. SHA256 is also used to calculate a block’s hash, and thus the proof-of-work for the block.

As the most well-known implementation of a blockchain, Bitcoin is a good starting point for comparing other implementations. But just because Bitcoin is the most famous implementation does not mean it is the best. There are many other cryptocurrency blockchains that make some improvements to the original protocol. In addition, blockchains that are not focused on cryptocurrency are more applicable to other use cases.

1.1. Summary of Bitcoin

Table 1 – Summary of Bitcoin blockchain

	Bitcoin
Purpose	Cryptocurrency
What kind of data can be stored?	Cryptocurrency transactions, plus some additional data in coinbase or OP_RETURN transactions
Scripting Languages	Script
Is the ecosystem open?	Yes
How can one participate?	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.
Native Currency	bitcoin (BTC or B or ₿)
Who are the registration authorities?	N/A
Is decision making transparent?	Yes
Does it use a managed PKI?	No
Who manages PKI?	N/A
Block-release Timing	10 minutes
Transaction Size	200 bytes minimum, 250 bytes avg.
Transaction Rate	3 tx/sec. avg., 7 tx/sec. theoretical maximum
Consensus Model	Nodes verify blocks and transactions, and select blockchain with the most blocks.
Mining	Proof-of-work

2. Alt Coins

The invention of Bitcoin inspired the creation of alternative decentralized currencies known as *alt coins*. One example is Litecoin, an alt coin derived from Bitcoin's source code. This makes Litecoin a software fork of Bitcoin, as is the case with most but not all alt coins. Litecoin runs on a separate blockchain that contains the entire history of all Litecoin transactions. As a software fork, Litecoin has many similarities with Bitcoin. For example, as with Bitcoin, Litecoin is an open-source protocol that specializes in cryptocurrency. Litecoin uses the same scripting language as Bitcoin, so transactions can store additional data using an OP_RETURN output script. In addition, the Litecoin blockchain can be accessed in the same way as Bitcoin: by trading for Litecoin's currency, *litecoin* (LTC), or by running a node.

However, Litecoin also has many differences from the original Bitcoin protocol. First, the time to calculate a proof-of-work for a block on the Litecoin blockchain is around 2.5 minutes, compared to 10 minutes in the Bitcoin network. Consequently, Litecoin can theoretically confirm transactions four times as fast as Bitcoin, making the theoretical maximum transaction rate 28 transactions per second. Next, Litecoin uses a different proof-of-work algorithm. Bitcoin uses SHA-256, whereas Litecoin uses *scrypt* (not to be confused with *script*, the scripting language of Bitcoin). It is said that scrypt is quicker and much simpler than SHA-256, but that it could be less secure. Furthermore, there are not as many lightweight client nodes for Litecoin as there are for Bitcoin. Litecoin has the Electrum Lightweight Litecoin Client (<https://electrum-ltc.org>), but it is still in the beta phase and few servers are running this client. This makes it slightly more difficult to join the Litecoin network. However, there are plenty of online services for storing litecoin, so it is not necessary to run a full node to obtain and use litecoin.

Yet another software fork of the Bitcoin protocol is Dogecoin. Dogecoin is also an open-source protocol specialized for cryptocurrency and runs on its own blockchain. As with Bitcoin, it is possible to join the Dogecoin network with a full node, a wallet client, or a lightweight client. The most significant difference between the Dogecoin and Bitcoin protocols is that it is much easier to generate new currency in Dogecoin. Whereas Bitcoin will stop creating new bitcoins once 21 million are created, the supply of *Dogecoin* (DOGE), the protocol's currency, is uncapped. Another difference to Bitcoin is that it takes much less time to create a proof-of-work for a block on the Dogecoin blockchain, taking only 60 seconds. This means that the Dogecoin network can theoretically confirm up to 70 transactions per second. Furthermore, Dogecoin uses scrypt for its proof-of-work algorithm, similar to Litecoin.

While the Litecoin and Dogecoin networks can confirm transactions faster than Bitcoin, they are still limited because they are specialized for digital currencies. Their transactions use the same fields as Bitcoin to store metadata. Saving only a small amount of metadata in transactions might not provide enough storage to leverage for non-cryptocurrency applications.

2.1. Summary of Alt Coins

Table 2 – Summary of alt coin blockchains

	Litecoin	Dogecoin
Purpose	Cryptocurrency	Cryptocurrency
What kind of data can be stored?	Cryptocurrency transactions, plus some additional data in coinbase or OP_RETURN transactions	Cryptocurrency transactions, plus some additional data in coinbase or OP_RETURN transactions
Scripting Languages	Script	Script
Is the ecosystem open?	Yes	Yes
How can one participate?	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.
Native Currency	litecoin (LTC)	Dogecoin (DOGE)
Who are the registration authorities?	N/A	N/A
Is decision making transparent?	Yes	Yes
Does it use a managed PKI?	No	No
Who manages PKI?	N/A	N/A
Block-release Timing	2.5 minutes	60 sec.
Transaction Size	(unknown)	(unknown)
Transaction Rate	28 transactions/sec. theoretical maximum	70 transactions/sec. theoretical maximum
Consensus Model	Nodes verify blocks and transactions, and select blockchain with the most blocks.	Nodes verify blocks and transactions, and select blockchain with the most blocks.
Mining	Proof-of-work	Proof-of-work

3. Ethereum

Not all blockchain implementations are specialized for cryptocurrency. The Ethereum protocol, for example, is a decentralized blockchain platform for running smart contracts. Smart contracts are small executable programs that are run only if certain conditions are met. Not only can Ethereum create new cryptocurrencies, but also it can create digital tokens representing something like a physical asset, a virtual share, or a proof of membership. It is also possible to run a virtual organization through Ethereum by controlling processes in the organization with smart contracts. While the Bitcoin blockchain contains a list of transactions, Ethereum tracks the state of different accounts. There are two types of accounts in the Ethereum network. The first is an *externally owned account* (EOA), which is controlled by a private key. The other kind is a *contract account*, which are snippets of code usually written in some high-level coding language. A contract account is controlled by its contract code and can only be activated by an EOA.

Ethereum runs on its own blockchain and has its own cryptocurrency, called “ether,” to use as payment for submitting transactions. Ether is used to pay for “gas,” which is the internal pricing for running a

transaction or contract in Ethereum. The amount of gas needed to commit a contract increases with the number of computational steps needed to run the contract. An important aspect of Ethereum is the Ethereum Virtual Machine (EVM), the runtime environment for smart contracts. All nodes in the network run the EVM and communicate with each other to execute the same instructions and maintain consensus. The primary language for writing contracts is Solidity, an object-oriented language designed to compile code for the EVM. Contracts can also be written in Serpent or Low-level Lisp-like Language (LLL), but these are less supported than Solidity.

As with Bitcoin, the Ethereum protocol is open-sourced on GitHub (<https://github.com/ethereum>), and the Ethereum network is open to anyone who can run a full node or use a wallet application. The Light Ethereum Subprotocol (LES) is currently in development to allow lightweight clients to run in Ethereum. There is also a public testnet blockchain for testing Ethereum applications, and the protocol allows you to create a private testnet blockchain for internal testing.

Blocks are added to Ethereum's blockchain similarly to Bitcoin. Ethereum uses its own algorithm called *Ethash* to create a proof-of-work for blocks. Transaction latency for Ethereum is only 12 seconds per block, which is significantly faster than the Bitcoin network. Furthermore, the maximum size of a block or transaction is limited only by the amount of gas in circulation. Since the block gas limit is about 3 million gas, the maximum block or transaction size is around 89 kB. The amount of gas necessary to perform a single transaction can vary greatly, but it cannot exceed 3 million gas.

While Ethereum does have its own cryptocurrency, the main purpose of the network is to run smart contracts. This makes it applicable to a larger set of use cases than Bitcoin. However, recent events have created some doubts about the stability of Ethereum. Refer to section 3.2 for details on these events. More information on the Ethereum protocol is available on the main webpage (<https://www.ethereum.org>) and the documentation webpage (<http://ethdocs.org>).

3.1. Summary of Ethereum

Table 3 – Summary of Ethereum blockchain

	Ethereum
Purpose	Run smart contracts
What kind of data can be stored?	Cryptocurrency, digital assets, smart contracts
Scripting Languages	Solidity, Serpent, LLL
Is the ecosystem open?	Yes
How can one participate?	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.
Native Currency	ether (ETH or ETC)
Who are the registration authorities?	N/A
Is decision making transparent?	Yes
Does it use a managed PKI?	No
Who manages PKI?	N/A
Block-release Timing	12 sec.
Transaction Size	Theoretically no max (actual max: 89 kB)
Transaction Rate	Theoretically no maximum
Consensus Model	Similar to Bitcoin, but uses Ethereum Virtual Machine
Mining	Proof-of-work using Ethash algorithm

3.2. Ethereum Network Hack

As usage of a blockchain network goes up, the chance of it becoming the target of hackers increases. An example is the recent hack on the Ethereum blockchain. On June 17th, 2016, a hacker stole over \$50 million dollars' worth of ether (3.6 million ether) from the Decentralized Autonomous Organization (D.A.O.), an experimental virtual currency project that had raised \$160 million in the form of ether.

It appears that the hack was not caused by any vulnerability in the Ethereum codebase. Rather, it was caused by a software vulnerability in the D.A.O.'s code. On June 9th, the vulnerability that was exploited on the 17th was noted by Petter Vessenes in a blog post. Furthermore, on May 27th, a group of computer scientists released a paper that explained many vulnerabilities in the D.A.O. Emin Gün Sirer, a co-author of this paper, noted that it's easy to make mistakes when using Solidity to code smart contracts. In short, the exploited vulnerabilities were public information before the hack even occurred.

A few hours after the hack, the price of ether dropped from \$21.50 per unit to \$15 per unit. A brief time later, Vitalik Buterin, the founder of Ethereum, proposed a *soft fork* in the Ethereum network to prevent the attacker from using the stolen funds for another 27 days. This was followed by a *hard fork* on July 20th to allow users to recover their ether. Buterin assumed that the majority of nodes would conform to the rules of the new hard fork, but many nodes remained separate from the hard fork. As a result, there are now two competing Ethereum blockchains: the first is simply called Ethereum (ETH), and the second is known as Ethereum Classic (ETC). The funds stolen were returned to the D.A.O. in the ETH network, but were not returned in the ETC network. Some people believe that having two networks will allow for "replay attacks." This means that when a user tries to send a transaction to one of these networks, they might accidentally send the same transaction to the other network as well.

This event demonstrates the importance of security in blockchains. When evaluating a blockchain, one should make sure the network is secure, and that there are no obvious flaws. However, the security of user clients, wallets, exchanges, and scripting languages must also be carefully engineered and monitored. In the case of this hack, the fault was not even in the codebase of Ethereum itself. Yet the network was still affected because many people implemented the hard fork to protect the investors of the D.A.O. As a network becomes popular and the value of transactions increase, it will become more and more attractive to cyber thieves and activists.

It is also important to look at the current state of the blockchain, regardless of whether it has been hacked. The hack on the Ethereum network resulted in two competing versions of the Ethereum blockchain, making it much harder and more confusing to participate. It is currently unclear which version will be more widely adopted. Thus, it may be beneficial to avoid blockchains like Ethereum that have been damaged by attacks.

4. MultiChain

An application running on the Bitcoin or Ethereum blockchains must conform to their respective protocols in order to participate. However, it is possible to obtain more implementation freedom with MultiChain. MultiChain is a platform for creating private blockchains. Users define the parameters of the blockchain they create, and they can record multiple kinds of assets on their blockchain. The platform is available to download and install on Linux or Windows machines, and the source code is available on GitHub (<https://github.com/MultiChain/multichain>). MultiChain extends the Bitcoin APIs, and has a

similar protocol and transaction format. A node set up through MultiChain can also act as a node on the Bitcoin network or the Bitcoin testnet network.

For blockchains created using MultiChain, the protocol allows creators to determine what permissions a new participant will have without receiving them from an administrator. Before the blockchain is initialized, the creator determines the initial set of administrators, as well as whether anyone can connect to the network without restriction. Administrators can also dynamically control permissions to the blockchain for specific users while the blockchain is active. Such permissions include the ability to send, receive, or create assets, and the ability to create blocks. Decisions to alter permissions are made via consensus among administrators. The proportion of permitted administrators who must agree to modify a user's privilege is set before the blockchain starts running. This value can be fine-tuned for every kind of permission on the network.

Creators also control how fast the network moves and the size of the data in the blockchain. Again, these parameters are specified before creating the blockchain. The creator can specify the target for the average time to add a block, the maximum block and transaction sizes, and the maximum size of an OP_RETURN metadata output. They can also change the mining difficulty, how frequently the difficulty is updated, and even whether a proof-of-work is required to add new blocks. Regardless of these settings, the blockchain will use a randomized round-robin system to add blocks. The protocol relies on another parameter called *mining diversity*. This determines the minimum proportion of miners required to participate in round-robin mining in order to render a valid blockchain. When a miner adds a new block to the blockchain, it must wait for this minimum proportion of other miners to add blocks before it can add anymore blocks. This may significantly increase the Byzantine fault tolerance (BFT) of MultiChain blockchains in comparison to some other blockchains.

MultiChain offers a lot of flexibility in designing new blockchains, but it has some downsides. A significant limitation of MultiChain is that it does not implement smart contracts, nor does the company have immediate plans to implement smart contracts. Furthermore, although the source code for MultiChain was recently made public, the project is in beta at the time of this writing. The main Multichain webpage (<http://www.multichain.com>) contains materials with more details on the protocol.

4.1. Summary of MultiChain

Table 4 – Summary of MultiChain blockchains

	MultiChain
Purpose	Provide a platform for creating your own blockchain.
What kind of data can be stored?	Any digital asset you want to store.
Scripting Languages	N/A
Is the ecosystem open?	Configurable
How can one participate?	Install MultiChain app, and follow online instructions to make a blockchain.
Native Currency	N/A
Who are the registration authorities?	Configurable
Is decision making transparent?	Configurable
Does it use a managed PKI?	No
Who manages PKI?	N/A
Block-release Timing	Configurable

	MultiChain
Transaction Size	Maximum size configurable
Transaction Rate	Configurable
Consensus Model	Fixed ratio of admins approves privilege changes. Longest valid blockchain adopted as global consensus.
Mining	Round-robin system; proof-of-work requirement is configurable

5. Hyperledger

Many companies have come together to form Hyperledger, a Linux Foundation project whose goal is to advance blockchain technology to benefit a variety of business use cases. There are many projects associated with Hyperledger. Subsequent sections of this document will describe three of the Hyperledger frameworks: Hyperledger Fabric, Hyperledger Sawtooth, and Hyperledger Iroha. These projects provide users with the tools to deploy their own blockchains. A summary of all projects can be found at <https://www.hyperledger.org/projects>.

5.1. Hyperledger Fabric

Hyperledger Fabric is a blockchain implementation designed by IBM for industry use cases. A blockchain deployed using Hyperledger Fabric stores data in the form of *chaincode*, a programmatic code on the network that functions similar to smart contracts on other blockchains. The network currently supports Golang as the language for chaincode, and progress is being made to enable Java as well. By default, the network does not include its own native cryptocurrency. However, users can implement a cryptocurrency through chaincode.

The source code for Hyperledger Fabric is available on GitHub (<https://github.com/hyperledger/fabric>), and includes instructions on how to create a new blockchain using this implementation. A blockchain implemented using Hyperledger Fabric is a permissioned network. Thus, new participants must register with a proof of identity to the network membership services. Transactions sent by each user include derived certificates that cannot be linked to the sender. Each transaction has its content encrypted so it cannot be viewed by unintended participants. The initial registration authorities are set in a configuration file before the blockchain begins running. This file also determines which initial users can assign additional registration authorities while the network is running. The registration authorities are a part of the blockchain's membership services.

Membership services also include roles that implement a PKI. These roles include an enrollment certificate authority (ECA), a transaction certificate authority (TCA), and a transport layer security certificate authority (TLS-CA). An ECA issues enrollment certificates (ECerts) to network participants, a TCA issues transaction certificates (TCerts) to users with an ECert, and a TLS-CA issues TLS certificates to secure communication channels. Certificate authorities must be initialized before the network starts running so that new nodes will connect to the CA.

The performance of Hyperledger Fabric blockchains has not been extensively tested. However, it is expected that they can process transactions at a rate greater than 10k transactions per second using a BFT consensus model. Hyperledger Fabric has a pluggable consensus framework, which provides multiple options for the consensus algorithm used by the blockchain. Currently, the source code provides implementations for two different algorithms. The first implementation is for the Practical Byzantine

Fault Tolerance (PBFT) consensus protocol. The second is a “dummy” consensus protocol that doesn't perform consensus but still processes all consensus messages. The latter is simply for development and test purposes, and to provide an example of how to create a consensus model plugin.

Hyperledger Fabric has many benefits in theory, but it is still in development. Still, this establishes that people are working on other ways to implement blockchains. The success of this project would provide users another way to implement a blockchain and give them control over parameters in the blockchain's network. The documentation page (<http://hyperledger-fabric.readthedocs.io>) contains more details on the project.

5.2. Summary of Hyperledger Fabric

Table 5 – Summary of Hyperledger Fabric blockchains

	Hyperledger Fabric
Purpose	Enable the creation of blockchains for industry use cases.
What kind of data can be stored?	Chaincode (i.e. smart contracts)
Scripting Languages	Go (golang), Java (in progress)
Is the ecosystem open?	No
How can one participate?	Create a blockchain: Download source and follow instructions. Join existing network: Register with a proof of identity to the network membership services.
Native Currency	N/A
Who are the registration authorities?	Defined before blockchain is initialized, and more can be assigned while running.
Is decision making transparent?	(unknown)
Does it use a managed PKI?	Yes
Who manages PKI?	One or more entities in membership services.
Block-release Timing	(unknown)
Transaction Size	(unknown)
Transaction Rate	> 10k tx/sec.
Consensus Model	Pluggable consensus framework; 2 plugins provided: PBFT, and “dummy” plugin
Mining	N/A

5.3. Hyperledger Sawtooth

Hyperledger Sawtooth is a blockchain implementation published by Intel. The goal of the project is to provide companies a means to deploy their own blockchains. The type of data stored in transactions is determined by what the project calls a *transaction family*. Users can create their own transaction families to store any data they want to save in a blockchain deployed with Hyperledger Sawtooth. The source code includes three implementations of transaction families. One of these transaction families, called “MarketPlace,” enables buying, selling, and trading digital assets.

The protocol supports both permissioned and permissionless implementations. Intel has released the source code on GitHub (<https://github.com/hyperledger/sawtooth-core>), and provides instructions on how to deploy a blockchain using Hyperledger Sawtooth. Transactions are transparent by default, but the

network can optionally use an administration key for sending certain messages. The amount of time needed to create a block of transactions is also configurable to any length of time.

The current source code includes implementations for two different consensus protocols. The first is called Proof of Elapsed Time (PoET). It is a lottery protocol that builds on trusted execution environments (TEEs) provided by Intel’s Software Guard Extensions (SGXs). The protocol is based on the consensus algorithm used by Bitcoin. The other protocol is called “Quorum Voting,” and is an adaption of the consensus protocols used by Ripple and Stellar. The current release of Hyperledger Sawtooth includes software that simulates the PoET algorithm, but this implementation of PoET is not secure because it runs outside of Intel’s SGXs.

At this time, Intel warns against using Hyperledger Sawtooth for any security sensitive applications, as the project is still in the experimental phase. It will probably be a while before Hyperledger Sawtooth can be used to create secure blockchain applications. Hopefully, this project will provide another means to create blockchains. Documentation for this project is available at <http://intelledger.github.io/>.

5.4. Summary of Hyperledger Sawtooth

Table 6 – Summary of Hyperledger Sawtooth blockchains

	Hyperledger Sawtooth
Purpose	Enable companies to deploy their own blockchains.
What kind of data can be stored?	Anything that can be defined by a “transaction family”.
Scripting Languages	Python
Is the ecosystem open?	Configurable
How can one participate?	Download the source code from GitHub and follow their instructions.
Native Currency	Initially provides the MarketPlace transaction family, which can track assets.
Who are the registration authorities?	None, unless optional administration key is used.
Is decision making transparent?	Yes; transaction transparency is default
Does it use a managed PKI?	No
Who manages PKI?	N/A
Block-release Timing	Configurable
Transaction Size	(unknown)
Transaction Rate	(unknown)
Consensus Model	Provides two: Proof of Elapsed Time (PoET), and Quorum Voting
Mining	N/A

5.5. Hyperledger Iroha

Hyperledger Iroha is a distributed ledger project designed to be easy to integrate into infrastructural projects. The main goal of this project is to provide C++, mobile, and web development environments to Hyperledger contributors. Few details have been provided on how this implementation works, but there is some information within the wiki pages of the GitHub repository for Hyperledger Iroha (<https://github.com/hyperledger/iroha>).

A blockchain deployed with Hyperledger Iroha can store two types of data: objects and functions. Stored functions are known as *chaincode*, and they can be set by users who have sufficient permissions. *Sumeragi* is the name of the Byzantine fault tolerant distributed consensus algorithm used by Hyperledger Iroha. Most of the algorithm is based on the B-Chain consensus algorithm. In *sumeragi*, consensus is performed on individual transactions and the global state resulting from all transactions. When a transaction is submitted to the blockchain, $2f+1$ signatures are needed to confirm a transaction, where f is the number of Byzantine faulty nodes the blockchain's network can handle. The order of nodes which validate a transaction is determined by a server reputation system called *hijiri*. The *hijiri* reputation system calculates the reliability of each server based on the time they were registered with membership services, the number of successful transactions they've processed, and any failures that are detected on the server. *Hijiri* performs the following tests on each server:

- A data throughput test
- A version test
- A computational test
- A data consistency test

The developers of Hyperledger Iroha claim that the platform will provide transaction finality within two seconds.

In short, Hyperledger Iroha will provide developers with more options for environments through which they can contribute to Hyperledger.

5.6. Summary of Hyperledger Iroha

Table 7 – Summary of Hyperledger Iroha blockchains

	Hyperledger Iroha
Purpose	Provide tools that integrate easily into existing environments.
What kind of data can be stored?	(unknown)
Scripting Languages	(unknown)
Is the ecosystem open?	(unknown)
How can one participate?	(unknown)
Native Currency	(unknown)
Who are the registration authorities?	(unknown)
Is decision making transparent?	(unknown)
Does it use a managed PKI?	(unknown)
Who manages PKI?	(unknown)
Block-release Timing	2 seconds
Transaction Size	(unknown)
Transaction Rate	(unknown)
Consensus Model	Sumeragi
Mining	(unknown)

6. Steem

Steemit is a blockchain-based social media platform where users earn rewards for posting content that the community considers meaningful. The network stores posts, votes, comments, profiles and follows on the *Steem* blockchain. Steem is built on Graphene, a software platform for deploying application-specific

blockchains. Source code for the Steem network is available on GitHub (<https://github.com/steemit/steem>), and there are instructions on how to download the blockchain locally. Anyone who simply wants to make posts on Steemit can create an account for free that is linked to their Facebook or Reddit account. Posts can be viewed at <https://steemit.com>, and the blockchain can be viewed at <https://steemd.com>. The latter site allows users to view all votes on a specific post or comment, and who made what votes. Users can interact with Steem in more ways by obtaining one or more of the cryptocurrencies underlying the network.

Steem has three different cryptocurrencies: Steem (STEEM), Steem Power (SP), and Steem Dollars (SMD). STEEM is the main unit of currency in the Steem network. The main purpose of STEEM is to easily convert into either SP or SMD. STEEM allows users to maintain liquidity for short periods of time. Steem Power is used to vote for or against content posted in the Steem network. Steem Dollars are meant to be pegged to the US dollar, and can be used to pay for goods or services. The Steem website also allows users to buy STEEM or SP using bitcoins. Users can convert between STEEM and SMD at the current exchange rate on a cryptocurrency exchange website, including Steem's own site at <https://steemit.com/market>.

When a post is made, voting periods are set by the blockchain, after which a payout is made to the user who created the post. The payouts are 50% SMD and 50% SP. From the wallet page on the Steem website, users can convert STEEM into SP, which is known as “powering up”, or convert SP into STEEM, which is referred to as “powering down”. Powering up can happen instantaneously, but when powering down, SP is converted to STEEM over two years via 104 equal weekly payments.

Steem uses a Delegated Proof-of-Stake (DPOS) consensus model. In this model, blocks are produced in separate rounds, each of which has 21 witnesses selected to create and sign blocks. Nineteen of these witnesses are chosen by approval voting, an additional one is selected by a computational proof-of-work, and the last is timeshared by every witness that wasn't in the top 19. Each round, a miner is taken from a queue of miners and added to the active set of witnesses. The miner can then earn a reward if they produce a block while they are a scheduled witness. Like most blockchain networks, the miner must create a proof-of-work to create a block of transactions. The calculation uses the SHA256 hashing algorithm twice. Below is an algorithm showing how the proof-of-work is calculated:

Let H = Head Block ID

Let H2 = SHA256(H+NONCE)

Let PRI = Producer Private Key

Let PUB = Producer Public Key

Let S = SIGN(PRI, SHA256(H))

Let K = RECOVER_PUBLIC_KEY(H2, S)

Let POW = SHA256(K)

POW is the value of the proof-of-work necessary to mine a block. Once a miner creates this value, they earn a reward in the form of SP.

The Steem network outperforms Bitcoin in terms of transaction rate and transaction size. Steem is designed to create a block every three seconds, as opposed to 10 minutes for Bitcoin. Also, because Steem is built on Graphene, the same technology used by BitShares, Steem can process over 10,000 transactions per second, which is much faster than the Bitcoin network. Furthermore, the average size of a transaction is 250 bytes in Bitcoin, but is only 100 bytes for Steem.

Steem provides a unique implementation, supporting a use case for a targeted audience. While Steem is limited in the type of application data it can store, it provides more options than Bitcoin. The Steemit application is still in beta, yet it is already in wide use. This network shows great promise for usability with a variety of applications, but there could be some doubts for reasons that will be explained next. Refer to the whitepaper, *Steem: An incentivized, blockchain-based social media platform*, for more information on the Steem protocol. More details on the Steem project can be found at <https://steem.io/>.

6.1. Summary of Steem

Table 8 – Summary of Steem blockchain

	Steem
Purpose	Social media platform that rewards users for meaningful posts.
What kind of data can be stored?	Posts (text and pictures), votes, comments, profiles, and follows
Scripting Languages	(unknown)
Is the ecosystem open?	Yes
How can one participate?	Steemit: Make an account online and obtain assets. Blockchain: Download source code and follow their instructions.
Native Currency	Steem (STEEM), Steem Power (SP), and Steem Dollars (SMD)
Who are the registration authorities?	N/A
Is decision making transparent?	You know who upvotes what, but forking process isn't clear.
Does it use a managed PKI?	No
Who manages PKI?	N/A
Block-release Timing	3 seconds
Transaction Size	100 bytes avg.
Transaction Rate	10k transactions/sec.
Consensus Model	Delegated Proof-of-Stake
Mining	Proof-of-work

6.2. Steemit Network Hack

Users of the Steem network were recently hacked. On July 14th, 2016, around 260 accounts were compromised, with nearly \$85,000 worth of Steem Dollars and Steem stolen.

The day of the hack, some users noticed that Steem funds from their accounts on Bittrex, a US-based cryptocurrency exchange, were mysteriously transferred to an unknown Bittrex account. According to Steem CEO Ned Scott, "...the Steem blockchain was never hacked. Likewise, our servers were never hacked. Instead, the hacker exploited browser-side vulnerabilities..." (3). After the hack, Steem programmers worked on preventative measures against this kind of attack. But immediately after this issue was resolved, Steem servers were hit by a DDoS attack, at which time the Steemit site was taken down to mitigate the damage.

On July 17th, the Steem producers released a post describing a new approach to securing user accounts through multi-factor authentication. The post was likely in response to the attack on Steem's network. Then on July 19th, the Steem producers described a process to recover accounts for users whose accounts were affected by the July 14th hack. The process required users to remember a password that was valid in their account within the last 30 days.

This hack demonstrates that relatively new networks can be a target for attacks. Steem is still considered to be in beta at the time of writing this report, and it has already been targeted by two related attacks. In the case of the first attack, the flaw was not in the Steem network itself. Rather, it was reportedly in the browser interface with Bittrex. This is similar to the Ethereum hack, where a flaw in an organization utilizing Ethereum was compromised. Thus, it goes to show that whether the blockchain network itself is secure is meaningless if the software that uses the network is insecure. The possibility of a system exploit should be taken into initial consideration when deciding on a blockchain implementation.

7. Sidechains – Elements Project

On October 22nd, 2014, a group of people released a paper titled *Enabling Blockchain Innovations with Pegged Sidechains*. The paper proposes an innovative technology called *pegged sidechains*, blockchains that are linked to another blockchain. In the paper, they define a *sidechain* as “a blockchain that validates data from other blockchains,” then they define a pegged sidechain as “a sidechain whose assets can be imported from and returned to other chains; that is, a sidechain that supports two-way pegged assets” (4). Later that year, the same group of people formed the company Blockstream, whose goal is to enable the creation of sidechains. Then on June 6th, 2015, Blockstream announced the release of Sidechain Elements, a project that includes working code and a testing environment for creating pegged sidechains. This was later renamed as the Elements Project.

There are two ways to work with the Elements Project: one can either join the Alpha experimental sidechain, or they can create their own sidechain. Both can be done with the source code available on GitHub (<https://github.com/ElementsProject/elements>). The Elements Alpha sidechain focuses on cryptocurrency transactions, and uses a scripting language like that of Bitcoin but with several improvements. The Alpha sidechain is pegged to the Bitcoin testnet, as is any sidechain created using Elements. Currently, the only assets that these sidechains will track are its own native currency called *hostcoin*. The ability to pay using other assets is being developed.

The performance of either the Alpha sidechain or any pegged sidechains that users can create has not yet been tested. However, when creating one's own sidechain, by default the time to create a block is 60 seconds. This time can be modified in the appropriate configuration file.

As mentioned on the Blockstream website, the Elements Project is intended for research and development. Thus, this project should not be used with real-world assets, and is not yet fully secure. While this project is not yet complete, it demonstrates how sidechains could eventually become another way to implement a blockchain. The Elements Project is described in more detail on its main page (<https://www.elementsproject.org>). The specifics of the protocol are described in the whitepaper, *Enabling Blockchain Innovations with Pegged Sidechains*.

7.1. Summary of Elements Project

Table 9 – Summary of Elements Project blockchains

	Elements Project
Purpose	Create blockchain applications that utilize Bitcoin blockchain
What kind of data can be stored?	Cryptocurrency transactions
Scripting Languages	Language similar to Bitcoin, but with enhancements
Is the ecosystem open?	Alpha sidechain: Yes Sidechain you create: (unknown)
How can one participate?	Download the source code from GitHub and follow their instructions. Then join Alpha sidechain or create your own sidechain.
Native Currency	hostcoin
Who are the registration authorities?	(unknown)
Is decision making transparent?	(unknown)
Does it use a managed PKI?	(unknown)
Who manages PKI?	(unknown)
Block-release Timing	Configurable (default: 60 sec.)
Transaction Size	(unknown)
Transaction Rate	(unknown)
Consensus Model	(unknown)
Mining	(unknown)

8. Lisk

In addition to the development of sidechains through the Elements Project, the Lisk Foundation is also working on sidechain technology using the blockchain implementation they created, known as Lisk. The goal with Lisk is to enable users to create their own applications on separate sidechains that are connected to the main blockchain, or *mainchain*, of Lisk. The mainchain simply records LSK transactions between accounts. Any additional functionality must be programmed within a sidechain linked to the mainchain. According to Max Kordek, co-founder and CEO of Lisk, “This will give millions of developers the ability to create their own sidechains, particularly around consumer applications, including games, social networks, and the Internet of Things, but the same core functionality can also be used to develop and scale business applications” (5). The platform is written in JavaScript, and utilizes NodeJS and PostgreSQL in the backend. New users can participate by logging in through the main website and buying the native currency, denoted as LSK. The website also has downloads for a program to send and receive transactions on the Lisk network, and another program to enable API calls. Furthermore, the source code is available on GitHub (<https://github.com/LiskHQ>), along with a wiki that includes documentation about the platform.

Consensus and mining is done through a Delegated Proof-of-Stake model like the Steem network, only slightly different. The Lisk network is protected by 101 active delegates, each of which is elected by the stakeholders of LSK. All delegates are ranked by the number of votes they have received, and only the top 101 delegates are considered active. Only active delegates can generate new blocks. The consensus process contains multiple rounds. At the start of each round, the order in which delegates can forge relative to each other is determined. Then, each delegate forges exactly one block, which can include up to 25 transactions. Unlike Bitcoin, a proof-of-work is not required to create blocks. A delegate earns a

fixed amount of LSK if they successfully submit a block that is accepted by the system. In addition, participants in each round earn a portion of the transaction fees spent on transactions submitted during that round.

Blocks are created within 10 seconds, and the network can support up to 250 transactions per 10 seconds, or 25 transactions per second. The type of a transaction determines its maximum size, which can be up to 1223 bytes.

Lisk uses the Edwards-curve Digital Signature Algorithm (EdDSA) for hashing, which the developers believe is faster than the ECDSA used by Bitcoin. A user generates a private key-public key pair to participate in the network. The private key allows the user to sign transactions, while the public key is included in transactions the user sends and verifies the validity of the signature. For additional security, Lisk allows users to register a second pass phrase associated with their public key, thus requiring all subsequent transactions to be signed using the second pass phrase to be considered valid.

Since the Lisk platform is still in development, its full capabilities are still unknown. Hopefully, the protocol will progress to a point where users can create a wide variety of applications on sidechains.

8.1. Summary of Lisk

Table 10 – Summary of Lisk blockchain

	Lisk
Purpose	Allow users to create their own blockchain apps on sidechains.
What kind of data can be stored?	Mainchain: Cryptocurrency transactions. Sidechains: (unknown)
Scripting Languages	JavaScript
Is the ecosystem open?	Yes
How can one participate?	Create a new account on Lisk webpage, and buy LSK. Then create new sidechain by following instructions in documentation.
Native Currency	LSK
Who are the registration authorities?	(unknown)
Is decision making transparent?	Delegates (those who can mine currency) are known
Does it use a managed PKI?	(unknown)
Who manages PKI?	(unknown)
Block-release Timing	10 seconds
Transaction Size	1223 bytes maximum
Transaction Rate	25 tx/sec.
Consensus Model	Delegated Proof-of-Stake
Mining	Part of the consensus model; no proof-of-work

9. Summary of Networks

After reading through the previous descriptions of various blockchain implementation, one should generally understand how each blockchain functions and the pros and cons of each. The following tables summarize the important aspects of all blockchains discussed in this paper.

Table 11 – Summary of all blockchains: Bitcoin, Litecoin, Dogecoin, Ethereum

	Bitcoin	Litecoin	Dogecoin	Ethereum
Purpose	Cryptocurrency	Cryptocurrency	Cryptocurrency	Run smart contracts
What kind of data can be stored?	Cryptocurrency transactions, plus some additional data in coinbase or OP_RETURN transactions	Cryptocurrency transactions, plus some additional data in coinbase or OP_RETURN transactions	Cryptocurrency transactions, plus some additional data in coinbase or OP_RETURN transactions	Cryptocurrency, digital assets, smart contracts
Scripting Languages	Script	Script	Script	Solidity, Serpent, LLL
Is the ecosystem open?	Yes	Yes	Yes	Yes
How can one participate?	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.	Download the source code from GitHub, and follow their instructions. Obtain currency from online trading service.
Native Currency	bitcoin (BTC)	litecoin (LTC)	Dogecoin (DOGE)	ether (ETH or ETC)
Who are the registration authorities?	N/A	N/A	N/A	N/A
Is decision making transparent?	Yes	Yes	Yes	Yes
Does it use a managed PKI?	No	No	No	No
Who manages PKI?	N/A	N/A	N/A	N/A
Block-release Timing	10 minutes	2.5 minutes	60 sec.	12 sec.
Transaction Size	200 bytes minimum, 250 bytes avg.	(unknown)	(unknown)	Theoretically no max (actual max: 89 kB)
Transaction Rate	3 transactions/sec. avg., 7 transactions/sec. theoretical maximum	28 transactions/sec. theoretical maximum	70 transactions/sec. theoretical maximum	Theoretically no maximum
Consensus Model	Nodes verify blocks and transactions, and	Nodes verify blocks and transactions, and	Nodes verify blocks and transactions, and	Similar to Bitcoin, but uses Ethereum Virtual Machine

	Bitcoin	Litecoin	Dogecoin	Ethereum
	select blockchain with the most blocks.	select blockchain with the most blocks.	select blockchain with the most blocks.	
Mining	Proof-of-work	Proof-of-work	Proof-of-work	Proof-of-work using Ethash algorithm

Table 12 – Summary of all blockchains: MultiChain, Hyperledger Fabric, Hyperledger Sawtooth, Hyperledger Iroha

	MultiChain	Hyperledger Fabric	Hyperledger Sawtooth	Hyperledger Iroha
Purpose	Provide a platform for creating your own blockchain.	Enable the creation of blockchains for industry use cases.	Enable companies to deploy their own blockchains.	Provide tools that integrate easily into existing environments.
What kind of data can be stored?	Any digital asset you want to store.	Chaincode (i.e. smart contracts)	Anything that can be defined by a “transaction family”.	(unknown)
Scripting Languages	N/A	Go (golang), Java (in progress)	Python	(unknown)
Is the ecosystem open?	Configurable	No	Configurable	(unknown)
How can one participate?	Install MultiChain app, and follow online instructions to make a blockchain.	Create a blockchain: Download source and follow instructions. Join existing network: Register with a proof of identity to the network membership services.	Download the source code from GitHub and follow their instructions.	(unknown)
Native Currency	N/A	N/A	Initially provides the MarketPlace transaction family, which can track assets.	(unknown)
Who are the registration authorities?	Configurable	Defined before blockchain is initialized, and more can be assigned while running.	None, unless optional administration key is used.	(unknown)
Is decision making transparent?	Configurable	(unknown)	Yes; transaction transparency is default	(unknown)
Does it use a managed PKI?	No	Yes	No	(unknown)
Who manages PKI?	N/A	One or more entities in	N/A	(unknown)

	MultiChain	Hyperledger Fabric	Hyperledger Sawtooth	Hyperledger Iroha
		membership services.		
Block-release Timing	Configurable	(unknown)	Configurable	2 seconds
Transaction Size	Maximum size configurable	(unknown)	(unknown)	(unknown)
Transaction Rate	Configurable	> 10k transactions/sec.	(unknown)	(unknown)
Consensus Model	Fixed ratio of admins approves privilege changes. Longest valid blockchain adopted as global consensus.	Pluggable consensus framework; 2 plugins provided: PBFT, and “dummy” plugin	Provides two: Proof of Elapsed Time (PoET), and Quorum Voting	Sumeragi
Mining	Round-robin system; proof-of-work requirement is configurable	N/A	N/A	(unknown)

Table 13 – Summary of all blockchains: Steem, Elements Project, Lisk

	Steem	Elements Project	Lisk
Purpose	Social media platform that rewards users for meaningful posts.	Create blockchain applications that utilize Bitcoin blockchain	Allow users to create their own blockchain apps on sidechains.
What kind of data can be stored?	Posts (text and pictures), votes, comments, profiles, and follows	Cryptocurrency transactions	Mainchain: Cryptocurrency transactions. Sidechains: (unknown)
Scripting Languages	(unknown)	Language similar to Bitcoin, but with enhancements	JavaScript
Is the ecosystem open?	Yes	Alpha sidechain: Yes Sidechain you create: (unknown)	Yes
How can one participate?	Steemit: Make an account online and obtain assets. Blockchain: Download source code and follow their instructions.	Download the source code from GitHub and follow their instructions. Then join Alpha sidechain or create your own sidechain.	Create a new account on Lisk webpage, and buy LSK. Then create new sidechain by following instructions in documentation.

	Steem	Elements Project	Lisk
Native Currency	Steem (STEEM), Steem Power (SP), and Steem Dollars (SMD)	hostcoin	LSK
Who are the registration authorities?	N/A	(unknown)	(unknown)
Is decision making transparent?	You know who upvotes what, but forking process isn't clear.	(unknown)	Delegates (those who can mine currency) are known
Does it use a managed PKI?	No	(unknown)	(unknown)
Who manages PKI?	N/A	(unknown)	(unknown)
Block-release Timing	3 seconds	Configurable (default: 60 sec.)	10 seconds
Transaction Size	100 bytes avg.	(unknown)	1223 bytes maximum
Transaction Rate	10k transactions/sec	(unknown)	25 transactions/sec.
Consensus Model	Delegated Proof-of-Stake	(unknown)	Delegated Proof-of-Stake
Mining	Proof-of-work	(unknown)	Part of the consensus model; no proof-of-work

Conclusion

The comparison of different blockchain implementations reveals that no two networks are exactly alike. Separate networks may differ in purpose, performance, governance mechanisms, security algorithms, or some other ways. These differences can be good or bad depending on the requirements for the application utilizing a blockchain. But even now, most implementations focus on cryptocurrency transactions. While these blockchains obviously work well for monetary applications, they may not be optimal for storing other kinds of data. In these cases, it may be necessary to dive into the source code of the implementation to see how one can work around the blockchain's limitations. However, if the code is not open-source, it will be very difficult to make changes that work against the protocol's original purpose.

If none of the implementations listed is a good fit for an application, then it might be time to consider implementing a new blockchain. Building a blockchain can provide control over all factors that were considered in this paper for existing implementations. For example, it would enable the creation of new permissioned or permissionless networks, depending on the application's intended audience. It would allow modification of important performance factors, especially the block creation time. Many people dislike the fact that bitcoin has a 10-minute block creation time, so avoiding this issue is very significant. One could even implement their own PKI tied to the blockchain network so they can provide improved security in the network, and have the power to fix any problems in the infrastructure. Creating a blockchain would allow fine-tuned control over each user's permissions in the network. There may also be other factors that are important for some applications but were not considered in this paper. A blockchain created from scratch can be flexible enough to accommodate such factors. However, the robustness of the security model for a custom blockchain can be hard to assess. Moreover, leveraging open source solutions provides the potential of leveraging expertise of a community of developers akin to crowdsourcing.

The cable industry should consider leveraging blockchains to improve their work. It is currently under investigation which use cases could apply to the cable industry and utilize blockchains. According to Steve Goeringer, there are three use cases that may be a very good fit: “improving trust in content distribution, streamlining complex service delivery in the medical industry by leveraging cable, and providing improved secure digital content production and distribution” (1). There could be more applications that use blockchains and are relevant to the cable industry, but finding such applications will require more research.

Abbreviations and Definitions

1. Abbreviations and Definitions

1.1. Abbreviations

API	application programming interface
BFT	Byzantine fault tolerance
BTC	bitcoin
D.A.O.	Decentralized Autonomous Organization
DOGE	Dogecoin (the currency, not the blockchain)
DPOS	Delegated Proof-of-Stake
ECA	enrollment certificate authority
ECDSA	Elliptic Curve Digital Signature Algorithm
EdDSA	Edwards-curve Digital Signature Algorithm
ECert	enrollment certificate
EOA	externally owned account
ETC	Ethereum Classic
ETH	Ethereum
EVM	Ethereum Virtual Machine
LES	Light Ethereum Subprotocol
LLL	Low-level Lisp-like Language
LTC	litecoin
NIST	National Institute of Standards and Technology
PBFT	Practical Byzantine Fault Tolerance
PKI	public key infrastructure
PoET	Proof of Elapsed Time
RIPEMD160	RACE Integrity Primitives Evaluation Message Digest 160
SGX	Software Guard Extension
SHA256	Secure Hash Algorithm 256
SMD	Steem Dollars
SP	Steem Power
SPV	simplified payment verification
STEEM	Steem (the currency, not the blockchain)
TCA	transaction certificate authority
TCert	transaction certificate
TEE	trusted execution environment
TLS-CA	transport layer security certificate authority

1.2. Definitions

application programming interface	A set of subroutine definitions, protocols, and tools for building application software.
block-release timing	The minimum time to create a new block on a blockchain.
chaincode	Programmatic code on a blockchain, similar to a smart contract.
contract account	An account on the Ethereum blockchain containing a snippet of code that gets triggered by transactions or messages received from externally owned accounts.
Ethash	The proof-of-work algorithm for Ethereum.
Ethereum Virtual Machine	The runtime environment for smart contracts on the Ethereum blockchain.
externally owned account	An account on the Ethereum blockchain that can send transactions and is controlled by private keys.
fork	An event where there are two candidate blocks competing to be part of the longest blockchain.
hash	A digital fingerprint of some binary input.
hard fork	A change in the blockchain protocol where a subset of previously invalid blocks and/or transactions are made valid, or vice-versa.
hostcoin	The native currency of sidechains created through the Elements Project.
mainchain	The main blockchain within a blockchain network.
miner	A network node that finds valid proof-of-work for new blocks, by repeated hashing.
mining	The process of adding transaction records to a blockchain ledger.
mining diversity	A parameter in a MultiChain blockchain that sets the minimum proportion of miners required to participate in round-robin mining.
nonce	A counter used for the proof-of-work algorithm.
OP_RETURN	A script opcode used to mark a transaction output as invalid.
proof-of-stake	A type of algorithm by which a cryptocurrency blockchain network aims to achieve distributed consensus.
proof-of-work	A piece of data which is difficult to produce but easy to verify.
simplified payment verification	A method for verifying if specific transactions are included in a block without downloading the entire block.
soft fork	A change in the blockchain protocol where a subset of previously valid blocks and/or transactions are made invalid.
software fork	A copy of source code from one software package that is modified independent from the original software.
stack-based language	A programming language that relies on a stack machine model for passing parameters.

Bibliography & References

- Goeringer S. A Simple Overview of Blockchains: Why They Are Important to the Cable Industry. Cable-Tec Expo 2017. 2017 Oct.
- Antonopoulos A. M. Mastering Bitcoin. Sebastopol (CA): O'Reilly Media, Inc.; 2014. Chapter 4, Keys, Addresses, Wallets; p. 66.
- Scott, N. Second Update to July 14 Security Announcement from Steemit CEO Ned Scott. Steemit. 2016 July 18. [accessed 2016 Sept 19]. <https://steemit.com/announcement/@midou05/second-update-to-july-14-security-announcement-from-steemit-ceo-ned-scott>.
- Back A, Corallo M, Dashjr L, Friedenbach M, Maxwell G, Miller A, Poelstra A, Timón J, Wuille P. Enabling Blockchain Innovations with Pegged Sidechains. Blockstream. 2014 Oct 22. [accessed 2016 Aug 1]. <https://blockstream.com/sidechains.pdf>.
- Lisk Releases First Modular Cryptocurrency with Sidechains. Bitcoin PR Buzz. 2016 May 24. [accessed 2017 July 14]. <http://bitcoinprbuzz.com/lisk-releases-first-modular-cryptocurrency-with-sidechains>.

IT Data Security in An MSO Environment

A Technical Paper Prepared for SCTE•ISBE by

Robert Gyori

Group Vice President
Information Technology, Security & Compliance
Charter Communications
13736 Riverport Dr.
St Louis MO. 63043
314-388-8820
Robert.gyori@charter.com

Introduction

Are you in the process of launching a Security Program to support your Cable Company or are you looking to improve your existing program? This whitepaper and its supporting presentation will help you get started. This document should help you identify what needs to be protected and some basic protection measures.

IT Data Security in an MSO Environment

This discussion will walk you through some of the basic questions you should consider if you are standing up a security program.

1. Do you know what types of data to protect and how to protect it?
2. What are your risk measurement criterion?
3. Are you using a standards based approach to secure your data and critical infrastructure?
4. How are you protecting and validating your compliance landscape?
5. How are you handling Identity Management?

This discussion will walk through some of the basic concepts around protecting critical data assets in our complex MSO networks...

What do you need to secure your data?

Answer: People, Process, & Technology...

With a solid foundation of policies/standards...




Figure 1 - People Process Technology

Take-away concept: These are foundational building blocks for securing your information...

The “Why”

Why do Cyber events happen in the MSO space?

What are the motives?

1. In the Service Provider Space
 - 1.1. Sample Event: DDOS Attacks against the backbone
 - 1.2. Anger
 - 1.3. Hacktivism
 - 1.4. An attempt to disrupt communications (internet access, news etc.)
 - 1.5. An attempt to penetrate laterally to the Crown Jewels...
2. In the IT and/or Enterprise Support Space
 - 2.1. Money is THE motive
 - 2.2. Sample Event: Corporate Financial Instrument theft, Customer/Employee PII/PCI/Sensitive Data (incl: CC)

Take-away concept: *Focus on items that lead to the most \$ (lost or stolen)*

Discussion Point: *what is the half-life of PII?*

The “What”

What Cyber events can provide the most risk and/or impact to an MSO?

1. In the Service Provider Space
 - 1.1. Any interruption/degradation in the service we provide to our customers.
 - 1.1.1.E.g. DDOS Attacks against the backbone
 - 1.2. Any interruption/degradation in our ability to service our customers.
 - 1.2.1.Attacks against DNS, Customer Provisioning etc. (BACC/RDU etc.)
2. In the Enterprise Support Space
 - 2.1. Any Incident/Breach that impacted or resulted in the breach of :
 - 2.1.1.Corporate Financial Instruments
 - 2.2. E.g. Large scale unauthorized wire transfers etc.
 - 2.2.1.Customer PII/PCI/Sensitive Data
 - 2.2.2.Employee PII/PCI/Sensitive Data

Internal Vs. External: Background Stats...

1. Vast majority of threats (bad actors) resulting in incidents are external ~%75 (VZ 2016 & 2017 DBIR)
 - 1.1. Internal threats are important and DO exist....
 - 1.1.1.Internal threats have the potential to be devastating
 - 1.1.2.Internal threats can lead to or facilitate an incident from external threats
 - 1.2. Especially in orgs that are heavy in Intellectual Property or DoD based
2. Common thread in most incidents = Employee/Vendor Credentials (stolen and/or weak passwords)
 - 2.1. 63 % in 2016 (per Verizon 2016 DBIR)
 - 2.2. 81 % in 2017 (per Verizon 2017 DBIR)

Take-away concept: *So just how important is Identity/Credential Management to your security model?*

What Data Should you protect?

1. The Crown Jewels! Or... Soo much data so little time...
 - 1.1. MSO's store, process, and transmit data at all layers of their infrastructure. From STB data travelling upstream to billing and conditional access systems, to customer facing “.net” & eCom platforms.
 - 1.2. With so much data how do you know what should be the most protected?
 - 1.3. Or...What data would be the most valuable to a bad actor?
 - 1.4. Hint...Start with a Data Classification Policy applied to your data

Take-away concept: *This is one of the many reasons that having sound policies is foundational....without a Data Classification Policy, it's hard to know what data is most valuable.*

How should you protect your data/systems?

1. Start with a standard based framework approach
 - 1.1. NIST Framework for Improving Critical Infrastructure Cyber Security
 - 1.2. NIST 800-53 Security & Privacy Controls
 - 1.3. ISO 27001/27002
 - 1.4. SANS Top 20
2. Write your policies/standards to map to your Framework
 - 2.1. Most of the controls from one framework can be mapped to the others
 - 2.1.1. For Example, SANS Critical Control #1:
 - 2.1.1.1. Inventory of Authorized & Unauthorized Devices
 - 2.1.1.2. Maps to > NIST 800-53 (CM-8, a,c,d,2,3,4 & PM 5, PM 6)
3. Implement your tools & develop your process based upon your framework
 - 1.1. Throw in some Capability Maturity Modelling (CMM) for leavening...

Take-away concept: *Most standards have significant over-lap and can be cross-referenced etc.*

What about that Data Classification Policy?

Some Data Classification Samples: (Every company must define their own policy)

- **Public:** Information that does not fall within one of the more restrictive categories and that can be or has been made available to the public without any financial, legal or other implications to the Company
 - Examples (nonexclusive): Information in the public domain, on public websites, released press releases, published marketing materials, published annual reports, publically filed documents, etc.
- **Internal Only:** Information that is not Restricted or Sensitive and which is not approved for general circulation outside the company, where its disclosure would inconvenience the company, but is unlikely to result in significant financial loss or serious damage.
 - Examples (nonexclusive): internal memos, internal project reports, minutes of meetings, unreleased press releases, unpublished marketing materials, competitive analysis, internal non-proprietary policies, processes or procedures.
- **Restricted:** Information that is not Sensitive and which is considered critical to the organization's ongoing operations and could seriously impede or disrupt them if disclosed without authorization or made available to the public.
 - Examples (nonexclusive): accounting information, business plans, Personally Identifiable Information (PII) about customers or employees, etc.
- **Sensitive:** Any highly confidential internal information about customers or employees or other strategic or financial information which the loss of confidentiality, integrity, or availability could be expected to have an adverse effect on the company. The highest levels of integrity, confidentiality, and restricted availability are vital.
 - Examples (nonexclusive): customer or employee social security or tax identification numbers, driver's license or state issued identification numbers, financial or payment card information, impending mergers or acquisitions, investment strategies, etc.

Applying your Policy to specific use cases

Data Access Policy Sample for Storage

PHYSICAL STORAGE/ACCESS: defines how information may be stored when in a physical format. Inclusive of paper records and when electronic information is stored on a physical Medium (e.g., backup tapes, CDs, etc.).

Table 1 - Sample for Storage

Classification	Policy
Public	No Restrictions
Internal Only	<ul style="list-style-type: none">• Protect from inadvertent or unauthorized disclosures• If physical media contains electronic data, then it also must be protected• Storage under lock and key• Only authorized users may have access
Restricted	<ul style="list-style-type: none">• Access list must be reviewed periodically by business owner.
Sensitive	<ul style="list-style-type: none">• Access to or removal of Information may only be granted with management approval.• Must be kept under double lock and key• Attempted or actual unauthorized access, use or disclosure must be immediately reportable to compliance teams• Appropriate entry controls must be used to limit and monitor physical access to areas containing physical media containing payment card data.

Data Access Policy Sample for Database Access

ELECTRONIC STORAGE/ACCESS: how information may be stored or accessed when in an electronic format (such as in a database)

Table 2 - Sample for Databases

Classification	Policy
Public	No Restrictions
Internal Only	<ul style="list-style-type: none"> • May be stored in unencrypted format. • Must have individual access controls where possible and appropriate. • Only authorized users may have access.
Restricted	<ul style="list-style-type: none"> • May be stored in unencrypted format. • Must have the following access controls: <ul style="list-style-type: none"> • Must have individual access controls that restrict access to active Users and active User account only • Must assign unique IDs and passwords, which are not vendor supplied defaults, to each Users • When being accessing via a public network, logon credentials may not be passed in clear text • Access list must be reviewed periodically by the business owner.

Data Access Policy Sample for Database Access (Continued)

Classification	Policy
Sensitive	<ul style="list-style-type: none"> • Information must be encrypted • Must have the following access controls: <ul style="list-style-type: none"> • Individual access controls that restrict access to active User account only • Restrict access to those who need information to perform their legitimate job duties • Block access to any User ID after multiple unsuccessful attempts to gain access • Assign unique IDs and passwords, which are not vendor supplied defaults, to each user • When being accessing via a public network, logon credentials may not be passed in clear text • Access must be logged for a minimum of 30 days. • Access list must be reviewed on a quarterly basis by business owner. • Attempted or actual unauthorized access, use or disclosure must be immediately reported. • Within the Payment Card Industry (PCI) payment card environment, users, system administrators, and vendors accessing systems that store data classified as restricted or sensitive from outside the network must use additional criteria for authentication. • Such criteria are commonly referred to as “two factor authentication”. <ul style="list-style-type: none"> • Something you know, such as a password or passphrase • Something you have, such as a token device or smart card • Something you are, such as a biometric • For systems within the PCI payment card environment, only authorized users may query databases directly. • Standard user accounts must not have the ability to directly query databases that house sensitive or restricted data.

Policy Sample for Data Retention How long data is retained in its readable state before being subject to deletion or destruction. Data Retention is not defined by classification or categorization. Data retention policies are defined case-by-case for data stored in a database for the purposes of the application. The retention policy is defined to be greater than the legal minimum and the lesser of:

1. Legal maximum data retention requirement
2. Application requirement to have the data available




Figure 2 Data Retention 1




Figure 3 - Data Retention 2

Table 3 - Data Retention

Group	Responsible for
Legal	Definition of the data retention policy for the data in question
Application Owner	Non-functional requirements pertaining to the retention of data
System Administrator	Enforcement, monitoring, and reporting of the data retention policy

Policy Sample for Data Disposal/Destruction

Table 4 - Data Disposal

Classification	Policy
Public	No Restrictions
Internal Only	Electronic data must be erased or rendered most likely unreadable.
Restricted	Electronic data must be completely erased or rendered difficult to retrieve.
Sensitive	Electronic data must be completely erased and overwritten or rendered reasonably unrecoverable.

Protect Databases from Attacks

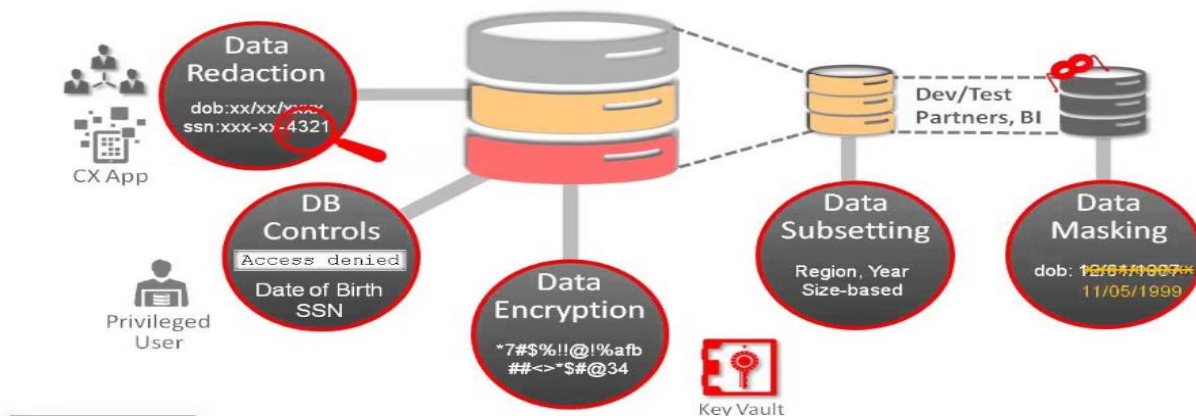


Figure 4 - Database Protection

Protecting databases from attacks

Sample Data Protection Policies/Standards

Table 5 - Data Protection

Technology	Description	Platform(s)
Encrypted Database	All files in the DMBS (database management system) are encrypted with keys known only to the administrators. If the database is separated from the application or the installation, the data is unrenderable	Oracle, MSSQL, MySQL, Teradata
Encrypted Tablespace	All tables in a defined tablespace are encrypted with keys known only to the administrators. If the data in the table is separated from the tablespace, the data is unrenderable	Oracle, MSSQL, MySQL, Teradata
Encrypted Column	All data in a defined column of a table is encrypted with keys known only to the administrators. If the data in the columns is separated from the table, the data is unrenderable	Oracle, MSSQL, MySQL, Teradata
Key Vaults	Repository with encryption keys for real-time data access via certificate, used to unencrypt data for use. This prevents applications from having native keys.	All
Access Controls and Role based privileges	Logins with defined permissions for data access. Roles are by default "deny all" and permissions to view data in databases must be <u>added</u> .	All
Encrypted Storage	Encryption of the data at rest in the storage servers. This is storage technology, not database technology, but applies to many of the databases in a N-tier architecture (where compute is separate from storage). Encryption of the stored data at-rest renders data unreadable without the separately stored keys.	EMC, Hitachi, NetApp, Violin

Sample Practices to implement Data Protection Policies/Standards

Table 6 - Data Protection Practices

Preventative	Detective	Administrative
Encryption Using technology to cypher data in such a way that it cannot be rendered without the cypher keys	Monitoring Using technology to monitor behaviors of data users, and alerting administrators of abnormalities	Governance Tracking and cataloging of data access granted to users. Regular audits of users and applications
Data separation The practice of storing partial data sets in separate data stores, resulting in the need for multiple data breaches in order to lose meaningful data	Firewalls Devices with abilities to detect abnormal traffic to and from data stores (databases and storage)	Key Management Storing the keys for encrypted data separate from the data or application. This practice essentially means you have to steal the database and the DBA to compromise the data.
Data classification Defined classifications for data sets for the purposes of defining the appropriate data protection level	Human Intelligence Anonymous portal for solicitation and documentation of events with potential security ramifications	User Management Enforcement of best practices for password rotation and complexity
User Controls Minimalist approach to data access. Individual users for every person or application with legitimate data access needs. Pre-defined roles for categories of users. Regular rotation of passwords.		Configuration Patching and upgrades of environments to prevent data compromise

Layered Security along the path...Protecting Data in Motion

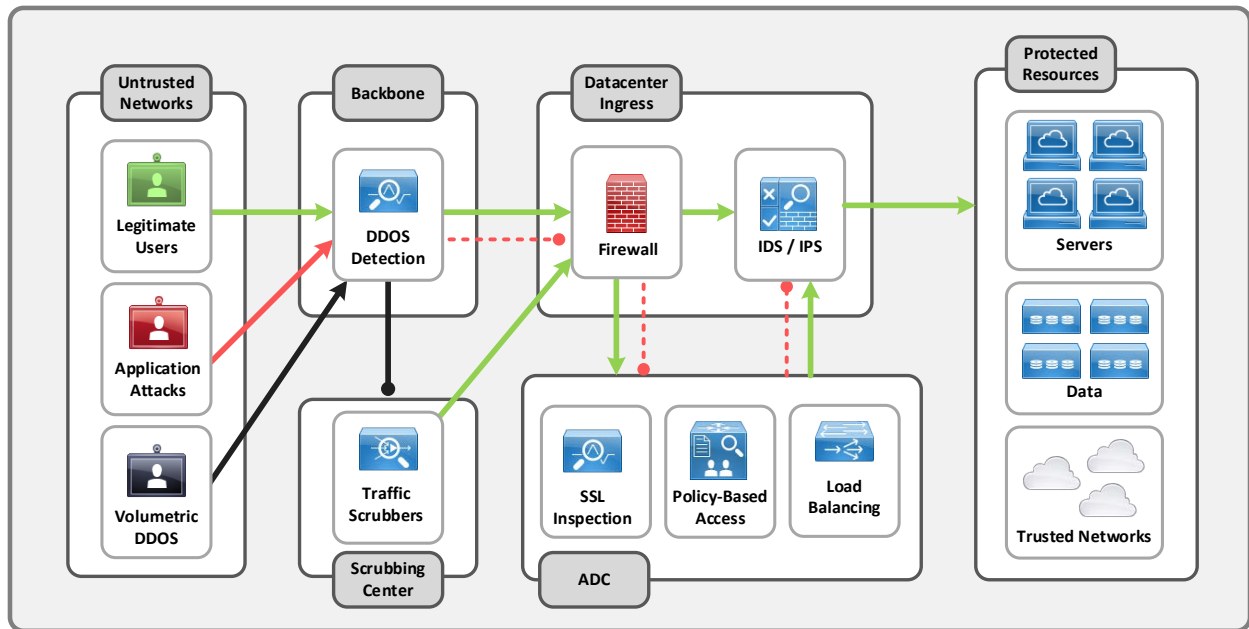


Figure 5 - Layered Security

Threat Prevention & Mitigation Sample Tools

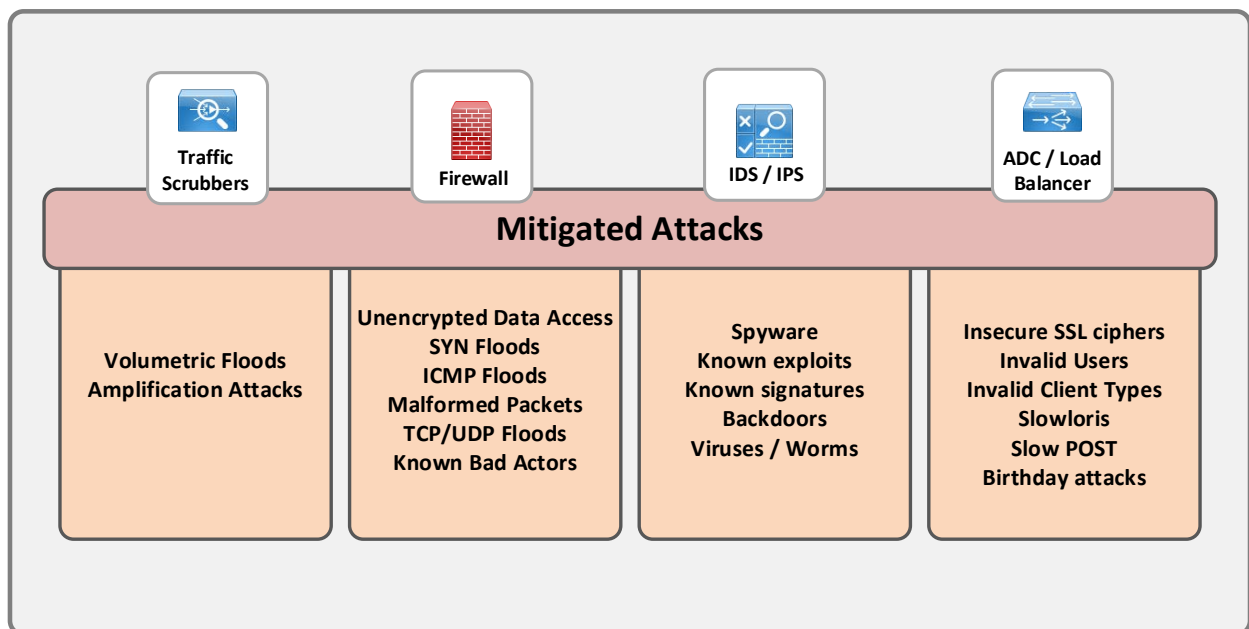


Figure 6 - Threat Prevention & Mitigation

Using Protocol Discretion

Table 7 - Protocol Discretion

Usage	Insecure Protocols (prohibited)	Secure Protocols (allowed)
Web Services	HTTP, SSL, TLS1.0	HTTPS, TLS1.2/1.1
File Transfer	FTP, RCP	SFTP, SCP
Remote Shell	Telnet, SSH1	SSH2
Remote Desktop	VNC	RDP

Compliance Obligations

Sectoral Approach to Privacy and Data Security in the U.S.

1. Sector-specific federal legislation (financial services, health care, and education) and marketing restrictions.
2. State laws fill gaps or raise standards (e.g., consumer privacy, breach notification, and data security).
 - 2.1. E.g., State information security laws requiring “reasonable” security (i.e., Massachusetts information security regulations).
 - 2.2. Secure data disposal and SSN protection laws.
3. Industry standards, voluntary codes, and government guidance also play key role.
4. Various state and federal agencies enforcing privacy and data security laws and regulations, including the Federal Trade Commission (FTC), the Federal Communications Commission (FCC), etc.

Sometimes Compliance Obligations may drive your security program

In our regulated business, compliance is usually not an option...

1. Sarbanes-Oxley Act of 2002 (SOX)

2. Health Insurance Portability and Accountability (HIPAA) and Health Information Technology for Economic and Clinical Health Act of 2009 (HITECH)
3. Customer Proprietary Network Information (CPNI)
4. Payment Card Industry (PCI)

SOX

Sarbanes Oxley

1. Focus on strengthening financial reporting in publicly-traded companies in response to large-scale financial scandals, such as Enron & WorldCom.
2. Provides corporate governance guidelines for companies, their boards, and their auditors.
 - 2.1. Does not explicitly mention data security issues; however, the auditing standards emphasize that upper-level management is responsible for ensuring the integrity of controls on data privacy and security.
 - 2.2. Controls on data stored with 3rd parties significant part of a SOX-compliant audit report

HIPAA / HITECH

Health Insurance Portability and Accountability Health Information Technology for Economic & Clinical Health

Protecting Employee data

1. Together, HIPAA and HITECH require implementation of data security requirements to protect the privacy and security of “protected health information” (PHI).
2. HIPAA’s information security requirements apply to “covered entities” (i.e., health plans and healthcare providers), and HITECH’s amendments expanded application to “business associates” that perform business services for the covered entity.
3. HIPAA also requires that covered entities conduct a security risk analysis to assess the potential risks and vulnerabilities of all electronic PHI created, received, maintained, or transmitted.
4. Need coordination between IT/NetOps etc. (employee data vs customer data)

CPNI

Customer Proprietary Network Information

1. The Federal Communications Commission (FCC) imposes detailed regulations that limit access, use and disclosure of CPNI in the context of voice services.
 - 1.1. CPNI is certain types of information (i.e., subscription information, call detail records, etc.) collected by telecommunications carriers related to their subscribers.
2. The CPNI regulations require carriers to take reasonable measures to discover and protect against attempts to gain unauthorized access to CPNI including:
 - 2.1. The use of proper account authentication; and providing notification to customers when certain account information activity occurs (i.e., after a customer's password, security question and answer, online account, or address of record is created or changed).
 - 2.2. Rules for CPNI in the ISP are still not finalized. Per FCC guidance MSO's should exercise reasonable good faith efforts to comply with the general nature of CPNI rules.

PCI

Payment Card Industry

1. PCI Data Security Standard (DSS) first took effect in June 2005; current version 3.2 came into effect in April 2016
 - 1.1. Requires merchants and credit card transaction processors that store, process or transmit cardholder data to build and maintain a secure computer network, maintain a vulnerability management program, and regularly monitor and test networks.
 - 1.2. The technical/operational requirements include restricting access to data, encrypting sensitive data transmitted over public networks, the use of firewalls, current virus software, and other data security measures.
 - 1.3. The requirements extend to application development and security processes to ensure PCI is protected throughout the applications lifecycle
2. Note that some states (e.g., Nevada and Washington) have codified compliance with PCI DSS requirements into state statutes.
3. In Scope vs. Out of Scope

- 3.1. Category 1 = Systems that Transmit, Process, Store Credit Card Data
- 3.2. Category 2 = Systems connected to Cat 1
- 3.3. Concept of “infected”
- 4. Drives businesses to improve security
- 5. Drives business to improve security specific to In Scope systems
 - 5.1. Build & Maintain a Secure Network & Systems
 - 5.2. Protect Cardholder Data
 - 5.3. Maintain a Vulnerability Management Program
 - 5.4. Implement Strong Access Control Measure
 - 5.5. Monitor and Test Networks
 - 5.6. Maintain an Info Sec policy.
- 6. PCI drives concepts that should already be in place in a mature security program
 - 6.1. Logging
 - 6.2. Change/Configuration management
 - 6.3. Encryption
 - 6.4. Patch & Vulnerability Management

Do you want to limit your systems considered in-Scope for PCI?

- 1. Don't: Transmit/Process/Store Cardholder Data!
 - 1.1. This actually does have some merit...
- 2. Network Segmentation
 - 2.1. Not a PCI requirement BUT does help reduce scope
- 3. Tokenization
 - 3.1. Replace any CC data in transit with tokens

Reducing your PCI scoped systems, can reduce your PCI cycles/costs

Identity Management & Access Control

- 1. It's difficult to extract/steal sensitive data without a legitimate credential...
- 2. Going back to the stats:
- 3. Common thread in most incidents = Credentials (stolen and/or weak passwords)
 - 3.1. 63 % in 2016 (per Verizon 2016 DBIR)

- 3.2. 81 % in 2017 (per Verizon 2017 DBIR)
- 4. Password Fatigue: Users have to maintain credentials for multiple systems.
 - 4.1. Show of hands....
 - 4.2. Enterprise Single Sign-On (ESSO) helps reduce Password Fatigue BUT makes it even more critical that your identity management/credential management system is rock solid.
 - 4.2.1. But with ESSO a single stolen password could be used to access multiple systems...
- 5. Implement/Upgrade your Identity Management Systems
- 6. Consider a dedicated IDM not keyed to an email account
 - 6.1. Many enterprises only use Active Directory...
 - 6.2. Multiple vendors in the identity management space
- 7. Use IDM as the source of truth instead of HR Systems
 - 7.1. Place IDM at the center and feed other systems
 - 7.2. Your unique ID should never change or be reused
- 8. Manage/Model your IDM Lifecycle
 - 8.1. Identities may start with an HR or Recruiting system
- 9. Upgrade your Active Directory footprint
 - 9.1. Red Forest etc.

How much does a breach cost?

Table 8 - Cost of a Breach

	Anthem (Source: Cnet)	Home Depot (Source: Fortune)	Target (Source: Target)
Damages	\$115,000,000	\$ 179,000,000	\$202,000,000
Customers/Records	80,000,000	50,000,000	40,000,000
Cost per/Record (Calculated)	\$ 1.44	\$ 3.58	\$ 5.05

Conclusion

Maintaining secure systems and staying compliant in a MSO environment requires a multi-faceted cross-functional approach. Active collaboration between Information Technology, Network Operations/Engineering, and Legal is critical to the success of a security program. It should also be stressed again that the approach needs to contain a good mix of People, Process, Technology, and Policies/Standards.

Abbreviations

CC	credit card
CPNI	Customer Proprietary Network Information
DBIR	Data Breach Investigations Report
DDOS	distributed denial of service
ESSO	Enterprise Single Sign-On
FCC	Federal Communications Commission
FTC	Federal Trade Commission
HIPAA	Health Insurance Portability and Accountability
HITECH	Health Information Technology for Economic & Clinical Health
ISO	International Standards Organization
NIST	National Institute of Standards & Technology
PII	personally identifiable information
PCI	Payment Card Industry

Bibliography & References

NIST 800-53 Rev. : National Institute of Standards & Technology, Security and Privacy Controls for Federal Information Systems and Organizations

NIST Framework for Improving Critical Infrastructure for Cybersecurity, Version 1.0

PCI DSS 3.2: Payment Card Industry Data Security Standard, Version 3.2

Assuring Security in the IoT

Implementing a Behavioral Analysis Approach to Thwart IoT Attacks

A Technical Paper prepared for SCTE•ISBE by

David Yates

VP Product Line Management
Guavus, A Thales Company
1800 Gateway Drive, Ste. 160
San Mateo, CA 94404
650-823-0674
David.yates@guavus.com

Introduction

The “Internet of Things” (IoT) is here. Manufacturers across industries are reinventing consumer products—from refrigerators to thermostats—by incorporating smart, connected capabilities to capitalize on the global IoT market. IoT is growing exponentially; Experts forecast upwards of 50 billion connected devices by 2020. The industry estimates that 8.4 billion connected things will be in use worldwide in 2017, up 31 percent from 2016. Total spending on endpoints and services will reach almost \$2 trillion in 2017. While consumers purchase more devices, businesses spend more. In 2017, in terms of hardware spending, the use of connected things among businesses will drive \$964 billion.

These businesses hope to ride the IoT wave to achieve competitive differentiation, cultivate additional revenue streams, deliver new monitoring capabilities that will transform maintenance services, and/or enhance and personalize the customer experience. Yet in the race to accelerate time-to-market for these next-generation products, many companies are giving short shrift to IoT-specific security challenges. IoT security concerns are comparable to enterprise IT concerns but complicated by scale, given the sheer number and variety of endpoints. Too often, organizations fail to account for security considerations during initial product development and test cycles. Alternatively, they opt to release IoT-enabled products quickly, only to address security issues after the fact. Both scenarios set the stage for possible breaches that can impact a company’s brand reputation and bottom line.

Concerns about IoT security are escalating as the number of IoT-enabled products multiplies. Experts have raised the possibility that any device connected to the Internet is at risk for hijacking. They underscore the unique security challenges of IoT, specifically that the small size and limited processing power of connected devices could impede encryption and other robust security measures.

It’s not just about the “things” under ownership or allowed to connect to the network. Imagine a situation where a flaw in a common network connected device was used to launch a Distributed Denial of Service (DDoS) attack—there have already been cases of home cable modems being used in botnets. This means systems may be affected by the insecurity of “things” in general. Similarly, if “things” are designed insecurely and have no mechanism for patching, they can act as a pool of malware capable of attacking other systems. So, the security of IoT is a concern for the entire Internet community, not just device owners. Those devices can affect the entire community, making for a common interest in the fundamental security of them.

Experts advise organizations to start building consumer trust in IoT devices by prioritizing security and by adopting “security by design” practices that integrate security capabilities at the earliest stages of product design. However, companies are still a long way from instituting security-by-design best practices. Most firms in the planning stages of IoT product development remain focused on connectivity, although organizations that have already released IoT-enabled products are grappling with more complex issues, including security.

Given the scope of possible issues, it’s increasingly clear that any business entering the IoT fray needs to consider security at the outset, not after the journey is underway.

Content

I. Business Trends and Pressures

At a high level, there are many business trends and pressures that will impact IoT security in the enterprise. IoT will represent a significantly complex system with low-capability users, creating a rich and evolving risk environment. Consumer and industrial ecosystems are different, in terms of technology, development, and priorities. There will be differing threats, actors, and reasons for targeting IoT. And, there will be additional considerations because of the technology's scale, adoption rates, and regulatory environment.

Cyber and physical worlds converge

The industrial product lifecycle encompasses cyber-physical systems where security equals personal safety. While there are complex governance and management processes, industrial systems often have long-lived and out-of-date information management systems due to underlying hardware. This represents a significant operational impact and financial investment. And includes a concern about connecting industrial systems: The industry may fail to learn lessons from the PC world about the importance of patching, upgrade lifecycle, and built-in security.

The consumer product lifecycle has a faster lifespan—focused on new features and quick time to market with minimal cost, which ultimately causes rapid obsolescence. The short support and upgrade periods means large numbers of devices will be left out there unpatched and forgotten. Some new business entrants won't even survive the length of time their products are in use. Also there will be a low bar to entry for IoT creators, which means there will be less experienced developers who aren't necessarily knowledgeable about security and privacy. The consumer environment makes use of crowd-sourced product development and common libraries; this results in systemic homogeneity that may increase the severity of widespread compromises like Shellshock and HeartBleed.

Different threats, actors, and reasons

The New Style of Business means there are new models to protect. Unfortunately, the criminals are smart and always seemingly one step ahead. With the changing landscape of technology and applications in the connected world, threat actors and attack vectors are expected to morph as well. In the beginning, threat actors will most likely be motivated by fame, focused on the newly interesting, novel technology. They'll want to showcase their hacking expertise and expose vulnerabilities.

But as IoT adoption spreads, the technology will be attacked or compromised based on the value to the attacker—monetary, ideology, or business disruption. Since real-time interactions are key to IoT value, actors may use jamming and interference of communications. This may include misrouting of information, impersonation, or flooding and draining of resources, which could cause a distributed denial of service (DDoS) or at least confusion. Many legacy industrial security controls assume the "protected" perimeter with walled environments. The connectivity into broader networks may inadvertently impact the physical security.

II. IoT Security Pain Points

IoT is ushering in whole new categories of products and services, but it is also creating novel opportunities for cybercriminals and other malicious actors to infiltrate networks and gain unauthorized access to data and systems. IoT hardware has access to sensitive information via a network connection—meaning it must be safeguarded just like any other enterprise endpoint. Unlike standard IT equipment, however, IoT products typically lack the processing power and memory to run antivirus software, firewalls, or other widely used enterprise IT safeguards. Similar to enterprise IT security, threats to IoT-enabled devices come in many forms and flavors. Following are some additional critical areas to consider as IoT evolves:

Outdated security model

Traditional IT security policies and controls will be untenable. The security model for it will need to transform to support all of the new aspects of operational technology security and transition to a data-centric aspect. Security will need to be automated, distributed, context aware, and real time.

Unauthorized Access and Control

By far the most pressing concern among companies that are building IoT products surrounds unauthorized access to an IoT device and/or restricting an unknown entity from taking control of its functions. Consider an early IoT product as a case in point. A smart refrigerator, which supported email and social networking applications, lacked the appropriate security features. As a result, it was compromised as part of a security scam that leveraged 100,000 consumer devices to send 750,000 malicious emails and online attacks. Researchers from the security firm Proofpoint discovered that the refrigerator's email capability served as a conduit for an attacker. Why? Despite its similarity to traditional endpoints, the fridge lacked the antivirus software and firewalls used to protect desktops, laptops, and servers.

Other common methods for hacking into IoT devices include spoofing, when a device pretends to be something or someone it's not, and identity forgery, when an unauthorized or unauthenticated user takes control over a device, or stolen credentials are used to gain access to data.

Also in this category are: **DDoS attacks**, which block a system from providing service by making it unusably slow, consuming available storage, or crashing. As with enterprise IT breaches, denial-of-service attacks in the IoT world are far-reaching, potentially impacting individual devices, vendor device accounts, third-party partners, and even end-user applications.

Compromised Data

A major point of IoT is to collect a treasure trove of data—about the product's physical state, how it's used, and environmental conditions, for example—which can later be mined for insights. If hackers gain unauthorized access to an IoT product, they can alter or maliciously damage the collected data. At best, this mitigates the usefulness of the data; at worst, it puts the company at risk for lost revenue opportunities, unhappy customers, and/or brand damage.

Man-in-the Middle Attacks

As in the enterprise IT world, an unauthorized entity can intercept communication and gain unintended control over an IoT product, such as a smart thermostat or home security system. The intruder could then

engineer a malicious action, such as a home invasion, or turn an IoT device into a vector for stealing payment information.

III. IoT Security Guiding Principles

As the security landscape continues to evolve, so will the threat actors. Currently, there are highly capable threat actors, capitalizing on the prolific black market to buy and sell capabilities and information. This will only continue to grow as additional devices and data sources come online. The growing volume and exchange of data require new technology to protect the user device and data entity. And, the expanding threat landscape and sheer number of devices— some smart, some not—will require adaptive, self-defending, autonomous capabilities.

In the future, there will still be fundamental quality and security requirements for solutions, systems, and devices. This isn't so different from current solutions, but there will be greater emphasis on beginning with the end in mind, because mitigating at the end becomes impossible with the distributed, massive scale of IoT. So, “things” on the Internet need to be designed for security, upgradability, and resiliency. IoT systems need to be safe and reliable with the following underlying attributes:

- **Secure access management**—The things and systems in the IoT ecosystem need to be identified and managed in the same way traditional enterprise systems are controlled. Key processes, which include identification, authentication, and authorization, will become more important because of the sheer quantity and variety of IoT systems. Trust mechanisms will be based on context and value scales, not simply a binary choice.
- **Self-protection**—IoT also needs self-protecting and self-healing systems. These attributes are important since systems will no longer have the advantages of a defined perimeter or enterprise-class managed environment. Some devices may also be specialized gateways or intermediaries that provide additional services and protections that can't be included in low power or small form-factor “things.” Security solutions will need to leverage the added value of crowd-sourcing and peer intelligence to help form a self-protecting mechanism. These mechanisms will be the basis for resiliency at the device level.
- **Privacy controls**— Because data will be created in increasing quantities and situated everywhere, it's imperative that solutions give clear control of the data to the owner or source. Ownership will be complicated due to the distributed nature of the systems and complexities of the governing environments. Security and privacy will need to be addressed directly at each device and interaction—transaction and communication.
- **Embedded security**—Security will need to be deeply integrated in hardware and application software layers. The diverse functionality and small form factors won't be able to withstand generalized, bolted-on security mechanisms. The technical designs will need to use contextual awareness, adaptive security that senses and responds to a range of trust mechanisms.
- **Real-time information processes**—Information will be pervasive, seamless, and integrated across the whole IoT ecosystem. Solutions will need mechanisms to process and leverage the enormous sets of data into information for safe and reliable operations. Information analytics in

IoT will need to be predictive, proactive, and near real time to operate with resiliency. Always-on operations require continuous security features and controls.

IV. Information Analytics in IoT Security

A major consideration in assessing the model of analytics needed for IoT is the various types of analytics and sources of data used in an application. In a traditional model, descriptive and diagnostic analytics (sometimes grouped as historical analytics) are often developed independently and have multiple connection points to the various sources of data. The structured, semi-structured, and unstructured data that is often stored in different data warehouses and logical locations is connected independently and requires multiple connectors to consolidate all the relevant information. This is time and cost prohibitive and makes it difficult to build IoT Analytics applications quickly and meet business imperatives for timely action. It also significantly delays time to value and is not scalable from an economic point of view.

The first step in designing a new approach is to simplify the process by integrating all the data for an IoT application. That includes all the structured, unstructured, and semi-structured data in the picture. This range of data must be integrated for the analytics that will be run on the data. Better business outcomes are achieved when these silos are removed and analytics are used across a broad spectrum of valuable data.

The second key step in the streamlining process is to unify the analytics layer. In the traditional model, descriptive and diagnostic analytics made the problem challenging because of the “siloe” approach to data access. This issue will multiply rapidly in scale and become much more serious with the addition of predictive and prescriptive analytics. The problem is more acute and unworkable for IoT applications. This traditional heterogeneous and one-off approach to types of analytics will not suffice for IoT because it will take significant time and effort for data management vs. focusing on delivering outcomes based on the analytics. The explosion of data in all forms in IoT requires a more robust and broader lens in order to enable smarter timely actions and better outcomes.

All the types of analytics must be unified into a single engine to ensure scalability and real-time performance. This includes historical analytics (descriptive & diagnostic), real-time streaming analytics, predictive analytics, and prescriptive analytics. In addition, a design philosophy of openness and unification is needed to help customers get results rapidly. Businesses looking to deploy IoT applications cannot be expected to “rip and replace” their existing investments, and need approaches to leverage their existing analytics and data investments and migrate them into a larger unified framework. The payoff is that users will now be able to spend more time on insights and business outcomes that matter most and avoid the time and distraction of creating or managing a complex infrastructure.

The approach to analytics outlined above is a good first step for IoT. However, it is the ability to execute analytics (real-time, on-demand, streaming, historical, predictive, and prescriptive) with relevant contextual and situational data that addresses the critical “last mile” for timely outcomes. This is then combined with the ability to take the steps below in any particular scenario that creates the greatest value.

- Ingesting data at speed and volume sets the stage for additional processing.
- Real-time Analytics processes incoming streams of data from IoT sensors and devices.
- This refined data is then correlated with contextual and historical data to provide a baseline for advanced analytics.

- The next step is to predict failures, anomalies, or patterns using predictive analytics that are based on machine learning over historical and situational data such as external events like weather.
- The final step is to apply prescriptive analytics to determine the next best action to take.

The important point is that specific actions based on a rich understanding of history and context must be taken NOW in order to capture that value. New tools are needed to achieve this ambitious goal for IoT.

V. An Example of a Modern Security Analytics Platform

Through behavior modeling, contextualization, machine learning and reasoning, at big data scale, Guavus' security analytics platform is designed to empower business operations to effectively deliver business outcomes that address IoT business imperatives.

This new platform offers a novel conceptual, machine intelligence approach to analytics and its associated software architecture. It provides 360° visibility across data silos (L3 (network), L7 (application), Threat Intelligence, Data Lakes, Cloud, BYOD and IoT) and opens up data models for threat hunting through its Security Analytics toolkit and modules built ground up for security (on-demand, streaming, real-time).

The focus of the platform is on addressing the challenge of rapidly delivering better business outcomes and value in IoT initiatives and projects. It accomplishes this goal by providing a platform that:

- Delivers faster analytics in real-time with a unique methodology that ingests data (streaming, historical, predictive, and prescriptive) with relevant contextual and situational data to improve the quality of actions that lead to better business outcomes and results.
- Accelerates application development via a set modules and automation that empowers analysts to create faster analytics in minutes vs. months. The platform's faster analytics provide a rapid path to insights and actions that empower organizations to take smarter actions that lead to faster and better business outcomes.

Another important capability is an integrated graph-relational view of identity-asset-network-adversary model which enables persistence, retrieval, search, analytics and visualization across all data sets and data. This powerful capability with pattern and anomaly detection enables analysts to rapidly detect threats which significantly accelerates time-to value for IoT projects. These capabilities empower analysts to be nimble to react to business challenges and create solutions with the platform that enables timely action, implementing complex analytics faster in minutes, not months, and thereby improve business outcomes faster.

By unifying ingestion, all types of analytics, real-time contextual and situational awareness, with behavioral modeling and threat intelligence, Guavus' security analytics platform enables organizations to build applications that will meet the challenges of IoT scenarios. This unification is important not only for performance reasons, but also because the unification between each of the layers requires careful design and engineering to meet the demands of real-time business. Guavus' security analytics platform was built with these imperatives in mind.

Conclusion

For IT security leaders, it's a brave, new world—where it's necessary to step out of the traditional role of compliance and embrace the risk-reward of new IoT business models. This includes building bridges to and skills in the consumer ecosystem and operational technology domains of the organization. Leaders should seek to collaborate with manufacturing and physical security leaders. They should expand more deeply with enterprise risk management, which goes beyond IT systems. Security leaders should also take the lead in raising board-level visibility and protecting the brand. Now, more than ever, information technology is the business, so information security is tied to the brand.

Existing process standards for managing, monitoring, and upgrading should be leveraged. However, this needs to be balanced with a risk/reward approach, not in a universal manner. Details about operational technology will matter in terms of technology and process. Especially in security and risk management, there will be expert shortages, aging assets, and the need for automation and capital discipline. All these should be considered while running the security capability like any other business.

Future-proofing security operations is about getting away from “prevent,” and moving past “detect and respond” to security foresight. The key is to: Focus on people, process, and reporting and close integration points between software and tools.

The futures for security in terms of security intelligence and insight include three areas of focus:

1. Advanced protection platforms: information-centric protections, endpoint activity monitoring and self-healing, advanced forensic capabilities
2. Predictive intelligence: advanced sharing capabilities, scalable threat intelligence vetting, feed-based to adversary-centric intelligence
3. Security analytics: detect the unknown with Big Data analytics, create advanced visualizations, establish proactive, counter-intelligence capabilities—hunt teams.

A new Security Intelligence Platform is needed to address this IoT world that demands rapid implementation time-frames and systems that enable intelligent real-time actions that deliver business outcomes quickly. It offers a careful and intelligent balance of unique and powerful security and Data and Analytics Engines that are the core of a broader and Intelligence platform that will work with a wide range of software and databases in place today. It provides powerful module-driven environment with visualization that accelerates time-to-value for even the most complex IoT applications.

This platform offers much more than just new technical approaches or faster “speeds and feeds.” It is a new kind of platform for business operation managers that accelerates projects through analytics, behavior modeling, contextualization, machine learning and reasoning and delivers better business outcomes faster for IoT initiatives and applications.

To keep your enterprise safe and secure, adopt a proactive approach that enables you to secure your information while improving its flow throughout the enterprise—enabling innovation, improving collaboration, and increasing competitiveness.

Abbreviations

IoT	internet of things
DDOS	distributed denial of service
IT	Information technology

Bibliography & References

Gartner (2017). *Gartner Says 8.4 Billion Connected "Things" Will Be in Use in 2017, Up 31 Percent From 2016*. [Press Release]. Retrieved from <http://www.gartner.com/newsroom/id/3598917>.

Russell, Brian and Van Duren, Drew (2016). *Practical Internet of Things Security*. Birmingham, B3 2PB, UK: Packet Publishing Ltd.

Adapting Proven Technology to Counter IoT Threats

A Technical Paper prepared for SCTE•ISBE by

Petr Peterka

Chief Technology Officer
Verimatrix
6059 Cornerstone Ct. W
San Diego, CA 92121
+1 858-677-7800 x4001
ppeterka@verimatrix.com

Introduction

The Internet of Things (IoT) is emerging fast from hype to reality across homes, enterprises, cities and infrastructures, creating massive opportunities in multiple sectors. But inevitably, given the associated proliferation in IP connected objects and services, it generates new security threats at different levels from minor nuisances to major national security threats. This has created a view that radically different technologies and strategies are needed to counter threats in this new security landscape. Verimatrix challenges this notion by arguing that although some of the threats may appear novel, they involve many techniques around theft of credentials and denial of service that are already well known. As a result, existing technology well proven in other spheres, especially pay TV revenue protection, can be adapted to counter these threats. While new threats are of course arising all the time and require constant vigilance on the part of security providers to counter, the required innovation is already taking place. Security firms such as Verimatrix are investing in AI, Machine Learning and other advanced techniques designed to provide early warning of emerging attacks so that they can be anticipated in advance, or at worst, countered as they unfold before significant damage has been done. Above all, the key to protecting the IoT lies in renewable security which is essential to stay ahead in the arms race against hackers and pirates.

Not Radically New Technology

1. Education Needed for Secure IoT

1.1. IoT Security Falling Behind Expansion

The IoT is expanding so quickly that security is lagging behind, both in deployment and understanding of the risks. There have already been some high-profile attacks, as well as demonstrations of vulnerabilities by security professionals that are potentially even more serious. Such attacks or demonstrations have varied in seriousness, from causing a nuisance or minor economic damage, to major threats to national security, as in the case of the now infamous hack of the Ukrainian power gridⁱ.

Even when vulnerabilities are unearthed by security researchers, the “good guys,” it can prove costly for the manufacturers or service providers involved. This was the case for Chrysler when forced to recall 1.4 million Cherokee Jeepsⁱⁱ after they were hacked in 2015 by two researchers demonstrating a complete remote takeover of the vehicles’ digital control systems. Such cases emphasize problems caused by lack of attention to security during design and development of core IoT components or subsystems, or even the whole infrastructure.

1.2. Public Awareness of Threats Growing

Even the general public has become aware of the threats posed by the growing connectivity between objects, including their cars and devices in their own homes. This has been brought on by several well publicized cases, some of which have implications for personal safety or privacy.

Although many of these cases have involved demonstrations rather than actual hacks, experience tells us that where vulnerabilities exist it is only a matter of time before they get exploited. The potential to cause injury or even death by taking over connected cars has already been demonstrated. On the privacy front, one of the most infamous cases involved toy internet-connected stuffed animals manufactured by

CloudPets, which were hackedⁱⁱⁱ in February 2017, exposing personal information of over 800,000 customers to eavesdropping.

Such cases highlight the need to educate the public clearly over the risks and often relatively straightforward measures that can be taken to guard against these threats. At the same time, there is a need for guidance over where responsibilities lie for DDoS or other large-scale attacks launched by recruiting botnets comprising domestic IoT devices.

1.3. Key Players Often Lack Understanding of Threats and Impact on Value Chain

Although there may be growing awareness of IoT security risks in principle, even providers of core components and services often fail to understand the full ramifications. Just as a joined up IoT opens new horizons for adding value and creating new business opportunities, so it also expands the threat landscape. While key players may have a good grasp of security threats to their own products or service domain, they may not appreciate implications for other IoT domains to which they are now connected. A manufacturer of smart talking dolls might address possible risks to children posed by malfunctions but fail to appreciate that a burglar might take it over to instruct a voice-controlled personal assistant such as Amazon's Alexa to open the front door. The main point then is to consider that the joined up connected nature of the IoT presents opportunities for many sorts of wrong doing across multiple domains. This needs to be communicated in particular to relevant IT departments so that the security nuances of the IoT can be taken into full account during software development. Providers of both components and services also need to be brought on board, given that a major factor making the IoT such an attractive target for hackers is that many devices are shipped with insecure defaults and exploitable code. Furthermore, they are rarely upgraded, usually lacking the capability.

1.4. False Perception that IoT Requires Security Revolution

It might seem natural to assume that because the IoT is a new era for telematics, opening up new vistas for existing and emerging players, it must also require radically new security technologies to counter threats that will arise in this different landscape. This is a serious misapprehension because although the IoT does undoubtedly introduce new contexts and modes of transmission, as well as greatly increased scale and opportunity for attack, the underlying methods are fundamentally the same. These include theft of content, hijacking multiple devices to launch DDoS attacks and injection of malware to disrupt activities, eavesdrop data or launch ransomware attacks.

1.5. Proven Methods Can Be Adapted to Counter Many IoT Threats

Although the consequences have the potential to be felt more widely, the IoT has thus far elicited threats that have been similar to previous attacks to the traditional IT infrastructure. This means that methods and technologies already developed in other spheres, including those that have protected of billions of dollars of revenue in the pay TV industry, can be adapted for the IoT and are already being deployed by Verimatrix and others. The great advantage is that such methods are already mature and well proven so that IoT service providers can deploy them with the confidence that they will work in their environment, provided they have been properly adapted. Section 5 explores how pay TV encryption and key management, authentication, entitlement management and other established processes can be adapted for the IoT.

1.6. Some New Tools Needed but Already Under Development

It is true that some new tools and techniques will be needed to counter emerging threats, but again, these are not unique to the IoT. Just about all telematics sectors face the common challenge of having to monitor for threats, some of which cannot be anticipated in advance, and be able to deal with them as they arise. Section 4 describes how Verimatrix designed its approach to counter threats based on four pillars of IoT security, with a fundamental requirement being that it can be renewed as required, not just to keep pace with the evolving threat landscape, but to stay one step ahead where possible.

1.7. Revenue Protection Vendors Well Placed

Given their experience combating piracy, content theft and various forms of attack, revenue protection providers are already armed with many of the tools and technologies needed to protect the IoT. They have long-standing experience securing the IP set-top box (STB), which was an early example of an Internet-connected thing, and more recently have had to adapt to online content distribution. This has required the extension of protection against content and service theft to many other connected devices, including tablets, smartphones, gaming consoles and cast dongles. Upgradeability has become essential for pay TV security, so providers in that field have become skilled not just at keeping their own software up to date through transmission of regular updates, but also other critical third-party components that can be vulnerable to attack if the latest fixes have not been applied.

IoT Threats

2. Four Threat Levels

2.1. Level 1: Nuisance

There are significant variations in impact even within the category of threats that might be defined just as a nuisance because there is no injury, loss of life or disruption on a large scale. It includes attacks on IoT components such as domestic refrigerators, with potential to cause upset and economic loss to individuals. It also includes threats to confidentiality and personal data which, while not causing physical harm, can still lead to significant distress in the event of identity theft, for example. The scope for such low-level threats will increase as the IoT becomes more inter-connected across domains, which is another reason for taking this category seriously.

At the same time, the IoT is attracting new forms of malware designed specifically to exploit the lack of security to cause malfunctions or deny service. A new malware strain called BrickerBot was detected in March 2017^{iv}, targeting IoT devices by corrupting their storage capability and reconfiguring kernel parameters. This can result in permanent denial of service (PDS) since the devices can be crippled to the point they either need replacing or factory restoration.

2.2. Level 2: Threats to Business and Brand

Attacks on businesses have become more common and larger scale as a result of IoT proliferation. This has increased the scale of DDoS attacks and also made them easier to mount by presenting large numbers of unsecured connected devices. This is a particular threat to smaller enterprises for which the economic or reputational damage could be terminal. The DDoS attack on the news site KrebsOnSecurity^v was such

a case where large numbers of routers and surveillance cameras were recruited, although fortunately that was thwarted by prompt action from CDN vendor Akamai.

Apart from DDoS, the IoT also gives greater scope for malware attacks against businesses, which can be motivated by an individual grudge and are increasingly common for extortion. Ransomware attacks have been encouraged by some large payouts made by firms desperate to restore critical systems in the event they fail to recover compromised systems. South Korean Web host Nayana admitted paying just over \$1 million in Bitcoin^{vi} after being unable to recover data stored on 153 Linux servers and 3,400 customer websites when it had been maliciously encrypted by ransomware attack.

This category overlaps with level 4 because many large-scale attacks, including both DDoS and malware, target multiple companies as well as national infrastructure. The widespread attacks in late June 2017 involving the Petya ransomware^{vii} afflicted both infrastructure and individual enterprises, with victims including the world's largest advertising agency WPP.

2.3. Level 3: Threats to Life or Limb

The third threat level embraces incidents threatening personal injury or death, rather than an enterprise or infrastructure. The connected car is the most obvious target under this category given there are now 112 million vehicles around the world with direct access to the internet, set to more than double by 2025 according to Gartner. There have been no proven attacks against connected cars that have caused injury, but the risks have been demonstrated by researchers under realistic conditions. This has exposed scope not just for targeting individual cars to disable say a braking system, but also for remote commandeering of a large number of vehicles. As fully autonomous driving comes closer, potential for causing serious accidents by taking over a vehicle's electronic control unit (ECU) will increase.

Cars will also be just as susceptible as other IP-connected systems to ransomware from attackers seeking to exploit vulnerabilities in ECUs themselves or infotainment systems to obtain money from the owners. In anticipation of such threats, several industry initiatives have sprung up reaching towards a coordinated approach to IoT security, such as the Automotive Security Review Board (ASRB) launched by Intel, alongside Aeris and Uber, in October 2015. This has staged several workshops in which engineers, cryptographers and security researchers from around the world are collaborating on an Intel and Linux-based in-vehicle infotainment (IVI) simulation platform.

Robotics is another obvious sector where there is potential for causing serious harm through malicious takeover and this field is growing just as fast as the connected car. While the growth is mostly concentrated in the enterprise and particularly manufacturing sector at present, robots are set to enter the domestic realm on a significant scale within the next few years. They too will be IP connected, over Wi-Fi or cellular networks, with optimism over their utility being tempered by fears over security. A recent report^{viii} found that robots were just as prone to hacking as other connected systems and noted that there had already been instances of injury and in one or two cases death caused by malfunctions that also demonstrated the scope for malicious damage.

2.4. Level 4: Threats to National Security and Critical Infrastructure

Threats under this category have naturally aroused greatest concern among governments and security agencies. This partly reflects the great potential scale of disruption but also the fact that several major attacks have already occurred. One positive aspect is that these attacks have galvanized coordinated

responses around the world and ensured that from now on, IoT security will be taken much more seriously by makers of components and providers of services, as well as infrastructure companies.

Just as for Level 2, these large-scale attacks can involve DDoS or various forms of malware, which as we have seen are now being tuned specifically to exploit IoT vulnerabilities. The first large botnets recruited for DDoS attacks involved coopting consumer broadband routers but have come to include surveillance cameras, webcams, digital video recorders, cable TV or other connected set top boxes and, most recently, new types of consumer IoT devices. The threat to critical infrastructure was demonstrated by the DDoS attack on US DNS service provider Dyn^{ix} in October 2016.

Even greater concern caused by malware and DDoS occurred two months later when the Ukrainian national power grid was subject to its second coordinated attack within a year, leading to a widespread two-hour blackout. The attack, orchestrated by multiple groups working together, was more sophisticated than the first known power outage that happened a year prior and resulted in a blackout for 225,000 households in the capital city Kiev. The event wasn't intended to cause serious damage, but it did serve as a training lesson for future attacks.

Architectural View

3. Three Alternative Architectures

The IoT as a whole covers a huge variety of infrastructures, services, use cases and devices, so it is not surprising that there is not just one underlying design architecture. Three alternatives have emerged, the first being the case of IoT devices connected to the wide area infrastructure, or cloud, via some form of intelligent gateway/hub. The second option, sometimes considered a variant on the first, still involves an intermediate unit between end devices and the cloud, but in this case, it is dumb and confined largely to aggregation, routing and protocol conversion. The third option is for devices to be connected directly to the cloud so that they participate as end points in an IP network overlaying local radio protocols. The three options are suited to different situations and vary in the security risks they pose. There is no one-size-fits-all approach to IoT security.

3.1. Device to Gateway to Cloud

Under this model, a centralized hub or gateway sits between the IoT devices and the associated service resident in the cloud. These can be regarded as network edge devices converting between local radio protocols on the client side and IP broadband into the cloud where the service is hosted. The gateway should also provide a connectivity layer locally above the radio protocols, enabling devices to interoperate irrespective of which protocol they support. This makes the service seamless, giving freedom to install devices whatever low power IoT radio protocol they support, whether ZigBee, ZWave, Bluetooth, Wi-Fi HaLow or other.

The gateways may also be capable of running applications to perform local actions that may involve coordination between different IoT devices, but which do not need reference to the cloud. Gateways will play a useful function in partitioning IoT services, filtering data, analytics processes and applications to avoid overloading the host in the cloud.

On the security side, the gateway will, to some extent, insulate devices from the cloud and protect the links on that side. It will also play a role preventing rogue devices from disrupting the wider service, with

the ability to shut them down. Crucially though, the gateway cannot provide end-to-end security by dint of its physical and logical position in the hierarchy. Its position as a gatekeeper with processing capability could make the gateway itself a target of attack itself from the cloud. Indeed, by being an edge device between the internet and the local wireless domain the gateway is a logical point of entry for any threat vector. Therefore, the service may require overlying security above the gateway to ensure end-to-end security at the application level.

3.2. Device to Bridge to Cloud

This model still imposes a form of gateway between devices and the cloud but here it is dumb and so best defined as a bridge, which will be confined largely to protocol conversion and aggregation. This model has emerged for situations in which local intelligence is not required, but when there is a need for operation at longer range than in a typical home. As with the device-to-gateway-to-cloud model described in 3.2 this connects devices by short range radio to an IP end point of the cloud, in this case a dumb bridge. The function of the bridge is to translate protocols that are low bit rate but often longer range than say ZigBee to a higher capacity wide area network. A typical scenario could be in agriculture in which multiple sensors may send data on environmental variables such as temperature or humidity intermittently to a dumb bridge or aggregator up to a few miles away, which, in turn, would forward these into the cloud for processing.

Protocols suited to this model include low-power wide-area network (LPWAN), which in turn is based on the LoRa chirp spread spectrum (CSS) radio modulation technology, optimized for very low power and bit rate but intermediate range.

This bridge model has also been proposed for some forms of home and factory automation using another protocol, 6LoWPAN designed specifically for IPv6 over Low Power Wireless Personal Area Networks. This, in turn, underpins the Thread protocol designed primarily for the home which may become more prominent with ongoing roll out of IPv6 replacing the original IPv4 which has address space that is all but exhausted. The argument here is that IPv6 also brings other benefits, including auto-configuration and end-to-end routing, which eliminate the need for an intelligent gateway. This makes it possible to implement a distributed approach based on dumb bridge devices just performing low level protocol conversion within the home.

E-health is another possible use case for the bridge model when mobile monitoring devices may connect to the service via the user's smartphone. In this case, the smartphone would act as a bridge between diagnostic sensors and the cloud-based center where data would be stored, monitored and analyzed. E-health is also a candidate for the full gateway model residing on a more powerful laptop or tablet, which would then perform the data preprocessing and even potentially diagnosis in less critical cases. Security of the data to ensure confidentiality would of course be critical, calling for full end-to-end tunneling, possibly using the HTTPS (secure protocol) between each sensor and the cloud server through the bridge or gateway.

3.3. Device to Cloud

There will be many instances when the best model will cut out an intermediate gateway and connect IoT devices directly to the cloud. This could be the case for services for which devices are not contained within the range of a static gateway, as in the connected car or container monitoring. In the case of the connected car, there will be some form of gateway, protocol converter or aggregator within the vehicle;

however, this could be treated as an IoT end point from the service perspective, usually communicating over cellular networks.

This model will also be favored for some IoT services within homes or enterprise premises where the direct communication with the cloud could will be via narrowband IoT (NB-IoT). This has been designed primarily for indoor coverage as part of the 3GPP suite of protocols within the LTE spectrum and is capable of running an IP protocol stack. It allows mobile network operators (MNOs) to allocate some of their existing spectrum to these IoT applications.

A key point about this model is that it runs the full IP protocol set end to end. This makes end-to-end security more straightforward to deploy. The service can exploit the security and privacy features already provided by the mobile network, including user confidentiality, device authentication and data integrity.

3.4. Pros and Cons of Three Models

The three models have evolved to suit different IoT services or use cases in terms of mobility, device capability and requirement for local computation or data analysis. Security has not really been considered and must adapt to the architecture as well as the varying levels and nature of the threats.

3.4.1. *Intelligent Gateway*

The intelligent gateway approach has an obvious advantage where there is a need for local decision making and processing of data that could overwhelm both the network and centralized resources if offloaded to the cloud. Such a dedicated IoT gateway can provide extra storage and processing services, allowing the end nodes to be as power efficient and cost-effective as possible. The gateway can also participate in link level security within the local IoT domain.

On the downside, there is uncertainty over optimal design of the gateway, which if dedicated could become an obstacle to rapid IoT innovation, just as legacy STBs can be in pay TV. The gateway can also be a single point of failure, as has already become apparent to users of smartwatches and wearable fitness or health monitors that are paired with the user's smartphone and cannot communicate when that is unavailable. For that reason, various architectures that allow devices to pair with any smartphone or other mobile connected device within range have been proposed, but these bring obvious security concerns, especially for domain services such as E-health where data confidentiality is critical. Dedicated gateways also present targets for attack, by virtue of their computational resources, which can be vulnerable to physical tampering, extraction of private keys, spoofing and even "man in the middle attacks." These can all be countered through strong end-to-end security but deter some service providers from this model.

3.4.2. *"Dumb" Gateway*

The bridge or "dumb gateway" model lends itself more to sensor networks where little more than polling and aggregated data collection are required. It may also be applicable for IoT in the home for monitoring domestic appliances such as freezers, fridges, toasters, kettles and water meters as communication may be intermittent and the level of processing required is small enough to be handled in the device itself or the cloud.

One advantage is that a simple bridge is less of a hostage to fortune than a dedicated gateway, which is partly why hybrid models recruiting smartphones, tablets and other mobile devices as intelligent hubs have been proposed. Such hybrid models can score by providing the processing required for edge

computing as is enabled by the dedicated gateway model, while avoiding single points of failure or reliance on a static device that may not scale well or adapt to future IoT services. The dumb gateway model fails to provide the local intelligence and data filtering that will be essential for many IoT scenarios.

3.4.3. Device-to-Cloud

The direct device-to-cloud model offers the big advantage of running a full IP protocol stack end to end, avoiding need for protocol translation and bringing a richer set of tools at the network level. End-to-end application level security can be deployed readily on top of the stack with less concern over vulnerabilities associated with intermediate gateways. This approach is well suited to applications for which roaming is required without a need for local intelligence beyond the end device itself. It is also applicable to a range of applications that can be served by suitable protocols that work within the mobile spectrum, such as NB-IoT.

Technology View

4. Four Pillars

IoT security should be built around four pillars that cover all aspects and components of IoT services, including devices, data, the service and the network infrastructure. The pillars do not define particular threats because these are constantly evolving and cannot be countered by any specific measures. The point is that the four pillars provide a flexible framework that can be expanded and renewed in the field to keep up with the evolving IoT security landscape and be ready to counter new threats as they emerge. These pillars have not come out of thin air and have their roots in well proven security in pay TV and other sectors. This section examines how each of the pillars maps onto recognized security practices, including the CIA Triad (confidentiality, integrity and availability), not to be confused with the U.S. Central Intelligence Agency sharing the same acronym. Another widely recognized and now venerable model is the IEEE AAA, for authentication, authorization and accounting.

4.1. Device Integrity

The first pillar of IoT security ensures that devices and the software they are executing have not been compromised by any means at any stage in their lifecycle. This corresponds closely with the “I” of the CIA Triad detecting attempts to hijack the device in some way and preventing pirates from succeeding. It requires firstly ensuring integrity of the bootstrap process by which devices or their users obtain key material and configuration information, among other parameters, to allow them to be authenticated for operation within an IoT domain.

Secondly, integrity of the updating process must be assured to avoid devices being subsequently compromised during operation. It is Verimatrix’s contention that integrity of both bootstrap and update can be safeguarded by existing proven mechanisms.

4.2. Device Authentication

The second pillar is essential to protect the wider IoT network or service from intrusion by unauthorized clients or users. This requires assurance that only devices explicitly or directly identifiable are allowed to join a given IoT network. That, in turn, prevents entry of spurious data into the IoT collection network

or

access to systems requiring authorization. This can be achieved by embedding unique authentication keys into protected areas of a chip. However, simple cryptographic solutions will be needed for small IoT devices such as sensors that operate at low energy with minimal computational capabilities.

On the other hand, some IoT devices will be operated by users, in which case authentication may be better associated with the individual concerned, who may have multiple clients accessing a given IoT network. In such situations, there is growing interest in the concept of virtual device authentication with ideas such as transferrable credentials like virtual car keys that can be carried around on mobile phones. The underlying point is that IoT device authentication is important but requires a flexible approach to take account of the highly diverse hardware and use cases involved. It maps naturally to the first A of the IEEE AAA, but goes further than what was envisioned at the time that model was developed to cater for the vast uncharted scope of the IoT. Device integrity and authentication, as well as integrity of communication, contribute to DDoS attack prevention and thus to availability of the overall service.

4.3. Integrity of Communications

Integrity of the communications between devices and the IoT network or hub is the third pillar of security and protects data from interception or alteration during transit. Rather than physically protecting a link, this involves the creation of secure tunnels to avoid eavesdropping or corruption of data. This should also prevent spoofing through falsification of data to masquerade as an authorized device or user.

Since the secure tunnel is enforced by encryption, communications integrity clearly relies on the first two pillars, device integrity and authentication, as well as security within the cloud hosting an IoT service, to be sure that it really does offer end-to-end protection of an IoT data path.

This pillar derives directly from the “C” of the CIA Triad for confidentiality, achieved by encryption, and can be built from proven technologies in pay TV for which integrity is essential to prevent theft of video assets during transmission. In fact, communications integrity and confidentiality have become even more critical for video service providers as they expand into analytics and rely on sensitive customer data for decisions relating to quality of service. As a result, they depend increasingly on their customers’ trust to obtain personal information and this is also becoming a requirement for many IoT-related services.

4.4. Security of Data

Security of the data collected by a connected device is the fourth pillar of IoT security. Similar to the third pillar, the objective is to protect against the corruption or faking of data, along with other possible malfeasances. The difference is that the focus is on the whole data lifecycle rather than just transmission. At this level, policy rules and privacy regulations should be enforced since they are intimately bound up with the data being collected. This pillar relies on the other three to protect against threats to data posed by rogue devices or events during transmission.

This relates to the C of the CIA Triad and also to the “accounting” component of the IEEE AAA because both of these rely on end-to-end security of the data. If the data is compromised at any stage, there can be no guarantee of confidentiality and the veracity of information to generate billing is uncertain.

4.5. Pillars to Extend Entitlement Management

Entitlement management is one vital aspect of security that builds on the four pillars and is applicable in most security domains, including pay TV where it evolved in the context of digital rights management

(DRM,) as well as many parts of the IoT spectrum. It equates to the “authorization” component of the IEEE AAA model by defining precisely what end devices are allowed to do through access control lists.

In the pay TV context, it could mean determining which channels a particular user can watch on a particular device, or which on-demand content can be accessed. The IoT is moving towards a similar model because it is becoming clear that devices on the network cannot generally be trusted and therefore must be restricted in their wider capability. In pay TV, this separation between local and remote access has long been executed in the context of the STB. Users are allowed to view channel guides for example but only the pay TV operator can change that guide’s contents.

Similarly, in the IoT, a surveillance camera can be configured to only be accessed remotely by designated members of the household or, perhaps, to grant temporary access to the fire department during cases of emergency and then revoke that access immediately thereafter. In the case of systems that have potential for serious harm in the event of malicious intervention or takeover, as in the case of autonomous cars or domestic robots, restrictions could be imposed on the actions taken. Cars could be prevented from taking actions that would risk injury to all parties, including occupants of other vehicles and pedestrians. Ramifications of this are discussed further in Section 5.5.

4.6. Renewability and Upgradeability Critical to Counter Emerging Threats

The ability to renew security remotely and almost transparently to the user has become well established across multiple telematics domains, including enterprise data centers, personal computing and pay TV. It is just as essential for the IoT, where clearly security must keep pace with evolving threats without requiring devices such as sensors or light bulbs to be returned to base for upgrades. This raises some new challenges given the very limited processing and storage resources available in many IoT devices and involves matching the renewability process to the use case. Remote environmental sensors do not pose the same threat as connected cars and so obviously do not require the same level of security. In that case, it may be sufficient to keep the aggregating bridge or gateway up to date on their behalf with the focus on the integrity of the data. For most use cases though, it will be imperative that security of end devices can be renewed directly, including electric domestic appliances that can be switched on or off with potentially adverse consequences, at the very least unnecessary consumption of power.

Renewability has played a fundamental role in the arms race against piracy in pay TV since the dawn of digital transmission with the focus on protection of the STB. With IP connectivity, scope for upgradeability has increased, with the aim of ensuring that pirates cannot get at the video content directly by first circumventing the box in which decoding occurred.

Such measures begin with secure boot and other techniques to make sure that when the system starts up, it is loading known authenticated software and not some malware or Trojan invading from the internet. Revenue protection vendors have developed technology for updating the software securely during operation.

Such techniques have also been deployed on PCs and have been carried across to connected devices for secure video delivery. Meanwhile, these techniques have been extended to cater for the fact that these connected devices, unlike STBs, do not have security components like DRM and now watermarking pre-integrated in the factory because they are not built for a single pay TV service. This has been addressed with the help of device makers and their system-on-chip (SoC) providers by deploying trusted execution environments (TEEs), enabling more vulnerable aspects of a pay TV service to be isolated from the underlying operating system and therefore from the apps running on top.

It is becoming possible to replicate the secure software updating, long available for the STB, on mobile devices such as smartphones, via standards for TEEs such as the GlobalPlatform TEE management framework (TMF) and the open trust protocol (OTrP.)

4.7. Extending Connected Security and Renewability to Lower Power IoT Devices

By itself TEE technology does not address the problem of protecting many IoT services because it requires significant processing power not available in small devices like sensors. Such devices will tend to run on small embedded chips like ARM Cortex M cores rather than, for example, the powerful quad-core processors found in many smartphones.

To bridge this gap in processing and storage capability, the OTrP was formed in July 2016. OTrP is really more than a protocol since it extends the security techniques already proven on smartphones and tablets to IoT devices, incorporating the same sort of trusted code management. Verimatrix supports OTrP and believes it provides the foundation for extension of code isolation and secure update to low power and resource-limited devices.

Although this still leaves important issues to resolve, these concern consumer awareness, assigning responsibility for breaches and financing of IoT security in general, rather than the underlying technology. The key point here is that the work now being conducted around the OTrP is leading towards a framework in which IoT can be shielded by the same protection as pay TV services for which there is also a need to satisfy third parties over security, in that case rights holders. Indeed, it is because the experience of pay TV security transfers naturally to IoT that Verimatrix has identified this as an important sector for its technology in years to come.

4.8. Security Options for IoT Protocols

Most IoT services will require end-to-end security at the application level to ensure there are no points of weakness. It is true that many of the components of a given IoT service will have some level of security built in, but this cannot be relied on to provide comprehensive end-to-end protection against all possible threats.

There is also security embedded in some of the IoT wireless protocols themselves, which can play a useful role in protecting the link between components and bridges or gateways as part of the overall solution.

Mapping Pay TV Security to IoT Threats

5. Adaptation of Existing Methods to IoT

The thesis of this paper is that new security challenges for telematics as a whole are often just old ones exploiting different vulnerabilities that arise in given domains. It is true the IoT does introduce some novel threats and uncertainty where the future security landscape looks highly unpredictable. But when the threats are peeled down to reveal the points of vulnerability that allowed them to arise as well as the attack vectors used, they are remarkably similar to those already experienced in the pay TV world, especially more recently as video services have embraced web and IP delivery. So although the future of IoT threats is unpredictable, that holds equally for other domains including pay TV. Indeed, it is precisely

because the future landscape is uncertain that renewable security has been developed firstly for the STB and then other connected devices for viewing video services.

As a result, methods developed for pay TV revenue and service protection can readily be adapted to the IoT. This section explains how the principle threat categories in pay TV map onto the IoT and can be met by adapting proven methods.

5.1. Encryption and Key Management

Encryption and key management have a vital part to play in many IoT domains, with the common theme being protecting data, whether from eavesdropping, deletion, theft or tampering. Data can be of different types, comprising valuable content as in video services, measurements as in environmental monitoring, or control of critical functions as in robotics as well as many other domains. On the surface, threats may look different and yet all involve similar techniques for attacking data beneath the bonnet.

In pay TV, encryption and key management are absolutely essential for protecting valuable video assets against piracy or theft of service, while in the IoT, the focus may be on protecting data during transit, whether from end devices to gateways or within the cloud. There may be a privacy aspect, as in medical applications where a user would not want an unauthorized third-party access to data about personal health, for example. There will also be safety considerations in the eHealth sector, given the potential to take over remote control of pacemakers, insulin pumps and internal units designed to administer a drug at a controlled level.

Another big and fast-growing area common to many sectors is big data analytics, harnessing information from multiple sources, some of which is confidential. Exploiting such data relies on building trust with consumers, whether in pay TV for serving recommendations, or domestic appliance monitoring to gain valuable information about usage. In all such cases, encryption and key management can protect data against interception and unauthorized access.

In pay TV, revenue security vendors are uniquely placed not just to protect analytics data but also obtain it by virtue of their privileged position as custodians of the service. Similarly in the IoT arena, they will be able to assist service providers in this dual role of data protectors and generators.

5.2. Authentication and Protection of Software Integrity

Many attacks on IT systems and more recently IoT devices have involved infiltration with viruses or malware that cause alien functions to be executed that are very different from those for which the system was designed. This has been the cause of most incidents to date involving the IoT, primarily DDoS attacks resulting from recruitment of Botnets in this way. Such attacks are not new, but the IoT, by making unprecedented numbers of devices available, is increasing their scale and potential for disruption greatly. The IoT enables much greater data volumes to be focused against individual targets with the ability to cripple even the web sites of major enterprises for up to several hours in some cases before the attacks are defused.

Techniques already widely deployed, including secure boot, download and upgrade, ensure that sources and software integrity are verified before any execution is allowed. In such a controlled environment, it should be difficult to install any unauthorized software on the devices and therefore hard for viruses or malware to come in and cause chaos.

5.3. Protection Against Cloning Attacks

Cloning attacks are common to pay TV and the IoT, while having different motivations. In pay TV, cloning emerged soon after the introduction of smart cards as the device holding the user's credentials for decrypting authorized content in the STB. By creating a clone that looks like it belongs to a paid subscriber, it was possible to access channels free of charge.

In the IoT world, cloning might have different motives, to mimic a device that enables certain actions to be performed or to give a service the false impression that it is operating normally to disguise malicious actions. A lot of attention has already been paid to cloning of radio frequency identification (RFID) devices widely used for inventory control, object tracking during transportation and security badges for employees to enter work premises. This has brought obvious risks of unauthorized access to buildings as well as theft of valuable items in transit by cloning relevant RFID tags. As a result, various schemes have been proposed to identify and counter RFID cloning.

However, RFID only supports one-way wireless communications and a greater concern for the IoT might be the cloning of systems which support near-field communication (NFC.) This operates at a shorter range than RFID at distances up to just 4 inches but with two-way communications.

It is possible that NFC will become a major medium for configuring IoT devices via smartphones, given that these are already internet-connected and have security built in. By tapping a device, a smartphone could automatically configure a new IoT device via NFC and admit it to the service.

Many applications will require somewhat longer range than NFC but still local communications, such as virtual keys for opening cars and possibly gaining access to buildings as well. This may well use Bluetooth Low Energy for the exchange of credentials, which makes that a possible target for cloning attacks to steal a car for example. There are tools readily available on the web that claim to enable such attacks.

However, such attacks can be countered by methods already well deployed in pay TV where this has been an issue for two decades.

5.4. Extending Entitlement Management to Pay TV

Even when devices have been verified and checks have been made to ensure they are not running any malicious software, it is still possible for them to perform unexpected actions as a result of direct physical access, malfunction or misconfiguration. That is where entitlement management comes in by defining exactly what each device can and cannot do. In pay TV, entitlement management has long been deployed for DRM to restrict access to given channels, often since device, location and time of day. It can also control actions such as storing a piece of content on a personal video recorder (PVR).

Entitlement management can be carried across to the IoT for a wide range of functions, some relatively trivial but still important like ensuring that your own light bulbs but not your neighbor's are connected to your network. It can also bear down on DDoS and other attacks by applying business rules to data flows to and from devices. For example, it can ensure CCTV data is just transmitted to local DVR or cloud video storage and avoid it being sent to unknown web sites. It can also discriminate between devices according to their security capability. It may be that some devices have the full-blown protection of a TEE while others with less computational resource just have an embedded key. In that case, the former device

may be allowed to access external services while the latter is confined to communications within the local IoT domain.

6. What New Technology is Needed

6.1. IoT Increases Uncertainty

As previously distinct and isolated IoT domains become connected, well-defined security boundaries will break down, making it even harder than before to predict how threats will emerge. A threat that might appear to be confined to one domain can affect others. The case of the smart toys discussed in Section 1.2 demonstrated how an IoT device could have unintended consequences by threatening privacy—something the developers probably had not envisaged.

Such cross-domain security risks will extend well beyond privacy and become almost impossible to spot in advance. They have already been quite widely discussed in the context of voice-driven personal assistants like Amazon’s Alexa and Apple’s Siri as they increasingly interact with various IoT domains, including home environmental control and indeed security for operating doors and windows. The prospect of malicious takeover of such assistants to wreak mischief or even perpetrate a crime such as a break in is no longer just speculation but presents real threats that can be demonstrated.

6.2. IoT Security Must Be Proactive and Adaptive

Given this unpredictability, IoT security needs two qualities—renewability and intelligence. The role of renewability in upgrading security to address emerging threats was discussed in Section 4.6, but on its own, it is not sufficient because in some cases the damage will already have been done. It is essential that attacks can be anticipated as they unfold through recognition of associated unusual patterns of activity. This is particularly vital in attacks exploiting unexpected vulnerabilities that have not been covered, when it may be necessary to take some emergency action like temporarily shutting down a particular server or segment of a network.

Machine learning and AI come in by providing the capability to recognize and respond intelligently to unusual patterns. These have already been deployed in pay TV and again can be adapted for the IoT in different contexts. For example, a water meter registering a sudden massive increase in consumption might indicate a leak or a hack. Application of AI might help distinguish between the two by analyzing the precise pattern of consumption data.

6.3. Need for IoT Device Birth Certificate

Apart from cross domain threats, another unintended consequence of the IoT could arise when devices are redeployed or relocated, or even when there is a change in homeownership. In such cases, an IoT device could end up under new ownership, and if it has not undergone a proper factory reset, there is the possibility of unauthorized access to personal or confidential information. In the absence of a full inventory of all IoT devices ever built, which is unlikely to happen, there is a need for some form of birth certificate associated with each to track its lifetime history. This would identify where the device was manufactured and certified, as well as which service providers deployed it during its lifetime and with which customers.

The blockchain system has been suggested as a possible solution given its increasingly wide deployment. Although designed to secure financial transactions and most closely associated with the internet Bitcoin

currency, its operation as a distributed peer-to-peer ledger resistant to data modification makes it ideal for inventory management and for logging the life history of IoT devices.

In fact, these same properties look likely to involve blockchain in many IoT services that require tracking and traceability either of transactions or objects. Real-time tracking to monitor equipment in the field or packages in transit can be implemented over a distributed blockchain network, allied to public key infrastructure (PKI) to facilitate secure information transfer between components. This will enable a range of new applications, some of which are already being trialed.

Blockchain can combine security with distributed operation and time stamping, which will enable the traceability essential for many IoT sectors, including the lifecycle management of devices themselves.

6.4. Responding to Successful Attacks

One golden rule of security is that it cannot succeed in blocking all attacks because there will always be some new or undiscovered points of weakness that can be exploited. Inevitably, there will be some attacks that get through all perimeter defenses and the measure of a good security system lies in how well it can cope with these and minimize damage.

Pay TV revenue protection specialists such as Verimatrix are already deploying machine learning and AI to protect customers' video services and are extending these to the IoT both for proactive monitoring and post-attack response. For example, machine learning can be applied to detect unusual patterns that indicate a potential hack. A well-trained machine learning model should be able to identify connections or nodes where unusual activity or data traffic patterns are occurring and shut those down, while leaving others open so as to minimize overall impact on a service. While at present, machine learning in cyber security is often confined to providing warnings upon which human analysts then act, in time they will take over more and more of the ultimate decision making. This is important for the IoT whose proliferation will expose even more the shortage of skilled human security analysts, as well as the lack of time available to make decisions. In the end, a machine can process information and act much faster than a human, provided it has acquired the contextual intelligence.

Indeed, as monitoring becomes more sophisticated, it will be more likely to pick up attacks early or even sniff them out before they occur. The potential for this was demonstrated by the case of the Cherokee Jeeps hack discussed in Section 1.1. The two security researchers pointed out it had taken months of trial and error before they succeeded in taking over the vehicles. This would have left traces detectable by an intelligent monitoring system.

Conclusion

Two key take home messages can be extracted from this paper.

1) When the application and service specific aspects of the IoT are stripped down, the same fundamental threats and attack vectors common to many telematics domains are revealed, including enterprise data centers and pay TV services. These threats can be countered by security tools and services that are already mature and well proven in these sectors (e.g. secure boot and upgrade, device authentication, secure protocols, etc.).

2) Advanced techniques based on AI and machine learning are being developed by existing security specialists and these will be applicable in the IoT. They will be increasingly capable of detecting attacks either before they occur or very quickly afterwards by identifying unusual patterns of traffic or activity. They will enable much faster diagnosis of such activity, which may result just from a faulty device but could indicate that an attack is unfolding. As these techniques mature, they will become capable of taking over from human security experts in making critical decisions such as shutting down parts of a service in response to attacks. This can save vital time in handling incidents as well as freeing up human experts to take more strategic roles in IoT security.

Abbreviations

ASRB	Automotive Security Review Board
CSS	Chirp spread spectrum
DRM	Digital rights management
ECU	Electronic control unit
IoT	Internet of things
IVI	In-vehicle infotainment
LPWAN	Low-power wide-area network
MNO	Mobile network operators
NB-IoT	Narrowband internet of things
NFC	Near-field communication
OTrP	Open trust protocol
PDS	Permanent denial of service
PKI	Public key infrastructure
PVR	Personal video recorder
RFID	Radio frequency identification
STB	Set-top box
SoC	System-on-chip
TEE	Trusted execution environment
TMF	Trusted execution environment management framework
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

ⁱ <https://www.technologyreview.com/s/603262/ukraines-power-grid-gets-hacked-again-a-worrying-sign-for-infrastructure-attacks/>

ⁱⁱ <https://www.wired.com/2015/07/jeep-hack-chrysler-recalls-1-4m-vehicles-bug-fix/>

ⁱⁱⁱ http://www.huffingtonpost.com/entry/cloudpet-hack-recordings-messages_us_58b4aef0e4b0a8a9b7857b45

^{iv} <https://www.bleepingcomputer.com/news/security/new-malware-intentionally-bricks-iot-devices/>

^v <https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/>

^{vi} http://www.nayana.com/bbs/set_view.php?b_name=notice&w_no=961

^{vii} <http://www.wired.co.uk/article/petya-malware-ransomware-attack-outbreak-june-2017>

^{viii} <https://ioactive.com/pdfs/Hacking-Robots-Before-Skynet.pdf>

^{ix} <http://www.computerweekly.com/news/450401576/Dyn-DDoS-attack-highlights-vulnerability-of-global-internet-infrastructure>

Device Risks to Network Operators from IoT

Exploring the Critical Aspects of Onboarding, Authentication, Authorization and Accountability

A Technical Paper prepared for SCTE•ISBE by

Brian A. Scriber
Principal Architect, Security
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
b.scriber@cablelabs.com

Introduction

The promise of the Internet of Things (IoT) is to develop systems, sensors, and rules to help automate our environments in a way that brings connectivity between people, security and protection to individuals and families, and allows us to bridge gaps in geography. Security in this realm has been widely discussed, often bemoaned and still in need of improvement.

Grouping like concerns around IoT Security results in seven groups of closely related concepts: Identity, Onboarding, Confidentiality, Integrity, Availability, Lifecycle Management, and Future Security. Each of these can be further broken down into subcategories, and each comes with its own particular threats as well as techniques to help secure it.

Content

1. Identity

Device identity is the foundational aspect of the security model for the IoT environment. Without a solid device identity model, spoofing attacks can allow malicious devices to masquerade as trusted entities. When mitigation efforts attempt to find the source of traffic or commands, a device capable of spoofing its identity can simply change to another one, complicating and delaying efforts at neutralizing threats which endanger devices, people in the world those devices operate, or even the network itself.

Device identifiers are often an attribute akin to a serial number which the manufacturer grants to the device, sometimes it is a Media Access Control (MAC) address, and in some cases it is a collection of attribute values from the device which offer some semblance of uniqueness. The problems begin if these addresses are not *unique*. If two devices can share an identifier, then the identifier is useless from a security context. When manufacturers opt to duplicate MAC addresses in products intended for different networks (or even different parts of the world), we lose the promise of uniqueness, and we lose the ability to hold devices accountable for their actions within a network.

When an identifier can be changed, when it can be switched between different values, or when it is based on aggregating attributes which can be modified, we also lose the ability to drive toward accountability within a network. If the device identity is not stored in a protected read-only location, if buffer-overflow attacks can reach this memory, or if the software delivering this address to requesting parties can be corrupted, we have no way to trust the device is what it claims to be or to trust what it will do within or beyond the network. When we see the ease of which MAC addresses can be spoofed, and the difficulty in tracing back to compromised devices participating in botnets, it becomes clear that the *immutability* of the device identifier is critical to security.

Even if a device identifier cannot be modified, and it appears unique on the network, if that device can switch between authentication or authorization credentials, it can still do a great deal of harm and may appear to be a normal device, operating within expected usage definitions, but still harbor malicious code which could be capable of acting on other devices across the network. It is for this reason that we must be able to *algebraically attest* that the device identity is uniquely associated with the credential used for both authentication and authorization.

All devices on a network, whether they can be classified as “IoT” or not, must have identifiers which are simultaneously *unique*, *immutable*, and *attestable*. With all three attributes satisfied, only then can additional trust be built upon the relationship these devices have within the network.

There are privacy concerns that arise from devices that broadcast their IDs prior to network onboarding, and one way to support identity as well as Confidentiality is to use a pre-onboarding technique to generate temporary IDs. Once a device has been onboarded and provisioned within network, it should be using only its unique, immutable and attestable identity.

While there are different ways to satisfy the three requirements of identity, certificates backed by a Public Key Infrastructure (PKI) provide a solution to all three.

2. Onboarding (AAA)

During onboarding and provisioning, a device transitions from a position of no trust within the network to an entity with at least some level of trust. Three aspects of this process include Authentication (confirm the device is what it says it is), Authorization (establish what the device is allowed to do in the network), and Auditing (sometimes referred to as Accountability). During this process, both the device and the network are vulnerable to different types of malfeasance, revoked or expired credentials, network credential exfiltration, and others, but there are specific concerns around Man in the Middle (MitM) attacks where either the device or the network is being misrepresented to the other by an intermediate party. Reducing MitM attack opportunities can be accomplished through using aspects of the Identity section, above.

Identity is obviously a key part of Authentication, proving you are who you say you are and not a MitM imposter. Using a certificate and public key is one way in which the device can prove to the network that it is in possession of a private key. That private key and certificate combination can also be used to validate the chain of Certificate Authorities (CAs) that have permission to grant certificates on behalf of the root authority for the ecosystem. Each of these can be validated algebraically and, using specific asymmetric cryptographic protocols such as a Diffie-Hellman Exchange, cannot be spoofed by a MitM attacker. Additionally, services offered by the CAs include validation not only of the provenance of the certificate being validated, but also confirms the validity period for that certificate and can verify that it has not been revoked since it was issued.

The recommendations for Authentication are to

- 1) Use strong authentication backed by an attestable authority
- 2) Guarantee unique credentials (do not share or rely upon “default” credentials)
- 3) Verify the credentials against a Certificate Revocation List (CRL), against an Online Certificate Status Protocol (OCSP) responder, or against a blockchain registration authority
- 4) Use similar tools/techniques to confirm the issuance of the credentials and match those to the expected device type being onboarded.
- 5) Confirm (again through the CRL or OCSP) that the credentials being used are currently valid.

Ecosystems and like-devices that wish to talk to each other have developed communication protocols and norms for discovery, onboarding, interoperation, and with that they have authorization schemes. These strategies are what grants permissions to devices to access data on other devices, to monitor other devices,

and to send data or commands to other devices. One common way to manage authorization is to create Access Control Lists (ACLs) on devices allowing specific relationships with other devices. Some ecosystems enable not only inbound ACLs, but also outbound ACLs, others use roles and assign roles in a way that enables or disables access control. Ecosystems should

- 1) Ensure their access control is restrictive rather than permissive (default to no access, rely on explicit grants to enable access)
- 2) Protect credentials and valuable resources on the device
- 3) Restrict proxy behavior (device A using device B to talk to device C)
- 4) Severely limit unauthenticated discovery or introspection of devices

Auditing and Accountability are areas which are severely lacking in many IoT devices today, the argument used against it usually revolves around the devices being constrained in terms of storage/power/computation/cost, not having a way to easily interact with a user, or component libraries not supporting auditing. Each of these arguments balance cost and capabilities against the benefits of an audit record, and when systems are compromised, the best way to learn how to protect other systems in the future is to understand the chain of events that occurred during the compromise. While that can help with retrospective analysis, audit records and accounting can also help to recognize deviations from expected behavior to trigger real-time responses a la Intrusion Detection or Prevention Systems (IDS/IPS).

Audit records should be in a standardized format to allow for automated reading and responding to events or actions recorded within the logs. Each action should have an auditable link between itself and the Authentication of who triggered the action and the Authorization of which ACL or role allowed the action to take place. Every log should be immutable; even if it gets overwritten with a new log after a period of time, the log, while it exists, must allow only the event logging mechanism to modify it. In a perfect world, the log would be distributed (to ensure multiple copies in case one is corrupted/compromised), private (encrypted to capture both the privacy as well as the log message integrity), and they should trigger alerts which a human could interact with to identify potential network penetration concerns.

3. Confidentiality

This aspect of security is the part most people think of when IoT Security is discussed. The topic that ends up coming up in those conversations is usually encryption and this tool is hailed as the way to solve whatever security problem is being encountered. Encryption is certainly an important tool for security, and shouldn't be overlooked, but it relies critically on the Identity and Authentication aspects discussed above. Identification of sensitive information is the first step in protecting it, this information could be Protected Health Information (PHI) or Personally Identifiable Information (PII), it could be credentials used by the device, or operating data. Each individual grouping of data may not be PII on its own, but it may become protected in certain regulatory environments if it is stored with other groups of data which in aggregate become PII or PHI. Protecting data can be done in three states of the data: at rest, in use, and in transit. When the data is at rest, it can be encrypted and kept in hardened areas of the device. When important data like private keys or credentials are being used, they should be operated upon from within a protected storage module like a Trusted Platform Module (TPM) where the storage should allow certain operations (e.g. signing a message) to take place within the module, ensuring that the keys never need to

leave their storage where they would be subject to exfiltration attacks. Data in transit should be protected using application-level (also known as “end-to-end”) encryption for traffic.

4. Integrity

Once Identity has assured, Authentication established, Authorization approved, Accountability ensured, and Confidentiality protected, attention is turned toward making sure the device itself is protected and that it performs within the boundaries set forth for it. For critical devices such as those which act as hubs within the home, or those which bridge between an ecosystem and the larger network (or “cloud”), certain elements (listed further below) help to mitigate threats. The threats at this level come from attackers with physical access to the device which “top” the chips, which solder connections and use tools to extract data from device memories. The attacks at this level also come from those trying to upload malicious code into the devices, they could also be scans that look for open ports with known software (libraries or custom code) that have vulnerabilities, scans that look for web servers with default passwords to enable full access to the control states of the device, and more. The recommendation to focus on the critical devices should not be interpreted to mean that all devices are not worthy of protection, but rather assistance in the prioritization of efforts.

The recommendations to protect against these types of advanced attacks are to:

1. Use AAA to confirm that device identifiers, the execution environment, configuration data and communications are all authorized and appropriate.
2. Harden the device by developing a Secure Execution/Executable Environment (SEE), use a TPM, follow the Joint Interpretation Library (JIL) or Federal Information Processing Standard (FIPS) guidance.
3. Minimize the attack surface by closing unnecessary ports and disabling unnecessary services, particularly those services typically used for engineering access, but which are occasionally left installed (and in some cases, enabled) on devices.
4. Use a secure bootloader to ensure that attacks on the state transition, on the software update, or on the configuration data are stymied.
5. Validate configuration data; if this isn’t encrypted, it should at least be signed by an authority that the device trusts (using certificates allows for the root to be persistent in the device trust store).
6. Use non-repudiation for critical communications. Non-repudiation is the act of requesting and receiving a receipt acknowledging a communication sent to another device and having that device sign the receipt with its private key or network credential. The act of returning a signed receipt goes to Auditability and proves that the receiving device did receive the message and cannot claim ignorance of it, or act in such a way as to deny that the message was delivered.

5. Availability

The Availability aspect of IoT Security can be broken into two parts, the availability of the actual device, and then the availability of shared resources (such as the internet) when groups of infected IoT devices participate in botnet attacks. From the perspective of the former, when you go to your garage door and it doesn’t open, it’s unavailable. When you can’t access the status of that garage door from your phone, it’s unavailable. If the power is out in your house or garage, it’s also unavailable. Obviously, there are some

aspects of unavailability that aren't as related to security (e.g. the power being out), but how the device reacts and recovers from them is critical to IoT Security. From the device availability perspective, manufacturers need to expect jamming attacks (frequency flooding, TCP/IP traffic attacks, etc.), there should be a process for how to deal with a loss of power or a loss of network connectivity. From either of these outages, auditing and recovery, notifications to device owners, and potentially re-onboarding, are all recommended. These are processes manufacturers should not only know, not only share, but physically test these as well to confirm that a jamming attack doesn't allow for the front door to be opened without any alerts, alarms, or notifications ever being sent to the owner of the home. Ecosystems need to limit the anonymous requests that can be made, particularly those which are multicast and have the ability to trigger each device to do work (such as introspection or discovery). Finally, every outage must be audited, recorded, and communicated. In a perfect world, any change that occurred from one of these recoveries would be triggered as suspect and notifications sent as appropriate.

The set of devices now known as IoT wasn't always as connected as they are now, they weren't as capable, didn't include as many vulnerable libraries as they do now, and often have default usernames/passwords used as credentials which enable root-level access to the devices. Because of this legacy, Distributed Denial of Service (DDoS) attacks are possible through the use of armies of corrupted IoT devices known as botnets.

As of June 2017, botnets on the dark web could be rented for DDoS attacks at a rate of \$5 for a 1-hour sustained 100Gbps attack and \$10 for a 200Gbps attack¹. With the Mirai botnet capabilities, and the release of the code for exploiting device vulnerabilities to enlist them into a Mirai-type botnet, attacks close to the terabit per second range are no longer just theoretically possible, but we have seen 600-800 Gbps attacks in 2016.

To help mitigate against botnets (as opposed to mitigation of DDoS, not covered in this paper), there are a few recommendations. First, use restrictive (rather than permissive) access control and default network traffic. Second, monitor for inappropriate/unusual traffic patterns (e.g. if the lightbulb suddenly wants to send tremendous amounts of data to a site online with questionable bona fides. A third option is to segment internal traffic and devices into subnets and manage those, keeping an eye on boundary traffic and confirming that it flows within expected norms.

6. Lifecycle Management

Manufacturers of short-lived devices (cell phones can have a 1-2 year support period) that are moving into IoT devices with differing support periods are encountering challenges to their business models. When a consumer purchases a lightbulb there may be different durability expectations than when a consumer makes a purchase of a refrigerator. When both of these devices are "smart" devices, when they connect to the home network as well as the internet, they require support because security isn't a fire-and-forget industry. Security engineering is different than some other types of engineering; security engineers aren't concerned exclusively with weather, stresses, friction and time, they face intelligent, active, adversaries who are constantly searching for weaknesses. With such adversaries, security engineers need to adjust practices (encryption, algorithms, hardening, identity, protocols) to ensure continued confidentiality and integrity of the devices they work to protect.

A "smart" light bulb is different from a standard (classic?) light bulb in that it has computational power, memory, a radio, and, importantly, it has network credentials. It also has an operating system, drivers,

firmware, and a networking stack which all rely on libraries that have weaknesses. Those libraries may have been updated against new threats and improved against unknown threats, but the likelihood that the firmware in the light bulb has been updated is small. For many devices, updates aren't even available from the manufacturer (who may not be incented to spend time or money on supporting a product after sale). If they are available, there may be only complicated methods to install the updates – some involving the use of a soldering iron. With knowledge that nearly half of the homes in the United States do not have working batteries replaced in their smoke detectorsⁱⁱ, the idea that these homes are going to update the firmware in their light bulbs seems to be a stretch.

From the lifecycle management perspective, consumers should expect the following technical elements from their devices:

1. Secure, standardized, automated updates
2. Devices that have clearly defined end of life functionality (such as the light working from within the local network, but perhaps after the support period ends the light can no longer be activated over the internet).
3. Devices should allow for credential renewal and revocation

From a procedural perspective of lifecycle management, we have inconsistent support for disclosure of:

1. Vulnerabilities and remedies
2. The support period to consumers prior to purchase
3. Functionality of the device available after the support period ends

7. Upgradeable Security

Regardless of the security algorithms or encryption schemes shipped with the IoT device, adversaries may find weaknesses in embedded libraries or in the operating systems before the product even hits the shelves. A method for secure software update is necessary, but the device itself must be designed to support the update procedure and be prepared to address the changes. If longer key lengths are required in the future, there must be enough secure storage to retain those keys and if algorithms need to change, devices must have the capabilities to have cryptographic libraries updated. Adversaries may also change over time, as could their motivations, the security principles (or principals, for that matter) that protected devices when they were designed may not be the same as when the device is used in practice.

When considering security for devices, it's easy to fall into the trap of thinking that some devices are in need of less security than others, but the manner in which attacks take place is often to target the weakest part of the perimeter of a network, and continue using credentials of a weakly protected device to extend the attack. No device with computational power, memory, a radio, and network credentials should be weakly protected, and the credentials should be protected behind a hardened hardware element such as a TPM if possible.

Category	Description	Goal for Every IoT Device
Device Identity	Require an attestable, immutable, and unique identifier for each device.	Use a secure certificate-based approach with centralized management (PKI) to provide an attestable, immutable, and unique identifiers for each device; certificate issuance; lifecycle management; and revocation.
AAA/Onboarding	Require the use of strong Authentication, Authorization, and Accountability (AAA) methods for provisioning and management of devices.	Use the established device identity to enforce strong authentication (identifying the user or device), no shared default credentials; use authorization enforced through established mechanisms to provide access network and device resources, using techniques such as ACLs, RBAC, etc.; incorporate accountability mechanisms that associate device actions with authentication and authorization.
Confidentiality	Protect data so that they are not made available or disclosed to unauthorized individuals, entities, devices, or processes.	Identify sensitive information (PHI, PII, credentials, etc.) and protect that data appropriately. Use of encryption for locally stored sensitive information. Provide protection of sensitive information (e.g., private keys) while in use. Use mutual end-point authentication and application-level encryption (end-to-end) for sensitive data in transit. Provide an ephemeral identifier for anonymous discovery requests and limit information available to anonymous introspection and reflection requests.
Integrity	Assure the device is trustworthy and the processes, data, and communications associated with the device are accurate.	Confirm that the device identities, execution environment, configuration, and communications are authorized and appropriate using the AAA methods. Harden the device to minimize the attack surface by closing unnecessary ports, disabling unnecessary services, and using a secure bootloader with configuration validation. Consider use of non-repudiation methods for critical communications.
Availability	Safeguard devices and associated communications for proper functioning.	Use restrictive, rather than permissive, default network traffic policies to limit communications to expected norms, guarding against both unintended as well as malicious denial of service attacks. Plan for appropriate device behavior in the event of network or radio communication failures, overloads, or outages (e.g., jamming).
Lifecycle Management	Support sufficient secure operation, update, and communications throughout the life of the device.	Provide for secure, automated, update mechanisms during the disclosed support period and publicly disclose vulnerability remedies, EOL functionality changes, and credential renewal and revocation.
Upgradeable Security	Plan for security improvements required to support equivalent device and network security in concert with Lifecycle Management.	Include support for longer key lengths, stronger cryptographic algorithms/cipher suites, and hardware based security over the supported life of the device.

IoT Security Categories

Conclusion

The seven categories of IoT Security are intended to help guide thoughts around protection of devices and the networks upon which they operate. The following table provides an overview of these principles and the key elements of each. These are intended to be a good place to start and not a comprehensive list of every security vulnerability that could ever arise. They are also intended as a second step on the security

front, after engineering portals are eliminated from devices, after default passwords have been removed, and after open ports allowing tools like telnet to connect to devices have been closed.

Abbreviations

AAA	Authentication, Authorization and Accountability
ACL	Access Control Lists
CA	Certificate Authority
CRL	Certificate Revocation List
DDoS	Distributed Denial of Service (attack)
FIPS	Federal Information Processing Standard
IDS	Intrusion Detection Systems
IoT	Internet of Things
IPS	Intrusion Protection Systems
ISBE	International Society of Broadband Experts
JIL	Joint Interpretation Library
MAC	Media Access Control
MitM	Man in the Middle (attacks)
OCSP	Online Certificate Status Protocol
PHI	Protected Health Information
PII	Personally Identifiable Information
PKI	Public Key Infrastructure
SCTE	Society of Cable Telecommunications Engineers
SEE	Secure Execution Environment
TPM	Trust Platform Module

Bibliography & References

ⁱ CableLabs security research on the Tor-addressable network known colloquially as the “dark net”, 2016 and 2017.

ⁱⁱ <http://www.nfpa.org/news-and-research/fire-statistics-and-reports/fire-statistics/fire-safety-equipment/smoke-alarms-in-us-home-fires>

Security of Open Distributed Architectures

Yet Another SDN and NFV Security Paper

A Technical Paper prepared for SCTE•ISBE by

Steve Goeringer
Principal Architect
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
s.goeringer@cablelabs.com

Dr. Indrajit Ray
Professor
Computer Science Department, Colorado State University
Fort Collins, CO 80528
Indrajit.Ray@colostate.edu

Introduction

Our networks are continuously evolving and there are lots of neat ideas out there on how this evolution should be done. There are lots of ideas including software defined networking (SDN), network functions virtualization (NFV), development operations (devops), containers, virtual machines, just to name a few. All of these complex ideas are ultimately about automating the deployment and delivery of network services over open distributed architectures. Network operators are avidly reinventing themselves using open distributed architectures to reduce cost, accelerate service velocity, and enable new types of services. However, these technologies also present new security challenges — our traditional ways of addressing network security must evolve. This paper briefly reviews the unique security challenges and opportunities network operators face as they apply these technologies. After a brief introduction, the paper discusses some of the emerging challenges. These are organized into two parts: larger and obfuscated attack service; and new risks. This is followed by discussion of the opportunities operators can leverage to use these technologies to improve security. The paper concludes with a review of proactive actions currently available to operators.

Security of Open Distributed Architectures

1. Frame of Reference

Software defined networking, network function virtualization, virtual machines, continuous integration, containers, development operations... There are so many terms and buzzwords used to describe the ideas being applied to evolve networks today. In aggregate, these technologies create open and distributed network architectures that support services programmatically. The result is complete reinvention of our networks where now we architect network factories that churn out services at an unprecedented rate.

The technologies network operators use to achieve open distributed architectures are complimentary. First, we started with programmatically controlling the information flows of our networks by applying software defined networking (SDN) [ONF 1]. SDN used controllers as an interface to application logic to programmatically implement flows between switching elements. The most common SDN implementation uses OpenFlow® [ONF 2]. Then, we started to deploy and manage entire virtual resources programmatically using network functions virtualization (NFV). Through NFV, we programmatically orchestrate deployment of virtual network functions (VNFs) which comprise of one or more workloads distributed across multiple general purpose servers [ETSI NFV 1]. The common thread between SDN and NFV, synergistic technologies, is automation of network engineering and operations on an industrial scale, completely transforming the way network operators fulfill their customer's needs. An emerging technology, network slicing, promises to go one step further – to beginning orchestration of deployment and maintenance of entire virtual overlay networks [NGMN 5G].

It's important to recognize the IT revolution behind SDN, NFV, and slicing. Agile engineering, continuous integration, development operations, virtual machines, containers, and other evolutions in computer science disciplines and tools fuel the evolution of *open distributed architectures*. In synthesis, this all combines to create entire, automated network supply chains, binding operators and their technologies to vendors and software companies closer than ever realized in the past. This paper will use the phrase open distributed architectures to reflect the whole technologies that are emerging to allow our networks to be more adaptive and agile in programmatic ways. The focus of this paper is discussion of gaps in securing developing open distributed architectures. Certainly, new and emerging architectures

must represent new risks – the fabric of developing networks is fundamentally different than that of our legacy “big iron switches and routers”. So, it must follow they have different vulnerabilities.

2. What is security?

First, let’s readdress why we need to address security at all and then reaffirm what security goals must be, at least from the perspective of a service provider. This may seem obvious, but perhaps there are important subtleties that are often overlooked.

Network operators and their supply chain partners work very hard to engineer very specific experiences for their users – experiences that are repeatable and provide unique value to their users. The infrastructure and software used to provide great experiences has inherent value. Usually, the components upon which services are built are general purpose and can be used to support many kinds of service functions. And, of course, they must be interconnected to other components and ultimately to the end users. Therefore, it is axiomatic that network resources have value and, to at least some degree, those resources must be exposed to be usable. Since these resources have value and must be exposed, others – our adversaries – seek to harness the value of network resources to provide services other than intended.

This is a fundamental subtlety that seems largely overlooked. When hackers or other cyber criminals attack and compromise a network resource, they change the experience of end users. Often, this may be done in a way that end users may not even realize. Sometimes, the impacts are quite apparent – the services may not work, users’ confidentiality or privacy are compromised, or worse. On a macro scale, network hacking really hijacks, or at least subverts, network operators’ supply chain – it redirects the entire value for which a service was intended to benefit another part.

What then should the goal of the security engineer be? Certainly, it’s not to implement perfect security. Perfect security – making it impossible to misuse any resource by any party that is not authorized – is impractical (and probably impossible). To have value, networked resources must be exposed. The less exposed they are, the harder they tend to be utilized. Rather, the goal should be to make using a networked resource expensive to misuse, while being inexpensive to use by authorized users for the intended purposes. This orientation recasts the focus of security on managing the total life cycle cost of delivering a targeted experience.

3. Security challenges of open distributed architectures

The security of open distributed architectures is challenging on several fronts. There are many, many stakeholders – customers, service providers, vendors, and in some cases regulators and other government agencies (such as law enforcement who require support for lawful intercept). Even within a service provider, multiple organizations and even business units may be engaged. And, of course, there is lots of complexity which is easily seen in papers from the Open Networking Foundation (ONF) [ONF Primer] and the European Technical Standards Institute (ETSI) [ETSI NFV 1]. Multiple stakeholders responsible for lots of complexity create an environment with a much larger attack surface than legacy networks possessed.

However, the programmatic nature of evolving networks must also be considered. The attack surface can be obfuscated because SDN and NFV networks enable a great deal of abstraction. Abstraction makes complex things appear less complicated. One example of abstraction is the use of application programming interfaces (APIs). When the application programming interfaces (APIs) used to access

layers of abstraction are well implemented, the APIs can significantly improve security. However, they can also hide security problems and make monitoring the security of running processes hard to measure. Moreover, these network tools can be applied recursively. In SDN, for example, you may have an SDN controller of multiple SDN controllers which may in turn control other SDN controllers. In NFV, VNFs can nest VNFs which in turn may be distributed across multiple containers or virtual machines. The result is that security vulnerabilities may be recursively exposed.

The programmatic nature of these network technologies is driven by various forms of data models. How this works varies per implementation, but there may be model driven service abstractions, templates, scripts, configuration files, etc. Consider that now our networks can be hacked simply by hacking the models used to create them.

A major differentiator between data center oriented and network virtualization is how service state is distributed across so many elements. At a minimum, service may involve a cable modem, a set top box (which may have a cable modem integrated into it for non-linear video requiring broadband functions), an access point (particularly for shared or public WiFi), the cable modem termination systems (CMTS), a variety of service and management servers (for configuration, software download, address grants, etc...), and more at the core level. Often, the nature of service state and management across such a complex access architecture is forgotten. Transitioning this architecture to a virtualized model and ensuring predictable and manageable service will be challenging.

SDN poses specific challenges [ONF 3]. At a minimum, the attack surface is increased by the introduction of the SDN controller. The server and application software implementing the controller present new vulnerabilities. For example, the Southbound OpenFlow® and NETCONF [RFC 7803] interfaces and the Northbound RESTful APIs are required to enable meaningful services. Moreover, the model driven service adaptation level (MD-SAL) introduces new types of data elements (files and entries in files) [ODL]. In addition, legacy management and operations interfaces must still be supported, along with any proprietary interfaces. Control plane and data plane separation are changed, and partial or full data plane packets may be sent to the SDN controller for analysis and processing to create new flow table entries.

NFV also presents specific challenges [ETSI NFV 2]. At the most fundamental level, NFV transitions functions that used to be achieved using dedicated hardware to software running on general purpose computers (at least, that's the theory). This means that our service infrastructure is now software based. Consequently, software vulnerabilities become infrastructure vulnerabilities. Also, NFV is very complicated – it introduces lots of new functions and approaches to the way we implement network services. The attack surface, like with SDN, becomes larger relative to legacy network implementations as we introduce additional controllers, orchestrators, hypervisors, virtual machines or containers, application stores (and more). Management and control of this complicated architecture requires lots of layers of abstraction, which can obfuscate security vulnerabilities. Moreover, these abstracted vulnerabilities can cascade through the entire architecture as, for example, a compromised template for virtualized network function (VNF) can create vulnerabilities in forwarding graphs which in turn may introduce vulnerabilities to additional VNFs that are part of service chains [Lee].

Unfortunately, focusing solely on SDN and NFV really misses the point. The point of network evolution seems to be focused in automation. So, open distributed architectures also incorporate IT technologies for agile service development and deployment including continuous integration and deployment, live builds, development operations, and more. The result is a new network infrastructure which allows operators to rapidly integrate services as a highly automated supply chain. This supply chain bridges both

development and production networks. This means now we have to secure software development and integration components as part of our production network for service delivery. This includes our tools sets for bug and feature tracking, software build environments, code integration, package management, configuration, testing, and more.

4. Some critical needs

Given the evolutionary path we are on, what must we achieve in terms of security? Our goal is simply to ensure all things done with future networks are done securely. This sounds trite, of course, but that is the goal. The challenge is the scope is audacious and the depth is daunting. In terms of scope, we want to ensure all activities related to developing and integrating our virtual network components and the infrastructure on which those components will be deployed is secure. The images, models, and data repositories used to store configurations must be secure and the loading of those resources into runtime must also be secure. Then the run time of virtual network functions themselves must also be secure. Again, the scope is audacious (see Figure 1).

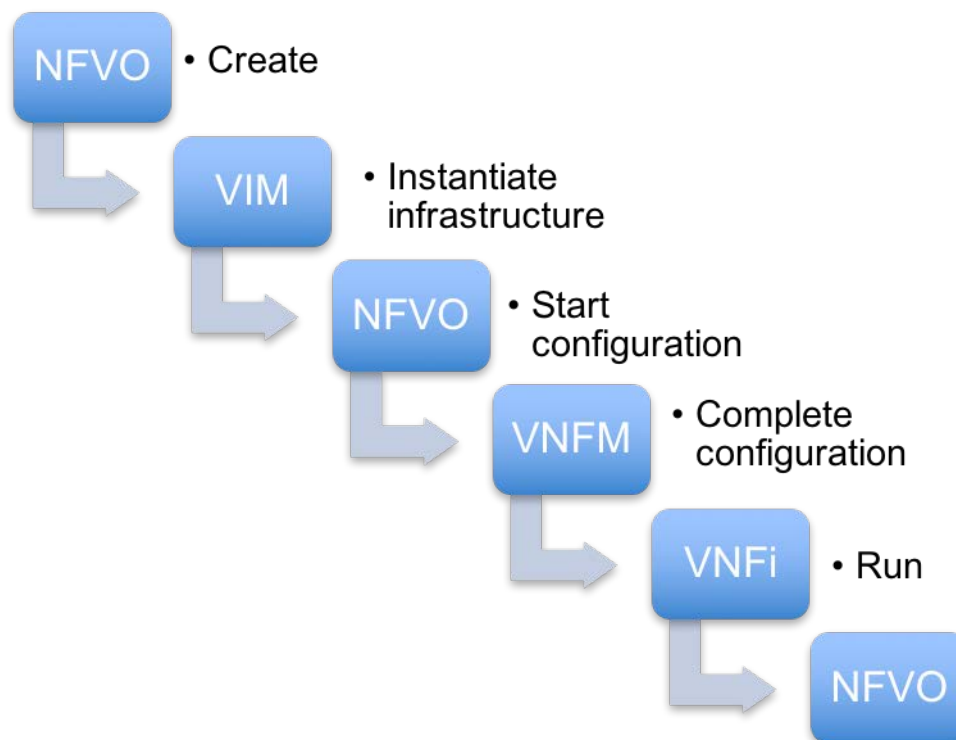


Figure 1 - VNF development and deployment process

Of course, “everything” involves significant detail. Hardware must be discovered and identified, as must software elements. The hardware must be jumpstarted and made available as infrastructure components, whether as part of a pod, cluster, cloud, or whatever architectural concept is used. The VNFs themselves must be developed and once well integrated and understood, stored in reference implementations (gold configuration) or images in catalogues. When customer demand justifies a new network element, onboarding a VNF must occur, bootstrapping the new VNF from the catalogue or image repository, configuring it for the specific need at hand, and then integrating into the actual routing and switching

flows of the current running network. Once the VNF is up and accepted, life cycle management processes such as maintenance, migration, restarting, and recovering must be supported. Once the VNF is no longer needed, it must be terminated (turned down). And all of this may need to have corresponding customer facing activities through operational and business support systems. And it must all be done in a secure way, including strong attestation (which ultimately means the run time result is provably secure).

In efforts to pursue creation of an end-to-end proof-of-concept of this entire vision have been challenging. When considering OpenStack, OPNFV, ONAP, and FD.io as possible solutions, much of their security focus has been on vulnerability management of the code base themselves. This is not to undervalue the good work their various security working groups have labored to achieve. They have worked hard to provide a good basis for security. However, the result is largely flows from the Linux legacy on which these solutions are built. They're all Linux based. Consequently, security settings and implementation options are distributed as flags and fields in lots of files supporting dozens – sometimes hundreds – of processes that implement any given VNF. Encryption and authentication keys are often stored unencrypted in these files in addition to various access control indicators and other security controls. While OpenStack™ and OPNFV do have excellent security guides, they are insufficient as hardening guides. They advise what security features exist, but don't advise how to configure security against specific threats. Moreover, no configuration and security validation tools are optimized for validating security of SDN or NFV infrastructure. Never-the-less, a diligent and persistent security engineer probably can reasonably lock down SDN and NFV infrastructure and applications – it's just going to take a great deal of effort and incremental validation.

There are three critical areas where open distributed architectures are falling critically short. When SDN, NFV, and slicing are combined with continuous integration and development operations, the result is a dynamic supply chain. Supply chains are networks in and of themselves. We tend to think of them in terms of trucks and warehouses, but in the information age, they use telecommunications networks and servers to develop and deliver digital goods. These digital supply chains reach into both our development and production networks. Consider MANO, OpenDaylight and the model driven service delivery it considers, our devops tools like Puppet and Juju and dozens of others, SDN and NFV application stores, and VNF catalogues. These resources reach into our networks, but few security engineers have fully considered how cyber supply chain risk management must be approached differently than traditional network security risk management. How do we patch this kind of architecture? How do we assert a consistent security policy across this environment? Code validation and automated network security management must be integrated. Also, open distributed architectures should apply a zero trust model for all network functions.

Consider the ETSI NFV reference model illustrated in Figure 2. The model shows how virtualized network functions will overlay an infrastructure layer and how all the various NFV elements must use a wide range of reference point connections for inter function communications. This includes connectivity to OSS/BSS support and Management and Orchestration functions. Three types reference points are shown – execution, other, and main NFV reference points. Only the NFV reference points are specified, with execution and other reference points waiting to be specified or judged out of scope. Even, so, the result is extremely complicated (see Table). This does not even include other interfaces to support connectivity SDN controllers, the VNF catalogue, instrumentation and telemetry services, container stores, devops application, and more. How can these multiple connections be sufficiently specified and architected to support a zero trust architecture? Simplification.

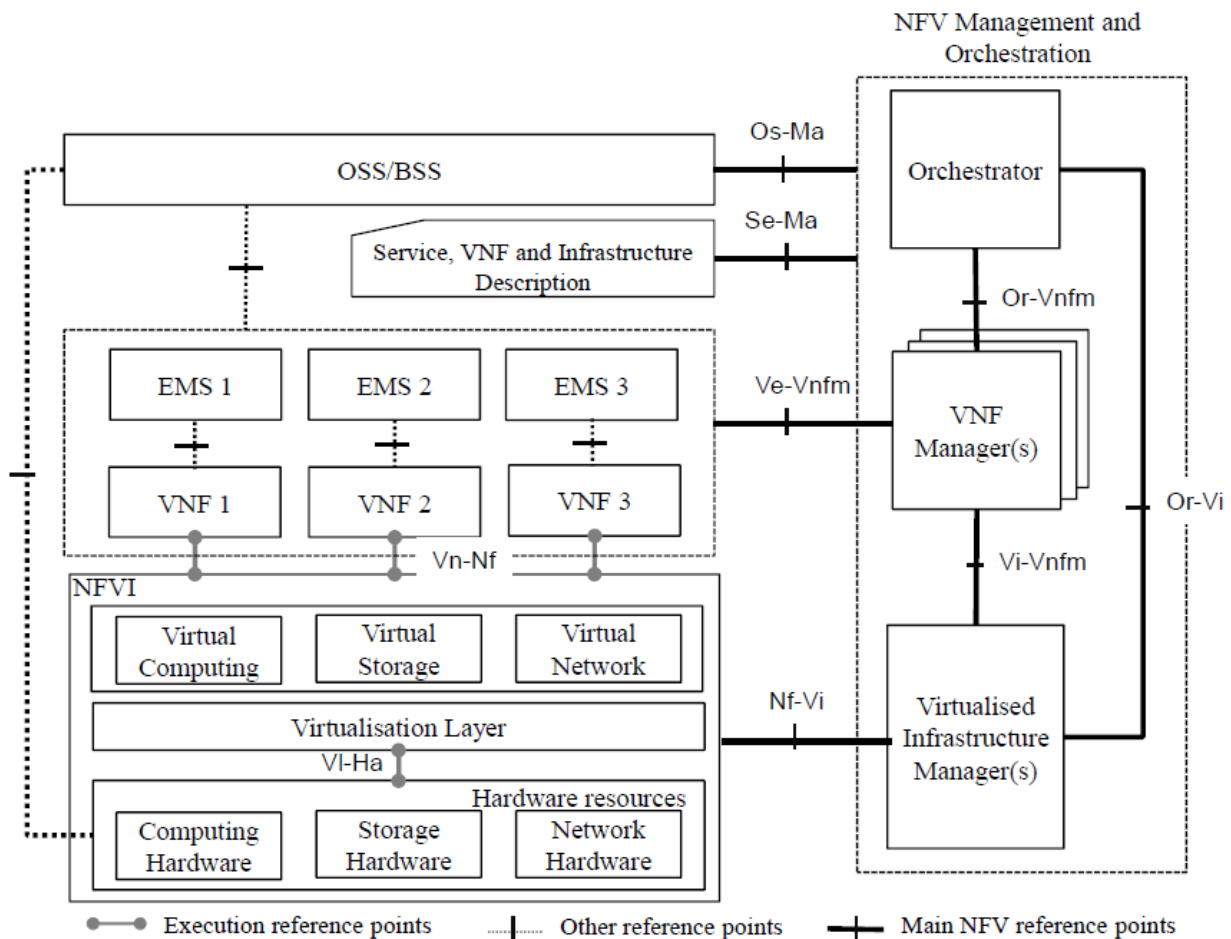


Figure 2 - ETSI NFV reference architectural framework

Three types reference points are shown in the ETSI NFV reference architecture framework – execution, other, and main NFV reference points. Only the NFV reference points are specified, with execution and other reference points waiting to be specified or judged out of scope. Even, so, the result is extremely complicated (see

Table 1). This does not even include other interfaces to support connectivity SDN controllers, the VNF catalogue, instrumentation and telemetry services, container stores, devops application, and more. Ensuring secure implementation of all these interfaces can be difficult.

Table 1 - ETSI reference points

Reference point classification	Reference point	Terminating entities	
Main NFV reference points	Os-Ma	MANO	OSS/BSS
	Ve-Vnfm	VMF-Manager	EM or VNF
	Nf-Vi	VIM	NFVI
	Or-Vi	VIM	NFV Orchestrator
	Vi-Vnfm	VIM	VNF-Manager
	Or-Vnfm	VNF-Manager	NFV Orchestrator
Execution reference points	Vi-Ha	Hardware resources	Virtualisation layer
	Vn-Nf	VNF	Network Function
Other reference points	Not specified	EM	VNF
	Not specified	OSS/BSS	EM/VNF
	Not specified	OSS/BSS	HW resources

Open distributed architectures need to focus on establishing security associations between functions and devices rather than simply interfaces (as reflected by a standard reference point). Such a security association must be:

- Based on strong identity: Identity is the basis for any meaningful trust system. Identity should be based on a secret paired with a unique identifier. The identity must be attested by a certificate or equivalent by signing or equivalent cryptographic operation. The certificate may contain other information (but not any information that should be changed such as software versions).

- **Authenticated:** Each security association must be verified when the association is requested using a cryptographic challenge.
- **Authorized:** Once entities have validated their mutual identities, their resource or activity accesses must still be authorized. Authorization should be based on a system or service wide policy system. The policy system should assume a least privilege orientation and assure separation of duty and function. Implementation may use a policy lookup or token grant approach.
- **Isolated:** Isolation of network, storage, and compute resources used for specific workloads must be assured. There are a wide range of obvious security risks that are managed this way, however, it is equally important from a performance perspective. Specifically, workloads or process should not impact other workloads or processes unless allowed by the operator. Isolation may be achieved by network segmentation (through secure addressing or encapsulation) and various virtualization tools for ensuring workload isolation in memory, CPU, and storage.
- **Confidentiality:** Data and communications should be kept private. The isolation functions discussed above may achieve sufficient confidentiality. However, encryption will ensure even stronger confidentiality, assuming adequate protection of encryption keys.
- **Attested:** Finally, all the security controls that implement a security association and protect it must be provably untampered. This is traditionally done using accounting and logging mechanisms. There are improvements in trusted computing systems that allow secure boot and run time monitoring to improve on legacy approaches. Whatever specific strategies are used, the goal must be to verify that the infrastructure and the security associations implemented to interconnect both hardware and software components are, indeed, what they are expected to be.

These ideas are illustrated in Figure 3. When implemented as a whole, they provide a secure communication channel (security association) between elements. The security association itself is attestable and can be managed according to network wide policies. Every interface in an open distributed architecture – whether link layer (physical) or network layer (logical) should be implemented as a security association.

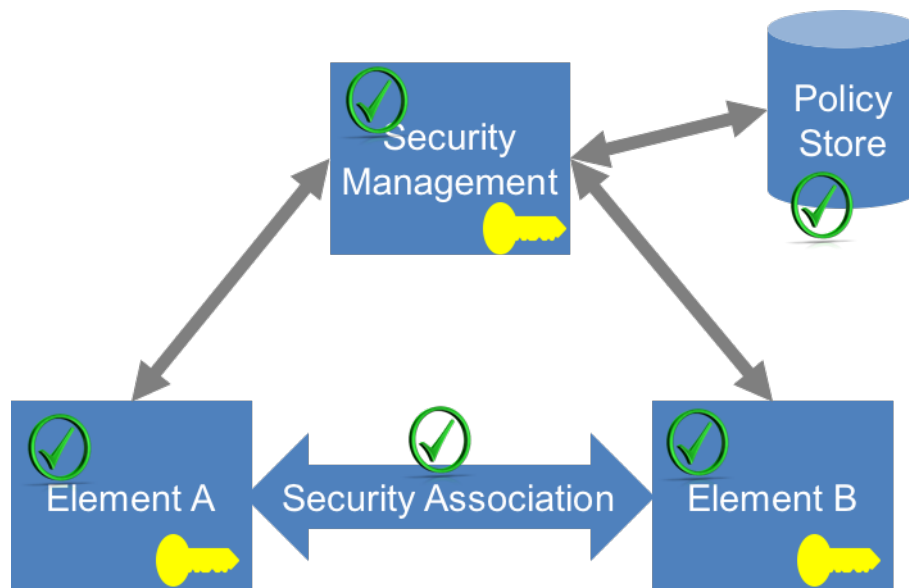


Figure 3 - Security Association

The characteristics outlined above are left intentionally non-specific. Details will depend on whether the trust system uses PKI, symmetric keys, or tokens. Further details may depend on how secrets are issued and stored (such as through use of a Trusted Platform Module or Hardware Security Module).

Applying identity to evolving technologies is proving challenging. Security professionals working on IoT, connected healthcare, and virtualized network technologies are all finding that establishing a strong, evolvable and scalable identification mechanism suitable as a basis for ecosystem wide trust is challenging. A secure identity is comprised by at least three elements. First is a secret known to the entity against which the identity is associated. This is most often a private key. This secret is the basis for trust in the ecosystem in which the entity participates. In asymmetric cryptographic trust systems, the private key will have an associated public key and the private key should not be known to any other element in the system. In symmetric cryptography trust systems, the key must be known by at least one other element (such as an authentication center). This secret must also be immutably bound to a unique identifier (within the scope of the ecosystem). This immutable binding may be achieved by digitally signing the identifier and other attributes and information using a public key certificate [X509][RFC 4158] or an external validator such as an authentication center.

One of the reasons identity is challenging is that to be most useful identities must be persistent. Persistence, however, comes at a price. It usually means identity must be issued by some central authority, which can create friction in the supply chain that supports distribution of the devices or software elements being identified. The central authority can also represent a risk factor in cyber security and supply chain terms. Moreover, there must be some way of revoking persistent identities, which requires processes and resources that ultimately are overhead. Multiple schemes are available to fluidly and dynamically assign identity, but these are problematic as they increase the attack surface of the ecosystem the identity supports. Dynamic identities also are hard to associate to security policies. Token

based systems attempt to mitigate this deficiency. However, while many token systems are reasonable authorization schemes, they are insufficient for authentication. Tokens must be validated as authentic before they are used to authorize actions or use of resources. Otherwise, tokens can be misused by third parties. A superior identity system suitable for ecosystem trust might require a combination of public key (asymmetric), symmetric (such as specified by 3GPP/GSMA™), and token based technologies. The challenge is implementing such a hybrid with minimal complexity – otherwise, the hybrid will likely increase the attack surface of the ecosystem without sufficient gain in security functionality.

5. Opportunity

The scope and depth of new threats and attack vectors related to open distributed architectures should be neither exaggerated nor understated. There is both risk and opportunity. There is the opportunity to improve security relative to legacy infrastructure, mostly as a consequence of the programmatic nature of emerging networks. The improved potential of automation provides the opportunity for more consistent execution of security processes and controls. Moreover, automation and more consistent execution of interfaces and management protocols may make it easier to upgrade and patch the network as security threats evolve and develop.

Both SDN and NFV also promise specific advantages. As NFV is based largely on general purpose hardware, it may be much easier to enable cryptographic functions. Pervasive support for encryption may be possible much more cost effectively than in the past. Also, using general purpose hardware may provide improved support for consistent and well implemented authentication. Finally, both SDN and NFV promise to support very granular control of packet monitoring and again more cost effectively than in the past.

Conclusion

This paper has reviewed some issues to realizing secure virtual infrastructure and software based services. While there are challenges, there are also opportunities that promise a future where our networks are easier to secure. However, we need to address at least three gaps. We need a consistent approach to identity management, starting with a method of asserting identity in virtualized infrastructure. We need a similarly consistent way to secure the myriad of interfaces our future networks require. This paper advocates applying the notion of security associations in a consistent way to all interfaces. Finally, we need to consider how emerging distributed architectures allow service providers to implement supply chains that turn out networks as a service. Our operational security practices need to extend into development and integration and even to the open source organizations and vendors supply both hardware and software used in our infrastructure.

Abbreviations

API	Application programming interface
CMTS	Cable modem termination systems
devops	Development operations
ETSI	European Technical Standards Institute
HSM	Hardware security module
IoT	Internet of things
MD-SAL	Model driven service adaptation layer

NFV	Network function virtualization
NFVO	Network function virtualization orchestrator
ONF	Open Network Foundation
PKI	Public key infrastructure
SDN	Software defined network
TPM	Trusted platform module
VFNM	Virtual network function manager
VIM	Virtual infrastructure manager
VM	Virtual machine
VNF	Virtual network function
VNFi	Virtual network function infrastructure

Bibliography & References

Software-Defined Networking (SDN) Definition. Open Networking Foundation. Online. Downloaded July 11, 2017. <https://www.opennetworking.org/sdn-resources/sdn-definition>.

OpenFlow. Open Networking Foundation. Online. Online. Downloaded July 11, 2017. <https://www.opennetworking.org/sdn-resources/openflow>.

[ETSI NFV 1] ETSI GS NFV 002 V1.1.1 (2013-10): Network Functions Virtualization (NFV); Architectural Framework. European Telecommunications Standards Institute. 2013. Online. Downloaded July 11, 2017. http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.01.01_60/gs_NFV002v010101p.pdf.

[NGMN 5G] NGMN 5G White Paper. NGMN Alliance. 2015. Online. Downloaded July 11, 2017. https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf.

[ONF Primer] SDN Architecture – A Primer. Open Networking Foundation. Online. Downloaded July 11, 2017. <https://www.opennetworking.org/images/stories/downloads/sdn-resources/7-26%20SDN%20Arch%20Glossy.pdf>

[RFC 7803] RFC 7803: Network Configuration Protocol (NETCONF). Internet Engineering Task Force. 2011. Online. Downloaded July 11, 2017. <https://tools.ietf.org/html/rfc6241>.

[ODL] OpenDaylight Controller: MD-SAL. OpenDaylight Wiki. A Linux Foundation Project. Online. Downloaded July 11, 2017. https://wiki.opendaylight.org/view/OpenDaylight_Controller:MD-SAL.

[ONF 3] ONF TR-530. Threat Analysis for the SDN Architecture. Open Networking Foundation. Version 1.0. July 2016. Online. Downloaded July 11, 2017. https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/Threat_Analysis_for_the_SDN_Architecture.pdf.

[ETSI NFV 2] ETSI GS NFV-SEC 001 V1.1.1 (2014-10): Network Functions Virtualization (NFV); NFV Security; Problem Statement. European Telecommunications Standards Institute. 2013. Online. Downloaded July 11, 2017. http://www.etsi.org/deliver/etsi_gs/NFV-SEC/001_099/001/01.01.01_60/gs_NFV-SEC001v010101p.pdf.

[Lee] Resource Management in Service Chaining. IETF Internet-Draft. draft-irtf-nfvrg-resource-management-service-chain-0. July 2015. Online. Downloaded July 11, 2017. <https://tools.ietf.org/id/draft-irtf-nfvrg-resource-management-service-chain-01.html>.

Enhancing Public WiFi Security

A Technical Paper prepared for SCTE•ISBE by

Ivan Ong
Principal Engineer
Comcast
1701 John F Kennedy Blvd
Philadelphia, PA 19103
215-286-2493
Ivan_Ong@comcast.com

1. Abstract

Approximately 15 million Xfinity WiFi public hotspots are available today domestically, the total tonnage or usage for the month of Jan 2017 was 174 Petabytes and there were 1.79 billion sessions. More than 7 million apps were downloaded on the network in that month. Ensuring the security of the users on public hotspots will minimize threats and provide an overall improved user experience. This paper will explore EAP-TLS mechanism and its implementation approach on public WiFi.

2. Overview

At a minimum, a typical hotspot broadcasts an open and secure SSID today, with EAP-TTLS and EAP-GTC mechanisms to ensure service compatibility primarily for Android and iOS systems. A user has the option of associating their mobile client to either SSID, however, a profile is generated for secure SSID association as it offers some advantages over a non-secure SSID. A mobile client associating to a secure SSID today will have the ability to generate and download a profile that will ensure connectivity to a valid trusted service provider hotspot due to the server authentication that occurs in the EAP inner mechanism that is employed.

The majority of Android mobile clients are accommodated by the EAP-TTLS method which performs a two-phase authentication: the outer authentication, from server to client, mandates the client to authenticate the server certificate. Once validated, a TLS tunnel is established; the inner authentication, from client to server, will then exchange encrypted information typically based on a simple non-TLS authentication method. In the case of Xfinity WiFi, username and password credentials are exchanged.

The majority of iOS mobile clients are accommodated by the EAP-GTC method, a basic EAP standard that utilizes token management for authentication. This method was employed due to the fact that iOS requires an Apple based application to generate a profile; without relying on the user to download an application that may seem intrusive to the inherent connection manager, this method was the most viable approach. This is used as a stepping stone to generating the profile remotely and subsequent associations will leverage the EAP-TTLS mechanisms. EAP-GTC does not protect the authentication data, both text challenge and reply are sent as clear text.

Combine with EAP-PEAP, MSCHAP, for windows mobile clients, EAP-TTLS and EAP-GTC mechanisms do offer some form of security assurance, however, the path to improve user experience led us towards the employment of EAP-TLS which offers improvements over these EAP methods. EAP-TLS mandates server and client certificate based authentication. In brief, certificates are used instead of a subscriber's username and password as credentials. The content within the certificate typically consists of various attributes that are encrypted, the subscribers username and password are typically not included within. While this requirement makes it more secure than most other EAP methods, it is more challenging to deploy as it requires certificates to be generated and there is additional cost incurred in using a trusted

Certificate Authority (CA) to produce these certificates. There are benefits of using a trusted Certificate Authority as most of the Global Certs are already embedded within various devices operating systems and applications. This enables the client to server/AP authentication to occur seamlessly and to ensure the subscriber is associated with a trusted AP and not a rogue AP.

To understand the improvements that are proposed, let's explore the specifics around EAP and IEEE 802.1x.

3. EAP Primer

EAP, Extensible Authentication Protocol, is an authentication framework developed to function at the link layer. It is defined in RFC 3748 and updated in RFC 5247. There are many authentication mechanisms that function within the EAP framework, they were developed by various vendors to accommodate different operating systems.

EAP is designed to function within the link layer, negotiation occurs between the supplicant (end user mobile client) and the Authentication Server (typically a RADIUS server) via an Authenticator (typically an Access Point with RADIUS client). RADIUS attributes are required to be supported for each new authentication mechanism that is developed. Figure 1 depicts the high-level flow of an EAP transaction

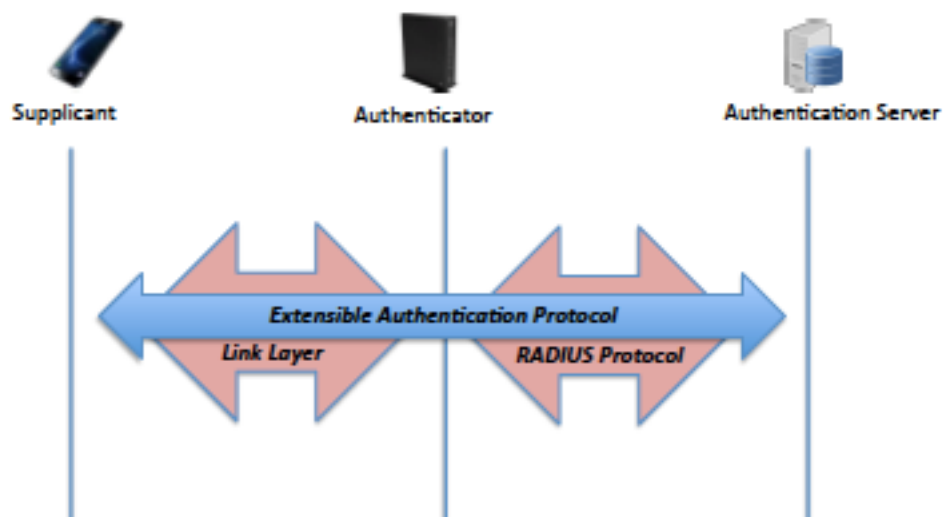


Figure 1 - High Level EAP exchange between Supplicant and Authentication Server via Authenticator

4. IEEE 802.1x and EAPOL

The IEEE 802.1x standard and EAP over LAN (EAPOL) are typically coined as the same term, they referenced one another within the standard. IEEE 802.1x is found within the 802.11i standards where security attributes such as WPA2 (WiFi Protected Access version 2), TKIP and AES are derivatives.

IEEE 802.1x is an IEEE standard for port based Network Access Control, what this translates to is access managed by means of a port. Until authentication is validated, access is blocked by the port to a LAN or WLAN. It does so by encapsulating EAP protocol over IEEE 802, also known as EAPOL.

Figure 2 depicts the EAPOL frame and Figure 3 the EAPOL high level flow, as mentioned previously, the EAPOL frame contains destination and source MAC address but no network frame blocks as it functions on the Link Layer

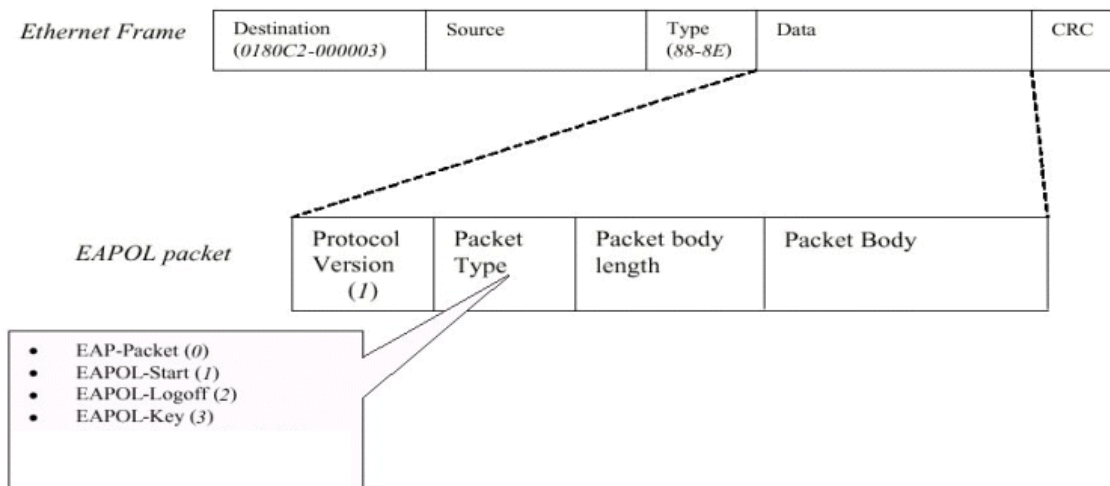


Figure 2 - EAPOL Frame structure

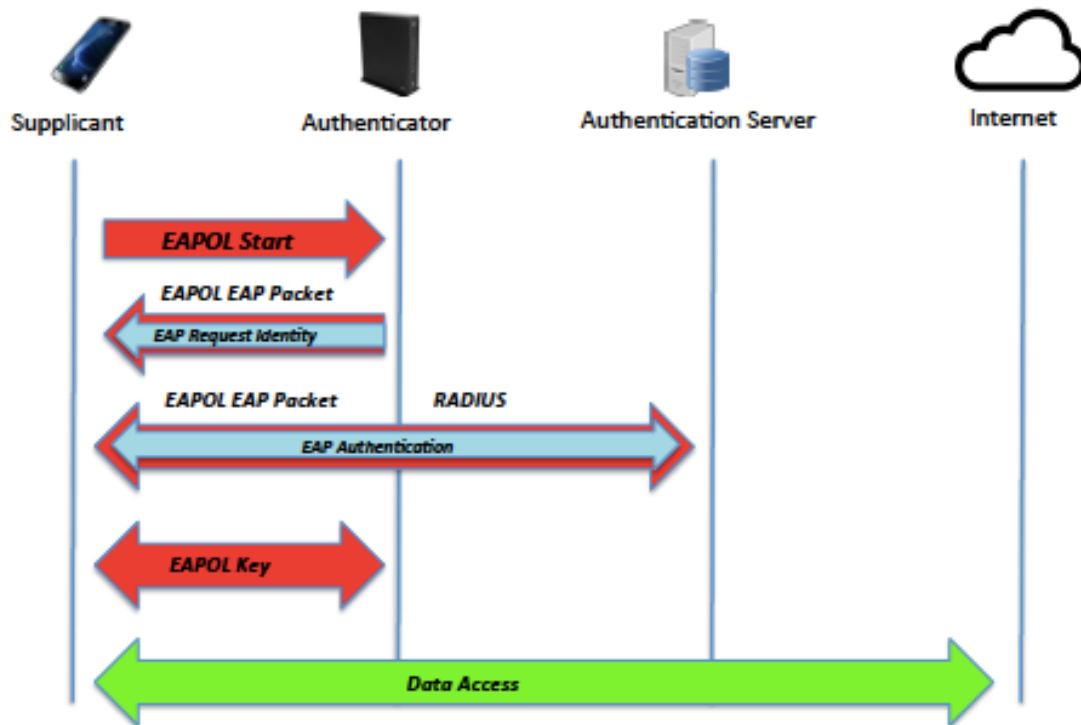


Figure 3 - EAPOL High Level Flow

5. EAP-TLS Primer

EAP-Transport Layer Security (TLS) is the authentication mechanism selected to provide a secure experience. EAP-TLS, as defined in RFC 5216, is an IETF open standard that requires the use of both client and server certificates, preferably from a trusted Certificate Authority (CA). This is known as mutual authentication where certificates from both entities are validated prior to enabling access. Certificates adhere to the x.509 standard where the process of invoking a certificate is defined clearly with specific attributes such as encryption type, method, organization name, expiration date, extensions, among others are supported by most Certificate Authorities in their Public Key Infrastructure specifications. Please refer to Figure 4 for the structure of an x.509 certificate:

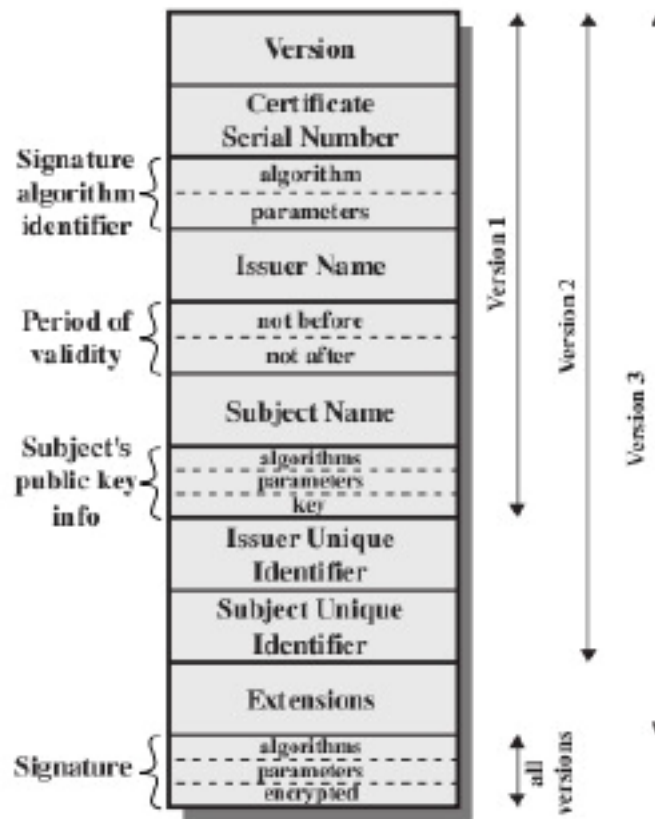


Figure 4 - x.509 Certificate Structure

Attributes within the certificate are defined by the organization and issued by the Certificate Authority. A public and private key is established as part of the certificate generation process, the certificate(s) are then chained to a trusted root certificate of the CA. The server and client certificate will contain the private keys that are only known to itself and the CA, it is used to decrypt the public key that are then distributed to other devices/applications. An analogy would be a user accessing their banking website today, the mobile device employs the use of a browser that contains certain certificates (private key) that are inherently trusted and embedded. When the banking website is accessed, the SSL transaction would be called upon and check the certificate (public key) of the website to ensure it is a legitimate website. With EAP-TLS, the exchange occurs on both ends with the client and the server validating both certificates. If the server certificate is not who they claimed to be, then the authentication fails, this would be akin to preventing connectivity to a rogue AP with the added benefit of validating the client as well. If the server certificate passes but the client certificate fails, that could be translated as a user who is no longer a valid subscriber due to a number of reasons.

Figure 5 depicts the EAP-TLS flow, it is based off the EAPOL message flow if compared with Figure 3. Digital certificates are used in place of username and password as credentials. Mutual authentication exchange certain RADIUS attributes between the mobile client (supplicant) and the Authentication Server (RADIUS server) via the AP (NAS). A series of RADIUS messages are exchanged after the initial EAPOL transaction is initiated, a RADIUS Access-Request will be invoked as a result of an EAP-Response/Identity call. Common RADIUS attributes in this flow include NAS-Id, NAS-Port, Calling-Station-Id. The RADIUS server in turn will send a RADIUS Access-Challenge message to the AP which produces an EAP-Request message to the mobile client. Common RADIUS attributes in this flow includes Session-Timeout, Service-Type. The mobile client then responds with an EAP-Response [containing the certificate] to the AP which will produce a RADIUS Access-Request message where AAA will then validate the certificate. If certificate is valid, AAA will return a RADIUS Access-Accept message and the session keys to the AP, and in turn an EAP success.

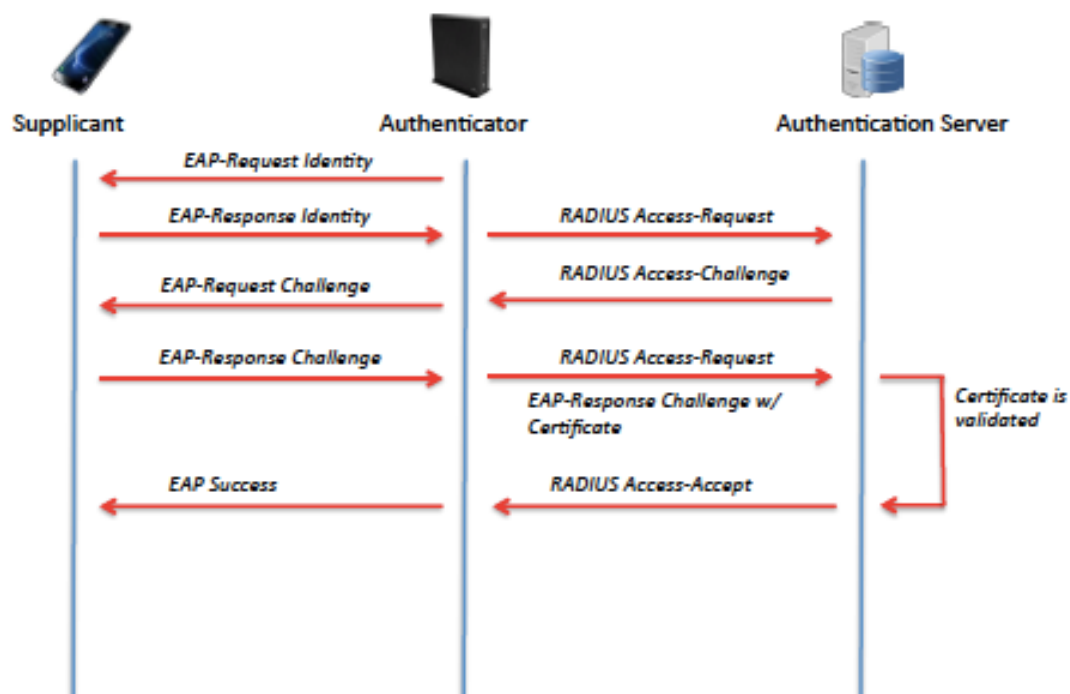


Figure 5 - EAP-TLS message flow

6. Certificates

Essentially, two forms of authentication will occur in EAP-TLS: network will be authenticated against client, then client will be authenticated against network. Instead of using username and password, certificates are the common elements involved. On the network side, the certificate may reside on the AAA or any component playing the role of the AAA server. Figure 6 displays a sample certificate on the network side:

```
Certificate:
Data:
  Version: 3 (0x2)
  Serial Number:
  00:
  Signature Algorithm: sha256WithRSAEncryption
  Issuer: C=US, O=DigiCert Inc, CN=DigiCert SHA2 Secure Server CA
  Validity
    Not Before: Jan 31 00:00:00 2017 GMT
    Not After : Feb  5 12:00:00 2020 GMT
  Subject: C=US, ST=PA, L=Philadelphia, O=Comcast Corporation, OU=xfinitywifi, CN=
  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    RSA Public Key: (2048 bit)
      Modulus (2048 bit):
        00:c8:c3:3f:4b:4a:67:b9:17:35:34:e6:0f:65:a2:
        00:0b:0c:1d:16:7b:c1:c1:c1:c1:c1:c1:c1:c1:c1:c1:
        55:30:4d:40:0c:20:c7:03:ac:ad:10:9f:04:cc:37:
        01:51:39:4e:2d:0e:c1:c0:04:ac:54:0d:c2:e3:0d:
        10:17
      Exponent: 65537 (0x10001)
  X509v3 extensions:
    X509v3 Authority Key Identifier:
      keyid:0f:0
    X509v3 Subject Key Identifier:
      A5:6
    X509v3 Subject Alternative Name:
      DNS:
    X509v3 Key Usage: critical
      Digital Signature, Key Encipherment
    X509v3 Extended Key Usage:
      TLS Web Server Authentication, TLS Web Client Authentication
    X509v3 CRL Distribution Points:
      URI:http://
      URI:http://
    X509v3 Certificate Policies:
      Policy:
      CPS:
      Policy:
    Authority Information Access:
      OCSP - URI:http:
      CA Issuers - URI:
    X509v3 Basic Constraints: critical
      CA:FALSE
  Signature Algorithm: sha256WithRSAEncryption
  4b:ce:3a:7f:75:57:e4:4a:ee:d4:50:06:9b:46:f5:08:ac:1e:
  e3:f1:cc:09:5c:ee:00:4d:b9:3b:05:7d:db:9f:6d:15:a0:06:
```

Figure 6 - Server Certificate

Some of the common attributes defined within a server cert are the Country, State, Locality, Subject, Organization, Organizational Unit, Common Name. The certificate encryption may be RSA or ECC based, typically with a 2048 modulus bit encryption applied. Ultimately, the server cert is chained to the trusted CA Root Cert that is prevalent on most device operating systems root store.

Figure 7 displays a sample certificate on the client side which is generated based on certain user attributes. Ultimately, the certificate is chained to a trusted Root CA.

```

7fLtsCCUjGdZvgv0eDmKXzKkRQaOHC8W7FSDStYtGqI2ziNkj/*S1TGW1/T02
ykMQuo/ZH4lQtVrJgCt2/DpAPBRFXLAvdS1BFvPkMUegctiKtayGnxAT3demnQt
+79LiRJ0BJZBwJVvAx/Ae02M/l/IRvpM4ua20tjtanyUgCQZASq06uLyEaqAgwQ
h41bLWLP6204ATEatlicUoP+aBr/5m2+hglj4sokS0d4ZIKwT0ju/M1ltqb/S5A49
qS9m43Wrkk2DYkcTAJsq9nMTSFonSHcC800FQ+ONHDyJ/rft8nu3SDsvCW/R90axMLa
hfweXy5jU15c0lcFYG0dgF0TR0/Nz7X3wiJ/eOPPo608nM0d460HaFvyfQjmb90C
L1HLB17ws1vcN/XwS8QgmAOSGgcY2RRWLxYmfkcIPdrMnopYCi fefQ==
-----END RSA PRIVATE KEY-----
PKCS7 Encrypted data: pbeWithSHA1And40BitRC2-CBC, Iteration 2000
Certificate bag
Bag Attributes
    localKeyID: 01 00 00 00
subject=/C=US/ST=PA/O=Comcast Corporation/OU=[REDACTED]/CN=XFINITYWiFi
issuer=/C=US/O=Comcast Corporation/CN=Comcast Internet Security Corp., Inc
-----BEGIN CERTIFICATE-----
MIIFDTCCA/WgAwIBAgIQCT2GRVPpcMUR4RKTOuhM40DANBgkqhkiG9w0BAQsFADBR
MQswCQYDVQQGEVZlU2EcBHoGA1UEChNTQ29tYTZFdzdCBDb3Jwb3JhdGVlbjEkMCIG
A1UEAxMBT29tYTZlbnVzdCBUCnVzdGVkIXZlZXIgdUNBIEYENBM4XDTE3MDVxMjAwMDAw
MFoGQ2Vkd2M2YXMyEyMDAwMFowZDZELMAkGA1UEBhMCVVMxChZAJBgNVBAGTA1BBMRww
GgYDVQQKEXdNb2J1ajYNNDIENYcnBvcmlF0aW9uMScwJQYDVQLXLE4xyNjA2MzcwMzAz
MDEuMDEuMTIwY2Vkb2J1YXN0LWN1c3QvFDASBgNVBAMTC1hGU5U1VJFLXAuUZpMIIBiJA
NBgkqhkiG9w0BAQEFAEQCAQAEAQAMICBCGKAQAQAAQ1XwouUMeTPLTDQFzpxLe9gg3WM1
Gtn+PEPdxMuLS1532mb0csi on1dmxfgKmEtUm4PMJP5U8QvmQrtskvXH87dz14g
oRSMUpj4Dng1fyfMYNy5AL5Xce9/GYN4B1dSVkv8kt+3XMd7iUPhgW3C219b411
kshHG158rm000+pJpfdnElxa08kdTD/VBRZ+yapbfq/K6RUMEgnH25TtQ1n9/Nj
yhrroWKFP4MYGS4DvsyGzqCD/DkZxI2XyHytu7HSWCiNGZYbVVSa46We9DG14Qa
szf+tdcllwjCs/M981h01/n94y7B/rF9shgBLExiYJJePGXiFCJ6rSivGQIDAQA8
0+FBU+t11awTjYwYDR0j8BgwFoAUDI7q23BSlx937wCKvhoFGPInboHQYDVR0B
BBYEFPs2b/ben/+Mc91v0SLBWUK2wko3MAwGA1UdEWEB/wQCMAAwDgYDVR0PAQH/
BAQDAQwGBMMGA1UdQMMAAGCCSGAQUFBwMCMEEGA1UdIAQMDOmgNY9IVTZIAyb9
bAYwKjAOBggrBgEFBQDCARycaHR0cHM6Ly93d3cuZGlmaGlnaW9uZWY2LnNQUCZCB
```

Figure 7 - Client Certificate

7. Implementation

The benefits of employing certificates mutually between server and client and vice versa ensures a secure and seamless on-boarding experience for the user. To achieve this, many factors come into play such as certificates from trusted Certificate Authorities, core network infrastructure upgrades (such as captive portal development), policy changes (primarily to AAA), AP security features support (to support EAP mechanisms), and app development for various client operating systems (to provide an intuitive user interface)

The user interface will essentially redirect a user to a captive portal to assist in generating and installation of a client certificate based on certain required parameters. This is the initial encounter which will not be repeated unless in the event of an expired certificate or when deemed necessary to reissue a new certificate. Subsequent client associations will not require any user input, the EAP-TLS exchange would ensure a transparent, secure connectivity between the client and AP. Even in a roaming scenario between APs, the certificates exchange would occur in the background and help authenticate and validate the mobile client on their respective service provider public WiFi network.

8. Looking Forward

The security implementation of EAP-TLS will help ensure the proper framework is in place to accommodate other EAP mechanism as additional services are offered that may take advantage of the public WiFi network such as cellular service. The underlying foundation are in place to also support other industry standard specifications such as Hotspot 2.0. As the implementation help evolve the public WiFi network to a carrier grade level, the users will only stand to benefit from these changes.

9. Abbreviations

Acronyms	Description
AAA	Authentication, Authorization and Accounting
AP	Access Point
CA	Certificate Authority
EAP	Extensible Authentication Protocol
EAPOL	Extensible Authentication Protocol over LAN
GTC	Generic Token Card
IEEE	Institute of Electrical and Electronics Engineers
LAN	Local Area Network

10. Bibliography & References

802.11 Wireless Networks: The Definitive Guide, 2005, O'Reilly, Matthew Gast

R10.0 WiFi Authentication with EAP. LCW442H-V3.0-SG Edition 1, NokiaEDU

Figure 2: http://www.zyxeltech.de/SNoteZW5_362/app/8021x.htm

Figure 4: <https://www.slideshare.net/koolkampus/chapter-ns4>

Service Theft in DOCSIS Networks

Identifying the Hidden Leaks in Your System

A Technical Paper prepared for SCTE•ISBE by

Egbert Westervelt
Sr. Security Engineer
Irdeto
Eindhoven, NL
EWesterveld@irdeto.com

Edward Florendo
Service Delivery Manager
Irdeto
Hoofddorp, NL
edward.florendo@irdeto.com

Dave Belt
Technology Evangelist
Irdeto
Conifer, CO US
(303) 653-7647
dave.belt@irdeto.com

Introduction

The following describes a technical proof of concept (POC), developed by the authors, and performed for a major U.S. cable operator. This effort looks at the prevalence of service theft on broadband networks, focusing initially on Data-over-Cable Service Interface Specifications (DOCSIS), but the technologies applied herein can extend to other networks including digital subscriber line (DSL), fiber and Wi-Fi.

The effort incorporated a two-tiered approach. The first component consisted of extensive forensic research via the Internet, using open sourced intelligence (OSINT) techniques, identifying known attack vectors directly from the hackers themselves. This qualitative intelligence was then used to develop a service monitor appliance designed to actively monitor for service theft.

The results of this comprehensive theft assessment are quite significant with regards to the level of theft involved. When one looks at the lost revenue, as well as the increased overbuild costs, it quickly becomes impractical to not address this hidden issue.

The Problem

The high-level concepts and technologies implemented within this effort are applicable to many different broadband technologies. For the sake of keeping the topic focused, we will focus solely on DOCSIS network technologies herein.

Theft of service has been a well-known problem within DOCSIS networks and there have been some significant efforts by the industry to address the issue. The release of the DOCSIS 3.0 security specification (Cable Television Laboratories) is one case in point. By introducing the use of public key infrastructure (PKI) device authentication, each device can now uniquely pair to a cable modem termination system (CMTS). This specification had a significant impact on system security, so much so that it has spawned a black market for pre-DOCSIS 3.0 modems on the Internet. Many operators have pushed forward to move their entire networks to DOCSIS 3.0 technology, which would provide a higher level of scrutiny and security, however this effort involves the elimination of all older devices from the network, a significant consumer premise equipment (CPE) expense.

Assuming that DOCSIS 2.x devices have flawed security, the efforts described herein are focused primary on DOCSIS 3.x devices. While the DOCSIS 3.0 specification outlines the use of PKI certificates for device authentication, it says little about the implementation on the devices themselves. As a result, the protection of the credentials, as well as the overall protection of the device platform varies greatly and leaves significant gaps for theft to occur.

Project Initiative

The POC described herein was performed for a major North American operator with the intent of quantifying the theft described. A general gap was identified between the CMTS and actual reconciliation with the operator's billing systems. A major portion of the POC was intended to bridge this gap. Simultaneously, the vendor's cybersecurity team performed a comprehensive search of the "dark web" to identify known attack vectors against the operator's network. This information was then fed back into the reconciliation system to identify these attacks occurring and mitigate in real time.

Approach

1. System Assessment

The initial effort consists of inventory and assessment of the operator's system. CMTS as well as CPE hardware are inventoried and current configurations noted. These are compared to the manufacturer recommended security configurations and changes are suggested where necessary.

2. Online OSINT Research

For the up-front research, the solution vendor performs an online deep dive using the inventory information obtained during system assessment. Research into the Internet is performed looking specifically for attacks against the CPE utilized by the operator, as well as targeted attacks against the operator's network itself. This information yields specific attack vectors that are then implemented into the service monitor appliance utilized for modem reconciliation, continually improving its accuracy of detection and removing rogue devices.

3. Modem Reconciliation

Modem Reconciliation consists of the integration of a service monitor appliance into the operator's data center. This appliance consists of a Hadoop database that identifies all devices connected to the network via CMTS and reconciles these devices with a valid billing address. Failure to reconcile indicates one of two issues: either the device was not registered properly via the operator's processes, or a case of real theft is occurring. In the former case, it is in the operator's interest to reconcile these processes so that an accurate measurement of theft can be made in the latter.

Once the initial reconciliation is complete, the system continues monitoring the usage patterns of the connected devices based on the attack vectors identified in the system assessment. Once a positively identified rogue device is identified, it can be mitigated per the operator's policies. It is important here to err on the conservative side as removal of valid devices clearly creates a poor customer experience.

Ultimately it is the continuous feedback loop of the online research coupled with the service monitor appliance that makes this solution a success. As new mitigations are implemented, the cybersecurity research team can monitor the hackers' reactions online, staying in sync with them as the continual cat and mouse game plays out.

Findings

The following provides a summary of the findings from the POC.

1. Online OSINT Research

Multiple exploits were identified through online research, providing a very distinct picture of the avenues used by hackers into an operator's system.

1.1. Hacking Exploits

A significant body of knowledge has been dedicated to the direct hacking of DOCSIS modems on an operator's network. Access to the DOCSIS endpoint enables free access to service as well as potential access to the operator's backend systems on more sophisticated gateway appliances. Some of the more significant forums and discussions identified were as follows.

- Discussions on how to directly hack into CPE – Step by step instructions on how to hack directly into DOCSIS devices are prevalent, not only on dark web sites, but also on regular web sites. Exploits are usually on a specific vendor and device basis and frequently utilize the Joint Test Action Group (JTAG) interface and/or soldering of leads and interfaces directly onto the circuit boards. Instructions include detailed circuit board diagrams indicating access points for a “noob” hacker.
- Discussions and techniques on how to circumvent security measures – Similar to device hacking, these discussions are at more of a system level and focus on how to gain access to the network, primarily through a hacked device.
- Discussions of security measures and countermeasures by multiple system operators (MSO) – Frequently as soon as an operator rolls a new piece of firmware with a new security measure in place, a slew of exploits dependent upon the previous vulnerability become unavailable. To a potential piracy service this is the equivalent of an operator outage and as such the hackers are immediately looking for the next exploit. The cat and mouse game of measures and countermeasures are well documented within these hacker forums.
- Trade of modem device files – Frequently an exploit is enabled, or supplemented by the modification, ingestion or spoofing of modem data. Sharing of these memory images enables a fellow hacker to recreate a particular exploit.
- Trade of medium access control (MAC) addresses and PKI certificates –The DOCSIS standard ties the device authentication to the MAC address so these are frequently traded together. MAC address cloning is a well-known attack vector on DOCSIS networks. So long as two devices with the same MAC address don't reside on the same CMTS, these devices can coexist peacefully on an operator's network.
- Trading of modem configuration and boot files – Modification of device configuration and boot files enables a potentially lower tiered subscriber to obtain a higher level of service. Once a modification is made, the files along with upgrade instructions are shared openly.

1.2. Sale of Theft Related Products

Designed for the lazy hacker or non-technical person who merely wants cheap service, pay services provide a number of exploits packaged into an easy to use product. Examples are as follows.

- Sale of hacked modems – There has been since the turn of DOCSIS 3.0 a black market for DOCSIS 2.0 modems due to their inherent lack of security features. As operator's cut off these devices from their networks, the market is now turning to pre-hacked DOCSIS 3.x devices. The exploits described above are applied in a production manner to a block of devices and sold at a premium enabling free or upgraded service.
- Sale of activation services – A cable customer with a specific piece of CPE can send this to an activation service to have the exploits described above applied to their device for a fee. Alternatively, an exploit package and instructions are sometimes offered in order to gain free or upgraded service.
- Sale of hacking and modification services – For the customer with a device without a pre-defined hack, custom services are offered to open the device. The device is sent to the hacker and a general tool box of exploits is applied to find entry into the device. Once a known exploit is defined, it may be resold as one of the other services above.

1.3. Internal Theft

Internal theft is as described, theft from within the organization itself with the goal of external monetization.

- Bad contractors – These are self-identified internal contractors to a particular operator who sell enablement services. A customer with a cable modem provides its MAC address and for a fee the modem is registered within the operator's backend service.

1.4. Social Engineering

In addition to the dark web and sheer persistence, information is also obtained through social engineering techniques, also shared on the dark web. Examples include the following.

- "Play Dumb" Techniques – Methods for convincing a customer service agent (CSA) to reauthorize a device or remove a device from a blacklist.
- "Pumping" for Information– Hackers will often try to engage a CSA or a higher level service technician in conversation so they may extract additional information regarding an operator's security measures
- Names of Internal Systems –internal system names are identified and published to support hackers in their conversations with technicians when they attempt to extract security information

Based on this research it is clear that the attack vectors as well as markets created around them are numerous. Much of the information gleaned from this online search was then used to develop the service monitor appliance discussed below. This appliance allows the operator to quantify the theft occurring on a particular network.

2. Service Monitor Appliance

The service monitor appliance performs the modem reconciliation described in the Approach section. By reconciling every modem with a billing account, a consistent accounting of these devices on the network is obtained. Additionally, cross checking of firmware versions and configurations ensures that all devices on the network are authorized to be there. Rogue devices are removed based on the trust level of the detection. The following outlines some of the major theft identified on the operator's network analyzed for the POC.

2.1. No Billing Account

Devices on the network with no billing account can have multiple reasons for being unregistered. These can be via stolen credentials, internal registration or simply slipping through the operator's provisioning processes. A full 2% of the devices on the network per day fell into this category identifying the largest loss within the POC.

2.2. Billing Status Disconnected

These devices have an account associated with them; however, their billing status is set to disconnected. The reasons for this status can include internal registration, modification via social engineering or gaps within the operator's system. Of the devices within the POC, 0.4% per day were identified with a disconnected status.

2.3. Restricted Boot File

Devices with a restricted boot file have an alternate configuration than provisioned or firmware that is not authorized for use on the device. Reasons for this variation indicate some form of firmware tampering. Of the devices within the POC, .2 - .3% per day were identified with rogue firmware.

2.4. Cloned MAC Address

One of the most understood attacks is also apparently one of the most addressed. MAC address cloning on the network accounted for only .06% of theft.

Additional theft cases were monitored, however with diminishing returns moving forward. All told, the level of loss on the network is significant enough to warrant a closer look by the operator. The loss in revenue due to free service, along with the additional operational expenditure and build out costs quickly justify closing this gap.

Conclusions

We have presented the results of a POC researching DOCSIS service theft performed with a major U.S. operator. The first phase of the effort searched the dark web as well as regular web sites for information on hacking of the operator's CPE and network. This search identified numerous attacks against both, providing qualitative intelligence for the development of second phase.

The second phase consists of the deployment of a service monitor appliance to reconcile modems on the network with registered billing accounts. This implementation provided a definite accounting of rogue, unaccounted and otherwise lost devices. This appliance identified a consistent network loss of an accumulated 2.5%. This type of loss amounts to real numbers when looking at lost revenue as well as additional operational expenditure and build out costs.

DOCSIS networks have provided us with increasing levels of security with each version of the DOCSIS specification that is published. Even so, there are still gaps remaining in the implementations creating significant opportunities for theft of access.

Abbreviations

CSA	customer service agent
CMTS	cable modem termination system
CPE	consumer premise equipment
DOCSIS	Data-over-Cable System Interface Specifications
DSL	digital subscriber line
JTAG	Joint Test Action Group
OEM	original equipment manufacturer
OSINT	open sourced intelligence
MAC	medium access control
MSO	multiple system operator
PKI	public key infrastructure
POC	proof of concept
SCTE	Society of Cable Telecommunications Engineers
STB	set-top box

Bibliography & References

Cable Television Laboratories, Inc. (2016, June). CM-SP-SEC3.0-I16-160602, data-over-cable service interface specifications, DOCSIS 3.0 security specification. Louisville, CO: Author.

Automated Detection for Theft of OTT Services and Content

Identifying Your Content Out in the Wild

A Technical Paper prepared for SCTE•ISBE by

Lucas Catranis

Product Marketing Director
Irdeto
Hoofddorp, NL
lucas.catranis@irdeto.com

Brian Yuan

Product Owner
Irdeto
Ottawa, Ontario CA
brian.yuan@irdeto.com

Dave Belt

Technology Evangelist
Irdeto
Conifer, Colorado US
303 653-7647
dave.belt@irdeto.com

Introduction

Theft of video content has been an issue since the dawn of pay television (TV). Since that dawn there has been a continual security cat and mouse game between operators and pirates, with the technology actively evolving with it. Most of this technology has, however, focused primarily on protecting the video pipeline with the assumption that control ends at the playback device.

Within this paper we look at content beyond the device, after it has been played back, pirated and distributed over the Internet. We'll first look at the common methods of obtaining pirated content; then, we'll look at some of the state of the art techniques be used to combat these.

Indeed, the next generation of content protection must move beyond the device to understand and mitigate the path of active online piracy.

The Problem

As mentioned, theft of premium video content has and is a continual problem. Loss of content revenue undermines the business models of all participants of the video production and delivery pipeline. This theft is now reaching a new unprecedented level due to the ubiquity of the Internet coupled with the ready availability of tools, devices, and piracy services available to the layman. The piracy services themselves are now becoming active and imminent competitive threats to the operator's business model, many looking like legitimate over-the-top (OTT) services to the novice.

Content itself has a natural decay with regards to its value over time. In general when we speak of "premium content" we are frequently referring to the age of the content itself or the quality of the encoding and delivery. As such, live sporting events, early release movies and ultra-high-definition (UHD) content, are classified as premium content. This content has sufficient demand for a pirate to successfully monetize and also serves as the greatest loss to the content owners and operators.

Many of the techniques discussed herein are aimed primarily at premium content due to its value to the providers.

Content Theft Methods

1. Multicast

In multicast systems, a hacker can go through the exercise of compromising the set-top box (STB) via Joint Test Action Group (JTAG) interface hacking, side channel attacks and the like. If one is merely interested in obtaining high quality video though, the easiest route is via the high-bandwidth digital content protection (HDCP) port.



Figure 1 – Workflow using HDCP Ripper

As shown in Figure 1, an HDCP ripper is plugged into the high-definition multimedia interface (HDMI) port of an STB. The ripper simulates a valid HDCP client and is able to decrypt the content delivered over the port. The decrypted stream is then captured via a personal computer (PC) and re-broadcast over a streaming service.

The ease with which HDCP rippers are obtained and the simplicity of such a system make online piracy detection the only mitigation to such a scenario.

2. Over the Top

Theft of OTT services offers a different set of opportunities. If one has a 10' viewing device in the form of an internet protocol (IP) STB, one can again use the HDCP ripper described above. Other attacks involve the sharing of credentials which are necessary for authentication on unmanaged devices. Hacking of an operator's application itself can enable free service to a large population via distribution of the application.

While we do our best to secure the delivery of the content up to the viewing screen, content can and is getting out. To plug this hole in the revenue model we need to secure content from the outside in.

Content Identification Methods

Due to the prevalence and dynamic nature of content piracy, correctly identifying content in a timely and cost effective fashion is key to significantly reducing the amount of content theft which leads to subscriber churn. Various methods are deployed to identify pirated content with levels of information varying based on the technique. To get a highly accurate picture of the theft of a particular piece of content, multiple complementary methods should be deployed.

1. Metadata Analysis

The first step in identifying content is to comprehensively sort through the myriad of content freely available on the internet to the content which is specifically infringing on the content provided by the service provider. Step one typically uses keywords and metadata analysis in a variety of search modes; from search engines to social media, the internet should be scoured for potential streaming/content downloading sources. Common techniques include multi lingual keywords (both positive and negative), metadata analysis for validated users, and metadata analysis based on posted data such as file name, type, and run time. Depending on the content life cycle, these keywords are aggressively applied by automated systems to rapidly identify potential new targets of interest.

2. Video Fingerprinting

Well architected metadata analysis still typically generates hundreds to thousands of potential real time infringements for each broadcast event. Sorting through each of these with human analysts is neither timely nor cost effective, so the solution is to use automated video identification techniques such as video fingerprinting.

Video fingerprinting typically takes a real time or video on-demand (VoD) file for an upcoming event and analysis of the video for unique characteristics to create a unique digital signature of the upcoming event. This digital signature is, in turn, compared automatically against the gamut of infringements found via metadata analysis and sifted into three categories: full match, partial match, and no match.

The key to effective fingerprinting is the accuracy of the fingerprinting technique. For example, there are many common video obfuscation techniques employed to circumvent the fingerprinting techniques of common sites such as YouTube. And while sites like Facebook employ highly effective (proprietary) implementations of audio fingerprinting, the current state of their video fingerprinting technology is still quite easy for content pirates to circumvent. For proper content identification, a layered identification approach is recommended to ensure that content is correctly characterized and identified in a timely fashion.

3. Deep Learning Image Recognition

There is a great deal of information which can be learned from piracy data; source identification, subscriber demand, cryptographic integrity, and source leakage are just a few of the pieces of information which can be obtained by carefully analyzing piracy data. As an example, the logo can tell you which broadcast infrastructure sourced the leak. This information, in turn, can tell you the efficacy of their anti-piracy/anti-tampering technology which may be correlated to the technology used within your existing network.

Traditionally, this was done via human analysts to ensure proper tagging of the data for information ranging from video quality to broadcast information. Nowadays, this is performed via machine learning/image recognition which dramatically speeds up and improves the overall accuracy and quality of the piracy information provider service providers with an even more accurate picture of the piracy landscape.

4. Watermarking

Watermarking is the deliberate injection of information into a piece of content, with the intent of identifying its source upon later identification. Watermarking can occur with varying layers of granularity, with session based watermarking being the most useful. With each session identified with its own identifier, it is a simple forensic exercise to trace a piece of content back to the leaking device and mitigate accordingly.

5. OTT Credential Theft Monitoring

There are a variety of OTT credential theft mechanisms.

First is the sharing of OTT credentials between friends and family members. This takes a variety of forms from the relatively passive form of simply sharing user identification and password information to actively adding the devices of friends and family to one's active device list.

Second is the sharing of credentials for potential profit. For example, a university student may "rent" out OTT credentials to others to earn a little money on the side as they know their own parents/grandparents will not use any of the provided OTT credentials.

Third is the most nefarious form where hackers attack and obtain credentials via hacks of applications or via man-in-the-middle attacks via various public wireless hotspots. These fraudulently obtained credentials are, in turn, offered via dark web websites to members of the public.

All of these result in a major loss of potential subscribers to Operators in addition to churn which can be attributed to web streamed content re-broadcasts.

Addressing this specific mode of content loss requires a variety of tools. Restricting the number of concurrent streams per household via Rights Management concurrent stream control is one popular method. This is often supplemented by adding geo restrictions and controlling the number of registered devices per account to control access. To further protect against unplanned credential sharing, these measures are supplemented via application security such as cryptographic protection and communication as well as big data research to better track, visualize, and react to changes in OTT content usage to better react in real time to unwanted content theft.

6. Peer to Peer Monitoring and Analysis

P2P technology is typically used to share non-live content such as TV shows and films, but it can be used to also share live content via technologies such as Acestream and SOPCast. As a result, P2P technology applies to both nonlive and live content.

P2P non live content is typically shared via Torrent index sites. Bit Torrent is the dominant protocol in the P2P space so users search these Torrent index sites to obtain P2P infohashes. Once they have obtained this information, they use a Bit Torrent client to join a "swarm" and obtain all the necessary file parts from peers in their network.

P2P Live content is shared via SOPCast and Acestream. These addresses are shared via a variety of mechanisms from websites to social media. Social Media platforms such as Twitter and Reddit are fast becoming a preferred method for distributing this real time P2P piracy information.

To cover both forms of P2P piracy, crawlers for both the Torrent index sites as well as social media sites must be built. Using metadata analysis, the number of potential infringements must be quickly sifted so that an optimal number of swarms can be joined for video capture to enable video fingerprinting identification. This technology is key for enabling rapid cost effective anti piracy services for both real time and VoD.

In addition, P2P data can be collected to uniquely understand demand on an aggregate and geographic basis. This business intelligence information, in turn, can be used by operators to determine the potentially elasticity of piracy in their regions of interest in response to marketing, pricing, or product campaigns.

Mitigation

Mitigation of content can take many forms, including legal, based upon the nature of the detection. Upon the detection of a known piece of content, the issuance of a takedown notice is a common approach. In the case of real time content such as sporting events, these notices must be issued quickly otherwise they have little effect.

On a more technical level, assuming techniques such as watermarking are used, individual devices can be targeted for disablement. If an operator can identify the source of a leak, cutting off the client is the most direct solution. This approach is somewhat risky as the client must be identified as a legitimate pirate point. Cutting off legitimate users from their service never ends well for the operator.

Conclusion

We've historically been living with piracy in parallel with our video delivery systems and making the value judgement of how much security is enough. The easy availability of tools and devices make piracy easier than effort and is feeding a monetized piracy industry that is mounting a real challenge to video operators moving forward.

In order to stem this hemorrhaging, we must now look beyond the device and identify our own content in the wild and take appropriate actions to mitigate its theft. Luckily, as discussed herein, the technologies and techniques exist and are being deployed on operator- wide scales.

Abbreviations

HDCP	high-bandwidth digital content protection
HDMI	high definition multimedia interface
IP	internet protocol
JTAG	Joint Test Action Group
OTT	over-the-top
PC	personal computer
STB	set-top box
TV	television
UHD	ultra-high-definition
VoD	video on-demand

Bibliography & References

Motion Pictures Laboratories, Inc. (2010, May). levels of verification for P2P scanning, version 2.0.1. San Francisco, CA: Author.

AUTHOR INDEX 2017 Fall Technical Forum

Acke, Willem.....	162	Fagadar, Mihai.....	162
Akkala, Srimi.....	226	Fantuzzi, Doug	368
Al-Banna, Ayham.....	569, 773	Fiorenzo, Mariela.....	1140
Albano, Claudio.....	1059	Flask, Robert J.	77
Alrutz, Mark	1048	Flesch, J.R.	807
Andreoli-Fang, Jennifer.....	935, 961	Florendo, Edward	1797
Ansley, Carol.....	1190	Gamble, Darren	866
Ariesen, Jan	1625	Gauvreau, Jean-Louis	1560
Arnold, Brent.....	585	Gayton, Jim.....	1
Baldry, Jon.....	56	Ge, Shengbo	1572
Baron, Gregory.....	737	Ghuman, Harj	991
Baselice, Michael.....	737	Gibellini, Emilia	1140
Bastian, Chris	536	Gibson, Mark.....	19
Bellini, Marc.....	491	Glapa, Martin J.	1419
Belt, Dave	1797, 1804	Goeringer, Steve	1681, 1774
Bernstein, Alon.....	282	Gold, Ken	1
Brittingham, Mark	238	Gosko, George.....	737
Brownell, David	1206	Greene, Todd	355
Brzozowski, John Jason.....	238	Gyori, Robert.....	1718
Bugajski, Mark	1486	Hamzeh, Belal	807
Burg, Bernard	536	Harris, Steven	1537
Burroughs, Steve	914	Hatambeiki, Arsham.....	388
Campos, L. Alberto	1008	Hawkins, John	43
Carro, Gabriel.....	1140	Hayes, Erin	464
Catranis, Lucas	1804	Hayes, Keith R.....	64
Chamberlain, John.....	1048	Hernandez-Valencia, Enrique.....	1419, 1648
Chapman, John	1353	Hickey, Wayne	1434
Chapman, John T.....	935, 961	Hintzman, Zane	1691
Cheevers, Charles	807, 1190, 1249	Holobinko, John	355
Cheikhrouhou, Anis.....	557	Howard, Daniel.....	708, 737
Chen, YuLing	282	Howe, Jeff.....	773
Chow, Hungkei.....	1419	Jia, Zhensheng	1008
Cloonan, Ruth.....	1080	Jin, Hang.....	1353
Cloonan, Tom.....	569, 773, 1080, 1165	Jones, Doug	690
Coomans, Werner	1419	Joseph, Jean-Philippe	1648
Cooper, Mike.....	585	Kahn, Brian.....	307
Cuffaro, Angelo.....	1560	Kapauan, Prudence	1460
Cunha, Gary.....	412	Kelkar, Anish.....	557
Dai, Yuxin	647	Knittle, Curtis	1008
Davenport, Jim	557	Kocks, Chris	1501
De Arca, Florencia.....	1140	Kogan, Ira	368
Dharkar, Supriya.....	737	Kourlas, Tony	1446
Dolan, John.....	708	Krapp, Steven	1667
Dorairaj, Sanjay	536	Lartey, Franklin	506
Downey, John	673	Lewandowski, Benny	585
Early, David S.	325	Lintz, Chris	1340

Liu, Tong	267, 793
Loeffelholz, Todd	660
Lopez, Brionna	914
Luke, Chris	238
Manuga, Ajay	104
Maricevic, Zoran	1059
Martel, Etienne	124
Marut, Daniel.....	737
Masoud, Fady	112
Matatyaou, Asaf.....	181
Mathur, Tushar	1080
McCarthy, Sean T.....	1589
McCluskley, Michael.....	1249
McManus, Tanner.....	737
Meador III, Guy	307
Meng, Jiayou	629
Menon, Narayan	1560
Mersch, Todd.....	1560
Mezhoudi, Mohcene	1648
Migueluez, Phil.....	585, 894
Moroney, Paul	1486
Mukhopadhyay, Amit.....	1460
Murphy, Arnold	708
Mutalik, Venk.....	585
Naqvi, Salman	1206
Nickel, Ken.....	708
Noll, Kevin	914
Ochoa, Fernando Rodrigo.....	1140
Ong, Ivan	1787
Ovadia, Shlomo	250
Pavlich, Bryan	807
Peterka, Petr.....	1747
Pinckernell, Nicholas.....	536
Pralle, Darren.....	9
Prodan, Richard S.....	1375
Ramos, Hugo Amaral	1059
Ravisankar, Arun.....	1519
Ray, Indrajit.....	1774
Righetti, Claudio.....	1140
Romano, Carlos Germán Carreño.....	1140
Rothschild, Keith Alan	307
Santitoro, Ralph	151
Schauer, Paul E.....	325
Schnitzer, Jason K.	325
Scriber, Brian A.....	1765
Shapiro, Joseph.....	1434
Sharon, Hadar	368
Six, Erwin	162
Smargon, Dave	708

Spear, Gregory.....	124
Stengrim, Chris.....	1008
Stevens, J. Clarke.....	1581
Sun, Wendell	877
Sundaresan, Karthik.....	442, 1109
Sundelin, Andrew	344
Szczesniak, Mark.....	197
Tang, Benjamin Y.....	1648
Tang, Haibin	629
Tavares, Joseph.....	1059
Thottian, Jenson.....	250
Touze, David	1606
Troll, David	470
Tyre, Jeffrey	877
Ulm, John	1059, 1080, 1165
Vale, R. J.	1419
van de Kerkhof, Leon	1606
van Niekerk, Riebeeck.....	737
Vanderstraeten, Hans.....	162
Viorel, Dorin.....	807
Virag, David	807
Wang, Jing	1008
Wellen, Jeroen	1460
Westervelt, Egbert	1797
Widrevitz, Ben.....	1080
Williams, Tom.....	612
Yang, Wenle	209
Yarborough, Sean	91
Yates, David	1738
Yin, Zheng.....	238
Yuan, Brian.....	1804
Zhao, Zhuo	629
Zheng, Ruobin	209
Zhu, Jay	442
Zimmerman, Ron.....	355