

THE COMPLETE
TECHNICAL PAPER PROCEEDINGS
FROM:



A Deep Learning Approach for Detecting RF Spectrum Impairments and Conducting Root Cause Analysis

A Technical Paper prepared for SCTE by

Kevin Dugan

Data Scientist

Comcast

1800 Arch Street, Philadelphia, PA 19103

719.493.2600

kevin_dugan2@cable.comcast.com

Justin Evans

CoreTech

Comcast

1800 Arch Street, Philadelphia, PA 19103

719.493.2600

justin_evans@comcast.com

Maher Harb

Director, Data Science

Comcast

1800 Arch Street, Philadelphia, PA 19103

215.990.8376

maher_harb@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Background	4
2.1. Proactive Network Maintenance (PNM).....	4
2.2. FBC Data Characteristics.....	4
2.3. Downstream Wave Impairments	5
2.4. Network Topology as a Graph	6
3. Model Architecture	6
4. Model Training.....	7
4.1. Labeling.....	8
4.2. Transformations	8
4.2.1. Full Spectrum Smoothing.....	9
4.2.2. Truncate	9
4.2.3. Binning	9
4.2.4. Guard Band and Pilot Removal	10
4.2.5. Vacancy Removal	10
4.2.6. Final Data Form	11
5. Model Performance	11
5.1. Model Predictions.....	12
5.2. Model Improvement	13
6. Root Cause Analysis	14
6.1. Methodology.....	14
6.2. Annotating the Plot.....	15
6.3. RCA Results per Impairment	15
7. Conclusion.....	18
Abbreviations	18
Appendix	19
Bibliography	19

List of Figures

Title	Page Number
Figure 1 – An Example of a Normal Frequency Response	5
Figure 2 – Examples of Wave Impairment Frequency Responses.....	6
Figure 3 – The Neural Network Architecture.....	7
Figure 4 – The Labeling User Interface for Assigning Impairments	8
Figure 5 – Full Spectrum Smoothing	9
Figure 6 – Spectrum Truncated to Downstream	9
Figure 7 – Preprocessed Spectrum Binned to 2,000 Values.....	10
Figure 8 – Preprocessed Spectrum with Guard Bands and Pilots Removed	10
Figure 9 – Preprocessed Spectrum with Vacancies Removed.....	11
Figure 10 – Confusion Matrix for Resonant Peaks and Standing Waves.....	12
Figure 11 – Classification Examples (True Positives, False Positive, False Negatives).....	13
Figure 12 – Model Recalibration Cycle	13
Figure 13 – Standard F1 Score Calculation for Network Elements	14
Figure 14 – Color Scheme Legend for RCA Graph Plots	15

Figure 15 – Resonant Peak RCA Graph Plot	16
Figure 16 – Standing Wave RCA Graph Plot.....	17
Figure 17 – Suck-out RCA Graph Plot.....	17

List of Tables

Title	Page Number
Table 1 - Example Binary Classification	7
Table 2 – Model Performance Metrics	12

1. Introduction

In proactive network maintenance (PNM), a goal of downstream spectrum analysis is to identify customer devices with impairments in RF (Radio Frequency) spectrum. Identifying the type of impairment is useful in determining the general causation, which is leveraged to infer the geolocation of the impairment's origin in the network. This paper describes the implementation of an automated, end-to-end solution that analyzes millions of sets of spectra and delivers PNM opportunities for the organization.

Utilizing wideband RF frequency response data from full-band captures (FBC), spectral impairments in the downstream are detected using a signature matching algorithm implemented as a 1-D convolutional neural network (CNN). The signature matching algorithm evaluates the set of spectrum data for customer devices on the cable plant. Certain impairments that originate in the cable plant, such as resonant peaks that can occur at an amplifier, impact the RF signal for multiple customers further downstream of the impacted component. One of the goals is to identify when multiple customers are experiencing the same signature, pointing to an issue in the network that can be resolved without individual visits to affected customers.

The impacted modems feed into a root cause analysis (RCA) algorithm that overlays the impaired devices onto a graph representation of the network topology. Through methods rooted in graph theory, the RCA algorithm narrows down the geolocation and system component(s) for the probable network device(s) where the issue originates. The impact of this workflow is an automated capability to identify not only customer-impacting issues, but also opportunities to proactively resolve issues before they become impacting across the entire network.

2. Background

2.1. Proactive Network Maintenance (PNM)

Utilizing full band capture data for the purpose of PNM initiatives to detect impairments in the downstream frequency response has been discussed in significant detail in multiple preceding works. Those contributions have directly influenced this automated system, by serving as a general roadmap for the planning, design, and implementation. The scope of this document starts with these works as the base knowledge and describes an automated system to enrich PNM opportunities [1-10].

2.2. FBC Data Characteristics

The basic unit of data for this system is the FBC of RF spectrum for customer devices where each spectrum sample consists of 8,704 values spanning 6 MHz to 1026 MHz. Spectrum features important to this work include downstream SC-QAM channels, guard bands, vacant spectrum, and pilots. Figure 1 illustrates a sample of a normal frequency response from a modem that maintains consistent power for occupied spectrum at the appropriate levels (approx. -17.1 dBmv).




Figure 1 – An Example of a Normal Frequency Response

Downstream SC-QAM channels are defined as 6 MHz wide spans separated by guard bands. Along with pilots, guard band power values are not particularly useful in detecting impairments in many cases. Vacant spectrum represents a span where no services occupy one or more channels and can be identified programmatically by the power values and channelizing the spectrum. By reducing the spectrum only to occupied channel signal power, the signatures become more amenable to machine learning techniques that detect anomalous spectra.

2.3. Downstream Wave Impairments

Certain impairments in downstream SC-QAM channels can be identified in the FBC. For the scope of this implementation, the focus is on four impairments – standing waves (amplitude ripples), water in the cable, resonant peaks, and suck-outs. These patterns will be referred to as the ‘wave’ patterns. Figure 2 shows examples of each type of wave impairment. Within each plot, the solid green line at -17.1 dBmV and solid red line (bottom line) at -33 dBmV represent an appropriate range for frequency response values of occupied spectrum.

Standing waves, caused by impedance mismatches, are periodic in nature and generally extend across the entire downstream spectrum (Fox, et al., 2021). When water is introduced to the cable, an aperiodic wave is produced due to random attenuation and may also be associated with a negative tilt (Fox, et al., 2021). Resonant peaks are significant, narrow spikes in the spectrum caused by any number of reasons (i.e., cold solder joints or loose modules) on network devices (Cable Television Laboratories, Inc., 2016). A suck-out is represented as “a concave notch with sinusoidal boundaries with attenuation in amplitude/power caused by impedance mismatches evenly distributed through the network” (Cable Television Laboratories, Inc., 2016).




Figure 2 – Examples of Wave Impairment Frequency Responses

2.4. Network Topology as a Graph

A convenient way to analyze network topology is through graph theory. Comcast has mapped the “access network from the CMTS to the customer premise equipment (CPE)” into a graph structure “while incorporating all of the physical and logical elements that form part of the network” (Harb, Subramanya, Narayanaswamy, Walavalkar, & Rice, 2021). The graph facilitates the application of algorithms, such as lowest common ancestor (LCA), to network elements and their attributes. For example, clustering RF impairments on the graph gives the organization comprehensive knowledge about the network elements involved and a view of the common experiences amongst multiple customers. In turn, this knowledge is leveraged to deploy the correct resources to a specific physical location for resolution. In combination with network monitoring tools, an opportunity arises to automate the workflow from the current manual process.

3. Model Architecture

The adopted neural network for classifying RF impairments is a four-layer CNN that makes binary classifications for each of the wave patterns. Figure 3 represents the architecture diagram for building the pattern detection data model. Each of the 1-D convolutional layers uses a kernel size of five and a rectified linear unit (ReLU) activation. The convolutional layers are followed by a down sampling operation via max pooling. The max pooling operation calculates the maximum value in each section of the feature maps, pointing to the most present features. The increasing number of filters, as the CNN grows in depth, is attributed to the larger number of pattern combinations in each subsequent layer and using an increased number of filters allows the capture of more abstractions from the signal data.

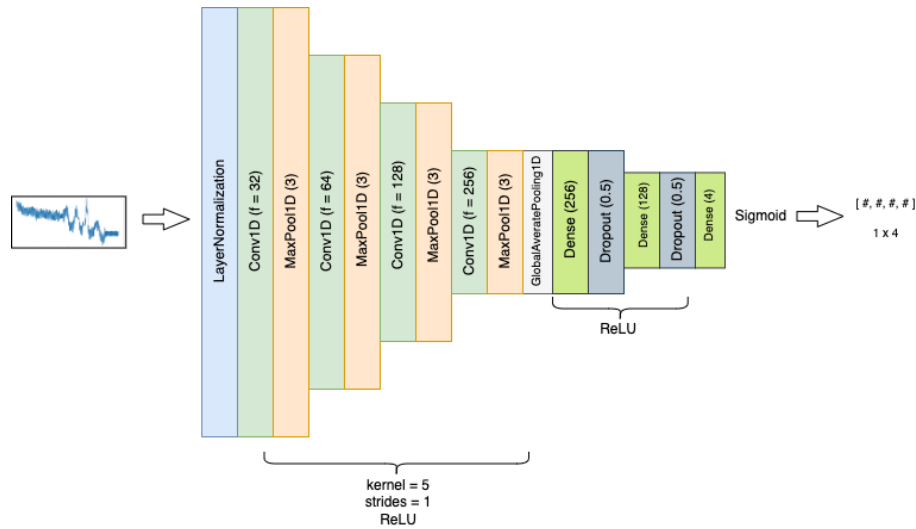


Figure 3 – The Neural Network Architecture

Multiple impairments may manifest on a single capture of spectrum. For example, for water to get into the cable, it needs an access point, such as a nick/chew in the cable, which would commonly cause standing waves (Fox, et al., 2021). Therefore, in certain cases it is difficult to differentiate among them, and their classification may impact the type of maintenance that would be referred to resolve the issue. It is more common to observe water in the drop cable or tap near the home; however, it may be observed in components for a larger group of customers and thus may be an issue that requires a network technician (Fox, et al., 2021). Being able to differentiate the impairments guides downstream decision-making processes.

Instead of structuring the model to calculate a single classification per spectra, the model predicts a probability for each type of wave. Taking this approach gives additional opportunities for analyzing the signatures of wave types within any given group and gives insights that improve the root cause analysis. Table 1 demonstrates possible classifications of multiple impairments.

Table 1 - Example Binary Classification

Sample Identifier	Water	Standing Wave	Resonant Peak	Suck-out
A	T	F	F	F
B	T	T	F	F
C	F	F	F	T
D	F	F	T	T

4. Model Training

The training dataset consists of 3,170 samples of labeled spectra from a population of 10,000 sets of spectra. The validation dataset consists of 500 hold-out samples. The labels were manually entered by a group of subject matter experts over the course of four weeks. The full collection of impaired and non-impaired spectra was acquired from an existing threshold-based detection algorithm that served as the basis for addressing this problem with machine learning methods.

Due partly to the limited training dataset and available resources, applying transformations on the data to reduce the noisy data points allows the CNN to learn appropriate classifications under the constraints. A downside of the transformations is the extended processing time for the feature extraction portion of the pipeline, processing millions of spectra for each iteration in production. As part of the model re-training cycle and growing the training dataset from validated classifications, we are optimistic that the preprocessing steps will be reduced in the future.

4.1. Labeling

To derive accurately labeled data, SMEs were presented with a random selection of both impaired and non-impaired sets of spectra via a user-interface (UI) adapted from a CableLabs initiative. Each spectra sample could be assigned any number of 13 labels. Figure 4 is a screenshot of the UI with a sample containing a standing wave. The two horizontal lines at -17.1 dBmv and -33 dBmv serve as visual indicators of the tolerable range in DOCSIS protocol to facilitate user interpretation.

At least two different users would be presented the same frequency response, with the usable training data samples meeting the criteria that more than one user assigned the same label. While this approach expedites much of the label validation effort, the overall number of samples is reduced because of requiring multiple labels from different users. Based on initial modeling experimentation, it was discovered that high-quality labels would be more useful than a few thousand additional labels whose quality was not checked as stringently. A follow-on experiment determined that by adding samples with a single label degraded the model's classification capability.



Figure 4 – The Labeling User Interface for Assigning Impairments

4.2. Transformations

The transformations applied to the raw data fall into one of two categories – data reductions and signal smoothing. Reduction logic reduces the number of data points per spectrum, while signal smoothing methods reduce the noisy parts of the spectrum's signal including guard bands, pilot signals, and vacant spectrum.

4.2.1. Full Spectrum Smoothing

An initial smoothing algorithm is applied to the spectrum based on the channelization of power values. The static nature of downstream SC-QAM channel characteristics enables a smoothing algorithm that considers the signal between guard bands but leaves the guard bands intact. Since the data source for the smoothed channel data is from a visual tool, it maintained spectral characteristics for presentation in a UI.




Figure 5 – Full Spectrum Smoothing

4.2.2. Truncate

The impairments under investigation all primarily reside in the downstream portion of spectrum. As a standardized transformation for all spectra, the usable spectrum for model training is defined as the span from 113 MHz to 748 MHz and may include OFDM spectrum values. As this applies to all spectra straightforwardly, any downstream services above 748 MHz are not considered for the wider bandwidth devices. Each sample in the transformed dataset now contains 5,420 values of smoothed spectrum.




Figure 6 – Spectrum Truncated to Downstream

4.2.3. Binning

As a further reduction of data volume, the samples of 5,420 values each are placed into 2,000 evenly spaced bins. For each bin, the mean of the values inside becomes the updated spectrum values. This not only reduces volume, but also irons out excess data points that ultimately can be supplanted without degradation to the primary patterns in the spectrum.




Figure 7 – Preprocessed Spectrum Binned to 2,000 Values

4.2.4. Guard Band and Pilot Removal

In early modeling experiments with labeled data, the data model would misinterpret guard band and pilot artifacts as characteristics of impairment types in many cases. Since this created a significant number of false positive results, it became necessary to smooth the values at these locations along the spectrum.

To remove the artifacts, the values are replaced with averaged values of the power immediately around the artifact – essentially creating a short linear regression line. A second approach to this for guard bands specifically would be using the channelized characteristics to replace the spectrum values. Both have proven to work well, and both achieve the desired transformational outcomes.




Figure 8 – Preprocessed Spectrum with Guard Bands and Pilots Removed

4.2.5. Vacancy Removal

The final transformation removes the spans within spectrum considered to be vacant – meaning there are no services in that span for any number of reasons. Vacancies are programmatically detected using thresholding techniques with the spectrum values being updated similarly to guard band and pilot artifacts. A linear regression is calculated between where the vacancy starts and ends; this matches the general trend of the spectrum at that location.




Figure 9 – Preprocessed Spectrum with Vacancies Removed

For very wide vacant spectrum, this approach may introduce a signature that resembles characteristics of one or more impairments. While the false positive rate is very low for these situations, discovering improved vacancy removal algorithms remains an important task. Again, the channelized spectrum data provides another strategy to identify vacant spectrum with high confidence and is the next step in this evolution.

4.2.6. Final Data Form

The final data form consists of 2,000 values, with noise removed for the model while retaining the signature's primary characteristics. From a scaling perspective, the transformations facilitate faster model training and inference operations, but at the expense of requiring more processing resources prior to involving the model.

The training data was lastly augmented by simply reversing the order of each sample, doubling the number of samples that retain the same pattern signatures but in different locations. Prior to training on the CNN, the data were passed through a normalization operation to improve consistency.

5. Model Performance

With a limited set of data to train the model, a 5-fold cross-validation architecture is implemented for estimating its generalization capabilities. The top performing cross-validation model is selected based on validation and training loss results, and then re-trained on the complete set of training data.

The model's ROC-AUC curve is 90% or greater for both the test data and for the hold-out validation dataset. The most confident classifications are the resonant peak and water signatures. Standing waves are the most difficult impairments to classify as certain signatures are similar to water or resonant peak signatures in some cases. Suck-out misclassification occurs primarily when artifacts remain from the vacancy removal process in which the leading or trailing edge was not properly cleared.

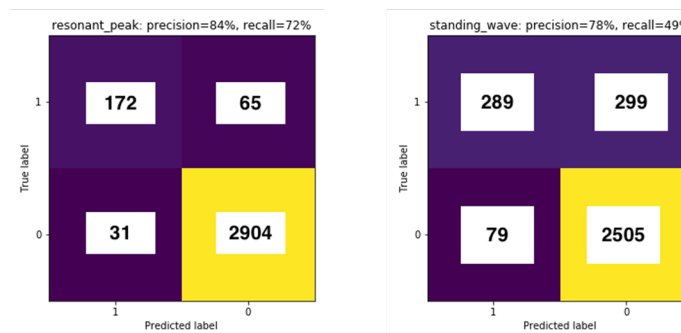


Figure 10 – Confusion Matrix for Resonant Peaks and Standing Waves

To calculate model lift, a dummy model (all samples labeled as no impairment) was evaluated in which the accuracy ranged from 81% for standing waves to 93% for resonant peaks. The impairment detection model outperformed the dummy model and created significant lift for each impairment in consideration.

Table 2 – Model Performance Metrics

Impairment	Training / Test AUC	F1 Score	Dummy Model Acc	Lift	Validation AUC
Resonant Peak	98%	0.78	93%	12.1	98%
Standing Wave	90%	0.60	81%	4.1	91%
Water	97%	0.74	89%	7.6	97%
Suck-out	94%	0.70	87%	6.7	91%

5.1. Model Predictions

Figure 11 highlights an example from each impairment with a correct prediction (TP), a false positive prediction (FP), and a missed prediction (FN) – reading column-wise. The red text in the middle row indicates the label applied to the sample during the labeling process described above.

The model performs exceptionally when the impairment type is straight-forward and mostly follows the definitions of the impairment. The model is less-confident when a frequency response indicates the presence of complex waves – that is, when multiple impairments or other conditions show characteristics of multiple anomalies. An example is the FP sample for standing waves shown in Figure 11. The frequency response has strong characteristics of a standing wave and weaker characteristics of a water wave, giving rise to a complex wave type. Since water waves are akin to standing waves, an assessment of this spectra is that there is moisture present in a standing wave, but not enough to induce drastic random attenuation. This determination impacts the root cause analysis algorithm, explained later in this document.

Another factor influencing model behavior is the volume of data and the quality of labels. The limited dataset lacks signature diversity, making it difficult for the model to interpret frequency responses that are considered complex. The labeled data also contains contradictory labels for similar-looking samples, demonstrating that interpretations of spectra can differ among individuals.




Figure 11 – Classification Examples (True Positives, False Positive, False Negatives)

5.2. Model Improvement

To obtain a larger training dataset for model recalibration moving forward, a system was established that collects random samples of predictions from each pipeline iteration, feeds them into a UI where a user validates the prediction, and the sample gets added to the training data. Figure 12 represents the workflow for retraining the model with additional samples.

The diversity in the initial training dataset was limited due to the low volume of available samples. To correct the model for any shortcomings related to this, obtaining validated predictions is critical for growing the training data. Additionally, the validation system can be filtered to a particular impairment so that any imbalances in the training data may be addressed by validating more quality samples of specific impairments.

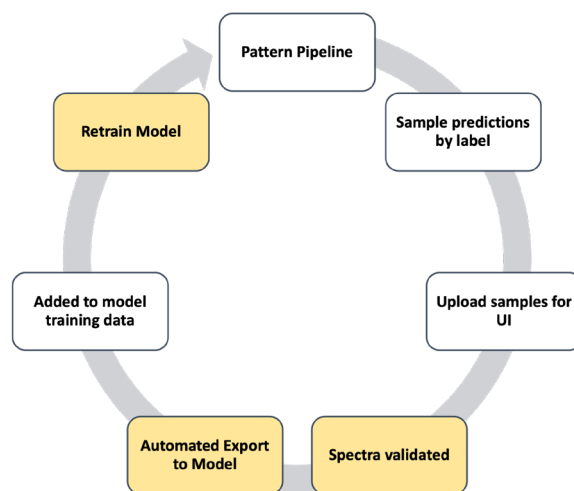


Figure 12 – Model Recalibration Cycle

In the validation process, a user can observe the original frequency response in addition to the data passed through the model. This has been helpful in catching edge cases with the preprocessing in which unwanted artifacts remained under certain conditions. In the vast preponderance of cases with an unexpected prediction, the artifact had a direct impact on the visual representation of the sample as it passed through the model.

6. Root Cause Analysis

An RCA algorithm was developed that leverages the graph representation of the network topology. Conceptually, the RCA examines a group of impaired modems within the context of seeking common network elements in the topology and calculates probabilities that each common network element may be the origin of the impairment. The algorithm surfaces results of the RCA to network monitoring tools, so that as PNM opportunities arise, the appropriate action can be taken to improve network conditions and reliability.

6.1. Methodology

The inputs into the algorithm consist of the following primary elements:

1. Type of impairment;
2. Impaired device list; and
3. Total device list in the grouping that reported frequency responses (a fiber node, for these purposes).

The graph of the topology for the node is extracted and a lowest common ancestor (LCA) search is performed for the impacted devices. The LCA makes different determinations on network elements depending on the type of impairment. Only amplifiers are considered when evaluating resonant peaks. For water in the cable, only customer drops are aggregated – meaning each LCA will point to a specific customer location and not a network element. Standing waves and suck-outs do not restrict which network elements are considered.

Standard F1 scores are calculated for each of the network components available in the topology as shown in Figure 13.

$$Precision = \frac{\# \text{ impaired devices on network element}}{\# \text{ devices reporting spectra}}$$

$$Recall = \frac{\# \text{ impaired devices on network element}}{\text{total } \# \text{ impaired devices}}$$

$$F1 \text{ Score} = \frac{(2 * Precision * Recall)}{(Precision + Recall)}$$

Figure 13 – Standard F1 Score Calculation for Network Elements

The LCA scores are ranked in descending order, with the possibility that multiple vertices (network elements) have the same score. This occurs in situations in which a sequence of vertices has the same number of impaired devices and the same population of devices that reported spectra – i.e., when multiple networks elements are connected with no additional modems in between them. To break the tie, the element with the longest path back to the CMTS is selected, as that element is the one closest to the impacted devices.

6.2. Annotating the Plot

When plotting the data for visual inspection, both physical and logical network elements are represented using the color scheme shown in Figure 14 (left). The figure also describes the highlighting scheme that represents the status of individual customer devices.

Vertex Class		Highlighting	
●	clamshellRx (Node)	a	Impaired Spectrum (Resonant Peak)
●	RfActivePort (Bus-leg)	a	Missing Spectrum
●	RfActive (Actives – i.e., Amplifier)	a	Normal Spectrum
●	RfPassive (Passives – i.e., Splitter)	a	Root Cause
●	RfTap (Taps)		
●	device (Modems, STBs)		

Figure 14 – Color Scheme Legend for RCA Graph Plots

The device highlighting provides visual clarification for where the impairments exist within the larger collection of devices available in the topology. Not all devices report spectrum, and since the LCA scores consider only the reporting devices, identifying them on the plot assists in validating the RCA results. High scoring vertices will contain few instances of normal spectrum and have more devices that are impaired or did not report spectra. The lowest common ancestor (highest scoring vertex) is highlighted in a unique color and has an annotation attached that describe the type of network element and the unique identifier.

6.3. RCA Results per Impairment

The initial problem statement for developing the RCA algorithm using the network topology was an issue commonly seen in amplifiers that causes resonant peaks to present in FBC data. For this reason, the RCA algorithm only calculates the score for each of the amplifiers found in the node's topology. In Figure 15, the modems impacted by resonant peaks are clustered on the left side of the graph (highlighted in pink), with no modem in the cluster reporting having normal spectrum. Therefore, the amplifier highlighted in the figure has an RCA score equal to the maximum possible.




Figure 15 – Resonant Peak RCA Graph Plot

Standing waves are a result of impedance mismatches caused by damaged cables, breached cable jackets, improper connections, and animal chews, to name a few. Standing waves may be manifested either in the plant or in a customer's home. The RCA algorithm considered only the standing wave events that likely originate in the plant or at a multiple dwelling unit (MDU) by considering a minimum number of impacted devices. In Figure 16, the highest scoring vertex was a passive network device under which the preponderance of modems showed impairments alongside only a few with normal frequency responses.

The topology edge data includes properties about cable lengths between vertices. While not yet implemented in this system, the capability exists to calculate the length to the voltage reflection in the cable from a starting point. With the known length, it is then possible to further refine the LCA calculation by considering only common network elements with a minimum cable length of the known distance to the reflection in the line.




Figure 16 – Standing Wave RCA Graph Plot

Suck-outs are commonly caused by “mechanical or grounding issues in active or passive network elements such as seizures, connectors, lids, or fittings” (Cable Television Laboratories, Inc., 2022). Similar to the diagnosis of standing waves and resonant peaks, a minimum number of impacted modems were required to pass through the RCA. The identified vertex in Figure 17 was a tap under which some modems had normal frequency responses, and some showed the suck-out impairment.

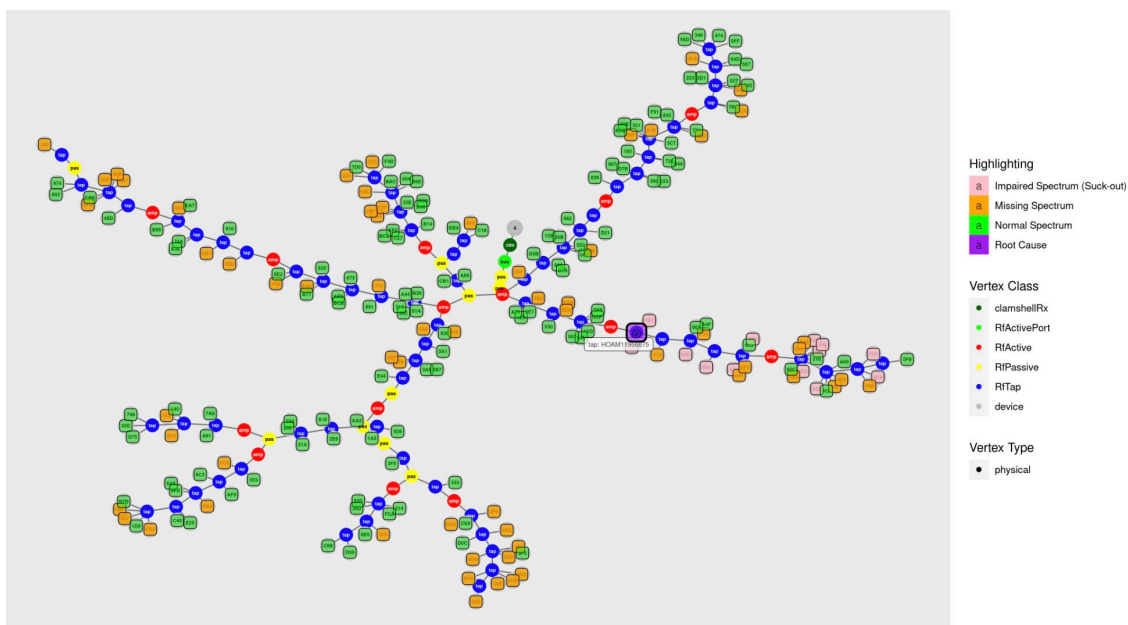


Figure 17 – Suck-out RCA Graph Plot

Water waves are most identified in customer drop cables and taps (Fox, et al., 2021). The RCA algorithm here only considers the drop locations to group the devices in the home/MDU into a single event. This was an intentional implementation decision while the clustering algorithms for water waves remained under development and will be used in future iterations.

7. Conclusion

Frequency responses from FBC data have been used extensively to identify certain impairments that impact network reliability. Additionally, multiple initiatives have surfaced deep learning models designed to automate detection of signal impairments in different types of signals. This implementation takes an additional step by automating the triangulation amongst network elements to narrow down the origin of the impairment through a network topology graph representation.

The deep learning model described here is but one form the model can take. In the development of this system, multiple models that produced quality results were evaluated. This detection model is an improvement over currently known threshold-based algorithms, often finding anomalous frequency responses where threshold-based algorithms miss them. The advantage to the previous algorithm is that it reliably finds moderate to severe cases; however, it is unable to differentiate among some wave types so they are categorized under one label. This impairment detection model not only differentiates the waves, but also has demonstrated the capability to detect both obvious and subtle impairment signatures.

By leveraging network topology as a graph, more advanced analytics are possible to understand the scope and impact of impairments on collections of modems. The lowest common ancestor algorithm is one of several forms of analysis that can be done. Through an understanding of RF impairments and their causes, a root cause analysis performed on network topology generates confident assessments of where within the plant the issue originates. The impact is better visibility into network conditions and improved support for the technicians that deploy to resolve issues in the network and at customer's homes.

Abbreviations

CNN	convolutional neural network
CPE	customer premise equipment
dBmv	decibels relative to one millivolt
FBC	full-band capture
FN	false negative
FP	false positive
LCA	lowest common ancestor
MDU	multiple dwelling unit
OFDM	orthogonal frequency-division multiplexing
PNM	proactive network maintenance
RCA	root cause analysis
ReLU	rectified linear unit
ROC-AUC	receiver operator characteristic – area under curve
RF	radio frequency
SC-QAM	single-carrier quadrature amplitude modulation
SME	subject matter expert
TP	true positive
UI	user interface

Appendix

Network Topology

Network topology can be accurately described as a directed acyclic graph or DAG. This allows for easier and more efficient algorithms to be used without the need to cycle check. A DAG is a special type of graph that allows traversal in one direction without the chance for a connection to “backtrack” to a previous vertex. The reason for this is plant topology does not loop back on itself and upstream edges will always direct to the devices closer to the CMTS in the plant.

The two primary plant topologies revolve around having an analogue CMTS doing all the work vs having a distributed access architecture where the load is balanced across the node. These two topologies do pose a challenge when creating the topology map as different devices have similar functionality yet different naming schemes. Nodes with vCMTS (virtual Cable Modem Termination Stations) have no amplifiers and scale incredibly well leading to easier changes in the plant.

Creating the graph that this work depends on took many teams a lot of effort in order to overcome the challenges required. In order to flesh out the graph, multiple data sources needed to be pulled from and synced. To achieve this the team used a property-graph architecture as well as Apache Tinker pop to perform the aggregation queries, look-ups ext. The graph database ROCI is still not fully mature but has proven extremely valuable in many projects including our own.

Bibliography

- Cable Television Laboratories, Inc. (2016). *PNM Best Practices: HFC Networks (DOCSIS 3.0)*.
- Cable Television Laboratories, Inc. (2022). *Primer for PNM Best Practices in HFC Networks (DOCSIS 3.1)*.
- Ferreira, J., Harb, M., Subramanya, K., Santangelo, B., & Rice, D. (2020). Convolutional Neural Networks for Proactive Network Management. *SCTE 2020 Fall Technical Forum*.
- Fox, K., Rupe, J., Willimas, Tom, Zhu, J., Hranac, R., Zedan, N., . . . Wolcott, L. (2021). Water Can Run, But It Can't Hide. *SCTE 2021 Fall Technical Forum*.
- Harb, M., Subramanya, K., Narayanaswamy, R., Walavalkar, S., & Rice, D. (2021). How Network Topology Impacts Rf Performance: A Study Powered By Graph Representation of the Access Network. *SCTE 2021 Fall Technical Forum*.
- Hranac, R., Campbell, C., Fish, R., Kolze, T., Kristoffersen, E., Medlock, J., . . . Wolcott, L. (2020). Full Band Capture Revisited. *SCTE 2020 Fall Technical Forum*.
- Virang, D., Chari, S., & Kraiman, S. (2021). Machine Learning for RF Impairment Classification. *SCTE 2021 Fall Technical Forum*.
- Vishnyakova, A., Mahajanam, R., O'Dell, M., Merkle-Tan, M., Hay, C., & Pham, L. (2021). Right Technician at the Right Time. *SCTE 2021 Fall Technical Forum*.

Volpe, B., & Ottlik, B. (2021). Machine Learning and Proactive Network Maintenance: Transforming Today's Plant Operations. *SCTE 2021 Fall Technical Forum*.

Wolcott, L., O'Dell, M., Kuykendall, P., Gopal, V., Woodrich, J., & Pinckernell, N. (2018). A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact. *SCTE 2018 Fall Technical Forum*.

Zhu, J., Sundaresan, K., & Rupe, J. (2020). Proactive Network Maintenance using Fast, Accurate Anomaly Localization and Classification on 1-D Data Series. Louisville: IEEE. Retrieved 2022, from <https://arxiv.org/abs/2007.08752>

A Demuxed State of Mind: Transforming Content Ingestion and Distribution in an OTT World

A Technical Paper prepared for SCTE by

Yasser Syed, PhD

Distinguished Engineer
Comcast

Comcast Center 1701 JFK Blvd. Philadelphia, PA 19103
1-215-286-1700.
yasser_syed@comcast.com

Alex Giladi

Fellow
Comcast

Comcast Center 1701 JFK Blvd. Philadelphia, PA 19103
1-215-286-1700.
alex_giladi@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Production Media Workflows.....	3
3. Carriage of Media in Distribution Service Workflows.....	4
4. Real-Time Considerations in Media Workflows.....	6
5. Adapting to Newer Technologies: IP Interfaces/ Cloud Processing.....	7
6. Content Processing and Content Experiences.....	9
7. Demuxing Content for Improved Workflows.....	10
8. What Further Work is Required?.....	11
9. Conclusion.....	12
Abbreviations.....	13
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 – Production Workflow.....	4
Figure 2 – SDI Frame.....	4
Figure 3 – MPEG-2 TS Service Distribution.....	5
Figure 4 – MPEG-2 TS Muxplex Composition.....	5
Figure 5 – Adaptive Streaming HTTP Service Distribution.....	5
Figure 6 – Stream of Frames.....	6
Figure 7 – Quality-Latency-Bandwidth Tradeoffs.....	7
Figure 8 – Factors in Playback.....	7
Figure 9 – IP Encapsulated SDI.....	8
Figure 10 – Adaptive Streaming using HTTP.....	8
Figure 11 – Roles of Cloud in Media Workflows.....	9
Figure 12 – Content Processing and Content Experience Processing.....	10

1. Introduction

Content is comprised of video, audio, and text or auxiliary information(e.g., closed captioning). An absence of any media component can lead to an incomplete media experience. A non-synchronized playout of the media components can lead to an incomprehensible media experience. Yet each media component has a separate creation, production, distribution encode and client decode path in its workflow chain with each step traditionally requiring all media components to be received together before processing such that complete and synchronized media playout is ensured.

The speed of each step is dependent on the processing of the slowest media component of the content. Processes can be both automated and manual with creative processes being usually slower than automated processes. Automated processes can still add latency to the workflow, dependent on the amount of compute power needed for operational processes or scene analysis and scaled to match expected quality output. And yet, with the introduction of cloud processing, this can be adjusted.

How can we create evolutionary changes to these workflows? It may be through adapting media workflows to be more like data workflows and taking advantage of better bandwidth and adjustable compute power. Through IP interfaces, increased network capacity, and volume-efficient cloud server processing, media workflows can be more efficient, and deliver better quality along with additional media experiences. But these improvements require us to separate (demuxed) media components to be handled with better optimizations in today's technology environment. Demuxing content allows for better handling of processing demands but reassembly is also important and should add additional information for resynchronization of media components in order to make the playout of the content feasible.

This paper will provide an overview of these processes in the context of traditional media workflows, what can be done today with present technologies, and what could be done in the future with newer technologies like over-the-top (OTT) delivery and cloud processing.

2. Production Media Workflows

Content origination starts with production and distribution workflows (see Fig 1) from capture to editing to distribution by service providers. This can be in the form of scripted content (e.g. movies and TV shows) or live events (e.g. a football game televised live). Some operations that happen are more automated such as chroma subsampling (4:4:4/4:2:2 to 4:2:0), bit depth reductions (16/12 to 10 bits), audio formatting, or lookup table (LUT) conversions for high dynamic range (HDR)/standard dynamic range (SDR) conversions. But there are a lot of creative operations happening here as well, such as editing, shading, translations, and subtitling, all of which can take longer but could also be parallelizable. Ultimately this is received by the service provider in the form of contribution linear feeds (e.g. see SCTE 277) or mezzanine assets. Multiple feeds or files may result from this process.

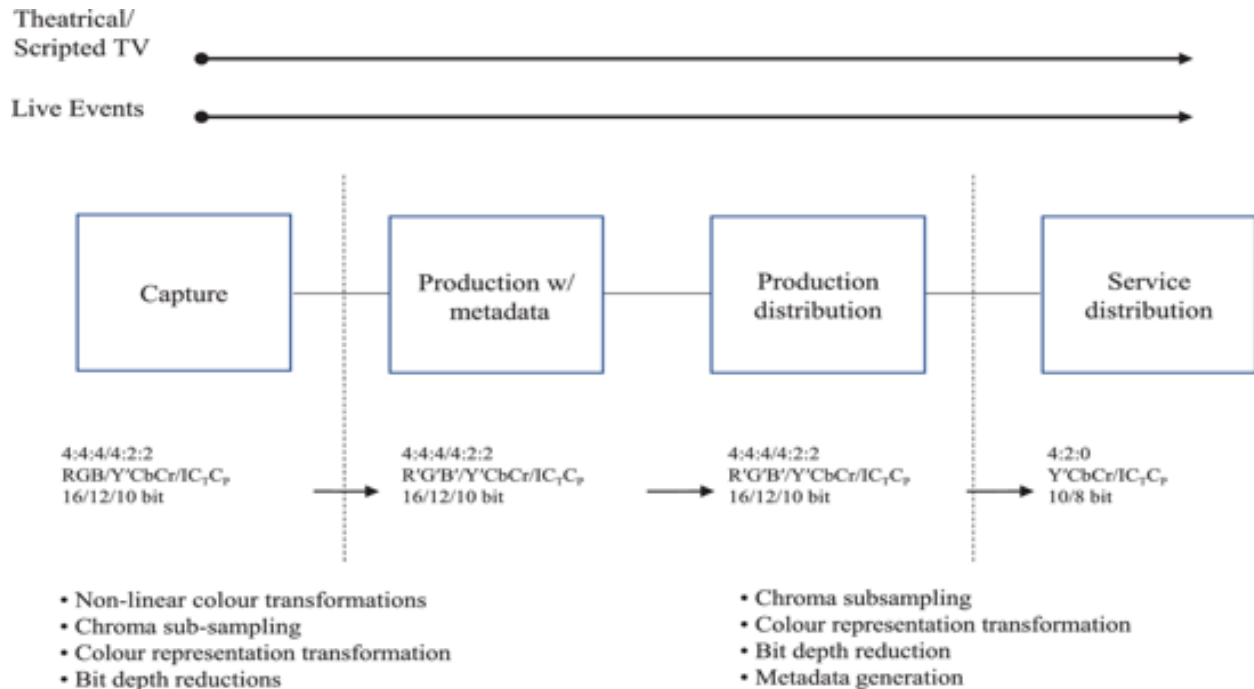


Figure 1 – Production Workflow¹

In terms of moving assets or feeds around in this domain, SDI (Serial Digital Interface) (see Fig 2) is widely used (e.g. SDI, HD-SDI, 3G-SDI, 12G-SDI) providing a baseband interface to allow for frame editable content to be distributed within a local area. An SDI frame allows for a timing frame to encapsulate video, audio (16 channels), closed captioning, and timecode within a serial stream.




Figure 2 – SDI Frame

The limitations in this type of setups are that content is only locally moveable and that the cabling infrastructure is designed specifically for moving SDI signals around to production and post-production sites. With the developments from ST 2022-6, which encapsulates SDI over IP, it facilitated production workflows to begin to move beyond the local plant.

3. Carriage of Media in Distribution Service Workflows

Once the studio passes content to the service delivery workflow, the content is mostly formed and it is more of a question of delivering the content in the right format for the client player for playout. The

¹ From ITU-T H. Supp 19/ISO 23091-4

workflow still follows the live and file-based workflows that are then encoded for distribution to create a decodable bitstream in which compression of content to utilize bandwidth efficiently is essential to operate within the capacity and scalability limitations of the distribution network. Most of the media operations are done on a single device on which the content is unwrapped, the components are processed, and finally the components are reassembled back into an integrated content format again to be carried on the network for device distribution.




Figure 3 – MPEG-2 TS Service Distribution




Figure 4 – MPEG-2 TS Muxplex Composition²

For traditional one to many broadcast QAM delivery over MVPD networks, an MPEG-2 transport stream is used to carry content (see Fig 3). The transport stream packetizes each media component into elementary streams, which are grouped together as a program stream. The MPEG-2 packetization allows for the media component packets of the same timeframe to be interspersed in the data stream (see Fig 4), such that a decoder with coded picture buffer can receive the packets and output a time-continuous stream while taking advantage of temporal compression strategies to reduce the bandwidth demands of delivering the content.




Figure 5 – Adaptive Streaming HTTP Service Distribution³

² From SCTE 223 – Adaptive Transport Stream

³ From IEEE ICIP Tutorial Ali Begen/Yuri Reznick Sept 19th, 2021

With recent adaptive streaming technologies that are used in OTT delivery, service functionality moved from a single encoding device to distributed functions (see Fig 5). Because OTT services deliver content over the internet using unicast internet protocols, the operations of encoding, packaging, placement, and delivery were spread across the network; additionally, media components did not have to be packaged together for delivery, but ultimately delivery still was constrained by the same decoder buffer restrictions on outputting a continuous stream of frames.

4. Real-Time Considerations in Media Workflows

Video is a series of frames displayed to the eye at a certain rate (see Fig 6), allowing a simulation of the what the brain would perceive in the real world. The rate to achieve this is known as frame rate (frames-per-second [fps]) and can vary from approximately 24fps for feature film content to 60fps for live production and most video sources. There are also some sources at 30 fps and new high frame rate video formats at 120 fps that are better at reproducing fast action.⁴ For real-time processes, frames need to be created, edited, produced, delivered, encoded, packaged, distributed, buffered, and then decoded before the next frame is outputted. Causes that will add time to the workflow would be compute time for media processing at different points in the chain, or network bandwidth needed to deliver bits across to different points in the workflow, or simply any creative handling of the content.




Figure 6 – Stream of Frames

At a constant frame rate 60 fps mode, successive frames need to be outputted every 16.7 ms (1/fps). This means that for a glass-to-glass delivery, all processes need to be completed within this time. Traditional mechanisms to achieve this are to introduce startup delay such that the player buffer could be filled before an initial frame would be outputted, then the duration of the buffer could extend the time that the next frame would need to be sent to the buffer to keep it filled. This adds latency to the video, but through approaches like this, continuous video can be achieved at a manageable bandwidth for the network.

Other ways to address real-time demands would be to reduce temporal compression by creating a simplified group of pictures (GOP) structure such as all I-frames or forward predicted pictures (FPP). Another way would be to increase the network bandwidth such that created content packets could be delivered faster than real-time; this would work more for file-based workflows and less for live feeds.

⁴ Typical framerates in use also includes fractional framerates such as 23.97, 59.94, or 119.89 which can add some complexities to timing and synchronization factors, but conceptually remains the same. For this paper, integer number frame rates examples will be used for simplicity while discussing concepts.




Figure 7 – Quality-Latency-Bandwidth Tradeoffs

But real-time processes are not the only factors to consider when delivering video. Depending on the service, the video delivered also must be of acceptable quality and of acceptable bandwidth for network distribution (see Fig 7). Increasing quality or bandwidth may result in increased delay and designing workflows must consider and balance all three of these factors in delivering content (see Fig 8).




Figure 8 – Factors in Playback⁵

To avoid stalls, the decode buffer must have new frames coming in at regular intervals. To have fewer resolution switching factors, managed bandwidth network connections or smaller bandwidth requirements must be maintained while handling more bits to deal with higher bandwidth or color space demands. The need to reduce latency factors also brings into play reduced player buffer capacity, which puts increased pressure on real-time computation demands.

5. Adapting to Newer Technologies: IP Interfaces/ Cloud Processing

Carriage of Media workflows as data across IP interfaces bring many advantages, and the content production ecosystem is gradually adopting new IP based approaches. For instance, in production workflows, the same source point can be sent to multiple end points in a local network so that editorial or

⁵ From IEEE ICIP Tutorial Ali Begen/Yuri Reznick Sept 19th, 2021

production tasks do not always have to be serially performed but could be performed in parallel (see Fig 9). Through SMPTE ST 2022-6, SDI video can now be encapsulated and carried through an IP distribution. A benefit of this is the reutilization of existing production equipment with only an SDI interface while using IP to carry the signals and IP routers to switch and distribute them. With SMPTE ST 2110, the IP carriage takes a further step to natively carry the video, audio, and data as independent essences that are sent separately, thus allowing for individual points to deal with specific media components without sending the entire content. With moving to a demuxed delivery of media components, a timing synchronization is now required to reassemble the media components into the complete content. In SMPTE 2110, this is achieved through using a Precision Time Protocol (PTP) timing mechanism.




Figure 9 – IP Encapsulated SDI

In adaptive streaming delivery, a similar evolution occurred from what initially was IP carriage of MPEG-2 TS streams (e.g. HLS and DASH). This evolved to utilizing an ISO base media file format (ISOBMFF) delivery that allowed each media component (independent essence) to be requested from a set of options (adaptationSets) and delivered separately using segment timelines as the method to synchronize media components (see Fig 10). Further refinements with CMAF allow segments to be broken up and delivered separately as chunks; this assists in providing lower latency delivery.




Figure 10 – Adaptive Streaming using HTTP

Cloud processing is now becoming utilized in many places for high performance computing, AI/ML analysis, or on-premise offloading. The modes of service in the cloud can be infrastructure as a service (IAAS), platform as a service (PAAS), software as a service (SAAS), or serverless. For content media workflows, cloud processing can operate in several areas for real-time captioning services, basic SDR/HDR conversion processes, scene detection for more efficient guided encoding, third party processing, or manifest manipulation. For better integration of cloud processing into content media

workflows, it requires more efficiencies in bandwidth workflows to handle even editable workflows and IP addressing to multiple source delivery points (see Fig 11). To make it more practical, compute latency, storage latency, and network latency need to be reduced. For compute latency, more pipelining and parallelization strategies as well as more powerful processors need to be available, depending on the processing task. For storage latency, better caching and reading/writing strategies already exist but need to significantly improve to reduce latency. Lastly, network latency needs to be improved through the analysis and adoption of new strategies in regard to compute placement (e.g. edge compute).




Figure 11 – Roles of Cloud in Media Workflows

6. Content Processing and Content Experiences

In content media workflows, processes can be categorized as automatic or manual (see Fig 12). Automatic processes typically can be done without manual work; examples in production are up/down sampling, tone mapping, and some speech recognition/captioning; while examples in distribution include such functions as transcoding and packaging. Manual processes require more hands-on work and can describe functions like editing, shading, audio description, or third person signing. With cloud processing and artificial intelligence (AI)/machine learning (ML), some of the manual tasks can be done or accelerated with assistance by trained AI/ML approaches. With post-production, developing a format that allows for IP distribution while allowing for frame accurate editing is needed to reduce the bandwidth demands and compute demands to allow for this. Through developments in I-Frame HEVC/VVC or JPEG-XS or J2K, future image/video coding techniques may address these needs in the future.




Figure 12 – Content Processing and Content Experience Processing

Another type of category is processing of existing content or adding to experiences of the content asset (see Fig 12 again). Processing of content type of operations does not change the version of the content but may change format, or appearance of the content. These can include adding more subtitle language tracks, or new dubbed languages in audio, or simply adding audio description tracks to the content. Most of these types of processes are more automated tasks rather than manual tasks. Content experience processing usually alters the version of the content and provides additional personalization of the content such as additional languages, alternate versions of content, adding announcers to sporting events, or adding audio descriptions to events. Many of these added experiences are manual processes but with new types of cloud processing, IP carriage, and bandwidth efficiencies, some of these tasks can reduce latency through assistance using these new technologies. An example of this that is now seen frequently is the real-time creation of subtitles from the audio track.

7. Demuxing Content for Improved Workflows

Demuxing content into separate media components can benefit media workflows in both production and distribution by reducing the amount of content that needs to be wrapped, transported, unwrapped, processed and re-wrapped many times in the overall media workflow. For production, demuxing can help with real-time operational processing demands such as real-time generation of closed captioning through implementation of an audio track or in translation an existing closed caption track. With third party signing, it may require generating a video proxy linear stream to send with the audio track. Content processing transformations (see Figure 12) can also be done offsite but requires some sort of frame accessible video format that is protected and is transportable across IP networks instead of the local network. In addition to processing content, information can also be gathered for scenes that can be beneficial for distribution transcoders to do higher quality encodes. Lastly, some additional experience tracks can be generated during this production process.

Using separate media components for processing provides some step up advantages to that triangle of quality, latency, and bandwidth considerations (see Fig 7). For bandwidth, there are fewer bits to be sent across IP networks to process media through the production and distribution workflows. For latency, the ability to send information in a parallel fashion earlier in the media workflow and the ability for the cloud to assign elastic processing power to what the task demands are strong advantages. Additionally, developing AI/ML techniques that may be continually refined would be helpful to get better quality

distribution encodes (including pre-processing techniques) at lower bit rates. Lastly, third party contributions to content assets and feeds for such content enhancements as home/away announcers, third party signing, or real-time indexing for more searchable content could be done across IP networks as well.

8. What Further Work is Required?

Some additional development that will be required – some of which already is happening – to better handle demuxed workflows falls into five categories:

- Automation in production and distribution workflows;
- Synchronization across media components;
- Integration of third party created media components;
- Placement of new content experiences across adaptationSets in manifests; and
- Playability of content assets and linear feeds in an OTT system.

To pull off media components of the content for separate processing, metadata is needed to automate job processes. This metadata, which can include metadata generated by AI and ML processes, can aid in other future processes in the workflow such as enhanced quality encodes. This information when generated or provided needs to be stored in a format that may originate in production but may be needed upon distribution. A format needs to be developed that can provide this time sensitive information in both the production and distribution domains and may need to be stored as a metadata track retrievable in the manifest.

For automation in production, an editable format of the media component needs to be created that can span across IP networks rather than be limited to a local LAN, at a reasonable bandwidth. Some formats being looked at are JPEG-XS, I-frame only AVC/HEVC/ VVC, or J2K. Audio Formats may not need to be compressed since audio file and feeds are smaller but may require video proxies to be sent along with them, but there is also a need also to carry some of the distribution signalling metadata for audio back to baseband carriage. For third party components such as signing, audio description, or subtitling, these type of files or feeds may need to be fed into the distribution workflow rather than returned to the production workflow.

Synchronization between media components is also a factor that needs development. In one area it is more conforming additional tracks to the content asset or file; this basically means each media component should have an aligned timeline that can be used to generate synced timelines on newly created assets. In the production domain and with the SMPTE 2110 format, the PTP timing format is used. In the distribution area, however, time is more aligned with segment timelines and alignment between both approaches needs to be realized. Another area is synchronization between audio and video; in OTT the video segment determines the edit point but there are instances in which the audio segment that goes along with the video may naturally be slightly lagging the video; in those cases that slight lag needs to be maintained.

To integrate third party media assets and linear feeds to the content, IDs need to be created and matched such that manifests can add these new adaptationSets and authentication systems can validate that content. URLs involving the retrieval of content need to consider content distribution network (CDN) design and whether one or more CDNs are involved. There also needs to be a way to identify tracks that are being processed versus tracks that are finished and retrievable by the client player. Gaps in timelines of media components need to be considered and handled as well so mechanisms such as content failover can be considered. Additionally, player behavior needs to be redefined such that new content experiences can be added to an existing media layout which may be added as the manifest is updated. Even when the

concept of a linear feed or a content file being complete no longer exists, the playability of the asset needs to be known.

9. Conclusion

Content is nowadays not a static asset or channel. It is evolving according to the expectations of the viewer and can involve multiple parties creating the overall content experience. More formats, better quality, multiple types of playback devices, and more ways to experience content is an expectation. The platform for media workflows is also evolving to incorporate faster and intelligent processing not limited to all work being done in a specific location. Media workflows for production and service workflows have traditionally been linear, but with the advents of higher bandwidth, IP interfaces for both uncompressed and compressed workflows, and access to cloud processing, linear workflows can evolve. Production efforts can be outputted directly into OTT feeds and into CDNs to be ingested by OTT players as they become aware of what new options they have to experience the content, and authorized third party contributors can also provide new experiences to the content without delaying the workflow. Cloud processing can accelerate the workflows by just adapting processor power to the job demand or by pre-processing analysis to aid in encoding of the content. But these evolutionary efforts do require some adaptation of the current system to make it easier to conform and synchronize separate media assets to the content and to create adaptationSets, IDs and URLs to incorporate these workflows and to make players aware of new choices in the manifest.

SCTE has a number of working groups that are already involved in these areas. In the Digital Video Services committee, there are two related working groups that are involved in these areas. Working Group 1 (Video/Audio) just recently published SCTE 277 (Linear Contribution Encoding Specification) which defines ingestion of signals that originate from production/post-production services. Additionally, WG1 also defines the Video and Audio codec streaming constraints for consumer distribution including recent modifications for adaptive streaming to consumers. On the consumer distribution side, WG7 (adaptive streaming /DASH) works on the suite of SCTE 214 specification which define manifest and segment constraints for HTTP IP delivery of content through MVPD networks.

Abbreviations

ADP	ancillary data packet
AI	artificial intelligence
AVC	advanced video coding
CDN	Content Distribution Network
CMAF	Common Media Application Format
FPP	forward predicted picture
fps	frame per second
GOP	group of pictures
HANC	horizontal ancillary data
HBI	horizontal blanking interval
HDR	high dynamic range
HEVC	High Efficiency Video Coding
HTTP	Hypertext Transfer Protocol
IASS	infrastructure as a service
ID	identification
I Frame	Intraframe
IP	Internet Protocol
ISOBMFF	ISO base media file format
J2K	JPEG 2000
JPEG	Joint Photographic Experts Group
LAN	local area network
LUT	lookup table
ML	machine learning
MPEG	Moving Pictures Experts Group
OTT	over the top
PAAS	platform as a service
PTP	Precision Time Protocol
QAM	Quadrature Amplitude Modulation
RAP	random access point
SAAS	software as a service
SCTE	Society of Cable Telecommunications Engineers
SDI	serial digital interface
SDR	standard dynamic range
SMPTE	Society of Motion Picture and Television Engineers
Url	uniform resource locator
VANC	vertical ancillary data
VBI	vertical blanking interval
VVC	versatile video coding

Bibliography & References

ISO/IEC 13818-1:2020, Information technology - Generic coding of moving pictures and associated audio information: Systems.

ITU-T H-Series Supplement 19| ISO/IEC TR23091-4 Usage of Video Signal Type Code Points

ANSI/SCTE 277 2022 Linear Contribution Encoding Specification

ANSI/SCTE 223 2018 Adaptive Transport Stream

SCTE 214-1 2022 MPEG DASH for IP-Based Cable Service Part 1: MPD Constraints and Extensions

ANSI/SCTE 214-4 2018 MPEG DASH for IP-Based Cable Service Part 4: SCTE Common Intermediate Format (CIF/TS)

Ali C. Begen and Yuriy A. Reznick, Advances in Multimedia Streaming: Algorithms, Standards, and Optimization Techniques IEEE ICIP Tutorial September 19th, 2021

SMPTE ST 2022-6: 2012 SMPTE Standard - Transport of High Bit Rate Media Signals over IP Networks (HBRMT)

SMPTE ST 2110-20: 2017 SMPTE Standard – Professional Media over Managed IP Networks: Uncompressed Active Video

SMPTE ST 2110-30: 2017 SMPTE Standard – Professional Media over Managed IP Networks: Uncompressed PCM Audio

SMPTE ST 2110-40: 2017 SMPTE Standard – Professional Media over Managed IP Networks: Ancillary Data

AWS Media Blog Part 1: Background and key benefits of SMPTE 2022-6 on AWS Elemental Live, Aug. 16th, 2021, <https://aws.amazon.com/blogs/media/awse-part-1-background-key-benefits-smpte-st-2022-6-aws-elemental-live/>

AWS Media Blog Part 1: Background and key benefits of SMPTE 2110 on AWS Elemental Live, Nov 3rd, 2020, <https://aws.amazon.com/blogs/media/part-1-background-and-key-benefits-of-smpte-st-2110-on-aws-elemental-live/>

A Necessary Journey Towards an AI-driven Operation

Telecom Argentina perspective

A Technical Paper prepared for SCTE by

Claudio Righetti

Chief Scientist

Telecom Argentina S.A.

crighetti@teco.com.ar

Mariela Fiorenzo

Tech Expert Scientist

Telecom Argentina S.A.

mafiorenzo@teco.com.ar

Horacio Arrigo

Tech Scientist

Telecom Argentina S.A.

hgarrigo@teco.com.ar

Table of Contents

Title	Page Number
Abstract	4
Content	4
1. Introduction	4
2. What it means to be AI-driven operation	7
3. Our first steps without AI using math, statistics and small data	8
3.1. Dimensioning of our VoD system	8
3.1.1. Erlang B statistical model	8
3.2. Characterization and impact of user behavior of OTT services	10
3.2.1. Measuring performance for decision making	10
3.2.2. Forecasting subscribers and traffic	11
3.3. IP traffic dimensioning based on patterns	12
3.3.1. Dimensioning tool based on simulations	14
3.3.2. COVID-19 and HFC traffic growth	14
3.4. Real-time analytics for IP video multicast	16
3.4.1. Multicast gain formula	17
3.4.2. K-means clustering applied to the selection of multicast channels	17
4. Evolving with AI and Big Data	18
4.1. Network capacity and Machine Learning	19
4.1.1. Principal Component Analysis (PCA)	19
4.1.2. Artificial Neural Network (ANN)	20
4.2. Optimizing video customer experience with Machine Learning	21
5. AI Ops	22
5.1. Ada	24
5.2. Customer claim prediction for HFC network	25
5.3. Intelligent agents and Autonomous Networks	26
6. Learned lessons	28
6.1. Define and communicate AI-driven operations	28
6.2. Workforce	28
6.3. Models, algorithms and Explainable AI	29
6.4. Cloud and Data/AI platforms	30
7. Next steps	30
Conclusion	30
Abbreviations	30
Bibliography & References	33

List of Figures

Title	Page Number
Figure 1 – Percentages of Network Activities that are automated, 2021	5
Figure 2 –From CSP to DSP	6
Figure 3 – Hype Cycle for Artificial Intelligence, 2022.	7
Figure 4 – Relationship between n° STB and Blocking (%) for different scenarios according streaming codification and quality.	9
Figure 5 - Monthly evolution and forecast of Netflix users among our subscribers.	12

Figure 6 - Linear Regression of the average bandwidth per subscriber. Blue line represents the adjusted and estimated current values. Green and orange lines, the forecast for the next two years.....	14
Figure 7 – Change in daily HFC traffic patterns.....	15
Figure 8 - Evolution and forecast for Downstream traffic, 2020.....	16
Figure 9 – Evolution and forecast for Upstream traffic, 2020.	16
Figure 10 - K-means clustering applied to the access frequency per channel by day for the Legacy system. Algorithm used to classify the signals between multicast and unicast.	18
Figure 11 - PC1 vs PC2 and its relation with traffic management.	20
Figure 12 – Our ANN scheme.....	21
Figure 13 – VMAF Algorithm.....	22
Figure 14 – AIOps cycle by Gartner.....	23
Figure 15 – AIOps automation path.	23
Figure 16 – Levels of Network AIOps functionality.	24
Figure 17 – Clustering and Forecast performed with Ada.	25
Figure 18 - High-level view of our project: inputs, outputs, outcomes and architecture.	26
Figure 19 – Intelligent agent.....	27
Figure 20 – Agent in AN source: Autonomous Networks Technical Architecture (IG1230). Source: TM Forum	27

List of Tables

Title	Page Number
Table 1 – Traffic increase for a particular date and maximum traffic increase registered.....	16

Abstract

More than a decade ago we began introducing analytics, machine learning, and finally artificial intelligence to our networks and services. Evolving our work teams from a data-driven culture to AI-driven. It was not an easy task, it involved great challenges and cultural changes, it really is an accelerated transformation process during this pandemic, and it continues.

In this technical paper, we go through the path we are transitioning in Telecom Argentina from data analysis and AI/ML perspective to achieve operational excellence. We present the challenges we went through along, difficulties, learned lessons, success stories and next steps.

Content

1. Introduction

In our networks and services, the Artificial Intelligence (AI) has the potential to change, the way we operate, and to become the foundation of the transformation that leads to the fourth industrial revolution. But this requires hard work, a long-term commitment, and a deep cultural change. That is why we present here our journey that we started to make our operations AI-driven.

In the industry, Analytics, AI, and Automation are often differentiated. Let us remember that in [5], we define:

- Data Analytics: monitoring data to look for patterns and anomalies (without applying intelligence) and applying those patterns towards effective decision making.
- Artificial Intelligence: the development of computer systems capable of performing tasks that normally require human intelligence; this includes visual perception, speech recognition, decision-making, and translation between languages.

In a survey regarding enterprise networks, automation was enquired. The poll results indicates that more than 65 percent of enterprise networking tasks are carried out manually (often referred to as "ClickOps"), indicating that own network automation underlies on servers' automation. Ansible, customized "DIY" scripts (usually based on Python), and single-vendor, network infrastructure-focused packages are among the most widely used network automation tools. It is worth noting that these scripts are totally deterministic, that is, they only perform repetitive tasks. AIOps is the term used when a decision-making process is automated using an AI algorithm.

In next sections, we introduce in more detail what we understand by AIOps, we may state:

$$\text{AIOps} = \text{Analytics} + \text{AI} + \text{Automation}$$

AIOps is part of the 5G ecosystem, since from its conception the knowledge plane has been included. That is, a layer within the architecture oriented to the operation and orchestration of networks and services [7]. Figure 1 shows the percentage of network activities that are automated according to Gartner.

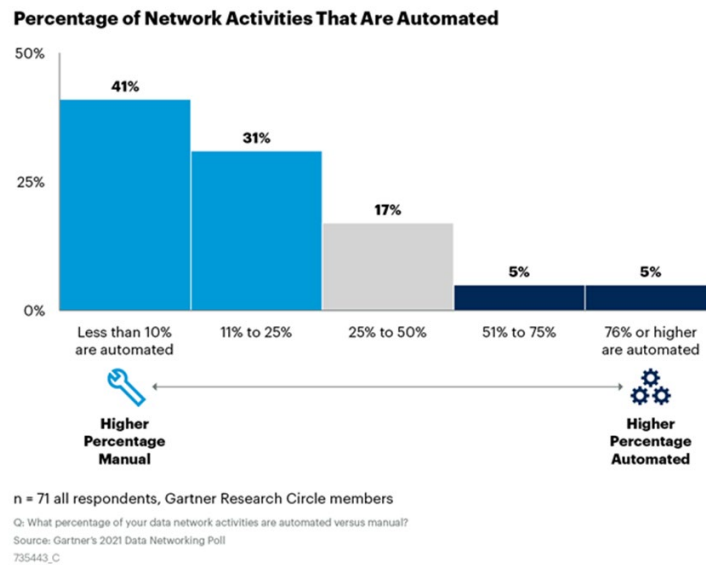


Figure 1 – Percentages of Network Activities that are automated, 2021.

Source: gartner.com

In this technical paper, we outline the path we've taken toward converting our AI-driven networks and services, along with a few use cases. The use cases development provided us with many lessons learned.

This document is organized as follows. After this section, we introduce Telecom Argentina and Financial impact of Analytics, AI, and Automation. At section 2, we expose AI-driven Operation. In section 3 we present the concept of the knowledge plane and our reference architecture. In section 4, we introduce the first use cases we did using AI and Big Data. On section 5 we present the AIOps use cases in Telecom Argentina and the state of the art in the evolution towards autonomous networks. In section 6 we outline the key lessons learned along this journey and lastly, at section 7 we describe the next steps we are considering.

1.1. About Telecom

Telecom Argentina is a company in constant evolution, which offers connectivity and entertainment experiences and technological solutions throughout the country. We boost the digital life of our over 30 million customers, with a flexible and dynamic service, in all their devices, through highspeed mobile and fixed access, and a live and on demand content platform that combines series, movies, gaming, music, and TV programs.

Our trademarks Telecom, Personal and Flow consolidate an ecosystem of platforms, and new businesses, a comprehensive and convergent experience for individuals, companies, and institutions across the country. We are present in Paraguay with mobile services, and in Uruguay, with cable television services.

Telecom Argentina has become a company that thrives in the digital world. It evolved from a traditional telecommunications company to consolidate itself as an ecosystem of apps and platforms that are based on connectivity as a differential quality value.

With the vision of going beyond connectivity, we are developing new 100% digital businesses, based on IoT, 5G, Fintech, entertainment, and Smart Home solutions, among others. With the most innovative

technology, in alliance with world class technology partners, and an investment which over the last five years reached USD 5 Bn, the company focuses on enhancing its infrastructure and systems and providing more and better services.

Our fixed-mobile network is the most extensive of the country. With more than 75,000 kilometers of FTTH, HFC and ADSL technologies, we are present in 18 provinces with over 60% of coverage of households in the country. Our 4G+ mobile network is the fastest in the country, and it is available in 100% of our infrastructure. We reach more than 1,900 locations and have a coverage of 95% of the population. In 2021 we inaugurated the first 5G network in Argentina, with 20 sites in the City of Buenos Aires, Rosario, and Costa Atlántica. We also bring connectivity to numerous towns with less than 500 inhabitants in different provinces, and in many cases, we are the only link they have with the rest of the country and the world.

The convergent and comprehensive operation of the network is one of the key challenges of the evolution of Telecom Argentina. Automation, analytics, and artificial intelligence, among other innovative technologies, will undoubtedly mark the path of our completely transformation from a Communication Service Providers (CSPs) to a Digital Service Providers (DSPs), Figure 2.




Figure 2 – From CSP to DSP.

1.2. AI-driven

Adoption of cloud computing, network function virtualization (NFV), and the development of software-defined networking (SDN) have all advanced considerably in recent years.

These developments have made it possible to create infrastructure that is more adaptable and to make storage and computing power more plentiful than before. We are compelled to use both artificial intelligence and machine learning techniques because of this progress and the growing requirement to enhance the management and administration of networks and services.

On the other hand, the next generation of 6G communications is already being investigated, where the devices, antennas and infrastructures are embedded with artificial intelligence software. Today networks cannot survive without artificial intelligence [1].

We are confident that these technologies will assist us in resolving and enhancing the current challenges with network efficiency so that our more than 30 million consumers have a better digital experience. Additionally, it will free up our specialized labor to work on more difficult tasks connected to emerging digital services and businesses.

These technologies will help us to:

- Reduce time to market for new products and services
- Predict and mitigate network and equipment service issues before they happen

Operating, managing, and provisioning future services with automation processes becomes essential to increase efficiency.

- Future Challenges: We are starting down the path towards the automation of our network operations. Using AI to efficiently manage the exponential traffic growth and the complexity and variety of new services that 5G will enable.

2. What it means to be AI-driven operation

When working on AI-driven initiatives, it is crucial to establish the minimum level of understanding inside our organization regarding the scope and restrictions of AI technology as it relates to the functioning of our networks and services.

There is no single definition for all AI technologies or framework. When we refer to technologies such as DOCSIS or 5G, there are no ambiguities since they are very mature in their standardization process.

AI-driven operations refer to the use of AI for the operation, planning and decision making in networks and services. To convert the operation of networks and services into an AI-driven operation, it is necessary to go through three stages:

Stage 1: Use of artificial intelligence techniques to develop applications and use cases.

Stage 2: Using AI to improve services, processes, or products.

Stage 3: Use of AI to help decision making.

In Figure 3 we present the Hype Cycle for artificial intelligence according to Gartner.




Figure 3 – Hype Cycle for Artificial Intelligence, 2022.

Source: [gartner.com](https://www.gartner.com) (July 2022)

3. Our first steps without AI using math, statistics and small data

In this section we present some data-driven decision-making use cases we performed as our first steps in this AI-driven operation journey.

3.1. Dimensioning of our VoD system

By the end of year 2012, Telecom Argentina launched its Video on Demand service. Due to certain recommendations from other operators, the original coverage was 900 HD set-top boxes (STBs) per service area (SA). However, when the service was launched it was not possible to achieve this (because of CAPEX restrictions), and many service areas were oversized up to 2,500 STBs. Hence, our customers experienced VoD service outages.

For the above, it was necessary to propose a statistical model that enables the resizing the SAs, establishing a balance between the system capacity, the number of customers and the *blocking probability* [2].

After studying and analyzing various models, we proposed a model based on the Queuing Theory applied to VoD traffic, the Erlang B formula, already used in telephony traffic [2]. In particular, the Erlang B formula allows us to relate three fundamental variables: the offered traffic, the number of available streams and the blocking probability, resulting in the appropriate number of STBs that the SAs must contain given a certain blocking probability that is considered acceptable.

Traffic definition is based on empirical assumptions, including peak service period, penetration per SA, average content duration, and average number of first-attempt requests. These last two depending on the quality of the content (SD or HD).

The data analysis was not only critical for the design of the model but, also contributed to establish an acceptable blocking probability for the service and to determine the phenomena that influence the performance of the service and their affectation degree.

3.1.1. Erlang B statistical model

The objective is to calculate the number of STBs per SA to obtain an acceptable blocking probability for our system.

For its formulation, certain empirical assumptions were considered after the characterization of the service. Such as, peak period, peak period duration, average duration per quality (SD/HD), average number of requests and penetration rate. We didn't consider blocked orders to be retried.

The traffic calculation is then formulated as follows:

$$A = \frac{h \cdot (\lambda_{SD} \cdot t_{SD} + \lambda_{HD} \cdot t_{HD}) \cdot p}{T}$$

Being,

- h : n° of STB per SA.
- λ_q : average number of first attempt requests, per STB and per period ($q = \text{SD or HD}$).
- t_q : average time spent on the system per period (min) ($q = \text{SD or HD}$).

- p : VoD service penetration per period.
- T : peak period duration (min).

Since we analyze the number of STBs needed for each SA, the goal is precisely h . It is obtained, using the Erlang B formula, based on the blocking probability (P_B) and the available streams during the peak period (N). They are dynamic depending on the bitrate of the content, which varies between 1.875 and 15 Mbps (depending on SD/HD quality and encoding).

The application of the formula is possible since the system satisfies the hypotheses of the model.

$$P_B = \frac{\frac{A^N}{N!}}{\sum_{i=0}^N \frac{A^i}{i!}}$$

Figure 4 shows graphically this formulation evaluate on different scenarios.

We sat a blocking probability (P_B) value of 3%, based on our observations and conditions of acceptance for quality of service. Then, from the model, for a MPEG2 content with a ratio of SD/HD of 70/30, our recommendation was 700 STBs per SA.

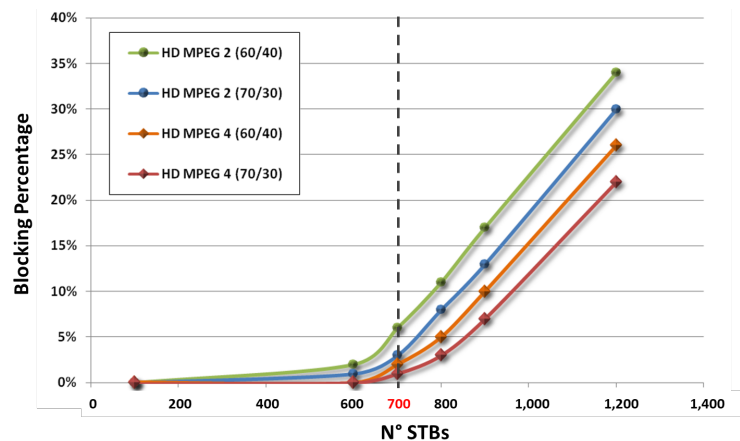


Figure 4 – Relationship between n° STB and Blocking (%) for different scenarios according streaming codification and quality.

The variables that most affect the model were also studied, these are:

- Very high rating content (peak time)
- Free content
- SD/HD service ratio
- HD Encoding (MPEG 2, MPEG 4)
- Holidays

Given the success of the model, we extended its use to assess the impact, in terms of probability of blocking, of different modifications to the service, such as: conversion to HD of CatchUp content, the addition of new services related to the 2014 World Cup, among others.

3.2. Characterization and impact of user behavior of OTT services

One of our major tasks was the study of applications' traffic. OTT video services represents more than 50% of our total downstream traffic. In terms of downstream traffic and client consumption, Netflix was the service with the highest utilization, especially during peak hours.

We found evidence that the link between subscriber count and traffic volume fits sub-exponential distributions. These long-tailed distributions can conveniently describe high variability. Historically, some long tail distributions have an origin in the income distribution, for example, Pareto and Log-Normal. The latter is one of the least understood and most widely used functions.

In Argentina, broadband users access streaming video content, mostly located in the USA. Thus, the impact of Round-Trip Time (RTT) in service performance is very important. A typical RRT in our service groups is around 120-150 msec. Therefore, CDN usage turned crucial.

Two server farms were set up in different company data centers to create a local cache of Netflix content. Each Open Connect Appliance is made up of three 12 Gbps-capable servers and were provided by Netflix. We had to connect to Netflix in Brazil via international peering since we were getting roughly 100 Gbps of Netflix traffic at its busiest each day. This allowed us to meet the demand for traffic that our CDN was unable to provide.

Since throughput is inversely related to RTT, which is smaller at shorter distances, an increase in throughput was obtained by connecting to Brazil instead of the USA. This was made possible by Netflix's adoption of the TCP-based DASH protocol.

3.2.1. *Measuring performance for decision making*

At the beginning of the Netflix CDN implementation, since it was only applied to the half of the service groups, we performed several tests where the traffic from Netflix users was compared pre- and post-implementation for each partition.

We observed the performance increase was noticeable for service groups using CDN, even more during the morning when workload was low. We have found the traffic generated by Netflix CDN users had doubled the traffic generated by those who were not using CDN.

Based on the above results we could estimate the total bandwidth traffic growth in the service groups. Assuming an increase of about 25% in traffic, since the implementation of the CDN, and maintaining the same number of active flows, the associated downstream Netflix traffic should have the same grow proportion. Let:

- N : "Previous Netflix Traffic"
- T : "Previous Total Traffic"
- N' : "Post Netflix Traffic"
- T' : "Post Total Traffic"
- Δ : "Traffic Growth"

So,

$$N' = N + \Delta N, T' = T + \Delta T \quad \text{and} \quad \Delta T = \Delta N \quad \rightarrow \quad T' = T + \Delta N$$

$$\frac{T'}{T} = 1 + \frac{\Delta N}{T} * \frac{T'}{T} = 1 + \frac{\Delta N}{N} * \frac{N}{T}$$

Thus, the percentage increase in total traffic is the product between the percentage increase in Netflix traffic and the percentage that Netflix represents of total traffic. Assuming a service group not using the CDN the relation N/T is 35% (calculated previously) and the percentage increase is 25%, we got:

$$\frac{T'}{T} = 1 + 0.23 * 0.35 = \mathbf{1.0805}$$

So, the total traffic growth of the service group will be 8.05%, supposing the number of concurrent subscribers remains invariant.

3.2.2. Forecasting subscribers and traffic

Having understood the Netflix's traffic importance on the capacity planning, we developed a time series model to forecast Netflix subscriptions, the amount of generated traffic and therefore, the impact on the access network.

We proposed a model based on time series, which provides powerful statistics. Commonly used in business and economics where data occurs in the form of successive values of a variable, in an ordered sequence in an equally spaced time interval. A stochastic model for a time series will generally reflect the fact that nearest observations in time, more closely relation than observation further apart.

Time series models allows us to:

- Understand underlying characteristics and structure that produced the observed data, through different analysis methods.
- Fit a model and forecast future values based on previously observed values for monitoring or even feedback and feedforward control.

Through statistical software, we observed that ARIMA was the best-fitting model. For Netflix users forecasting, we selected a confidence interval (CI) of 95% that, calculated from a given set of sample data, returns an estimated range of values which is likely to include an unknown population parameter. The results are shown in Figure 5.




Figure 5 - Monthly evolution and forecast of Netflix users among our subscribers.

Remarkably this model has been successfully applied to forecast the whole broadband subscribers' population and particularly for Netflix subscribers, by fitting properly ARIMA model parameters.

In following sections, we describe how these first forecast models evolved into a tool with much more features and characteristics.

3.3. IP traffic dimensioning based on patterns

In this work we presented a forecasting model for the Average Bandwidth per Subscriber parameter, as a novelty we have made clusters by Service Tier Levels and by subscriber's location.

Exploiting our DPI tool, we analyzed trends and Internet users' preferences. Through mining our data at different periods, we found common daily patterns among our subscribers. This allows us to focus on a single day to have a whole understanding of the network. Considering the different Service Tier Levels, we also found common patterns related to the most used applications. Only Video Streaming services showed a variation among the lowest tiers. Thus, increasing download access speed drives consumption much more in video applications.

The resulting statistical parameters of our subscriber's characterization are inputs for a network dimensioning tool we have developed to analyze traffic impact over the service group's QAM carriers and simulate different scenarios.

For the last years, we have been collecting and analyzing various statistical measurements including the average residential bandwidth (BW) traffic considering the access speed and we found there is a mathematical relationship between them.

Some technology vendors have different proposals to estimate bandwidth capacity. Cloonan et.al proposed in [3], a technic based on the average amount of bandwidth per subscriber during the busy hour (T_{Avg}) calculated as follows:

$$T_{Avg} = 8 \cdot \frac{B}{W * N_{sub}}$$

Where N_{sub} is the number of subscribers within a typical service group, B is the number of bytes passed into the service group within a given window of time (W) measured in seconds.

We introduce a different methodology, we collect 5 minutes of bandwidth traffic data from the whole network (global, per service and per tier) and the four service tier levels, using a DPI technology. Monthly, we take the maximum bandwidth traffic per service tier level and divide it by the total number of subscribers of each. Then, we model this data with Linear Regression algorithms.

We found at least three different applications of this methodology:

- Adjust current BW traffic values per Tier
- Estimate BW traffic values for future offered access speed
- Forecast BW traffic values for the following two years

We define

- Avg_BW_subs : Average bandwidth per subscriber per Tier.
- $Speed$: Max access speed for the Tier.
- j : Tier number.
- t : Number of months used to forecast.
- β, γ : Parameters to estimate.

To adjust and to estimate, our model takes the form:

$$Avg_BW_subs_j = \beta \cdot \sqrt{Speed_j} + \varepsilon, \quad j \in \{1, \dots, 5\}$$

To forecast:

$$Avg_BW_subs_j = \gamma_j \cdot t + \varepsilon, \quad j \in \{1, \dots, 5\}$$

Figure 6 shows the Linear Regression of the average bandwidth per subscriber. Blue line represents the adjusted and estimated current values. Green and orange lines, the forecast for the next two years.




Figure 6 - Linear Regression of the average bandwidth per subscriber. Blue line represents the adjusted and estimated current values. Green and orange lines, the forecast for the next two years.

3.3.1. Dimensioning tool based on simulations

We developed a tool that aims to obtain a set of simulations to evaluate the use of bandwidth by service group, traffic consumption, customer portfolio and its evolution over time. The tool outcomes allow us to visualize measurements, make estimates and use projected values to simulate consumption.

The traffic sizing criteria adopted suggests taking the maximum 95th percentile, within a 15-day period, to predict the estimated growth at a given time. We take these statistics as sample of the worst hour. Simulations are performed in 24x15 scheme to compare measurements and simulations using the same criteria.

Portfolio update scenario simulation entails some essential input parameters: service tier levels distribution per service group and *Avg_BW_subs* values obtained through linear regression. The latter allows us to analyze service group's load, considering the number of subscribers and their respective service tier. It also allows to estimate the impact on the network by adding new tiers to the customer base as well as a massive portfolio upgrade.

3.3.2. COVID-19 and HFC traffic growth

Since the lockdown started in 2020 because of the COVID-19 pandemic, totally disruptive changes in clients' behavior began to be observed, affecting network performance. Therefore, we began to analyze the evolution of downstream and upstream traffic on a weekly basis with the aim of finding new patterns of use, that allow the growth prediction in a completely uncertain scenario and take proactive actions in the network to alleviate the impact on the service.

We presented these changes graphically to briefly understand the suffered impact after lockdown, Figure 7; and numerically to enable the comparison between past situations and predict future ones.




Figure 7 – Change in daily HFC traffic patterns.

Traffic increases were calculated between peak of April 30th, 2020 (post-lockdown), and the peak of the reference week, Feb 29th to March 6th (pre-lockdown), for midday and peak hour. Midday is defined as 12 p.m. to 2 p.m., and peak hours are defined as 8 p.m. to 2 a.m. Peak hours remain to be when DS and US traffic peaks happen rather than midday. For US traffic, the post-lockdown midday spike outpaced the corresponding spike pre-lockdown peak-hour by 41%. Figure 8 and Figure 9, display the evolution of US and DS. SARIMA model was implemented to describe and predict the US and DS behaviors.

The Home Office, school online courses, and Internet-related recreation are the main drivers behind this transformation. The most popular categories are video meetings (Zoom, Webex, etc.), streaming (Netflix, Youtube) and gaming; file sharing also increased. The former calls for a heavy reliance on upstream, which has not been in great demand for years.

Traffic Increase April 30, 2020	DS	US
Mid-day	4%	85%
Peak Hour	28%	93%

Max Traffic Increase	DS	US
Mid-day	37%	90%
Peak Hour	48%	95%

Table 1 – Traffic increase for a particular date and maximum traffic increase registered.



Figure 8 - Evolution and forecast for Downstream traffic, 2020.




Figure 9 – Evolution and forecast for Upstream traffic, 2020.

3.4. Real-time analytics for IP video multicast

To understand the impact of multicast implementation, it was necessary to collect data on key indicators such as the number of concurrent streams, average bitrate, and average bandwidth to estimate bandwidth gain [13].

3.4.1. Multicast gain formula

We evaluated the multicast gain at service group levels as a percentage of the capacity needed under a 100% Unicast scheme, which we define as follows:

$$Capacity\ 100\%\ Unicast = \sum_{\substack{All \\ channels}} Concurrence \cdot Avg\ bitrate$$

When working with data from the Legacy system, we used the access frequency to approximate the concurrence, and we assumed that the average (Avg.) bitrate is 4 Mbps.

We examined different scenarios defined as follows: *Top "X"*, the "X" most popular channels are delivered to Multicast, and the rest remain Unicast.

The capacity needed is calculated as:

$$Capacity\ Top\ "X"\ Scenario = \sum_{\substack{Top\ "X" \\ channels}} Avg\ bitrate + \sum_{\substack{Other \\ channels}} Concurrence \cdot Avg\ bitrate$$

Finally, the multicast gain is:

$$Multicast\ gain\ for\ Top\ "X"\ Scenario = \frac{Capacity\ 100\%\ Unicast - Capacity\ Top\ "X"\ Scenario}{Capacity\ 100\%\ Unicast}$$

Gain is calculated using the overall concurrence for each channel. Additionally for service group level gain, internal service group concurrence is used.

In addition, a theoretical scenario is proposed, in which all the channels are transmitted via multicast to estimate the maximum multicast gain. This is helpful to determine whether the gain in other cases is nearly at its maximum or not.

$$Capacity\ 100\%\ Multicast = \sum_{\substack{All \\ channels}} Avg\ bitrate$$

The maximum gain is estimated as:

$$Maximum\ Multicast\ gain = \frac{Capacity\ 100\%\ Unicast - Capacity\ 100\%\ Multicast}{Capacity\ 100\%\ Unicast}$$

When there are approximately 10 channels delivered via multicast, it has been shown that the gain is approximately 50% during peak hours. The marginal gain tends to decline as more channels are multicast supplied. It was discovered that the benefit almost reaches its maximum with 25 multicast channels.

3.4.2. K-means clustering applied to the selection of multicast channels

K-means algorithm was used on ranking data to investigate the channel count that would be provided using multicast if it were an autonomous and unsupervised procedure, Figure 10. We discovered this technique categorized between 4 and 9 channels as the most popular after evaluating six months of worth data.




Figure 10 - K-means clustering applied to the access frequency per channel by day for the Legacy system. Algorithm used to classify the signals between multicast and unicast.

Real time analytics provides an efficient alternative for monitoring a policy-driven multicast strategy, due the advent of unconsidered special event which could drastically shift the ranking for a few hours.

It is proposed a continuous process, based on k-means clustering algorithm, executed every 10 minutes. The application searches for the channels with the highest viewing rates and determines if they are on the list of multicast channels. It sends an alert and reports the top list when it detects one or more channels that are being accessed massively and are not in the multicast list.

4. Evolving with AI and Big Data

Machine Learning has made significant advances in the telecommunications business, delivering numerous benefits. Given the tendency of strong virtualization, it simplifies operations.

However, we must not forget that the success of the application is dependent on the duties performed by people. In general, the tasks must be established to locate the learning models are as follows:

- Data: Split data into development and validation. Define instances, classes and attributes.
- Experimentation: Attributes selection. Performance metrics. Cross-validation.
- Model validation: Processes intended to verify that models are performing as expected, in line with their design objectives and business use case. It's the most important step in the model building sequence.

At Telecom, there are a few initiatives that we have considered within the AIOps framework. At STEM team one of our missions is to lead AIOps in our current and future networks. Our recipe is diversity, work in cells, agile mindset, and self-learning. In the following sections we present two of the initiatives we have been working on.

4.1. Network capacity and Machine Learning

We proposed the use of machine learning techniques such as Principal Components Analysis (PCA) and Artificial Neural Networks (ANN) to characterize the optical nodes that integrate our network and define the strategies the company will use to meet short and long-term demand [4].

We conducted different analysis at node level based on variables such as monthly consumption, households passed, traffic per port and downstream channels distributions, among others. We used a huge volume of data from different sources to obtain examples for the training sets used in the algorithms. Due the produced results are needed for a significant number of cases and on a regular basis, the process must be automated by applying machine learning and multivariate analysis techniques.

This analysis used traffic-data collected every Sunday during prime time, from all network's ports. One variable in the datasets contains the maximum traffic (Kbps) registered in each port. Our approach consisted in analyzing two key indicators:

- Average bandwidth traffic per residential subscriber at peak time.
- Ports usage.

The first one will provide information about the zones where there is a need for higher bandwidth, and the second 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 subscriber at peak time, the metric was 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:

$$\text{Max utilization [\%]} = \frac{\text{Max Port Traffic}}{\text{Port capacity}} \quad (2)$$

To gather data about how the two key indicators relates to other variables, we also included in our analysis: the count of segments or zones connected to the port, CMTS model, optical node classification according to the region of location, investment plan status, network capacity (1GHz or other), DOCSIS version, cable modems count, HHP, network extension (in Km) and total monthly downstream consumption.

4.1.1. Principal Component Analysis (PCA)

Principal Components are the underlying structure in the data. Their interpretation is 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 Sunday
- 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

We concluded there are two main PC, which can be explained as follows:

- PC1: It takes 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 takes higher values as the number 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.

PCA highlighting there are two port groups, Figure 11. One where the aggregation of more subscribers draws a substantial increment in the traffic, and another where the impact of adding subscribers is lower.




Figure 11 - PC1 vs PC2 and its relation with traffic management.

4.1.2. Artificial Neural Network (ANN)

We based our ANN on the following principles:

- 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 dataset has affected any step of the learning process, its ability to assess the outcome has been compromised.

We determine one of these four options to increase network capacity and, as a result, access speed: chassis upgrade, recombination, node segmentation, and node division

For the neural network training we first needed a sample of nodes to be classified by the expert team. Sample characterization was made by classifying nodes into three strata: the first one has the nodes in which their ports have a mean utilization below 50%; the second stratum contains nodes where mean utilization lies below or equal to 80%, the third groups the nodes with mean utilization above 80%. It is in our interest to have a faithful representation of the HHP variable in the sample.

For overfitting avoidance, sample data was split in training (60%), cross-validation (20%), and testing (20%). Training set was used to estimate the weights in the neural network. Cross-validation helped validating the model in terms of variables and optimization of the selected parameters. The testing set wasn't used in the construction of the neural network but to check whether there was overfitting or not, by measuring network classification performance with 'new observations. Quadratics terms were included to search for higher accuracy. The ANN scheme got is shown in Figure 12.

Network's first layer contains eight inputs, the four variables mentioned (represented as x) and the same variables at square (x^2). Second layer, also called hidden layer, contains eight data points too ($a^{(1)}$), and 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 no investment was needed at the time.

The optimal weights solution in the network threw an accuracy level of 96% with the training set, and an accuracy around 90% with testing set.




Figure 12 – Our ANN scheme.

4.2. Optimizing video customer experience with Machine Learning

For video service providers it is important to evaluate the processes that affect the quality of video considering the perception of customers. The subjective evaluation of the quality of the video measuring the opinion of human users is expensive and slow, although there are public databases with standardized results. To automate the evaluation of quality of video objective models that try to approach the subjective evaluation human are used. Recently, objective models emerged using Machine Learning (ML) algorithms which are trained using databases with subjective evaluations, to combine a variety of classical metrics. Classical metrics are much simpler to implement and at a lower cost but, they produce worse results that do not always fit the human perception. On the contrary, the metrics based on ML produces results very close to the subjective opinion of the customers, but they are more complex to implement and provide development opportunities.

Within the objective metrics based on ML there is a method named Video Multimethod Assessment Fusion (VMAF), an open-source method proposed by Netflix in 2016 and is a video quality metric that combines human vision modeling with ML to provide a great viewing experience to their members.

Telecom has its own IPTV platform, called Flow, that is based on unmanaged (second screens) and on managed devices (set top boxes). It provides different types of advance video services, including Linear TV, various flavors of On Demand services (VoD, CuTV, Reverse EPG, StartOver, network DVR, Pause Live TV, and Trick Modes), using different streaming technologies, Search and Recommendations. Thus, it's very important for the Company to develop VMAF as a tool for optimizing Flow customer experience and to equalize video quality with other existing video platforms.

In order to train VMAF for optimized Flow customer experience we defined a dataset following Netflix recommendations regarding the type of content. We selected 35 videos, each 10-sec long from Flow catalog. To make the distortions, each source video was encoded with 6 resolutions up to 1080p. Video characteristics were variable, and used a selection of videos with fire, water, nature, animation, close-up, action, crowd, among others.

We ran a subjective test through 6 different focus groups. Each group of about 15 subjects. Each subject sits in a living room-like environment and was instructed to watch an unimpaired reference video, then the same video impaired and give a rating on a continuous scale from “bad” to “excellent” (ACR methodology), then we translated the scale to a range from 1 to 5 and calculated the MOS.

Then, we trained several models with different sets of parameters for the Support Vector Regressor, the ML model used to perform VMAF (Figure 13), to avoid the underfitting and the overfitting.

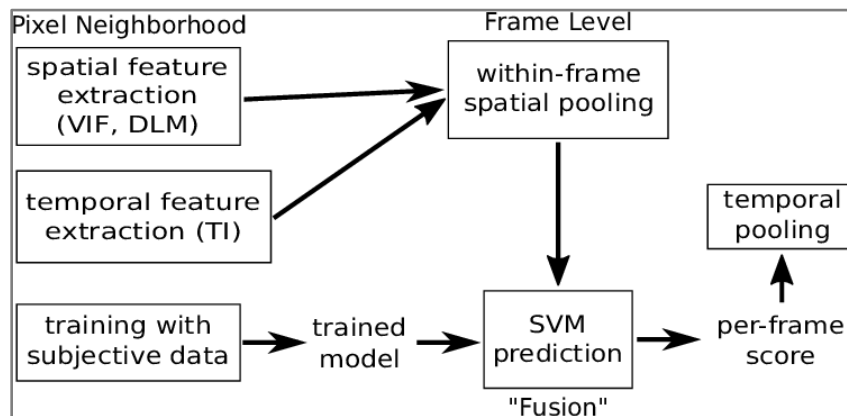


Figure 13 – VMAF Algorithm.

Once we had our model trained, we tested it with the testing dataset previously defined. After the testing, we calculated the performance metrics to measure the prediction accuracy (60%). Based on the results we obtained, we understood that we must continue improving the model with larger training and testing datasets or with other elementary metrics combination. In any case, the accuracy of machine learning models that use subjective variables does not usually exceed 70%.

5. AIOps

The term AIOps was coined by Gartner in 2016 and have pushed the concept into the marketplace. According to Gartner “AIOps combines big data and machine learning to automate IT operations processes, including event correlation, anomaly detection and causality determination”. And, according to Figure 14 has three main elements: observe, engage, and act.




Figure 14 – AIOps cycle by Gartner.

While AIOps was developed to give scalability to the management of IT systems, given their increasing complexity and framework, we can extend its applications to Telecommunications operations.

A first definition of AIOps in the telecommunications industry is the use of Artificial Intelligence for the operation of networks and services. The long-term goal is to achieve autonomous networks (AN), which in some ways, at the beginning, entails automating many of the operational procedures [5].

Automation in AIOps is understood to range from automating network capacity planning through some ML algorithm, automatically detecting anomalies in traffic flows through ARIMA time series, or even automating a process of adjusting a modulation profile in the form autonomous (closed loop) (Figure 15).

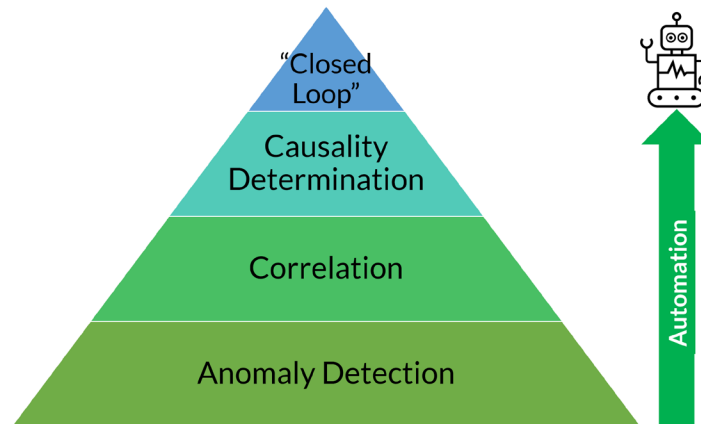






Figure 15 – AIOps automation path.

Gartner also defines four levels of Network AIOps, that we can see in Figure 16.

Four Levels of Network AIOps Functionality

Level 0	Level 1	Level 2	Level 3	Level...
				
Full Human Intervention Required. Reliance on templates, manually refreshed best-practice database, detects basic configuration errors	Substantial Human Intervention. Real-time detection and correlation of issues, limited or no automated issue remediation	Minimal Human Intervention. Detection and correlation of cross-domain issues, advanced issue remediation and real-time optimization	No Human Intervention Required. Automated topology mapping, resource orchestration, pervasive remediation and real-time optimization	Future AIOps Functionality yet to be defined

Source: Gartner
763626_C

Gartner

Figure 16 – Levels of Network AIOps functionality.

We adopt AIOps as a framework to continue our journey towards an AI-driven operation. While autonomous networking is our goal, we will only make a brief reference below.

In summary, AIOps refers to the application of AI technology to automate business processes and operations at CSPs.

5.1. Ada

At Telecom Argentina, we have been working on a network dimensioning solution since a long time ago. In 2017 we developed a tool that assisted the decision-making process for HFC network dimensioning, mentioned in a previous section. Later, we incorporate data from the radio-access network (RAN) and relaunched the project under the name Ada. The goal is to have a tool that combines machine learning and subject matter experts (SMEs) input to assess investment decisions on CAPEX and OPEX.

Historical data is collected on every HFC node, as well as for every cell on the mobile network. The concept is that for short-term decisions (two years forward), ML-based forecasts are offered, and the duty of experts is to identify priorities and act on them. Long-term decisions (5-10 years periods), on the other hand, ML-based forecasts would require a longitude of history that hasn't yet occurred. In our experience, however, it is possible to provide reasonable approximations if forecasts are supported not only by historical data but also by experts' knowledge. Hence, a ML clustering model has been added, which characterizes the nodes according to their traffic and other variables. Based on the short-term forecast, nodes characterization, and feasible tasks that can be carried out on the network (such as migration from HFC to FTTH), possible long-term scenarios are obtained.

Planning engineers used to research on potential scenarios and make a series of calculations to approximate what they thought was going to happen in the long-term, based on their knowledge of the average client and use cases. Working together, we built profiles for the operation area, so they can refine what to expect from a variety of use cases. For example, consider two cases where a household group mostly uses the service for social media and the income is low, with another group with heavy streamers and high income. We provide better information to planning teams at the beginning. An example of the Ada application is displayed on Figure 17.

We continue working on how to enable engineers to pass information to a model about what they expect the CAGR would be at different kinds of operating sectors, in the next 5-10 years. This will involve a combination of simulation and forecasting techniques. It will also require a certain level of automation that is not achievable outside a ML framework.

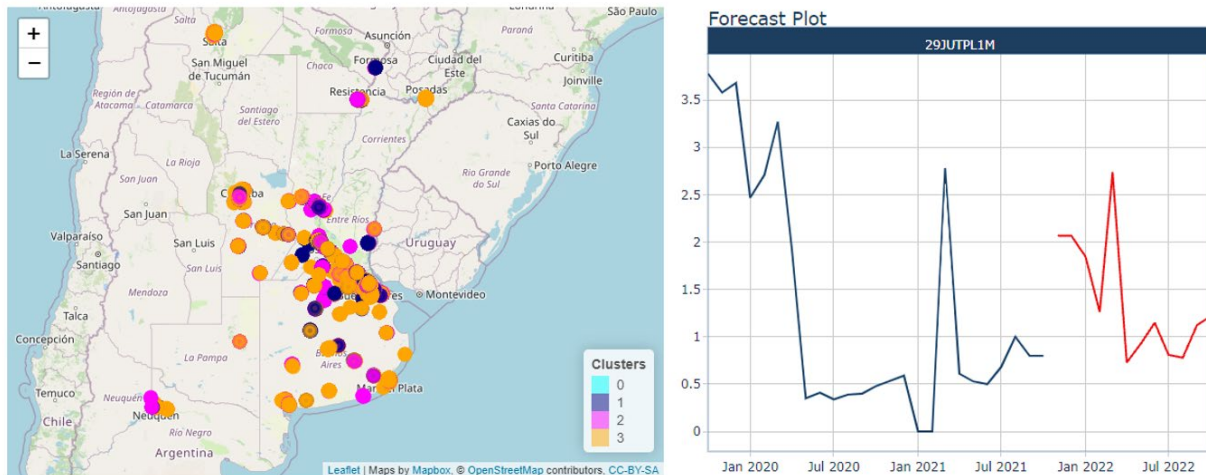


Figure 17 – Clustering and Forecast performed with Ada.

5.2. Customer claim prediction for HFC network

Within the domain of AIOps one of the most popular use cases is customer claim or tickets prediction. The main idea is to use information derived from different elements of the network, such as: CMTS, cablemodems, electronic devices, etc. to estimate the probability of a customer to generate a complaint.

Information is generally collected through OSS systems and ingested in a machine learning pipeline where preprocessing, analysis and ML model testing is performed. We are currently developing this kind of solution to ultimately increase customer satisfaction.

Using hourly collected information from over 3.5 million DOCSIS 3.0/3.1 cablemodems we are trying to anticipate customer complaints two days in advance. To handle this amount of data we have partnered with Google to develop and deploy this project using Google Cloud Platform (GCP) services. A high-level view of GCP implementation is shown in Figure 18.

Potentially relevant variables (e.g., signal to noise ratio, consumed bytes, average Rx, t3 and t4 time outs, etc.) have been identified and our Service Assurance team have been able to efficiently transfer these data, collected by our OSS systems, to GCP.




Figure 18 - High-level view of our project: inputs, outputs, outcomes and architecture.

Regarding technical challenges, two of the most difficult tasks have been dealing with an extremely unbalanced dataset and label noise. Although from a business perspective having a low proportion of customer claims is a good indicator, some ML models can struggle to learn from these kind of data sets. In our case, after filtering the target population, the positive class dropped to 0.1%, which added more complexity to the problem. To reduce the impact of this issue on model performance we applied hyperparameter tuning and a technique called SMOTE combined with an under sampling of the majority class. Both approaches led to improvements in model performance.

We have been able to overcome many challenges and finally the ML model selected was XGBoost. As a result, we obtain a daily list of customers with high probability to claim. From this list, with CX and Field Service teams, we can make proactive calls to solve customer problems remotely or to send a technician to their home if necessary. Then, we measured through surveys and NPS how was the experience of our customers and we found that this proactive action has a positive impact on their satisfaction.

Although we are going through an early stage of this project, we understand that the future result will be to increase the number of promoters of the company and consequently avoid churn. In addition, the truck roll and the costs associated with it would be reduced.

5.3. Intelligent agents and Autonomous Networks

Peter Norvig and Stuart Russell present in [6], eight definitions of AI, from which they define Intelligent Agents as “agents that receive precepts from the environment and take actions that affect the environment.” Agents’ actions change the world. Thinking (Cognition): Interpret sensory data. Then updates its environment (model of the world) and, decides on next best action. The autonomous networks that are being defined in the TM Forum and ITU are based on this definition.





Figure 19 – Intelligent agent.



**Figure 20 – Agent in AN source: Autonomous Networks Technical Architecture (IG1230).
Source: TM Forum**

There are two concepts we have been working on with our stakeholders and team in the CTO office.

- Autonomous Networks make decisions without human intervention (decisions made by Agent).
- Automation operates without human control (can be implemented without AI technology, as we have seen).

Another related idea is “intent-driven automation”. Intent-based networks enable service providers to define the behavior they expect from their network from service and business perspective. Intent was first introduced around 2015 in the context of SDN controllers. “Intent is the formal specification of all expectations including requirements, goals, and constraints given to a technical system”.

This definition is inspired by and compatible with the definition IETF has published in 2020 [10]. The definition associate's intent with goals, requirements and constraints provided in a declarative way. Intent constitutes and expresses knowledge about these concerns and enables sharing this knowledge between the originator and receiver of the intent.

“Autonomous networks are those that possess the ability to monitor, operate, recover, heal, protect, optimize, and reconfigure themselves; these are commonly known as the self-properties” [11].

The impact of autonomy on the network will be in all areas including planning, security, audit, inventory, optimization, orchestration, and quality of experience. At the same time, autonomy raises questions about accountability for non-human decisions that affect customers.

The growing virtualization and cloudification (Telco-Cloud) of our networks make possible the evolution towards AN. At Telecom, we use the AIOps framework since we know that it will lead to AN through automation [5].

6. Learned lessons

In this part, we'll highlight some of the lessons we've learnt from our experiences or those of other CSPs with whom we've collaborated on AI working groups for network and service operations. Over the years, adopting an AIOps framework and making our operations AI-driven have required significant culture change and teamwork, all in line with the company's technology strategy. The fact that AI is used in the process does not imply that we are AI-driven.

Operations + AI technology \neq AI-driven

We are training our engineers to focus on the operations-related tasks that are more challenging and where they can bring the most value rather than the less interesting and repetitive ones that can be handled by the systems. As indicated by the AIOps TM Forum framework, engineers working on AI-systems exception management and continuous optimization.

6.1. Define and communicate AI-driven operations

A lot of expectations and skepticism occur when AI is introduced because of the numerous movies that have been made over the past few decades. Because of this, it is crucial to carry out an internal communication and evangelizing work on what we mean by AI-driven operations. The various stages of the analytical component must be clearly discussed.

- Descriptive – What is happening now?
- Diagnostic – What happened and why?
- Predictive – What might happen?
- Prescriptive – What actions should be taken?

We need to plan out how we can effectively inform the various business teams about these AI-driven projects.

It's also important to mention that operations engineers view AI as a tool for augmented intelligence (AgI).

While AI builds machines that behave and function like people, AgI uses the same machines but takes a different tack to enhance human skills. AgI actually entails a collaborative effort between humans and machines that makes use of each party's advantages to boost overall commercial value. In other words, AgI's main objective is to enable people to work more effectively and smarter. [7].

Finally, it is also important to communicate what problems can we solve, and which ones cannot, their scope and limitations.

6.2. Workforce

Believing that a pair of data scientists can resolve a network operating or planning problem is one of the most frequent errors we have observed.

Teams using AI should be set up with professionals in data sciences, operations engineering, field service, customer experience and also with technicians to tackle operational challenges. We need to build internal AI teams.

Like many businesses throughout the world, our main challenge is finding experts who are knowledgeable in programming, telecommunications, and AI. We have junior and semi-senior data scientists on the STEM team because of this. Tech Scientist is the next level of seniority. The tech scientist has expertise in the technology that he uses in collaboration with the operations and field service engineers, in addition to data science and programming skills.

In several talks with colleagues from other companies, they have even raised the difficulty of recruiting data specialists. Because of this, digital training and recruitment are a priority in Telecom. Along with learning about communication technologies, our engineers at the CTO office continually train in data sciences. Our aim is to keep our work teams AI-driven learning continuously so they may continue to advance their careers and bring value to our clients.

Our recipe is diversity, work in cells (with the experts of service assurance, field service, customer experience, etc.), agile mindset, and self-learning.

To ensure that these people choose to work at Telecom despite the rising demand for these profiles in this hyperconnected world, we are creating a reskilling strategy for our employees as well as loyalty and retention programs.

Automation will also make it possible for operational departments to integrate technology to streamline procedures and improve our customers' digital experiences, which will further raise business efficiency.

To sum up, we are establishing a strategic talent plan for transforming our current teams into an AI-enabled workforce.

6.3. Models, algorithms and Explainable AI

Always keep in mind that "All models are wrong; some models are useful," and that the key is for the work team to agree on which model to adopt.

Once the useful model has been adopted to solve an operation or decision-making problem, we must somehow make the models transparent, understandable, and interpretable. Experts must comprehend how the findings were obtained and the degree of confidence that the model has before they can decide to adopt an AI tool.

The model's output shouldn't be what matters, but rather understanding why an algorithm produces a particular output. Because of this, explainable AI (XAI) is a new growing field. The goal of enabling Explainability in AI/ML, as stated "is to ensure that algorithmic decisions as well as any data driving those decisions can be explained to end-users and other stakeholders in non-technical terms" [7].

Explainability sits at the intersection of transparency (consumers have the right to have decisions affecting them explained in understandable terms), causality (it is expected of the algorithms to provide not only inferences but also explanations), bias (the absence of bias should be guaranteed), fairness (it should be verified that decisions made by AI are fair) and safety (reliability of AI systems) [8].

We know that many machine learning algorithms have been labeled "black box" models because of their inscrutable inner workings. What makes these models accurate is what makes their results difficult to interpret and understand. They are very complex. The discussion about audit AI is still open [7].

6.4. Cloud and Data/AI platforms

The democratization of data by hybrid clouds is another lesson learned. “No amount of AI algorithmic sophistication will overcome a lack of data (architecture), Data collection & preparation is the most time consuming and difficult part of AI” [12].

There is a project in our OSS systems evolution program that is solely responsible for gathering and organizing data. In complicated hybrid multi-cloud systems, we are employing platform services to tackle problems using data and AI. but before using platform services, one should give them a serious evaluation.

Platform services are increasingly being used by CSPs to accelerate the use of AI in business operations. We should take a hybrid approach to deploying platform services.

7. Next steps

Our next step is to define long-term roadmap towards Autonomous Networks. We anticipate finishing the surveys, definitions, frameworks, and scope for the strategic definition toward the autonomous networks paradigm by the end of the year. Virtualization and softwarization of networks are clearly included in this strategic framework. This ecosystem includes 5G and IoT.

In the AIOps framework, we began by identifying the claim root cause and its causality¹ using our model of customer claim prediction for the HFC network.

Conclusion

We offer the AI-driven initiatives that were significant turning points in terms of lessons learned since we started our journey. Right now, our focus is on developing within the AIOps framework.

AIOps' objective is to advance from automation to autonomous networks, but we must remember to see it from the standpoint of augmented intelligence (AgI). Together, people and robots may maximize the value of their own capabilities for the benefit of the organization.

Include operations engineers and technicians early in the use case, and keep in mind how crucial Explainability of AI models and their results are.

To reach full AN it is necessary to advance to the "Telco Cloud".

Cultural and process change is required. **"Think big, start small"**.

Abbreviations

4G	fourth generation wireless
5G	5th generation mobile network
6G	sixth-generation wireless
ACR	Absolute Category Rating
ADSL	Asymmetric Digital Subscriber Line

¹ We are trying to apply AI models with causal inference methods

AgI	Augmented Intelligence
AI	Artificial Intelligence
AIOps	Artificial Intelligence for IT Operations
ANN	Artificial Neural Network
AP	Access Point
ARIMA	Auto Regressive Integrated Moving Average
AVG	Average
bps	bits per second
BW	Bandwidth
CAGR	Compound Annual Growth Rate
CAPEX	capital expenditure
CDN	Content Delivery Network
CI	Confidence Interval
CMTS	cable modem termination system
COVID-19	Coronavirus disease of 2019
CSP	Communication Service Provider
CTO	Chief Technology Officer
CX	Customer Experience
DASH	Dynamic Adaptive Streaming over HTTP
DIY	Do It Yourself
DOCSIS	Data Over Cable Services Interface Specification
DPI	Deep packet inspection
DS	Downstream
DSP	Digital Service Providers
DVR	Digital Video Recording
EPG	Electronic programme guide
ETL	Extract Transform Load
ETSI	European Telecommunications Standards Institute
FEC	forward error correction
FTTH	Fiber to the home
Gbps	Gigabits per second
GCP	Google Cloud Platform
GHz	GigaHertz
HD	high definition
HD	High Definition
HFC	Hybrid Fiber-Coaxial
HHP	household passed
Hz	hertz
IETF	Internet Engineering Task Force

IPTV	Internet Protocol television
IT	Information technology
ITU	International Telecommunication Union
K	kelvin
Kbps	Kilobits per second
Mbps	Megabits per second
ML	Machine Learning
MPEG	Moving Picture Experts Group
NFV	Network functions virtualization
OKR	Objectives and key results
OPEX	Operational expenditure
OSS	Operational Support System
OTT	Over the top
PCA	Principal component Analysis
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
RTT	Round trip time
Rx	Reception
SA	Service Area
SCTE	Society of Cable Telecommunications Engineers
SD	Standard Definition
SDN	Software-Defined Networking
SME	subject matter expert
SMOTE	Synthetic Minority Oversampling Technique
SQL	Structured Query Language
STB	Set-top Box
SVM	Support Vector Machine
Tbps	Terabits per second
TCP	Transmission Control Protocol
TM Forum	TeleManagement Forum
Tx	Transmission
US	Upstream
USA	United States of America
VMAF	Video Multimethod Assessment Fusion
VoD	Video on Demand
W	Window time
XAI	Explainable AI

Bibliography & References

- [1] C. Righetti, E. Gibellini, C. Carreño and G. Carro, Can Future Networks Survive Without Artificial Intelligence? SCTE Cable Tec Expo 2019, 2019.
- [2] L. Takacs, "On Erlang's Formula", The Annals of Mathematical Statistics, vol. 40, no. 1, pp. 71-78, 1969.
- [3] T. Cloonan, J. Allen, M. Emmendorfer. Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements. SCTE Expo 2014.
- [4] C. Righetti, E. Gibellini, F. De Arca, C. Carreño, M. Fiorenzo, G. Carro, F. Ochoa, Network Capacity and Machine Learning. SCTE Cable Tec Expo 2017, 2017.
- [5] Claudio Righetti, Mariela Fiorenzo, Emilia Gibellini & Martin Juiz The Evolution Towards Autonomous Networks - 2021 FALL TECHNICAL FORUM PROCEEDINGS -SCTE-NCTA-Cablelabs
- [6] Stuart J. Russel, Peter Norvig 2003: “Artificial Intelligence, A Modern Approach“ Third Edition (2016)
- [7] C. Righetti, M. Fiorenzo, et.al., Augmented Intelligence: Next Level Network and Services Intelligence, SCTE NCTA CableLabs 2020 Fall Technical Forum, 2020.
- [8] S. F. M. H. J. K. S. V.-T. a. H. W. S. Barocas, "The FAT-ML Workshop Series on Fairness, Accountability, and Transparency in Machine Learning," 20 07 2022. [Online]. Available: <http://www.fatml.org>
- [9] H. Hagras, "Toward Human-Understandable, Explainable AI," IEEE Computer, vol. 51, no. 9, pp. 28-36, 2018
- [10] Intent-Based Networking - Concepts and Definitions March 9, 2020 <https://tools.ietf.org/id/draft-irtf-nmrg-ibn-concepts-definitions-01.html>
- [11] ITU-T (FG-AN), Available: https://www.itu.int/en/ITU-T/focusgroups/an/Documents/FG-AN_Terms_of_Reference.pdf
- [12] 2018 MITSloan ”Reshaping business with AI” <https://sloanreview.mit.edu/projects/reshaping-business-with-artificial-intelligence/>
- [13] C. Righetti, E. Gibellini, F. De Arca, M. Fiorenzo, G. Carro, Telecom Argentina, Real-Time Analytics for IP Video Multicast, NCTA, 2018

A New Model for Power Plant and Health Estimation

A Technical Paper prepared for SCTE by

Kang Lin, PhD

P&E RF Lab Manager

Comcast

1002 Cornerstone Blvd, Downingtown, PA 19335

480 430-8254

Kang_lin@comcast.com

Michael Nispel

Senior Principal Engineer

Comcast

1002 Cornerstone Blvd, Downingtown, PA 19335

610 952-3783

Michael_nispel@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background	3
3. Approach	5
4. Temperature Effect.....	6
5. Aging Effect.....	7
6. Discussion/Results	8
7. Conclusion.....	9
Abbreviations	10

List of Figures

Title	Page Number
Figure 1 – Department of Energy OE417 – Annual Summary.....	3
Figure 2 – System Average Interruption Duration Index (SAIDI) – EIA-0861, Annual Report	4
Figure 3 – Pad mounted Outside Plant with Batteries Shown	4
Figure 4 - Manufacturer's Published Battery Performance Data – with curve fit equation	5
Figure 5 – Tested Batteries for Temperature Effect	6
Figure 6 - Predicted Runtimes vs. Actual Runtimes	8

1. Introduction

Backup batteries are an essential part of the reliability equation for the cable outside plant architecture. When utility power fails, batteries are relied upon to power the network and helps ensure customers remain on-line. Being able to predict how long a battery can support the network is an essential planning and maintenance tool.

As one of the largest industrial Internet-of-Things (IOT) applications, Comcast has developed the unique ability to automatically test and monitor over a million installed batteries at well over a quarter of a million installations. This paper describes possible ways to use this continuous sensing of metrics from the power supplies of our outside plant network to begin developing a sophisticated diagnostic and planning tool. When finalized, this tool will use machine learning and artificial intelligence to be able to predict the expected runtime of a battery during a utility outage based on its actual load, and then to adjust this runtime prediction based on multiple key factors. The final desired result is the ability to know whether a battery is performing as expected, and to be able to track its degradation for preemptive maintenance purposes. This includes unexpected loss of battery performance from unknown factors as well as the expected losses from known factors. The status of this work is presented, showing the current predictive model, along with a brief discussion of the future work planned in this area.

2. Background

Battery systems have been used for many years to power the Cable network during a utility outage. Utilities themselves are undergoing a significant transformation, with increases in renewable generation, flexible load programs, decarbonization of commercial buildings and major capital deferrals. Throughout and sometimes due to these major strategic shifts, energy reliability and disruptions remain a significant issue that network providers must deal with. The disturbing and growing trend of electrical disturbances for the past 20 years can be clearly seen in Figures 1 and 2.




Figure 1 – Department of Energy OE417 – Annual Summary




Figure 2 – System Average Interruption Duration Index (SAIDI) – EIA-0861, Annual Report

It is relatively easy to detect an outage and to report on the duration of a battery's discharge event. However, this in-and-of itself is not sufficient to determine the health of the network. Battery manufacturers provide tables and curves for the expected discharge runtime vs. load of a new battery. Over time as the battery ages, the runtime performance of a battery will be expectedly reduced by numerous factors and conditions, while the battery remains healthy. The goal of this work was to use the existing metrics available from the power supplies to calculate the runtime of a new Outside Plant (OSP) battery, and to then adjusted the runtime based on the effects of the major, known factors. Future work will focus on identifying unhealthy batteries that have degraded faster than would be expected. This tool is expected to ultimately result in fewer truck rolls as well as enable a more effective OSP battery (see Figure 3) replacement strategy.



Figure 3 – Pad mounted Outside Plant with Batteries Shown

3. Approach

The existing set of metrics that are reported from the power supply was first examined to determine what gaps exist. This set includes the following key items.

- Power output of the power supply
- Current output of the power supply
- Individual battery voltage
- Ambient temperature
- Timestamps of all items
- Inverter status (charging, discharging, +)

Additional items known include the age of the battery, the manufacturer and the model, the number of battery strings, and the total number of batteries in the power supply. Using all of these metrics, we were able to craft a working predictive model.

The first step was to determine how the battery was expected to perform as new. Manufacturers generally provide a limited performance data set. This runtime data can be based on a constant current discharge or a constant power discharge. Using this data set, the discharge performance for each battery model in the network was plotted and curve fit to allow a calculated runtime for any load imposed on the battery. This generally follows the well-known Peukert's relationship, which is a generalized relationship between discharge load and runtime. Since the performance curve of all batteries vary based on their design, it was decided to individually curve fit the performance data for each battery model in the network. Using a standard exponential equation, a very accurate curve equation was then created for each battery model. An example of the manufacturer's data provided and the resultant curve equation for one of the battery models are shown in Figure 3.




Figure 4 - Manufacturer's Published Battery Performance Data – with curve fit equation

To use this curve equation, the actual direct current (DC) load placed on the battery is required. Unfortunately, this is not one of the measurements of the power supply. To resolve this, the efficiency of the power supply inverter was calculated using the measured alternating current (AC) voltage and current output of the power supply. By using this calculated inverter efficiency, the DC load on the battery could then be calculated. With this, we were then able to calculate a predicted battery runtime for any utility outage where the power supply transferred to battery backup.

The limitations of this preliminary equation were immediately evident. Using the performance curves from the manufacturers predicts their runtime as a brand new battery. All batteries will expectedly decline in performance over time as they naturally age. This degradation is not accounted for in the curve fit of the performance data. Additionally, temperature has a major effect on battery performance. Without accommodating for these factors, one is unable to ascertain if a battery that runs for a shorter than expected period of time is prematurely failing. A short-running battery may be defective, or it may be perfectly healthy but several years old. It could also be healthy but with a reduced runtime because it is being used in the winter in the Minneapolis area.

4. Temperature Effect

The power supply cabinets are environmentally uncontrolled, and the temperatures vary widely depending on the latitude and the season. As all students of chemistry are aware, virtually all chemical reaction rates are affected by temperature as described by Arrhenius many years ago. To complicate this temperature effect on the reaction rate, a battery's performance is the result of a two-phase reaction. The discharge reaction occurs in the liquid electrolyte phase, and then deposits onto the solid surface of the plates. Temperature will affect both rates differently depending on the battery design. The liquid electrolyte phase is further complicated due to temperature induced convective movements during extended runtimes. Finally, the surface reaction is affected by double-layer capacitive effects. For this first level model, a test program was initiated within the Comcast labs as shown in Figure 4 to experimentally quantify the total effect of temperature on the battery performance for each major battery model in the network. Using controlled environmental chambers, measured performance discharges were conducted on battery models under varying ambient temperatures.



Figure 5 – Tested Batteries for Temperature Effect

Based on these results, the Temperature Correction Factor (TCF) was determined to follow the general polynomial equation as shown in equation 1. Applying the TCF to the predicted runtime allows the overall effect of temperature for each battery model to be included in the runtime prediction.

$$\text{TCF} = A_1 * \exp(-0.05) * (\text{Temp, degrees Celsius})^2 + A_2 * (\text{Temp, degrees Celsius}) + A_3 \quad \text{Eqn. 1}$$

[A_n are experimentally derived constants]

5. Aging Effect

The second major effect on batteries is the aging effect. Industrial batteries are considered at end-of-life when their performance has degraded to 80% of their initial capacity. Beyond 80%, the change in degradation accelerates significantly, and the battery is in danger of unpredictably and sharply falling below the minimum requirements. There are two separate types of aging – calendar and cycle effects. Either of these will independently degrade the performance of a battery.

A battery has a predetermined calendar life. Even if never discharged, the internal side reactions that occur within the battery will corrode the positive grids and cause electrolyte loss. Numerous factors will affect the corrosion rate, including grid alloy, temperature, float voltage, electrolyte strength and AC ripple. These factors are generally balanced by the battery manufacturer, so that in an ideal setting, the battery will degrade to 80% capacity at the end of its published design life. Life claims vary by manufacturer, but a seven-year design life was used as an initial estimate. This is known to be a gross estimate, as the corrosion rate is known to be strongly affected by latitude/temperature. For our initial model, a straight-line estimation was used, assuming the battery loses 2.8% of its capacity each year, thus hitting its 80% end-of-life capacity at the end of 7 years.

$$\text{Calendar Aging Factor (AF)} = - (1 - 0.80) * (\text{year})/7 + 1 \quad \text{Eqn. 2}$$

The second major aging effect is due to cycling. In addition to the grid corrosion occurring during calendar aging, any discharge/recharge cycle will degrade the positive and negative plates within the battery and reduce its capacity. This is understood, but not yet implemented. The issue is that a generic model cannot be used, as the ability to cycle is dependent upon the specific design of each battery model. The types of batteries used in the network commonly are designed for approximately 200 ‘deep’ cycles, but can easily vary from 75 to 500 depending upon the internal components and method of construction. Additionally, the ability of a battery to withstand shallow vs. deep discharges varies tremendously in a non-linear manner, making it necessary to know the cycle-life curve for every battery model. This information is typically not provided by manufacturers in sufficient detail and lab testing is expected to be required. Once characterized in our labs, the number and depth of the discharge information will be collected and compiled from the field, with the intent of adding this element to a future revision of this model.

$$\text{Cycle Aging Factor} = \text{function}(\text{depth of discharge}) + \text{function}(\text{cycle quantity}) \quad \text{Eqn. 3}$$

6. Discussion/Results

The predictive model as described takes the known informational points of the power supply listed below. It calculates the predicted runtime of the power supply, based on the output load, the battery model and battery quantity. This calculation is then modified based on the temperature and the calendar age of the battery to provide a final predicted runtime.

- Battery manufacturer - recorded
- Battery model - recorded
- Number of battery strings - recorded
- Total number of batteries - recorded
- Date code - recorded
- AC voltage output - measured
- AC current output - measured
- AC power output - measured
- Temperature - measured

Actual field data was collected and compiled over a 6 week period and plotted in Figure 5. The predicted, adjusted runtimes were calculated and plotted vs. the actual runtimes. Obvious sources of error were excluded, such as missing dates, errant currents and runtime that were too short or too long to be meaningful. This is an area of ongoing work as the errors found continue to be cleaned and corrected. In an ideal world, all points would lie on a straight, 45 degree line (shown as the thin black line in Figure 5), which would indicate the model was perfectly accurate in predicting battery performance over a wide range of runtimes.




Figure 6 - Predicted Runtimes vs. Actual Runtimes

This is a very preliminary snapshot of the data, and refinements and adjustments are continuing. However, despite the substantial scatter in the data, even this preliminary snapshot is extremely encouraging. A distinct trendline shown in red in the plot shows that a correlation is evident in the runtime calculation. The deviation from the ideal 45 degree slope indicates that as the runtimes extend longer, the deviation grows from the expected. Since the model is based on the manufacturer's published data, the prediction is only as accurate as this data. One possible cause of this error shift is based on the author's experience that the published manufacturers data is often progressively more conservative for shorter runtimes. This will be verified in future planned lab testing to see if the actual performance data does indeed deviate from the published data. Another possible cause of error may be the temperature rise in the battery from ohmic heating during a discharge. During longer discharges, it is possible this internal heating could progressively improve the performance of a battery. The current model calculator neglects this effect as it corrects for only the initial temperature measurement.

In Figure 5, all points below the 45 degree line indicate batteries that ran shorter than predicted. The points very close to the horizontal axis indicate batteries that failed immediately. This could be from batteries with defects, improper connections, broken cables, etc. There currently are no predictive elements in the present calculator to identify such defects. However, from knowledge of battery failure modes, there are suspected metrics that are currently being reviewed that may identify defective batteries prior to a discharge. For instance, a large standard deviation for the batteries within a string could conceivably indicate one prematurely failing battery and could thus be a predictor of a shorter runtime. Another example is that a battery not fully charged will perform poorly. To identify batteries that are not at a full state-of-charge, it is planned to track the time and depth of previous discharges to determine the recharge efficiency. Additionally, it is expected that a battery in this condition could be identifiable by its high internal resistance, which could conceivably be seen in a low initial voltage at the onset of the outage. For this example and other suspected causes of poor performance, there are now efforts underway to use machine learning as a tool. This will allow correlation of possible groups and trends with identifiable causes.

7. Conclusion

Comcast has developed the significant ability to automatically test and monitor over a million installed batteries through its Outside Plant network. Using the data collected from this immense installed base, a runtime calculator was created to predict the runtime of outside plant batteries. This calculator is based on Peukert's law, which approximates the non-linear change in battery capacity due to changes in discharge rate. An exponential equation was fitted to each battery model's performance curve to allow an accurate prediction of the runtime of a new battery at a constant, 77 degrees F temperature. This model was then adjusted for the two major, known causes of variations in battery performance - temperature and calendar aging. With this new model, preliminary comparisons to actual field outage data shown a promising correlation between predicted and actual runtimes. Refinement of this preliminary model is continuing to identify and reduce the sources of error present. Future work will focus on additional laboratory testing of batteries to more fully characterize their performance over the range of field conditions. This includes the addition of a cycling degradation factor due to accumulated discharge events of different depths. In addition, future efforts are also planned on machine learning to associate patterns in the data with identifiable causes. (see Acknowledgements) It is optimistically expected that as this model evolves and become increasingly sophisticated and accurate, the ability to quantify and predict battery degradation will prove to be an immensely valuable tool for predictive maintenance and asset planning.

Acknowledgements

The authors would like to acknowledge our colleagues that were major contributors to the success of this effort. Matt Stehman and Chris D’Andrea were instrumental in setting up and developing the tools used in the data collection and analysis of this project. Their work has been greatly appreciated and they are expected to continue to play a major role in the machine learning efforts as this exercise continues into the future. Interested readers are encouraged to read the details of their machine learning efforts “Machine Learning and Telemetry Improves Outside Plant Power Resiliency for More Reliable Networks”, by Stephanie Ohnmacht and Matt Stehman, presented at this Expo.

Abbreviations

IOT	Internet of things
OSP	Outside Plant
DC	Direct current
AC	Alternating current
TCF	Temperature correction factor
AF	Aging factor

A Roadmap for Cable Access Reliability

A Technical Paper prepared for SCTE by

Jason Rupe

Principal Architect

CableLabs

858 Coal Creek Cir, Louisville, CO 80027

303-661-3332

j.rupe@cablelabs.com

Ron Hranac

rhranac@aol.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Definitions.....	3
3. Goals for cable network and service reliability.....	4
3.1. Measure to Manage	5
3.2. Setting Service Level Agreements	6
3.3. Service Assurance	7
3.4. Fault Management	8
3.5. Repair and Supply Chain Optimization	9
3.6. Removing Degraded or Poor Quality Components	9
3.7. Vendor and Contract Management.....	9
3.8. Network Design	9
3.9. Technology Tradeoffs and Lifecycle Management	10
4. History is Our Foundation	11
5. Conclusion.....	12
6. Acknowledgements	12
7. Appendix	13
7.1. Failure Modes, Effects, and Criticality.....	13
7.2. Service Reliability	14
7.2.1. Reliability of a service is availability	14
7.2.2. Reliability of a service is really performance.....	15
7.2.3. Customers are a mystery	15
7.2.4. Service performance.....	15
7.2.5. Measuring service performance.....	15
7.3. A proposed measurement framework for cable	16
7.3.1. Features	16
7.3.2. Goodput	17
7.3.3. Latency.....	17
7.3.4. Jitter	18
7.3.5. Packet Loss.....	18
7.3.6. Availability	18
Abbreviations	19
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1 – Various types of repair cycles coexist in operations.....	8
Figure 2 – A sample of the draft FMECA currently being built.....	13
Figure 3 – A depiction of how network performance states relate, from perfect function to intolerable or failed service.	14

1. Introduction

As the speed of access networks increases and become less of a bottleneck to providing quality service, customers turn their concerns toward reliability. But they refer to service reliability, not network reliability. Still, network reliability is a key component of a reliable service. So, what is a cable operator to do?

As the cable industry turns more attention toward reliability, we have the opportunity to lead. The reliability engineering discipline is many decades old, and has a lot of tools, knowledge, and practices that we can start from, along with our own cable industry history of successful reliability engineering. Now, service usage is different, expectations are higher, networks are built and services are provided in new and different ways, and the technology we use today is rapidly evolving. The way we assure reliability has to be different too.

This paper provides a roadmap for addressing network and service reliability for the cable industry. Instead of a complete answer, it is a roadmap for the work ahead. There are many routes to take depending on where and how far the service provider wants or needs to go. CableLabs[®] and the Society of Cable Telecommunications Engineers (SCTE) can provide the van pool for part of the journey, and there will be vendors at the rest stops to help, but the journey is for the operators to take.¹

2. Definitions

There are many sources for finding definitions of reliability and the many related terms. A few simple ones are offered here, and hopefully explain away some of the sources of confusion. Unfortunately, some of these terms have use in marketing, engineering, and non-technical contexts with different meanings. Even in an engineering use, there are often assumptions being made that make it difficult to know just what is being defined and under what context. As we apply a little focus on these definitions for our specific purpose, consider that these definitions also are the desirable properties of networks and services.

- *Reliability* as a word by itself is ambiguous, context dependent, and can mean a lot of different things depending on the situation. Consider first the perspective of the user of the word, and the context they use it under.
 - Customer – Whatever it is, it must work as I want it when I want it, without repair action on my part, so that the system is invisible to me when I use the service – this is service or use case reliability.
 - Provider – Sometimes a service provider uses this word to mean availability, suggest a lower repair rate, infer fewer customer calls, or other operational costs – this is operations reliability, or network reliability [1].
 - Academic – A more precise definition of reliability is the probability that something functions as intended up to time $t > T$ given it works at time $T = 0$. This is the reliability at time t . Note the reliability function is a decreasing function over time. Note also this says nothing about networks and services, which are repairable.

Availability is better suited for repairable items, though reliability is still relevant.

- *Availability* is the long-term percentage of time that a repairable system works. In other words, availability is the ratio of time that a service, device, or network is available for use over the total time, usually expressed as a percentage of the total time. Equally, it can be expressed as the

¹ Get it? A Roadmap for Cable Access Reliability (CAR). Did you expect a General Path Solution (GPS)?

probability that a repairable system works at some far future time. As such, availability considers the uptime, downtime, and repair time issues of the system or network as a whole. It can't tell you whether failures are frequent or infrequent, or repairs are lengthy or fast, but tells you the proportion of time that something can be counted on to work.

- *Maintainability* is the ease with which something can be maintained. Often this feature is determined by maintenance time estimates, sometimes through time and motion studies. Maintainability can include repair, but does include planned maintenance.
- *Repairability* refers to how easily and quickly a component or system can be repaired, or the property of being repairable. This term focuses on repair instead of planned maintenance, though the distinction is not always clear.
- *Survivability* is the ability of a system or network to operate under attack, and provide service in the presence of failures. Parts can fail, but the system or network still functions and provides service.
- *Resiliency* refers to failure recovery and fault tolerance, and the ability to provide service under degradation, over a broad range of demands. Degradation exists, but service functions.

Note that survivability and resiliency are related, but different in that the former refers to surviving a partial failure such as a lost link in a mesh network, while the latter refers to functioning under degradation such as ingress interference in a DOCSIS[®] network.

- *Performability* is the convolution of the performance function and the probability function of the system. It's a complicated concept, but let's think of it like this: 90% of the time, my bike works great; but 9.9% of the time, the tires are low and it is hard to pedal the bike; and the rest of the time, the bike is in the shop. If I consider that the bike with low tires performs at 50% while the fully functional bike is performing at 100%, in this simple example, the performability overall is $(0.9 * 1) + (0.099 * 0.5) + (0.001 * 0.0) = 0.9495$, which is less than 95%, even though availability is 99.9%. Think of performability as a state probability weighted performance measurement. Then realize that the probability function of the possible states, translated to the probability function of the possible performance levels, is a better measure of the experience than simple availability. If you replace the performance states with a continuous performance function, the concept still works though the math gets more complicated. But keep in mind that a single number representing performability is not as useful as the full function representing probability of performance.

Referring to the customer's definition of reliability, see that all these factors contribute to a user's perception of service or use case reliability. The unreliability of a service can be impacted by a number of performance measures as they relate to the usage or use cases associated with the service. Users are all unique, but they reveal their preferences through their product choices and willingness to pay for features; this information translates well to their perception of service friction² and thus reliability.

3. Goals for cable network and service reliability

In support of our industry's 10G Platform goals, operators have a lot of well-informed tasks to accomplish. Categorically, some of these tasks include:

² The concept of service friction in this context is new; we use it here to represent any impedance a customer experiences from using a service as intended in the desired manner.

- Assure service reliability primarily, which requires network and system reliability, availability, maintainability, and appropriate resiliency and survivability. But it also requires reliable processes, procedures, management, and more.
- Build a foundation of understanding, linking customer experience to system and network events, so operations, design, and upgrades all provide the best service possible.
- Design reliable network and service solutions, with degrees of freedom to manage service reliability, that are also reliable in executability, obtainability, etc.
- Select reliable, repairable solutions and components for given deployments.
- Create and maintain fault management that is reliable, inexpensive, and maintainable. That includes proactive network maintenance (PNM), which identifies and can be used to fix faults before customers are impacted.
- Develop operations tools that are inexpensive, reliable, understandable, and useful for proactive and reactive maintenance.
- Build intelligence to enable micro-financial decisions for preventive maintenance, technology replacement, resiliency, operations planning, etc.

3.1. Measure to Manage

Service and network reliability require well defined measures of performance that can enable management for effective results. This requires well understood service performance measures, network performance measures, and operations performance measures. When a customer experiences any service friction, that should be reflected in a key performance indicator. Aligning the measures to the customer experience is most important. It is not acceptable to answer a customer complaint with an “everything looks fine” because that suggests either you are blind to an important aspect of service, your measurements are insufficient, or your customer is wrong. The latter option is not a helpful assumption. The other two tell you improvement in your operations is needed. High levels of “no trouble found” point to the need for improvement as much as repeat trouble tickets do.

Doing all of this well requires knowledge of the failure modes, effects on networks and service, and criticality of the failure modes. A useful tool for capturing and referencing this knowledge is a failure modes, effects and criticality analysis (FMECA). Considering fault management in networks and complex systems, faults should be included with failures.

To obtain and maintain this knowledge, effective collection of components and failure modes is required. This enables analysts to determine useful corporate knowledge including what manufacturer or lot of components are not performing to specifications, what parts are wearing out, which failure modes must be addressed quickly to defend service, etc.

The most convenient example with direct application to our cable industry happens to be in this year’s Cable-Tec Expo. See [2] for this year’s Fall Technical Forum paper on applying FMECA to cable faults. Also see the Appendix of this paper for an example with explanation. This work is based on expert knowledge and is generalized for hybrid fiber/coax (HFC) networks.

But operator specific knowledge is necessary to support reliable services, so that problems specific to certain plant designs, aging or degradation, or even poorly performing components (hardware or software) can be found and addressed.

To use our cable industry’s strength of sharing knowledge and energy toward common goals, we could develop standard methods for coding repair tickets to capture failure mode and component details so

operators can fully benefit from this knowledge, and apply it to assure service. But the implementation and use of the result will still be operator specific.

The industry could also benefit by standardizing how service and network reliability are measured. Fortunately, we're well on our way in an effort through the CableLabs PNM Working Group, which is sharing the output with several SCTE Working Groups, too.

But the work is just beginning, and the industry can benefit much by continuing the effort further. We should work to specify standard ways for measuring service and network reliability including

- the measurement definitions,
- how they relate to service and network reliability,
- how to track statistics and interpret them, and
- how to set control limits, perhaps setting specification limits, too.

See the Appendix for a starting framework that could serve as the foundation. But it is only a start. As you will see in the rest of this paper, we need equivalent, supportive measurements from all aspects of network operations to fully support service reliability.

3.2. Setting Service Level Agreements

Based on existing service performance information, service level agreements (SLAs) for high end customers can be set with confidence, and even rebates can be offered at net profit. When new technology is involved, models of the resulting performance may be needed, and appropriate SLAs should be set based on the network providing the service. Fortunately, simple mathematical models are often sufficient for setting and designing services for SLAs.

SLAs should be based on customer use cases but translated to service and network measures of performance. Define the service missions and translate the measurements defined to the customer use cases. For example, consider the use case of watching a movie through video streaming, including pausing a few times, requiring several functions to work when needed for the duration; what is the resulting experience, and how does it vary by customer or network condition or resource utilization?

The SLAs must be set rationally, so that they are achievable, and demonstrable. Achievability can be validated through a model, and the model fed with field data when available. Demonstrability can be achieved through data collection and translation to the customer experience. The translation again can be achieved through a use case model. For example, the movie use case just mentioned requires high availability from the network and supporting systems, and reliable performance of the network and functions for the duration of the use. If the network availability is 99.99% (equally 0.9999), and the probability of successfully delivering the movie and needed functions for the two-hour duration is 0.99999, then the overall probability of success for that mission is approximately $0.9999 * 0.99999 = 0.99989$. If a user has this use case once a week, then the probability of not experiencing a failed attempt to watch a movie in a year is approximately $0.99989^{52} = 0.9943$; there is a good chance (0.0057) that quite a few customers (more than two in a thousand) will not be able to watch a movie at least once a year, even with these seemingly high reliability and availability targets!

3.3. Service Assurance

Low friction, high reliability service is delivered through reliable networks and systems supported by reliable, efficient operations. From a network perspective, reactive, proactive, and predictive maintenance all play a role, along with fault management.

- *Reactive*: Fast restoration first then repair, prioritized by severity of impact.
- *Proactive*: Timely repair, cost efficient, prioritized by severity and opportunity, afforded due to resiliency, with no restoration needed. Also, proactive maintenance can be thought of as fault management, as it is a mechanism to manage faults before they become failures that must be reactively addressed.
- *Predictive*: Planned maintenance to address degradation before service is impacted in any way. Predictive maintenance occurs before a fault impacts network or service performance, so it can happen ahead of proactive maintenance. For example, detecting a trend in early degradation of a particular component type can lead an operator to predictively replace those components based on useful life prediction. Prognostics and Health Management is an emerging field of research which addresses this need. But predictive maintenance can also follow from proactive maintenance, such as when additional damage is observed in the proactive repair, leading to further maintenance planning. Well planned maintenance can minimize operations costs.

Standard methods for coding repair tickets to capture failure mode and component details for service and network assurance, as mentioned previously, would help operators gain full benefit from that knowledge for superior service assurance.

Note that reliable operations can play a most important role when customer facing, because operations usually faces the customer in response to service friction. The first touch point for a customer when they experience friction is usually the call center; today that is supplemented with a software application. Behind these touch points resides all the network operations tools and back-office systems, all of which are a part of the service provided, and must reliably reduce friction for that customer. A poor experience is a failure in service, so must be addressed through rapid reactive repair. Likewise, service can be proactively and predictively repaired, too, through early detection of risk (security, privacy, fault, and failure), and continuous improvement of systems and processes.

See Figure 1 for a depiction of the various types of repair cycles which complement and assure effective operations. Note that predictive management of services includes planning and engineering, including information technology (IT), functions that engineer reliability into the solutions that deliver service, as well as predictive maintenance to replace failed systems well before they have a chance to degrade other parts or impact network functions, far ahead of impacting service. But if you wait or don't detect the problems that become faults and failures, then you can still stay ahead of service impact through proactive management, which includes fault management and PNM, plus other forms of proactivity. But if you wait further, service is impacted because the faults and failures are felt by the customer through their service experience. Reactive management requires fast, and often expensive, restoration and repair; but sometimes the repair is not as fast as everyone wants because other resources, processes and systems are reactively taxed. Spare parts supply chains may extend the restoration and repair time, as might technician availability. When service is impacted, severity should determine the restoration priorities, and repair to follow that. Note that, with proactive and predictive maintenance, restoration is not necessary. Reactive repair requires more work, higher stress, higher cost, and results in less customer happiness.

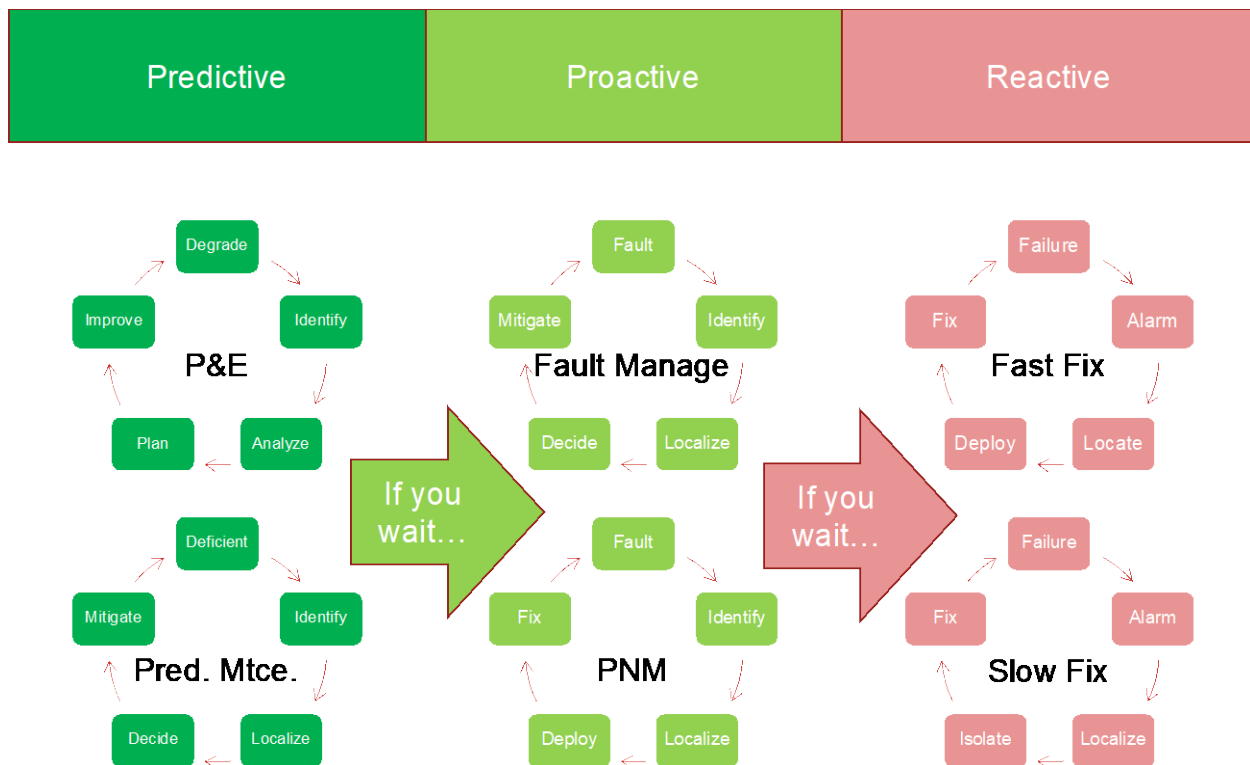


Figure 1 – Various types of repair cycles coexist in operations.

3.4. Fault Management

Knowing how a network can fail, the faults that lead to failures, and how those faults and failures are revealed in network performance can let an operator automate fault management and operate using PNM. Identifying the faults that impact service and where they come from is an important first step. How they relate to failures is important, too. An important goal in fault management is to automate as much of the fault identification, localization, and isolation as possible. And to do so reliably, which includes low false positive and very low false negative occurrences.

Based on event frequency, effect on service, and ability to test or monitor, set the policy based on established goals. Faults that are automatically mitigated can be ignored by repair technicians, but may need to be monitored by systems if they are indicators of other problems. Faults that require intervention can be handled with the appropriate repair cycle. Efficient fault management, like repair, requires an effective way to translate telemetry into action, such as the ProOps framework available from CableLabs [3], [4], [5] which provides a framework to observe (collect telemetry and information), orient (add context, assess the information, and potentially collect more information to assess), decide (translate information into faults and failures, then identify and localize faults and failures), and act (take appropriate action based on the assessment and information, with consideration of resources, priorities, schedules, etc.).

3.5. Repair and Supply Chain Optimization

With a strong handle on the priorities and planning for maintenance, operators can optimize their repair operations in many ways. Planning repairs on a longer schedule allows optimization of the travel time and distance required with maintenance, avoids unproductive technician time, and minimizes outage impact on service.

In addition, spare parts can be optimized to reduce held inventory and assure spare parts are on hand when and where needed to never adversely impact service, and never require expensive expedited shipping of parts. Critical parts necessary to correct critical failures should be readily available. Well-designed spare parts inventories can be created with lowest cost and appropriate spare parts availability. There are many applicable mathematical models available that can help operators set optimal inventory levels and policies for given targets of delivery time and probability of shortage.

3.6. Removing Degraded or Poor Quality Components

Technicians who deal with the plant all day know that some components wear out sooner than others, and some have specific faults in their design or manufacture that results in failure modes that emerge earlier or uniquely to these components. Sometimes environment has a strong influence on the early emergence of these failure modes. Temperature cycling, humidity, exposure to water, and even dry climates can impact network components differently. But even in controlled environments, poorly designed, selected, or built components can exhibit early failures which need to be addressed predictively. Early warnings from a few components can foretell the emergence of failures in the rest. As a result, tracking failure events by component type, manufacturer, age, location, and other factors can allow the operator to predict early issues and address them with predictive maintenance programs, instead of waiting for one-by-one replacement at failure. It is far cheaper to replace soon-to-fail components while doing other maintenance, to save on truck rolls and unproductive time. If such a program is required earlier than expected, a vendor management issue may need to follow, including perhaps warranty assisted replacement.

3.7. Vendor and Contract Management

Once an operator sets their goals for service, and can articulate how the network and its components translate to meeting those goals, they can align their contracts toward the goals, and even manage vendors to meet their contribution to the goals.

Component and system testing assures functionality, which reduces friction in the user experience. Testing for design and features is well established in our industry. Testing for basic features and functionality is a necessary foundation. Testing for capabilities necessary to provide specific services and features is important, too. Because long duration testing of hardware-software integrated systems is not feasible in most cases, it is important to test software well, life test hardware, and design-in system health monitoring and management capabilities for what can't be assured otherwise. Measuring early and useful life performance of components and system parts allows prediction of problems and validation of vendor performance.

3.8. Network Design

Networks should be built with performance goals in mind, and that performance should include reliability concerns as well. Doing this requires modeling of network behavior, including protection and restoration and resiliency mechanisms, for hardware, software, systems, and even people.

Network architecture will dictate allocations of service and network reliability to provide a given level of service, based on the measurements specified. How much friction-degradation and/or downtime can be allowed at, say, an optical backbone link as compared to the cable modem termination system (CMTS), cable modem (CM), or access network? Operators who are targeting service and network reliability will be collecting information and modeling to assure good decisions get made at the point of system and network design. This purposeful design enables management of network sections and knowing where to focus resources for network and service health.

Network operational cost-benefit modeling is an important component of this ability. Start with a framework for modeling the needed tradeoffs and making decisions around improvements.

This work can apply to operations design. Should a technician be sent to fix a proactive problem today, or should we wait a week in case there are more issues that can be solved with the same truck roll? If we are not sure whether a particular fault is caused by a failure mode in the home or in the yard, which technician type should be sent to keep costs lowest, and have the best chance of fixing the problem the first time?

This work can apply to decisions about the customer, too. For example, it may be worth modeling the impact of an uninterruptable power supply (UPS) in the home, or conduct a cost-benefit analysis of providing long term evolution (LTE) backup in the gateway.

But most obviously, architecture choice applies to the network decisions too. Should the operator consider media access control (MAC) manager redundancy architectures, or is a single hardware solution good enough if software and/or state are maintained redundantly? Should nodes be daisy chained in a particular deployment scenario, or is an optical ring truly necessary for the level of service we need to provide?

All these decisions need to be made with data and analysis considered, not just a gut feel, or a first-cost-driven approach. For some examples which come from our own world and are simple to use, see [6], [7].

3.9. Technology Tradeoffs and Lifecycle Management

Operators and vendors both need to benchmark the performance and reliability of existing deployed technology. This allows us all to set goals for future technology based on needed improvements or stability of reliability, availability, maintainability, survivability, and performance. Operators can model the comparison in deployed areas against the goals set by the company, and then enforce the component performance to assure goals are met as the new technology is deployed. Some high-level steps to follow:

- Benchmark existing technology
- Set goals for architectures as deployed
- Set goals and requirements for components
- Deploy and measure performance

Network components wear out. Replacing versus repairing is a decision that should consider costs, useful life, and impact to service.

At some point, an entire system or network may need to be replaced, because it has been used to the end of its useful life. This limit happens when the network or system can no longer meet its intended function in a reasonable way, or the requirements of the system or network have shifted so it can no longer meet

the current set of necessary use cases.³ See [8] for an appropriate model and treatment of the problem. Planning for wear out is important for budgeting, operations planning, supply chain management, and more.

4. History is Our Foundation

The cable access network community has given attention to reliability for decades, with considerable success. Aside from the various papers mentioned in the previous section, there are several other noteworthy works worth mention, study, and utilization.

In the late 1980s the cable industry began upgrading its networks from all-coax tree-and-branch to what is today known as HFC. Around the same time, the industry became interested in network reliability. Operators, equipment vendors, and others worked together to determine just how reliable cable networks really were and what it would take to improve their reliability. Of particular interest was whether cable networks could meet the old Bellcore “four nines” availability spec. More on that in a moment.

The topic of network reliability and availability is introduced in the context of cable networks in [10]. In chapter 20 of that book, the topics of benchmarking, definitions, calculations, redundancy, and network analysis are all discussed.

In 1992, CableLabs and several cable operators organized an Outage Reduction Task Force to “address the issues that stem from cable system outages.” The task force studied and reported on key topics relating to reliability in the cable industry [11]. CableLabs published “Outage Reduction” as a summary of the task force’s work, with chapters covering seven major topics in a large three-ring binder:

- Customer expectations, detection, and tracking
- Reliability modeling of cable TV systems
- Plant powering in cable TV systems
- Outside plant and headend protection
- Service restoration
- Cable TV system power supplies
- Power grid interconnection optimization

“Outage Reduction” was accompanied by a computer diskette with a Lotus 1-2-3 based reliability model. In addition to the published document and reliability model, CableLabs conducted half-day training workshops for member companies on the subject matter in the document’s first four chapters.⁴ Among the many recommendations in “Outage Reduction” was a critical threshold of no more than two outages in a three-month period (0.6 outages per month per subscriber) be a target for operators to achieve and maintain.

While the aforementioned guidance was considered suitable at the time for an entertainment model, any movement to telephony and data services required a higher performance threshold – hence the interest in the Bellcore four nines (99.99%) Standard Application Grade availability spec [12]. That parameter translates to no more than 53 minutes of outage time per year. Studies and analyses in the 1990s

³ Arguably, DOCSIS technology was born out of the need to meet the new set of use cases that the current network technology could not; but the network could be augmented to allow it to meet the new use cases, reusing coax.

⁴ The first four chapters of “Outage Reduction” were also published in the December 1992 through March 1993 issues of *Communications Technology* magazine.

confirmed that cable networks could meet four nines, assuming certain network architecture design criteria, device and component cascade limits, backup power and redundancy, and so forth.

While much has changed since the cable industry's earlier work in reliability and availability, some of the methods and knowledge collected form a useful foundation for today. Now that we are in a DOCSIS access network world, some of that work should be revisited.

Alberto Campos [13] in 2011 presented a paper that laid another foundation for evaluating the quality of experience (QoE). He tied performance metrics that impact QoE to the events that operators experience in the network, and the reliability of several of these features. He identified a large number of factors that contribute to the customer experience, and highlighted the importance of key elements by proposing a service availability metric. This proposed approach gathered in one place the many factors that influence service reliability and quality, plus it provided a convenient way to pull it all together into a single quantity for management. With some updating, a useful standard or operational practice could be created; with additional tailoring, operators can have a strong foundation of measurements to manage with.

Thankfully, SCTE has a new working group on Network and Service Reliability which should be the right place to tackle the new challenges, building on the foundations noted in this paper, and the papers and resources referenced by these works.

5. Conclusion

If you are an operator, you probably have been thinking while reading this paper that you already are doing all these things. You may have even participated in some of the noted foundational work. But there are at least two questions each of us should ask:

- Are we designing and executing these activities toward improved service and reliable networks and services? and
- Are we maintaining our reliability management and knowledge with changes to service, customer demand, technology changes, competition, and factors outside our control?

Operations survive by being cost focused. But that focus should be a long-term focus. And when it is, designing your operations and services toward appropriate reliability goals is your friend, and serves as the lenses for you to keep your eye on that long-term focus of managing cost as well as revenue and the drivers of both.

Once you can answer the two previous questions, you are ready to join us at SCTE's Network Operations Subcommittee Working Group 8 (NOS WG8): the Network and Service Reliability working group. See you there!

6. Acknowledgements

The authors wish to thank the numerous people who have contributed to the development of this work, especially including the hard working operator and vendor members participating in the FMECA work.

7. Appendix

7.1. Failure Modes, Effects, and Criticality

FMECA is a proven methodology for analyzing a system, process, or network for ways it can fail, determining the effects of failure, and assessing the criticality of each failure modes. The applications of this method are broad, but generally allow for appropriate design of technology to meet the requirements. An existing deployed solution is often a source of information when conducting an FMECA, either for augmenting the existing solution with improved operations, telemetry, fault management, etc.; or for designing the next generation solution for optimal performance.

A sub-team from the PNM Working Group at CableLabs has been working for many months on an FMECA that focuses on physical layer failures from the headend out to the customer, the access network. A sample of that is provided in Figure 2.

FMECA				Layer 1 - PNM first, make the layer cable later									
System	Subsystem	Component	Failure Mode	Sub-effect									
				Network Eff. Service Effect									
				Probability									
				Detection method									
Cable Access	Back office	ONT server PGP server TRIP server Config server Config files TRIP proxy											OS Spectrum Analysis
	Headend or Hub	OMES											
		Headend Combiner	Adjacent Channel Power Alignment incorrect filters, attenuators, etc. failed slope control - in line eq failed bad solder joints misconnected Cross talk - isolation										
		Headend Combiner - connector	bent mis-thread loose corroded crimp, poor fittings mechanical failure wrong type, model, poor fit poor EMC - poor connection quality weatherproofing failure or missing incorrect slope length										
		Headend Combiner - Amp	bent mis-thread loose corroded crimp, poor fittings mechanical failure wrong type, model, poor fit poor EMC - poor connection quality weatherproofing failure or missing										
		Optoelectronics											
		Local TV											
		Satellite	Source problem (satellite, uplink, programmer, etc.)										

Figure 2 – A sample of the draft FMECA currently being built.

In the figure, see hardline, connector, and part of the adapter failure modes; these components are part of the outside plant subsystem of the cable access system. Component and subsystem effects are described under the sub-effect heading, where we include several degradation causes and detectable impairment types. Under the heading of network effect, we indicate the effect each failure mode can have on the network from accelerating degradation, through signal impedance and capacity loss, to network separation. The service impact is indicated under service effect, and depicted in greater detail in Figure 3.

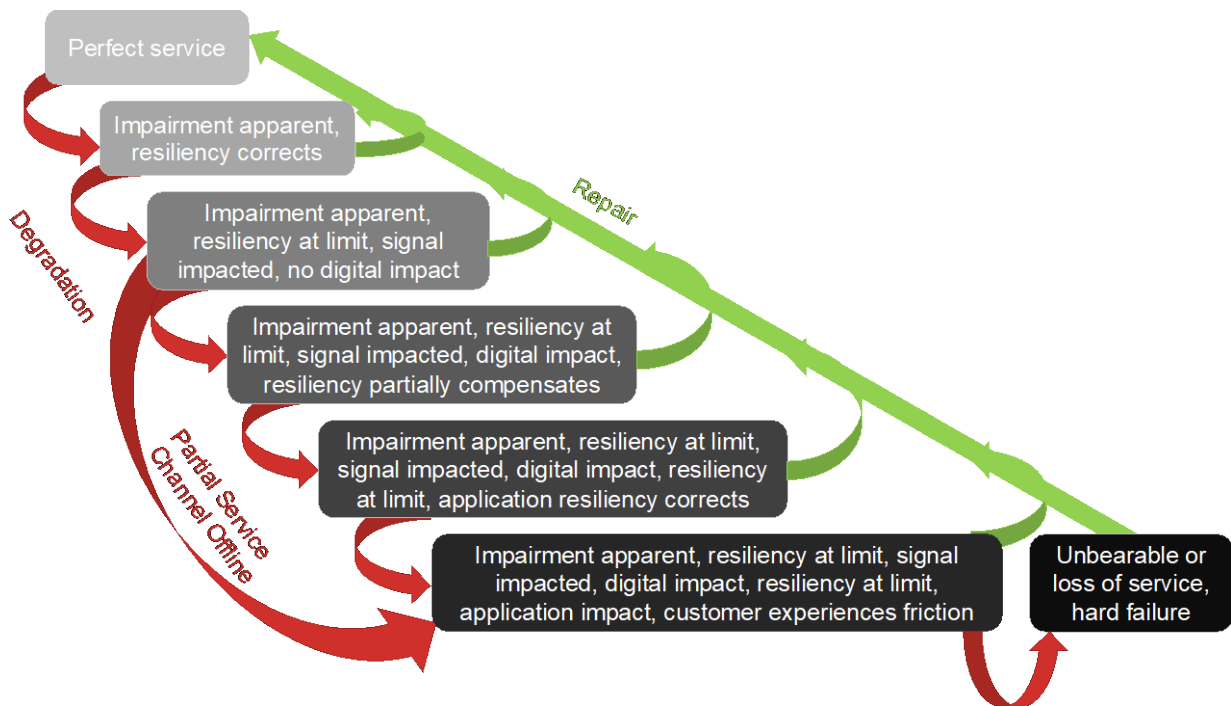


Figure 3 – A depiction of how network performance states relate, from perfect function to intolerable or failed service.

The FMECA here is focused on the physical network as it supports the mission of customer service. But some failure modes can be detected early, and some may accelerate degradation which eventually impacts service.

Because some of the failure modes can have no immediate effect on service, the FMECA documents some effects that impact the network and its components as well. PNM has identified and cataloged several signal impairments that when not too severe do not impact service, but can impact network RF bandwidth or at least foretell of future service issues.

As this work continues, we should be able to show how the repair actions relate to the failure modes, and thereby find new opportunities for improving fault management (identification, localization, and removal) in the access network.

7.2. Service Reliability

7.2.1. Reliability of a service is availability

Reliability is the probability that a system or component is working at a future time when needed. Replace component or system with product, and the intent is close but not accurate.

A service can experience downtime or degradation, but is repairable. Reliability is a non-increasing function which does not describe a repairable thing like a service. What we really want to consider then is availability, which is the probability of a service being in a functional state at some future time. And usually that future time is not defined, so we usually mean a long-term steady state of the system.

Put another way, managing service reliability requires an availability measure of the service.

7.2.2. Reliability of a service is really performance

But a service is a complicated mix of use cases and capabilities, any one or more of which could be available at different times. The service is not a single thing necessarily, and depending on how a customer wants to use it, may or may not work as intended. And if not as intended, is there value in it working in an alternate way or a degraded way? Think in terms of an email that does not go through right away, or a streaming video that takes a few seconds to buffer; does the customer notice or care, or not?

To address this issue, the concept of performability was developed many decades ago. It is a functional convolution of the performance probability distribution function with the value achieved at each performance level possible. While complicated, it does deliver a single measure of performance.

Put another way, service availability actually needs to be a performance measure of the service that includes each possible degraded performance level including complete failure, and the utility that a customer gets from the given performance level. Network reliability and performance both contribute to service availability, but they do not represent it.

7.2.3. Customers are a mystery

Note, however, that each customer is different, and the impact to their perception of performance levels varies by their situation, tolerance, emotions, and the value they put on aspects of the service. A CEO trying to close an important deal might value video conferencing much more than a student doing homework. Further, tolerance for a degraded condition might depend on the person's tolerance to previous outages, expectations of the overall quality of service, and other factors.

Put another way, the impact of service performance on individuals is highly variable and complex. The utility they get, and their overall tolerance of the experience being poor, are not easy to quantify. All an operator can do is provide the best level of service they can for the use cases known, at the price point customers are willing to pay for it.

7.2.4. Service performance

Because simple is an important goal when developing measurement systems, and recognizing that all services are a three-legged stool of cost, performance, and reliability, with cost being understood by the customer, our measure should be centered around performance and reliability, which as just described is really the aspects of performance that are delivered in each available state.

In other words, we can quantify the probability of the service delivering given levels of performance, which is what we can manage. We can seek to understand the customer, what they care about, what they are willing to pay for, and how they see competitive options. But first we must measure what we are providing in terms of service: the performance of that service as a probability space, not just an average, not an average and standard deviation, but as a probability function.

7.2.5. Measuring service performance

The task here is to identify the features of a service that describe the utility of the service to a customer. If latency is not important to, say, a webpage load, then latency measures are less important. But if the service is also being used for video conferences or video games, then latency matters, and a solid latency measurement is important to service reliability.

Examining the use cases and types of service we offer in our industry, a few basic performance measures are obvious.

1. *Goodput* – bits per second, throughput of the useful service-delivering bits on the network, assuming digital data delivery (not analog video, for example).
2. *Latency* – how long it takes to deliver each bit of data.
3. *Jitter* – packet delay variation, or how stable is the latency, and to a certain extent the goodput of the service.
4. *Packet Loss* – data that does not reach its destination.
5. *Availability State* – what is the state of performance of the service in terms of its capability?

Note that, at a packet or bit level, packet loss may be a considered measurement for either availability or as a factor for goodput or latency.

While necessary for measuring service reliability, these measures are not sufficiently described yet, and are not the end of the task for providing service.

Each performance measure statistic must be based on sufficiently detailed measurements to assure sufficient resolution of the differences in service performance levels, and measure all aspects of the performance measure. For example, measuring performance once a day at the same time every day is neither sufficient resolution nor unbiased. Measuring from the CM to the node is not an end-to-end measurement so does not represent the service experience. For understanding service reliability in sufficient depth, service providers have to design the measurement system thoughtfully, to meet their goals of continuous improvement and maintain a focus on service assurance.

But also knowing there is a problem is only the beginning; it takes more information to know the cause, locate it, and remove it from impacting service. That is the work of network operations, or network reliability, which is a key part of service assurance.

7.3. A proposed measurement framework for cable

Each service should have requirements in terms of required goodput, latency, jitter, packet loss for performance, and availability, if at all possible. Lacking a complete set of requirements, it is still incumbent on the provider to measure the service delivered. This section proposes a service availability measure based on telemetry that can be collected from the cable network.

Many of the measurements suggested here are a part of the proposed FCC 22.7, which was announced in late January of 2022. That proposal makes addressing this issue an urgent one, but also supports much of what is addressed in this document, the first draft of which formed in late 2021, with this version acknowledging what is known about the FCC proposal.

First, we need to consider several aspects of the service from a measurement feature point of view. Then we can treat each measurement in that framework.

7.3.1. Features

Several features of service reliability measurements were suggested earlier.

- Use of the measures – Depending on the various uses of the measure, the remaining features must be sufficient to address all needs.
- Bias in the measures – Because service usage varies by time of day, day of week, etc., any point measurement must be taken over a sufficiently large amount of time, with sufficient frequency, and of a sufficient sample size of the traffic to be measured.
- Resolution of the measures – The frequency of sampling must be sufficient to provide proper resolution. For example, an estimate of availability found by sampling daily will not provide good resolution for a highly available service for quite some time.
- Service level – The applications should provide the estimates when and where possible. But that is not always possible. So, service specific measures of performance at the end devices are close and sufficient for many uses. And because the operators do not manage the applications in many cases, but only the service classes as defined in DOCSIS, we should rely on these service classes first, and augment with application specific measurements when possible.
- Actual or surrogate – In some cases, we use special measurement packets to estimate actual service performance. But this method is known to be highly inaccurate and relies on a translation model that is not ideal. It is best to avoid this approach, and favor measurements on the actual traffic.

Each of these features need to be applied to each measurement. The measurements in the set are complimentary, so a complete set is needed.

7.3.2. Goodput

Throughput in terms of end user useful data is goodput. If a goodput measurement is not possible, then a throughput measurement by service type is a useful approximation because goodput can be estimated from this throughput by modeling for overhead.

When the data rate needed exceeds the capacity of the link, interface, or other component, packet queueing and congestion happen, unless discard is the only option. When the purchased data rate is not supported in the grants given by the CMTS to the CM, then applications experience latency. These understandings lead to secondary measurements for throughput.

In many cases, this measurement is used to guard against network congestion. In the access network, a simple network utilization may be sufficient. But when considering that there are customers who may be impacted by impairments and low signal-to-noise ratio (SNR), or high performing customers who also are high bandwidth users, individual CM-level throughput is important to estimate.

Recommendation: One measurement of network utilization, one measurement of bandwidth requests made and granted requests by CM, and one measurement of utilization by profile by CM. All separate for upstream and downstream.

7.3.3. Latency

DOCSIS has defined a latency measurement to support low latency DOCSIS (LLD). This measurement is taken at the CM supporting DOCSIS measurement, which is a subset of the actual experience, but a useful one nonetheless. Using this measurement for all service traffic is an excellent starting point.

Recommendation: Use the LLD latency measurement already defined for DOCSIS, and apply it to all service classes. Report by CM and service class. Augment with application-level sampling of packet delay when possible.

7.3.4. Jitter

Jitter, or packet delay variation, is the variation in the arrival of data. Some applications handle this factor through buffering, but not all applications can be made insensitive to jitter. A measurement of jitter for service types sensitive to it would be important to define. Jitter is strictly defined already, but there are alternatives that we could develop that would meet the needs for our industry.

The time between the arrival of packets would provide useful data for estimating a jitter-like measurement. A mechanism that provides the packet delay variation directly is useful if it is well defined, testable, and validate-able.

Recommendation: Use packet-level jitter measurements already defined in specifications. Augment with application-level sampling when possible.

7.3.5. Packet Loss

While packet loss is not permanent in reliable transmission protocols, thus would be reflected in terms of latency and jitter and goodput at the application layers, it is included here as it is a proposed measure in FCC 22.7.

Applications that rely on unreliable transmission protocols will not experience packet retransmission, so packet loss is an important problem and should be measured.

Applications that are latency-impacted may discard packets that are late, resulting in the same impact as packet loss. Therefore, some application consideration is important for packet loss measurement.

Forward error correction (FEC) statistics are included in DOCSIS and would be an important supportive measurement to include here, and for a DOCSIS reporting point of view would surely be more than sufficient as a measurement which can generate appropriate statistics.

However, we may need to report FEC statistics by service class or application to provide a useful measure of service reliability-availability.

Recommendation: Rely first on FEC statistics, particularly uncorrectable codeword errors. Each of these represents lost packets or data which require either application layer or protocol layer retransmission or re-requests. For reliable protocols, measure discarded packets and retransmissions. For unreliable protocols, measure lost packets. Augment with application-level sampling of packet loss when possible.

7.3.6. Availability

The overall availability of the network is an obvious, important component of service reliability. Network availability should consider cases where the user wants to use the service, but it is not available. Estimates can be obtained through polling logs from the CM, or polling state from the CMTS, or through ping-response approaches, or likely a combination.

Timeout statistics would be a useful contributor here, but there are known issues with timeouts being inaccurate as estimates of availability due to various contributing factors. However, it may serve as a

surrogate measure that could be translated into an availability estimate through a translation model, or as a contributor toward an estimate that incorporates logs and other traffic data.

Recommendation: Provide timeout statistics, augmented with logs from the CMTS and CM to estimate network availability. More detailed assessment is needed to develop the models here. Augment with application-level or device specific sampling when possible.

Overall recommendation: Measure or estimate the service experience; when insufficient, drill down toward the cause, and address the fault.

Abbreviations

CEO	chief executive officer
CM	cable modem
CMTS	cable modem termination system
DOCSIS	Data-Over-Cable Service Interface Specifications
FCC	Federal Communications Commission
FEC	forward error correction
FMECA	failure mode, effect, and criticality analysis
HFC	hybrid fiber/coax
IT	information technology
LLD	low latency DOCSIS
LTE	long term evolution
MAC	media access control
NOS WG8	[SCTE] Network Operations Subcommittee Working Group 8
PNM	proactive network maintenance
QoE	quality of experience
SCTE	Society of Cable Telecommunications Engineers
SLA	service level agreement
SNR	signal-to-noise ratio
UPS	uninterruptable power supply

Bibliography & References

- [1] R. Hranac, “Service Availability,” Communications Technology, December 2007. Available at <https://scte-cms-resource-storage.s3.amazonaws.com/07-12-01%20service%20availibilty.pdf>
- [2] M. Spaulding, L. Wolcott, J. Rupe, “Improving Operational Intelligence for Maintaining Cable Networks,” SCTE Expo 2022.
- [3] J. Zhu, K. Sundaresan, J. Rupe, “Proactive Network Maintenance using Fast, Accurate Anomaly Localization and Classification on 1-D Data Series,” 2020 IEEE International Conference on Prognostics and Health Management (ICPHM), 2020.
- [4] J. Rupe, J. Zhu, “Kickstarting Proactive Network Maintenance with the Proactive Operations Platform and Example Application,” SCTE Expo 2019.

- [5] J. Rupe, J. Zhu, “Comparison of RxMER Per Subcarrier, Bit Loading, and Impairment Driven versus Measurement Variability,” SCTE Expo 2020.
- [6] J. Rupe, “A General-Purpose Operations Cost Model to Support Proactive Network Maintenance,” SCTE Expo 2019.
- [7] N. Foroughi, J. Rupe, “Distributed Gain Architecture: Increased Performance, Decreased Power Draw,” SCTE Expo 2020.
- [8] J. Rupe, “Optimal Maintenance Modeling on Finite Time with Technology Replacement and Changing Repair Costs,” Annual Reliability and Maintainability Symposium (RAMS), 2000.
- [9] M. Spaulding, L. Wolcott, J. Rupe, “Improving Operational Intelligence for Maintaining Cable Networks,” SCTE Expo 2022.
- [10] W. Ciciora, J Farmer, D Large, M. Adams, “Modern Cable Television Technology: video, voice, and data communications,” Elsevier, 1999, 2004.
- [11] Outage Reduction Task Force, “Outage Reduction,” CableLabs Technical Report 1992.
- [12] Telcordia Technologies Generic Requirements, “Reliability and Quality Measurements for Telecommunications Systems (RQMS-Wireline),” GR-929-CORE, 2002.
- [13] A. Campos, “Holistic Approach to Evaluating Quality of Experience,” SCTE Cable-Tec Expo 2011.

A Unified GitOps Continuous Deployment Approach for Telco Hybrid Workloads

A Technical Paper prepared for SCTE by:

Stephan Salas

DevOps Engineer

Comcast, Inc.

1800 Comcast Technology Center, Philadelphia, PA, 19103

+1 267-260-0881

stephan_salas@comcast.com

Ruibing Hao, Ph.D

Distinguished Engineer

Comcast, Inc.

1800 Comcast Technology Center, Philadelphia, PA, 19103

+1 267-260-0881

ruibing_hao@comcast.com

Table of Contents

Title	Page Number
1. Abstract	3
2. Introduction.....	3
3. A High-Level Deployment Architecture for Multiple Infrastructure Platforms.....	5
3.1. GitOps Architecture	5
3.2. Multi-Platform Deployments with CI/CD and GitOps	7
4. CI/CD Listeners Implementation	8
5. GitOps Listeners Implementation.....	9
5.1. ArgoCD Implementation.....	10
5.2. Argo Workflows Implementation	11
6. Custom Kubernetes Operators	11
7. Openstack Deployment Orchestration Architecture.....	13
7.1. Core Deployment Orchestration.....	15
7.2. Auxiliary Deployment Orchestration.....	16
8. Advanced Deployment Capabilities for Web-scale Workloads.....	20
9. VoIP Stack POC Deployment using ArgoCD/Argo Workflow	22
10. Impact & Caveats of Unified Deployment Strategy.....	24
10.1. Operator Design Pattern	24
10.2. Argo Workflows Design Pattern	24
10.3. Resource Savings using Unified Platform Approach	25
11. Conclusions.....	25
12. Acknowledgements	26
Abbreviations	26
Bibliography & References.....	27

List of Figures

Title	Page Number
Figure 1 - Six Key Challenges Facing Software/Systems Delivery Teams	4
Figure 2 - CI/CD GitOps Multi-Platform Deployment Architecture.....	7
Figure 3 - CI/CD for Multiple Platform Types	9
Figure 4 - ArgoCD High-Observability Architecture	10
Figure 5 - Custom Operator Reconciliation Components	12
Figure 6 – Example Kubernetes Resource Specifications with Heatstack CRD	14
Figure 7 - Kubernetes Resource Specifications for Heatstack Deployment.....	15
Figure 8 - OpenStack Kubernetes-Operator Architecture Components	17
Figure 9 - Traefik Http-Route Integration with Auxiliary OpenStack Operators.....	19
Figure 10 - A Simplified Advanced Deployment Process for Traefik Orchestration	21
Figure 11 - Freeswitch-Openstack High Level Architecture	22
Figure 12 - Freeswitch-Openstack Deployment Process	23

List of Tables

Title	Page Number
Table 1 - Traditional DevOps vs GitOps Key Attributes.....	6
Table 2 - OpenStack Deployment Architecture Components	17

1. Abstract

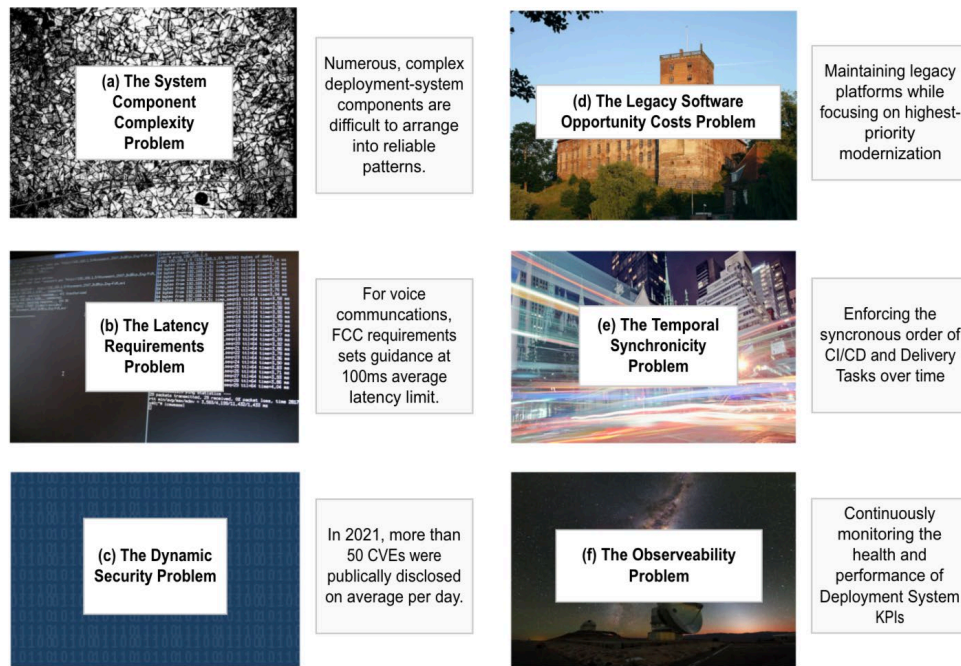
The telecom industry has moved toward a hybrid of cloud-native and virtualization technologies without a single, unified deployment approach for a variety of DevOps needs. While containerization and virtualization have both been used to solve a wide set of technical challenges in our industry, it is estimated that at least 30% of workloads worldwide still leverage virtualization technologies such as OpenStack [1]. For instance, while containerization might be advantageous for certain Layer 7 Workloads, it may be non-performant for Session Initiation Protocol (SIP) and Real-time Transport Protocol (RTP) processing needs. This trend will continue due to long-term investments to sustain both current operations and embrace more modernized ways of operating with new types of applications and infrastructure. This difference in needs across telecom organizations has led to the use of a diverse and complicated set of continuous integration and continuous delivery/deployment (CI/CD) tools and infrastructure arrangements. Unfortunately, the tendency of increasing technical-tool diversity is reflected by an increased division of organizations by technical expertise, which in turn can often-times prevent widespread adoption of modern CI/CD technologies among these organizations [2].

In this paper, we propose an approach and a framework to expand GitOps-based deployment orchestration automation into the virtualization stack, by leveraging customized Kubernetes Operators, ArgoCD, and Argo Workflows, Open Container Initiative (OCI) Containers, and Packer [3-7]. We demonstrate the feasibility and practicality of this approach on OpenStack with the help of an open source, full-stack voice over internet protocol (VoIP) implementation and Traefik HTTP Load Balancer [8]. The combination of these technologies enables several advanced deployment capabilities for OpenStack such as canary deployments and scaled rollouts. This solution has the potential to converge our industry toward a unified and modern CI/CD approach for DevOps teams and smoothen the transition towards cloud-native platforms, while helping to prevent the disorganized “tool-sprawl” [9] required to sustain both legacy and modern tech-stacks.

2. Introduction

While open-source software is “eating the world” [10] by increasing the capability of technology organizations to improve their product-competitiveness, some estimates put the proportion of costs attributed to software maintenance and sustainment at a whopping 60% of overall software project expenditures [11]. From the first monolithic inventions of a burgeoning telecommunication industry to today’s highly digital approach to scaling by means of distributed services, the burden of managing increasing levels of complexity as a result of fast-changing technologies and needs continues to be a significant cost-driver for telecom organizations [12].

DevOps professionals across our industry continue to struggle with a diverse set of organizational and technical challenges in sustaining existing operations while also looking toward future development platforms. This constant struggle to balance the two priorities can cause organizations to lose track of the core issues that are initially involved in software delivery lifecycles (SDLCs) and to get distracted by day-to-day problems. Thus, we have found that it is helpful to reframe the causes of this phenomena as a narrow set of key challenges facing modern DevOps teams:



[13]-[21]

Figure 1 - Six Key Challenges Facing Software/Systems Delivery Teams

These kinds of challenges have been denoted by organizational academics as “technical strategic bottlenecks”, which drive increased reliance on a particular cross-section of expertise within a company to solve [22]. While these bottlenecks can often be viewed as both opportunities and challenges within organizations, there are certain problems that are inherently more difficult to solve than others, such as Fig. 1(a) and Fig. 1(d). A common tactic in mitigating these issues for deployment teams is to outsource these problems into separate systems, tools, and SMEs within an organization – all of which take a non-trivial amount of productive engineering time [23]. It is thus our view that the goal of every DevOps team should ideally be to manage the widest possible problem-set with the least and simplest possible tools.

Unfortunately, some analysts estimate that software-delivery teams interact with somewhere between 20-50 different tools daily [24], which often causes the all-too-familiar “too many tabs” issue for everyday engineers. The brain-drain of having to organize individuals or teams to manage this deployment tooling complexity can cost organizations a significant amount of focus in order to silo teams by expertise and can further hamper innovation due to lack of shared understanding of deployment platforms [25].

Multiple deployment platform availability can be a blessing in disguise -- Surveys of IT leaders demonstrate that organizations tend to shy away from using multiple deployment systems due to increased costs related to learning curves and lack of expertise. As of 2020, the percentage of organizations leveraging multi-cloud technologies was nearly half of those that chose to stick with on-premises solutions [26]. Even from a CI/CD software standpoint, most workloads as of 2020 run on older, more proven tools such as Jenkins [27] and TravisCI [28], with a minority of organizations choosing newer and more powerful open-source technologies [29].

As a strategic approach to the challenge of multi-platform management and the key challenges listed in Fig. 1, this technical paper proposes a proof-of-concept system built upon OpenStack and a Kubernetes

Administrative Cluster (KADC). The central goal of our approach is to share a helpful set of solutions with engineering teams in our industry that are tasked with ubiquitously deploying, testing, and validating changes across disparate, complex systems.

Using common industry use-cases, we demonstrate both web-scale workload and VoIP workload deployments on the OpenStack platform using Git as a declarative information store. We also leverage custom Kubernetes Operators as service abstractions of core and auxiliary deployment logic. By providing a framework for how engineers might abstract deployment details in two key telecom use-cases, we demonstrate a possible solution for managing both legacy and newer forms of workloads using a single set of tools.

Such a solution is even more important today considering the tremendous amount of investment going toward infrastructure management in the telecom industry in tandem with decreased margins of traditional telecom products [30]. In response to increased technological innovation from both entrenched and burgeoning competitors, process automation is being touted as the top factor of cost effectiveness in our industry because it drives increased organizational agility and responsiveness to telecommunication customer needs [31]. With our proposed approach, we hope that we may help engineering teams deploy software more quickly and reliably, which will in turn contribute to increased resource efficiency and increased value within our larger industry.

3. A High-Level Deployment Architecture for Multiple Infrastructure Platforms

While there are many CI/CD methodologies in the open-source community that could accomplish the goals we have outlined in our introduction, the option we chose is the GitOps methodology [32]. As a self-contained approach to infrastructure management, the GitOps methodology seeks to provide high observability and re-useability of deployed state. The key initial considerations for choosing this approach were its simplicity, popularity among the industry, and ability to tie into many different deployment platforms.

3.1. GitOps Architecture

GitOps is an opinionated deployment framework with numerous valid setups used throughout our industry, however there are a few key attributes that most GitOps implementations have in common. These attributes are more clearly defined than a traditional DevOps setup that can result in a wide variety of unintended side effects and outcomes for operations teams:

Table 1 - Traditional DevOps vs GitOps Key Attributes

Deployment Question	Traditional DevOps Setup	GitOps Setup
Where is Deployment State Stored?	State is stored in databases and procedural deployment scripts.	State is implemented and stored in a declarative fashion in Git.
How is Deployment State Stored?	State may or may not be stored with its version history depending on the implementation.	State is stored in a way that supports immutable versioning and retains a complete history of changes.
How Often do Deployment State Processes Trigger?	A variety of custom-created, procedural deployment systems are typically leveraged only once to perform deployment process updates.	Software agents continuously compare a system's actual state to its desired state in order to enforce eventual consistency.
Who interacts with the Deployment State?	DevOps Engineers will typically be the primary Operators of a deployment system due to its complexity.	Developers and Code Reviewers interact with a Git interface (the "Git" in GitOps) through pull requests as a security measure to approve and commit final deployment state.

Based on our analysis, the benefits of using GitOps as opposed to traditional DevOps methods are threefold:

1. *Unification of Deployment State into a Single Location*

Compared to many of the attributes of a traditional DevOps scheme, GitOps provides a much more streamlined and unified store of application and infrastructure state. The benefits of using the Git platform as opposed to others for this purpose is perhaps the most impactful reason why the methodology is becoming increasingly popular today, with an estimated 84% of developers considering themselves as active contributors to open-source tools [33] and 92% preferring Git as their primary source control software [34]. As a single data-store of infrastructure state, the open-source Git platform also serves the principal goal in this paper of reducing excessive DevOps tooling management.

2. *Declarative Deployment State*

In addition to its unifying characteristics, GitOps also aids with separating minor, unimportant procedural details from state using the framework's declarative design. In contrast with traditional DevOps methods, specifying a group of declarative file manifests as state aids software engineers in organizing their deployments more effectively into logical units, and helps in increasing observability of what is currently deployed across different parts of their organizations.

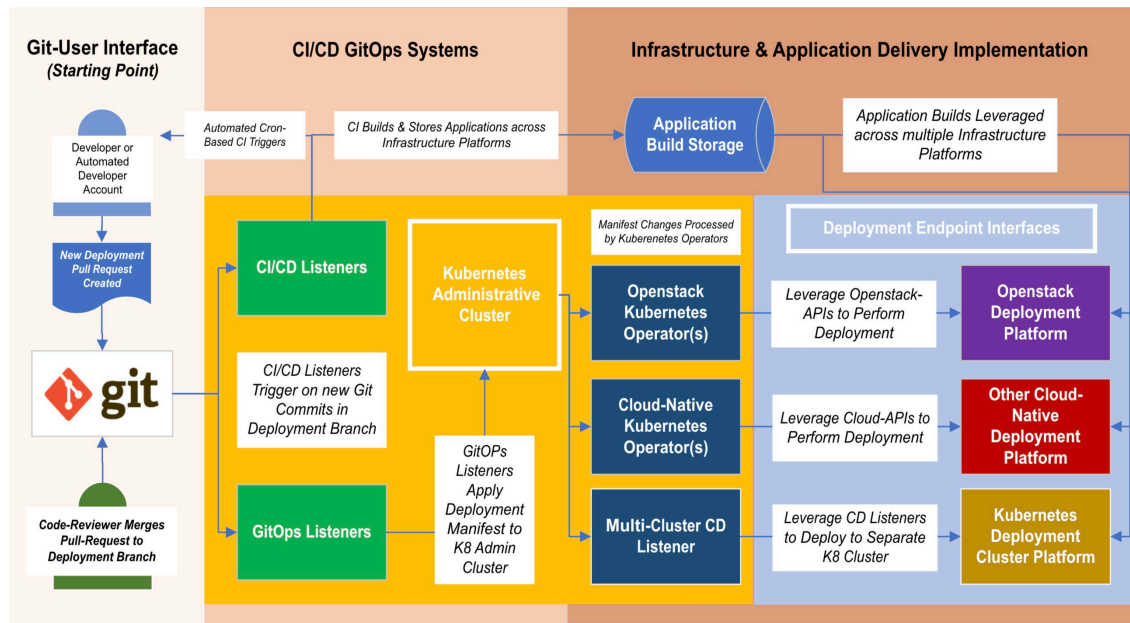
3. *Secure Deployment State*

From a security standpoint, most Git repository providers enable enterprise-grade functionality to log into repositories, pull down commits, and push new pull requests – just to name a few of the potential scenarios. In our case, by requiring pull-requests for changes to both infrastructure and application manifests, approvers can institute a code-review process that enforces certain requirements to deploy to different platforms or environments.

These three benefits serve to help mitigate many of the problems outlined in the introductory section's Fig. 1:

- (a) **System Component Complexity Problem:**
 - GitOps assists with reducing the number of components that store state in a deployment system.
- (c) **The Dynamic Security Problem:**
 - GitOps enables DevSecOps [35] via transparent declarative handling of secrets used in deployments.
- (e) **The Temporal Synchronicity Problem:**
 - GitOps can be used to aid in enforcing order of operations in deployments due to the time-based nature of Git and specifically because timestamps are associated with Git commits.
- (f) **The Observability Problem:**
 - GitOps enables high observability of deployment state in a Git repository using the Git CLI and related tooling.

3.2. Multi-Platform Deployments with CI/CD and GitOps



[36]

Figure 2 - CI/CD GitOps Multi-Platform Deployment Architecture

Combined with the popular Kubernetes Platform, GitOps provides consistent and highly observable deployment arrangements. In our proposed architecture, we introduce the unifying concept of the KADC, which is our continuous deployment orchestrator that leverages the following key components:

Git-based User Interface

The Git Interface serves as the deployment data-store that controls who has access to various repositories. Some examples of Git interfaces include on premise, cloud-hosted, and self-hosted systems. In our proposed Git interface, we envision the best-practice of developers using pull requests to commit infrastructure and application deployment changes after careful review from a group of authorized reviewers.

Continuous Integration/Continuous Delivery

The CI/CD Listeners' two responsibilities are to process application-builds within a storage context and to commit back to Git for cronjob-based automated deployment changes. For application builds, a CI/CD listener triggers build scripts upon commits to specific Git repositories. The build process may trigger testing and validation steps to verify that it is ready for deployment. Once ready, the listener pushes those built containers or images to an application build storage location.

Continuous Deployment using GitOps

GitOps Listeners enable continuous, asynchronous changes to infrastructure within targeted deployment environments based on Git repository commits. Engineers can program these systems to listen for certain changes and take automated actions. In our case, we designed the GitOps listeners system to deploy manifests (can be either YAML or JSON) into the KADC for further processing. If the manifests are Kubernetes native resources, they will be deployed directly in the target Kubernetes cluster; otherwise, these manifests describe resources deployed to non-Kubernetes platforms such as OpenStack or public clouds, and in this case KADC is used as a proxy.

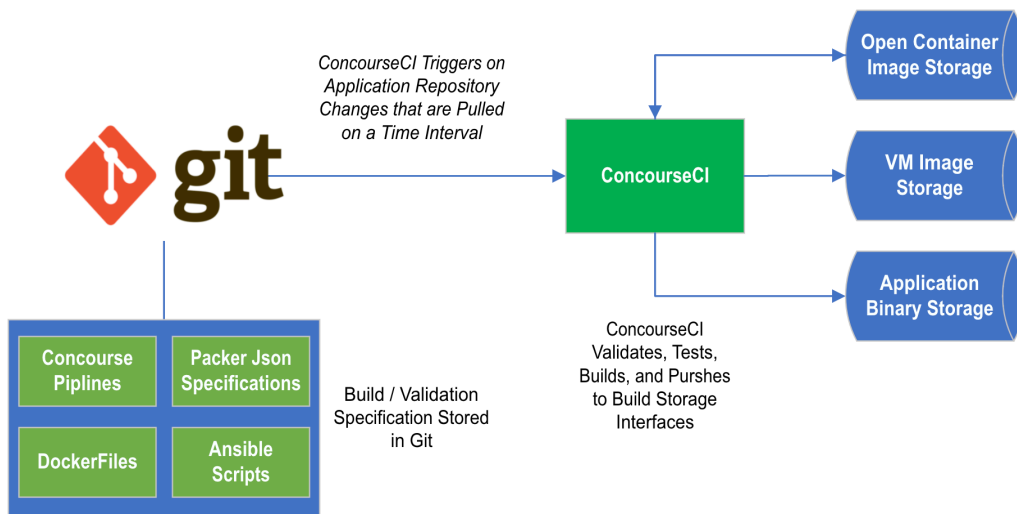
Application Delivery Implementation

The Infrastructure and Application Delivery Implementation contains the core logic to deploy manifests to various environments. This functionality is implemented with custom Kubernetes Operators in the case of OpenStack or public cloud deployments, and the built-in standard Operators in the case of Kubernetes. Depending on the *type* of manifest provided, either Kubernetes standard Operators or custom third-party Operators will handle the submitted manifest.

In all cases, Kubernetes resources hide the complexity of custom deployment logic with more simplified, declarative manifests that are easier to read than traditional procedural scripts.

4. CI/CD Listeners Implementation

While there are numerous options available that enable multi-platform CI/CD, ConcourseCI [37] and Tekton [38] were chosen for evaluation due to their maturity in the open-source ecosystem. In our implementation, we have narrowed down our scope to a ConcourseCI instance. We installed an on-premises instance of ConcourseCI inside our KADC, and also leveraged a shared instance installed on AWS. This setup acted as the “CI/CD Listener” noted in Fig. 2, and it accomplishes a variety of common application build, validation and integration tasks. The purpose of using this Kubernetes native tool is to further unify deployment state on top of a single platform. By implementing three key application-build types, as well as their corresponding testing and validation steps, we were able to integrate software build processes across the platforms listed in Fig. 2. (Kubernetes, OpenStack, and Cloud-Native) using industry-standard tooling:



[36]

Figure 3 - CI/CD for Multiple Platform Types

1. Open Container Image (OCI) Storage
 - Utilizes Dockerfiles [39] to create/push OCI container images
2. VM Image Storage
 - Utilizes Packer and Ansible Scripts [40] to create/push OpenStack VM images
3. Application Binary Storage
 - Storage of Application-Specific binaries (i.e. “.jar” file for Java [41], “.lib” for Golang [38]) that can be used by either a container or an virtual machine

Behind the scenes, ConcourseCI pushes the built container images into the OCI storage and pulls the necessary images from this storage to run each step of its pipeline. Therefore, there is a bidirectional link between ConcourseCI and Open Container Image Storage.

5. GitOps Listeners Implementation

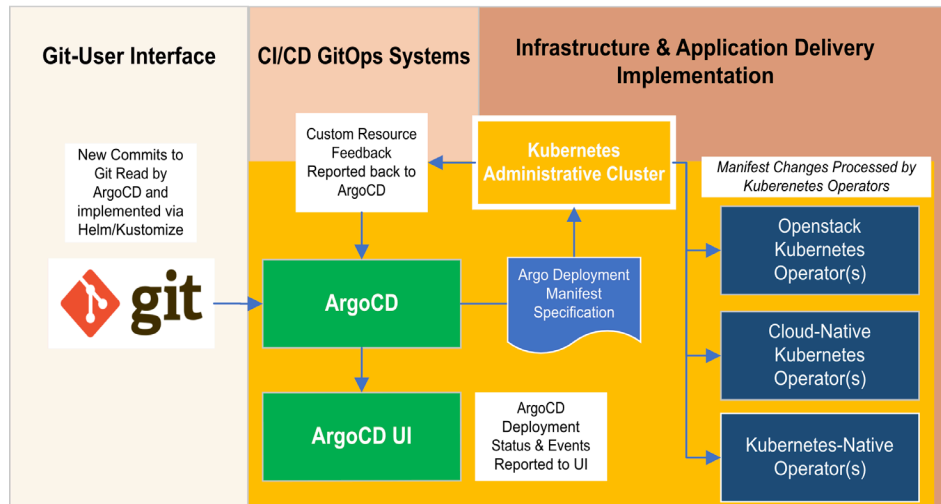
For the GitOps functionality abstractly mentioned in Fig. 2, two key opens-source DevOps components were chosen: *Argo Workflows* [5] and *ArgoCD* [4].

Argo Workflows accomplishes the orchestration of a full stack deployment in which there can be dependencies between stack components. For instance, the passing of information from one stack component to the other so the latter can be configured properly is one example of this dependency. This tool is also ideal for scheduled and repeatable deployment tasks that would otherwise burden teams with manual steps, such as scheduling the scaling up of the stack in anticipation of peak hour traffic or scaling down in the inverse case.

In comparison, ArgoCD was specifically chosen to perform the deployment of a single stack component for continuous deployment integrations. Dependencies between deployment steps are handled in this case by Kubernetes Operators that implement custom logic in a manner that is more complex than is practical to implement in Argo Workflows.

The settings for both software packages are stored and controlled by the KADC, and this design furthers our overarching goal of deployment platform unification.

5.1. ArgoCD Implementation



[36]

Figure 4 - ArgoCD High-Observability Architecture

ArgoCD Deployment Manifest Resource

An ArgoCD Application is a Kubernetes Custom Resource (CR) that reacts to changes within a specific Git repository. Within an ArgoCD Application CR, the most important attributes are: 1) the Kustomize/Helm directory to listen on, 2) the number of times to retry a CR change before declaring failure, 3) whether to auto sync changes, and 4) which deployment customization approach to use.

Kustomize [42] and Helm [43] are the two most popular open-source deployment customization approaches to use within ArgoCD at the time of this writing. While Helm is a templating solution for Kubernetes that allows for major deployment details to be highly-reusable, Kustomize is more of a patching solution that allows you to replace specific fields without a template on a more case-by-case basis.

The two CRs status fields for an ArgoCD Application are the sync and health attributes. Because an ArgoCD Application refers to a Git repository for either a Helm or Kustomize deployment, its health and sync status are “all or nothing”, meaning that for the Application to be considered fully deployed and healthy, all resources need to be successfully deployed *and* fully up to date with latest Git commit events. This information is propagated from the KADC into ArgoCD components for observability purposes.

ArgoCD UI

ArgoCD UI is the main user interface typically used to interact with for DevOps continuous delivery tasks. This interface provides a single place for both developers and DevOps engineers to view the status of their deployments of Kubernetes CRs.

The ArgoCD UI displays three key pieces of information that are useful to the end-user:

1. The health and status of each component of an ArgoCD Application
2. Kubernetes “Info” and “Warning” events associated with each component of an ArgoCD Application
3. The health & status of the overall ArgoCD Application

ArgoCD Health Checks

ArgoCD Health Checks are either Kubernetes-native (supported by ArgoCD “out of the box”), or custom-made with Lua [44] scripts for non-Kubernetes-native resources. For instance, we create several custom health checks for resources managed by our proposed OpenStack Operator.

ArgoCD Kubernetes Event Reporting

Kubernetes events are propagated to the ArgoCD UI for each CR deployed within a single ArgoCD Application. These events are published by a Kubernetes Operator and are meant to help users troubleshoot deployment issues and give more visibility into error details logged by KADC Operators.

5.2. Argo Workflows Implementation

Argo Workflows is an open-source container-native workflow engine for orchestrating tasks on Kubernetes. It is implemented as a set of Kubernetes custom resource definitions (CRDs) and its own custom Operator. The core primitive of Argo Workflows is the workflow resource, wherein each task of the workflow is implemented by a container, and the workflow itself contains a sequence of tasks with dependencies between tasks captured in a directed acyclic graph (DAG).

Argo Workflows was initially designed to run compute intensive jobs for machine learning or data processing but has been adopted to orchestrate continuous delivery tasks as well. In our proposed GitOps architecture, a workflow will be used to capture the steps required to deploy a stack of applications using a DAG. Each step of the DAG can have one or multiple success conditions that make sure this step is only considered as complete when its resources have been fully deployed and readily available. Each step is also typically responsible for the deployment of one component of the full stack, or a subcomponent of a complex component in the full stack.

Each workflow is captured as a Workflow CR in YAML format and can either be deployed to the KADC using an ArgoCD Application or directly into the KADC using a Kubernetes interface. The first approach is more appropriate when the component manifest requires much more information than is made available during deployment time. The second approach is more appropriate when the component manifest is relatively static and does not change often over time.

6. Custom Kubernetes Operators

KADC Custom Operators provide computational and logical separation of concerns for deployment to various platforms. These Operators come in the form of third-party software, and in the case of this paper, a set of custom Kubernetes Operators we developed that integrate with OpenStack. The “Operator design pattern” as described by the CNCF Whitepaper, splits functionality of CRs into controllers, which continuously reconcile changes from requested state to desired state in order to accomplish a deployment:




Figure 5 - Custom Operator Reconciliation Components

Operators can contain one or more controllers, which in turn typically manage one CR per controller. The controller is a code-bound component of a Kubernetes Operator and can be written in various languages that are compatible with the Kubernetes runtime. In our case, we chose the open source OperatorSDK Framework [45] and Golang Language [46], although any Kubernetes-compliant Operator implementation will work as well.

The fundamental primitive of a controller is the control loop, which reacts to state-change events until either the desired state is achieved (“reconcile” in Fig 3.) for a particular CR, or until it reaches a final error state. As the runtime-manager of CRs, the control-loop within a controller act as a time-bound polled reconciler of changes. It can also publish events that give a DevOps engineer further information about deployment checkpoint status or error details:

Custom Resource Definition Specification (CRD)

A custom resource is a conceptual representation of an object within Kubernetes. Included in a CRD specification are the structure of the object, the variables within the structure, and the datatype of each variable. Thus, this CRD primitive is highly configurable, and much like object-oriented programming, demands its own design considerations when developing custom controllers to manage them.

Event Types

Within a Kubernetes controller, there are three major events that are reconciled for the current state: 1) Create; 2) Update; and 3) Delete. The controller must have logic that handles each of these cases in a graceful manner for both happy-path and error-path situations so that it is reliable and feature-complete.

Current State

The current state is the latest requested state committed into the KADC. For instance, when the current state is updated, an update event is sent to the Kubernetes Operator along with the next state to be reconciled.

There are two forms of state that can change in a CR: 1) The specification field, and 2) the status field. The change of one of these fields will trigger the *Reconcile Loop*.

Reconcile Loop

The reconcile loop is a function that processes events in order to converge these events toward a desired state. Within a reconcile function, there are three possible outcomes:

1. Successfully process state event and don't requeue
2. Requeue the event to be processed again later in time due to an error scenario
3. Stop requeuing due to error scenario (with exponential backoff being an option for repeated errors)

With these three options, programmers can code controllers to be resilient to faults that may occur in the event of network issues or one-time errors, while also handling repeated errors gracefully with an exponential backoff option.

Operator services are hosted within the Kubernetes runtime as Deployments, Replica Sets, and Pods, and are easily configurable with the settings of these common Kubernetes resources. The main idea of hosting these Operators within the KADC runtime is to utilize high availability (HA) deployment capabilities inherent within the Kubernetes control plane that contributes to increased reliability of deployments.

7. Openstack Deployment Orchestration Architecture

OpenStack is a complex virtualization platform with many possible arrangements and use-cases. For deploying different kinds of workloads – namely VoIP and Web-Scale, it is important to first decide which API integration we wanted our Kubernetes Operators to interact with within the OpenStack Ecosystem. We could then decide how Operators should interface with this integration.

Upon careful exploration of available options, we decided to integrate with the popular Heat Orchestration Template (HOT or HEAT) APIs [47] because they leverage declarative resource templates that are more easily compatible with our chosen GitOps approach:

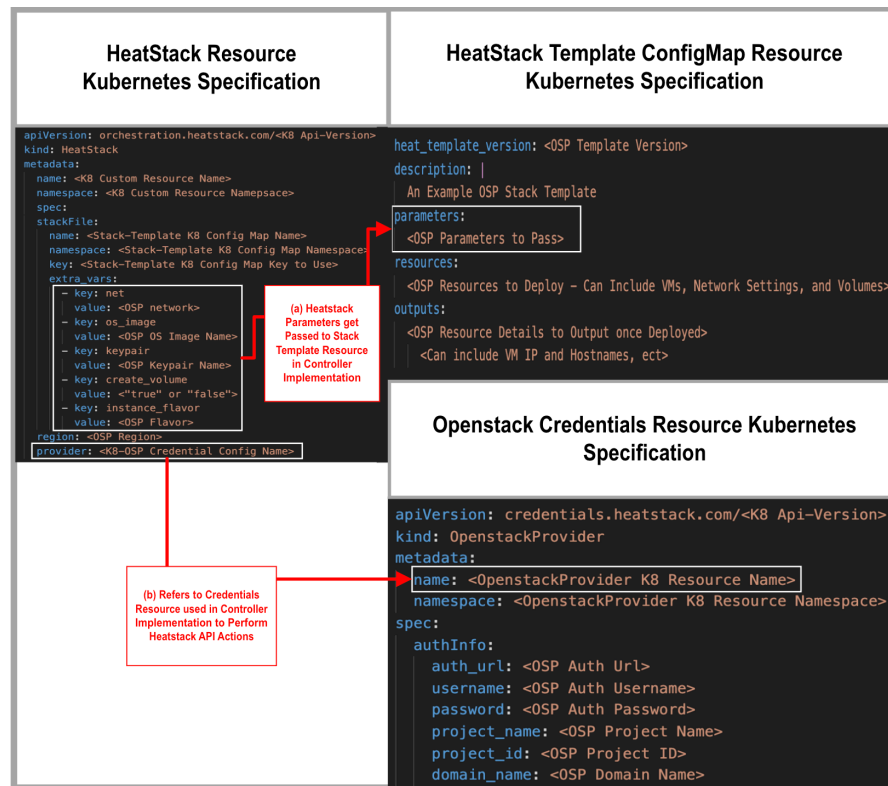


Figure 6 – Example Kubernetes Resource Specifications with Heatstack CRD

There are two essential components within our custom resource specification for OpenStack: 1) An orchestration template, and 2) An authentication template.

The resource scheme developed with the declarative specification listed in Fig. 6 allows for a high amount of flexibility in deploying key OpenStack resources. For instance, rather than having a set schema that declares which variables can be passed to a HEAT template, the “extra_vars” field in Fig. 6(a) can have an arbitrary number of parameters that work with a wide variety of Heat templates. The Heat-Template abstraction is then meant to be a highly flexible schema that serves a wide variety of use cases on the OpenStack platform.

With the “OpenstackProvider” resource listed in Fig. 6(b), we can further configure the authentication settings that we are using to interact with the HEAT APIs, which are needed for creation, update, and deletion of Heatstack CRs. This template-based approach is also applicable to other platforms such as cloud-native and other virtualized setups as well, as declarative API specifications have become popular within the IT Industry in general.

KADC resource specifications can thus be used to convert very broad requirements into specific ones, without a high amount of setup effort for simple deployments. In this example of a HEAT Template API interface, the outputs of the Heatstack Template are returned by the API, and these fields are populated as state in the Heatstack CR.

7.1. Core Deployment Orchestration

In this section we propose an opinionated way of utilizing the primitives we have designed in Fig. 6 so that we can create a base set of API interfaces with OpenStack for the HEAT Template primitive:

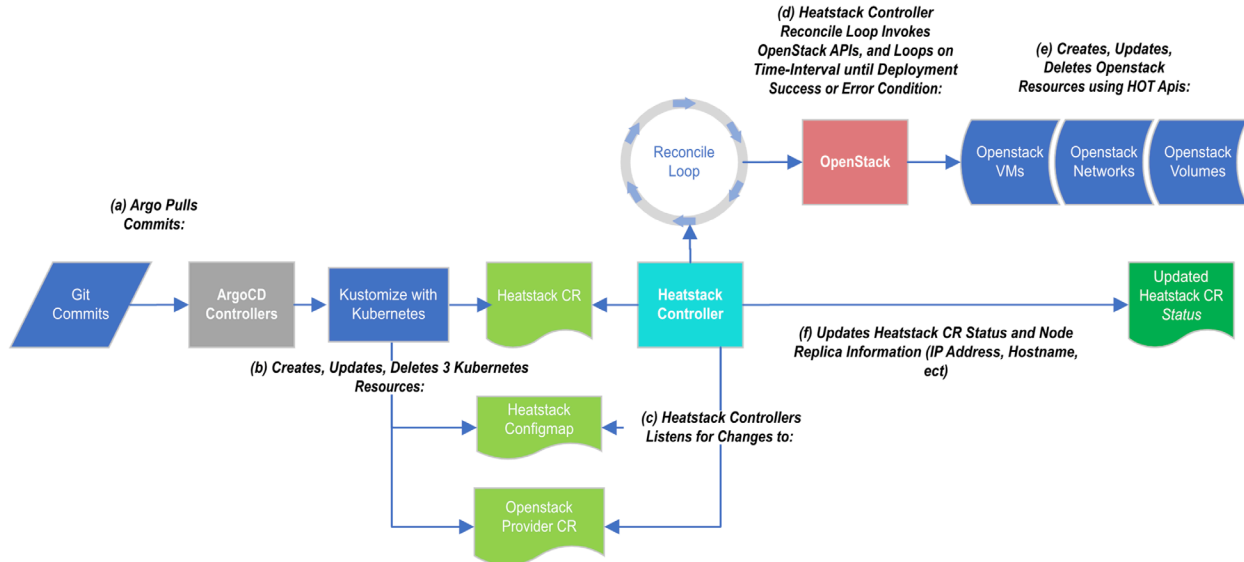


Figure 7 - Kubernetes Resource Specifications for Heatstack Deployment

With Fig. 7(a) and Fig. 7(b), we utilize ArgoCD's Kustomize interface (although Helm is also possible) to submit updates to each of the three Kubernetes CRs listed in Fig. 6. The controller then listens for change events for each of the custom resources and reacts to those changes in steps Fig. 7(c)-(e).

OpenStack API Control-Loop Logic

In the event of a bulk *create* event on resources, the Heatstack Controller reacts to the create events in Fig. 7(c), which will create the Heatstack CR, the Heatstack Config Map, and the OpenStack Provider CR as resources in Kubernetes. In step Fig. 7(d), the controller invokes an initial API call with the OpenStack API, however finalization of this API integration takes time. For example, a large-sized VM deployment can take a few minutes to a few hours to complete depending on the magnitude of scale you are targeting. Thus, it was important for us to design the Heatstack controller to poll the OpenStack resource it just created in order to validate the health of the deployment over time and report its state to our GitOps Listeners.

Heatstack Controller Status Update Implications

This idea of polling the HEAT-API for details on deployment state is fundamental to the implementation of resources that rely on the Heatstack, such as Load Balancers, TLS Certificates, and DNS entries, because each of these features rely on an up-and-running deployment. Even without these auxiliary features, reporting the status of the deployment back to ArgoCD via health checks enables better visibility of the deployment logic within the ArgoUI interface detailed in Section 5.

The status field of the Heatstack resource serves as the fundamental way that health checks are implemented, with a status field having to be “complete” for the Heatstack resource to be considered healthy in ArgoUI. Other details in Fig. 7(f) are also updated within the status field, as such:

1. deploymentStatus

- Can be either: IN_PROGRESS, CREATE_COMPLETE, UPDATE_COMPLETE, or ERROR

2. deploymentStatusReason

- a. A string field that indicates success or error reasons

3. outputs

- A list of key value pair objects that store critical data from HEAT deployments.

For each successful create, update and delete event, the deployment status gets updated with a simple tag that aids us in tracking the deployment status. The “outputs” field is updated in Fig. 7(f) with multi-VM IP and Hostname attribute details following a completed change event, which aids in health checks and further controller processing for auxiliary features.

Resilience Features in the Heatstack Controller Design

The control-loop approach for CR Status updates improves overall deployment resilience. Using control-loops, controllers can continuously integrate the latest changes committed to Kubernetes via ArgoCD while also validating previous changes or discarding them depending on the situation. This level of runtime control within our deployment implementation also allows for proper handling of errors. With exponential backoff capabilities within the control loop, we can eventually stop processing changes that are causing repeated and sustained errors over time, while also informing users through the “deploymentStatusReason” field of the underlying issue.

7.2. Auxiliary Deployment Orchestration

Operators leverage create, read, update, and delete (CRUD) APIs to orchestration HEAT-template resources. In our case, we have chosen to use a single controller within our OpenStack Operator called the “Heatstack-Controller” in order to manage these resources, while using other controllers as auxiliary integrations around this fundamental controller to supports dependent features.

Overall, the 5 key areas of solutions we incorporated into our OpenStack integration design via various APIs were:

Table 2 - OpenStack Deployment Architecture Components

Resource Deployed	API Used by Operators	Key Objective Accomplished	Leverages OpenStack API?
Heatstack(VM, Storage, Network Management)	OpenStack HEAT, aka “Heatstack Template” APIs	Orchestrate / Deploy VMs, Volumes, and networks to OpenStack using Images	Yes
Application Management	OpenStack Image APIs	Managed Packer-Built Images used by OpenStack VMs	Yes
DNS Management	VinylDNS API	Manage DNS Records in VinylDNS System	No
Certificate Management	Certificate Manager API	Manage Certificates tied to DNS Records	No
Load Balancer Route Management	Traefik Kubernetes CRD API	Manage Traefik-LB Routes Exposed on various HOT VMs	No

At the core of the deployment is the VM, Storage, and Network resource management solution, while several additional open-source ancillary components (VinylDNS, Certificate Manager, Traefik Kubernetes CRD provider) were chosen to demonstrate additional functionality [48-50]. These additional features were chosen because they are typically challenges that are faced by engineering teams in getting their applications to production and in managing complexity of common deployment setups on OpenStack. It is also important to note that these additional components may be replaced within this design with other software that has similar API functionality to support interchangeable components.

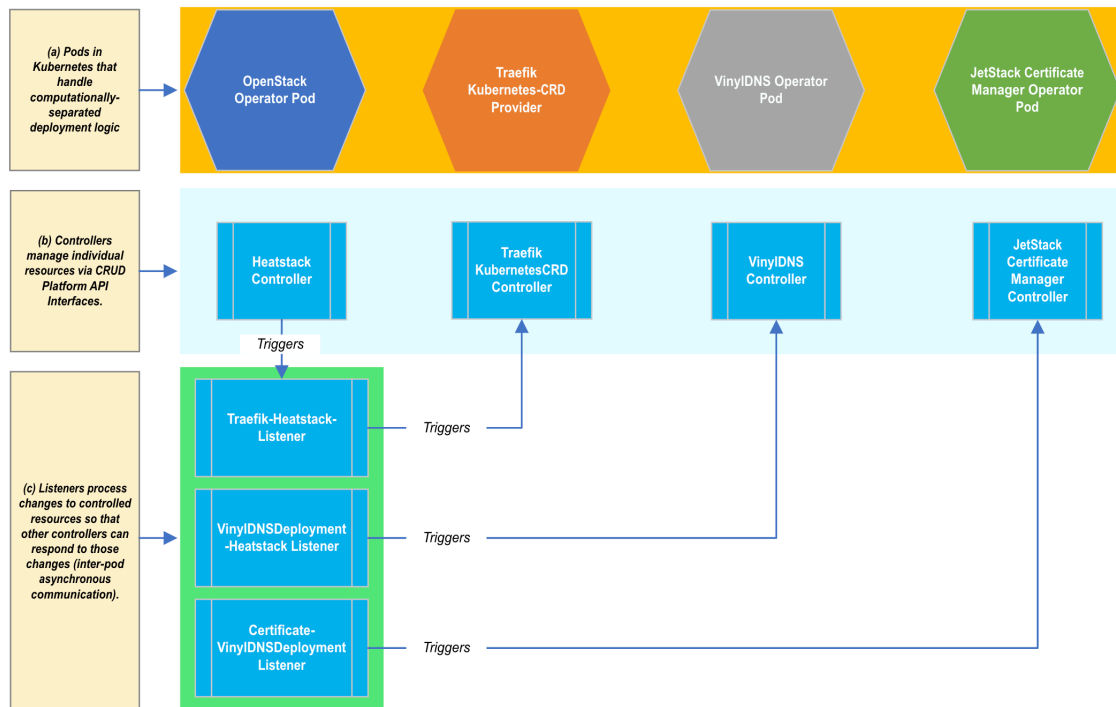


Figure 8 - OpenStack Kubernetes-Operator Architecture Components

There are two essential relationships between the resources listed in Table 2 and the Operator Architecture listed in Fig. 8, which are the *controller-resource relationship* and the *listener-resource relationship*.

In the *controller-resource relationship*, custom resources in Kubernetes are processed by various Operators to create new platform-specific implementations via their various API endpoints. In our case, the core resource being deployed is the Heatstack (described in Table 2) which supplies the declarative specification of resources supported by OpenStack. In the case of a Heatstack CR, a single controller will implement this relationship. In a similar way to the HeatStack Controller, the ancillary controllers - VinylDNS, CertificateManager, and Traefik, enable the management of DNS, TLS Certificates, and Load Balancing Routes.

In the *controller-listener relationship*, controllers listen to changes on *controlled* resources and react to those changes based on additional information captured through Kubernetes CR annotations to implement synchronous and orderly deployments. For example, the “Traefik-Heatstack-Listener” waits until a Heatstack has been fully deployed before exposing it through a Traefik Load Balancer Route using settings specified within the Heatstack CR annotations / metadata fields.

For each pod in Fig. 8(a), the controllers and listeners underneath match with the pod from a service perspective. This architecture separates concerns on both the computational level and from a logical standpoint. Using this clear separation of concerns, we built an architectural scheme with 4 pods that interfaces with OpenStack, Traefik, VinylDNS, and Certificate Manager:

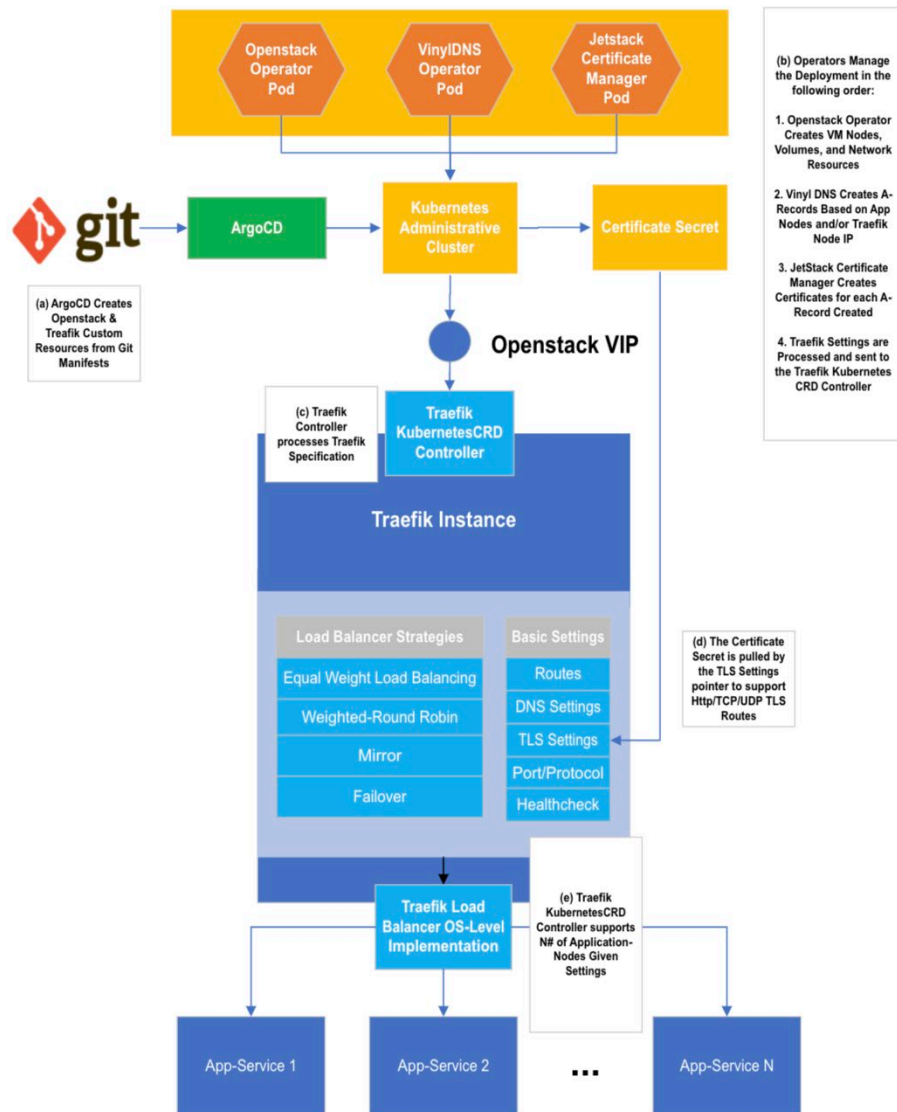


Figure 9 - Traefik Http-Route Integration with Auxiliary OpenStack Operators

Our GitOps implementation supports the synchronous processing of core and auxiliary components of a web-scale deployment as described in Fig. 9(a) and Fig. 9(b), and with all the components listed in Table 2. As mentioned earlier in this section, we leverage KADC Operators to ensure that core components are created before auxiliary components are processed. The key example in Fig. 9 is that Fig. 9(c) and Fig. 9(d) are processed after Fig. 9(b)1-3. This ensures that infrastructure VMs, VinylDNS and Certificates prerequisites are all created before services are exposed via the Traefik Load Balancer. After this initial work is performed, step Fig. 9(e) triggers with the settings passed to the Traefik KubernetesCRD provider-controller which exists on the Traefik Load Balancer instance itself, and this step processes the Traefik Route custom resources created by the Traefik-Listener controller in the KADC in order to expose groups of OpenStack VM services to load-balancer routes. The configurable settings within Traefik are listed in the Fig 9(e), split into load balancer strategies and basic settings for our ease of understanding. With a single exposed route, you can typically choose a single strategy to work with

depending on your needs, however virtually unlimited routes can be exposed with a single manifest specification in the KADC.

In addition to enforcing order in the bulk-create-case, in the case of a scalability update scenario the design also ensures zero-downtime deployments. This is accomplished with careful coding of custom KADC Operators for these common scaling cases:

1. **Scale up VM nodes:** Traefik listener will not process VM node scale-ups until the action is complete. Scaling up with the OpenStack HOT API does not delete existing nodes or recreate them, and thus this setup enables zero-downtime deployments.
2. **Scale down VM nodes:** Traefik immediately removes the necessary VM(s) from exposure in the event of a scale-down before the Heatstack-Controller deletes them. This also ensures zero-downtime deployments as well.

8. Advanced Deployment Capabilities for Web-scale Workloads

A web-scale workload is a component in many telecom products that use REST, gRPC, and other popular HTTP-based communication protocols. To demonstrate the augmentation of web-scale services with advanced load balancer strategies, we propose leveraging the weighted-round-robin strategy (listed in Fig. 9) within Traefik to allow assignment of different integer weights to groups service nodes, such that certain nodes can proportionally receive more traffic than others. Our aim in implementing this functionality is to demonstrate that our design can take advantage of the following advanced deployment capabilities not widely available in virtualization infrastructures and typically reserved for cloud-native/Kubernetes platforms:

- Blue/Green Deployment [51]
- Canary Deployment [52]
- Scaled-Rollout Deployment [53]

The general process by which the weight changes are leveraged is using the previously mentioned integration with Git in Fig. 2, which is accomplished with either manual Git commits, or with automated Git commits using a service account triggered by Argo Workflows or ConcourseCI:

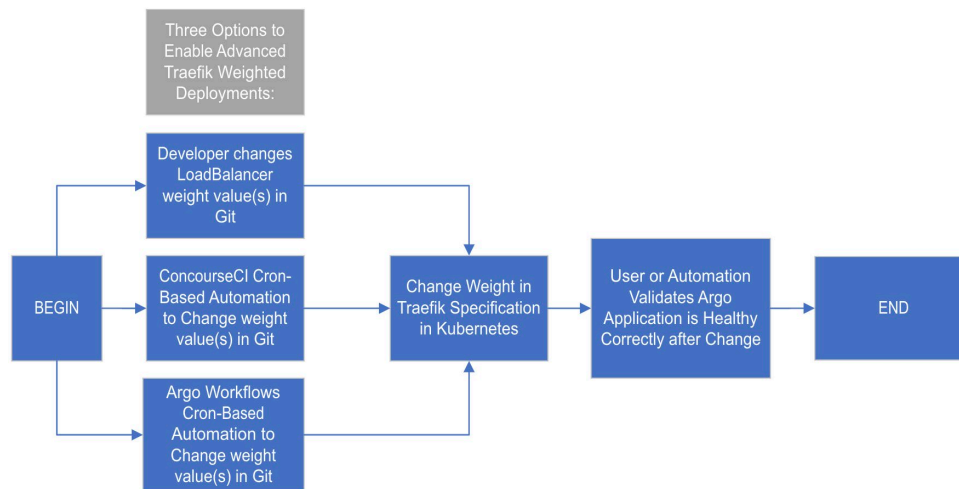


Figure 10 - A Simplified Advanced Deployment Process for Traefik Orchestration

Blue-Green Deployment

For a blue-green deployment, two separate weight changes for two groups of applications occur. Within the Traefik Load Balancer specification, we can group applications under label “A” and label “B” and give one group a weight of “1”, while the other gets a weight of “0”. This is implemented as a Kubernetes annotation on two separate HeatStack CRs, which correspond to the applications “A” and “B”. Using an atomic switch functionality implemented with our Heatstack controllers, we can ensure that A and B switch weights via the Traefik configuration in a single transaction, such that Traefik immediately switches over from A to B with zero downtime. Using the Operator framework and a Config Map lock, we can successfully process both weight changes via the Traefik Listener such that it is transactionally atomic, and thus accomplishes the goal of integrating Heatstacks with Traefik on OpenStack, while keeping processing scaling of each Heatstack independent of this functionality.

Canary Deployment

For a canary deployment, a similar technical scheme is used as the blue-green strategy in order to change the proportion of traffic going to services. As opposed to blue-green deployment where we perform an immediate switch-over from one application to another, in this case a new web application is introduced and validated over time with increasing levels of traffic so as to reduce risk of application issues in the event of a blue-green deployment.

As mentioned in Fig. 10, the process of committing weight changes to Git can be either performed through manual Git commits by a developer or Git commits via planned automation. In the case of canary deployments, it is preferable to use a platform such as ConcourseCI or Argo Workflows to perform the canary deployment changes so that incremental traffic changes can be automatically applied over time without human intervention. This was demonstrated in Fig. 10 with the “Automated Cron-Based CI Triggers”, which make it easier for planned changes to be continuously integrated based on a pre-scheduled change.

Scaled-Rollout

Scaling up / down actions can also be combined with load balancer weight changes to perform even more complex and useful arrangements of deployment schemas such as scaled-rollout strategies. As discussed earlier in Section 7, order is enforced in scale up and scale down situations between core and auxiliary components in Fig 8. This augmented functionality enables us to perform a scaled-rollout scenario with zero downtime, similar to how Kubernetes performs this same action for Replica Sets and Deployments.

With many open-source tools available at our disposal that are alternatives to Traefik, there are a whole host of different load balancers that could be very easily supported in similar ways. In fact, the popular open-source load balancers Nginx [54] and HAProxy [55] also provide Kubernetes manifest interfaces that would allow for similar scale up / down functionality, although the only caveat is that this support would require significant investment in Operator development to expand your load balancer option.

9. VoIP Stack POC Deployment using ArgoCD/Argo Workflow

As we called out in the introduction of this paper, most telecom workloads are still deployed on private on-premises cloud and running on virtualization solutions such as OpenStack. In this section, we will introduce a representative VoIP stack that is fully consisting of open-source implementations and mirroring of what a typical telecom provider might have in their network, and explain how to use the proposed CI/CD GitOps architecture to achieve end to end automation.




Figure 11 - Freeswitch-Openstack High Level Architecture

This VoIP stack consists of the following components: 1) FreeSWITCH SIP feature server [56]; 2) Kamailio SIP Proxy [57]; 3) Jambonz CPaaS solution [58].

FreeSWITCH is an open-source modular SIP feature server that can be configured in different ways to fulfill roles such as a call feature server, a voicemail server, a multi-party audio/video conference server, a media server, a transcoding SBC, or a WebRTC gateway.

Kamailio SIP Proxy is a very popular open-source SIP load balancer that can be used to front the FreeSWITCH and perform different kind of SIP load balancing. It is often deployed in a HA setup to allow for local redundancy.

Jambonz is an open source CPaaS platform that exposes Webhooks and RESTful APIs layer for invoking communication network capabilities such as call control, digit collection, or voice interaction. Underneath it is utilizing FreeSWITCH with additional modules to integrate with public cloud offering for text to speech, speech to text, or even voice dialog solution such as Google DialogFlow.

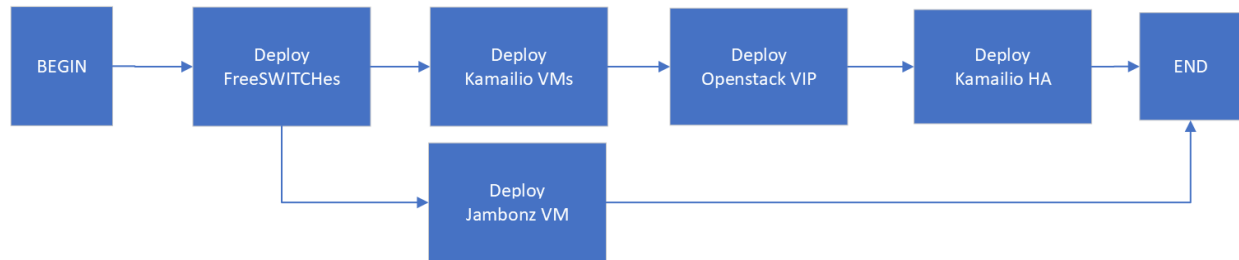


Figure 12 - Freeswitch-Openstack Deployment Process

The above Argo Workflow describes a desired deployment sequence of the VoIP stack and the dependencies between components.

- 1) The set of FreeSWITCHes will be the first group of components to be deployed; Those FreeSWITCHes will be created in Openstack using the same FreeSWITCH packer image we generated in the application CI stage but will be instantiated with the proper configuration depending on which role the VM install is going to play, for example, a voicemail server might load a voicemail FreeSWITCH configuration for it to load the required modules and the correct dialplans;
- 2) Once the IP addresses and SIP ports of FreeSWITCHes are known, we can deploy the two Kamailio SIP Proxy VMs using the Kamailio SIP Proxy packer image, each with the proper configuration to load balance SIP requests to the above FreeSWITCH instances.
- 3) Once the two Kamailio SIP proxy VM are created, the deployment process will need to acquire a VIP resource from the underlying Openstack infrastructure and modify the Kamailio SIP VM network port configuration so the VIP can be honored by those two network ports. The interaction with Openstack is done through Openstack CLI client running within an Argo Workflow container. This step is required before the next step HA configuration for Kamailio SIP Proxy.
- 4) With the VIP generated from Step 3, the deployment process can install keepalived on those two Kamailio SIP proxy VMs with the proper keepalived to monitor each other so they form a HA pair. Only one VM will be claiming the ownership of the VIP at a time, the other will only take over when the current one fails to respond to keepalive pings beyond a defined threshold.

- 5) In parallel to step 2-4, a separate deployment task will be used to kick off the Jambonz CPaaS API VM deployment using a Jambonz Packer image, and instantiated with the proper configuration to point to the FreeSWITCH IPs and Ports.

10. Impact & Caveats of Unified Deployment Strategy

In this paper, we presented two key methods of tackling the “Temporal Synchronicity Problem” in Fig. 1:

1. The Operator Design Pattern
2. The Argo-Workflow Design Pattern

The primary outcome of our efforts with these two patterns was the compression of complexity into manageable abstractions that help simplify the continuous deployment process.

10.1. Operator Design Pattern

With the *Operator Design Pattern*, we solve the “Temporal Synchronicity Problem” in Fig. 1 with independent controllers separated across fault-tolerant pods within Kubernetes. Operator architecture leverages clear separation of concerns as a key aspect of this solution, which was demonstrated in the Traefik Load Balancer example in Section 8. Combined with the GitOps methodology, the Operator Design Pattern also shines in its capability to *continuously* integrate with platform endpoints to ensure that actual state converges with latest state in Git over time. This enables easier handling of complex interactions between components, such as the interaction between OpenStack Infrastructure and DNS allocation logic where one step is dependent on another.

While it is favorable from an engineering standpoint to compress deployment complexity into the Operator Design Pattern to solve the issue in Fig. 1(a), there is admittedly a significant fixed cost in setting up the Operator infrastructure to support a new platform. In addition to this fixed cost, there are some variable costs to maintaining a deployment KADC platform Operator, however our evaluation is that this cost is minimal compared to maintaining a diversity of different DevOps tooling. The general rule of thumb we have discovered in designing and developing Kubernetes Operators for custom needs is that if you need a highly complex and continuously validated API-based integration with a new platform, Operators are probably your best option.

10.2. Argo Workflows Design Pattern

Whereas the Operator design pattern enforces deployment order through sub-patterns such as the resource-listener architecture, the *Argo Workflows Design Pattern* does so via its native “direct acyclic graph” compatibility, or DAG for short. DAGs provide a very powerful way to orchestrate both synchronous and asynchronous actions based on containerized workloads so that order may be easily implemented. Compared to the *Operator Design Pattern*, Argo Workflows does not require nearly as much custom coding and setup, as it is a template-based solution.

As most notably demonstrated with the VoIP Stack deployment in Section 9, Argo Workflows is leveraged in order to deploy various components within the proposed VoIP Stack. Due to the simplicity of this use-case where VM configuration needs to be updated, Argo Workflows shines with a short, containerized script that accomplishes a small set of tasks.

10.3. Resource Savings using Unified Platform Approach

Revisiting all the issues tackled in this paper listed in Fig. 1, the end-goal of solving these problems is to aid engineers within telecom organizations in ultimately saving time and effort in performing complex deployments. Having experimented with both manual and automated approaches in designing the proposed systems outlined in this paper, we can comfortably report those automating deployments with our proposed GitOps-based architecture speeds up our deployments significantly in the two examples we explored:

1. Web-Scale Deployment on OpenStack with Traefik
 - a. Without GitOps Automation: 1 Work Day Average
 - b. With GitOps Automation: 2 Minutes Average
2. VoIP stack on OpenStack
 - a. Without GitOps Automation: 1~2 Week Average
 - b. With GitOps Automation: 30 Minutes Average

In addition to the quantitative resource time-savings, we were also able to unify and significantly augment our current deployment capability from a qualitative standpoint. With improvements to the workflow of deployments and increased observability via GitOps Operators such as Argo Workflows and ArgoCD, we demonstrated a straightforward and streamlined method of deploying resources across platforms, while also abstracting away key details of deployment procedures from the deployer.

The OpenStack Operators listed in this paper further allowed us to augment our currently available deployment approaches with advanced methodologies such as blue-green, canary, and scaled-rollout strategies. Our goal in augmenting the OpenStack platform with the Traefik Load Balancer is to demonstrate that advanced capabilities are possible on virtualized platforms and can be reasonably implemented with speed and efficiency in mind. On the other hand, by using Argo Workflows as a multi-stack orchestrator, we demonstrated how resources could further be updated on a scheduled basis to remove critical manual steps from routine deployment situations.

11. Conclusions

One of our industry's most burdensome software-development trends is a diversity of application requirements that will continue to cause major strategic bottlenecks in deploying new types of workloads, while also driving increased long-term costs of sustaining older legacy apps of an assorted variety. In this paper, we have proposed a proof-of-concept solution that seeks to solve these problems by enabling increased platform deployment diversity and velocity by using a single administrative control-plane for both web-scale and VoIP workloads, as well as across different deployment use-cases.

With our proof-of-concept web-scale and VoIP Stack deployment approach for the OpenStack platform, we demonstrate one possible implementation for a variety of common telecom industry-specific scenarios. With a GitOps methodology for deploying OpenStack resources, we established that it is possible to create opinionated deployment abstractions that compress complexity into fault-tolerant Operator-pattern primitives, while allowing for extensibility and reusability of these primitive in an object-oriented manner. With a scheme of custom-resource organization, we implemented recognizable and easily understood constructs with general implementations of design-patterns to deploy compliant infrastructure and software across multiple platforms. By extending this approach into the realm of cloud-native and other more modern types of workloads, it is also easy to imagine adding similar ways to deploy to newer and more experimental, cutting-edge platforms through similar designs and architectures.

While the IT Industry moves toward GitOps as a popular methodology, we believe the key lesson for our telecom organizations is that it may be difficult to onboard complex applications such as VoIP Stacks to most platforms without better forms of deployment orchestration. While GitOps Operators such as ArgoCD and Argo Workflows provide a good starting point for abstracting deployment listeners, most of the DevOps work in our proposal resides with custom Operator development and sustainment for highly complex scenarios, while Argo Workflows shines in simple DAG workflow cases. This methodology provides a basis for future development with Kubernetes Operators or other chosen organizational constructs that are practical for software teams to adopt over time.

The decision to move toward multi-platform deployments will no doubt require careful thought and investment in compressing key implementation details of deployments into manageable abstractions. Within the realm of web-scale and VoIP workloads, it is important to appreciate the complexity of deployment logic, the tools available to solve common problems experienced by DevOps Teams, and the proposed solution's architectural tradeoffs. While the industry continues to move toward increased complexity of newer, more modern and powerful platforms, we should consider from a resource standpoint that managing all these systems can become unduly burdensome and subject to human error. Our hope for future research in this area is that it continues to find increasingly efficient and simplified ways to use GitOps, Kubernetes and similar tools that augment the overall DevOps experience.

12. Acknowledgements

We thank our colleagues who have contributed to our GitOps innovation work: Chris W., Johan C., Pawan T., Jaipal K., Eugene N., Robert S., and Sachin P. We also want to thank the Red Hat Open Innovation lab team for their constructive input and guidance with OperatorSDK. Finally, we want to thank John D., Arvind K., John G. and Brett S. for their continued sponsorship of the GitOps innovation work inside Comcast Communication Engineering organization.

Abbreviations

API	Application Programming Interface
CI/CD	Continuous Integration & Continuous Delivery
CLI	Command Line Interface
CR	Custom Resource
CNCF	Cloud Native Computing Foundation
CRD	Custom Resource Definition
DAG	Directed Acyclic Graph
HA	High Availability
REST	Representational State Transfer
RPC	Remote Procedure Call
RTP	Real-time Transport Protocol
SDLC	Software Development Lifecycle
SIP	Session Initiation Protocol
VoIP	Voice Over IP

Bibliography & References

- [1] Kurek, Tytus, “Openstack is Dead? The numbers speak for themselves.”, *Ubuntu*, <https://ubuntu.com/blog/openstack-is-dead>, 3 March, 2022.
- [2] Tozzi, Chrisopher, “DevOps Tools: Why We Don’t Need More CI/CD Suites”, *ITProToday*, <https://www.itprotoday.com/devops-and-software-development/devops-tools-why-we-don-t-need-more-cicd-suites>, 7 July, 2020.
- [3] <https://kubernetes.io/docs/concepts/extend-kubernetes/operator/>
- [4] <https://argo-cd.readthedocs.io/en/stable/>
- [5] <https://argoproj.github.io/argo-workflows/>
- [6] <https://opencontainers.org/>
- [7] <https://www.packer.io/>
- [8] <https://github.com/traefik/traefik>
- [9] Piscaer, Joep, “DevOps tool sprawl: is ‘tool tax’ just the tip of the iceberg?”, *Azmatic*, <https://amazic.com/devops-tool-sprawl-is-tool-tax-just-the-tip-of-the-iceberg/>, 12 November, 2020.
- [10] Andreessen, Marc, “Why Software is Eating the World”, *Andreessen Horowitz*, <https://a16z.com/2011/08/20/why-software-is-eating-the-world/>, 20 August, 2011.
- [11] Wood, David, “Metadata Foundations for the Life Cycle Management of Software Systems”, *David Wood PhD Thesis*, https://www.researchgate.net/publication/43496613_Metadata_Foundations_for_the_Life_Cycle_Management_of_Software_Systems, December, 2008.
- [12] “Digital transformation for 2020 and beyond -- A global telecommunications study”, *Ernst & Young*, https://assets.ey.com/content/dam/ey-sites/ey-com/en_gl/topics/tmt/tmt-pdfs/ey-digital-transformation-for-2020-and-beyond.pdf, 19 February, 2019.
- [13] “- complexity [4]” <https://wordpress.org/openverse/image/a346ff48-ea6b-4f5b-80d6-ecd9316c9fe4> by nerovivo is licensed under CC BY-SA 2.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.” 6 March, 2007.
- [15] “4.29M/s and 4ms latencies” <https://wordpress.org/openverse/image/ed1ad5c1-3b0c-4206-bb3d-780d5209b248> by Kai Hendry is licensed under CC BY 2.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>.” 13 December, 2008.
- [16] Pg. 23, “FCC Fact Sheet - Measuring CAF Recipients; Broadband Performance – Order on Reconsideration – WC Docker No. 10-90”, *Federal Communications Commission*, 4 October, 2019.

- [17] "File:Koldinghus - Old castle in Kolding - Denmark 017.jpg" <https://wordpress.org/openverse/image/b7e220a0-4208-467b-be5a-7521f3954d15> by S.Juhl is licensed under CC BY-SA 3.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/3.0/?ref=openverse>.
- [18] "Computer Security - Protect Data - Computers" <https://wordpress.org/openverse/image/464bebf1-9575-4018-b8b6-990953dd0cb0/> by perspec_photo88 is licensed under CC BY-SA 2.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/2.0/>, 10 June, 2015.
- [19] Henriquez, Maria, "2021 Breaks the record for security vulnerabilities", *Security Magazine*, <https://www.securitymagazine.com/articles/96668-2021-breaks-the-record-for-security-vulnerabilities>, 9 December, 2021.
- [20] "New York City - Timelapse on Vimeo by stimul" by Retinafunk is licensed under CC BY-SA 2.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/2.0/?ref=openverse>.
- [21] "La Silla Dawn Kisses the Milky Way" <https://wordpress.org/openverse/image/a7619e1e-96e1-4cf1-88d2-e9a95988c2b4/> by European Southern Observatory is licensed under CC BY 2.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/2.0/?ref=openverse>, 22 March, 2016.
- [22] Baldwin, Carliss, "Design Rules, Volume 2: How Technology Shapes Organizations: Chapter 7 The Value Structure of Technologies, Part 2: Technical and Strategic Bottlenecks as Guides for Action", *SSRN Electronic Journal*, 10.2139/ssrn.3270955, January, 2018.
- [23] Baldwin, Carliss Y., Design Rules Volume 2: Chapter 16—Capturing Value by Controlling Bottlenecks in Open Platform Systems. Design Rules Volume 2: How Technology Shapes Organizations, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3482538, 10.2139/ssrn.3482538, 7 August, 2021.
- [24] McKendrick, Joem "The snags holding back DevOps: culture, delivery, and security", *ZDNet*, <https://www.zdnet.com/article/devops-is-a-mixed-bag-so-far/> 15 March, 2021.
- [25] Baldwin, Carliss Y., Design Rules Volume 2: Chapter 16—Capturing Value by Controlling Bottlenecks in Open Platform Systems. Design Rules Volume 2: How Technology Shapes Organizations, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3482538, 10.2139/ssrn.3482538, 7 August, 2021.
- [26] Pg. 4, "CNCF Survey 2020", *Cloud-Native Computing Foundation*, https://www.cncf.io/wp-content/uploads/2020/11/CNCF_Survey_Report_2020.pdf, 17 November, 2020.
- [27] <https://github.com/jenkinsci/jenkins>
- [28] <https://github.com/travis-ci/travis-ci>
- [29] Pg. 12, "CNCF Survey 2020", *Cloud-Native Computing Foundation*, https://www.cncf.io/wp-content/uploads/2020/11/CNCF_Survey_Report_2020.pdf, 17 November, 2020.
- [30] Gaibi, Zakir, Jones, Gareth, Pont, Pierrem, and Vaidya, Mihir, "A blueprint for telecom's critical reinvention", *Mckinsey & Company*, <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/a-blueprint-for-telecoms-critical-reinvention>, 28 April, 2021.

[31] Pg. 12, “Digital transformation for 2020 and beyond -- A global telecommunications study”, *Ernst & Young*, https://assets.ey.com/content/dam/ey-sites/ey-com/en_gl/topics/tmt/tmt-pdfs/ey-digital-transformation-for-2020-and-beyond.pdf, 19 February, 2019.

[32] <https://www.gitops.tech/>.

[33] Grams, Chris, “3 charts that show how open source developers think”, *opensource*, <https://opensource.com/article/20/6/open-source-developers-survey#:~:text=The%20vast%20majority%20of%20developers,be%20active%20open%20source%20contributors>, 6 June 2020.

[34] Lindberg, Erica, “Global Developer Survey reveals need for more collaborative workflows”, *Gitlab Blog*, <https://about.gitlab.com/blog/2016/11/02/global-developer-survey-2016/#:~:text=TL%3BDR%3A%20New%20survey%20shows,work%3B%2092%25%20prefer%20Git,>, 2 November, 2016.

[35] <https://www.devsecops.org/>

[36] “File:Git-logo.svg. (2022, February 14)”, *Wikimedia Commons, the free media repository*. <https://commons.wikimedia.org/w/index.php?title=File:Git-logo.svg&oldid=629775061> is licensed under Attribution 3.0 Unported (CC BY 3.0), To view a copy of this license, visit <https://creativecommons.org/licenses/by/3.0/>, 14 February, 2022.

[37] <https://concourse-ci.org/>

[38] <https://tekton.dev/>

[39] <https://docs.docker.com/engine/reference/builder/>

[40] https://docs.ansible.com/ansible/latest/user_guide/playbooks_intro.html

[41] <https://www.java.com/en/>

[42] <https://github.com/kubernetes-sigs/kustomize>

[43] <https://github.com/helm/helm>

[44] <https://github.com/lua/lua>

[45] <https://github.com/operator-framework/operator-sdk>

[46] <https://go.dev/>

[47] https://docs.openstack.org/heat/rocky/template_guide/hot_guide.html

[48] <https://github.com/vinyldns/vinyldns>

[49] <https://cert-manager.io/>

[50] <https://doc.traefik.io/traefik/reference/dynamic-configuration/kubernetes-crd/>

- [51] Fowler, Martin, “BlueGreenDeployment”, *MartinFowler* <https://martinfowler.com/bliki/BlueGreenDeployment.html>, 1 March, 2010.
- [52] Fowler, Martin, “CanaryRelease”, *MartinFowler* <https://martinfowler.com/bliki/CanaryRelease.html>, 25 June, 2014.
- [53] <https://kubernetes.io/docs/concepts/workloads/controllers/deployment/#scaling-a-deployment>
- [54] <https://kubernetes.github.io/ingress-nginx/>
- [55] <https://haproxy-ingress.github.io/>
- [56] <https://github.com/signalwire/freeswitch>
- [57] <https://github.com/kamailio/kamailio>
- [58] <https://github.com/jambonz>

Advanced Workflows for Onboarding

A Technical Paper prepared for SCTE by

Peter Cline

Principal Engineer II
Comcast
1800 Arch St. Philadelphia, PA
215-917-2645
Peter_cline@comcast.com

Priyasmita Bagchi

Sr. Manager Software Engineering
Comcast
1800 Arch St. Philadelphia, PA
priyasmita_bagchi@cable.comcast.com

Nirav Dave

Principal Architect II
Comcast
183 Inverness Dr W, Englewood, CO
720-236-5223
Nirav_Dave@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Onboarding Process	3
3. Leased Gateway Orchestration.....	4
3.1. Observability.....	5
3.2. Orchestration Modification	5
3.3. Workflow Testing and Deployment	6
3.4. Workflow Composability.....	6
4. Workflow Orchestration Architecture.....	6
4.1. Architectural Constructs	7
4.1.1. Task	7
4.1.2. Workflow Definition	8
4.1.3. Workflow Engine	8
4.1.4. Workflow Integrator.....	8
4.2. Benefits	9
4.2.1. Metrics and observability	9
4.2.2. Reusability.....	10
4.2.3. Development agility.....	10
4.2.4. Ease of programmatic client integration	10
4.2.5. State management.....	10
5. Case Study: Mesh and Advanced Security Firmware Agent Enablement.....	11
6. Unified Client Interface	13
7. Conclusion.....	14
Abbreviations	15

List of Figures

Title	Page Number
Figure 1 – Leased Gateway Workflow Orchestration	5
Figure 2 – Workflow Orchestration Architecture	7
Figure 3 - Workflow Orchestration Architecture Implementation	9
Figure 4 – Conditional Agent Enablement in the Monolithic Orchestration	11
Figure 5 – Agent Enablement with the Workflow Orchestration Architecture.....	12
Figure 6 – Channel, Product Specific Client Onboarding Interfaces	13
Figure 7 – One Onboarding Hub for all Channels and Products	14

1. Introduction

Every brilliant network deserves a brilliant onramp, one which makes it simple and easy for customers to get quick, ready- access to the services for which they are paying. This first interaction with services and products will leave a lasting impression that will be difficult to change if it isn't positive. At Comcast, teams are intently focused on ensuring that this onboarding or first-time user experience (FTUE) is frictionless and positive for our customers. We look to minimize customer interactions stemming from difficulties with onboarding and to direct as many folks as possible into the self-install installation route. This paper examines how we are using cloud native such as workflow orchestrators and Functions as a service (FaaS) to realize this goal. We will examine how previous paradigms employed for onboarding provided a foundation for the new workflow orchestration architecture presented here and helped propel us in that direction. From the perspective of software development, we wish to develop platform services that are robust, highly observable, scalable, quick and easy to modify and deploy, while providing the best customer experience.

2. Onboarding Process

The process of onboarding IP (Internet Protocol) gateway devices has evolved with changing technologies and business opportunities. Customers used to rely heavily on technicians to help onboard their equipment. Interactive Voice Response (IVR) systems were available to help customers who needed it. There has long been a web interface used by both customers and technicians to facilitate the onboarding process. With the advent of mobile applications, new business opportunities arose. Comcast introduced mobile applications and recognized the potential of the superior user interface there to improve the onboarding experience. This was also about the time when our advanced xFi gateways running the Resource Development Kit – Broadband (RDK-B) firmware were introduced. These gateways would necessitate changes to the existing onboarding process, as additional backend services were now involved. Initially only IP gateway devices that were leased to customers ran the RDK-B firmware and consequently our focus was solely on these devices. Onboarding functionality was implemented as part of the single Application Programming Interface (API) supporting most of the functionality for what was then called the xFi app. This paradigm served us well for a long time.

The success of the onboarding process in the xFi app encouraged the business to seek additional opportunities. Soon discussions were underway about how we could support customers who chose to bring their own device, so-called customer owned and managed (COAM) devices, when subscribing to Comcast High Speed Data (HSD) service. A growing percentage of Comcast broadband customers opt for this route and ideally their onboarding experiences are as consistently positive as those experienced by customers who lease gateways from Comcast. Concurrently, software engineers began to recognize the drawbacks of operating software as large applications performing many different functions and the era of microservice architectures dawned. As we contemplated adding support for COAM devices to the xFi application, we began examining how we could leverage the benefits of the emerging microservice architectures at the same time.

The decision was made to build the COAM activation as a collection of FaaS components that would then be orchestrated by an external workflow engine. Separating concerns in this way affords us a host of benefits which we shall examine in detail.

Onboarding requires asynchronous execution of a series of tasks to activate service, update core device configurations such as wireless fidelity (wi-fi) radio credentials, provision or update status of the device in a central device repository, and conditionally apply other device configurations. All of these functions could potentially be performed by

discrete functional units as suggested by microservice architectures. The decision was made to build the COAM activation as a collection of FaaS components that would then be orchestrated by an external workflow engine. Separating concerns in this way affords us a host of benefits which we shall examine in detail.

The COAM onboarding workflow was successfully implemented in this fashion and introduced into the xFi app. We were now able to provide the superior mobile application onboarding experience to both leased and COAM customers, so when the business value of facilitating onboarding of gateways for customers in Multi Dwelling Units (MDUs) via the mobile application became apparent we were well positioned to tackle that work.

3. Leased Gateway Orchestration

The introduction of leased gateway orchestration to the xFi mobile application was a large success. We were able to gain more insight into how the onboarding process was performing, identify opportunities to improve the customer experience and the orchestration of all the requisite back-office processes. Most significantly, customers could onboard their devices out-of-band, meaning without being connected to the wi-fi network broadcast by the IP Gateway itself. This opened many opportunities for an improved user experience in the xFi app and facilitated the increased use of Self-Install Kits (SIKs) for IP Gateway onboarding. SIKs meant fewer technicians visiting homes to facilitate the onboarding process. The leased gateway onboarding functionality was part and parcel of the platform API supporting the xFi mobile application. This single large software application was in line with how most folks were building software and made it easy to deploy and operate. Figure 1 illustrates this leased gateway workflow orchestration.

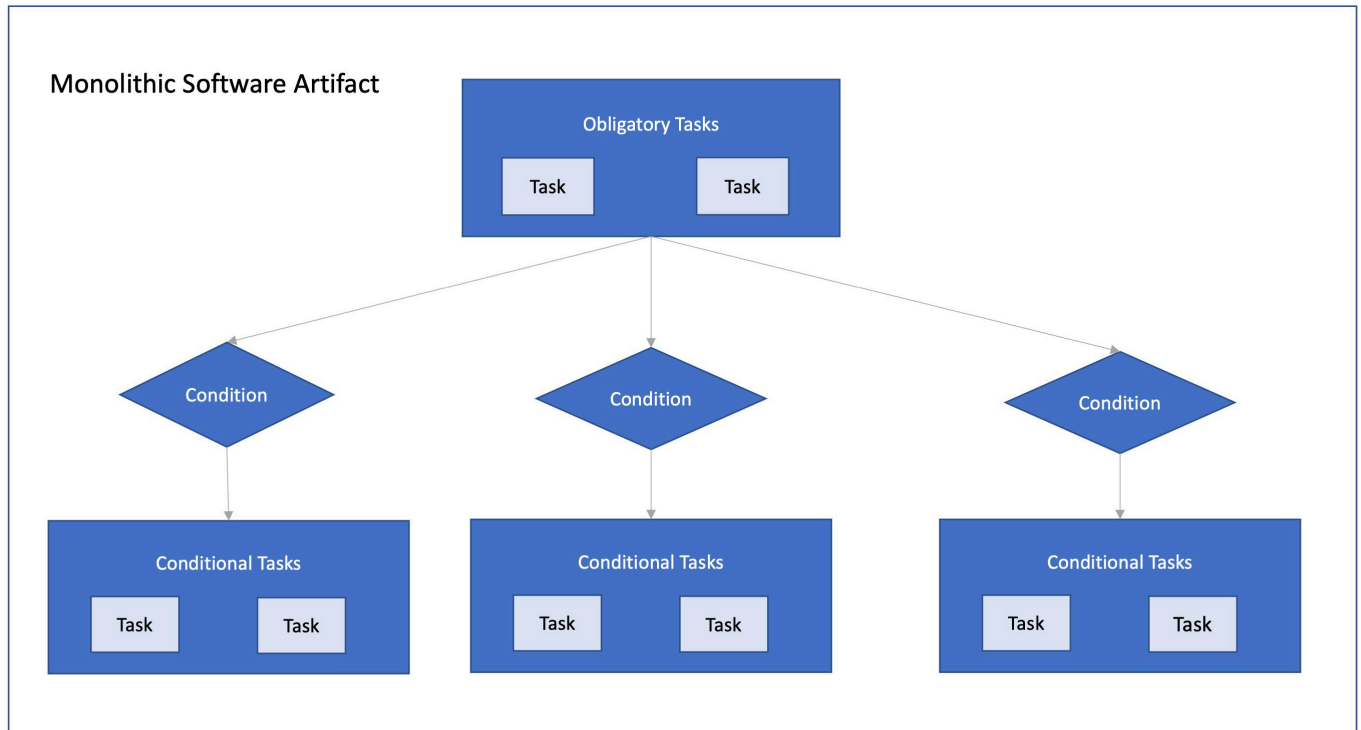


Figure 1 – Leased Gateway Workflow Orchestration

Our leased gateway onboarding implementation provided some great benefits as detailed here.

3.1. Observability

Because we were orchestrating the whole onboarding process from a single software workflow, we had great visibility into the whole process. A home-grown system for distributed tracing, named Money, which Comcast later open-sourced, allowed us to follow a request from the mobile application client through our orchestration to the various back-office services involved in the onboarding process. As long as the leased gateway flow was the only one supported this process worked well and gave us much insight into the onboarding workflow. As we added additional workflows for COAM and MDU devices we realized we had an opportunity to do better still and gain insight into how common tasks performed in aggregate across all the onboarding workflows.

3.2. Orchestration Modification

Once established, the leased gateway onboarding workflow changed infrequently. Within the xFi platform API, the code for the workflow was well encapsulated and could be modified when need be to support evolution of the onboarding process. The addition of the COAM workflow provided the impetus to extract the workflow definition from the code itself so that it could be managed and evolved independently of the code that implemented the business logic described by the workflow definition.

3.3. Workflow Testing and Deployment

While the leased gateway onboarding workflow was the only one we supported, testing and deployment was straightforward. Changes would be made, and the workflow would be tested via a mixture of automated testing and manual testing using the client interface. As we began to consider COAM onboarding, we looked for opportunities to manage the two workflows separately so they could be tested and deployed independently and changes to one workflow would not necessitate any testing of the other if the changes did not apply.

3.4. Workflow Composability

In the monolithic orchestration in Figure 1, the conditional tasks shown are responsible for additional configurations that fall outside of the primary onboarding function as previously discussed. The gateway device is onboarded and functional from the customer's perspective after the primary tasks in the orchestration have been successfully completed. The advent of microservices architecture and the emerging architecture for COAM onboarding would provide the opportunity to treat these conditional workflow tasks as independent workflows whose execution could be done independently of the primary onboarding workflow. This separation would allow each of the workflow to be given the retry semantics they required and provide greater clarity as metrics could be separated and tallied individually for each workflow.

4. Workflow Orchestration Architecture

The introduction of COAM onboarding in the xFi mobile application gave us the opportunity to implement the architecture we'd been formulating, built upon an external workflow engine and independently deployable FaaS components. This paradigm decouples the workflow orchestration from the task implementation. An externalized workflow engine handles orchestrating tasks for which implementation isn't coupled to the workflow engine in any way. Tasks are implemented as discrete deployable units that are likewise free of dependency on each other.

The implementation of a common task only needs to be completed once, and it joins a library of functionality from which engineers can draw as they build additional workflows. These workflows are expressed in a domain specific language (DSL) wherein the interactions between the tasks are described along with instructions about conditional execution of tasks, how to handle error conditions, and when to retry task executions. These workflow definitions are consulted by the workflow engine and used to determine which tasks need to be executed and in what order. This has allowed us to quickly support new device classes or types as they are introduced, and to build workflows that handle some ancillary concerns around gateway device configuration that often accompany onboarding, such as

The introduction of COAM onboarding in the xFi mobile application gave us the opportunity to implement the architecture we'd been formulating, built upon an external workflow engine and independently deployable FaaS components.

the conditional enablement of gateway agents according to business and product requirements. This paradigm also allows for greater visibility into workflow performance so that potential issues can be identified and understood quickly and addressed with minimal disruption to our customers. Additionally, it provides deeper insight into how each task of the workflow is performing, both in the context of a single workflow and in aggregate across all the workflows in which it is used.

A high-level view of our new architecture looks like the image below, where service is used generically to refer to independently deployable functional units. In our specific implementation we rely on FaaS:

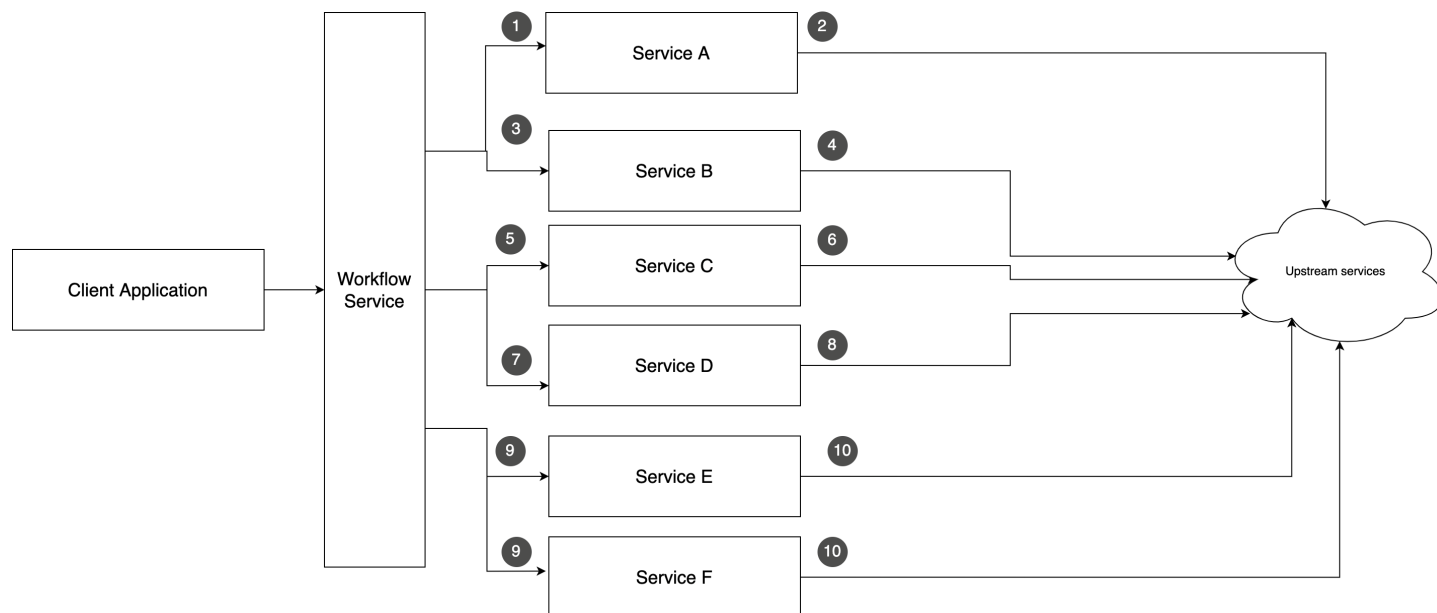


Figure 2 – Workflow Orchestration Architecture

4.1. Architectural Constructs

There are a few primary architectural constructs employed by our workflow orchestration architecture which are described here. From these basic building blocks, we can compose and execute new workflows, and reap other benefits, all of which we will examine in detail.

4.1.1. Task

A task is a discrete functional unit of work. In the onboarding domain this most often equates to a Hypertext Transfer Protocol (HTTP) interaction with an external service used to fulfill some specific, often repeated function. Examples include interactions to provision a gateway device in a centralized device repository, or to associate a gateway device with a set of configurations in a cloud database. Tasks should be generic enough to be reused by multiple workflows. If more than one workflow requires the same piece of functionality, they will ideally use the same task implementation to accomplish this work. The difficulty here lies in making tasks reusable without making them too large or generic. If tasks aren't granular enough, some of the benefits of the workflow orchestration architecture are lost, particularly the observability and state management functions we will look at shortly. In our architecture, task

implementations are typically done as FaaS. Serverless FaaS functions are cost-efficient as they only consume the computing resources they require. During periods of general inactivity, such as in the very early hours of morning when most people opt for sleep over onboarding of their gateway devices, little to no compute resources will be consumed. We've also done a good amount of work to standardize the request and response payloads each task uses so that communication between tasks and with the workflow engine is facilitated.

4.1.2. Workflow Definition

Workflows are defined in a DSL wherein the set of tasks comprising the workflow, the order of their execution, number of retries and retry semantics for each service, the inputs and output fields for the entire workflow, and the input and outputs required by each task are specified. The workflow definition may specify for each task, a fixed number of retries repeated at a fixed interval, an exponential back-off strategy in which each subsequent retry is delayed by an order of magnitude more time than the last, or even that no retries are warranted. The workflow can also specify that tasks be executed concurrently or serially and provide conditions that must be met before a task is executed. Exit criteria for a workflow may also be found in the definition. Certain task failures should result in termination of the workflow, while in other cases workflows may be able to continue after failure of a task that is optional or not essential to the overall workflow success. Our architecture uses a serverless cloud-based workflow engine as described in the next section. Each workflow definition can be managed via the cloud console or defined declaratively as JavaScript Object Notation (JSON) files and managed independently of the cloud console. This allows for automated deployment of workflow modifications. They can also be visualized using the tools provided in the cloud console.

4.1.3. Workflow Engine

Our architecture needs an engine to drive workflows. This engine reads workflow definitions to receive its marching orders and then executes the instructions defined in the workflow definition. It is the engine that orchestrates the task executions and applies the retry, concurrency, and conditional rules specified in the workflow definition. The workflow engine manage failures, retries, and parallelization as described in

the workflow definition so developers needn't be concerned with these ancillary functionalities and can focus instead on where they can produce the most value, namely in implementing the business functionality required by the onboarding process. All this functionality was provided by the initial leased gateway onboarding, but now it is handled as a separate component that can orchestrate many workflows and focus on its primary functionality without being bogged down with the details of task implementations or the business logic embedded in the tasks. Since the tasks use standardized request and response payloads the workflow engine simply feeds the output from one task to the next task in the workflow.

4.1.4. Workflow Integrator

The workflow integrator is a component that provides a means for client interaction. It performs several key functions that fall outside the purview of the workflow engine, the tasks, or the workflow definitions. Chief among these is authorizing clients wishing to initiate workflows, mapping client requests to workflow definitions, starting workflows via the workflow engine, and reporting status on currently

executing workflows to clients who request it. The workflow integrator maintains a mapping of APIs to workflow definitions. This mapping will also include data about what inputs are required for each workflow. The client application makes a call to an API exposed by the workflow integrator which consults its mapping, authorizes the client, verifies proper workflow inputs have been supplied, and starts the Step Function State Machine that correlates to the client request.

Our specific implementation of the workflow orchestration architecture looks something like this:

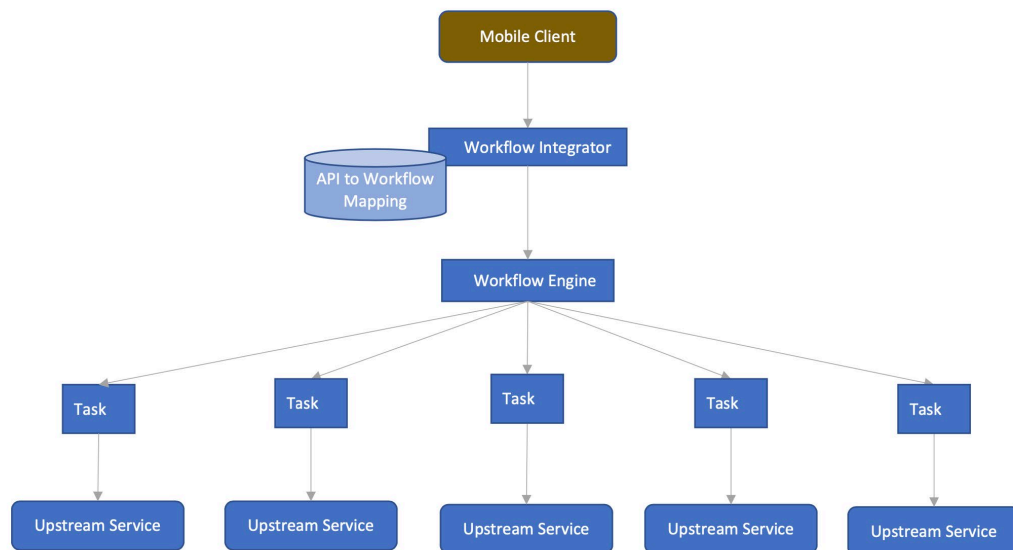


Figure 3 - Workflow Orchestration Architecture Implementation

4.2. Benefits

The workflow orchestration architecture has delivered handily on its promise. It is currently in use for both COAM and MDU onboarding workflows. Here we examine this benefits in detail.

4.2.1. Metrics and observability

With the modular units in our new architecture, logs are emitted to their own buckets in our log aggregator, making it easy to aggregate logs for a given task, or get details of how a given task is performing. Additionally, the workflow engine itself produces a wealth of data about the workflows it has executed. These metrics have been collected and exposed via dashboards in various observability tools, supplying great insight into workflow executions, and allowing us to find and address issues promptly.

4.2.2. Reusability

Separating orchestration from business logic implementation in the services has helped us reuse the same tasks for different onboarding experiences. For example, most onboarding experiences would involve associating cloud-based configuration with a gateway device. Since each task is concerned only with its specific business logic and does not have other dependent services, we can use the same configuration-to-device-association task in multiple workflows without the need for duplicated efforts.

4.2.3. Development agility

Given that we can reuse tasks very efficiently, writing new workflows, in the best-case scenario, has been reduced to writing a new definition file with references to the existing tasks that have already been developed and deployed. This makes it easy to update an existing workflow. As an example, if an existing workflow needs to change the order of the tasks it executes, the change is limited to changing a workflow definition file and deploying it to production with no code change involved. Conversely, shared tasks can be updated to implement some universal change, like a new endpoint or authentication mechanism for an upstream service or a tweak in the business logic without having to modify the workflow definition.

4.2.4. Ease of programmatic client integration

The clients who call our platform API which today include several back-office processes, and web interfaces in addition to the original mobile application client, have an easy intuitive point of integration. All workflows are initiated through API endpoints exposed by the workflow integrator and specified using well understood open-sourced standards like OpenAPI. We are able to generate the client code needed to interact with the platform API thereby eliminating a significant chunk of work the client development organization would otherwise need to undertake. Additionally, since clients do not interact with the workflow engine directly, workflows can evolve independently of and transparently to the clients so long as the responses provided or inputs required don't change. This allows for the tasks to add more features and functionality without impacting clients. If the contract between the workflow integrator and client application stays intact, there is no change needed on the client application and hence no updated version of the app to be released. Mobile application evolution is complicated by the fact that customers we wish to support may still be using older versions of the application. Being able to drive down new features to customers without a client application update helps us provide a more frictionless experience.

4.2.5. State management

With the new workflow-based architecture, the workflow service tracks the state of the workflow as it orchestrates the calls between multiple tasks. This allows clients to resume from the last successful task execution in an earlier attempt; customers need not start the entire flow from the beginning and repeat work they already completed. With the modular breakup in the new architecture, customers can resume from the place they had left off in their previous attempt and not repeat the steps that they had already done.

5. Case Study: Mesh and Advanced Security Firmware Agent Enablement

After a comprehensive discussion of the evolution of onboarding processes and the benefits of the workflow orchestration architecture, we look at a specific example of how the move from one paradigm to the other improved an important business process. The onboarding process involves the enablement of firmware agents to support certain valuable features of our wi-fi product, specifically mesh networking and advanced security. In our initial leased gateway onboarding, if conditions were met indicating the need for agent enablement, this would be tried as an optional part of the onboarding orchestration. It was optional in the sense that should these task executions fail, these failures weren't reported as such to the client, but rather as warning that these portions of the workflow had not completed successfully. Because of the single orchestration, it was not possible to apply different retry semantics to this conditional agent enablement or to allow customers to resume the workflow at these optional tasks so that any failures were left to be dealt with by other external systems outside the context of onboarding. This was expedient in that it allowed customers to accomplish their primary goal of getting access to their HSD service and allowed any trouble in the ancillary configurations to be dealt with independently without requiring action on the part of the customer or delaying their use of the HSD service. Figure 4 illustrates this process.

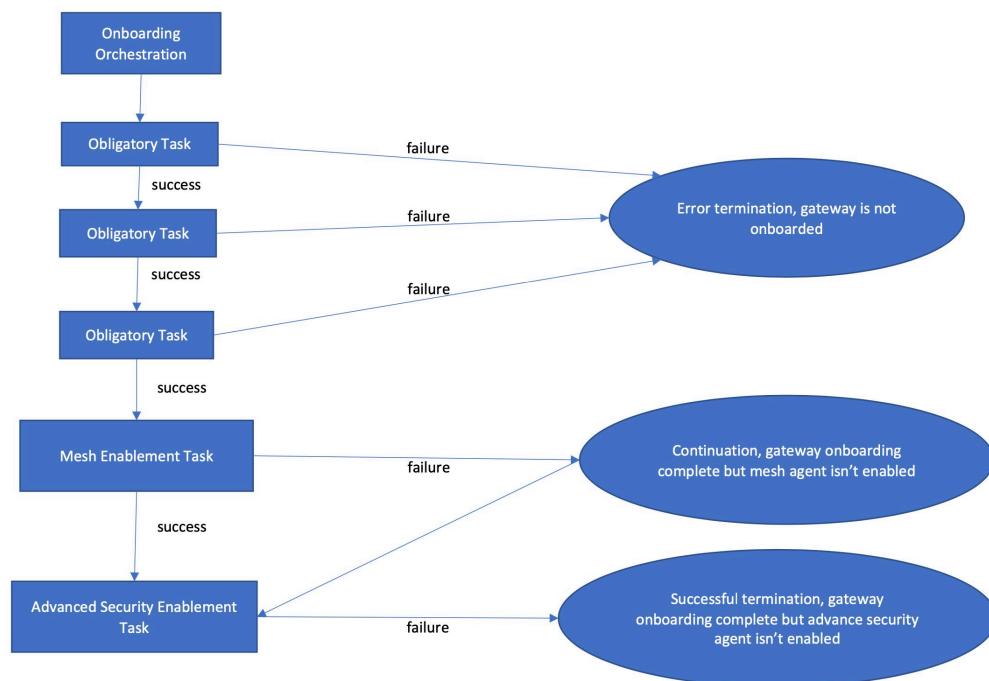


Figure 4 – Conditional Agent Enablement in the Monolithic Orchestration

The workflow orchestration architecture allows us to recognize the conditional enablement of agents as independent workflows, with their own retry semantics and definitions of success or failure. This still accomplishes the primary objective of avoiding a dependency on ancillary configuration before customers

can use their HSD service. The workflow orchestration architecture however, allows us to intelligently handle failures with the ancillary configurations and address them without reliance on external services. When the primary onboarding workflow has successfully completed, it produces an event to a message

Using the new workflow orchestration architecture for mesh and advanced security firmware agent enablement has resulted in better than 99.8% success for each of these processes.

bus which is then used to trigger the subsequent workflows to perform the agent enablement; the primary workflow is completely decoupled from the subsequent workflows that perform the agent enablement. Each can report success or failure on their own terms, and each can employ their own retry semantics. We can define appropriate retry semantics and ensure the enablement completes successfully. Further decoupling is achieved by inserting a message bus between the primary workflow and the agent enablement workflows. The primary workflow does not initiate the subsequent workflows directly. Each workflow has an initializer component that listens for events on the bus and initiates the workflow in response. This allows for externalized retries in addition to the ones defined for the workflow itself. This is particularly useful in a world where gateway devices may be activated prior to shipment to customers,

but agent enablement requires the gateway device to be present on the network. The independence of these agent enablement workflows has also allowed us to introduce incremental improvements to them while leaving the primary workflow untouched. Figure 5 illustrates these improvements facilitated by the workflow orchestration architecture. Using the new workflow orchestration architecture for mesh and advanced security firmware agent enablement has resulted in better than 99.8% success for each of these processes.

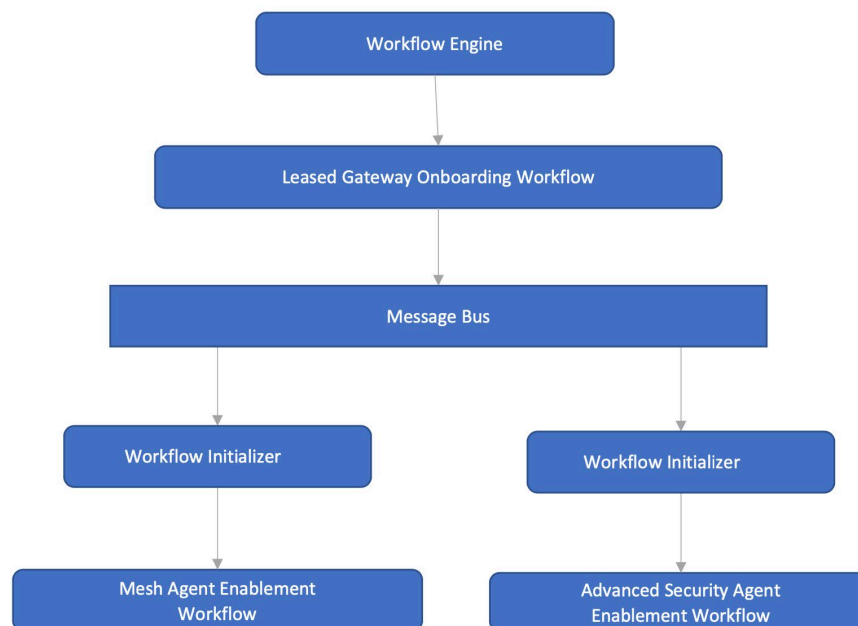


Figure 5 – Agent Enablement with the Workflow Orchestration Architecture

6. Unified Client Interface

As the onboarding process evolved, client interfaces needed to evolve along with them and means of interaction with the platform in presenting the onboarding experience to customers varied. Because of the nature of business process and software evolution, the differing onboarding processes discussed in this paper have and will continue to live concurrently for some time. This diversity in client interfaces is exacerbated by the independent evolution of business and product requirements for different channels, such as mobile clients used by customers versus the web interface used by technicians, or the tools employed by care agents to aid customers with onboarding trouble. Two disparate activation and onboarding platforms have consequently emerged, with some level of interdependency. As new products and devices are introduced, each channel needs to be modified to satisfy the latest requirements. While this has allowed each channel to deliver the appropriate experience, the workflow orchestration architecture gives us the chance to improve this situation. A single client interface for all activation and onboarding needs across all products and devices would be preferable. Having different platforms also increases the potential for customer experience inconsistencies and variance in how the different channels achieve the onboarding process. Figure 6 depicts the crisscrossing interactions that result from the current path.

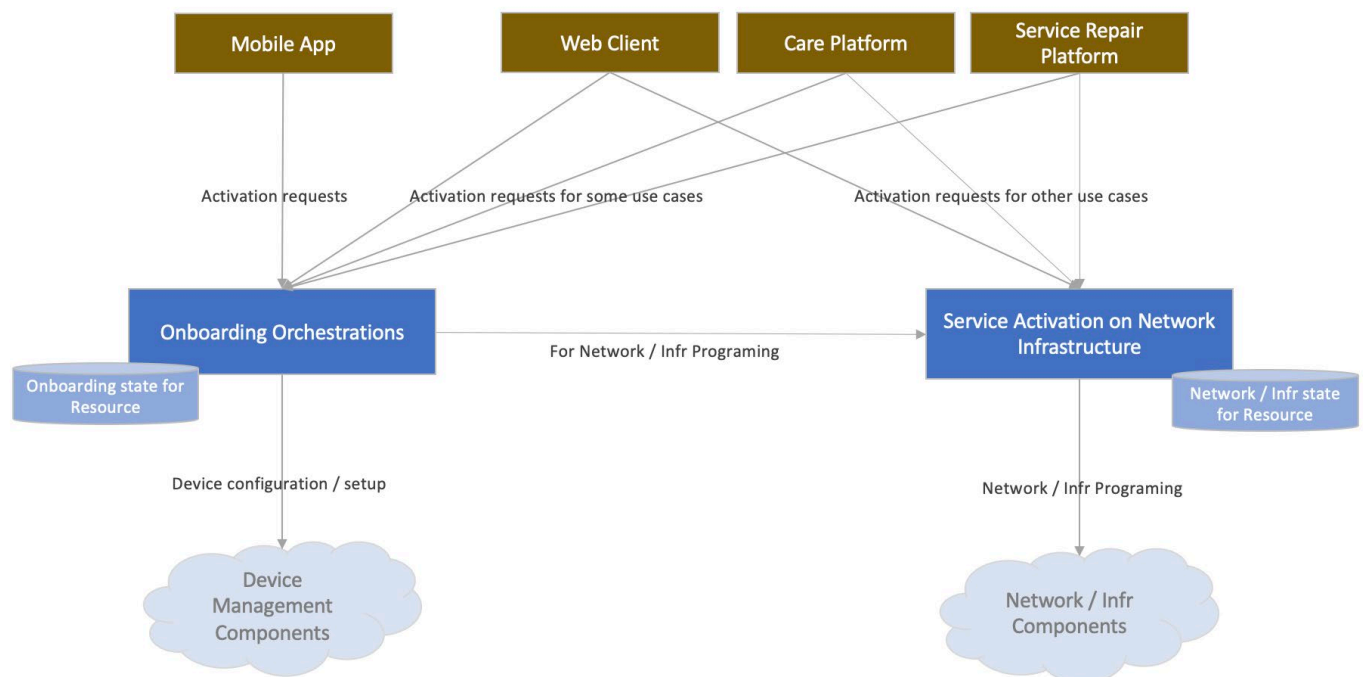


Figure 6 – Channel, Product Specific Client Onboarding Interfaces

To realize the opportunity that now presents itself, the teams responsible for these two platforms have carefully crafted a plan to launch one onboarding hub for all products and devices that can be used by all channels. While offering unified client interfaces, the hub will allow for channel-specific onboarding

considerations while relying on single components for common functionality. Figure 7 illustrates how the divergent platforms coalesce to provide the desired common client interface.

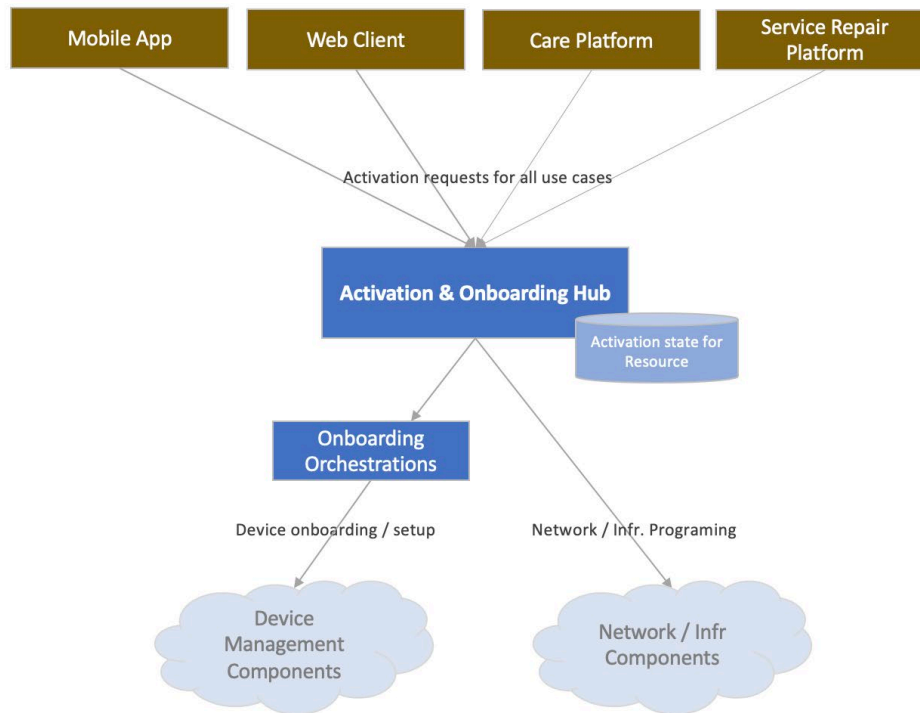


Figure 7 – One Onboarding Hub for all Channels and Products

7. Conclusion

The workflow orchestration architecture we have introduced for onboarding processes has paid great dividends. We have gained great insight into the functioning of workflows and the tasks of which they are comprised. We have a library of discrete functional units that are used to compose new workflows as required. We can quickly and easily modify these functional units apart from the workflow definition itself and vice versa. Workflow specifications can be read and understood apart from the code. We can implement different functional units in whatever programming language is most appropriate for the specific functionality they provide. All of this is helping us to deliver the best possible FTUE to our customers. We have seen steep improvements in onboarding success for COAM customers who can now leverage the Xfiniy App for onboarding. The process is continually evolving but we have the proper tools at our disposal to ensure the onramp to the network shines as brilliantly as possible.

Abbreviations

API	application programming interface
COAM	customer owned and managed
DSL	domain specific language
FTUE	first-time user experience
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
IVR	interactive voice response
JSON	JavaScript Object Notation
MDU	multi dwelling unit
Wi-Fi	wireless-fidelity

Agile High Technology Course Development in the Enterprise

A Technical Paper prepared for SCTE by

Oliver Jojic

Distinguished Researcher
Comcast
1701 JFK Blvd, Philadelphia, PA
703 677 6067
Oliver_Jojic@comcast.com

Mike Baker

AI Program Director
Comcast
1701 JFK Blvd, Philadelphia, PA
267 207 5865
Mike_Baker@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Dialectic of Teaching Alternatives	3
2.1. Legacy Teaching Structures	3
2.1. Attempt at Formalizing the Alternatives	4
2.2. Sync / Async dimension	5
2.3. Short lessons / Whole course dimension	5
2.4. In-house / External dimension.....	6
3. Student's Practice	7
3.1. Capstone projects	7
3.1. Mentors	8
4. The Agile Teaching Process	9
4.1. Administrator journey	9
4.2. Leadership journey	9
4.3. Student journey	9
4.4. Teacher journey	10
4.5. Mentor journey	11
5. Pilot Results	11
5.1. Ability of in-house domain experts to create teaching material.	11
5.2. Ability of in-house domain experts to teach	11
5.3. Ability of the agile teaching process to scale to more students	12
5.4. Retention	12
6. Challenges	12
7. Business Impact	13
8. Conclusion.....	13
Abbreviations	14

List of Figures

Title	Page Number
Figure 1 – Dialectic of Teaching Alternatives.....	4
Figure 2 – A Capstone Project and its Soft Prerequisites.....	7
Figure 3 - Example Interaction between a Student and a Mentor	8

List of Tables

Title	Page Number
Table 1 – The two Legacy Programs	4
Table 2 – Size of the Catalog of short Lessons	11
Table 3 – Student NPS	12
Table 4 – Students per Year	12

1. Introduction

This paper describes an effective agile process for teaching *technical* subjects in the workplace. For example, data security, machine learning, cloud computing are technical subjects.

With the advent of Big Data and with today's breakneck speed of technical innovation it is more important than ever to provide the technical workforce with continuous education.

But many questions need answers: how frequently should a full-time employee be diverted from work due to their education? What budget should be assigned to the continuous education of technical employees? How can the teaching material be always kept up to date and relevant to the company?

As we were tasked to teach machine learning to a large portion of the software engineering workforce, we had to ask ourselves these questions and more, and after a few iterations we reached our current process which seems to “just work”: no extra budget and an average student's NPS (net promoter score) nearing 100%. The process uses in-house technical experts to design and teach short lessons, and it uses past graduates to act as mentors for new students. In addition, the process has the originally unplanned benefit of favoring networking among employees from different parts of the company.

Encouraged by our successes we recently applied the same teaching process to new technical domains: Data Science and Full Stack DevOps, again, it seems to just work. No extra budget and an average student's NPS in the 90s.

In this paper we answer the above questions and more, we detail our teaching process, and we share some quantified results.

2. Dialectic of Teaching Alternatives

As we embarked on our task of teaching machine learning to a large fraction of the software engineering workforce, we had to evaluate the then current legacy teaching structures. This evaluation led us to attempt a somewhat formal description of teaching alternatives, with the purpose of identifying what would work best for our task.

2.1. Legacy Teaching Structures

Two very distinct teaching programs existed when we started.

The first was akin to a university course but only for the selected, high-performing employees. Over the span of nine months, a carefully selected few employees are gathered in a classroom with an external adjunct teacher with a curriculum that is optimally designed to match Comcast needs. The rather trivial problem with this program is that it is not scalable, not only on the “horizontal” view of fulfilling the education of many employees, but also on the “vertical” view of the continuous education of one employee. Nevertheless, this is a successful and coveted program, and it is being continued today.

The other was open to all employees: it is a selection of online teaching materials. The problem with this totally async approach is that students that can learn on their own don't necessarily need this pre-selection, while students that would need some help do not get it.

Table 1 – The two Legacy Programs

	“Selected Employees” program	“Open to All” program
Instructor led	Yes	No
Duration	9 months	Async
Budget	Expensive	Cheap
Scalability	Low	High
Effectiveness	High	Low
Resilience to changing needs *	Ok	Chaotic

** The AI domain is rapidly evolving. How quickly can a program be updated to include the teaching of new concepts or tools*

The stark contrast between these two legacy programs incited us to attempt to formally describe teaching alternatives.

2.1. Attempt at Formalizing the Alternatives

This section attempts to be somewhat formal in listing teaching possibilities. Having a clear view of the pros and cons of all alternatives is helpful in designing the right teaching process.

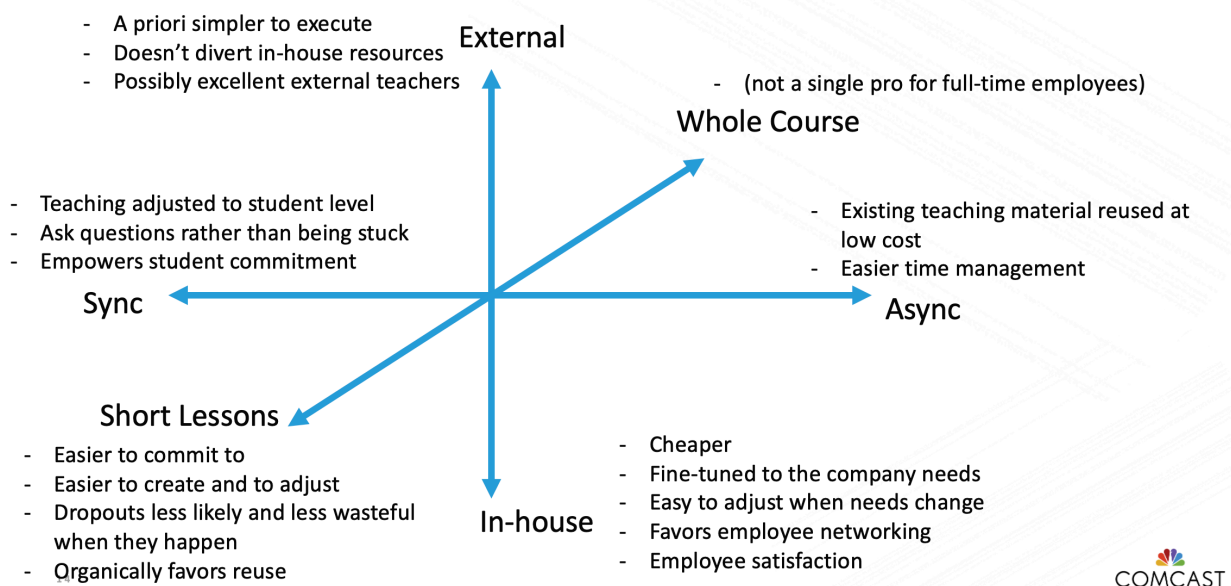


Figure 1 – Dialectic of Teaching Alternatives

2.2. Sync / Async dimension

Technical subjects are complex, and students will benefit from all the help they can get.

Sync means teacher and students are interacting in real time. Async means the student is consuming a previously prepared material. The async version may include some domain expert helping the student, for example by answering questions on a messaging app. We will call such an async helper a “soft mentor”, or a mentor for short.

Note: in our use of sync/async *the interaction* between the teacher and the students is key. If the information flows in only one direction (teacher to a large class that is not expected to interrupt) then it practically matches async teaching.

Sync: the teacher can adjust what is being taught to the level of the students.

Sync: students do not get stuck because they can ask the teacher for clarification.

Sync: empowers student commitment.

Note: student commitment is key, otherwise continuous learning risks of being a farce, a mean for the student to check some educational boxes. The teacher empowers student commitment by interacting with every student in the class, by asking short, simple questions in some round-robin fashion.

Async: time-management is easier for full-time employees.

Async: existing teaching material can be reused without extra costs.

Conclusion: both sync and async have merits. Possibly both can be used for the successful continuous education of full-time employees: a short sync class followed with a mentored async practice.

2.3. Short lessons / Whole course dimension

Throughout our younger lives we have all been educated with *whole courses*: long duration, well designed material that teaches enough facets about a new subject for the learner to get a good logical understanding of the whole. Yet, long duration courses are clearly not ideal (even incompatible) for full-time employees. So, we explored the alternative: *short duration lessons*.

An example of a whole course would be “Machine Learning”. An example of a short lesson would be “Random Forests”.

What is short? Our experiments showed that up to 8 hours spread over one week is favored by full-time employees.

Short lesson: easier to commit to for full-time employees.

Short lesson: easier to create and to adjust.

Note: we mentioned earlier that the AI field changes at breakneck speed. For example, it is easier to create a new lesson about Attention and Transformers rather than adjusting a well-designed coherent whole course with the new content.

Short lesson: student dropouts are less likely and are less wasteful when they happen.

Short lesson: organically favors reuse.

Note: a whole course about machine learning and a whole course about data science might both have a section about Python. The content about Python could be almost identical, but there is no easy way to reuse here. Instead, with a collection of short lessons, the Python lesson is already there and available: organically favoring reuse.

Short lesson: scalable.

Note: spawning the same short lesson multiple times throughout the year to reach more students is doable because it is short. Spawning a whole course multiple times is hard and limited because it is long duration.

Conclusion: For full-time employees we did not find a single *pro* for the *whole course* alternative. In the enterprise, short lessons are always to be favored instead of whole courses. The disconnected nature of a long list of short lessons can be resolved by having the company's educational website somehow link the available short lessons into coherent wholes.

2.4. In-house / External dimension

The course material and its teaching can be developed in-house, or it can be selected from external sources.

Note: The analysis and conclusions about this In-house/External dimension depend heavily on the prior selection between *short lesson* and *whole course* (described in the previous section). Here we assume that the teaching solely consists of short lessons. A few weeks per year an in-house domain expert will have to spend up to 8 hours teaching. This is done without much impact on the employee's regular activities. The design of the short lesson is somewhat more involved, but it happens only once. Sometimes the domain expert selects an existing online source material, in which case the design of the short lesson is even less time consuming.

External: a priori simpler to execute.

In-house: cheaper.

External: doesn't need to divert in-house resources.

In-house: fine-tuned to the company needs.

In-house: easy to adjust the teaching material when the needs change.

In-house: favors networking among employees from different parts of the company.

External: external adjunct teacher can be excellent at teaching.

In-house: sense of satisfaction for the domain expert employee thanks to being selected to teach to the whole company.

Note: after running our agile teaching process for a while, we noticed this sense of pride forwards to the employee's leadership as well.

Conclusion: while there are benefits to both, a company should favor the in-house alternative. For this to be feasible the teaching material should be restricted to short lessons.

3. Student's Practice

To be a practitioner of a new technical skill, a student learns the relevant technical material, but also spends significant time practicing the new skill. Anything long duration that is imposed on full-time workers is best implemented in async mode. Therefore, after completion of a short sync lesson, we do encourage the employees to practice their new skills, but the frequency and total duration of the practice is up to the employee.

Whenever possible, a short sync lesson should be accompanied with async practice exercises (sometimes this can be a simple link to already existing online material).

3.1. Capstone projects

A Capstone project is a bigger, more significant mean of exercising a new skill. Each student decides when and how frequently they will work on their Capstone project. There is no need for a time limit.

The completion of a Capstone project should be a formal event so that the employee rightfully feels a sense of accomplishment. This can be done with a ceremony, grouping together multiple recently completed projects, where each graduate describes their project to the group. At the end a diploma is handed to each employee.

Capstone projects also give structure to an otherwise long list of disjointed short lessons.




Figure 2 – A Capstone Project and its Soft Prerequisites

Each oval is a short lesson. The arrows represent soft prerequisites. Soft means that an employee does not have to enroll in a short lesson if they already know the subject. On the company's educational website, a Capstone project is roughly equivalent to a short lesson. The entry should list all the soft prerequisites.

Students can enroll to the Capstone project and their duty is to complete the project and to create a presentation of their work to be shared during the graduation ceremony.

While practicing, whether with a Capstone project or with simple exercises, the student can get help from mentors.

3.1. Mentors

As with any new knowledge, students should practice what they learned if they want to become proficient. Practicing alone can be hard, as the chance of getting stuck is high when dealing with something new. Having mentors to help students practice is extremely useful. The word mentor here should be understood as a soft mentor, or a learning assistant.

This is great in principle, but how is scalability resolved? With continuous learning there is a continual flow of active students. Our solution to the scalability of mentors is to ask select graduates if they are willing to mentor future students. Through time, as the number of students grow, the number of available mentors grow as well, resolving the scalability issue.

A company-wide messaging app (Slack, MS Teams) is used to interact asynchronously between practicing students and mentors. Domain experts (i.e., the teachers) sometimes join a conversation as well.

This process has the originally unplanned benefit of favoring networking among employees from different parts of the company, of building a community that is excited by the same technical domain (machine learning in our case).



i have tried doing pca, the explained variance is so l



Tsung-hsiang Hsueh 🏠 10:55 AM

Yes, but you can see that some attributes have much more variance than others. Are these normalized?



Prasad Menon 10:56 AM

yes i used standard scalar before going with pca



Tsung-hsiang Hsueh 🏠 11:03 AM

Some experiments you should try:

Do you get different results if you use for example t

Figure 3 - Example Interaction between a Student and a Mentor

4. The Agile Teaching Process

The Agile Teaching Process implements the best practices discovered in earlier sections:

- Short lessons rather than long duration courses
- Sync teaching followed with async practice
- In-house domain experts encouraged to teach their expertise
- Past graduates encouraged to mentor new students

The process is best explained by following the journey of all parties involved: administrator, leadership, students, in-house domain experts, and mentors.

4.1. Administrator journey

The administrator typically belongs to the education side of the company. The administrator role is key to ensure that the agile teaching process is successful. As such, the administrator, directly or indirectly, must be in contact with all participants. Administrator responsibilities are highlighted in the subsequent “journey” sections and are not being duplicated here.

4.2. Leadership journey

Leaders must be convinced that it is worthwhile for their expert employees to be sometime diverted from their daily activities to teach their expertise to the rest of the company. This radical meshing of activities across orgs with well-defined boundaries is best started by the administrator introducing the new process as a long-term pilot.

The administrator should provide the leaders with a quarterly update.

4.3. Student journey

Employees are made aware of the available short lessons by browsing the company’s educational website. The administrator ensures the website is kept up to date.

How frequently should a short lesson be instantiated? This depends on its popularity: the administrator monitors the number of students that are registered to a class. To maximize the effectiveness of sync teaching a class should be limited to a maximum of roughly 25 students. The administrator will increase or decrease the frequency of a class based on the number of registered students. But, prior to this adjustment, the administrator should verify that the number of registered students indeed reflects the popularity of the class rather than being due to the students not noticing the class; this is particularly true for new classes. The administrator should maintain a growing mailing list of potential students. When the number of students being registered to a class is low, before reducing its frequency, the administrator will email a reminder about the class. In our experience, this often results in many more students registering to the class, proving that the lack of popularity was due to lack of employee awareness rather than a class being instantiated too frequently. This description implies that the teachers must show some flexibility with their calendar planning; this is an agile teaching process, and the administrator should from the start make everyone aware that some agility is to be expected.

For full-time employees, to be able to benefit from sync teaching, the class duration must be short. From our experience, a maximum duration of 8 hours spread over one week works well. Always favor spreading the learning hours over multiple consecutive days. For example, teaching 2 hours during 4

consecutive days is to be preferred to teaching 8 hours within a day. There is only so much a brain can learn in a day.

Courses are recorded allowing students to play or replay the lessons in async mode.

To learn new skills, practice is key. Many short lessons should be accompanied with exercises for the student to complete asynchronously after the class is over. The student is free to do it at any time. Jupyter notebooks are perfect mediums for holding the exercises (and for holding the course as well). If they have difficulties during the completion of the exercises, a student can ask mentors for help. A company-wide messaging app is used for the interaction between students and mentors. The student asks a question and asynchronously a mentor, or more, answers it.

Students can more thoroughly practice their newly learned skills by completing Capstone projects. Capstone projects have their own entries in the company's educational website. A student must register to a Capstone project. There are no teachers for Capstone projects. Instead, the Capstone project's entry in the educational website describes the required knowledge by listing the relevant short lessons. These lessons are only "soft" prerequisites, e.g., if an employee already knows Python, they do not have to register to the Python prerequisite class. There is no time limit for the completion of the Capstone project, it is up to the student. Students working on a Capstone project will frequently communicate with mentors through the messaging app. In addition to completing the project, a student must also prepare a presentation describing how this was done. Once enough Capstone projects are completed, the administrator organizes a graduation ceremony where each student describes their project to the group and at the end receives a diploma.

4.4. Teacher journey

In this agile teaching process, all teachers are in-house domain experts. It is the domain expert, in coordination with their leaders, that decides what new lesson should be developed. They know what knowledge is needed for the company technical employees. They are the first to notice the advancements in the state-of-the-art or the emergence of new, better, tools. The administrator is responsible for reminding domain experts and their leaders of the company's technical educational needs and of their duty within this agile teaching process. Our experience shows that most domain experts are excited about the opportunity of teaching their art across the company and their leaders are honored by their group being acknowledged as an important source of knowledge within the company. This motivation can be boosted further by explicitly listing the teaching of one's art in the yearly goals of all technical employees above a certain level. We do believe that this process brings a sense of purpose, builds a community, and ultimately strengthens retention of employees.

Once a new lesson has been agreed upon, the administrator helps the expert commit to the design of the lesson and of the accompanying exercises by setting up a schedule to which the expert can comfortably adhere to. The total duration of this schedule varies but it takes around 3 months on average from the original idea down to its completion.

During the course design a specialist helps the expert with good education principles. Two of the most important educational principles are 1) the teacher is not in the class to bombast his knowledge but rather to adjust it to the current level of the students and 2) the teacher maintains students' attention by continuously forcing them to participate with short simple questions in some round-robin fashion.

In our experience failed teachers are rare, as proven by an average NPS from students that is above 90%. When it happens, the administrator gently discards the failed expert teacher and finds a replacement.

4.5. Mentor journey

When practicing new skills, it is crucial to have some assistance to help avoiding “being stuck”. Our use of the term *mentor* in this paper means exactly this: a practice assistant.

The agile teaching process ensures that every student has helpful mentors during practice. How is that scalable?

After graduating a student may be contacted by the administrator about the possibility of becoming a mentor to future students. It is the administrator’s responsibility to notice which students have superior skills, as poorly done mentorship could be counterproductive. The administrator will consult with teachers to confirm the appropriateness of a student to mentorship.

Once selected, and willing, a past graduate is promoted to mentor. The duty of a mentor is simple: monitor the company-wide messaging app for student questions and answer the question and be generally helpful. Thus, the total number of mentors keep growing, resolving the scaling issue.

While not necessarily intuitive, this mentorship by past graduate, in our experience, is working very well. A few star mentors are quickly answering most of the questions (even though the interaction is not expected to be real time) and the students are very happy with the help they are receiving.

In our experience, most contacted graduates were excited to become mentors.

5. Pilot Results

The last 4 years were exciting as we kept fine-tuning the methodology for teaching machine learning at Comcast. In this section we share some numerical results.

5.1. Ability of in-house domain experts to create teaching material.

The pilot switched to the current version (in-house experts preparing short lessons and teaching them) just 2 years ago.

Table 2 – Size of the Catalog of short Lessons

Number of Short Lessons in the Catalog (as of this writing)	Average time to add a new Lesson
25	< 3 months

5.2. Ability of in-house domain experts to teach

The table below shows the average students’ net promoter score in the last 3 months. The score is quite stable over time, between 90 and 94%.

Table 3 – Student NPS

	This Agile Process	Legacy “Democratic”	Legacy “Best and Brightest”	Education global
Student NPS	91%	Not tracked. Mainly constructive feedback	85%	55%

5.3. Ability of the agile teaching process to scale to more students

Yearly course completion or students per year:

Table 4 – Students per Year

	This Agile Process	Legacy “Democratic”	Legacy “Best and Brightest”
Students per year	500+	100+	30

5.4. Retention

Anecdotally the fraction of mentors and teachers that leave the company is less than the average of all employees.

6. Challenges

For convenience we gathered here the challenges to this agile teaching process.

Administering an in-house teaching process is demanding. The tasks include: consulting with business leaders and experts about courses that should be developed. Finding an expert that is willing to design and teach the course. Tracking the progress of the course design. Ensuring that employees are aware of the list of available courses. Sharing results with business leaders.

In house expert resources are diverted from their regular duties. This is facilitated by leadership agreeing that knowledge sharing is an official part of an expert’s duty.

Assumption that domain experts are apt at teaching. This was a major unknown at the beginning of our pilot, but the results are positive. At least in technical fields, most experts can teach efficiently (based on the student NPS scores for experts vs. external faculty classes). Still, we had to cancel about 10% of experts due to their inability to teach.

A minority of employees prefer async training. An easy remedy is to record the classes and let each employee decide whether they prefer to join a sync class or to listen to the async material.

What happens to their class when an expert leaves the company? It happens, and the administrator must find a replacement expert to teach the same class. The replacement expert is made aware that the course material is now theirs and that it is within their rights to adjust the material to their teaching style.

7. Business Impact

For convenience we gathered here the impact to the business of this agile teaching process.

Mastering data driven decision making. Some parts of our big organization have business leaders that push for more data-driven decision making. The progress toward that goal was slow and the reason might have been due to the engineering workforce, on average, having a lacking knowledge in AI/ML. After three years of applying this agile teaching process we have boosted the knowledge of machine learning to more than 500 employees. Many of them are now genuinely excited by this technology, and now push for data-driven decision making from the bottom-up as well. Compared to a few years ago, our company is noticeably more agile in extracting knowledge from data. It is hard to quantify how much of it is due to our application of this agile teaching process.

Low dollar cost. Not having to pay for external faculty is a cost saving. Even async teaching material can be expensive. With this agile teaching process, the dollar cost is replaced with the in-house experts being diverted from performing their regular duties. This “cost” is welcome if leadership agrees that knowledge sharing is an official part of an expert’s duty.

Scalability with no cost increase. If it is in the business interest to spread some knowledge quickly, or simply if a class is very successful, the expert will teach more frequently, say 8 hours per month instead of 8 hours twice per year.

Expert employee retention. While anecdotal (this one pilot cannot be considered statistically significant) it seems that designing and teaching a course has beneficial effects for expert employee retention. If true, it is likely due to the combination of two psychological effects: the pride of being selected to design and teach complex material to the rest of the company; as well as the significant increase in connections, sometime even new friendships, between a teacher and their many students.

Significant gains in NPS performance. Based on the students’ NPS ratings, their least favorite method of learning is asynchronous, where the company selects some async training material (sometime quite expensive) and the student is left on their own. Next up, with a significant NPS jump, is external faculty teaching. Lastly, the agile teaching process described in this document with its in-house expert-based teaching, has the highest NPS scores.

8. Conclusion

In this paper we described an effective agile process for teaching technical subjects in the workplace.

This process is scalable, it closely tracks the educational needs of the business, it empowers the continuous education of technical employees, it builds a community with the same passion, and it does not require any additional budget.

How? By leveraging the expertise that already exists within the company.

In-house domain experts design and teach short lessons. Because the teaching material is being built by the company’s employees, it is always fine-tuned to the company needs. Because each lesson is short, it can easily be scaled up to more iterations for classes in high demand.

Students are grateful because they have a sync teacher and can ask questions but also because the course material is cut in small lessons and easier to fit within their busy schedule. But learning is only half of the path toward proficiency. The other half is practice, and practicing new skills is hard. When to practice is completely up to the student. Yet, the student can still get async help from mentors through a company-wide messaging app. Select graduate students are asked in turn to become mentors, making the whole process scalable. Student's satisfaction is clear given that their average NPS is above 90% promoters.

Lastly, and this benefit was not originally planned, we found that this process forces the interaction between experts, students, and mentors and organically builds a community with the same passion and ultimately is likely to strengthen employee retention.

Abbreviations

SME	Subject matter expert. For clarity, in this paper, we use “in-house domain expert” instead
Mentor	Learning assistant
Sync	Teacher and students are interacting in real time
Async	Student is consuming a previously prepared material
NPS	<p>Net Promoter Score:</p> <p>Students rate a class from 1 to 10.</p> <p>Total = number of student ratings</p> <p>Promoters = number of ratings equal to 9 or 10</p> <p>Detractors = number of ratings equal to 6 or below</p> <p>$NPS = (Promoters - Detractors) / Total$</p>

AI for IT Operations (AIOps)

Using AI/ML for Improving IT Operations

A Technical Paper prepared for SCTE by

Hongcheng Wang

Distinguished Engineer, Machine Learning
Applied AI & Discovery, Comcast
1325 G Street NW, Washington, DC 20005
332-301-5055
Hongcheng_Wang@comcast.com

Praveen Manoharan, Applied AI & Discovery, Comcast

Nilesh Nayan, Applied AI & Discovery, Comcast

Aravindakumar Venugopalan, Applied AI & Discovery, Comcast

Abhijeet Mulye, Applied AI & Discovery, Comcast

Tianwen Chen, Applied AI & Discovery, Comcast

Mateja Putic, Applied AI & Discovery, Comcast

Table of Contents

Title	Page Number
1. Introduction.....	4
1.1. What is AIOps?	5
1.2. The Difference from MLOps	5
1.3. Why AIOps?	5
2. AIOps Platform Architecture.....	6
2.1. Real-Time Data Processing	7
2.2. Scalable Architecture	7
2.3. Data Storage and Retrieval.....	7
2.4. Real-Time model training and model packaging.....	7
3. AIOps Use Cases.....	8
3.1. Proactive Monitoring	8
3.2. Smart Alerting.....	8
3.3. Topology Analytics and Root Cause Analysis (RCA)	8
3.4. Log and Trace Analytics.....	9
3.5. Cohort Analysis	9
3.6. Automated Remediation.....	9
3.7. Smart Orchestration	9
3.8. Release Management	9
4. Anomaly Detection	10
4.1. Anomaly Detection Algorithms	10
4.2. Anomaly Detection Platform.....	11
5. AIOps – Operators’ Perspective.....	12
5.1. Onboarding a tenant	13
5.2. Training a metric with model	13
5.3. Schedule a metric for real-time inference and prediction	13
5.4. Configure Alerts and RCA.....	13
6. AIOps Case Study: Connected Living Object Detection Operation	13
6.1. Connected Living Object Detection Operation.....	14
6.2. Log mining and Correlation Analysis	16
7. AIOps Case Study: Rex Browse and Search Operation.....	19
7.1. Dependency Graph	19
7.2. Root Cause Analysis with Dependency Graph.....	20
7.3. Failure Prevention and Auto Remediation	22
8. AIOps Case Study: Release Management for RDK Firmware	24
8.1. Machine Learning Model.....	24
8.2. Model Results.....	24
9. Conclusions.....	25
10. Acknowledgements	25
Abbreviations	26
Bibliography & References.....	27

List of Figures

Title	Page Number
Figure 1 – AIOps High Level Architecture Diagram.....	6
Figure 2 – Anomaly Detection Platform	12
Figure 3 – Metric in Object Detection Operation.....	14
Figure 4 – Anomaly Detection Predictions in Object Detection Operation	15
Figure 5 – Manual Log RCA in Object Detection Operation	17
Figure 6 – Automatic Log RCA in Object Detection Operation.....	18
Figure 7 – Manual Flow of Issue Root Cause Diagnosis by Rex Operations Team.....	19
Figure 8 – Dependency Graph Created in Rex Case Study.....	20
Figure 9 – AIOps Showing Automatic RCA Correlated Ranked List	21
Figure 10 – AIOps Showing Probable Root Cause Metric Time-series.....	22
Figure 11 – Percentage of Pods Having GC activity per 5-min Period.....	23
Figure 12 – Instances Showing Abnormal GC Count Over Time	23
Figure 13 – Precision Recall Curve for the model (Left Figure); 5G WiFi Signal Distribution for Old Firmware Version and New Version (Right Figure)	25

1. Introduction

The proliferation of microservices as a dominant IT architecture has created opportunities as well as challenges for operations teams responsible for maintaining software reliability. When production systems deviate from service-level objectives, operations teams must detect failures and discover their root causes to promptly resolve the issue. These teams are most often focused on minimizing mean time to resolution (MTTR) or mean time between failures (MTBF). The process of identifying, diagnosing, and resolving issues in cloud microservice architectures largely falls into two phases: anomaly detection (AD) and root cause analysis (RCA).

AD is the process of identifying anomalies that correspond to system failures. RCA is the process of determining the reason why an anomaly occurred and identifying the originating service or system. Anomalies are typically detected by defining margins of normal operation on key performance indicators (KPIs) and setting alerting thresholds that generate notifications. RCA is then typically performed by inspecting the system that generated the alert and tracing the problem back to its source. Operations teams use log, trace, or metric data sources, often displayed in dashboards to diagnose and debug problems.

However, there are several problems with this and related existing approaches:

- (1) ***Simplistic AD still dominates:*** State-of-the-art failure detection is still based on simple thresholding, which is prone to drift, and fails to capture low-frequency events such as weekends, holidays, or special events. Failure detection based on simple thresholding, misses opportunities to preemptively diagnose problems, thereby increasing MTTR.
- (2) ***Manual RCA still dominates:*** Root cause analysis is the most time-consuming step of issue resolution, largely due to a reliance on a human in the loop. When an alert is received, reliability engineers spend significant time identifying the root cause by looking at numerous plots, traces, and logs. This work is repetitive, tedious, and ripe for automation.
- (3) ***Excessive alert volume:*** Operations teams often receive a large volume of alerts, many of which are false or redundant. These are generated by rules that often remain unchanged for the life of the application. Further, the volume of services in a microservice application makes it difficult to know if a service has failed on its own or as part of a cascade.
- (4) ***New deployment challenges:*** When a new product or service is deployed, the operations team keeps a closer eye on the alerts, metrics, and system performance. The decision to move forward to 100% general availability (GA) or roll back to a previous version usually takes unnecessary lead time, which may create negative customer impact.
- (5) ***Institutional knowledge monopolies:*** Often the knowledge needed to quickly debug operational problems is held by a small number of individuals on the team. Root cause analysis can be time-inefficient except for the few individuals who hold a monopoly on that knowledge.

According to the report by Smartsheet [1], nearly 70% of employees say that automation reduces the time wasted on repetitive work, and out of which nearly 60% believe that if repetitive jobs were automated, they would save 6 or more hours (almost a full workday) each week. This brings huge opportunities for AIOps.

1.1. What is AIOps?

AIOps (AI for IT Operations) uses artificial intelligence and machine learning (AI/ML) and big data analytics to identify or predict IT operations issues in a timely manner and help the DevOps team quickly identify the root cause of the issues. A machine learning model can learn the conditions that lead to an alert and can predict when an alert is about to occur. When an alert does happen, ML based RCA is used to generate candidates of sources of failure to be diagnosed by a human operator. Thus, it can reduce MTTR by automating repetitive jobs, present failure, and provide better decision making.

AIOps has recently received extensive attention from industries and academia. According to a survey by Reportlinker – “AIOps Platform Market Forecast to 2028 - COVID-19 Impact and Global Analysis by Component, Deployment, Organization Size, and Vertical” [2], the AIOps platform market size is expected to grow from \$ 2.8 billion in 2021 to \$ 19.9 billion by 2028. It is estimated to grow at a compound annual growth rate (CAGR) of 32.2% from 2021 to 2028. Apart from the market growth, its impact can save much more than that by providing the predictive abilities which will lead to efficient utilization of resources so that companies are able to cut additional costs related to overprovisioning of their cloud and other resources. It will also help in better application maintenance resulting in a better overall customer experience, thus positively contributing to business revenues.

1.2. The Difference from MLOps

The terms “AI” and “ML” are interchangeable in many contexts. In this context, however, the meaning of MLOps and AIOps are significantly different. MLOps refers to machine learning model operations, from data acquisition to model development, testing, validation, and deployment. MLOps seeks to increase automation and improve the quality of production ML models, by focusing on the operation of ML models, and leveraging the continuous integration/development (CI/CD) practice of DevOps in the software field.

The AIOps methodology is applicable to any IT operations, including MLOps. AIOps could make ML model operations more reliable and cost effective. On the other hand, the MLOps pipeline could be leveraged to make the operation of AIOps itself more efficient and reliable.

1.3. Why AIOps?

AIOps aggregates data from multiple sources and provides context and insights when problems occur. It improves the visibility (observability), reliability, availability, and cost of IT operations. In general, AIOps provides the following key business benefits:

- (1) **Improved system availability:** AIOps improves availability by reducing MTTR in several ways. First, predictive AD can intelligently identify issues before they occur, automatically categorizing issue criticality, and preventing unnecessary escalation of issues. Second, automated RCA significantly narrows the scope of the problem on which a human operator must focus, greatly reducing the amount of time needed to reach resolution. Third, capturing institutional knowledge to a model means faster issue resolution, even if less experienced operators are on call.
- (2) **Reduced operational cost:** There are multiple ways AIOps reduces operational cost. First, as AIOps sends out fewer alerts and automates RCA, the resulting reduction in workload could potentially reduce the headcount of operations teams. Second, as ML models can predict the pattern of traffic from historical data, AIOps can help to orchestrate resources more intelligently for cost savings. Third, AIOps helps to quickly identify any potential issues within limited

deployments and provides better insights and decision making before promotion to GA, reducing unnecessary use of valuable human and computing resources.

- (3) **Improved employee experience:** The reduced number of false alarms creates more efficient (noise free) work for reliability engineers while lowering the overall work volume. Online learning models eliminate threshold drift and reduce manual effort. The AIOps system can act as a partner in a pair-debugging strategy, that enhances the capabilities of the human operator. The overall reduction in issue volume results in a happier, more productive operations team by eliminating pager fatigue, allowing them to focus on more meaningful tasks.

Our AIOps team is striving to help address the operational challenges with AIOps. We have explored and experimented on several use cases including:

- (1) **Intelligent Infrastructure Monitoring (IIM):** We built the state-of-the-art anomaly detection technology to alert the DevOps team with detected anomalies based on load and resource utilization to prevent application failure as early as possible.
- (2) **Root Cause Analysis (RCA):** When there is an anomaly or a failure in the operation, we correlate the system and application log data with the detected anomaly and help the team to quickly identify the root cause and recommend the correct actions to the team.
- (3) **Release Management (RM):** We use ML to help manage the release by quickly identifying when a gap occurs with a partial rollout.

In this paper, we will give an overview of AIOps use cases and summarize our findings from several practical case studies from IoT (Internet of Things), content discovery, and RDK (Reference Design Kit) applications.

2. AIOps Platform Architecture

In this section, we describe the high-level AIOps platform architecture we built, as shown in Figure 1.

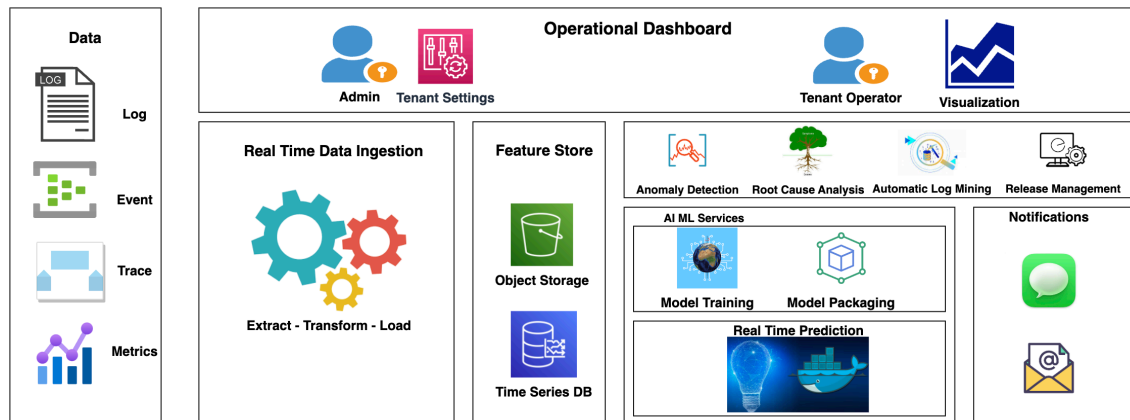


Figure 1 – AIOps High Level Architecture Diagram

AIOps architecture address end to end solution right from data ingestion to data transformation, data storage, model training, real-time prediction, fine tuning of model and notifications.

The data layer is composed of the data from different tools (e.g., metrics and log data sources) and systems, internal or external. This data may include a large scale of logs, events, traces, and

metrics data from applications onboarded onto the AIOps platform. Data transformation layer performs required transformation on the data and load to Feature Store. The data is ingested into the backend, and further engineered (e.g., transformation, aggregation) to store in the feature store (Object Storage or Time Series Database). AI/ML Service is the core of AIOps that enables model training, model tuning and real time inference for any selected operational metric from the feature store. Model training layer gives the ability to an operator to select, and tune required model to trial run and fine tune until the operator satisfies with required precision, recall against operational metrics and save the configuration for real time prediction. Model training can be supervised, semi-supervised, or unsupervised. The real time prediction layer performs prediction on the scheduled frequency against the scheduled time and stores the results. High level use cases like Anomaly prediction, Root cause analysis, automatic log mining, release management is performed. Notification layer alerts users on the configured channel like messenger, email etc. The data can also be visualized via an operational dashboard. The dashboard also provides a visualization of the output generated from the AI/ML engine and allows operators to provide feedback which can be incorporated in the models for better AI/ML predictions. With some business logic or high-level rules, we also allow generation of a report, which will alert the operations team via e-mail, SMS, or messenger. The major challenges in building AIOps architecture are listed in the following subsections.

2.1. Real-Time Data Processing

Real-time data processing is critical for AIOps, as timely prediction of application issues and outages brings strong value to the operators. This involves designing suitable architecture that can handle complex tasks including data ingestion, model prediction, and alerting in real time.

2.2. Scalable Architecture

Scalability is another key aspect in AIOps since the architecture must handle thousands of metrics coming from different application services, ingesting real-time metrics every second, and inferencing anomalies at seconds' level. For example, for a scale of one thousand metrics that are ingested every minute and inference also made at minute level, we are looking at more than 1 million model inferences per day.

2.3. Data Storage and Retrieval

As data is the key for any ML tool, access to historical data will help the ML model train over a longer duration which can lead to a better model fit and reduced bias. AIOps architecture facilitates data storage and retrieval for longer durations with the ability to retrieve data with minimum latency.

2.4. Real-Time model training and model packaging

With AIOps, operators can experiment with multiple machine learning models, train a model for a specific metric with historical data, package the model for inference instantly, and schedule the model for real-time inference on that metric. Operators can also provide feedback on the test inference and fine tune the model in real time. So, real-time model training and packaging is one of the core requirements fulfilled in our AIOps architecture.

3. AIOps Use Cases

Through AIOps, we are aiming to empower DevOps teams to perform their tasks efficiently. The following are the major use cases which we plan to incorporate in the AIOps platform for the benefit of operators:

3.1. Proactive Monitoring

Operation teams have been using monitoring system and observability platforms to configure alerts for their applications. The drawback of such an arrangement is that the alert threshold remains static, whereas the applications, or their usage, evolve over time due to various reasons requiring operations teams to regularly adjust their alert thresholds to avoid noise caused by false alarms. Also, there can be cases where certain problems in the system go uncaptured.

For example, batch jobs happening at a specific time can create a spike in application metrics. Spikes during that time would not be an anomaly since it's an expected behavior but at other times it could be an anomaly. Perhaps the batch job is taking a longer time than expected and if so, this is an anomaly to notify. In a traditional setup, this could be missed since thresholds are static and don't take into account such changes in data behavior.

AIOps learns from historical data. Especially in this case, the model would learn from the time-series data pattern and seasonality trends to capture these abnormal behavior and notify accordingly.

3.2. Smart Alerting

Alerting in conventional systems is based on static limits. In AIOps, we take advantage of AI to have dynamic thresholds depending on the trend of data. Apart from the metric threshold rules, we can also have rules based on model confidence. Confidence of model can be described as the certainty or strength in prediction done by a model on certain data, whether it is anomalous or not. This confidence is developed by the model learning from historical data over a period of time. This can even surpass human performance since sometimes it's not possible for the operators to continuously monitor such high loads of data manually, whereas ML models can do that easily.

This feature is useful in predicting potential system latency or downtime which normally wouldn't be captured using metric threshold-based rules. Thus, this feature helps DevOps with such cases and in turn, improves the customer experience with our applications.

3.3. Topology Analytics and Root Cause Analysis (RCA)

Topology analytics is used to establish a dependency topological graph of application, network, and infrastructure for a complex system, and to drill down to the root cause of the issue. Since AIOps is ingesting live metric and log data, AI models can predict anomalies in real time. These predictions are further used to correlate anomalies between various system metrics using statistical techniques such as Pearson Correlation Coefficient [6], and give a list of most probable root causes. It will be a much faster process compared to the operator doing RCA manually over thousands of metrics at a time. This feature will immensely help the DevOps team in automatic identification of problems and help them reduce MTTR in general.

3.4. Log and Trace Analytics

Each application may also have an immense volume of application-specific log and trace data. By leveraging Natural Language Processing (NLP) tools, we can extract the log metrics and their context from certain error/warning entities present in such data and perform anomaly detection against it. Additionally, we can also do correlation analysis to correlate the log anomaly with application metrics anomaly which can help the operations team to quickly identify the root cause of the issue. Similar analysis can also be done for trace data.

3.5. Cohort Analysis

Building on top of RCA, cohort analysis can provide further insights into system performance. It can point the team towards the probable issue in particular parts of their system even though they haven't received any specific alerts on them.

Multi-variate time-series analysis, i.e., time-series analysis done simultaneously on multiple metrics and clustering on anomalies as well as error logs, can correlate and identify such groups of metrics or systems that are causing problems, and can notify the team accordingly in advance for better maintenance of their application, all leading to better customer experience.

3.6. Automated Remediation

Consider a case where an application is deployed in a Kubernetes cluster, a group of nodes used in a Kubernetes deployment. Assume there is a 'garbage collection' system metric emitted from the cluster pods indicating one of these: application might go to a bad state or there's a problem with the pod itself. In general, this metric count is ignored since it's a normal behavior of the application. In some cases, though, it could actually be a problem with the cluster pods. Time-series models can detect such changes in data trend patterns and provide alerts to the team. Going one step further, the system can automatically take some actions (e.g., auto-restart) for remediation before it results in any application downtime or customer impact.

3.7. Smart Orchestration

Many teams use major cloud providers' auto-scaling features or their own customized rule-based scaling for their application resources. Through AIOps, we provide them with more intelligent auto-scaling abilities. Since the model will be able to forecast demands by learning the pattern and behavior through historical data, it will recommend the most efficient option available. This will ease the load on DevOps in their capacity planning for their cloud and other such resources, and provide significant cost-savings for the application team. In this manner, AIOps will be useful for an organization to reduce their software operational costs.

3.8. Release Management

When new software or firmware is deployed, it is often done in a phased manner. After a small portion of deployment occurs, AIOps can be used to analyze the difference between the new deployment and the previous version using machine learning models. If there are some significant changes, the model will identify which parameter or characteristic has changed. These insights would help the operations team to make a decision on whether to continue with the deployment or to roll back to a previous version.

4. Anomaly Detection

Anomaly detection is the key for proactive monitoring. It can be used for data quality monitoring, and fault detection of IT operations. Anomalies can be observed in tabular data, timeseries, or temporal data, as well as in graphical images. However, when it comes to fault detection in IT operations, data is typically in the form of time-series and so time-series anomaly detection is performed in AIOps.

Anomaly detection refers to identifying any abnormal behavior or pattern of data from the learned normal pattern from historical data. These anomalies may indicate a problem or an interesting event. The learned model can be used to detect anomalies with varying degrees of probability, and to predict future data with certain confidence.

4.1. Anomaly Detection Algorithms

When it comes to detecting anomalies in time-series data, the task can be performed either in a supervised, a semi-supervised, or an unsupervised fashion depending upon the dataset.

- (1) Supervised anomaly detection is possible when we have annotated data of anomalies available. We can split the data with anomalies into train and test data, and train a machine learning model as a binary classification problem, which means, training a model to predict whether a point is an anomaly or not using the labels available as feedback to train the model. This method can be performed using any machine learning model that can be used for binary classification such as Deep Neural Networks (DNN), Support Vector Machines (SVM), Random Forest, Gradient Boosted Tree (GBT), eXtreme Gradient Boosting (XGBoost), and others.
- (2) Semi-supervised anomaly detection can be performed when we do not have labelled anomalies, but we have a significant amount of data that is normal and do not have many anomalies that we can use for training a model. This model is trained supervised using the normal data and it learns the normal data behavior. It then detects anomalies when any deviations are observed. A DNN model like Autoencoder, which is a kind of neural network that learns a pattern and tries to replicate it, can be used for this. We can also use other DNN models like Long short-term memory (LSTM) and Deep convolutional neural network (CNN) (e.g., DeepAnT [3]). One-class SVM, Gaussian Mixture Model (GMM), Kernel Density Estimation (KDE), and others can also be used for performing semi-supervised anomaly detection.
- (3) When we do not have labelled data, as in most practical situations, we use unsupervised anomaly detection techniques. In this, the train or test data may or may not have anomalies. The models in this class generate an anomaly score for every data point and we can select a suitable threshold depending upon the data to classify points as anomalies or not. This can be accomplished using statistical methods like simple Moving Average model or complex ones like Prophet Forecasting Model [4] (uses Fourier series) or SARIMAX (Seasonal Auto-Regressive Integrated Moving Average with eXogenous factors, an extension of the ARIMA model - Auto-Regressive Integrated Moving Average). These approaches learn the periodicity or seasonality and the trend of the time-series data, and detect when the data deviates from the normal pattern. Density or distance-based outlier detection algorithms like k-Nearest Neighbors (kNN), Isolation Forest, or Local Outlier Factor can also be used to perform unsupervised anomaly detection.

With users' feedback, the detected anomaly can be labelled as true or false detection, and stored in the database. The historical annotations of ground-truth labels can be useful to evaluate the performance of the anomaly detection model in terms of evaluation metrics like Precision, Recall, and F1-Scores [7]. Also, they can be used to train a classification machine learning model for performing supervised anomaly detection. By getting feedback from the user (an

operator) for false positive detections, we can also adjust the threshold for anomaly detection, which is particularly beneficial to improve the accuracy for semi-supervised and unsupervised models.

4.2. Anomaly Detection Platform

We built an anomaly detection platform for proactive monitoring (Figure 2) with the following functionalities.

- (1) **Algorithm selection:** We allow the user to select among multiple algorithms, which run in real time by learning the seasonality and trend from data. We have options to perform either unsupervised, semi-supervised, or supervised anomaly detection depending on the use case.
- (2) **Special events:** There can be cases when irregularity is expected on days like deployments, holidays such as Christmas, or special events such as a key NFL game. Anomaly detection algorithms, like Prophet or SARIMAX, will consider these as special events. We have the option in our AIOps platform to provide such dates in advance to prevent false alarms as the model treats such days differently than a normal day.
- (3) **User feedback:** The tool can also take feedback from operator if they feel that certain prediction points are false positives, i.e., false alarms or shouldn't be anomalies from their subject knowledge.
- (4) **Intelligent alerting:** We provide three rules for the user to configure alerts. They are:
 - a. **Everytime** – To trigger alert notification everytime a model detects an anomaly.
 - b. **Interval Threshold** – To send alert notification only if the percentage of anomalies detected by a model out of all the data points present over a certain time period (10 minutes, 1 hours, daily) specified by a user, exceeds a certain threshold (0% – 100%) which is specified by the user as well.
 - c. **Score Threshold** – To send alert notification if a model predicts an anomaly score (based on model confidence) greater than a specified score threshold which is greater than or equal to the anomaly detection score threshold.

We allow the user to configure one or more such alerts for a single prediction model depending upon their needs. We also let the user choose the severity of the alert (low / medium / high) which they can configure accordingly.

- (5) **Messenger notification:** We send alerts to the operations team via messenger with a snapshot of the data (and potentially root cause analysis report) and a link to the data dashboard. We also provide 'Pause Alert' buttons at different durations for the operator to pause alerting for a specific metric prediction.
- (6) **Alert history:** We store all the historical alerts for all metric predictions for the user to view from the dashboard anytime. In addition, we also allow the operations team to collect the alerts as a metric data into their metrics endpoint, if required.

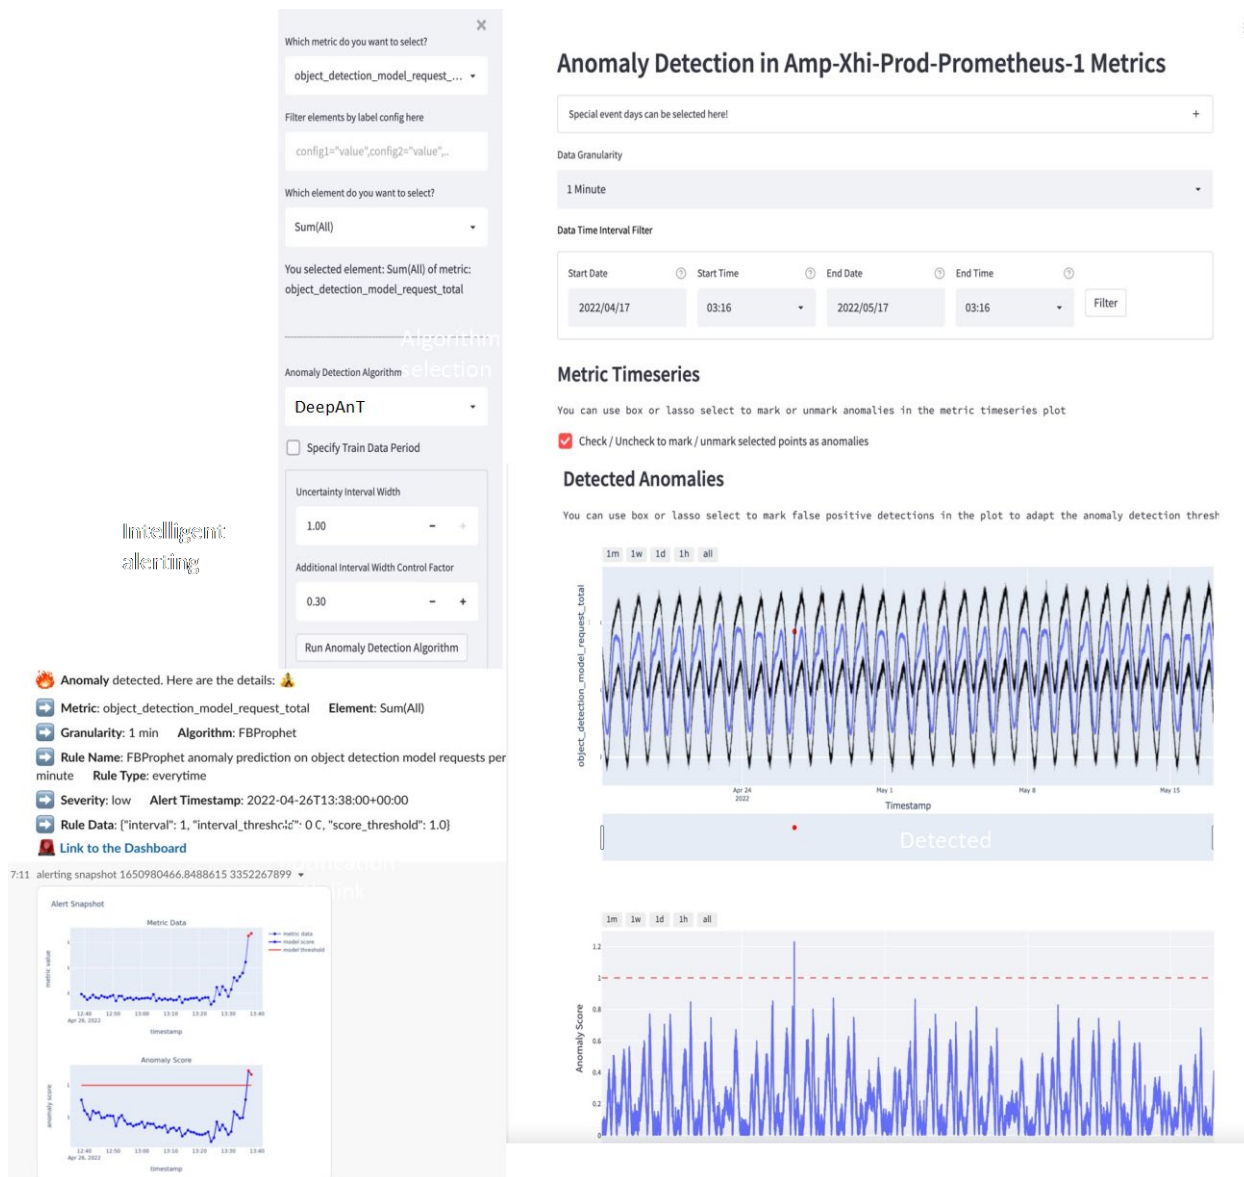


Figure 2 – Anomaly Detection Platform

5. AIOps – Operators' Perspective

This section describes AIOps from an operator's perspective as they will be the primary users of our platform, interacting with it in their routine work. This includes things that they need to know to onboard their application onto the AIOps platform, to configure anomaly thresholds, to provide feedback while evaluating a machine learning model against a specific metric, and to configure alerts for notification.

5.1. Onboarding a tenant

To onboard an application or a tenant, the AIOps tool needs read-only access to metrics data and log data sources. Log data is required if we are enabling automatic log mining to correlate with anomaly metrics. Operators can view the list of available metrics from an end point and start experimenting with specific metrics. After enabling the permission for AIOps to fetch data from their application tools, they need to provide a config file stating which metrics to pull and the endpoint configuration related details. Upon submission of this, the AIOps tool starts the data ingestion in real time.

5.2. Training a metric with model

The next step for an operator after onboarding an application is to configure an AI model. For that, they can choose the metric of interest and preview the metric trend over a selected time period. They can then preprocess the data, like aggregations, over a period and experiment with the choice of machine learning algorithms available in the AIOps platform.

The operator can perform model training live on the chosen data and preview the results on the fly. They can provide feedback like false positives or negatives to fine tune the model. After this, they can also validate the model with some test data on a certain duration of data. This process can continue until the operator is confident with the results for the data selected. Finally, if satisfied, the model can be packaged and deployed as it is ready for real-time inference.

5.3. Schedule a metric for real-time inference and prediction

With the AIOps platform, the operator can schedule one metric or a group of metrics for real-time inference configuring the frequency in any duration. Once configured, the AIOps platform will continue to run the selected model for metric anomaly analysis in real time.

5.4. Configure Alerts and RCA

Alerts can also be configured for the detected anomalies along with the severity and the considered interval or a model score threshold. The corresponding alerts can be notified to clients using messenger via their web APIs.

Root cause analysis can be performed automatically where the system performs correlation calculations on multiple anomalies detected on the metrics around pre-selected specific periods. This is performed by correlating the anomalies detected with the errors observed in log data in order to identify the most probable source error logs that could have caused the anomalies through a scoring mechanism which helps to rank the root cause errors. The AIOps tool recommends a possible hypothesis that corresponds to the most probable root cause of the problem which an operator can easily pick up for subsequent actions towards resolution. In this way, MTTR can be reduced greatly by using our AIOps platform.

6. AIOps Case Study: Connected Living Object Detection Operation

For our Connected Living business, we have millions of cameras in customers' homes. Our customers would like to get notification when objects (e.g., person, vehicle, and pet) or events (e.g., package delivery) of interest are detected from their cameras. The AI for Connected Living

team built the state-of-the-art house object detection algorithm which is used to efficiently detect person, vehicle, and pet.

The main challenge for the object detection operation are as follows.

- (1) There are many metrics to keep track of such as load (request per second), latency (upstream, downstream, and inference), CPU, and memory for each node.
- (2) The load (request per second) changes dramatically between day and night. There is a greater load in the daytime than in the night which is reasonable as generally human life is busier in daytime. The previously used static threshold rule-based alerting is not adaptive to address this issue.
- (3) The application has lots of log data. It logs the interactions of the object-detection module with the input metadata (from camera and the backend platform). Once there is an alert, the operations team often needs to dig into this log data to identify the root cause.

We onboarded this application onto the AIOps platform to help the operations team by addressing the above challenges. In that, we deployed time-series anomaly detection algorithm to alert the operations team in messenger, and then correlated the detected anomaly with log metrics anomaly to report what error messages are the probable root cause.

6.1. Connected Living Object Detection Operation

The metric data¹ of concern for this case study is the number of object detection model requests per minute. This metric captures the information of the load and has a well-defined daily seasonality pattern as we can see from the sample metric time-series plot shown below in Figure 3.




Figure 3 – Metric in Object Detection Operation

In the Figure 3, we can observe the metric time-series data plotted in blue. As we can observe, there are situations that lead to some sudden dips or spikes in the metric values that are of concern for the operations team to monitor. The points marked in red dots are the annotations

¹ We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

provided by a user by using a lasso or box selection tool available in the plot in the dashboard. By default, the metric data fetched from the metrics data endpoint will be unlabeled. Since this metric data is unlabeled and it has a well-defined seasonality, we chose an unsupervised anomaly detection algorithm for this case study. Prophet [4] is one algorithm that can be used for univariate time-series forecasting as it fits on the historical data and forms an additive model of the trend, daily, weekly, and yearly seasonality components, as well as holiday effects, and additional regressors that are available and learned from the data. By fitting such a model over this data for a training period of at least two weeks, we get an accurate model that learns the seasonality and forecasts with anomaly scores that can be derived from the uncertainty bounds generated by the model. The anomaly score is controlled by the interval width factor (which represents the percentile values) and an additional multiplicative factor controlling the width of the model prediction bounds. Since the model generates samples by Maximum A Posteriori (MAP) or Markov Chain Monte Carlo (MCMC) sampling techniques, along with the expected predicted value, we can also get the percentile values that are used for generating the uncertainty bounds and in turn, the anomaly scores. We can observe the predictions made by the model in the following time-series plot in Figure 4.

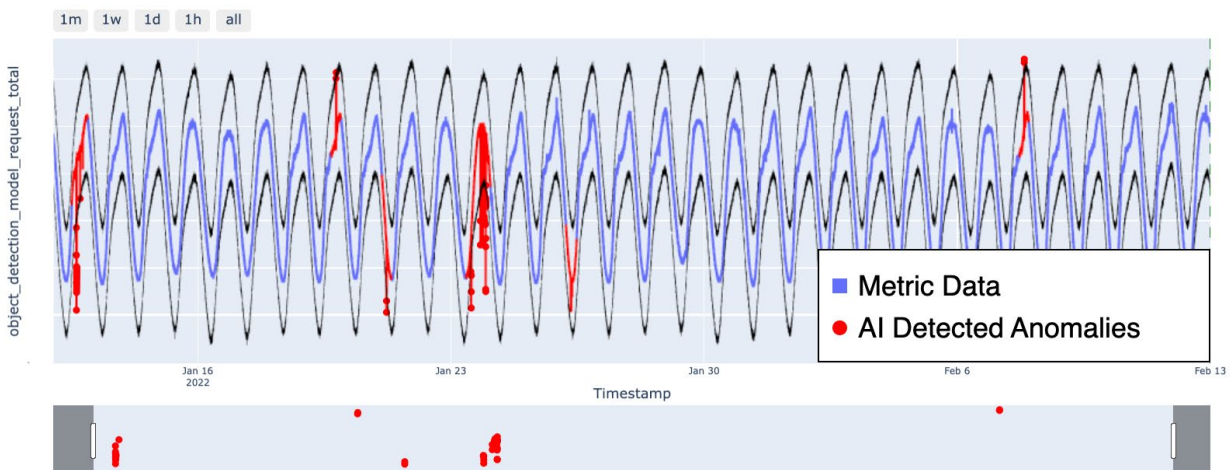


Figure 4 – Anomaly Detection Predictions in Object Detection Operation

We can see the points that are marked in red, and they represent the anomaly detections made by the model. Since the ground-truth labels are annotated in the metric plot, the sections of the time-series data where labels are provided have been marked with red lines while the other sections are in blue. These sections represent a tolerance interval to evaluate the model performance in terms of Precision, Recall, and F1 Scores. Having a tolerance is generally followed for time-series anomaly detection evaluation as forecasting anomaly in advance proactively, or after a pre-defined delay, is generally acceptable [5]. So, predictions happening within such contiguous tolerance intervals are all considered as precise predictions while evaluating a model's performance. Since this is followed in literature and by popular service providers, we have followed a similar approach.

The model evaluation is subjective to the dataset and what the operator chooses to mark as actual anomalies. Therefore, we cannot have evaluation metrics for all the models at all data periods. In

this section of the data shown in Figure 3 and Figure 4, however, the evaluation metrics are as follows:

- Precision: 1.0
- Recall: 0.87
- F1 Score: 0.93

Just as we can annotate labels in the metric plot shown in Figure 3, we can also annotate false positive predictions in the plot shown in Figure 4. That feedback is used to adjust the anomaly score thresholds to predict those points as normal. We fitted our model on this dataset and adjusted the thresholds to have a good prediction model that can accurately detect anomalies and alert the operations team for further investigation by setting an alert after the training and fine-tuning activities of the model are completed. The alerts are sent to a messenger channel with the information of the alert and the metric data and provides a snapshot of the metric and anomaly scores. There is also a shortcut link provided to the operator for them to look at the data dashboard directly from the message. As we can see from a sample snapshot of the messenger alert shown in Figure 2, the operations team was notified of the alerts in real time by the application and they were able to check why the object detection requests suddenly shot up. This real-time data pulling, model inference, and smart alerting capabilities that are provided by the application helped to ease the monitoring task performed by the operators for this case and provided timely alerts for them simplifying their operations.

6.2. Log mining and Correlation Analysis

Beyond anomaly detection and timely smart alerting, one of the common situations faced by an operations team when they encounter such alerts is finding the root cause for an anomaly. This operation is not straightforward, and the DevOps team will have to manually look over the log messages to identify the errors and warning messages in the logs within the period closer to the timestamp when the alert occurred. This manual operation is usually time-consuming and may also critically impact the businesses if the resolution or remediation cannot be taken within a certain time. A sample of such an operation has been shown in Figure 5.

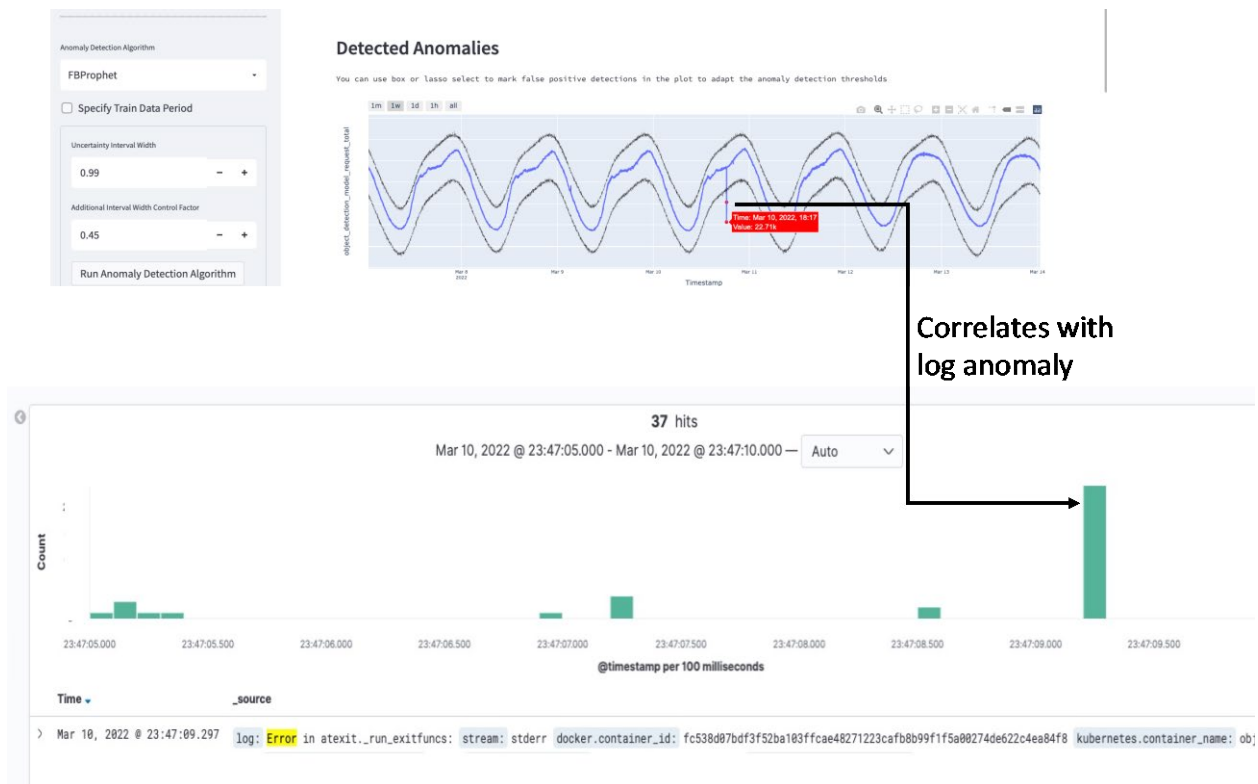


Figure 5 – Manual Log RCA in Object Detection Operation

In this Figure 5, we can see that an alert has been sent to the operations team of an anomaly detected in the object detection requests per minute metric. In order to find the root cause, the log messages in the dashboard were scrutinized and the team was able to identify a unique error log message that was logged near the timestamp when the anomaly was detected. As we can see that the requests had suddenly dipped, we observed that it was due to a deployment that had taken place at that time that was causing some errors from the Kubernetes pods running the object detection model containers. To reduce this manual effort, for this case study, we worked with the operations team to help them solve this problem using intelligent log mining and automatic root cause error analysis, in addition to anomaly detection and alerting.

In our application, we added the feature to pull and store log messages in addition to pulling and storing metrics data. We do not store every log message but only error and warning log messages through smart keyword searches for such terms in the logs. Apart from pulling and storing the logs, we also do intelligent text parsing to identify different kinds of log messages. With this method of segregating the error logs and labeling each type uniquely, we create error log count metrics for each error log message type, and store them as a separate metric.

So, the outcome of intelligent log mining provides us the error log metrics that can be used to individually train anomaly detection models to predict and alert whenever a pattern changes or an appearance of an error log message occurs. In addition to performing anomaly detection on them, these metrics can also be used for root cause analysis when anomalies are detected in metrics as we have a mechanism to correlate anomalies detected on a metric with other metrics over a selected data period and provide a ranked list of correlated metrics as probable root causes, as we'll see in the next case study.

Even though the error log metric creation gives us two potential use cases, the primary objective of log mining is to perform intelligent instant root cause analysis directly from the log messages when anomalies are detected in metrics, and to notify the operator of the alert along with the probable root cause error log messages in the messenger notification. For this, we have some scoring mechanism that is used to rank the error log messages that appear in the recent past period from the time the alert was identified, by comparing the frequency and distribution of the same over a much larger previous historical period. This helps us provide the operations team with the limited set of the most probable root cause errors instantly that they can quickly identify and perform remediation. The dashboard has the option to let the user select any alert that occurred in the past and view a more detailed root cause analysis report by showing the error log trends of each of the recently observed error logs as well the raw log messages as shown in Figure 6. This implementation reduced the time taken for root cause analysis tremendously and helped the operations team to be more productive.

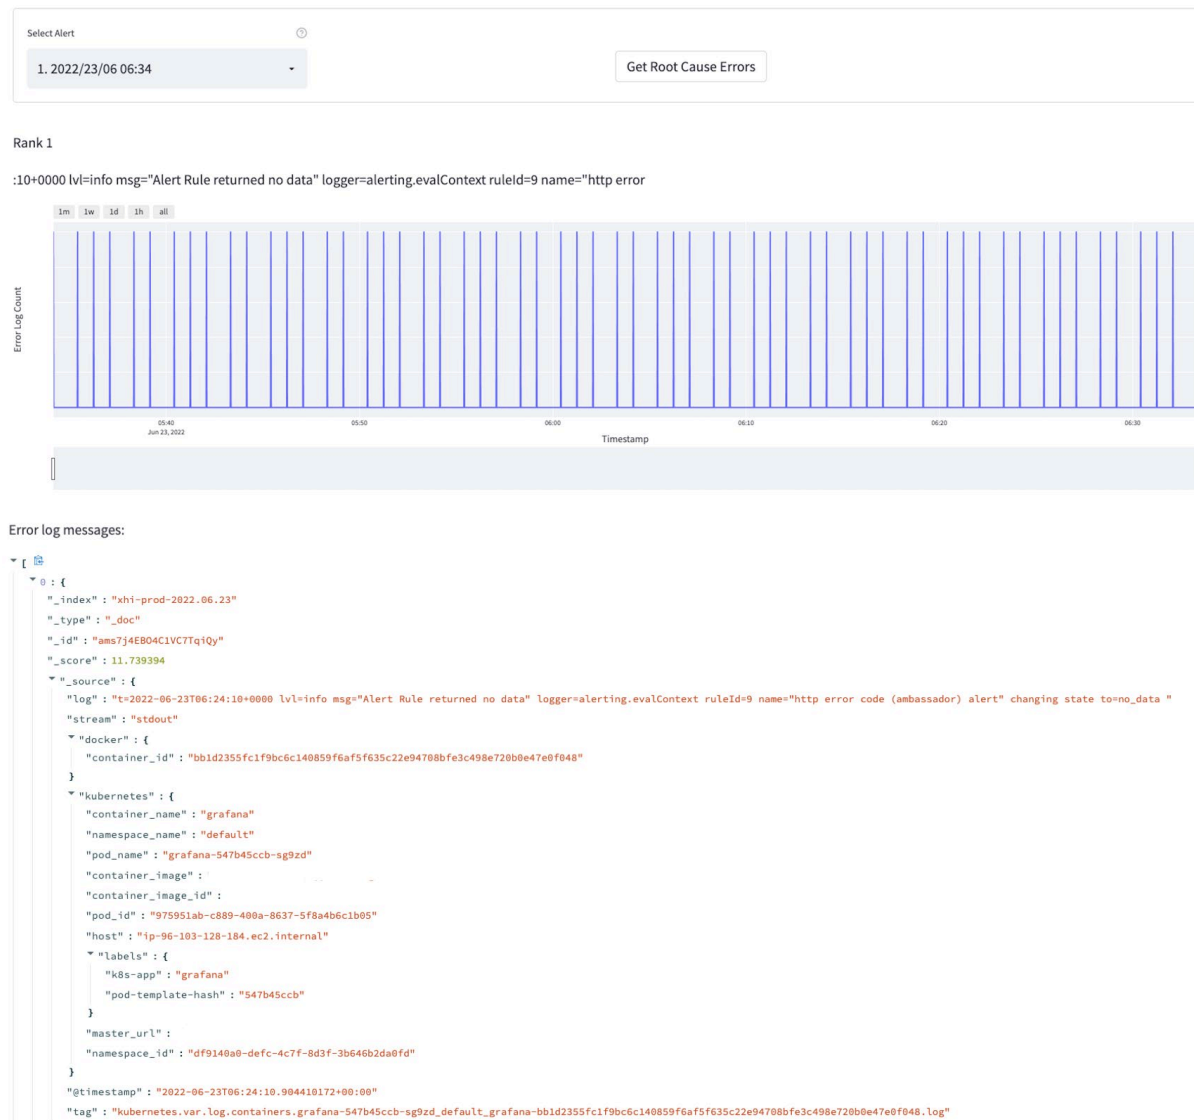


Figure 6 – Automatic Log RCA in Object Detection Operation

7. AIOps Case Study: Rex Browse and Search Operation

We have a large number of customers using an Xfinity box for cable TV and video streaming purposes. Here, the Rex platform is providing the video discovery features for X1 boxes. Specifically, Rex offers Keyword Search and Menu Browse services to X1 video clients. The platform also provides support for Video Recommendations, Personalization services and Video Metadata service.

Rex is deployed in several datacenters across the globe to cater to our business needs. Due to this, the DevOps team can handle significant complexity between various system metrics flowing from several sources. Their major pain point is performing RCA, especially in critical situations when there's time-constraint if the issue has a direct customer impact.

Let's take a look at an example using the below Figure 7.




Figure 7 – Manual Flow of Issue Root Cause Diagnosis by Rex Operations Team

Figure 7 shows the investigative steps performed by the operations team after a long diagnosis of an issue (dip in a system metric) affecting customers using Rex features. We can observe one of the regular flows of issue triaging and diagnosis performed by the Rex operations team in case of an issue. At each of those steps, the operator has to go through hundreds of metric graphs before drilling down to the next level of the problem. This process involves significant delays and inefficiencies. Because it is a customer facing application, time is of the utmost importance and MTTR converts to business value; the lower the MTTR, the better the customer experience. Through the AIOps tool, we provided features which greatly assisted the operations team in their routine work as we'll see in next sections.

7.1. Dependency Graph

The dependency graph of a system denotes the hierarchy in which different services inside the system are interlinked with each other. The network calls go from the top layered service to bottom layers according to the service dependencies. Let's look at a structure of a graph we had generated in a Rex case study. We have masked the names of the actual micro services used in Rex and have provided the contextual details alone for confidentiality concerns.




Figure 8 – Dependency Graph Created in Rex Case Study

In Figure 8, we can see the dependencies among different services (MS – Micro Service, used in the general sense of software engineering literature) and also the error flow denoted by color code (red – error, green – ok). This graph was constructed by our AIOps team using domain expertise from Rex technical operations team. We aim to automate the process of generation of such graphs in upcoming releases of the AIOps platform.

The benefits of having such dependency graphs for operations work are:

- This can be used by operations team to have a bird's-eye view on the entire system.
- Using such a hierarchical system topology i.e., hierarchy within system services, and applying RCA at each level, will help in swift identification of root cause. Also, the error flow captured will help the operator perform faster and detailed resolution.
- In this way, we are reducing the time and effort required by DevOps in their normal work, thereby reducing MTTR.
- Also, this graph can be used for efficient scheduling of on-call rotations since we know from the diagram what the affected systems are, and can estimate the efforts and expertise required to address those cases.

In the next section, we will see how we leverage AI techniques in our platform to help the Rex team with RCA.

7.2. Root Cause Analysis with Dependency Graph

As we observed in the earlier sections, manual RCA work required a great deal of time and effort by the operations team, increasing the time taken for tracing and resolving the issue. We have a feature for automating the RCA task in our AIOps platform. We performed anomaly detection on all the relevant metrics that we ingested into the platform. These are the metrics used for monitoring by the Rex operations team and are mostly custom configured metrics using queries. We then correlated the anomalies detected in one metric over a time period with those of other

metrics and generated a ranked list of probable root causes. Figure 9 shows the results of performing automatic root cause analysis using the AIOps tool for Rex case:

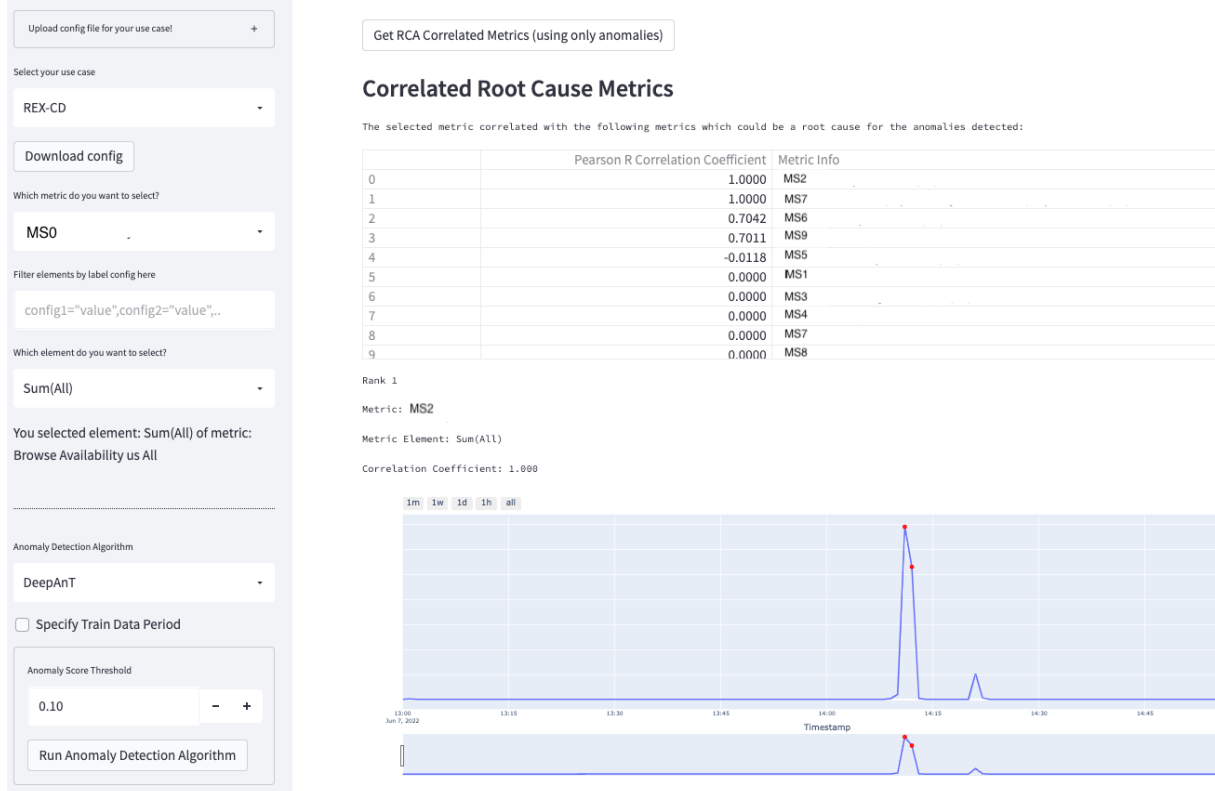


Figure 9 – AIOps Showing Automatic RCA Correlated Ranked List

As we can see from Figure 9, the tool showed the ranked list of metrics whose detected anomalies correlated with the anomalies detected on the selected metric. If we take a metric from one microservice, say MS0, when anomalies are detected and alerts are notified, the operations team can check for all other metrics from other microservices like MS1, MS6, and others where we observe correlated anomalies.

The AIOps tool shows the Pearson correlation coefficient score in the ranked list for the operator to understand how well the metrics correlate. In addition to showing the ranked list, the tool also shows the individual metric time-series plots of the correlated metrics in the same ranked order, for the operations team to quickly verify the correlations and drill down further. We can see the further list of correlated metrics in Figure 10.

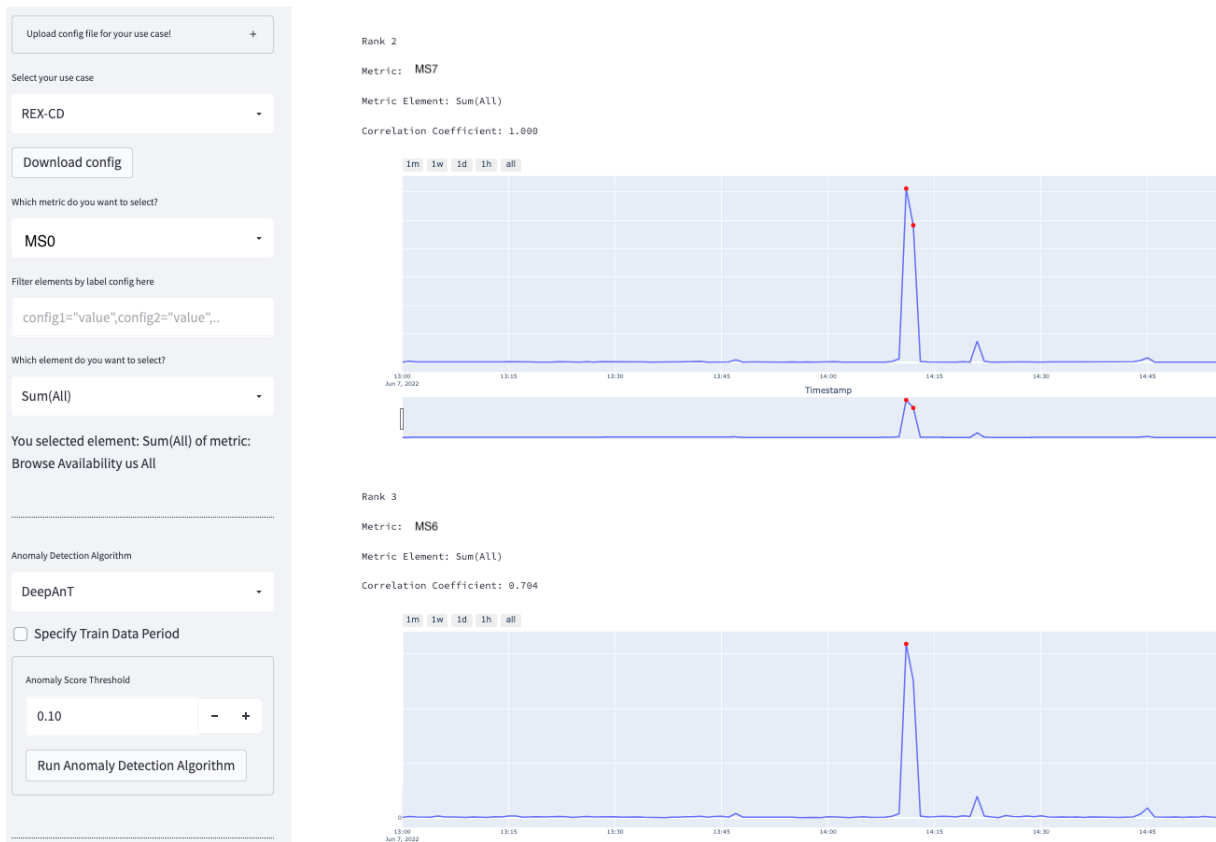


Figure 10 – AIOps Showing Probable Root Cause Metric Time-series

As per the manually identified dependency graph, we observed the root cause metrics to correlate well with the parent metric. As automatic dependency graph generation from user input is something we have in our pipeline, currently the correlations are shown across multiple levels of the hierarchical topology. However, we plan to improve the application to perform this root cause analysis at each level with only their dependent metrics (immediate child nodes) and after identifying the most probable root cause(s) with some threshold, automatically drill down further from that level and so on using dependency graph traversal as shown in Figure 8, instead of performing a correlation over all the metrics across all levels of hierarchy. This way, the operation will be optimized in the application when we have a smaller beam width for searching.

7.3. Failure Prevention and Auto Remediation

The Kubernetes pods run on Java Virtual Machine (JVM). They periodically collect garbage when requests sent by clients back up and physical memory runs low. Excessive Garbage Collection (GC) can also be a sign of the overall service going into a bad state. One or two pods collecting garbage at one time is not problematic, however GC running for hours unchecked in a self-reinforcing loop can infect other nodes. It is at this stage that drops in a metric like availability become noticeable. Any pod doing abnormal GC should be restarted after a preset period, but only if the outage is isolated to it. To accomplish this goal, we maintained a count of nodes doing GC at any given time. A two-pronged approach was used to highlight pods that can be restarted:

- First, we ruled out a system level outage by looking at the percent of pods doing GC as shown in Figure 11. The normal threshold can vary by the microservice and data center. Some data centers which are physical and running older, slower machines can be more failure-prone than others. Some microservices may be more memory intensive. Checking every such combination manually was not possible.




Figure 11 – Percentage of Pods Having GC activity per 5-min Period

- Second, assuming that there was no system level outage, we sought to find the pods with anomalous GC activity. This is a user input, defaulted to x standard deviation multiples of normal GC activity. The on-call messenger channel was alerted with the IP addresses of pods exceeding this threshold with the suggestion to reboot. The kill signal was also sent automatically as part of an auto remediation AIOps use case.




Figure 12 shows some of the malfunctioning pods. We had significantly narrowed down the list of potentially malfunctioning pods to around 2% of all pods, saving precious time during monitoring.

8. AIOps Case Study: Release Management for RDK Firmware

When a new RDK firmware version is released for Camera (RDK-C), Broadband (RDK-B) or video (RDK-V), we usually adopt a phased rollout. Usually, a small portion of customers with random selection (1%, or 5%) will be first deployed with the new firmware, and the operations team will closely monitor potential performance issues.

The challenges lie in several aspects:

- (1) There are often increasing device-specific issues, such as performance or stability, and at the same time there are hundreds of telemetries and parameters for the operations team to track. Some key parameters to monitor for each firmware version include VOD (Video on Demand), Linear, WHIX (Wi-Fi Happiness Index) [8], SpeedTest, CPU (Central Processing Unit), and Load among others.
- (2) Identification of new release issues are heavily dependent on Call-In-Rate (CIR) and high-level call movers/truck rolls, i.e., no-signal and no-block-sync. This results in long lead time for another iteration.
- (3) RCA and triaging need manual review of multiple boxes/examples. This is time consuming.

The RDK team deployed our release management AI component to drive release decisions in an automated manner. We can evaluate the impact on sub-populations (e.g., targeted/control for the specified segment such as region, CMTS (Cable Modem Termination System) version, HDCP (High-bandwidth Digital Content Protection) version, or accounts with pods) and firmware segments. Thus, the operations team can identify the anomalies with limited deployment, identify the root cause, and respond to the events quickly.

8.1. Machine Learning Model

We built a non-linear classification model to classify gateways with two firmware versions using one day's data where a new firmware version was deployed into the field. We focused on gateways with the same model to minimize hardware specific difference, and we additionally balanced the dataset such that the quantity of gateways with old version was roughly the same as those with new version.

The labels to the classifiers were the firmware versions. There are two major types of model features: counts of RDK-B telemetry key occurrence and certain measures on gateway usages (e.g., CPU usage, memory usage, Wi-Fi signal strength). Features were engineered from telemetry aggregated within a time window of 24 hours. Specifically, for a gateway with new version, features were collected within the 24 hours when the new version was deployed, and during the same 24-hour period, features were also collected for gateways with the old version. Our hypothesis was that if the version changes caused unexpected errors or performance change, then the classifier using RDK telemetry would be able to accurately differentiate the two firmware versions and indicate what telemetries were significantly impacted by the version change, and possibly point to the root cause.

8.2. Model Results

We built the model on a sample of approximately 200K gateways of a specific type. Our model was able to differentiate the two firmware versions with a high accuracy (Figure 13 Left). Based on feature importance score from the model, we were able to rank the features based on their

impact on the model performance. One of the top ranked features was the Wi-Fi signal strength on 5G band (5GRSSI_split). We further investigated that feature and found that there was a significant shift between overall signal strength distribution before and after the version change (Figure 13 Right). Here the old firmware version is 3.3p19s1, and the new one is 3.4p3s1.




Figure 13 – Precision Recall Curve for the model (Left Figure); 5G WiFi Signal Distribution for Old Firmware Version and New Version (Right Figure)

We confirmed the change with RDK-B team and learned that the shift in distribution was due to a bug fix on the reporting of Wi-Fi signal strength, i.e., old version tended to artificially report better signals with a 10 dB difference. The machine learning model we developed here not only helped us to gauge whether a new firmware release might cause significant difference in gateway performance, but also helped RCA for further investigations.

9. Conclusions

We presented an introduction to AIOps, and discussed our AIOps platform, which includes anomaly detection, root cause analysis, and release management, among many other use cases. We onboarded three applications using the AIOps platform and demonstrated the effectiveness of the platform regarding improvement in the experience and productivity of operation teams, and potential reduction in operational cost. As next steps, we would like to further evaluate the impact of AIOps quantitatively.

10. Acknowledgements

We would like to thank the following teams for the discussion and close collaboration with AIOps team to onboard the application: (1) AI for Connective Living team, especially Noam Krendel, Don Tolley, and Luke DeLuccia. (2) Rex team, especially Li Cao, Dirk Walter, Luciano Polo, and Albert Ogonevskiy for brainstorming the use case, and (3) RDK Analytics team, especially Matthew Long. We also thank Jan Neumann and Amit Bagga for their guidance.

Abbreviations

5G	5 th Generation
AD	Anomaly Detection
API	Application Programming Interface
ARIMA	Auto-Regressive Integrated Moving Average
AI	Artificial Intelligence
AIOps	Artificial Intelligence for Information Technology Operations
CAGR	Compound Annual Growth Rate
CD	Continuous Delivery
CI	Continuous Integration
CIR	Call-In-Rate
CMTS	Cable Modem Termination System
CNN	Convolutional Neural Network
COVID-19	Coronavirus Disease 2019
CPU	Central Processing Unit
DeepAnT	Deep Learning-based Anomaly Detection for Time-series
DevOps	Software Development and Information Technology Operations
DNN	Deep Neural Network
e-mail	Electronic Mail
GMM	Gaussian Mixture Model
GA	General Availability
GBT	Gradient Boosted Tree
GC	Garbage Collection
HDCCP	High-bandwidth Digital Content Protection
IIM	Intelligent Infrastructure Monitoring
IoT	Internet of Things
IP	Internet Protocol
IT	Information Technology
JVM	Java Virtual Machine
KDE	Kernel Density Estimation
kNN	k-Nearest Neighbors
KPI	Key Performance Indicators
LSTM	Long Short-Term Memory
MAP	Maximum A Posteriori
MCMC	Markov Chain Monte Carlo
ML	Machine Learning
MLOps	Machine Learning Operations
MS	Micro Service
MTBF	Mean Time Before Failures
MTTR	Mean Time To Resolution
NFL	National Football League
NLP	Natural Language Processing
RCA	Root Cause Analysis
RDKit	Reference Design Kit
RDKit-B	Reference Design Kit for Broadband
RDKit-C	Reference Design Kit for Camera

RDK-V	Reference Design Kit for Voice
RM	Release Management
RSSI	Received Signal Strength Indicator
SARIMAX	Seasonal Auto-Regressive Integrated Moving Average with eXogeneous factors
SCTE	Society of Cable Telecommunications Engineers
SMS	Short Messaging Service
SVM	Support Vector Machine
TSDB	Time Series Database
VOD	Video On Demand
WHIX	Wi-Fi Happiness Index
Wi-Fi	Wireless Fidelity
XGBoost	eXtreme Gradient Boosting

Bibliography & References

- [1] How Much Time Are You Wasting on Manual, Repetitive Tasks? Survey by Smartsheet
<https://www.smartsheet.com/content-center/product-news/automation/workers-waste-quarter-work-week-manual-repetitive-tasks>
- [2] AIOps Platform Market Forecast to 2028 - COVID-19 Impact and Global Analysis by Component, Deployment, Organization Size, and Vertical, Reportlinker, Apr. 2022.
- [3] M. Munir, S. A. Siddiqui, A. Dengel and S. Ahmed, "DeepAnT: A Deep Learning Approach for Unsupervised Anomaly Detection in Time Series," in IEEE Access, vol. 7, pp. 1991-2005, 2019, doi: 10.1109/ACCESS.2018.2886457.
- [4] Taylor SJ, Letham B. 2017. Forecasting at scale. PeerJ Preprints 5:e3190v2
<https://doi.org/10.7287/peerj.preprints.3190v2>
- [5] Hansheng Ren, Bixiong Xu, Yujing Wang, Chao Yi, Congrui Huang, Xiaoyu Kou, Tony Xing, Mao Yang, Jie Tong, and Qi Zhang. Time-series anomaly detection service at Microsoft. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pages 3009–3017, 2019.
- [6] Pearson correlation coefficient wiki page,
https://en.wikipedia.org/wiki/Pearson_correlation_coefficient
- [7] Precision, recall and F1-score wiki page, https://en.wikipedia.org/wiki/Precision_and_recall
- [8] Krithika Raman, Charles Moreman, THE WiFi Happiness Index (WHIX), SCTE Fall Technical Forum, 2011

Approaching the Cosmic Speed Limit:

Introducing Hollow Core Fibers for Low Latency and High Capacity

A Technical Paper prepared for SCTE by

Venk Mutalik

Fellow

Comcast

1800 Arch Street, Philadelphia, PA 19103

+1 (860) 262-4479

Venk_Mutalik@Comcast.com

Amarildo Vieira, Principal Optical Engineer

Dan Rice, Vice President

Bob Gaydos, Comcast Fellow

Elad Nafshi, Chief Network Officer

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background	3
3. Holey Fibers to Hollow Core Fibers	4
4. Latency, Capacity and Reach	6
4.1. The Latency-Capacity Continuum.....	6
4.2. Comparing SCF, LEO and HCF Latencies	7
4.3. Geographical Flexibility	9
4.4. Critical Infrastructure Implications	9
5. Features of Hollow Core Fibers	10
5.1. Taxonomy of Optical Impairments	10
5.2. Linear Impairments: Fiber loss.....	11
5.3. Linear Impairments: Fiber Dispersion	11
5.4. Linear Impairments: Fiber Splicing and Backwards Compatibility	12
5.5. Non-Linear Impairments:.....	13
6. Hollow Core Fiber in Comcast	14
6.1. Establishing Latency Improvement using a Hybrid Loop Test.....	16
6.2. Spider Diagram Grand Summary	18
7. Conclusion.....	19
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 – Illustrating different approaches to Hollow Core Fibers [3,4].....	5
Figure 2 – Hollow Core Fiber Loss over the decades [3].....	5
Figure 3 – The Latency - Capacity Continuum	6
Figure 4 – Solid Core, Low Earth Orbit and Hollow Core Approaches to Latency	7
Figure 5 – Comparing Latencies of the Various Approaches	8
Figure 6 – Increasing the Latency Envelope and Enhancing Geographical Flexibility	9
Figure 7 – Taxonomy of Optical Impairments	10
Figure 8 – Illustrating fiber loss vs. wavelength for various fibers [3]	11
Figure 9 – Illustrating dispersion and Latency envelopes for SCF and HCF.....	12
Figure 10 – Splicing HCF-HCF and HCF-SCF fiber cores	13
Figure 11 – Block Diagram Illustrating the Comcast Deployment	14
Figure 12 – Illustrating various direct detect and coherent spectra and the passives passband	15
Figure 13 – Bi-Directional Optical Wavelengths over the Comcast HCF Link.....	16
Figure 14 – The Hybrid Loop Test Illustrating Latency Reduction.....	17
Figure 15 – Illustrating Latency improvement with HCF (top) over SCF-SMF (bottom).....	18
Figure 16 – Qualitative Comparison of Fiber Attributes	19

1. Introduction

Folks already know that the speed limit in the Universe is the Speed of Light ... nothing is faster. However, many will be surprised to know that light travels slower in an optical fiber than it does in free space! Due to light-matter interaction, the speed of light, which is about 300 million m/s in free-space is 'only' 200 million m/s in glass. But in a world that prizes getting the latest information fast, this fundamental additional latency can sometimes be limiting for many applications. This was dramatized best in the Selma Hayek movie the "Humming-Bird Project", where she builds a set of line-of-sight microwave links to eliminate milliseconds worth of latency between the Chicago and New York bourses.

In this paper Comcast introduces Hollow Core Fibers and the cutting-edge work being done to reduce this fundamental latency and enhance capacity all at the same time. New fiber technology creates a fiber guiding mechanism that has a core that is hollow and consequently allows light to travel in the hollow of the fiber as-if in free space, but along the fiber path!

We describe a system which is a hybrid of hollow core fibers and standard single-mode fibers. Keeping in view that Comcast operates core and access networks this system has amplified and unamplified bidirectional Coherent systems of 400G and 100G wavelengths along with direct detect systems of 40G, 25G and 10G wavelengths all simultaneously on a single strand of fiber, in what we believe to be the first such system in the world.

In addition to describing our test system and initial results, we also analyze some of the benefits in addition to latency reduction that are related to reduced optical non-linearities and other effects and the wider spectrum and the consequent Capacity enhancements possible. Practical deployment strategy of fibers is essential at Comcast, the ability to use a portfolio of fibers for some of the emerging low latency market while also driving fiber technology deeper into the network is of the essence of our approach towards the industry initiative of 10G.

2. Background

Everyone likes to peer into the future, that most inscrutable of entities right in front of all of us. In real life, all of us generally uncover it along with everyone else, but efforts to glimpse it just before everyone else is a passion that drives great technology.

The fastest way for information to travel is at the speed of light in vacuum, that speed, set at 299,792,458 m/s is a universal constant. In other transparent media, the speed of light is slower due to light matter interaction defined as the refractive index of light. Refractive index of light is the ratio of the speed of light in the medium vs. the speed of light in vacuum. For water this number is 1.3, and the mismatch in the speed of light in air and water is what makes a straw in water look like it is shifted. Furthermore, the same effect also allows for reflections at the boundaries and these features are exploited to build optical waveguides to help guide light along non-rectilinear paths.

Optical fiber waveguides were conceived around 1965, and single mode optical fibers, a type of optical waveguide have become the premier means of data transmission for the past 5 decades or more. A careful tradeoff on the geometry, the refractive index differential between the core and the cladding gives rise to our modern Solid Core Single Mode Optical Fiber (SCF-SMF). Standard Single Mode Fiber typically consists of a slightly higher refractive index glass core no more than 9um in diameter surrounded by a slightly lower refractive index glass of 125um cladding. Many different types of optical fibers all more or less following the method outlined above, such as non-zero dispersion shifted fiber and large effective area fibers have been constructed and some of them are in use today as well. In the past 5 decades, every

single aspect of the optical fiber has been studied and analyzed, every material composition and build has been researched to create a fiber that has the lowest transmission loss, the highest micro and macro-bend performance and the greatest tensile strength for deployment across the world.

Around 1995, there emerged a revolutionary new concept of optical waveguides, in what looked like a solution looking for a problem. What if instead of solid waveguides guiding light, the light were guided within holes of an optical fiber! Among other things this was a radical concept because it was assumed then that light guiding could only occur when the core had a higher refractive index than the cladding, in this case the situation seemed to be the reverse.

These early fibers were called ‘Holey Fibers’ for the holes in the fibers. Soon reduced latency, higher power delivery possibilities came to be identified with these fibers. But before that, significant issues associated with fiber attenuation, micro and macro bend losses, and production techniques would need to be solved, and are discussed next.

3. Holey Fibers to Hollow Core Fibers

Fiber attenuation is always of prime concern for any optical fiber. Optical fibers have minor imperfections in the core composition, and these lend themselves to light scattering and reflections of very small magnitudes called Rayleigh Scattering in optical fibers. Over long links of fiber, these scattering effects diminish the amount of light that can get thru to the other end and are a primary cause of fiber loss [1,2].

When holey fibers came about initially, they were very hard to build, but in addition the built fibers had very high losses. Some of these fibers were 10 or 20dB/km of loss, while the standard SCF-SMF was then at around 0.25dB/km or less.

There was a new evolution in holey fibers that enabled this loss to become dramatically less, and it was based on creating a lattice like structure in optical fibers. These fibers are called photonic bandgap fibers (PBG) shown on the right in Figure 1. A large central hollow core followed by a lattice structure around would effectively confine light in the hollow of the core and enable light guiding in the transverse direction (i.e., along the fiber length). Such fibers when originally manufactured had lower loss than holey fibers. Furthermore, the micro and macro bending losses of these fibers were quite good. In principle, since the core confinement of light was good or in other words Confinement Loss (CL) is minimal, it was theorized that the fiber loss could be arbitrarily low. However, it was found experimentally that the surface scattering loss due to imperfections in the surface of the lattice structure was so great that the Surface Scattering Loss (SSL) would be a dominant mode of loss of the optical fiber limiting its minimal loss to around 0.65dB/km. This loss value although not very good by SCF standards is still very good for low latency applications and is in use today for various such applications which will be discussed later in the paper.




Figure 1 – Illustrating different approaches to Hollow Core Fibers [3,4]

Around the mid 2010s a new type of hollow core fiber has emerged that has the potential to offer lower loss. This is the Nested Anti Resonant Nodeless Fiber (NANF) shown on the left in Figure 1. In analyzing the SSL of the PBG fiber it was found that most of the entire lattice structure did not really do much to guide light, but actively participated in increasing SSL. And furthermore, it was found that having fewer non touching capillary structures guided light just as efficiently without spilling light. It is helpful to think of the light confinement as a series of Fabry Perot structures reflecting light and holding it in the hollow core of the fiber. A further innovation of using nested capillary tubes in the fiber improved core confinement and reduced Confinement Loss. By playing with the capillary thickness and placement further improvement occurred continuous single mode low loss operation across the O and the C bands.




Figure 2 – Hollow Core Fiber Loss over the decades [3]

Figure 2 is a graph of the impressive improvement in anti-resonant fibers in the past decade and a half. True to form, at the post deadline paper at the OFC this year, there were reports of a double nested anti resonant nodeless fiber (DNANF) of 0.17dB/km thus matching the loss of SCF-SMF fiber at the C-Band and at 0.22dB/km at the O-Band thus breaching the loss of the SCF-SMF [5].

As discussed before, the total loss of optical fibers are due to intrinsic mechanisms comprising the material purity and Rayleigh scattering loss, confinement loss and surface scattering loss and extrinsic losses due to micro and macro-bending loss due to cabling or installation. In the HCF, Rayleigh and surface scattering losses are very low, and confinement loss can be decreased by double nesting as well as by increasing the core diameter. But increasing the core diameter would probably increase the bend losses while multiple nestings could increase manufacturing complexity. But in any case, the idea is to let the industry know that such fibers are on the horizon.

4. Latency, Capacity and Reach

Our work on Hollow Core Fibers has primarily focussed on the NANF type of fibers. These fibers offer the potential for low optical loss, have low Chromatic Dispersion, have high power handling capability, a wide bandwidth covering the C+L band, all while having latency that approaches the fundamental limit of speed of light in vacuum. Of these, latency reduction is of fundamental importance, while other limitations such as higher losses or bandwidth could be overcome by optical amplification or higher allocated fiber counts.

4.1. The Latency-Capacity Continuum

Today it is common to see 100Gbps, 400Gbps and even 800Gbps type of transport rates per wavelength on core and access networks. But the rise in 5G and other such technologies has also focused our attention on latency or responsiveness of a network. With the understanding that latency and capacity form two independent metrics that define the quality of our network and there are many applications that require one or both metrics.

CAPACITY ==> High Low <==	File Transfer	Hollowcore Fiber-HFT/5G
		HSD-Com
		HSD-Resi
		LEO Sat
	GEO Sat	
	IOT	Ionosphere RF-HFT
	High <== LATENCY ==> Low	

Figure 3 – The Latency - Capacity Continuum

Consider Figure 3, in which Internet of Things (IoT) characterized by low traffic capacity needs and slow update requirements result in it being in the lower left quadrant. High Frequency traders and 5G with their need for high capacity and very low latency will fit in the upper right quadrant, these are best served by hollowcore fiber links. Businesses and campuses which do file transfers during the off-peak hours might benefit with very high capacity links but without the responsiveness of low latency needs. RF transmission utilizing the reflective properties of the Ionosphere for low capacity high frequency traders fit in the low capacity and low latency bucket in the bottom right.

Latency is a rather complicated topic depending upon multiple variables. At its most fundamental, it depends upon the time of flight of the physical medium, secondly, it depends upon the variopus electronic switches and routers that make up the circuit. Finally, it depends upon the various queues and buffers that form at the near and far side that gate information transfer. It is not surprising that industries with tight latency requirements have made great strides in reducing latency of their switches, thus some switches that are used by HFT are directly written in FPGA code and can approach ns latency. The latency of the the physical medium however is of fundamental importance and is described next.

4.2. Comparing SCF, LEO and HCF Latencies

As we have said before, light travel at roughly 300,000km/s or 300m/us in free space, but travel at ‘just’ 200,000km/s or 200m/us in SCF-SMF fibers. And the low Earth orbit call Leo is approximately 550 km above the surface of the earth. Therefore with this understanding we can now understand the latency implications of the three scenarios described in Figure 4. The left figure corresponds to time of flight latency of a conventional fiber, the middle corresponds to time of flight latency for a low earth orbit satellites free space links and the one on the right describes latency in a HCF system.

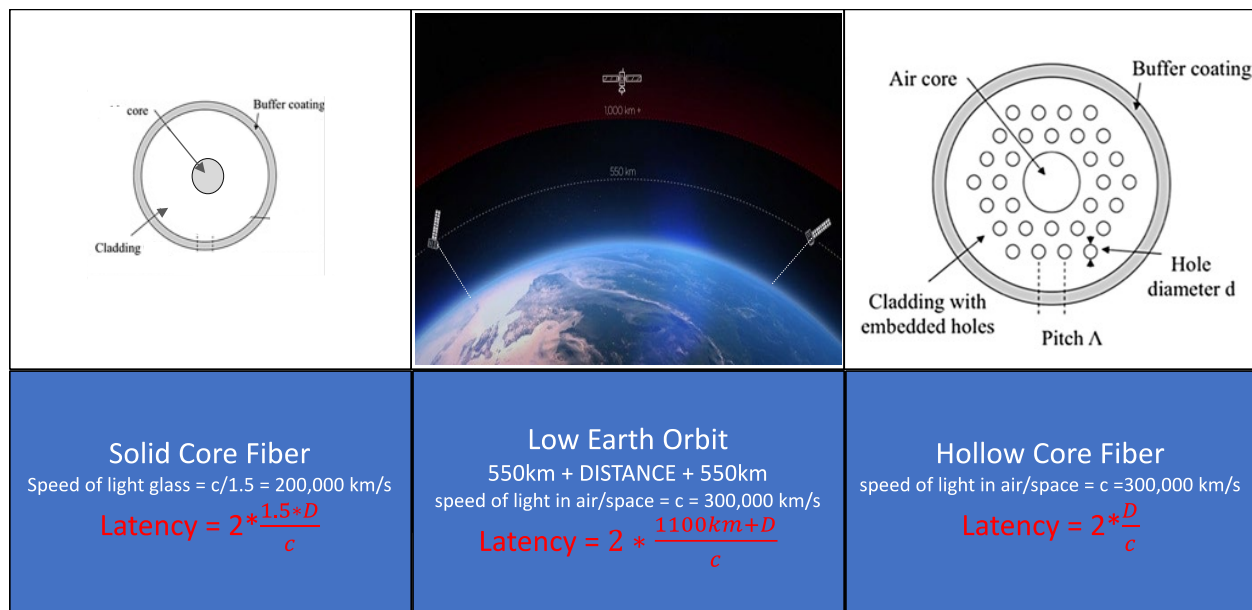


Figure 4 – Solid Core, Low Earth Orbit and Hollow Core Approaches to Latency

In a regular SCF-SMF, the two-way latency is just twice the distance divided by the refractive index of fiber material, which happens to be around 1.5 (actually ~1.4684 in the 1550nm region).

Many years back, satellite communication was rather latency prone since satellites were located at the geosynchronous/geostationary orbit around 36,000km above the earth accruing a delay of 250ms which

rather limited the scope of it in high-capacity communications. Recently, low earth orbit satellites have been launched which are just 550km away, the idea being to cover the earth in a sheath of LEO satellites to offer connectivity around the world. Furthermore, these satellites also can have inter-satellite communications, typically using 100G CFP2 type dual laser bi-directional coherent optics (a more detailed treatment is presented at a different paper in this Conference). Note that inter satellite transmission happens in free space where the refractive index is 1.0. In this case, the latency is significantly reduced, and in the limit may be as small as an additional transit time of 550km to go to one satellite and come back to the earth with the same 550km. Note there that we have not taken into account the processing time in the satellite opto-electronics.

Finally, use of HCF enables one to use the same geographical lines as the SCF-SMF, but with hollow core fiber. Thus, the time of travel is just twice the distance divided by the speed of light.




Figure 5 – Comparing Latencies of the Various Approaches

As summarized in Figure 5, HCF is ALWAYS the link with the least latency, but for very short distances, the delta between SCF and HCF time of flight is small, but always 50% faster in HCF than in the SCF. LEO Space communications, on the other hand start out with a higher latency, but by around 2500km, space communications becomes faster than SCF-SMF communications.

The implications of this are vast. If specific companies can figure out LEO sat communications that have very little optoelectronic delays in inter satellite communications, then they would have a huge ability to ‘see’ the future between the bourses of London and New York for example, which are more than 2500km apart. Within the US, for all locations West of Denver relative to Philadelphia, the LEO communications would have been the same! A hollow core fiber deployment on the other hand along straight line would always have a lower latency and would be at the universal speed limit and would never lead to such arbitrages.

While the case has been made for long haul transport, similar is the case within access networks as well. In the case of 5G front haul, where due to dispersion and latency, around 15km is an accepted distance between towers and central offices, the use of HCF can legitimately increase distances traveled to over 25km for the same latency envelope effectively doubling the area under reach for the cell towers connected to each central office. There is one other aspect of HCF that relating to fiber dispersion which will be described in a later chapter.

4.3. Geographical Flexibility

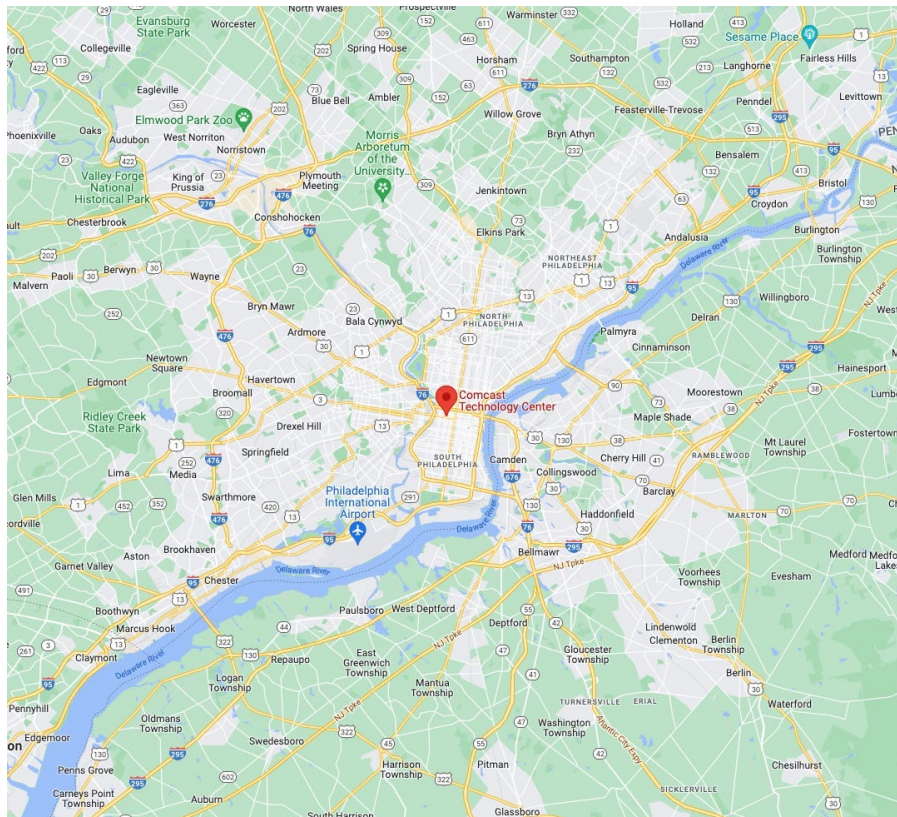


Figure 6 – Increasing the Latency Envelope and Enhancing Geographical Flexibility

There is a geographical flexibility case to be made here for the HCF as well. If we consider that a specific latency is needed from place A to place B, based on standard fiber, but the line between A and B is cut by additional natural or artificial barriers unanticipated at the time of design. In such cases, making a section of the path HCF could give a much-needed latency reprieve and enable the existing design to proceed naturally.

This concept of latency reprieve or latency trade can also be applied to a system in which the number of transactions is also a figure of merit. Here, the lower latency translates in quicker transactions, which in turn could enable - say - high frequency traders on HCF refine their positions continually many more times than the traders on SCF based on the much more rapid conversations.

4.4. Critical Infrastructure Implications

We previously discussed that the backward compatibility enabled the use of same terminal equipment. While this is true a potential additional benefit in the form of power savings might accrue on account of

the fact that individual DSP chips in pluggables may not need to work as hard to compensate for dispersion and optical non-linearities leading to less complex (and lower cost) pluggables and lower critical infrastructure costs.

5. Features of Hollow Core Fibers

The previous chapters have described in detail the evolution of Hollow Core Fibers and the most fundamental advantage of HCF, regarding the ability to offer the lowest latency possible. It turns out that in addition to this fundamental benefit, there are several secondary benefits that make the deployment of HCF (especially NANF) attractive in both core and access networks, what are discussed next.

5.1. Taxonomy of Optical Impairments

Presented below is the well-known ‘taxonomy of optical impairments’. Linear impairments such as chromatic dispersion and optical loss along with optical passives are presented in on the right-hand side, while non-linear impairments that depend upon the of light intensity in fiber are arranged to the left. Specifically, we distinguish between non-linear effects due to single wavelength and non-linear effects due to multiple wavelengths one below the other.

Describing effects of nonlinearities and their impact is an interactive process and will only be summarized here in the context of HCF. It must be emphasized that many of the linear and non-linear effects which are much more important for analog transmission of light that is characterized by high power levels and high optical performance requirements have been thankfully reduced in the migration to digital, but yet, many limitations unique to digital transmission still remain and may add up to significant power penalties reducing reach or capacity.

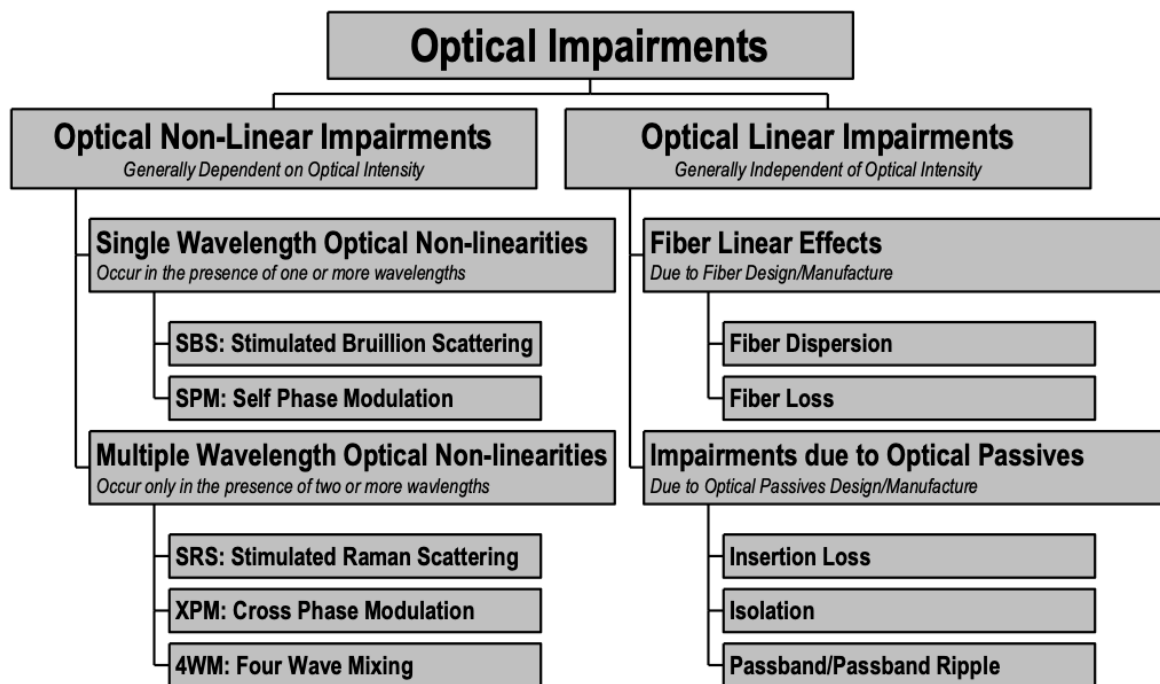


Figure 7 – Taxonomy of Optical Impairments

5.2. Linear Impairments: Fiber loss

Linear impairments are led by fiber loss concerns as fiber loss in present generation of HCF is around 1.5dB/km which limits deployment of long distances without amplification. The generation of dual nested fibers sport a loss of 0.17dB/km, which approaches the loss of SCF-SMF and will be quite suitable for long distance deployments. The current generation of fiber is good however for shorter distances and to proof out a number of other parameters that bear validating. It is important to note that the fiber has a continuous operational single mode region that cover the C and L bands. The bump in loss at the water peaks is attributable to the presence of water vapor (OH ion) in the hollow of the glass, one that may potentially be reduced by an appropriate purge of the core.




Figure 8 – Illustrating fiber loss vs. wavelength for various fibers [3]

5.3. Linear Impairments: Fiber Dispersion

Hollow core fiber has dispersion that is around 2.5ps/nm.km. For comparison, the SCF-SMF has a dispersion of around 17ps/nm.km in the C-band, yielding a 7x benefit of the dispersion parameter. In real life, this translates to longer reach for baseband transmission where the fiber reach decreases as the square of the ratio between baseband speeds as illustrated in this log-linear graph in Figure 9.




Figure 9 – Illustrating dispersion and Latency envelopes for SCF and HCF

For example, 10Gbps baseband signals routinely transmit over 80km, can now traverse 100s of km with minimal dispersion penalties on the HCF. Similar is the case crucially with 25Gbps which is an eCPRI standard was previously limited to 15km of transmission as that was what was possible over standard fiber, this can now transmit to over 80km as illustrated above. A point to note here is that the latency envelope of eCPRI can now be 25km rather than 15km with the HCF as well, thus indicating that HCF could bridge both the latency and reach envelopes as indicated above for various speeds.

In order to verify the benefits of the HCF fiber in access network applications, we verified error free performance of links using SFP28 DWDM CPRI and 40G QSFP DWDM optics over 20km of HCF currently deployed in the Comcast campus. In this verification optical amplification was required to overcome the HCF insertion loss. For comparison purpose we also ran performance tests with 10 and 20km of SMF and, as expected, we could verify the SMF dispersion limitation leading to system breakdown at 7 and 13km of standard fiber, respectively, without dispersion compensation. These results are presented in a later section, Wideband Direct Detect with HCF.

5.4. Linear Impairments: Fiber Splicing and Backwards Compatibility

It is common to find specialized fiber ribbon splicing equipment and software that can splice 12 or more fiber cores at the same time. Similarly, splicing equipment and software is available to field splice HCF.




Figure 10 – Splicing HCF-HCF and HCF-SCF fiber cores

But splicing a HCF to SCF is much more complicated as the geometry of the two fibers are different. Therefore, fiber adapter cables that have HCF on one end and SCF connectors at the other end are specially made in the factory and then used to light up fibers with conventional optical equipment in headends, hubs and nodes. This keeps ALL of the terminal equipment completely backwards compatible with all other equipment while enabling the use of HCF in the field.

In other papers in this conference, we have presented on new underground construction method called the CiC in which micro-fibers re-enter the conduit already holding RF cable and reuse the duct/conduit infrastructure already in use with minimum disruption. Most micro-fibers have 6 tubes within with 12 fibers each in each tube. In this case, one or more tubes could be dedicated to hollow core fibers while others for standard fiber and thus a portfolio of fibers may be installed when the conduit is re-entered. This type of innovative deployment may help bring fiber to the neighborhood where high-capacity latency sensitive endpoints may be located. This concept of a portfolio [6] of fibers is more general of course and may be used in other deployments as well in core and access fiber platforms.

5.5. Non-Linear Impairments:

Optical fibers provide excellent light guiding but in doing so also provide a long interaction length for light-matter interaction. When optical intensities increase in an optical fiber, this leads to non-linearities. All Non-linearities increase with fiber intensity. Some such as SBS stimulate the reflection of light back to the source, others depend upon the non-linear index of glass, which increase with optical intensity. Also called the Kerr effect, this has the unfortunate result that the light speed is modified by the light intensity. When this happens, either due to the self-channel intensity or adjacent/cross channel intensity, a range of non-linear effects arise. But the worst non-linear effect by far is the Four Wave Mixing effect. So called because the fiber itself generates ‘beats’ due to multiple wavelengths that can have a catastrophic effect on transmission. To make matters worse, these non-linear effects also interact with linear effects such as dispersion and polarization and can be intermittent. To tamp down on these effects requires a deep understanding of all effects and with it an effective wavelength plan and robust modulation format.

With fixed locations for amplification and the rise of flex grids and the desire to use C and L Bands together, the practical outcome of these non-linearities is unfortunately a rather severe limit on launch power. But low launch powers have the adverse effects of needing multiple amplification spots and the consequent degradation in signal quality. Furthermore, many non-linearities get exacerbated with the optical amplifiers themselves as these are also fiber-based devices. The net result is a delicate balance.

Not so the case in HCF. Since the core is hollow there is the elimination of light matter interaction. Furthermore, the non-linear index of air/vacuum is non-existent, therefore quite a large amount of light can be packed in to the HCF. The ability to launch larger amounts of light could lead to sparser amplifier spacing and in turn lead to longer reaches.

6. Hollow Core Fiber in Comcast

Late last year, Comcast acquired 20km of cabled HCF, of which 10km was cabled for inside plant performance and the other 10km for outside plant performance. These fibers were terminated in SC-APC adapters via the HCF-SCF adapter cables mentioned before. We began by ensuring that the fiber was reciprocal, by which we mean the loss tested identical from either side of the fibers. It turns out that both fibers had approximately 1.5dB/km of loss and the HCF was indeed reciprocal. Having done that, we deployed the fiber in Comcast.

The HCF was then deployed in a system comprising multiple direct detect and coherent systems. The configuration was 10km of the ISP HCF followed by various lengths of SMF and then finished up with the 10km of OSP fiber. All ends of the fibers are all connected with standard SC-APC connectors thus ensuring total backwards compatibility for this hybrid SMF and HCF system.




Figure 11 – Block Diagram Illustrating the Comcast Deployment

OTDR scanning of the HCF is a surreal experience as illustrated in Figure 11. Since there is no Rayleigh scattering in HCF, the entire section of fiber of 10km shows as if it is ‘cut’ on the OTDR trace. The only indication that the fiber is fine is based on the reflective peaks on the trace at the ends of the fiber due to the SCF-HCF adapter cables. The next 20km of fiber is unremarkable given that it is a standard SMF fiber with nominal loss of 0.25dB/km and its associated Rayleigh scattering. The last 10km of HCF fiber is similarly to the first 10km appears ‘cut’ save for the reflective peak at the end of the total 40km of cable. One other point to remember here is that the HCF although 10km long, appears to be just 6.8km! This is because the refractive index in the OTDR is set to that of 1.4684, if on the other hand it were set to 1.003, which is the refractive index of air, the result would have been that the HCF would show 10km, but the SMF would then have showed 14.7km.




Figure 12 – Illustrating various direct detect and coherent spectra and the passives passband

In a different paper at this conference, we have discussed Comcast plan to converge various wavelengths onto a single optical fiber in the access plant [7]. The picture above shows that the optical filters used in the access plants can handle all wavelengths up to 400Gbps a fact also tested in the system here.

A unique feature of our system is that it is bi-directional, where a single fiber was used to interconnect the two end points. The main reason for the BiDi approach is the need to understand how such systems perform using HCF with the possibility of preserving fiber, HCF or SMF. The system described in the figure below describes the BiDi system where coherent links are combined with direct detect system of 10G, 25G and 40G wavelengths to deliver different types of services simultaneously (DAAS, BERT testing, etc.) in the same string of fiber. We believe that is the first implementation of such system in the world. The optical spectrum analyzer capture depicted in Figure 13 shows the actual spectrum utilization for this trial.




Figure 13 – Bi-Directional Optical Wavelengths over the Comcast HCF Link

Towards that end, as the spectrum graph above shows, we tested bi-directional transmission on a single fiber comprising

- Coherent 400Gbps using 16QAM 63GB system
- Coherent 100Gbps using QPSK 32GB system
- Direct detect 40Gbps NRZ
- Direct detect 25Gbps NRZ to simulate CPRI data
- Direct detect 10Gbps NRZ

And were successful in closing all the links over 20km emphasizing the viability of Comcast vision over multiple fiber types.

6.1. Establishing Latency Improvement using a Hybrid Loop Test

To verify the latency improvements in our system we devised a test set-up using standard 10G test equipment. The approach consists in cascading the 10G traffic from a traffic generator back and forth using back-to-back the 10G optics connected to the CFP2 optics through an aggregation switch. The diagram below shows how the 10G traffic is propagated back and forth through the CFP2 to produce the 100G traffic. In our implementation we actually used a loop test with the same 10G traffic to go back and forth over 20km of fiber 40 times with the 80 10G SFP+ devices aggregated to 8 of the 100G CFP2 devices. Overall, in this hybrid fiber/electronic loop test [8], we actually propagated the 10G signal through

$2 \times 20 \times 40 = 1600\text{km}$. We then repeated this test with standard single mode fiber as well and used the test equipment to track packet errors, latency and jitter.

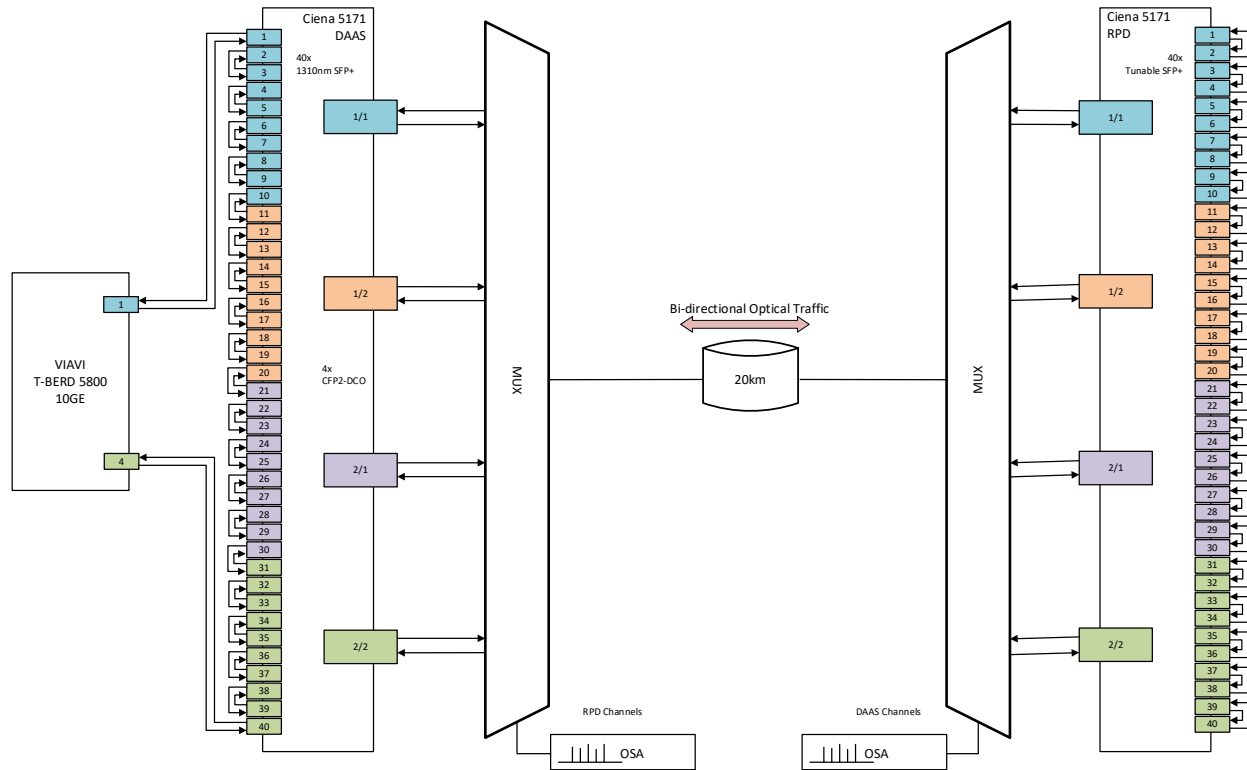


Figure 14 – The Hybrid Loop Test Illustrating Latency Reduction

Presented below is the test data collected and summarized. In it is seen that the latency on HCF is much lower than that on standard fiber. The 8972us delay on HCF compares to 11,330us delay in the SCF-SMF case. We also measured the latency of a link combining 20km of HFC and 10km of SCF-SMF, which led to 17,352us delay.




Figure 15 – Illustrating Latency improvement with HCF (top) over SCF-SMF (bottom)

6.2. Spider Diagram Grand Summary

Below in Figure 16, we present a Spider Diagram showing the key comparison points to summarize the main differences between the HCF and the SMF. The assumption here is to use the relative performance between the two fiber types. The key parameters used for the diagram vertices are purpose were latency, insertion loss, cross-talk (linear and non-linear), cost, and power consumption (green factor).




Figure 16 – Qualitative Comparison of Fiber Attributes

In addition to describing our test system and initial results, we also analyze some of the benefits in addition to latency reduction that are related to reduced optical non-linearities and other effects and the wider spectrum and the consequent Capacity enhancements possible. Practical deployment strategy of fibers is essential at Comcast, the ability to use a portfolio of fibers for some of the emerging low latency market while also driving fiber technology deeper into the network is of the essence of our approach towards the industry initiative of 10G.

7. Conclusion

In this paper we discuss the role of HCF in delivering different types of services over the access networks. We reviewed the history of HCF fibers and discussed some of the main attributes including a combination of minimum latency and non-linear impairments that help drive capacity and reach with continual improvements in HCF technology. We pointed out features that make HCF attractive not only in long haul networks, but also in access networks. While HCF can be a key tool for Comcast to provide low latency services in a cost-effective way. Its combination with traditional SMF provides the great flexibility to our designers to accomplish these goals.

Of particular note in this paper is our report for the first time in our knowledge of successful bi-directional transmission of optical signals across a hybrid HCF/SCF plant with a converging of multiple direct detect and coherent signals ranging from 10Gbps to 400Gbps on a single HCF fiber plant.

Abbreviations

AP	Access Point
Bps	Bits per second
BERT	Bit Error Rate Test
BiDi	Bi-Directional
C Band	Conventional Band (1530 to 1565 nm)
CL	Confinement Loss
CFP2	C Form-factor Pluggable 2 (100Gbps)
DAAS	Distributed Access Architecture Switch
DCM	Dispersion Compensation Module
DNANF	Doubled Nested Anti Resonant Nodeless Fiber
DSP	Digital Signal Processing
eCPRI	Enhanced Common Public Radio Interface
EDFA	Erbium-Doped Fiber Amplifier
FPGA	Field-Programmable Gate Array
HCF	Hollow core fiber
HFT	High Frequency Trading
Hz	Hertz
IoT	Internet of Things
ISP	Inside Plant
K	Kelvin
LEO	Low Earth Orbit
NANF	Nested Anti Resonant Nodeless Fiber
NRZ	Non-Return-to-Zero
O Band	Original Band (1260 to 1360 nm)
OFC	Optical Fiber Conference
OSP	Outside Plant
OTDR	Optical Time Domain Reflectometry
PBG	Photonic Band Gap
QAM	Quadrature Amplitude Modulation
QSFP	Quad Small Form-factor Pluggable transceiver (40, 100Gbps)
QPSK	Quadrature Phase Shift Keying
SC-APC	Standard Connector - Angled Physical Contact
SCTE	Society of Cable Telecommunications Engineers
SCF	Solid Core Fiber
SMF	Single Mode Fiber
SFP28	Small Form-Factor Pluggable 28 (25/28 Gbps)
SSL	Surface Scattering Loss

Bibliography & References

- [1] Hollow core optical fibres with comparable attenuation to silica fibres between 600 and 1100nm, Hesham Sakr et al., Nature Communications, 2020, <https://www.nature.com/articles/s41467-020-19910-7>
- [2] Hollow Core Fibers: Key Properties, Technology Status and Telecommunication Opportunities, David Richardson, ORC, University of Southampton and Lumenisity Ltd.
- [3] Figure used with Permission from Lumenisity, <https://lumenisity.com/>
- [4] Figure used with Permission from OFS, <https://www.ofsoptics.com/>
- [5] 0.174 dB/km Hollow Core Double Nested Antiresonant Nodeless Fiber (DNANF), OFC 2022, Technical Digest Series, paper Th4C.7
- [6] Comcast Underground: Innovative Fiber Deployments Over Existing Underground Critical Infrastructure, Venk Mutalik, Pat Wike, Doug Combs, Alan Gardiner, Dan Rice
- [7] Photon Avatars in the Comcast Cosmos: An End-to-End View of Comcast Core, Metro and Access Networks, Venk Mutalik, Steve Rupp, Fred Bartholf, Bob Gaydos, Steve Surdam, Amarildo Vieira, Dan Rice
- [8] First Demonstration of Field-Deployable Low Latency Hollow-core Cable Capable of Supporting >1000km, 400Gb/s WDM Transmission, A, Saljoghei et al, OFC, 2021.

Are Your Critical Facilities Ready To Be Managed With Big Data?

A Technical Paper prepared for SCTE by

Dustin Boyette

Product Manager, Intelligent Distribution

EnerSys

3767 Alpha Way, Bellingham, WA, USA

+1 360 603 0672

dustin.boyette@enersys.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. What is Big Data?.....	3
2.1. The Three Vs.....	3
2.2. A Dollar a Month	3
2.3. Of Warehouses and Lakes.....	4
3. Data Sources in Critical Facilities.....	5
3.1. Power	5
3.2. Environmental	5
3.3. Security/Access.....	6
4. Turning Data Into Information	6
4.1. First, You Need the Data.....	6
4.1.1. The Internet of Things (IoT)	6
4.1.2. “Smart” Infrastructure	7
4.1.3. Retrofitting Telemetry.....	7
4.2. Hands Off! Reducing the Human Burden of Data Management & Analysis.....	8
4.2.1. Modeling.....	8
4.2.2. Data Visualization	9
4.2.3. AI.....	10
5. Use Cases in Communications Facilities	10
5.1. Outage Prevention	10
5.2. Operational Efficiency	12
5.3. Capacity Planning	12
5.4. Automatic Provisioning.....	13
6. Conclusion.....	13
Abbreviations	14
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 - Historical Cost of Data Storage	4
Figure 2 - “Smart” DC Fuse Panel With Individual Branch Circuit Current Monitoring.....	5
Figure 3 - IoT Example Graphic.....	7
Figure 4 - Traditional Dry Contact Telemetry.....	8
Figure 5 - Process Temperature Overlay in a Factory.....	9
Figure 6 - Outage Postmortem	11
Figure 7 - Old-Fashioned Planning?	13

1. Introduction

Today, critical communications facilities are largely managed manually. While some facilities (headends, primaries, data centers) may have fairly extensive monitoring, the data provided still tends to end up on a 2D “dashboard” for humans to interpret and act upon as necessary. Some facilities—especially remote hubs—may have very little telemetry at all.

This paper introduces big data then presents it as key to enabling advanced management of critical facilities (CF) in the future. The importance of collecting broad and deep data—starting today—is discussed along with extant and future sources of this data from facility infrastructure. Common challenges of contemporary facility management will also be discussed in the context of how big data can help reduce the human burden posed by those challenges. Several example use cases enabled by big data—when combined with modeling, visualization, and artificial intelligence—will be presented.

2. What is Big Data?

As is typical for contemporary jargon, a concrete definition for big data is elusive; you’ll find slightly differing definitions depending on what organization is using the term. In the simplest sense, big data refers to data sets too complex and/or large for traditional data processing methods to take advantage of.

An analogy might be the spiral-bound paper notebooks of every student in a school district (replete with chemistry lab notes, doodles, and phone numbers of the student at the adjacent desk). In contrast would be the perfectly alphabetized contacts list in your smartphone.

2.1. The Three Vs

While there are many characteristics of big data, it is chiefly defined by “The Three Vs”:

Volume - Lots of unstructured data

Velocity - Data delivered rapidly, real-time or near real-time with little or no pre-processing

Variety - Many kinds of data available that don’t readily fit traditional, ordered database schemes

More characteristics (also conveniently starting with the letter V) have augmented the definition over time, but for the purposes of this paper the original three will suffice.

Additionally, the term “broad” can be synonymous with Variety here. The term “deep” can be synonymous with Volume over time; that is, historical data.

2.2. A Dollar a Month

The primary driver enabling big data is the ever-decreasing cost of mass storage. In the early days of computers, five megabytes of storage occupied the space of a washing machine and cost perhaps quarter of a million dollars. The idea of storing even a byte of data with unknown importance—data which may never be accessed again—was unfathomable.

Today, however, vast amounts of both mass storage and compute resources are available at trivial cost relative to the potential benefit. For example, one terabyte of infrequent-access storage through a commercial cloud storage service such as Amazon S3 is a little over ten dollars a month. And at the

lowest-cost “Deep Archive” tier, a terabyte is less than one dollar a month and becoming cheaper every year.




Figure 1 - Historical Cost of Data Storage

2.3. Of Warehouses and Lakes

In order to take advantage of any data sets, there first has to be place to store the data. The storage method must be reliable, extensible, and allow for quick access to subsets of data of interest.

An evolution of traditional data processing concepts, a data warehouse stores well-structured data. The schema for the data is predefined and data is transformed to fit. A customer contact database is a simple example of structured data: first name, last name, city, country, etc.

Moving this into the Critical Facility (CF) management area with which we are concerned, consider a database of facility equipment. Expected fields would be manufacturer, model, serial number, install date, rack location, and so on.

But what if data is available that doesn't fit the structure?

Power Input Module B Inlet Air Temperature 56°C at 03:02:01 on 2022-07-03

Where does that go?

It goes into a data lake, a vastly scalable centralized repository for raw data of virtually any type be it structured or not.

A data lake takes a different approach to data storage that is better suited to massive amounts of diverse data available today and in the future. The idea behind a data lake is to store data in its elemental, uninterpreted state. The data is minimally processed before storage. A value (the data type is less important than with structured data—integers, floating point, strings, even image data are welcome in the pool) along with some context—a source identifier and a timestamp—are all that are needed.

3. Data Sources in Critical Facilities

Even in a critical facility built without an aim towards monitoring/telemetry, there are many sources available for useful data. We are primarily concerned with data from the facility infrastructure itself and will expand on that below. However, it's important to note that to fully take advantage of big data even the management channel data from the revenue-generating equipment itself is important to gather.

3.1. Power

Virtually every piece of power equipment you will find in a communications facility—from an Automatic Transfer Switch (ATS) to a Surge Protective Device to a DC Plant—is available today with a means to provide data in real time or near real time to a monitoring system via a network.

Early network-enabled power systems provided only coarse data—bus voltage, total current, alarm status, and other data at that level. Today, power systems are available with the ability to measure consumption down to the smallest branch circuit.



Figure 2 - “Smart” DC Fuse Panel With Individual Branch Circuit Current Monitoring

Stationary batteries, a key component of high-reliability power systems, have traditionally offered little in the way of data. Monitoring systems are often seen in facilities, but they typically only offer jar-to-jar voltages and perhaps a few temperature measurement points. However, the intelligence mandated by lithium batteries in consumer and automotive applications is making its way into reserve power batteries. Imagine a large critical facility with dozens of strings of backup batteries, each cell able to provide detailed metrics on its status. There could easily be ten thousand pieces of useful data gathered from a facility's batteries per minute.

3.2. Environmental

Many facilities have environmental monitoring capability either as separate systems or as part of the active Heating/Ventilation/Cooling (HVAC) systems. Computer Room Air Conditioners (CRAC) have been available with network connectivity for some time.

Again, like many power systems, the data from these HVAC systems is usually basic and local to the unit itself. Perhaps there would only be a dozen temperature readings directly from the HVAC systems in a facility. However, the number of temperature and airflow sensors present on all the equipment in a facility is likely in the hundreds if not thousands. The low cost of digital temperature sensors means that granular temperature readings should be available for intelligent thermal analysis of an equipment rack, row, or zone.

In addition, new low-cost infrared (IR) sensors could soon enable equipment to detect local hot spots without the need for quarterly site visit by a technician with an expensive IR camera (the data from which is quickly stale).

3.3. Security/Access

Today's advanced security systems provide data far beyond simple contact closures triggered by magnets on doors. Key card and Radio Frequency Identification (RFID) integration allows the tracking of not only access by personnel, but the presence of tools and even equipment assets. Visual and IR cameras add another.

Looking beyond the dedicated security and access control systems, however, many pieces of equipment have sensors that can detect when a physical change has been made: module insertion/removal, access door open, rack door open, etc.

4. Turning Data Into Information

In traditional CF telemetry, there is priority given to quickly turn elemental data into important (or seemingly important) information. Formulae and scripts are written to take near-term data and either present it to an operator as information, or cause some immediate action: put a rectifier module to sleep, open a low-voltage disconnect, turn on an economizer (or trigger the fire alarm!) as examples. All important things at the time, to be sure.

However, often the raw data is discarded after being processed into immediate information, and is no longer available for analysis months or even years later.

4.1. First, You Need the Data

Obviously to take advantage of big data you first need data. What may not be obvious is the importance of rich historical data, and the importance of capturing that data as soon as possible.

4.1.1. The Internet of Things (IoT)

The Internet of Things, or more commonly just IoT, is a popular buzzword covering all manner of intelligent, network-connected widgets that will magically and seamlessly talk with each other and work together. Any of us who've wasted an afternoon trying to get a Hi-Fi music streamer to stay connected to wireless speakers can attest, it's not quite there yet.

However, the idea of every *thing* being connected is fascinating: from the expected like discrete current transducers and humidity sensors, to the mundane such as light switches and fixtures. If a light fixture happens to also know the temperature, why should it not report that data to the building network?




Figure 3 - IoT Example Graphic

4.1.2. “Smart” Infrastructure

The decreasing cost of sensors, microcontrollers, and ubiquitous network connectivity has made it possible for almost every piece of building infrastructure to be some degree of “smart.” Today, even doorbells have microcontrollers and wireless network interfaces!

In critical facilities, however, most of the infrastructure today is still fairly dumb. That is, it is not able to provide data via a network connection. Or if the equipment is able, it is not connected, or nothing is done with the data. Indeed, some operators do not see the benefit of smart infrastructure, but they are making that choice from the perspective of the tools available to process the data today.

To take advantage of big data in the future, smart infrastructure is necessary today.

4.1.3. Retrofitting Telemetry

Operators large and small are faced with a chicken-and-egg question when it comes to smart infrastructure: “Why invest in network-enabled fuse panels and air conditioners when the majority of my existing equipment lacks the capability to tie in?”

Some may prefer the “sidewalk” solution: municipalities often require sidewalks be built whenever new construction is undertaken even if the surrounding neighborhood lacks sidewalks. While this results in a few “where the sidewalk ends” jokes on shorter timescales, longer term it results in neighborhoods with nearly contiguous walkways.

Unlike the relatively static central offices of landline heyday, today's critical facilities change, evolve, and grow quickly. When old HVAC systems and fuse panels won't meet the needs of the latest revenue-generating network gear, the new systems should be smart.

For installation conditions that don't readily permit the migration of old infrastructure to smart infrastructure, intelligence can often be retrofitted. Temperature sensors and clamp-on current sensors are readily available to augment smart infrastructure and instrument critical facilities.




Figure 4 - Traditional Dry Contact Telemetry

4.2. Hands Off! Reducing the Human Burden of Data Management & Analysis

The aim of big data in most industries is to *reduce* the amount of manpower required to evaluate, plan, and act to accomplish a business objective. Unfortunately, as more data is made available, operators often find it more difficult to digest and act upon.

4.2.1. Modeling

The first step in making big data useful for managing critical facilities is to move ahead from the "Mission Control" dashboard. Facilities can be surveyed and physically modeled in 3D and presented virtually. Then, with sufficient data sets, the electrical and thermal behaviors of the facility can be modeled as well.

Instead of sorting through out-of-date photos, looking at rack inventory spreadsheets, or sending crews to survey sites, it's conceivable to view a current *model* of the facility. Operators will be able to virtually walk down a rack row and see what equipment is installed, where floor space is available, and even where fiber trough and cable tray is ready for growth.

4.2.2. Data Visualization

Layered on top of modeling, big data allows the ability to overlay important data directly over the virtual facility. Today, most of us use similar capabilities in everyday life without realizing it: the satellite view overlay of a map application, per-zone temperature overlay on a smart home, and so forth.

Now imagine being able to view the inlet air temperature gradient across the entire front of a critical equipment rack. Or go back in time and see what that same gradient looked like when an air conditioner failed last summer.




Figure 5 - Process Temperature Overlay in a Factory

4.2.3. AI

Despite being misunderstood and maligned in countless science fiction books and screenplays, artificial intelligence is at work today helping businesses and improving the lives of people.

What AI can already do today is impressive—from creating striking images prompted by just a few words (DALL-E) to finding the most efficient paths for rovers to explore other planets. But far more exciting is what AI will be able to do five and ten years from now.

In the CF space, it is readily conceivable that AI will allow us to “see the future”. In our example above, instead of viewing the rack inlet temperatures from last summer, what if we could view the rack inlet temperatures for next summer? After adding another 10kW of revenue-generating equipment in the row! AI makes that possible and can do so better than the current “snapshot” predicting through the use of massive amounts of historical data.

5. Use Cases in Communications Facilities

With a sufficiently broad and deep set of data, what can big data do for critical facilities?

5.1. Outage Prevention

Operators spend significant time and money to ensure their networks are robust, resilient, and reliable. Redundant generators, dual DC plants, and diverse wiring are all aimed at preventing customer-affecting outages. Yet outages still occur all too frequently.

In many post-mortem investigations, sifting through logs reveals that data was available that indicated an outage was possible if not imminent. Often this data was not given enough importance to make it into any alarm logic or status dashboard. Ultimately, it’s often a human who would have to react to the condition to prevent the outage. Or, if some automatic action is taken it’s not granular enough and results in mitigation instead of prevention, which still results in something less than normal delivery of service.

Instead, why not allow computers to search for patterns in both current and historical data, run myriad “what if?” scenarios against the facility model, and identify where trouble may be brewing. Days, weeks, or months beforehand.



Figure 6 - Outage Postmortem

5.2. Operational Efficiency

Most operators would likely admit that their critical facilities do not operate at peak possible efficiency. Priority is understandably given to ensuring equipment is working, service is reliable, and just keeping up with customer demand. Analysis of power conversion loss, HVAC efficiency, and such is performed as time allows—usually when a major facility upgrade is planned. Such analysis is usually done on “snapshot” or coarse-grained historical data, and while the result may improve efficiency somewhat, it cannot continually *optimize* efficiency.

Big Data and advanced tools operating upon it, combined with smart infrastructure, will enable facilities to self-optimize in near real-time. Consider load-shedding, power source switching (utility to photovoltaic solar (PV), for example), economizer activation, etc.—all driven by machine learning working off broad, deep data sets.

5.3. Capacity Planning

As the worldwide demand for high-speed data grows unabated, the need to expand critical communications facilities keeps step. In these buildings exist leading-edge routers and switches and optical equipment yet planning for future growth remains largely a manual process. Obsolete equipment inventories, stale usage data, and myriad spreadsheets are used to determine whether new equipment can be added to a facility to meet growing service needs. Consider just a fraction of the simple questions that must be answered in order to add a significant new piece of equipment:

Physical: *Is there sufficient rack space available?*

Power: *Is there sufficient ampacity available at the BDCBB? At the DC plant? The service entrance? Is there even a feeder fuse or breaker position available?*

Cooling: *Can the proposed zone support the additional thermal load?*

Big data could answer these immediate questions, and much bigger ones without an operator ever needing to open a spreadsheet:

When will the utility service no longer support the rate of growth at this facility?

What would be the first point of failure during a prolonged utility outage?

Would it be cost-effective to add alternate energy sources to this facility?




Figure 7 - Old-Fashioned Planning?

5.4. Automatic Provisioning

Taking capacity planning even further, a comprehensive CF management system could analyze customer trends, predict required service growth, then cascade the infrastructure requirements to match those needs. Imagine a system that places a purchase order for a new distribution panel along with all the properly-sized cables, circuit breakers, and mounting accessories. The order is placed taking into account delivery time of the various necessary components, and an installation technician can even be scheduled as soon as the equipment is slated to arrive.

6. Conclusion

The question posed by the title of this paper is rhetorical; your critical facilities are not yet ready to be managed by big data today, nor are the advanced modeling and AI applications yet developed enough in this field to do so. However, many facilities already have a surprisingly rich set of data available that could be valuable in future decision-making. Now is the time to collect the data already available and augment infrastructure where that data is missing. The future of critical facility management will be driven by big data.

Abbreviations

CF	critical facility
ATS	Automatic transfer switch
HVAC	Heating, ventilation, air conditioning
CRAC	Computer room air conditioner
IR	infrared
RFID	Radio frequency identification
BDCBB	Battery distribution circuit breaker bay
IoT	Internet of Things
PV	photovoltaic

Bibliography & References

Internet of Things (https://www.sas.com/en_us/insights/big-data/internet-of-things.html); SAS Institute

A Visualization is Worth a Thousand Tables: How IBM Business Analytics Lets Users See Big Data; Zif Davis 2014

Artificial Intelligence in Real-Time Video Encoding from Theoretical Promises to Operational Gains

A Technical Paper prepared for SCTE by

Jan De Cock
Director Codec Development
Synamedia
Luipaardstraat 12, 8500 Kortrijk, Belgium
+32 467 093721
jdecock@synamedia.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Applying ML to video compression: from ML to TinyML	4
3. Video encoding complexity (reduction)	4
3.1. A high-level view on encoding.....	4
3.2. Encoder complexity reduction	5
3.3. Reducing the complexity of ML inference networks.....	7
4. Rate control.....	8
5. Subjective improvements	9
6. Video quality measurement.....	11
6.1. From offline to real-time VQ measurement: tracking video quality	12
6.2. From rate control to quality control	13
7. Video quality monitoring.....	15
8. Conclusions.....	17
Abbreviations	18
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1. High-level overview of a hybrid block-based video encoder	5
Figure 2. Evolution of video coding standards.....	6
Figure 3. Example partitioning structure using VVC.	6
Figure 4. Simplified view on traditional rate-controlled encoding.....	8
Figure 5. ML-based rate control.....	8
Figure 6. Improvement in rate control prediction accuracy between traditional (left) and ML-based (right) rate control.	9
Figure 7. Example quality improvement between traditional (left) and ML-based (right) texture preservation.....	10
Figure 8. ML-based logo detection	10
Figure 9. Full-reference video quality assessment	11
Figure 10. Calculating VQ inside the encoder	12
Figure 11. ML-based VQ measurement inside the encoder, based on pre-analysis features.	13
Figure 12. Example networks used for ML-based VQ prediction.	13
Figure 13. Above a certain bitrate, adding more bits will no longer (or hardly) improve quality	14
Figure 14. Quality-controlled compression	14
Figure 15. Result of quality-controlled compression (target VMAF=90).	15
Figure 16. FR and NR quality measurement for VQ monitoring.	16
Figure 17. VQ monitoring at different points in the video delivery chain.	17

1. Introduction

Many books and articles have been written about artificial intelligence (AI) and machine learning (ML), in a variety of applications. ML is far from new, has an established theoretical foundation, and lots of different types of ML techniques have been introduced over the past decades. These techniques can be classified in different ways, but a full taxonomy is outside of the scope of this article. In this paper, we focus mostly on ML algorithms, as a subset of AI. Excellent introductions and overviews have been provided in e.g. [Goodfellow16, Bishop95].

Lots of successes have been claimed based on ML, and reports of AI intelligence are already the subject of ethical discussions [Google22]. Still, the powers of machine learning are not always a solution, and in many applications, even though they make for an interesting marketing statement, they do not lead to net gains or operational savings.

Machine learning has powerful applications in computer vision, image and video processing, and approaches using deep neural networks have become the center of academic and industry research. For example, residual neural networks have shown impressive results for image classification and recognition [Simonyan14, He16]. Still in most of these cases, very complex algorithms are needed, requiring e.g. deep neural networks containing dozens or hundreds of layers. While it's acceptable to have a very complex *training* stage (which needs to be executed once), it's primarily the complexity of the *inference* network (which needs to be repeated many times) that determines the feasibility of ML approaches¹. An important unit of expressing the complexity of ML inference networks is the number of *multiply-accumulate operations* (MACs). Some of the best-performing image recognition networks use millions of MACs per image.

Often, new approaches are deemed feasible when they can be run on state-of-the-art GPUs inside a server. In certain cases, this is acceptable, and the cost of a dedicated CPU or GPU is warranted. For real-time, cost-sensitive applications, however, this is not an option. In typical video encoding/transcoding set-ups, dozens or even hundreds of channels need to be processed on a single server, and the cost per channel is a crucial criterion. Furthermore, the latency of offloading decisions to accelerators (if they would be cost effective, which is not the case), would be prohibitive.

In this paper, we discuss the applicability of machine learning approaches in different areas of *real-time* video compression. We successively cover encoder complexity reduction, rate control, video quality improvements and video quality measurement. In each of these areas, we have studied ways to reduce the complexity of ML inference, to end up with algorithms that are applicable in real-time, cost-sensitive applications.

¹ While in this paper we're mostly concerned with the complexity of the inference stage (which needs to be run every time a decision is made), the cost of *training* these networks can still become prohibitive in some cases. Not only the impact on computation, but also on emissions needs to be considered. In the field of natural language processing, Transformer Networks such as GPT-3 have been developed, with 175 billion ML parameters, requiring tons of CO2e just for training.

2. Applying ML to video compression: from ML to TinyML

In contrast to what our marketing departments would like us to believe, applying ML to any problem is not a trivial task, and it doesn't work out-of-the-box. For any problem solved using ML, *domain knowledge* is required. In the case of video encoding, a lot of specialized knowledge about video compression and its components is essential.

When attempting to combine the difficulties of domain knowledge, ML knowledge and complexity aspects, it is easy to get stuck in the “trough of disillusionment”, where little or no net benefits of ML are reaped. In the end, product-grade encoders are usually mature, with smart human-designed algorithms. And indeed, in several areas, the gains are not immediately spectacular. Still, by pushing through, it is possible to obtain productivity and operational gains.

It is also tempting to assume that we can apply deep CNNs or similar computer-vision inspired techniques to video compression: detecting objects, humans, optical flow, or even apply a semantic meaning to each of these objects – after which we use that semantic information to help compression. While this is possible in theory, and a typical human reflex, such approaches are typically error-prone, inconsistent over time, and require a tremendous amount of computational resources.

In this paper, we focus on ML approaches that are feasible in real-time, with minimal impact on “channel density”, i.e., the number of video channels that can be processed on a single server. As a result, these are realistic techniques that lead to operational savings and efficiency increases.

As an analogy, the challenges encountered in ML-based video compression are similar to those in the research field of “Tiny ML” [Warden20] – even though we're working on servers with multi-core CPUs. But instead of running single tasks on a very low-power platform (in the mW range), we process hundreds of video streams on a single CPU, making millions of decisions per second. To continue the analogy with TinyML, every individual building block inside an encoder has only milliwatts of power available. This results in an exciting new combination of research on *low-complexity real-time ML inference for video compression*.

3. Video encoding complexity (reduction)

3.1. A high-level view on encoding

In this paper, we continue to focus on video encoding as application. Compressing and encoding video is an extremely complex process, comprising different steps including pre-analysis, mode decision, motion estimation, interpolation, intra/inter prediction, transform, quantization, entropy coding, in-loop deblocking etc. Executing each block is a time-consuming operation, but it's mostly search space exploration (i.e. evaluating the different encoding options such as partitions and motion vectors) that is expensive in encoders. A simplified version of a hybrid block-based encoder is shown in Figure 1.

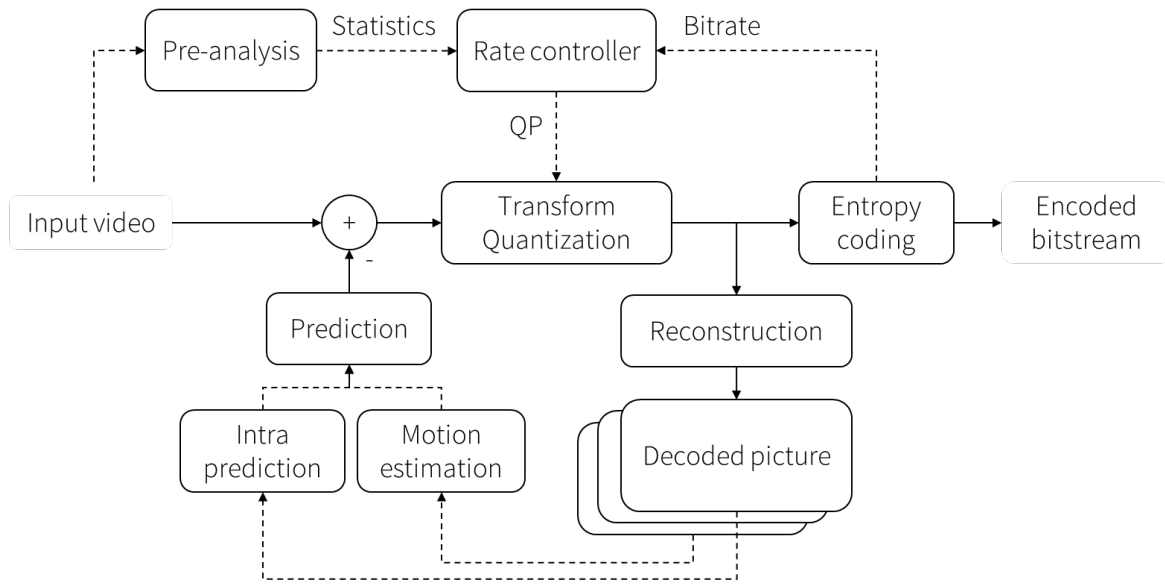


Figure 1. High-level overview of a hybrid block-based video encoder

Video encoding not only demands a lot of CPU power, but also memory (bandwidth). With higher resolutions (e.g. 8K) and higher frame rates (e.g. 120 fps), billions of pixels need to be processed every second. For lower-resolution streams, typically dozens to even hundreds of streams can be processed in parallel on a single server.

To cope with this complexity, lots of effort has been spent on developing hardware accelerators, FPGAs, ASICs and so forth. Still, encoding in software has many benefits, and brings maximum flexibility in deployment (on-prem, cloud-based etc), upgrades, and video quality improvements. For operational efficiency and flexibility, software encoding is often preferred, and will be the focus in this paper.

3.2. Encoder complexity reduction

To gain a competitive advantage, a low cost per channel is important. This is becoming increasingly challenging, as the complexity keeps increasing for newer compression formats. Each generation of video compression standards brings an increase in computational complexity (e.g. from MPEG-2 to AVC to HEVC), and newer standards are on the horizon (such as VVC). Typically, *decoder* complexity doubles with every generation, while jumps in *encoder* complexity can be even larger.

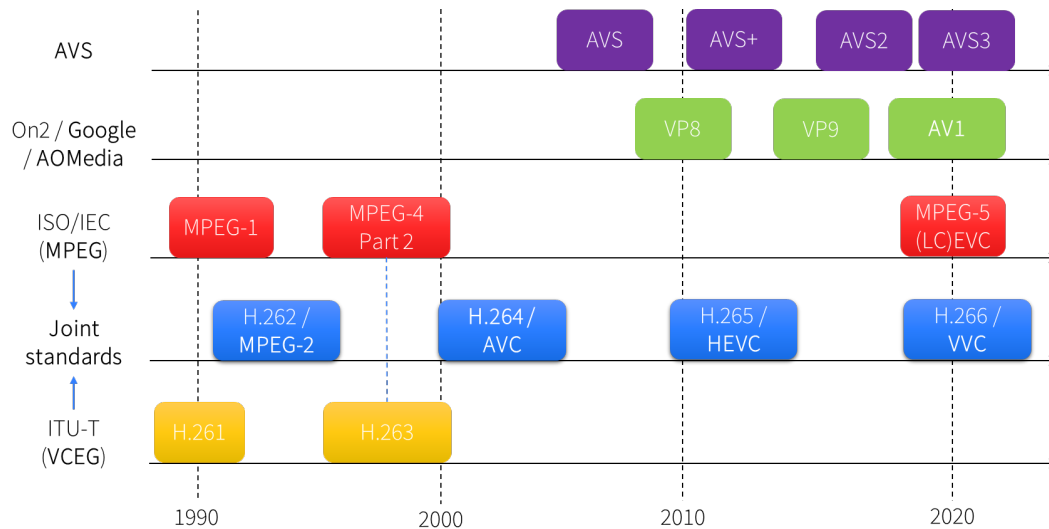


Figure 2. Evolution of video coding standards

Not only the resolutions are increasing, also the number of different ways to encode each individual frame or block is going up. While HEVC had about 80K different ways to partition a 64x64 block, the biggest coding units in VVC (128x128) can be partitioned in more ways than there are atoms in the universe².

128x128 coding unit in VVC




Figure 3. Example partitioning structure using VVC.

Given the potential for optimizations in this huge search space, finding efficient ways to make decisions inside the encoder is a popular research topic, and recent attention has shifted towards machine learning approaches. Plenty of references in this direction can be found in academic literature. Typically, however, the reference point in these papers is (extremely slow) reference software, and most of these gains cannot be transferred as such when applied to professional, real-time encoders.

Also, many publications on encoder complexity reduction focus on non-real-time (VOD-type) encoders. While this is useful as a starting point, those findings cannot be directly translated to *real-time* scenarios. Smart ‘production-grade’ encoders already make intelligent decisions, without exploring all options. In practical encoders, the promised complexity reductions are typically (way) lower.

² Every 128x128 coding unit can be split down to 4x4 coding units, with a recursive combination of binary, ternary or quaternary splits (or no split). I will leave the calculations up to the interested reader.

Still, interesting work in different directions has been performed resulting in ML networks with reasonable complexity, always while trying to limit the loss in compression efficiency:

- *Partitioning for HEVC and VVC.* [Liu16] presented a CNN-based CU partition size decision for HEVC with a reasonable complexity of 3,000 MACs, along with a hardware implementation. Among others, [Bhat21], [Wu21] and [Liu22] introduced CNN-based or SVM-based strategies for acceleration of VVC encoding, with a focus on the partitioning decisions, with encoding time reduction of roughly 30-80%, while limiting the impact on compression efficiency.
- *Intra prediction.* The work by [Santamaria20] presents NN-based intra prediction modes along with simplifications that lead to multiplications in the order of 100s up to 10,000s for 16x16 blocks. This builds on the work of [Pfaff18], and an interpretation analysis is run to come to simpler, explainable predictors that are easy to implement. The result is NN-based modes that are much closer to real-life usage.
- *Inter Prediction.* Although much of the recent work has focused on NNs, other ML techniques such as Decision Trees prove to be efficient ways to optimize encoder decisions, as in [Kim19], where inter prediction is accelerated for AV1.
- *Transform selection.* The transform search for AV1 is accelerated in [Su19], based on a neural network with one hidden layer. For transform kernel prediction, two shallow networks are used which are combined into a score for the 2D transform.

The message from these papers, along with our findings, are that fairly simple and shallow neural networks can produce accurate results, and at acceptable computational complexity.

3.3. Reducing the complexity of ML inference networks

Optimization techniques can help push the boundaries of what ML can achieve inside an encoder, and can help limit the cost of deeper networks. For example, *pruning* can be used to reduce the complexity of the networks, and to eliminate redundant MACs in the inference networks. Also, *quantization* allows to reduce the bit depth of operations, at the possible cost of some accuracy in calculations.

Specialized hardware *accelerators* can be tolerated in some applications, but might lead to a large overhead in latency, limiting their feasibility. Dedicated *instruction sets* (such as VNNI) provide a more convenient way to parallelize operations, and specialized matrix multiplication instructions are made available on the most recent CPU generations. Unfortunately, offloading ML inference to external accelerators or GPUs is usually not preferred, and would lead to unacceptable latency in the real-time applications we're discussing in this paper.

In all cases, a trade-off between accuracy and complexity needs to be found – again, domain knowledge is needed to find a good balance. That domain knowledge is also essential for the most important part, which is intelligent network design: shallow network and intelligent feature design.

4. Rate control

Rate control is another area where ML has proven to provide improvements. Video encoding can be considered as a resource allocation problem. Given a certain bit rate (=bit budget), the available bits need to be distributed in the best possible way across different frames and coding units, to reach the highest possible video quality. This is done by choosing the right quantizer (quantization parameter) for each block.

While this seems a fairly trivial task, it is actually an extremely difficult problem, given the multitude of options that every individual block can be encoded with. Furthermore, due to prediction, blocks are dependent on previously encoded blocks, further exploding the complexity of the problem. The power of its rate control algorithm is actually one of the biggest differentiators in the quality of an encoder.

In practical encoders, estimations are made to allocate quantizers to every block based on pre-analysis of the video content. Each block will be encoded with an estimated quantization parameter (QP), and the total bit rate for the frame needs to approximate the given bit budget as closely as possible. This pre-analysis and prediction stage is essential, and the prediction error needs to be as small as possible. The rate control process is depicted in Figure 4.




Figure 4. Simplified view on traditional rate-controlled encoding

Traditionally, prediction algorithms are based on human-designed heuristics, tweaking and testing. While this works well in general, there are cases where misprediction leads to fairly large bit estimation errors. This can lead to the rate controller over- or under-allocating bits for a number of frames. As a result, bits might be wasted, or quality might suddenly drop.

With machine learning, smart features can be calculated during the pre-analysis stage. With the resulting ML-based rate controller (Figure 5), we've noticed better resilience to a variety of content types, including sudden content or scene changes. As a result, we achieve better correlation between estimated and encoded bits, as illustrated in Figure 6.




Figure 5. ML-based rate control




Figure 6. Improvement in rate control prediction accuracy between traditional (left) and ML-based (right) rate control.

5. Subjective improvements

Reaching a high level of video quality is extremely important when offering premium video content. Over the past decades, major steps forward have been made, going from analog TV to HD digital TV, and now it is common to watch premium sports in UHD, with the first 8K channels becoming available. In any case, we’ve come a long way since analog or DVD-level quality, and viewers are getting used to high-quality and ultra-high-definition video, by watching popular VOD services.

Pushing the limits in video quality is important to offer a premium to viewers, and as a differentiator when selling video encoding services. When comparing encoders (in so-called shoot-outs), offering the best quality is one of the most important criteria. In the next section, we will discuss VQ measurement based on *objective* metrics (as calculated by algorithms). But in this section, the focus is on *subjective* quality, as perceived by the viewer.

Several elements are important to optimize the visual quality as perceived by viewer. In early videos circulating on the Internet, digital video was suffering heavily from blocking artifacts, blurring, mosquito noise, ringing etc. These are quality artifacts that should be avoided at all costs, and that are no longer acceptable (and fortunately, less common) in modern video distribution.

In recent years, per-title, content-aware (CA), or even shot-based encoding have become mainstream. Introduced by Netflix [Aaron15], CA encoding finds the best bitrate (or settings) for every segment or shot. Content-adaptive encoding can help boost the quality of every segment of a video sequence. In a “black-box” version, a wide range of settings can be determined and fed into an ML framework, e.g. as demonstrated by Facebook in [Coward16].

While these techniques were introduced for VOD-type encoders, and operating at a very high complexity (e.g. by evaluating multiple options before deciding on the final encoder settings), they can also be applied to real-time encoders. In this case, decisions need to be taken much faster, with low latency and limited lookahead, and before the entire shot is available.

Smart bit distribution inside encoders can make a substantial difference in video encoders, and is necessary to preserve detail in the right places (e.g. players on a football field, or the football itself), to reduce artifacts and preserve textures. ML is well-suited to make decisions in real-time, and with a high degree of content adaptivity. Coding tools such as adaptive quantization, sample adaptive offset (SAO) filtering in HEVC or VVC, ALF in VVC, are excellent candidates for smart decision making based on ML. Some examples of improvements that were obtained by moving from handcrafted algorithms to ML-based subjective decisions are shown in Figure 7.



Figure 7. Example quality improvement between traditional (left) and ML-based (right) texture preservation.

For detection of areas of importance, ML algorithms can be used for higher accuracy. For example logos, faces, football players etc can be better detected and protected in sports games.



Figure 8. ML-based logo detection

Also filtering and post-processing are excellent candidates for smart ML-based techniques. Several articles have been published in this direction, such as in [Yang17] and [Kuanar18]. Also deinterlacing can benefit from CNNs, as described in [Bernasconi20], by combining residual and dense neural networks.

6. Video quality measurement

To verify the effectiveness of VQ optimizations, a great deal of time needs to be spent on subjective quality assessment. This is a process in which viewers (experts and/or non-experts) provide feedback on the video quality. Subjective test methodologies have been designed and standardized to handle this process. And every codec development team will have a set of ‘golden eyes’ in house to guide this process.

In practice, it’s not feasible to verify the subjective quality of each and every video stream or channel, and subjective quality assessment is typically used only in specific occasions, e.g. during encoder comparisons, during the set-up or configuration of an encoder – or whenever an issue occurs. To reduce the high cost of human intervention, *objective* quality measurements have been introduced, to assist VQ measurement in an automated way, and as accurately as possible.

Objective VQ measurement is useful in applications such as encoder comparison and configuration, bitrate selection, ABR ladder (resolution, bitrate) optimization, real-time VQ measurement and monitoring, and in-loop quality control.

Different types of objective VQ measurement exist. In cases where the source video is available, a *Full-Reference* (FR) VQ comparison is possible (Figure 9). Examples of such metrics are MSE (mean squared error) or PSNR (peak signal-to-noise ratio), which calculate the difference between original and distorted pixels.

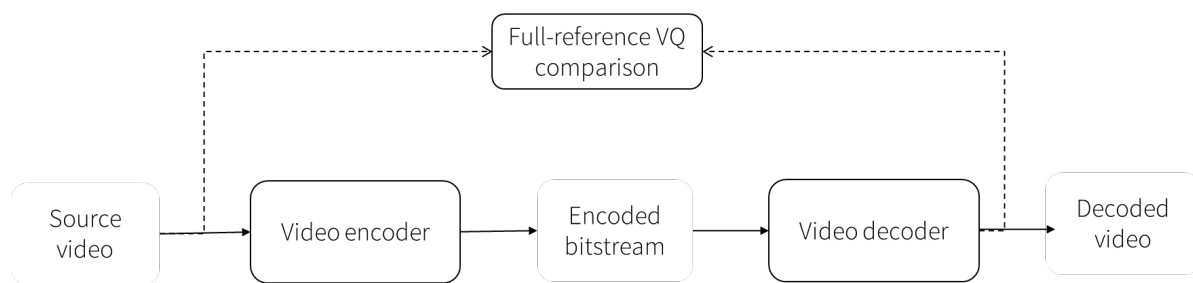


Figure 9. Full-reference video quality assessment

In other cases, a more difficult task is to evaluate the VQ without reference to the original (NR, no-reference). Here, indicators of e.g. the *naturalness* of images need to be calculated to get an impression of the overall quality. While NR metrics are extremely useful, the absence of a reference point makes NR scores less accurate. Inside the encoder, the encoded video can be compared to its input, and FR metrics can be used. We revisit NR metrics later on in the context of quality *monitoring*.

A multitude of FR metrics have been introduced over the past decades, including PSNR (peak signal-to-noise ratio), SSIM (structural similarity), MS-SSIM (multiscale SSIM), VMAF (video multimethod

assessment fusion) and others. The more complex metrics (such as MS-SSIM and VMAF) have shown to provide higher accuracy, while simple metrics like PSNR are not reliable enough for most purposes.

Machine learning has started to play a big role in VQ assessment and the creation of new metrics. An excellent example of an ML-trained metric is VMAF, which takes into account compression and scaling artifacts, and reaches a high accuracy on a variety of test sets. It is being used by many companies and is on its way to become a de facto industry standard. VMAF was trained on subjective data collected from human viewers, and uses a combination of underlying metrics as features. These features are weighted using SVM-based regression, resulting in a score between 0-100 to output the overall quality of the frames.

6.1. From offline to real-time VQ measurement: tracking video quality

A first obvious step to reach operational efficiency is to select the most efficient encoder and encoding configuration. Comparing encoders can be tedious, in particular when many different bitrates and parameters can be configured. Objective metrics can help identify the strengths and weaknesses of encoders, and can point to difference in encoder behavior over time or for different types of video content. Still, the question remains how the encoder will continue to perform, when an upgrade is applied, or when changes in configurations are made. Repeating extensive testing every time is an expensive and time-consuming task. Real-time VQ measurement is the preferred approach to track video quality over time.

As discussed above, ML metrics such as VMAF have been developed that work well for off-line measurement. For *real-time* measurement, however, we need metrics that are both accurate and affordable (meaning fast enough). VMAF's main downside is its computational complexity. Although efforts are ongoing to reduce the complexity of VMAF, its computational requirements remain high, especially when looking at the operational cost.

For operationally feasible VQ measurement, the cost of calculating VQ metrics needs to be reduced by several orders of magnitude compared to VMAF, and should be a fraction of the cost of the encoding itself. While simpler metrics like PSNR are often embedded inside encoders (as in Figure 10), they are not reliable enough for accurate VQ tracking.




Figure 10. Calculating VQ inside the encoder

Machine learning provides ways to calculate VQ metrics in a smarter way. Deep video quality metrics such as DeepVQA have been proposed [Kim18], or using dynamic receptive fields and CNNs [Kim20], which provide state-of-the-art correlation with subjective scores. Still, they require multiple convolutional layers to reach the final score, and as a result, a high computational cost. As an alternative, smart features can be

calculated inside the encoder (or even reused from the pre-analysis stage inside the encoder), as a more powerful input to ML networks. This process is illustrated in Figure 11.




Figure 11. ML-based VQ measurement inside the encoder, based on pre-analysis features.

Finding a balance between input features, network structure, network complexity and accuracy is a complicated process, which as mentioned before, requires a lot of domain knowledge. We have identified ML networks that have more than 90% correlation with subjective scores, and that provide a fast and reliable alternative to expensive VMAF calculation. Example ML networks are shown in Figure 12, where we started from the network on the left, but were able to reduce the number of MACs by 70% to reach the same accuracy with network on the right.




Figure 12. Example networks used for ML-based VQ prediction.

Once a VQ metric with acceptable accuracy and complexity is available in the encoder, it becomes possible to track video quality, and to verify 24/7 that a certain quality level is reached during encoding.

6.2. From rate control to quality control

Measuring video quality is one part of reaching efficient compression. An even bigger challenge is to actively *control* the VQ in real-time. In video compression, it is essential to reach a high video quality. But at the same time, for operational efficiency, you want to avoid overspending bits where they don't matter. At a certain point, bits can be added, but visual quality will no longer improve. This leads to waste, higher-than-necessary-bitrate, and delivery networks that are overloaded with redundant bits. Accurate video

quality metrics can help determine the saturation point and improve the intelligence of encoders. Figure 13 shows an example *rate-distortion* plot where the saturation effect is visible.




Figure 13. Above a certain bitrate, adding more bits will no longer (or hardly) improve quality

In the end, more important than constant bit rate, is to reach *consistent* quality. Note that *constant* quality is not feasible in practice. In all practical encoder systems, there is a maximum bitrate (cap rate) constraint, which will limit the allocation of bits. In case of very difficult scenes, quality might be limited because of that cap. Still, on average, in-loop quality steering will lead to less overspending (wasted bits) and less underspending (quality drops).

On top of VQ measurement, quality control poses an even more challenging problem, as VQ needs to be predicted before encoding. As a result, you end up with a chicken-and-egg problem. Fortunately, ML turns out to enable accurate prediction of encoded VQ, even before encoding. In this way, rate control can be turned into *quality* control. The result is a VBR stream which not only reaches a consistent quality, but also saves bits compared to traditional CBR rate control. Figure 14 shows how the different components fit together: ML-based rate control and VQ measurement, working alongside the ML-optimized video encoder.

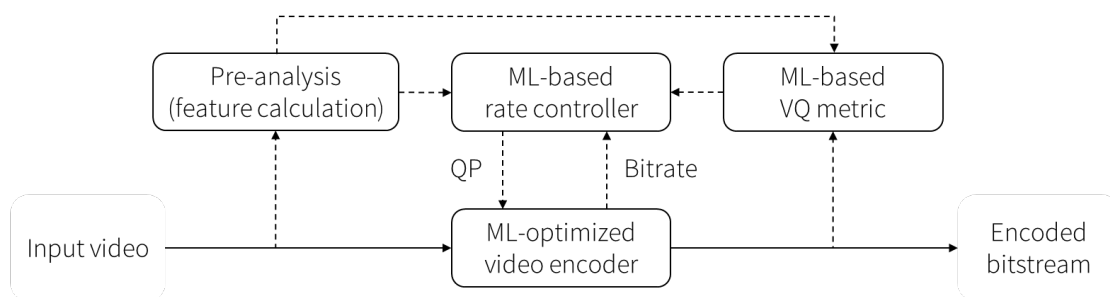


Figure 14. Quality-controlled compression

Figure 15 shows how quality-controlled compression can accurately reach a targeted quality value (in this case VMAF=90). Some variation is possible, and acceptable, when sudden changes in content occur, or when bitrate caps are reached.




Figure 15. Result of quality-controlled compression (target VMAF=90).

Quality-controlled compression can be extended across multiple streams, e.g. in *statistical multiplexing* scenarios. Even though statmuxed streams are inherently VBR in nature, traditionally they have not been optimized with a particular VQ in mind. By using quality control, target quality levels can be specified for each individual stream in a statmux bundle, while still satisfying the total constant bit rate.

Also, *adaptive bitrate (ABR)* encoding can benefit – not only while optimizing individual streams, but also across the whole ladder. Instead of statically defining an ABR ladder as fixed resolution/bitrate pairs, it becomes possible to define quality targets for bitrate ladders, avoiding the need to specify fixed bitrates or even resolutions.

7. Video quality monitoring

One step beyond video quality tracking is *VQ monitoring*. Automated monitoring reduces the need for human inspection to keep track of video quality, with the objective of detecting issues with input sources, transcoding, delivery, or unexpected quality loss. This is another area where machine learning can provide clear benefits.

While VQ measurement as described in the previous section can give a good view of the quality loss introduced by encoding or transcoding itself (by comparing output and input), it does not provide a view of the *absolute* quality of the signal, or whether a potential problem might have occurred. If a disrupted signal enters the transcoder, the FR metric might still report high output quality. Basically, garbage in leads to garbage out...

To allow more flexibility in monitoring, *no-reference (NR) metrics* are essential since they give an impression of the overall video quality. In recent years, NR VQ assessment has been a hot research topic, with successful introduction of new metrics, often using deep neural networks, such as:

- [You19] proposed 3D-CNN and LSTM networks to extract local spatiotemporal features from small cubic video clips.
- [Bosse17] presented a deep neural network approaches with 10 convolutional layers and 5 pooling layers, and 2 fully connected layers for regression, totaling 5.2 million trainable parameters.
- [Bianco18] discusses different design choices for CNN-based blind image quality assessment, where the best design reaches a correlation of 0.91 with subjective scores.

For a good overall impression of quality, a combination of metrics (or *indicators*) is recommended. In the Video Quality Experts Group, studies have been made of different NR metrics. One of the higher-accuracy metrics is Sawatch (v3) [vqeg22], which combines multiple underlying quality indicators to detect e.g. blurriness, blockiness, saturation levels etc.

Still, even the more complex NR metrics reach relatively low accuracies when compared to FR metrics. Also, when quality issues occur, they are more prone to false positives or false negatives when compared to full-reference comparison. For more reliable VQ monitoring, measurements can be taken in different places during delivery, and their results correlated in a central point (e.g. in a cloud service). This provides a reliable comparison point for VQ degradation, transcoding artifacts, or simply transmission errors. Whenever a problem occurs, captures taken from the devices can be uploaded for deeper inspection. Only in those cases, human intervention is needed.

By combining FR and NR metrics, different monitoring scenarios become possible, such as:

- Observing quality fluctuations at a *single point* in the delivery chain. For example, quality measurements can be taken inside an encoder or transcoder and monitored over time. For this use case, either FR metrics (comparing encoded output to its input, Figure 16(a)) or NR metrics (without comparison to a source, Figure 16(b)) can be used.
- Comparing measurements taken in *multiple points* in the video delivery chain (Figure 16(c)). For example in video transport cases, or in multiple processing steps, metrics can be calculated to track the evolution of quality in the delivery chain.

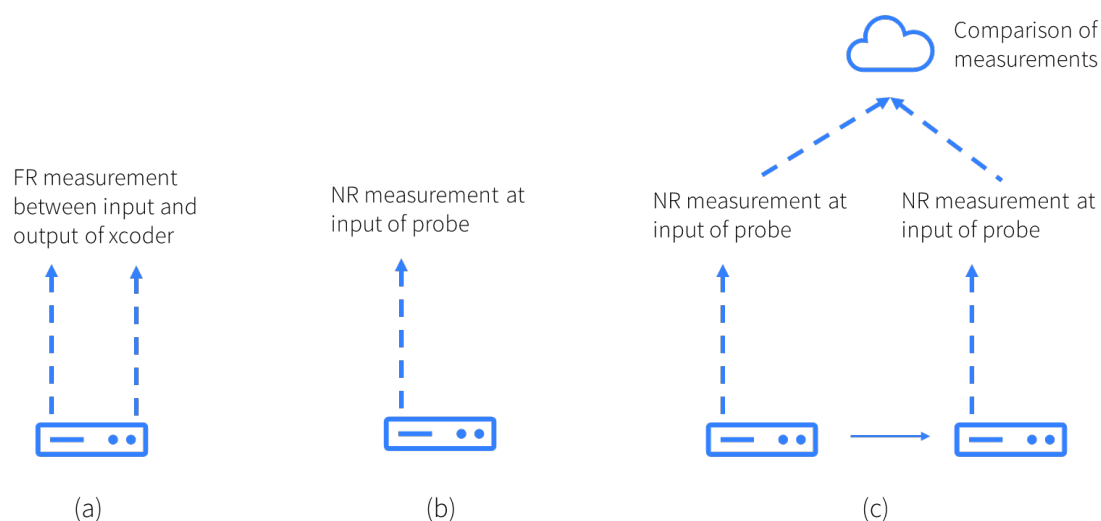


Figure 16. FR and NR quality measurement for VQ monitoring.

By combining these different measurement capabilities, monitoring across the whole delivery chain becomes possible, from contribution encoding to end delivery, as illustrated in Figure 17.

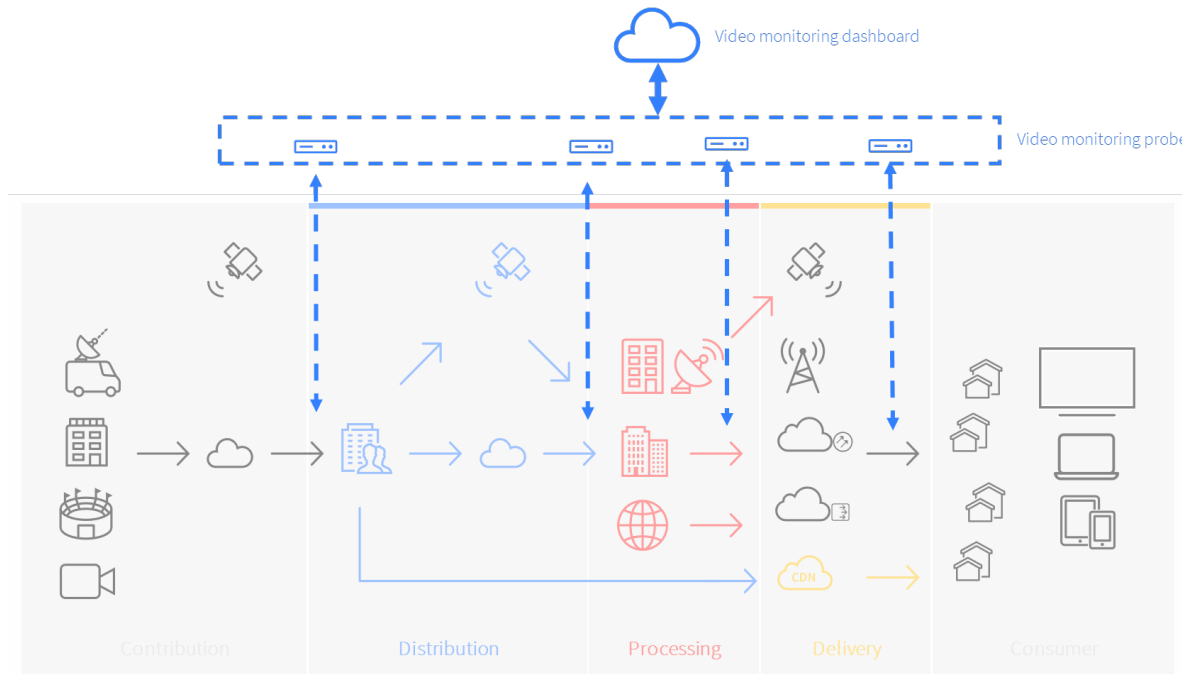


Figure 17. VQ monitoring at different points in the video delivery chain.

8. Conclusions

In this paper, we discussed several areas in video coding where machine learning has shown to provide benefits, including computational complexity reduction, rate control, and subjective video quality improvement. In contrast to popular deep neural networks for image recognition and classification tasks, less complex networks are needed in real-time scenarios, resembling the field of “TinyML”. Relatively shallow ML networks are both computationally acceptable and have been shown to lead to improvements.

The field of VQ measurement benefits from ML approaches, both for non-real-time (such as VMAF) and real-time metrics that can be embedded inside encoders. For VQ monitoring, the combination of FR and NR metrics can be used to analyze quality across the end-to-end delivery chain. By bringing together the ML optimizations for rate control, subjective improvements and VQ measurement, quality-controlled compression can be achieved.

Abbreviations

ABR	adaptive bit rate
AI	artificial intelligence
AVC	Advanced Video Coding
CBR	constant bit rate
CNN	convolutional neural network
DNN	deep neural network
DPB	decoded picture buffer
FR	full-reference
HD	high definition
HEVC	High Efficiency Video Coding
LSTM	long short-term memory
MAC	multiply-accumulate
ML	machine learning
MPEG	Moving Picture Experts Group
MSE	mean squared error
NN	neural network
NR	no-reference
PSNR	peak signal-to-noise ratio
QP	quantization parameter
SD	standard definition
SSIM	structural similarity
SVM	support vector machine
VBR	variable bit rate
VQ	video quality
VMAF	video multimethod assessment fusion
VNNI	vector neural network instructions
VOD	video on demand
VVC	Versatile Video Coding

Bibliography & References

- [Aaron15] Anne Aaron, Zhi Li, Megha Manohara, Jan De Cock and David Ronca, “Per-Title Encode Optimization”, <https://netflixtechblog.com/per-title-encode-optimization-7e99442b62a2>.
- [Andreopoulos22] Y. Andreopoulos, “Neural pre and post-processing for video encoding with AVC, VP9 and AV1”, AOMedia Research Symposium 2022
- [Barman19] N. Barman, E. Jammeh, S. A. Ghorashi and M. G. Martini, “No-Reference Video Quality Estimation Based on Machine Learning for Passive Gaming Video Streaming Applications”, in IEEE Access, vol. 7, pp. 74511-74527, 2019.
- [Bernasconi20] A. Bernasconi, A. Djelouah, S. Hattori, C. Schroers, “Deep Deinterlacing”, SMPTE 2020
- [Bhat21] M. Bhat, J. -M. Thiesse and P. L. Callet, "VVC partitioning decision driven by machine learning for a comprehensive hardware encoder," 2021 IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP), 2021.
- [Bianco18] S. Bianco, L. Celona, P. Napoletano and R. Schettini, "On the use of deep learning for blind image quality assessment", Signal Image and Video Processing, vol. 12, no. 2, pp. 355-362, Feb. 2018.
- [Bishop95] Christopher Bishop, “Neural Networks for Pattern Recognition”, Oxford University Press, 1995, ISBN 0-19-853864-2.
- [Bosse17] S. Bosse, D. Maniry, K-R. Müller, T. Wiegand and W. Samek, "Deep neural networks for no-reference and full-reference image quality assessment", IEEE Trans. Image Proc., vol. 27, no. 1, pp. 206-219, Oct. 2017.
- [Coward16] Mike Coward, “AI Encoding”, Video @ Scale 2016, <https://www.facebook.com/at-scale-events/videos/ai-encoding-at-scale/1682906415315789/>
- [Goodfellow16] Ian Goodfellow, Yoshua Bengio, Aaron Courville, “Deep Learning”, MIT Press, 2016.
- [Google22] <https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/>, June 2022.
- [He16] K. He, X. Zhang, S. Ren, J. Sun, “Deep Residual Learning for Image Recognition”, Computer Vision and Pattern Recognition (CVPR), 2016.
- [Kim18] Kim, W., Kim, J., Ahn, S., Kim, J., Lee, S. (2018). Deep Video Quality Assessor: From Spatio-Temporal Visual Sensitivity to a Convolutional Neural Aggregation Network. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science(), vol 11205.
- [Kim19] J. Kim, S. Blasi, A. S. Dias, M. Mrak and E. Izquierdo, “Fast Inter-prediction Based on Decision Trees for AV1 Encoding”, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2019, pp. 1627-1631.
- [Kim20] W. Kim, A. -D. Nguyen, S. Lee and A. C. Bovik, "Dynamic Receptive Field Generation for Full-Reference Image Quality Assessment," in IEEE Transactions on Image Processing, vol. 29, pp. 4219-4231, 2020.

- [Kuanar18] S. Kuanar, C. Conly and K. R. Rao, "Deep Learning Based HEVC In-Loop Filtering for Decoder Quality Enhancement," 2018 Picture Coding Symposium (PCS), 2018, pp. 164-168.
- [Li17] Y. Li, B. Li, D. Liu and Z. Chen, "A convolutional neural network-based approach to rate control in HEVC intra coding," IEEE Visual Communications and Image Processing (VCIP), 2017.
- [Liu16] Z. Liu, X. Yu, Y. Gao, S. Chen, X. Ji and D. Wang, "CU Partition Mode Decision for HEVC Hardwired Intra Encoder Using Convolution Neural Network", IEEE Transactions on Image Processing, vol. 25 (11), pp. 5088-5103, Nov. 2016.
- [Liu22] Y. Liu, M. Abdoli, T. Guionnet, C. Guillemot and A. Roumy, "Light-Weight CNN-Based VVC Inter Partitioning Acceleration," 2022 IEEE 14th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP), 2022.
- [Mao16] X.-J. Mao, C. Shen and Y.-B. Yang, Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections, Barcelona, SPAIN:NIPS, 2016.
- [Pfaff18] J. Pfaff, P. Helle, D. Maniry, S. Kaltenstadler, W. Samek, H. Schwarz, D. Marpe, and T. Wiegand, "Neural network based intra prediction for video coding," SPIE Applications of Digital Image Processing XLI, vol. 10752, 2018.
- [Santamaria20] M. Santamaria, S. Blasi, E. Izquierdo and M. Mrak, "Analytic Simplification of Neural Network Based Intra-Prediction Modes for Video Compression", IEEE International Conference on Multimedia & Expo Workshops (ICMEW), 2020.
- [Simonyan14] Simonyan, Karen & Zisserman, Andrew, "Very Deep Convolutional Networks for Large-Scale Image Recognition", 2014, arXiv 1409.1556.
- [Su19] H. Su, M. Chen, A. Bokov, D. Mukherjee, Y. Wang and Y. Chen, "Machine Learning Accelerated Transform Search for AV1," Picture Coding Symposium (PCS), 2019.
- [vqeg22] NRMetricFramework, <https://github.com/NTIA/NRMetricFramework>
- [Warden20] Pete Warden and Daniel Situnayake, "TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers", 1st Edition, O'Reilly, 2020.
- [Wu21] G. Wu, Y. Huang, C. Zhu, L. Song and W. Zhang, "SVM Based Fast CU Partitioning Algorithm for VVC Intra Coding," 2021 IEEE International Symposium on Circuits and Systems (ISCAS), 2021.
- [Yang17] R. Yang, M. Xu and Z. Wang, Decoder-side HEVC Quality Enhancement with Scalable Convolutional Neural Network, Hong Kong, China:IEEE, ICME, 2017.
- [You19] J. You and J. Korhonen, "Deep Neural Networks for No-Reference Video Quality Assessment," 2019 IEEE International Conference on Image Processing (ICIP), 2019, pp. 2349-2353, doi: 10.1109/ICIP.2019.8803395.

Bandwidth and Latency Analysis of 25GS PON

An Overview of 25GS Passive Optical Networks

A Technical Paper prepared for SCTE by

Edward Boyd

CTO

Tibit Communications, Inc

459 N Gilbert Rd, Suite A255, Gilbert, AZ 85234, USA

ed.boyd@tibitcom.com

Edward Walter

Director Member Technical Staff

AT&T, Inc.

123 Timber Mountain Dr., Boerne, TX 78006, USA

Ew8532@att.com

Fernando Villarruel

Chief Architect, MSO Practice

Ciena Corp

1120 Sanctuary Pkwy, Alpharetta, GA, 30004, USA

fvillarr@ciena.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Adding 25GS PON to the ODN.....	5
3. System Upgrade to 25GS PON	8
3.1. Large Chassis OLT	9
3.2. Pizza box OLT	10
3.3. Micro-OLT	11
4. Use cases for 25GS PON	12
5. Cost Analysis of 25GS PON	14
5.1. OLT/ONU	14
5.2. Optical Budget.....	15
5.2.1. Transmission Lasers.....	15
5.2.2. DSP, High Speed ADC, Coherent	16
5.2.3. Cost Summary	16
6. Bandwidth Analysis of 25GS.....	16
6.1. Downstream Bandwidth Calculation	17
6.2. Upstream Bandwidth Calculation	19
6.2.1. Burst Overhead	19
6.2.2. Static (Fixed) BW Allocation	20
6.2.3. Dynamic BW Allocation (DBA).....	21
7. Latency Analysis of 25GS	23
7.1. Downstream Latency	24
7.2. Upstream Latency	24
7.2.1. Start Latency	24
7.2.2. Continuation Latency	27
7.2.3. Queuing Delay	28
7.2.4. Balanced Delay	28
7.2.5. XGS and 25GS Sharing the Fiber.....	29
8. Conclusion.....	30
Abbreviations	30
Bibliography & References.....	31

List of Figures

Title	Page Number
Figure 1 – PON Wavelength Plan.....	6
Figure 2 – GPON and XGS PON Coexistence	6
Figure 3 – XGS and 25GS (UW1) PON Coexistence.....	7
Figure 4 – GPON and 25GS (UW0) PON Coexistence	7
Figure 5 – GPON, XGS, and 25GS (UW3) PON Coexistence	7
Figure 6 – Example of Chassis with XGS Cards and 25GS Cards	9
Figure 7 – Example of Chassis with Dual Speed XGS/25GS Line Cards	10
Figure 8 – XGS/25GS MicroOLTs in a multi-rate Ethernet Switch	11
Figure 9 – XGS and 25GS Coexistence Architecture.....	12
Figure 10 – Relative PON Transceiver Cost in Volume.....	15
Figure 11 – XGS and 25GS Downstream Framing	17

Figure 12 – Upstream Burst Overhead	19
Figure 13 – Upstream Framing Overhead	20
Figure 14 – DBA Upstream Bursts	21
Figure 15 – DBA Lab Results for 64 active ONUs	23
Figure 16 – DBA Lab Results for 1 active ONU on 64 ONU system	23
Figure 17 – Upstream Start Latency vs Polling	25
Figure 18 – Polling and the Single Transmitting ONU	25
Figure 19 – XGS Upstream Bandwidth versus Polling Interval	26
Figure 20 – 25GS Upstream Bandwidth versus Polling Interval	26
Figure 21 – Piggybacking and 64 ONUs transmitting	27
Figure 22 – 64 Active ONU Latency	27
Figure 23 – 64 Active ONU Bandwidth	28

List of Tables

Title	Page Number
Table 1 – Wavelength Definitions	5
Table 2 – ITU-T Optical Class Definitions	6
Table 3 - 1:16 Splitter, 25% 25GS PON, Decrementing 25GS	13
Table 4 - 1:16 Splitter, 25GS PON Capacity Met	14
Table 5 – Downstream Overhead	18
Table 6 – Downstream Bandwidth Available	19
Table 7 – Static Allocation Analysis for 20km, 64 ONU PON	21
Table 8 – Dynamic Bandwidth Allocation Analysis for 20km, 64 ONU PON	23
Table 9 – Sub Millisecond Delay for 20km, 64 ONU PON	29
Table 10 – Shared PON with XGS (up to 64 ONUs) and 25GS (up to 8 ONUs)	29
Table 11 – Shared PON with XGS (up to 64 ONUs) and 25GS (up to 16 ONUs)	30

1. Introduction

The fiber access market is largely dominated by passive optical networks (PON). A PON is a point-to-multipoint network supporting multi-service traffic across a single fiber that can be split up to 1:128 in the outside plant and run end-to-end passively (no amplifiers required). The passive splitting feature conserves trunk fiber and minimizes equipment in the central office, node, and headend.

The IEEE and ITU-T define standards for PON at multiple speeds. Ethernet PON (EPON) is specified by the IEEE 802.3 at 1 Gigabit per second (Gbps) in the upstream and downstream direction. The IEEE has also specified a 10 Gbps version called 10G-EPON deployed in Japan and with North American MSOs. The ITU-T defined the Gigabit PON (GPON) that provides 1 Gbps upstream and 2 Gbps downstream. Many operators in China deployed with GPON and selectively upgraded to a higher speed version called NGPON1 with 10 Gbps downstream and 2.5 Gbps upstream. The ITU-T later defined a 10 Gbps downstream and 10 Gbps upstream version of GPON called XGS-PON. Operators in Europe, North America, and other parts of the world started with GPON and upgraded to XGS. XGS has begun to dominate PON deployment and will be the favored solution for residential PON into the future.

Since user demand for bandwidth will continue to increase, it is inevitable that PON speeds beyond 10 Gbps will be needed. The IEEE most recently added a 25 Gbps version with the option of 50 Gbps by optically combining two 25 Gbps signals on the same fiber. Using the IEEE's 25 Gbps optical layer definition, a group of companies formed a Multi-Source Agreement (MSA) to create a 25 Gbps version of ITU-T's GPON. The ITU-T standard has been actively defining multiple versions of 50 Gbps PON. 50 Gbps downstream with 12.5 Gbps upstream or 25 Gbps upstream has been created. A 50 Gbps upstream is currently being defined by the ITU-T. Unlike the IEEE 50 Gbps, the ITU-T uses a single wavelength. Cable Labs has also started a 100 Gbps Coherent PON standard as well.

With multiple operators announcing field trials and devices becoming available, we wanted to write a paper on 25 Gbps PON to explore the bandwidth, latency, upgrade, and cost impacts. The three authors of this paper represent a component supplier, a system provider, and a major operator. From the component level, we want to understand the cost and technology differences between the 10 Gbps PON technologies, 25 Gbps PON technologies, and the future 50 Gbps/100 Gbps PON standards. From a system level, we will explain the coexistence and upgrade paths from lower PON speeds to 25 Gbps. From an operator perspective, we want to explore the use cases for a 25 Gbps symmetric PON solution. Finally, this paper explains the expected bandwidth and latency possibilities with a 25 Gbps PON system. PON has very significant overhead from forward error correction, framing, and burst overheads that will drop the 25 Gbps line rate. In the upstream direction, tradeoffs must be made between low latency and high bandwidth. This paper will present some simple models to explore those tradeoffs.

Since analyses of both the ITU-T and IEEE standards would be prohibitive for a single document, this paper focuses on the 25GS MSA standard that uses the ITU-T's XGS framework. (We believe that a significant amount of the analysis and conclusions would be the same for the IEEE's 25G EPON.) The 25GS MSA defines a downstream rate of 25 Gbps with upstream rates of either 10 Gbps or 25 Gbps. This paper will only consider the 25 Gbps symmetric system where both upstream and downstream are 25 Gbps.

The results in this paper are largely based on modeling of the XGS and 25GS PON systems with some spot checking of the model in XGS mode from lab tests. Since 25GS technology is still in development, large system testing results are not available.

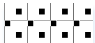

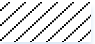


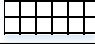







2. Adding 25GS PON to the ODN

Migrating between PON technology can be implemented via two methods:

- New fiber and splitters for each technology
- Wavelength coexistence via a single feeder fiber

For areas where an operator is fiber-rich and has capacity for splitters at the distribution enclosure, they may elect to use additional feeder fiber and additional splitters. In cases where fiber is constrained, an operator may elect to deploy a wavelength coexistence strategy. There are two methods for supporting coexistence. One method, which is internal to the OLT optic, is via a multi-PON module, in which multiple PON technologies (GPON and XGS, or XGS and 25GS PON) are integrated into a single module. The other option is an external coexistence element that combines multiple PON modules.

Table 1 – Wavelength Definitions

Legend	Wavelength Name	Center Wavelength (nm)	Wavelength Range (nm)	Specification
	ITU GPON US	1310	± 20	GPON "Reduced", ITU-T G.984.5
	ITU GPON DS	1490	± 10	ITU-T G.984.2
	ITU XGS PON US	1270	± 10	ITU-T G.9807.1
	ITU XGS PON DS	1577	+3/-2	ITU-T G.9807.1
	25G (UW0) US (MSA & 25G EPON)	1270	± 10	IEEE Std 802.3ca
	25G (UW1) US (MSA & 25G EPON)	1300	± 10	IEEE Std 802.3ca
	25G PON (UW3) US 25GS MSA Only	1286	± 2	25GS-PON MSA
	25G (DW0) DS (MSA & 25G EPON)	1358	± 2	IEEE Std 802.3ca
	25G (UW2) US (50G EPON - 2nd 25G WL)	1320	± 2	IEEE Std 802.3ca
	25G (DW1) DS (50G EPON - 2nd 25G WL)	1342	± 2	IEEE Std 802.3ca
	ITU 50G PON (12.5G/25G/50G) (Option 1) US	1270	± 10	ITU-T G.9804.3
	ITU 50G PON (12.5G/25G/50G) (Option 2) US	1300	± 10	ITU-T G.9804.3
	ITU 50G PON DS	1342	± 2	ITU-T G.9804.3




Figure 1 – PON Wavelength Plan

As illustrated in Figure 1, GPON and XGS PON upstream(US) and downstream(DS) wavelengths are separated which allows for coexistence across a single feeder fiber. The simple figure below illustrates how a passive coexistence element (CE) can bring XGS-PON US/DS Wavelengths together with GPON US/DS wavelengths. This would not be possible if the wavelengths overlapped.




Figure 2 – GPON and XGS PON Coexistence

In conjunction with the use of a CE, two other key factors must be considered. First, the ONTs must be able to filter out or block the downstream wavelengths for the coexisting OLT. In the GPON-XGS coexistence example above, this means that the GPON ONT must block the downstream XGS wavelength and the XGS ONT must block the downstream GPON wavelength. Second, since both PON variants share the same ODN, the OLT and ONT optics for both PON variants should be compatible with the attenuation range of the shared ODN.

Table 2 below from ITU-T G.9807.1 gives the attenuation range classes (B+, N1, etc....) for the ODN. May want to pull in the attenuation range info as a reference of ODN Classes.

Table 2 – ITU-T Optical Class Definitions

PON TYPE	ODN Class	Attenuation Range (dB)
GPON	B+ class	13-28
GPON	C+ class	17-32
XGS/25GS	N1 class	14-29
XGS/25GS	N2 class	16-31

PON TYPE	ODN Class	Attenuation Range (dB)
XGS	E1 class	18-33
XGS	E2 class	20-35

As identified in Table 2, B+ and C+ are ODN classes that supports a maximum fiber plant loss of 28db and 32db respectively. In addition, N1 (29dB), N2 (31dB), E1 (33dB), and E2 (35dB) are defined in ITU-T, while N1 and N2 are defined in the 25GS PON MSA specification.




Figure 3 – XGS and 25GS (UW1) PON Coexistence




Figure 4 – GPON and 25GS (UW0) PON Coexistence




Figure 5 – GPON, XGS, and 25GS (UW3) PON Coexistence

As seen in Figure 3, Figure 4, and Figure 5, there are three paths to facilitating coexistence between 25GS with GPON, 25GS with XGS, or 25GS with both prior PON technologies. The IEEE 802.3 standard developed the options for UW0 (1290-1310 nm) or UW1 (1260-1280 nm) with DW0 (1356-1360 nm).

The 25GS-PON MSA incorporated these into their specification and added UW3 (1284-1288 nm) with DW0 (1346-1350 nm).

UW0 and UW1 are 20nm wide, supporting the use of a DML for an N1 class ODN. UW3, on the other hand, is 4nm wide and will require an EML for an N1 class ODN. UW3 allows for triple coexistence between GPON, XGS, and 25GS but there are additional cost considerations that need to be weighed by the operator. With UW0 or UW1, an operator achieves double coexistence (GPON and 25GS or XGS and 25GS) with less expensive optics but loses the ability to coexist across three PON technologies. The 3rd PON technology must be placed on a separate fiber and splitter.

While this paper focuses on 25GS PON it should be noted that 50G PON (EPON or ITU-T) can coexist across a single feeder fiber with a prior PON technology.

3. System Upgrade to 25GS PON

Whether an operator is introducing PON for the first time or evolving a PON network to services supported by a 25G solution, a general consideration is how the new deployment implementation will support introduction to or coexistence with 25GS. Historically, PON technology evolutions involve evaluating customer behavior, preferred network topology, platform MAC scale and implementation granularity. These evaluations typically revolve around how and where an OLT is implemented.

There are three OLT model types: the large chassis, the pizza box, and the MSA pluggable OLT. In the following sections, we review the implications of these options with regards to introduction and evolution to 25GS.

3.1. Large Chassis OLT



Figure 6 – Example of Chassis with XGS Cards and 25GS Cards

The large OLT platforms, as seen in Figure 6, offer from the twenties to hundreds of PON ports. They are typically made up of modular, dedicated functional cards integrated into mainframes that distribute signal back and forth via a high-speed backplane. The platforms are physically large in the range of 10 rack units or more, and require locations with dedicated HVAC systems. OLT platforms of this type are quite efficient when dedicated in scale to the particular function or technology they were created for very large contiguous deployments of a particular PON technology, servicing a large number of customers within 20 km of a central office or hub. Their drawbacks are the sunk investment of a large platform when not used in scale and their otherwise inflexibility to work beyond the technology they were designed for.



Figure 7 – Example of Chassis with Dual Speed XGS/25GS Line Cards

The evolution of large OLT platforms to 25GS is implemented in the context of line card scale capability, their modularity, i.e., processing capacity and backplane throughput, with the understanding that if a box was created for, as an example, 100 PON instances of GPON, transitioning through XGS then to 25GS would likely result in a continual reduction of port counts to the point where it is no longer an efficient large scale box. Similarly, although to a lesser extent, evolving from XGS will also result in the PON instance capability being reduced. In one example, 25GS is supported by a new line card with half the number of PON ports. In the other example, the dual speed line card can select the speed by using a different transceiver. Large platforms can also be equipped to function as broadband network gateways (BNG) and collocated for large scale deployments.

3.2. Pizza box OLT

The “pizza box” refers to integrated platforms of a reduced size, typically one rack unit and generally less than four. The attractive quality for pizza box OLTs is flexibility. They can be used to service less dense population areas, and because of their size (and if temperature hardened) they can be placed in remote cabinets to service customers at distances much greater than 20km. They can also act in unison as a large OLT when racked together, interconnected by a top or rack switch and facilitated by a centralized controller.

With regards to evolution to 25GS PON, the flexibility of pizza box platforms is lost. Boxes that are not preconditioned to support a 25GS solution must be replaced with a new box. This is certainly a service-affecting change. If a 25GS option is available, it is most likely at a lower density than for prior PON technologies.

3.3. Micro-OLT




Figure 8 – XGS/25GS MicroOLTs in a multi-rate Ethernet Switch

As seen in Figure 8, the MSA pluggable OLT solution is also known as a micro-OLT (μ OLT). (MSA pluggable refers to SFP+, SFP28, CFP2, etc.) These μ OLTs facilitate one PON instance at a time through a single port on an Ethernet switch. In contrast to the large chassis and pizza box solutions which combine the MAC domain chipsets with the switching fabric but isolate the optics, the μ OLT technology combines the MAC domain chipsets with the optical interface, leaving the switching fabric separate.

The scalability of these platforms is determined by the host switch, which also determines the geographic deployment. The μ OLT can service sparsely deployed regions, including being temperature hardened and deployed in remote cabinets, in the most granular manner—one PON instance at a time. The flexibility is nearly boundless when creating PON deployments with μ OLTs.

A previously unattainable dimension of flexibility is also available when using μ OLTs. Because the PON service is limited to a dedicated switch port per the pluggable, combining other services on the same Ethernet fabric is now possible. Studies in recent years have looked at the convergence of services over the same access network, and the μ OLT facilitates service integration at the layer 2 and layer 3 levels. The considerations for deterministic bandwidth assignments are beyond the scope of this paper, but there is a whole field of study that expects to leverage μ OLTs as an easy method to enable service convergence.

With regards to the evolution or introduction to 25GS, the flexibility follows the granularity of the μ OLT. Transitions to 25GS would be a straightforward port change; μ OLTs have been generally available since 10G PON implementations. Note that the technical challenge is for pluggable solutions to maintain the development necessary to include additional capabilities, such as thermal growth into continually small form factors. One particularly interesting dimension of pluggable solutions is their ability to support the coexistence of different PON technologies on the same bridging domain. For example, a single switch can be populated with a mix of both 10G and 25GS services. Such flexibility

could potentially support large usage end-line customers, gradually transitioning their services to 25GS. This collection of system design qualities makes the μ OLT a very attractive option for 25GS solutions.

4. Use cases for 25GS PON

There are several factors that will determine how 25GS is used:

- Coexistence can be used as a mechanism to move from one PON technology to the next (e.g., GPON to XGS or XGS to 25GS) without having to place additional feeder fiber and splitters.
- XGS is now in scale deployment with many operators either offering or preparing to offer a multi-gig broadband service.
- Most of the cost of fiber-to-the-premises (FTTP) is fiber placement. If fiber is abundant, then it may be more cost effective to use a feeder fiber and place a new splitter. If fiber is constrained, then a coexistence strategy may prove to be more cost effective.
- ONTs are normally dedicated per unit/household and thus component costs are important to constrain (i.e., DML vs EML).
- OLTs are a significant cost but can be shared among multiple subscribers:
 - Large chassis: 128-256 PONs x 1:64 split (8K to 16K subscribers)
 - Pizza box: 24-48 PONs x 1x1:64 split (1.5K to 3K subscribers)
 - μ OLT: Varies from 1-48 PONs based on the switch x 1:64 split (64 to 3K subscribers)

So, the question is, what do you do with 2.5 times more capacity than XGS?

As shown in Figure 9 the architecture proposal will support consumer, business, and mobility applications across a single feeder fiber and splitter. This design will allow a mix of best effort multi-gig broadband services with a maximum latency of 4ms through the OLT, as well as business and mobility applications with 10G+ multi-gig broadband services with a “protected” (guaranteed) amount of capacity and sub-ms latency across the OLT.

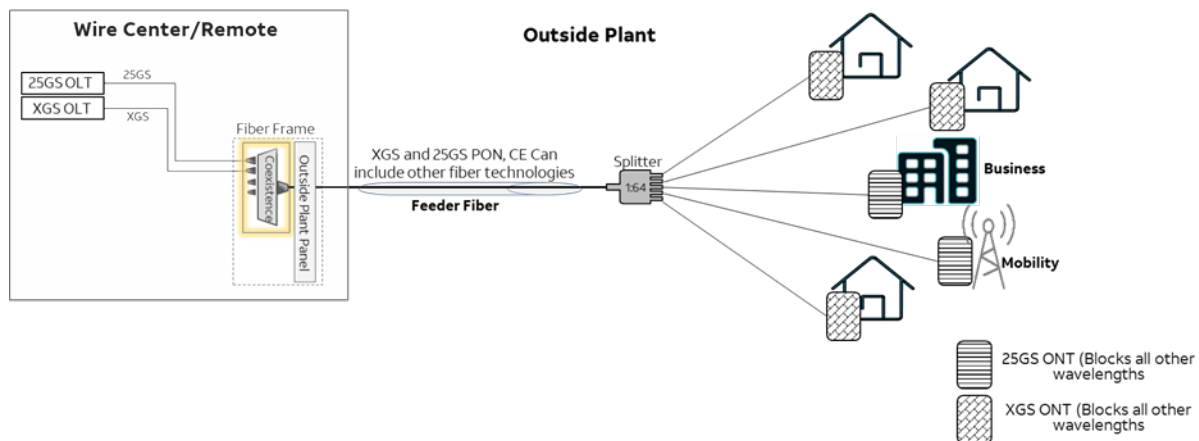


Figure 9 – XGS and 25GS Coexistence Architecture

XGS will continue to be a valuable technology, and currently supports consumer services through multi-gig broadband. Many users will continue to be supported over XGS for years to come. 25GS offers not only the ability to offer 10G+ broadband speed tiers, but also much more capacity at a small incremental increase in cost. 25GS could initially target business, work from home subscribers, mobility, and capacity failover for bandwidth intensive subscribers on XGS. Since consumer scale deployments on

25GS are assumed to be a future phase, a smaller percentage of the total is assigned to the splitter. For example, a splitter of 1:64 is usually capacity-managed to the total. In this design, however, two PON technologies are managed separately. XGS would have access to capacity of all 64 ports while 25GS would be set to capacity trigger at 25% (16 ports of a 1:64 splitter in this example). Placement of the XGS or 25GS ONT and assignment to the appropriate OLT PON port determine what is serviced on the splitter. An operator does not need to lock in specific ports of a passive splitter. Any port on the splitter can support XGS or 25GS as long as the capacity triggers are not exceeded. In this example, XGS can be deployed across more than 48 ports, and the capacity algorithm for 25GS would be decremented to maintain the total of 64.

The four initial use cases for 25GS mentioned above are described below:

- **Business:** 25GS has the capacity to support large numbers of customers behind a switch with an SFP28 25GS ONT or a purpose-built gateway. In this case, capacity, “protected” bandwidth, and sub-ms latency would benefit this market without driving significant cost.
- **Work from home:** an operator could use the additional capacity to offer consumer-grade broadband service-mapped interfaces (such as wireless SSID or physical LAN interfaces) while also supporting business services to other interfaces.
- **Mobility:** 25GS allows for more than 20Gbps x 20Gbps of usable bandwidth. “PON as a transport” for backhaul and mid-haul mobility services is possible with a sub-ms PON.
- **Capacity protection:** As new applications emerge that drive additional capacity requirements on XGS, 25GS offers a capacity failover from XGS and allows the operator to offer more tailored services to these high-end users. In addition, the XGS is managed for a better customer experience for the remaining subscribers. Note that with a coexistence strategy, the transition for an existing XGS subscriber to 25GS requires the shipment and customer installation of a new 25GS PON ONT or gateway.

Refer to the following Table 3 and Table 4 for details about PON capacity by splitter.

Table 3 - 1:16 Splitter, 25% 25GS PON, Decrementing 25GS

Splitter Port	XGS	25GS (Set at 25%)	
1	1		
2		1	
3	2		
4	3		
5	4		
6		2	
7	5		
8		3	
9	6		
10	7		
11	8		
12	9		
13	10		
14	11		
15	12		
16	13		Assigned to XGS and decrements 25GS

Table 4 - 1:16 Splitter, 25GS PON Capacity Met

Splitter Port	XGS	25GS (Set at 25%)	
1	1		
2		1	
3	2		
4	3		
5	4		
6		2	
7	5		
8		3	
9		4	Capacity of 25% on 25GS reached, new splitter
10	6		
11	7		
12	8		
13	9		
14	10		
15	11		
16	12		

5. Cost Analysis of 25GS PON

5.1. OLT/ONU

The highest equipment cost in a PON deployment is the customer-side device, the Optical Network Unit (ONU), also known as the Optical Network Terminal (ONT). These devices connect to an Optical Line Terminal (OLT) at the operator side. Since the OLT is designed to split a single fiber to 64 ONUs at 20 kilometers, adding only \$1 to the cost of the ONU adds \$64 to the cost of operating the OLT port. For 25GS to succeed as the next evolution from XGS, it must support a cost-effective ONU. This goal is currently attainable in most cases.

In 2020, the *Journal of Optical Communications* presented a relative cost comparison between components of the various PON speeds beyond XGS. The article noted that 25GS requires the same components as XGS, with minor upgrades for the higher speed. As a result, there is only a small cost increase for the 2.5x data rate. At speeds higher than 25 Gbps, additional components and major technology changes are required, which increase both the expense and the power requirements of the corresponding ONUs, as shown in Figure 10.




Figure 10 – Relative PON Transceiver Cost in Volume

Note that the specification for 50G/50G was not finalized at the time of this article. While we believe that costs will be significantly higher than for lower speeds, they will likely be less than the 14x shown above.

5.2. Optical Budget

PON speed increases require more optical budget. Going to a 25 Gbps line rate from 10 Gbps causes a 4 dB penalty. Forward error correction, or FEC (which sends redundant data to assist the receiving device with assessing errors) provides a portion of the needed gain. A larger block size and a low-density parity check (LDPC) FEC in 25GS provides additional gain over the Reed Solomon (RS) FEC used in XGS. 25GS FEC requires significantly more logic area in the PON OLT or ONU ASIC. The LDPC makes it difficult or cost prohibitive to use FPGAs for the OLT PON MAC. In ASIC form, the cost difference is not very significant. The 25GS FEC provides roughly 1.5 dB of optical gain over the XGS FEC.

In addition to the gain from FEC, the transmit optical power and receiver sensitivity must recover 2.5dB of the 4dB penalty. The proposed values for transmit optical and receiver sensitivity are near the limits of unamplified optics. For speeds above 25 Gbps in the upstream or downstream, a silicon optical amplifier (SOA) will be required. Adding the SOA can double the cost of the optics and consume an extra one watt of power. While the N1 optical budget doesn't require the SOA at 25 Gbps, the N2 optical budget is still undecided. It is easy to achieve N2 level with the SOA but it also seems very possible to achieve it with proper process control.

5.2.1. Transmission Lasers

PON systems use one of two laser types to transmit data across the fiber line:

- Directly Modulated Laser (DML) creates wider wavelengths with more frequency distortion. They are used for lower speeds and shorter distances.
- Externally Modulated Laser (EML) are temperature controlled with less distortion and narrower wavelengths. They are used for higher speeds and longer distances.

The EML is often 2.5x the cost of the DML. EML also requires up to one watt of additional power for temperature control. For the OLT transmitter, the EML is normally used for all current versions of PON and will likely be the solution for 25GS. As a result, migrating from XGS to 25GS will not result in significant OLT transmitter cost increase if the SOA is not added.

For the ONU transmitter, the DML can be used for GPON and XGS. For 25GS, the requirement is determined by the wavelength. DML can generate the wider UW0 and UW1 wavelengths, but EML is required for the narrower UW3 wavelength. For 50Gbps, EML will be required for all wavelengths.

5.2.2. DSP, High Speed ADC, Coherent

XGS and 25GS don't require a high-speed analog to digital converter and DSP to process the receive signal. Going to 50 Gbps and beyond will require these functions. These functions add significant cost, complexity, and power to the ONU receiver. In some estimates, 2 to 5 watts of additional power will be required to support these functions. Adding Coherent optics allows for higher split ratio, longer reach, and speeds beyond 100 Gbps. Unfortunately, Coherent optics will also increase the cost and power significantly over 25GS.

5.2.3. Cost Summary

With all things considered, 25GS has been estimated to be 1.5 times the cost of XGS in volume. Since the XGS ONU and 25GS ONU don't require an EML transmitter, SOA, or DSP, the increment in cost is reasonable for a 2.5 times speed increase. The 25GS ONU optics could be significantly higher cost if the SOA is required to reach the N2 optical budget and/or the EML transmitter is required to use the UW3 wavelength.

6. Bandwidth Analysis of 25GS

A 25GS delivers less than a full 25 Gbps worth of Ethernet bandwidth across the PON. Framing, forward error correction (FEC), and physical layer management eat into the total bandwidth, see Figure 11

6.1. Downstream Bandwidth Calculation




Figure 11 – XGS and 25GS Downstream Framing

To determine the actual downstream bandwidth, overhead costs will be removed, and variable overhead costs will be estimated. These overhead costs are described below and listed in Table 5:

- Like XGS, 25GS uses non-return-to-zero (NRZ) data without additional bits for line encoding. At a line rate of 24.8832Gbps, 25GS PON is exactly 2.5x the XGS line rate of 9.95328Gbps; no additional bandwidth loss is incurred by the higher rate alone. These starting rates are shown in the first row of the table.
- In both XGS and 25GS, the downstream contains a 24-byte physical synchronization block (PSBd) every 125μs. This block allows for downstream byte alignment, FEC block alignment, and frame alignment (these bytes are not covered by FEC). The PSBd impact is shown in the second row in the table.
- The most significant difference between XGS and 25GS is the new FEC (mentioned in the previous section). In XGS, a FEC block of 248 bytes contains 216 bytes of payload and 32 bytes of parity, reducing the data rate to 8.667648Gbps (efficiency ≈ 87%). In 25GS, a FEC block of 2144 bytes contains 1824 bytes of payload and 320 bytes of parity, reducing the data rate to 21.154304Gbps (efficiency ≈ 85%). The resulting data rates are shown in the third row in the table.
- After the PSBd, the 125μs downstream frame contains the framing sublayer (FS) header and FS trailer. For both XGS and 25GS, the FS trailer is a 32-bit interleaved parity over the FS data excluding FEC parity. The FS header has a minimum size of 4 bytes and a variable size for the physical layer OAM (PLOAM) and the bandwidth map (BWmap) carrying the upstream grants. The resulting data rates are shown in the fourth row in the table.
 - Each PLOAM message is 48 bytes long. For the bandwidth calculations, we assume one PLOAM every 125μs for both XGS and 25GS.

- Each upstream allocation in a BWmap is 8 bytes long. The number of allocations required by the upstream is determined by the polling and maximum burst size configuration.
- ITU-T-based ONUs are managed by ONU management control interface (OMCI) frames. These frames allow for setting up the ONU data path and monitoring the performance. The amount of OMCI required for XGS and 25GS should be the same, averaging two frames per second for each ONU; for a 64-ONU system, this is 128 frames per second. OMCI includes a baseline frame of 48 bytes and a variable extended frame size. Since many ONUs support only the baseline size, we will use this value in the calculations below. Finally, every OMCI frame includes an 8-byte XGEM header, bringing the total size to 56 bytes. The resulting data rates are shown in the 5th row in Table 5.

Table 5 – Downstream Overhead

	XGS (10G/10G)		25GS (25G/25G)	
	Deduction (bps)	Available BW (Gbps)	Deduction (bps)	Available BW (Gbps)
Initial Line Rate	-	9.953280000	-	24.883200000
<i>minus:</i> PSBd, FS Header/Trailer	2,560,000	9.950720000	2,560,000	24.880640000
<i>minus:</i> FEC Parity	1,284,096,000	8.666624000	3,727,360,000	21.153280000
<i>minus:</i> OMCI	57,344	8.666566656	57,344	21.153222656
<i>minus:</i> PLOAM	768	8.666565888	768	21.153221888
<i>minus:</i> BW Map	1,677,517	8.664888371	227,253,073	20.925968815
% User XGEM Data		87%		84%

The ITU-T-based PON system uses XGEM framing to carry Ethernet. The XGEM frame is an 8-byte XGEM header that encapsulates each Ethernet frame from the destination address to the CRC-32. XGEM frames do not include the 12-byte interpacket gap (IPG) or the 8-byte preamble found in the Ethernet frame. XGEM framing allows for segmentation of the Ethernet frames, but only at the 125μs frame boundary in the downstream direction. When comparing the XGEM bandwidth (BW) with the Ethernet Layer 1 BW, the per-packet overhead is 20 bytes (preamble + IPG) for Ethernet and 8 bytes for XGEM. This difference allows for a higher BW to be carried on the PON than on the Ethernet side of the network. The difference can be significant with small frames and less significant with large frames.

Finally, our calculations must include the application layer bandwidth available to customers running speed tests, FTP, etc. To calculate these values, we assume a layer 2 (DA/SA/TYPE/CRC-32) of 18 bytes, an IPv4 header of 20 bytes, and a TCP header of 20 bytes. Despite a drop in data rates from the table above, applications on XGS can still operate at download speeds above 8Gbps while 25GS can still operate at download speeds above 20 Gbps, as shown in Table 6.

Table 6 – Downstream Bandwidth Available

	XGS (10G/10G)	25GS (25G/25G)
XGEM Frame BW	8.659910 Gbps	21.146566 Gbps
L1 Ethernet BW (64B)	10.103228 Gbps	24.670994 Gbps
L1 Ethernet BW (1500B)	8.728822 Gbps	21.314841 Gbps
L7 Application BW (1500B)	8.280895 Gbps	20.381963 Gbps

6.2. Upstream Bandwidth Calculation

The upstream for 25GS uses the same 24.8832Gbps line rate (2.5 times the XGS line rate) as the downstream. The calculation of the possible upstream bandwidth is more complicated than for the downstream because the amount of bandwidth lost is determined by the variable size of the upstream burst. We will start by assuming 64 ONUs on the PON and a simple static granting to see the maximum possible performance. We will then add the overhead required to make the upstream a dynamic operation, so it can be shared on demand. Finally, although the PON can support multiple services in multiple upstream traffic containers or prioritized traffic in a single container, we will assume a single service and the same priority level for each ONU for both our bandwidth and efficiency analyses.

For these analyses, we will consider the three primary impacts on upstream bandwidth: burst overhead, static (fixed) BW allocation, and dynamic BW allocation (DBA).

6.2.1. Burst Overhead



Figure 12 – Upstream Burst Overhead

The upstream burst can be split into 3 parts, (as seen in Figure 12 and Figure 13): dead time between bursts, a preamble at the start of the burst, and the block of data. For ITU-T PON, the dead time and preamble are determined by the configurable parameters of guard time and preamble time. The guard time includes the ONU Laser ON time, the ONU Laser OFF time, and the dead time for any jitter in the upstream slot time. The preamble time includes the time required for the OLTs to perform gain control and clock recovery. The preamble is shown as the physical layer synchronization block for upstream (PSBu). Optical components for 25GS are still in development and they will certainly improve over time. For this analysis, we will assume the same time duration values as XGS. For XGS, a value of 256 bytes will be used for guard time and preamble time combined. For 25GS PON, we will multiply the XGS's 256 bytes by 2.5 to get 640 bytes for both values combined.




Figure 13 – Upstream Framing Overhead

After the PSBu, the FS burst is the FEC protected data. The FS burst starts with the FS header that is 4 bytes long without a physical Layer OAM (PLOAM) or 52 bytes with a PLOAM. The grant in the Bwmap from the OLT selects the inclusion of PLOAM, inclusion of bandwidth request (BWR), and the size of XGEM blocks within the FS burst. Bursts that only carry PLOAM and/or BWR can have a length of 0. The FS burst can contain grants to one or more allocation IDs on the ONU. For our analysis, we will assume a single allocation per burst. The standard allows for a maximum grant size of 9719 upstream slot times. In the case of XGS, upstream slots are defined as 128 bits or 16 bytes. In 25GS, the upstream slots are 2.5 times greater, so they are 320 bits or 40 bytes. In that case, the maximum grant of 9719 upstream slots equals 155,504 bytes in XGS and 388,760 bytes in 25GS. The FS burst ends with the 4-byte FS trailer that contains a parity for checking the data.

6.2.2. Static (Fixed) BW Allocation

With static allocation, the OLT sends a fixed-size grant to each ONU in a simple round robin. This is often called unsolicited granting in DOCSIS. In this analysis, we will assume that all ONUs receive the same fixed allocation.

Physical layer OAM (PLOAM) provides the initial configuration of the PON physical layer and key exchange for encryption. PLOAM in the upstream is infrequent after registration, so required bandwidth consideration is negligible. ONU management control interface (OMCI) is the management traffic from the OLT to the ONU. Data path configuration, firmware download, and statistic gathering are the primary tasks using OMCI. For OMCI frame bandwidth in the upstream, a fixed payload of 56 bytes is granted every 8ms. Discovery slots are 250μs of deadtime for ONUs to register (assuming a 20km PON). In this analysis, a discovery slot is granted every 3 seconds. With the maximum allocation of 9719 upstream slots and the overhead described above, we can calculate the maximum upstream bandwidth for both XGS and 25GS:

- For XGS, a grant of 9719 upstream slots would carry 155,504 bytes of payload data, 8 bytes of FS header/trailer, and 23,040 bytes of FEC parity.
- For 25GS, a grant of 9719 upstream slots would carry 388,760 bytes of payload data, 8 bytes of FS header/trailer, and 68,480 bytes of FEC parity.

It should be noted that while these grants are possible, they are not practical for a 64-ONU system, since the delay between grants would be very long. As an alternative, a static allocation of 20,000 bytes is included in the following table as an example of a viable lower latency configuration.

Table 7 – Static Allocation Analysis for 20km, 64 ONU PON

	XGS (10G/10G)		25GS (25G/25G)	
Line Rate	9.953280 Gbps		24.883200 Gbps	
Static Payload Size	20,000 Bytes	155504 Bytes	20,000 Bytes	388760 Bytes
Minus Preamble, Guard time, FS H/T	(-222,852,025 bps) 9,730,427,975 bps	(-29,233,025 bps) 9,924,046,975 bps	(-1,291,341,739 bps) 23,591,858,264 bps	(-69,785,265 bps) 24,813,414,735 bps
Minus FEC	(-1,256,075,050 bps) 8,474,352,925 bps	(-1,275,622,869 bps) 8,648,424,106 bps	(-3,507,347,934 bps) 20,084,510,330 bps	(-3,687,418,908 bps) 21,125,995,827 bps
Minus OMCI	(-38,912,000 bps) 8,435,440,925 bps	(-38,912,000 bps) 8,609,512,106	(-106,496,000 bps) 19,978,014,330 bps	(-106,496,000 bps) 21,019,499,827 bps
Minus Discovery	(-829,440 bps) 8,434,611,485 bps	(-829,440 bps) 8,608,682,666	(-2,073,600 bps) 19,975,940,730 bps	(-2,073,600 bps) 21,017,426,227 bps
L1 Ethernet BW (1518B packet)	8.507744902 Gbps	8.677271545 Gbps	20.08482188 Gbps	21.09803954 Gbps
L7 Application BW (1518B packet)	8.076272794 Gbps	8.237201857 Gbps	19.066215828 Gbps	20.028047937 Gbps
Latency	1.46 ms	9.46 ms	.762 ms	9.72 ms

The static allocation in Table 7, shows both a very high possible bandwidth (the second column for each rate) and the ability to achieve a lower latency with a smaller burst size consuming lower bandwidth (the first column for each rate). Note that using smaller burst sizes increases the amount of bandwidth required for the combined per-burst overhead (preamble, guard time, FS header and trailer), since smaller bursts require more bursts to be sent, and each burst contains its own overhead.

While static granting is helpful for a simplified overhead analysis, it is impractical for most PONs. For a 64-ONU system with the maximum burst size, the upstream Ethernet bandwidth per ONU is only 135Mbps in XGS and 329Mbps in 25GS. Since operators expect PON networks to meet their requirements for statistical gain, fewer ONUs and different allocations should be used in real-world scenarios to allow for grant sizes to be adjusted accordingly. The static BW analysis above shows the maximum possible upstream bandwidth and cost of the PON overheads.

6.2.3. Dynamic BW Allocation (DBA)




Figure 14 – DBA Upstream Bursts

Adjusting the upstream slot according to subscriber needs is facilitated with DBA in a request/grant methodology. The OLT sends a dynamic bandwidth resource unit (DBRu) in the upstream slot. The returning DBRu consumes 4 bytes in the upstream burst and indicates the amount of packet data remaining in the queue. If the ONU returns the DBRu with a value of 0 (indicating no data in the queue),

no grant will be sent from the OLT. Instead, the OLT will send another DBRu to the ONU after a fixed time interval. This process is known as polling.

If the ONU returns the DBRu with a non-zero value, the OLT then sends a variable-size grant to the ONU based on the reported data payload need. The data payload sent upstream from the ONU includes a DBRu reporting additional payload waiting. The granting of a follow-up DBRu with the data payload burst is known as piggyback granting. When a grant is piggybacked, polling is not needed.

Calculating the available upstream bandwidth for a DBA system is more complicated than for a static allocation system since there are endless possible burst sizes from different traffic scenarios and different SLAs. For this analysis, we will establish some basic parameters to model the most complex scenario, as follows:

- We will consider request/grant or status-reporting DBAs only, and not non-status reporting or predictive grants.
- All ONUs will be assigned the same SLA, with the goals of low latency and higher bandwidth for the entire PON.
- The status reporting methodologies used will be polling and piggybacking.

In a DBA PON system, the amount of overhead is determined by the number and size of the bursts. We will examine the extreme conditions to calculate the bandwidth and latency of the PON upstream in a PON of 64 ONUs, considering two scenarios:

- When a single ONU is requesting the maximum rate, it will use piggybacking, while the other 63 ONUs will use polling. The bandwidth for polling the 63 ONUs will decrease the amount of bandwidth available for the single piggybacking ONU. In this scenario, the key factor is the polling interval: a shorter polling interval will have lower latency for the 63 polling ONUs but consume more upstream bandwidth from the single piggybacking ONU, while a longer polling interval will increase latency for the 63 polling ONUs but save upstream bandwidth for the single piggybacking ONU.
- When all 64 ONUs are actively sending payload data and requesting the maximum rate, they will all use piggybacking exclusively. Since all ONUs are treated equally, the DBA will schedule the ONUs in a round robin based on the last time data was granted. The size of the round robin is the number of active ONUs times the maximum size of the burst. In this scenario, the key factor is the maximum size of the burst: a larger burst will be more efficient for the ONU that is sending its payload, but it will create a longer delay for the other ONUs.

Table 8 – Dynamic Bandwidth Allocation Analysis for 20km, 64 ONU PON

	XGS (10G/10G)		25GS (25G/25G)	
Line Rate	9.953280Gbps		24.883200Gbps	
Polling Interval	3.5 ms		3.5 ms	
Max PON Burst Size	60,000 Bytes		140,000 Bytes	
ONU Activity	1 active and 63 idle	64 active	1 active and 63 idle	64 active
Payload Burst Overhead	75,319,749 bps	75,920,301 bps	191,775,180 bps	193,589,029 bps
Payload FEC Overhead	1,259,482,114 bps	1,269,524,418 bps	3,634,877,255 bps	3,669,256,667 bps
Polling Overhead	80,064,000 bps	0 bps	232,128,000 bps	0 bps
L1 Ethernet BW (1500B packet)	8,539,654,048 bps	8,607,743,778 bps	20,761,985,282 bps	20,958,356,379 bps
L7 Application BW (1500B packet)	8,122,924,176 bps	8,187,691,169 bps	19,748,813,154 bps	19,935,601,462 bps
Latency	4.125 ms	3.95 ms	4.125 ms	3.80 ms

Table 8 above shows a practical configuration for XGS applied to 25GS. This configuration allows an XGS system to have a ~4ms maximum upstream latency. In both XGS and 25GS, polling inactive ONUs becomes significant when fewer ONUs need to use the upstream. The polling waste is reduced to 0 when all ONUs are active. XGS shows 8.5 Gbps of Ethernet bandwidth 8.1 Gbps of application bandwidth while 25GS achieves 20 Gbps of Ethernet bandwidth and just under 20 Gbps of application bandwidth. While the goal of 2.5 times the speed of XGS is not reached, the results are very close.

Main Port Traffic Statistics										
Name	TX L1 (%)	TX L1 (bit/s)	TX L2 (bit/s)	TX (pps)	TX (bytes)	TX (packets)	RX L1 (%)	RX L1 (bit/s)	RX L2 (bit/s)	RX (pps)
P-0-0-0	89.600	8,959,986,780	8,842,092,220	736,841	39,254,169,000	26,169,446	86.120	8,612,033,640	8,498,717,480	708,226

Figure 15 – DBA Lab Results for 64 active ONUs

Name	TX L1 (%)	TX L1 (bit/s)	TX L2 (bit/s)	TX (pps)	TX (bytes)	TX (packets)	RX L1 (%)	RX L1 (bit/s)	RX L2 (bit/s)
P-0-0-0	90.000	8,999,997,660	8,881,576,700	740,131	23,786,878,500	15,857,919	85.059	8,505,864,680	8,393,945,480
P-0-0-1	90.000	8,999,990,280	8,881,569,320	740,131	23,848,053,000	15,898,702	87.329	8,732,881,510	8,617,672,070

Figure 16 – DBA Lab Results for 1 active ONU on 64 ONU system

A lab test of a 64-ONU system in XGS with the sample configuration shown in the table able was performed to validate the model. With 64 ONUs transmitting 9 Gbps upstream, the “RX L1” of 8.612 Gbps received is very close to model’s prediction of 8.607 Gbps, as seen in Figure 15,. With a single ONU transmitting, the 8.505 Gbps “RX L1” on P-0-0-0 closely matches the model’s prediction of 8.539 Gbps, as seen in Figure 16. The downstream performance from the lab is also available on P-0-0-1 “RX L1”. The lab shows 8.733 Gbps which is very close to the downstream Ethernet L1 model’s predication of 8.729 Gbps.

7. Latency Analysis of 25GS

Latency has become a hot topic in the industry. Residential subscribers are looking for low latency for gaming and interactive experiences such as the metaverse. In this paper, we will look at the one-way packet latency between the OLT NNI port and ONU Ethernet UNI port. These delays assume layer 2 switching at the OLT and ONU. Routing devices, Wi-Fi interfaces, etc. will add more latency. This analysis focuses on worst case scenarios with a fully loaded PON. In most cases, the customers will see much better latencies due to a lower take rate or activity.

7.1. Downstream Latency

The downstream latency for PON is very low, and largely based on the functional characteristics of the network switches and the length of the PON fiber. On a 20km fiber, the flight delay is 100μs for both XGS and 25GS. The FEC block for 25GS is 2144 bytes versus 248 bytes for XGS, so the 25GS decoder will require more time for processing. However, the FEC decoder time is still minimal for 25GS, at approximately 1μs. This is a small value compared to the store and forward delays in the OLT and ONU switching. Finally, equipment delays are normally less than 50μs, so the worst-case delay on a 20km fiber is 150μs. This calculation also applies to both XGS and 25GS.

7.2. Upstream Latency

If operators don't have demand for the 20 Gbps of upstream traffic possible in 25GS, they can trade off the upstream bandwidth for a lower latency upstream. Since TCP/IP traffic downstream requires an upstream acknowledge, a lower latency upstream can also improve downstream throughput. By decreasing the interval for polling and reducing the maximum burst size, the latency for all ONUs in a 64-ONU PON could be reduced significantly. The inaccuracy of predictive granting, inflexibility of fixed granting, and other less predictable techniques can be avoided. The ability to lower the latency provides an opportunity to increase the downstream performance and throughput. Since the large downstream TCP/IP bursts require acknowledging in the upstream direction, minimizing the latency for these frames increases the downstream throughput and overall latency. Upstream latency on a PON can be broken down into 3 areas: start latency, continuation latency, and queue delay.

7.2.1. Start Latency

Whenever the upstream packet stream gives a DBRu of 0 to the DBA, it is considered idle and won't be granted until the next polling cycle. Since data arrives randomly compared to the polling cycle, the maximum wait time to be sampled is the polling interval. After the polling interval, the DBRu must be sent to the OLT/DBA to be granted. This transmit time can be 0μs for a 0 km distance ONU and 100μs for an ONU at the end of the fiber. An idle ONU should be at the front of the round robin so it will be granted quickly. The DBA will have a delay to issue the grant. This delay can be 125μs waiting for the start of the downstream framing or additional delay for software processing. In a software DBA, the cycle time is often used to define this time interval. In the example below, a hardware DBA is assumed that only waits for the 125μs downstream frame boundary. After the DBA issues the grant, the grant must travel to the ONU and back up the PON. This time is often referred to as the PON round trip time. On a 20km PON fiber system, it is 250μs. Because ranging an ONU delays the transmitter for closer ONUs, this delay is constant regardless of the ONUs position. The PON round trip time also sets the size of a discovery window to the same 250μs. If the polling or data grant to the ONU is needed after a discovery window has been requested, an additional 250μs of delay/jitter is possible. A small delay for the ONU and OLT hardware should be included as well. While not always the case, it will be considered a fixed delay in this analysis. When testing in the lab, the fixed delay will show up as the min delay on a long test and variable delay can be determined by subtracting the max delay from the min delay.

$$\text{Start_Up_Fixed_Delay} = \text{Upstream_Flight_Time} + \text{PON_Round_Trip_Time} + \text{ONU_HW_Delay} + \text{OLT_HW_Delay}$$

$$\text{Start_Up_Variable_Delay} = \text{Polling_Interval} + \text{DBA_Delay} + \text{Discovery_Window}$$

$$\text{Start_Up_Max_Delay} = \text{Start_Up_Fixed_Delay} + \text{Start_Up_Variable_Delay}$$

The start latency is a big factor in downstream TCP/IP performance. The upstream acknowledge frames are often spread further apart than the polling interval so they will see the start latency as the dominant

factor. Large bursts upstream will hit the start latency for the first packets but the delay for the tail of the burst determines the true latency of the transaction, so the start latency might not be the key factor in large upstream bursts.




Figure 17 – Upstream Start Latency vs Polling

The polling interval is the only factor in the start latency that can be easily modified. The fiber length, number of ONUs, and hardware delays are fixed inputs. Based on 20km fiber and estimates for the hardware delay, the start delay is $625\mu\text{s}$ plus the polling interval. Figure 17 shows the direct relationship between the maximum start latency and the polling interval. Both XGS and 25GS have the same start latency equation since the data rate is not a factor.




Figure 18 – Polling and the Single Transmitting ONU

Decreasing the polling interval can significantly reduce the bandwidth. Figure 18 shows the scenario with the largest polling penalty. 63 ONUs are idle and 1 ONU is requiring the full bandwidth. In this case, the polling grants are non-traffic carrying blocks of time that limit the bandwidth to the single ONU transmitting. A shorter interval limits the bandwidth. When more ONUs are transmitting, the bandwidth lost to polling idle ONUs decreases but it is minimal until a large percentage of ONUs are active.




Figure 19 – XGS Upstream Bandwidth versus Polling Interval




Figure 20 – 25GS Upstream Bandwidth versus Polling Interval

7.2.2. Continuation Latency




Figure 21 – Piggybacking and 64 ONUs transmitting

After an ONU has received the first grant of a burst, piggybacking will allow the OLT to accurately grant the remaining packets, see Figure 21. In this case, the latency is determined by the number of ONUs actively in the round robin at that time. In the best case, it is a single ONU. In the worst case, it is all ONUs requesting maximum burst sizes at the same time. In this case, the maximum latency is dominated by the number of ONUs and the maximum burst size allowed. Small bursts have a higher percentage of overhead to data (less efficient) with lower latency while large bursts have a lower percentage of overhead to data (more efficient) with greater latency. The continuation latency has similar equations as the start latency. The DBRu must travel upstream and the grant must traverse the entire PON. The big difference is the Round_Robin_Time that replaces the Polling_Interval and DBA_Delay. The DBA_Delay is often absorbed since the grant is known well before the Round_Robin_Time is available.

$$Cont_Fixed_Delay = Upstream_Flight_Time + PON_Round_Trip_Time + ONU_HW_Delay + OLT_HW_Delay$$

$$Cont_Variable_Delay = Round_Robin_Time + Discovery_Window$$




Figure 22 – 64 Active ONU Latency




Figure 23 – 64 Active ONU Bandwidth

7.2.3. *Queuing Delay*

Queuing delay occurs when the input bandwidth exceeds the upstream bandwidth available. In the PON upstream case, it can occur when the subscriber exceeds their SLA or congestion limits the bandwidth available to a smaller number. Queuing delay performs an important function in TCP/IP networks. The delay on the acknowledge frame causes the next block of data to be delayed and thus slowed to a lower rate. Alternatively, a queue that overflows will cause a timeout and retransmit of a data block. Dropping frames has a much more significant impact to the user than the queuing delay. The queues in the ONUs should be sized to absorb large bursts of data and avoid drops. Priority queues can be used to allow higher priority traffic to avoid congestion. It is important to downstream performance to have a stable and low latency upstream. Variable upstream delays can cause the downstream bursts to be delayed and thus lower the throughput. By configuring a priority queue in the upstream for small acknowledge frames, it is possible to minimize the latency jitter to the downstream bursts. Queuing delay can be calculated by determining the difference in data rate between the input and output. If the difference in data rate is multiplied by the queue size, the maximum queuing delay before discard can be determined. The latency in this paper focuses on flows that have adjusted to the data rate and thus are below the maximum capacity of the SLA or total capacity. Since the queuing delay is a short-term factor required for bandwidth adjustment, it isn't something that could or should be removed from the system. A low max delay and low jitter delay when the subscriber is below the maximum capacity is the focal point of this analysis.

7.2.4. *Balanced Delay*

If the maximum start latency and the continuation latency are roughly the same, a single maximum latency covers both traffic conditions. Balancing the two delays to a common value allows for a

consistent performance without wasting efficiency. For upstream traffic, the balancing of start latency and continuation latency is a good goal. To achieve a balanced maximum latency under 1ms in XGS or 25GS, the polling rate and maximum burst must be adjusted to smaller values. The amount of upstream bandwidth available is variable based on the number of ONUs bursting upstream. The table shows the lower of the two values for comparison.

Table 9 – Sub Millisecond Delay for 20km, 64 ONU PON

	XGS	25GS
Up Max Burst Size	9500 Bytes	22000 Bytes
Polling Interval	300 us	300 us
Start Latency	0.925 ms	0.925 ms
Continuation Latency	0.975 ms	0.975 ms
L1 Ethernet BW (1518B packet)	7.346 Gbps	16.739 Gbps
L7 Application BW (1518B packet)	6.992 Gbps	15.932 Gbps

Table 9 shows that reaching sub millisecond max latency for 64 ONUs is possible in both XGS and 25GS. In both cases, a significant amount of upstream data is used to guarantee the lower latency. In XGS, almost 7 Gbps of application bandwidth remains while 25GS has almost 16 Gbps of upstream application data. The 25GS can still provide a 10Gbps symmetric commercial service along with the more asymmetric residential service.

7.2.5. XGS and 25GS Sharing the Fiber

In the earlier scenarios, a PON was either 64 XGS ONUs or 64 25GS ONUs. In many deployment scenarios, XGS and 25GS will exist side by side on the same fiber. In this case, it is possible for an operator to selectively move subscribers to a higher speed PON. If subscribers purchasing 5 Gbps, 10 Gbps, or higher SLAs were moved from XGS to 25GS, the number of subscribers on 25GS would be limited and the need for very high upstream bandwidth on XGS would be lessened. For example, the number of high bandwidth SLAs on the PON could be limited to 8 or 16 subscribers. Therefore, the maximum number of 25GS ONUs would be 8 or 16. In this case, the penalty for achieving the sub 1 millisecond latency on 25GS is much less. Table 10 shows an example of up to 8 25GS ONUs on a PON. In this case, the burst size can be significantly increased since the round robin will be only 8 ONUs. Additionally, the number of idle ONUs polling at 300µs is limited to 8 on the 25GS PON. With these two factors, the 25GS can achieve a 20 Gbps upstream Ethernet bandwidth with a sub millisecond worst case delay. With a total upstream bandwidth of 28 Gbps, the combination of XGS and 25GS could have a long-term future in the market. If the number of 25GS ONUs is increased to 16, the upstream drops by 400 Mbps but still stays above 20 Gbps of Ethernet BW, as shown in Table 11.

Table 10 – Shared PON with XGS (up to 64 ONUs) and 25GS (up to 8 ONUs)

	XGS	25GS	Total
Up Max PON Burst Size	9500 Bytes	180000 Bytes	
Polling Interval	300 µs	300 µs	
Start Latency	0.925 ms	0.925 ms	
Continuation Latency	0.975 ms	0.924 ms	

	XGS	25GS	Total
L1 Ethernet BW (1518B packet)	7.347 Gbps	20.818 Gbps	28.165 Gbps
L7 Application BW (1518B packet)	6.992 Gbps	19.802 Gbps	26.794 Gbps

Table 11 – Shared PON with XGS (up to 64 ONUs) and 25GS (up to 16 ONUs)

	XGS	25GS	Total
Up Max PON Burst Size	9500 Bytes	100000 Bytes	
Polling Interval	300 μ s	300 μ s	
Start Latency	0.925 ms	0.925 ms	
Continuation Latency	0.975 ms	0.988 ms	
L1 Ethernet BW (1518B packet)	7.347 Gbps	20.417 Gbps	27.764 Gbps
L7 Application BW (1518B packet)	6.992 Gbps	19.421 Gbps	26.413 Gbps

8. Conclusion

The IEEE 802.3 and a 25GS MSA group of 50+ companies standardized a 25 Gbps symmetric speed for PON access. 25 Gbps is the last PON speed that doesn't require a DSP, SOA, or EML at the ONU so it can be cost effective and low power. 25 Gbps is a useful speed for that reason. 50 Gbps and 100 Gbps will be available in the future at a higher cost and power. The 25GS standard is essentially 2.5 times the speed of the ITU-T XGS standard with the LDPC FEC defined by the IEEE 802.3. 25GS can be mixed with GPON or XGS on the same fiber plant allowing for a simple upgrade path. Equipment vendors offer simple upgrade paths for 25GS and a way to co-exist in the same box. 25GS allows operators to offer 5 Gbps and 10 Gbps symmetric services to customers. With all overhead considered, operators can expect to get approximately 20 Gbps in the downstream or upstream application bandwidth with 25GS. In addition to higher speed tiers, operators may choose to use the additional upstream bandwidth to lower the upstream latency. It is possible to achieve sub millisecond worst-case latency for 64 subscribers on the PON. By mixing XGS and 25GS on the PON, it is possible to achieve sub millisecond latency and 28 Gbps of upstream bandwidth. Based on the cost, ease of upgrade, simplicity, and additional bandwidth, 25 Gbps PON will provide value to operators looking for higher bandwidth and lower latency.

Abbreviations

25GS	25 Gbps symmetric PON defined by 25GS MSA
bps	bits per second
FEC	forward error correction
Gbps	1,000,000,000 bits per second
GPON	ITU-T Gigabit Passive Optical Network
IEEE	Institute of Electrical and Electronics Engineers
ITU-T	International Telecommunication Union Telecommunication
LDPC	Low-density parity-check
MSA	Multi-source agreement

N1	XGS/25GS 29 dB loss budget
N2	XGS/25GS 31 dB loss budget
OAM	Operation Administration Maintenance
OLT	Optical Line Terminal. Carrier side PON device
OMCI	ONT Management and Control Interface
ONT/ONU	Optical Network Terminal/Unit. Subscriber side PON device
PLOAM	Physical Layer OAM
PON	Passive Optical Network
RS	Reed Solomon
SCTE	Society of Cable Telecommunications Engineers
XGS	ITU-T 10 Gbps symmetric PON

Bibliography & References

25GS-PON Specification: 25 Gigabit Symmetric Passive Optical Network, 10 August 2021, Version 2.0

Institute of Electrical and Electronics Engineers (IEEE) standards:

- 802.3ca-2020: IEEE Standard for Ethernet Amendment 9: Physical Layer Specifications and Management Parameters for 25Gb/s and 50Gb/s Passive Optical Networks

International Telecommunications Union – Telecommunications Sector (ITU-T) Standards:

- G.652: Characteristics of a single-mode optical fibre and cable,
- G.984.2: Gigabit-capable Passive Optical Networks (G-PON): Physical Media Dependent (PMD) layer specification
- G.984.5: Gigabit-capable passive optical networks (G-PON): Enhancement band
- G.9701: Fast access to subscriber terminals (G.fast) - Physical layer specification
- G.9804.3: 50-Gigabit-capable passive optical networks (50G-PON): Physical media dependent (PMD) layer specification
- G.9807.1: 10-Gigabit-capable symmetric passive optical network (XGS-PON)

Journal of Optical Communications and Networking, Sept 2020; IEEE/Optica Publishing Group

Bitcode Obfuscation

Protecting Software Without Source Code Access

A Technical Paper prepared for SCTE by

Rafie Shamsaasef

Director of Software Engineering
CommScope
6450 Sequence Dr, San Diego, CA 92121
1 (858) 404-2205
rafie.shamsaasef@commscope.com

Lex Aaron Anderson

Senior Security Architect
CommScope
PO Box 37-942 Parnell 1052, Auckland, New Zealand
+64 935 803 75
aaron.anderson@commscope.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Introduction to Software Obfuscation.....	3
3. Control-flow and Data-flow Obfuscation	5
3.1. Control Flow Obfuscation.....	5
3.2. Data Flow Obfuscation.....	5
4. Bitcode obfuscation techniques	5
4.1. What is bitcode?.....	6
4.2. Why bitcode protection?.....	6
4.3. How does it work?.....	7
5. Exploring the usage of bitcode obfuscation	7
5.1. Mobile applications.....	8
5.2. Cloud Server Applications and Web Services	10
5.3. IoT Applications.....	10
6. Conclusion.....	11
Abbreviations	12
Bibliography	13

List of Figures

Title	Page Number
Figure 1 - Encryption vs Obfuscation.....	4
Figure 2 - LLVM Architecture	6
Figure 3 - Cloud-based bitcode obfuscation use-case.	7
Figure 4 - iOS Mobile App Framework.....	8
Figure 5 - iOS Mobile App Framework with App Obfuscation	9
Figure 6 - Sample Cloud App with Obfuscation	10

1. Introduction

Software obfuscation techniques have become increasingly popular in recent decades due to their broad applicability toward malware threats and intellectual property protection. Reverse engineering and tampering attacks are prominent means for software piracy and exploitation. Obfuscation is a collection of techniques for securing software and protecting applications from harmful malware. The goal of these techniques is to increase the cost and feasibility for attackers to exploit security vulnerabilities and carry out successful attacks against software implementations.

Bitcode obfuscation is a form of obfuscation that operates on an intermediate code representation rather than the original source code or native binary. This technique is based on sound and proven mathematical principles that retain the security characteristics of native binary while also maintaining the portability and platform independence of source code. Obfuscating at the bitcode level enables developers to utilize obfuscation-as-a-service without exposing their source code or proprietary libraries, while achieving levels of protection not possible with source-code obfuscation. As such, bitcode obfuscation is a promising field for developing the next generation of software protection services in the cloud.

In this paper, we offer technical details of control-flow and data-flow obfuscation techniques based on the idea that no source code or other dependencies are required to apply strong obfuscation directly to the intermediate bitcode representation rather than the original source code or the native binaries. We conclude by providing insights into the usage of obfuscation in relation to the security requirements of connected and cloud software systems.

2. Introduction to Software Obfuscation

In software development, obfuscation is the act of generating source or machine code that is difficult for humans and automated tools to reverse engineer. The goal is to achieve maximally unintelligible code without introducing unacceptable levels of overhead. Many techniques have been described in the literature that have both heuristic and tractable security basis. A combination of techniques can provide synergy in terms of the work factor required to reverse engineer the resulting binary code.

Programs often use high-level programming language constructs, common design patterns, and reusable components. These software engineering practices make code easier to reverse engineer and exploit (Wikipedia, 2022). It is therefore not appropriate (nor sufficient) for programmers to deliberately obfuscate their code to attempt conceal its purpose (security through obscurity), since many advanced tools exist to deobfuscate code that has not been sufficiently protected (OWASP, 2016).

Obfuscation is not the same as encryption. There is a common misconception to think of obfuscated data or code being the same as encrypted ones. While they both manipulate the original data/code to a different form, they are fundamentally different processes. Encryption requires a key and deploys a well-known cipher algorithm such as AES to convert the data or code to an unreadable form; where a decryption process is needed to perform reverse operation and get back the original data/code. Obfuscation on the other hand utilizes algorithms which often do not require a key. The obfuscated data and code can be used as-is without a need to de-obfuscate them. In-fact the goal is to make the code hard to de-obfuscate while retaining the original purpose and minimizing overhead as much as possible.

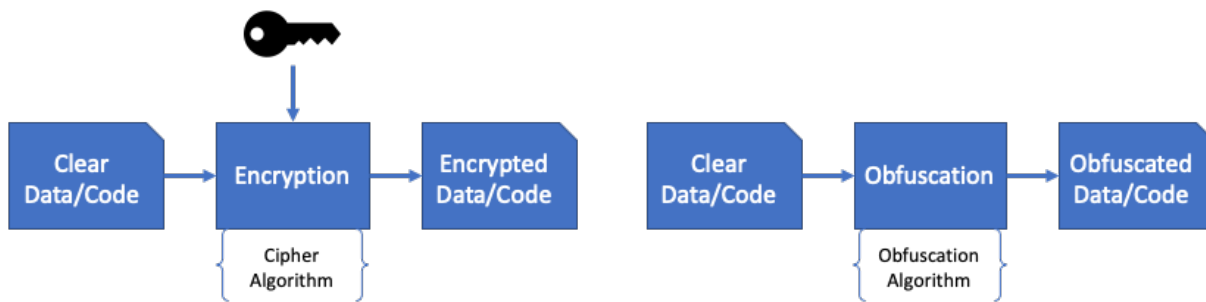


Figure 1 - Encryption vs Obfuscation

When it comes to source code or binary object encryption versus obfuscation, the encrypted object must be decrypted prior to execution rather the obfuscated object can be executed while remaining in the obscured form. It is possible for obfuscated elements to be captured and de-obfuscated, but the obfuscation presents a barrier to understanding and analysis (Arini Balakrishnan, 2005).

Several notions of obfuscation were shown in (Barak, et al., 2001) and (Balakrishnan & Schulze, 2005). The main result is that a strong notion of obfuscation cannot always be achieved. Over a decade later, Garg, Gentry, and Halevi (Garg, Gentry, & Halevi, 2013) gave the first candidate construction of an efficient general-purpose indistinguishability obfuscator, called multilinear jigsaw puzzles. While indistinguishability obfuscation promises secure general-purpose obfuscation, it remains an open question as to whether any practical real-world indistinguishability obfuscators can be implemented under this new model.

There are many mature obfuscation techniques offering a wide range of protections based on heuristic approaches to manipulate source code and even binary without altering the logic or the purpose of the code. The strength and weakness of obfuscation approaches can only be determined by effectiveness of the attacking and exploitation mechanism. Tools such as disassemblers and decompilers are often deployed to perform reverse engineering attacks by extracting sensitive information, adding malicious code and application cloning (Barak, et al., 2001). Although code obfuscation can thwart many attacks, given enough time and effort any of these techniques can be eventually overcome by reverse engineering process. Regardless, programs are obfuscated every day in the real world without any provable security guarantee. Nevertheless, obfuscation and diversification remain viable protection tools from the practical perspective (Garg, Gentry, & Halevi, 2013) (Goldwasser & Rothblum, 2007).

While obfuscation can be applied to programs written in any language, they are more effective for programs that are compiled to binary without the need for a virtual machine (VM). There are limited techniques that can be considered to obscure high level languages (i.e. Java and C#) instructions and control flow such that are heavily relying on their underlying VMs. Unlike C/C++, decompilation of Java programs is a much simpler task and therefore may be fully automated by attackers. Class hierarchy, high-level statements, names of classes, methods and fields can be retrieved from class files emitted by the standard javac compiler. Every current obfuscation product is easily circumvented by off-the-shelf de-obfuscation tools for java. The challenge gets even worse when trying to obfuscate a Python program. For Python, it basically comes down to renaming and hiding some of the instructions that can even be easily de-obfuscated (OWASP, 2016).

3. Control-flow and Data-flow Obfuscation

Obfuscation protects an application from reverse engineering, protecting proprietary source code, intellectual property and to increase the difficulty of exploitation. Obfuscation can be split into two main types: Control-flow and data-flow obfuscation. Together, these can provide a synergy that is similar in nature to how individually weak confusion and diffusion components when used together in cryptography lead to strong ciphers like AES (Anderson L. , 2015) (Worldwide Patent No. WO2018106439A1, 2018).

3.1. Control Flow Obfuscation

The aim of control-flow obfuscation is to make a program's execution difficult for an attacker to understand and hence reverse-engineer. Typical control-flow obfuscation methods include:

- **Instruction substitution** replaces assembly-level operations with randomly chosen code blocks that perform the same operation in different ways. The aim is to add resilience against both static analysis and dynamic analysis of the program code as well as automated attacks that look for code fingerprints.
- **Bogus-control-flow** obfuscation modifies a program's control-flow by adding entry points that evaluate complex expressions to determine the outcome of conditional jumps: either to jump to valid program code or to randomly altered "junk" code blocks.
- **Control-flow-flattening** rearranges a program's basic blocks in a randomized manner to give a program a uniformly random structure, where the original program's control-flow is only able to be re-established by the runtime computation of a control variable representing the state of the program.
- **Virtual-machine interpreter obfuscation** incorporates known hard mathematical problems in the computation of the control-flow to further increase the cost of reverse-engineering attacks.

3.2. Data Flow Obfuscation

Data-flow obfuscation is about randomizing the instructions that compute logical and arithmetic operations in a program (Worldwide Patent No. WO2018106439A1, 2018). Data-flow obfuscation methods include:

- **Randomized branch encoding** involves representing data-related logical and mathematical operations as a branching program composed of a sequence of permutations. Sequences of these branching programs are then concatenated together. When these programs are converted back to machine code, the result is uniformly randomized and unintelligible code that bears no resemblance to the original algorithm.
- **Randomized input, output encodings** can be used to make the obfuscated code even harder to reverse-engineer, as well as protecting constants and allowing seamless secure chaining to and from the obfuscate application.

4. Bitcode obfuscation techniques

Obfuscation (and other protection methods) can be applied directly to bitcode files. This enables application protection to be applied via a third-party service, and as such reduces the need for in-house tools and expertise in application security.

4.1. What is bitcode?

Bitcode is a platform-independent, universal low-level intermediate representation (IR) used by LLVM compilers and tools, such as Clang, XCode, Microsoft Clang-CL, Objective C/C++ Swift, GO, EM-Scripten, Rust, and many others. **Bitcode is the file format for LLVM IR** (LLVM Documents, 2003).

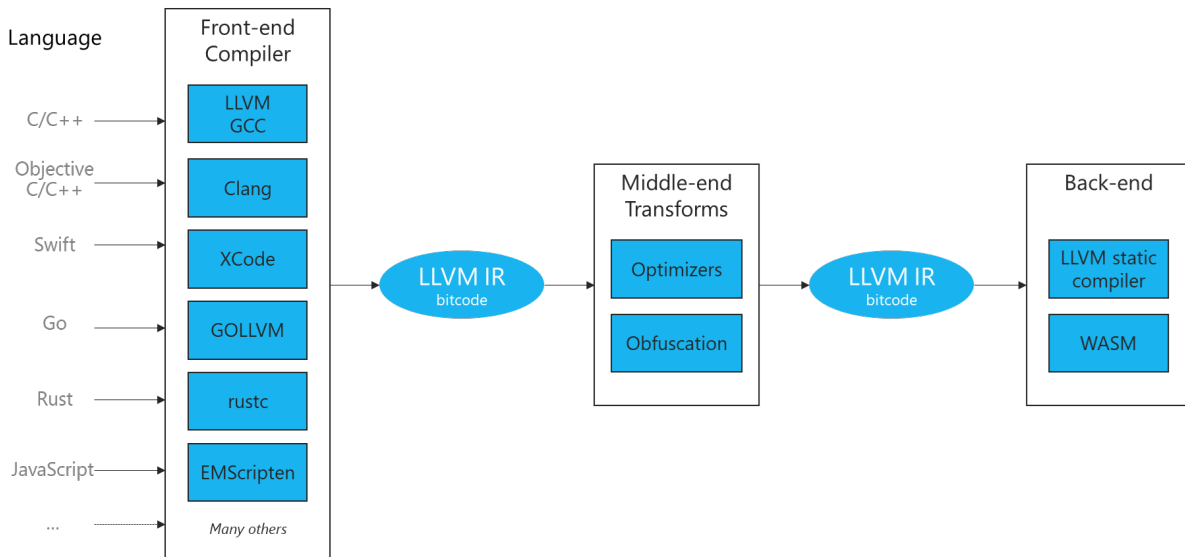


Figure 2 - LLVM Architecture

- Front-end compilers compile source languages to LLVM intermediate representation.
- The LLVM IR can then be optimized and transformed by middle-end tools that are front-end and source language agnostic. In addition to optimization, these transformations can include obfuscation and other protection techniques.
- Finally, back-end tools convert the LLVM IR to native machine language for any of the many supported platforms and architectures. These back-end tools are agnostic to both the front-end and the middle-end layer, thus facilitating broad platform and language independence of the LLVM architecture.

4.2. Why bitcode protection?

Developers are naturally focused on developing and delivering their features and services to their customers. As a result, security is often neglected to varying degrees at both management and operational levels. Security policies may be insufficient to identify unprotected, weak, or otherwise exploitable code, and under-tuned security parameters. Even if good policies are in place, expertise and resources required to manage and implement security may be insufficient to ensure adequate security protections are in place.

A well-designed cloud-based bitcode obfuscation solution can incorporate security policies and parameters to bridge the gap between in-house expertise and practices and robust security implementations.

- The cloud-based service would allow the definition of security policies and would then verify that these policies are correctly implemented during the application of bitcode obfuscation.

- Audit reporting of security coverage and any identified weaknesses can be provided to appropriate parties in the organization.
- Since the compilation of source code is done prior to the application of bitcode protection; and linking is done afterwards, confidential source code and proprietary libraries do not need to be exposed to the middle-end when applying bitcode protection to an application.
- Bitcode is platform and target agnostic, therefore protection can be applied across all LLVM supported source languages and target architectures. This will allow developers to set global security goals or customize security per target platform.
- Alternatives, such as source-code obfuscation methods can be quickly broken and circumvented with readily available de-obfuscation tools (GitHub, 2022). Bitcode obfuscation is not vulnerable to these same attacks.

4.3. How does it work?

Obfuscation is applied directly to bitcode files via a LLVM middle-end according to a set of protection parameters. No source code, header files, libraries or other dependencies are required in order to obfuscate the bitcode. In addition to the protected bitcode, the middle-end can generate audit reports and logs to assist with the monitoring, implementation, and management of the protected code.

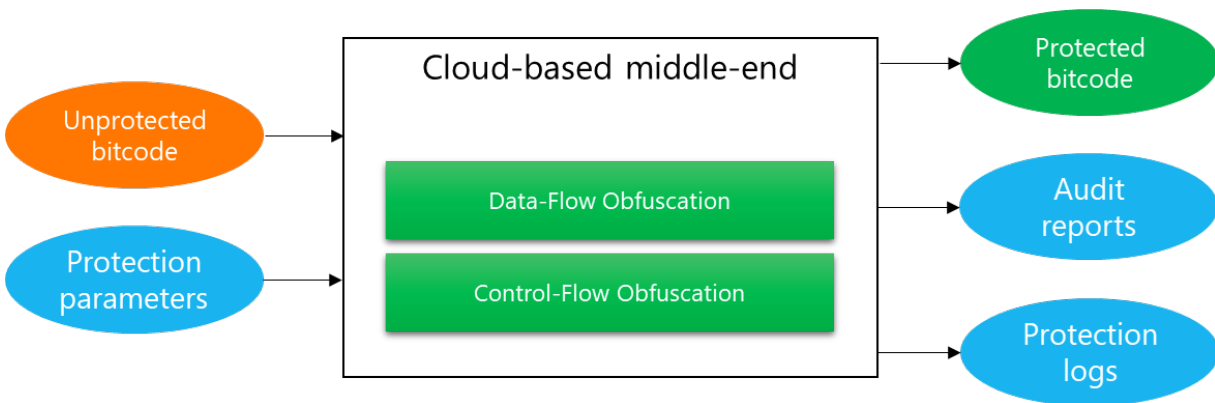


Figure 3 - Cloud-based bitcode obfuscation use-case.

- Unprotected bitcode is sent to a cloud-based middle-end that can apply dataflow and control-flow obfuscation to the bitcode according to a set of supplied protection parameters.
- The middle-end pass manager executes a series of passes to apply the control and dataflow obfuscating transforms. These passes are interleaved with each other and with optimization passes to ensure that the obfuscation complexity matches the tuning requirements specified in the protection parameters.
- The higher the complexity, the higher the runtime overhead, thus the importance of enabling developers to tune the amount of obfuscation applied by the middle-end. Further tuning is possible via LLVM function and inline attributes (LLVM Project, 2022), which can be embedded in the source code and remain readable from the bitcode by the middle-end transform passes.

5. Exploring the usage of bitcode obfuscation

Bitcode obfuscation offers strong protection against tampering and reverse-engineering attacks for software programs running in non-secure environments. These techniques allow tunability to achieve a balance between security and performance; and can be applied to a wide range of applications targeted to

various platforms and devices. In this chapter we explore how obfuscation can be utilized to protect specific types of applications in different industries.

5.1. Mobile applications

Mobile applications (running on iOS or Android devices) often process customer credentials and other sensitive data, while containing differentiated business logic. With stiff competition in the mobile application space as well as an increasing background security of exploits and attacks, it is crucial to protect applications against reverse-engineering by adversaries seeking to gain proprietary knowledge of the app. It is additionally important to prevent adversaries seeking to circumvent authorization and authentication, and to protect against the extraction of sensitive and confidential information.

Secret business logic embedded in the app's source code is considered intellectual property of the app owner and should not be exposed. Generally, all mobile code is susceptible to reverse engineering according to OWASP (OWASP, 2016). Even though security of iOS and Android mobile platforms has always been improving, access to low-level security components is not always available to apps. An attacker will typically download the targeted app from an app store and analyze it within their own local environment using a suite of different tools to statically analyze the code/binary and perform reverse engineering attack.

Let's take an iOS app for example and examine its structure to understand how obfuscation can protect it. The language of choice to develop an app for iOS devices is either Objective C or Swift. They are both compiled into LLVM bitcode binary suitable for Apple App Store submission. The following picture shows the application framework and underlying components of iOS platform (Lucideus, 2019).

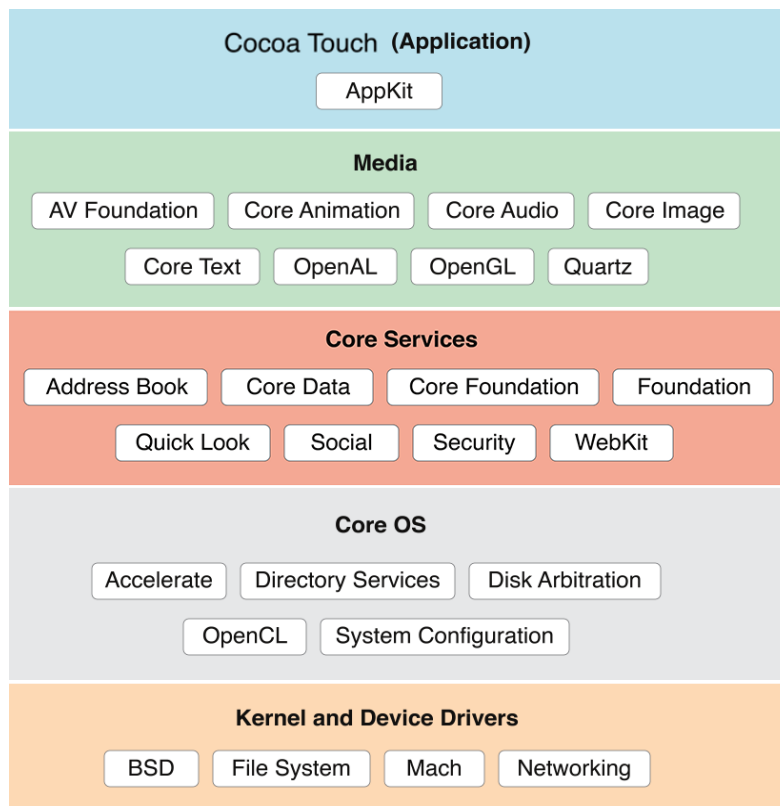


Figure 4 - iOS Mobile App Framework

The app utilizes iOS SDK APIs to access certain functionalities performed by low level libraries while containing a series of business logic and confidential codes within the app layer. App developers are essentially creating complicated puzzles of these APIs glued together with their app's specific codes to build a unique experience to serve the purpose of the app. While there is no secret about the usage of the iOS SDK APIs, the very experience created by the app becomes the intellectual property and requires proper protection. Furthermore, Apple App Store recommends iOS app developers to submit a bitcode version of their app for Apple to perform post-processing on the submitted app for thinning or optimizing purposes. Including the bitcode version of the app will allow Apple to re-optimize the app binary in the future without the need to submit a new version of the app to the Apple App Store (Apple Help, 2020).

An effective obfuscation should be at the bitcode layer to preserve protection against reverse engineering of the binary while still being compliant to Apple App Store. The final iOS app will be placed in a sandbox to create isolation from other apps running on the same device providing some level of security and protection. However, the app is still susceptible to reverse engineering attacks using static analysis offline. Applying obfuscation in the application layer and entry calls to iOS SDK APIs makes it hard for the advisories to identify and follow the app logic. Of course, any software obfuscation comes with an overhead associated with making control flow unpredictable. A bitcode obfuscation level can then be tuned to an acceptable level to balance desire app performance and protection.

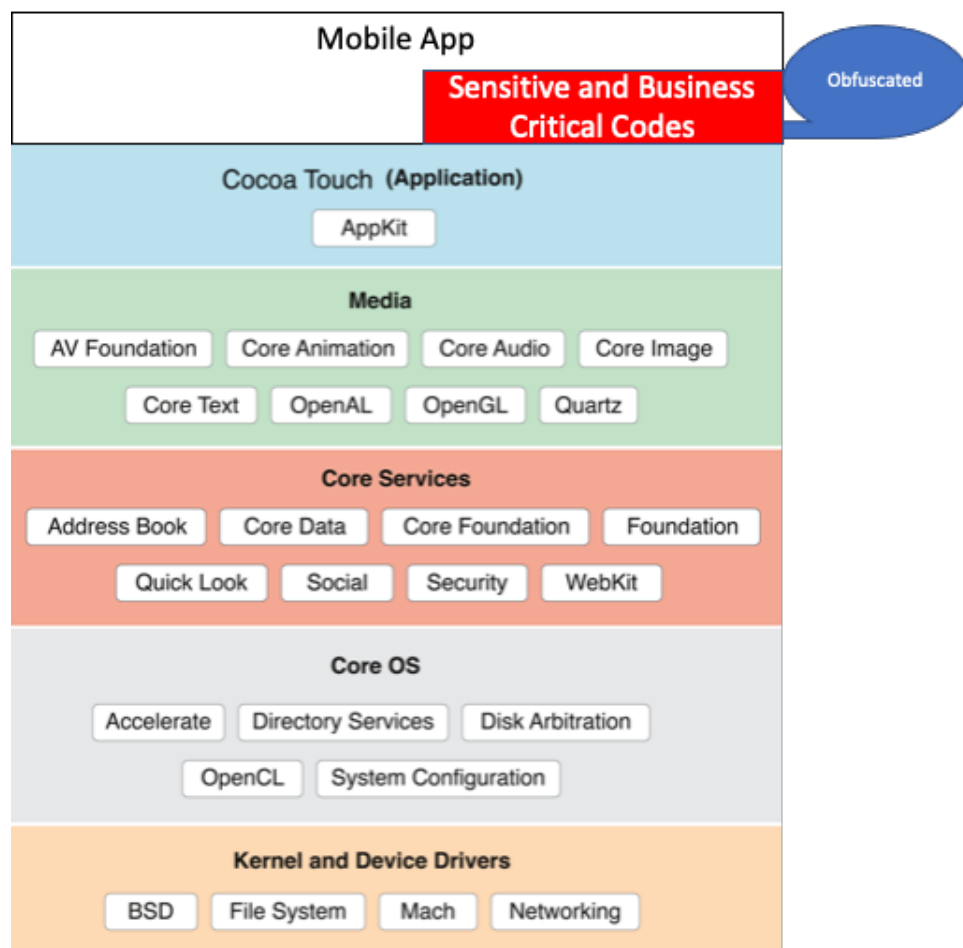


Figure 5 - iOS Mobile App Framework with App Obfuscation

The bitcode obfuscation provides a protection mechanism at the app layer that does not depend on underlying mobile platform security while conforming with Apple App Store requirements. It makes static analysis of the mobile app more complicated even with sophisticated tools.

5.2. Cloud Server Applications and Web Services

With the advent of cloud virtualization, more and more companies are pushing their server applications and web services to public cloud environments where they can clearly benefit from the scale, reliability, and availability of such services. This process requires a comprehensive analysis of cloud readiness of such applications. Companies typically rush to virtualize their server app to cloud without performing much needed due diligence evaluations. As a result, these server apps are posed to attacks when deployed in the cloud. Relying on cloud provider security alone is not an option when it comes to securing an app running in the cloud. Even though cloud providers typically come with a certain level of industry accepted secure environment, the security of data and application remains the ultimate responsibility of the cloud customers in every cloud model.

With exploitation of server apps using reverse engineering methods, an attacker can reveal information about back-end processes, steal intellectual property and gain intelligence needed to perform subsequent code modification (Dolan, Ray, & Majumdar, 2020). Web services are not exempt from these types of attacks either. A viable protection solution would require independence of the app protection from the underlying cloud platform. The notion of In-App protection means that the app is self-contained in terms of protecting its code and data within its binary regardless of its deployed environment.

Modern server applications and web services are typically written in high languages such as Java, JavaScript and GO. As discussed in the introduction, protecting such apps with obfuscation is not effective. Therefore, it's recommended to move all business critical and secret sauces in the app source code to native languages such as C/C++ for the obfuscation to be more appropriate. For example, if the cloud application is written in Java, there should be a Java Native Interface (JNI) layer to access C/C++ native code that are obfuscated as shown in the following diagram.

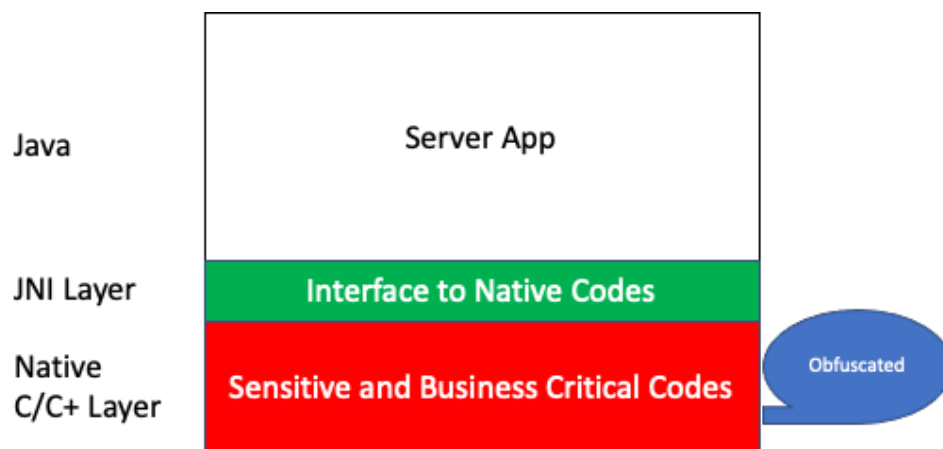


Figure 6 - Sample Cloud App with Obfuscation

5.3. IoT Applications

The rapid expansion of IoT devices creates opportunities for companies to be more innovative with the purpose and scope of their IoT applications. This large ecosystem comprises a variety of devices that are being used in diverse environments including healthcare, industrial control, and homes (Dolan, Ray, &

Majumdar, 2020). From SmartCity to SmartHome to Wearables, IoT devices are entering our world and the list of IoT apps will grow as technology evolves in the years ahead (Placeholder9) These IoT apps are edge devices collecting and processing sensitive data on end users' behavior, the user's devices, habits and their actions in addition to managing Personally identified information (PII).

Naturally the security of IoT apps remains the focus of various standard and industry bodies to regulate and advise. Limited security capabilities along with time-to-market pressure leaves IoT application developers with little or no provisions on securing the app. IoT devices typically come with no hardware security module, making them vulnerable to various attacks. These sophisticated apps analyze the collected data with rich business algorithms all within the IoT device. Even though most attack surfaces are runtime in nature and occur when the IoT app is in action, there is potentially a lot that can be discovered with statically analyzing and reverse engineering these apps.

As a result, the root of trust, device identity, keys and crypto operations conducted by IoT applications are exposed. Bitcode obfuscation can potentially reduce and even eliminate such concerns with IoT applications independent of the IoT device security.

6. Conclusion

In this paper, we introduced bitcode obfuscation as a very effective and powerful tool in protecting software application against reverse engineering attack. We explored different obfuscation techniques such as control flow and data flow obfuscation without accessing the original source code. The idea of operating on the intermediate binary without the need for the source code makes the bitcode obfuscation more practical to be offered as a cloud service (obfuscation-as-service). We also offered a few examples of how this protection can be applied to different software application domains and industries.

Abbreviations

AES	Advanced Encryption Standard
API	Application Programming Interface
GO	Google Programming Language
IoT	Internet of Things
JNI	Java Native Interface
LLVM	Refers to the LLVM compiler infrastructure project.
LLVM-IR	A platform-independent, universal low-level intermediate representation (IR) used by LLVM compilers and tools.
SCTE	Society of Cable Telecommunications Engineers
SDK	Software Development Kit
VM	Virtual Machine

Bibliography

- Anderson, L. (2015). A survey of control-flow obfuscation methods used in N-Mesh 2. (October).
- Anderson, L. A. (2018, 06 14). *Worldwide Patent No. WO2018106439A1*.
- Apple Help, A. (2020). *What is app thinning? (iOS, tvOS, watchOS)*. Retrieved from <https://help.apple.com/xcode/mac/current/#/devbbdc5ce4f>
- Arini Balakrishnan, C. S. (2005). Code Obfuscation Literature Survey. *University of Wisconsin, Madison*, <https://pages.cs.wisc.edu/~arinib/writeup.pdf>.
- Balakrishnan, A., & Schulze, C. (2005). Code Obfuscation Literature Survey. <https://pages.cs.wisc.edu/~arinib/writeup.pdf>, pp. 1-10.
- Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S., & Yang, K. (2001, apr). On the (im)possibility of obfuscating programs (<http://dl.acm.org/citation.cfm?doid=2160158.2160159>). *Journal of the ACM*, 59(2), pp. 1-48.
- Brekne, T. (2001). *Encrypted Computation*. Department of Telematics.
- Collberg, C. S., Thomborson, C., & Low, D. (1997). A taxonomy of obfuscating transformations.
- Dolan, A., Ray, I., & Majumdar, S. (2020). Proactively Extracting IoT Device Capabilities: An Application to Smart Homes. (A. Singhal, & J. Vaidya, Eds.) 42-63.
- Garg, S., Gentry, C., & Halevi, S. (2013). Candidate indistinguishability obfuscation and functional encryption for all circuits. *Proc. of FOCS*
- GitHub. (2022). *GitHub deobfuscation topic*. Retrieved from <https://github.com/topics/deobfuscation>
- Goldwasser, S., & Rothblum, G. N. (2007). On best-possible obfuscation. Springer.
- Hosseinzadeh, S., Rauti, S., Laurén, S., Mäkelä, J. M., Holvitie, J., Hyrynsalmi, S., & Leppänen, V. (2018). Information and Software Technology: A systematic literature review. 72-93.
- LLVM Doc, D. (2003). *LLVM Bitcode File Format*. Retrieved from <https://llvm.org/docs/BitCodeFormat.html>
- LLVM Project. (2022, 06 20). *LLVM Language Reference Manual*. Retrieved from <https://llvm.org/docs/LangRef.html#function-attributes>
- Lucideus. (2019, Jan 6). *Understanding the Structure of an iOS Application*. Retrieved from <https://medium.com/@lucideus/understanding-the-structure-of-an-ios-application-a3144f1140d4>
- OWASP. (2016). *OWASP M9: Reverse Engineering*. Retrieved from <https://owasp.org/www-project-mobile-top-10/2016-risks/m9-reverse-engineering>
- Wikipedia. (2022, 5 24). *Obfuscation (software)*. Retrieved from [https://en.wikipedia.org/wiki/Obfuscation_\(software\)](https://en.wikipedia.org/wiki/Obfuscation_(software))

Bringing the Mid-Split Factory Online to Rapidly Produce Terabytes

An Operational Practice prepared for SCTE by

Serge Kasongo

Director

Comcast, Field Operations
1717 Arch Street, Philadelphia, PA
215 286 7223
Serge_Kasongo@comcast.com

Dr. Robert Howald

Fellow

Comcast, CONNECT
1717 Arch Street, Philadelphia, PA
Robert_Howald@comcast.com

John Chrostowski, Comcast

Robert Thompson, Comcast

Table of Contents

Title	Page Number
1. Abstract	4
2. The What and Why of iHAT	4
2.1. Background	4
2.2. Purpose	4
2.3. Theory of Operation	5
2.3.1. Video Interference Risk	6
2.3.2. Drop Amplifiers	7
2.3.3. Issues with Legacy Systems	8
2.4. Evolution	9
2.4.1. iHATv1 – Proof of Concept and Functional Validation	9
2.4.2. iHATv2 – Ecosystem Automation and Transition to D3.1 OUDP Test Probe	10
2.4.3. iHATv3 – “Hitless” Implementation and Operation	12
2.4.4. Future Roadmap of iHAT	12
3. From Lab to the Field	13
3.1. Completing the Ecosystem	13
4. Integration with Production Ecosystems	15
5. Technical Operations Best Practices	18
5.1. Hardware	18
5.2. Software	20
6. The Terabyte Factory is Online	21
6.1. Mid-Split Activation	21
6.2. New Tools and Automation	22
6.3. Implications to Traffic Growth, Plant Upgrades, and Service Speeds	26
7. Conclusion	28
Abbreviations	28
Bibliography & References	29

List of Figures

Title	Page Number
Figure 1 – New Spectrum Split vs Standard Split Deployed Equipment	5
Figure 2 – The Two Basic RF Assessments Evaluated by iHAT	6
Figure 3 – Mid-Split Band Energy Isolation Path Across RF Splitter	7
Figure 4 – The iHAT “Black Box” Method and I/O	9
Figure 5 – Probe Signal Used in iHAT via DOCSIS 3.1 OUDP Feature	11
Figure 6 – iHAT as Part of a Production Upgrade Workflow	14
Figure 7 – Customer Journey During Upgrade	16
Figure 8 – Communications Strategy During Upgrade Cycle	17
Figure 9 – Customer Care Workflow.....	17
Figure 10 – Hardware Recommendations in the Home.....	19
Figure 11 – iHAT Definitions	21
Figure 12 – Mid-Split Channel Configuration of 4x SC-QAM + OFDMA	22
Figure 13 – Test Sample of an iHAT Dashboard	23
Figure 14 – OFDMA Health Statistics	24
Figure 15 – iHAT Execution Statistics.....	24
Figure 16 – Mid-Split Full Channel Configuration of SC-QAM + OFDMA	25
Figure 17 – Mid-Split Upstream Spectrum Capture Prior to Adding OFDMA.....	25
Figure 18 – Mid-Split Upstream Spectrum Capture Prior to Adding OFDMA.....	27

List of Tables

Title	Page Number
Table 1 – Communications Timelines and Tactics	18
Table 2 – Field Meters and Capabilites	20

1. Abstract

Cable operators are actively addressing upstream capacity with their most powerful tool – new spectrum. Mid-Split and High Split architectures satisfy long-term capacity and speed requirements, and are further empowered by DOCSIS 3.1 OFDMA. While capacity gains and speeds are straightforward to predict, the “how” of upstream spectrum migration requires careful planning.

At the 2021 SCTE Expo, in a paper entitled Executing the Upstream Makeover without Leaving Scars, a newly developed, remotely automated, in-Home Assessment Tool (iHAT) was unveiled. This tool was designed to manage large-scale activation of new Mid-Split spectrum while ensuring minimal customer disruption.

That was then! One year later, iHAT is fully integrated into production and technical operations workflows. The new spectrum is delivering up to 5 times more capacity. Technical training has been developed and deployed, and new care processes have been created and implemented. The iHAT tool itself has evolved – the production code is now ‘iHATv3’ – to make activation more efficient, automated, and less disruptive.

It is now timely to share the learnings along the path from tool development to production in scale. In this paper we will describe that journey. Technology, tools and processes for execution of large scale Mid-Split activation have been developed, launched, and scaled. Attendees will understand the spectrum activation journey from start-to-finish – system engineering, technical solutioning, operationalizing, and best practices – from the experts that made it happen.

2. The What and Why of iHAT

2.1. Background

Network upgrades of spectrum have historically targeted extending the Downstream from, say 550 MHz, to 750 MHz, 860 MHz, 1 GHz, or even 1.2 GHz. The outcome was fresh new fields of Bandwidth to seed with new and/or expanded video and data services. For the upstream, the spectrum has changed significantly less, based on consistently lower upstream demand and upgrades addressing upstream needs consisted primarily of adding additional DOCSIS channels and splitting nodes to prevent them from being congested.

Over the years, the spectrum has become fully occupied – albeit there is still capacity upgrade levers available by moving from DOCSIS 3.0 to DOCSIS 3.1. Furthermore, with the spectrum fully occupied, upstream traffic growth continuing, and mass CPE device swaps difficult and costly for modest capacity gains in the available upstream spectrum, node splits have become the primary tool of managing upstream, and as a result have been accelerating. The time to add more spectrum to the upstream has arrived, and Comcast, along with other MSOs, have made the decision to migrate to the 85 MHz Mid-Split architecture.

2.2. Purpose

Unfortunately, unlike the downstream, however, the “how-to” of spectrum migration gets more complicated in the upstream. The 42/54 MHz split has been in place for decades, and devices that adhere solely to it, particularly set-top boxes (STBs), are in many millions of homes. These devices support services that will remain active, and large scale swap outs of devices in homes, or

processes that require large scale physical visits to homes, should be avoided in order to cost effectively and efficiently deploy Mid-Split spectrum at scale.

Therefore, production-scale tools, techniques, and processes must be developed to ensure that a new, wider upstream path can be efficiently operationalized, while being transparent to customers. This the purpose of the iHAT tool. It is the foundational kernel of the Mid-Split activation process, allowing remote diagnosis of the drop-home environment necessary to activate seamlessly and create processes to activated ubiquitously and harmlessly to customers.

2.3. Theory of Operation

The problem statement for Mid-Split activation, highlighted in Figure 1, is to develop a way to unobtrusively discover the state of a home with respect to these two criteria:

- Potential for video interference
- Ability to support DOCSIS upstream pass-through in the Mid-Split band

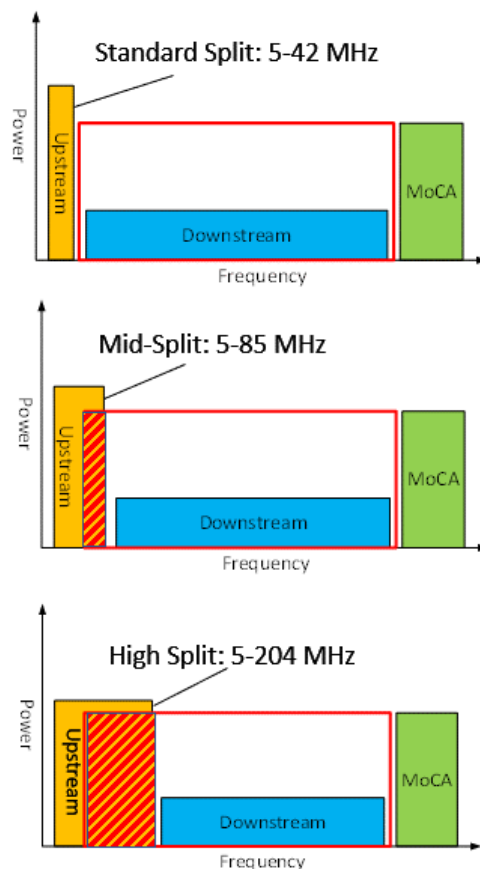


Figure 1 – New Spectrum Split vs Standard Split Deployed Equipment

To enable this home-by-home assessment in scale, an automated in-Home Assessment Test – aka iHAT – was developed to enable a seamless migration of capable CMs to utilize Mid-Split when conditions 1 and 2 above are satisfied, as shown in Figure 2




Figure 2 – The Two Basic RF Assessments Evaluated by iHAT

2.3.1. Video Interference Risk

With the decades of spectrum split in North America being set at 42 MHz/54 MHz, all QAM video STBs deployed are configured this way. They were built to receive video channels beginning at 54 MHz. When the network is NOT configured this way, and instead is upgraded to enable the duplex split to expand the upstream and activate new spectrum above 42 MHz, the QAM STB's point of view for Mid-Split or High-Split changes. This is shown in Figure 1.

The red cross-hatched areas in Figure 1 represent the spectral overlap imposed on a QAM STB by a Mid-Split capable cable modem when utilizing that band. Any signal energy that appears above 54 MHz can be seen by the STB downstream receiver, because it is built expecting to operate on downstream signals that begin at 54 MHz. Unfortunately for the STB, in a home that also contains a Mid-Split capable cable modem (CM), the CM sees that band as “eligible” for placing carriers when the CMTS is configured to allow CMs to use it.

Note that the STB is not acting on any specific signal type – it is simply adapting its Automatic Gain Control (AGC) function to deliver the ideal level to the A/D converter. AGC measures the total energy in the downstream band and doesn't care about its origin. Thus, if new Mid-Split upstream energy on the STB receiver is very high, the STB receiver will add attenuation. When this happens, the *desired* video channels will inadvertently be pushed lower through a phenomenon called “Adjacent Channel Interference” or ACI. If it attenuates too much, then the QAM video signals can be low enough to cause low SNR in these channels, and video pixelization could ensue.

Note the above description is a “static” or time-fixed snapshot view of upstream energy and spectral overlap with signals moving downstream to a STB. Actual upstream traffic is “bursty.” This is important because the AGC function has dynamic characteristics, but they tend to be slow acting. As a result, the duty cycle (off/on ratio) and burst duration is a factor that can impact the AGC implementations differently in different STBs.

As represented in Figure 1, the nature of the levels is not favorable – the downstream receive level is low (DS Rx), while the upstream transmit level (US Tx) is high. Until now, there was a diplex filter to separate them, but now, between 54-85 MHz, this is no longer the case. Fortunately, between a CM and a QAM STB there will be an RF splitter, the design of which will inherently isolate the port of a CM from the port of a STB by some amount. This scenario is illustrated in Figure 3, showing just one isolation path between modem (MTA) and a STB's RF inputs.




Figure 3 – Mid-Split Band Energy Isolation Path Across RF Splitter

How much energy leaks through to the STB? That question depends directly on the splitter and home wiring shown in Figure 3. Home wiring has a very predictable dB/loss per foot and is easily modeled. The most important factor with respect to the ACI phenomenon is the splitter(s) used to distribute RF to devices for video and data services.

2.3.2. Drop Amplifiers

In addition to OSP and traditional CPE devices that provide residential services that only know the 5-42MHz split, many homes also use drop amplifiers. These are also built with a Low-Split diplexer.

These devices must come out...eventually.... but because they may or may not be customer-impacting, they do not *necessarily* have to be tackled coincident with the activation of Mid-Split spectrum. From a capacity perspective, every drop amplifier that can be removed is good for capacity. The operations perspective depends on the percentage of homes that include a drop amplifier – estimated at 15-20% but can cluster depending on geography and practices. Methodically removing in-home drop amps over a period of time may make more sense than dealing with amplifiers transactionally as part of a service call or product upgrade. A proactive plan to address drop amps will eliminate the perpetual limbo state that is mixed-mode devices working in mixed-mode spectrum.

With capacity and product in mind, we can itemize home amplifier management into two buckets:

- Capacity-driven – All DOCSIS 3.1 CMs at Comcast are Mid-Split-capable, and all DOCSIS 3.0-only CMs are not. CMs are migrating to DOCSIS 3.1 steadily, so in time the majority will be installed and capable of Mid-Split upstream connectivity. However, even if a Mid-Split-capable CM is present, any home that cannot allow the spectrum to pass out of the home reduces capacity gain and upstream lifespan is compromised. Over time, these amplifiers must be removed to deliver on the full capacity promise. How quickly this must be done is a mathematical analysis of utilization vs penetration trajectory.
- Product-driven – One of the key benefits of the Mid-Split is the ability to deliver HSD speeds in the upstream such as 100 Mbps, 200 Mbps, 300 Mbps, and higher. Once such products are made available, customers with home amplifiers will be (self)-blocked from receiving them. Interest in speeds that require Mid-Split would trigger immediate action, to remove the blocking amplifier. The challenge is how to manage this efficiently and, more importantly, in a way that minimally impacts the customer. The good news is these blocking devices (amplifiers or any filtering within the band that may have been installed inline) can be discovered remotely and in real-time. While the customer cannot get the new upstream speed immediately, a rapid and transaction-based process can serve to notify the customer that additional steps are required to support the speed upgrade.

2.3.3. Issues with Legacy Systems

Some of the lessons learned from production deployments are that we now have the ability to uncover modifications made on the network over time to keep it operating and delivering services, but that have not been historically documented, or at least well-documented. Drop amplifiers are of course a case where documentation has been limited, but other scenarios exist:

- Noise and/or trap filters
- Frequency selective in-line equalizers
- RF Amplifiers added for plant extensions after initial build

Filters installed, such as for ingress, that impact upstream transmission bandwidth are somewhat obvious candidates to be found. AS scale is built, we are able to begin putting numbers around the likelihood of these types of devices effecting the ability to fully activate Mid-Split. The percentage is low, but not negligibly so. Incorporating processes to detect and locate these devices during the upgrade and activation will be incorporated. Not all cases need immediate remediation, but all cases where a customer is looking to achieve higher upstream speeds that only Mid-Split can provide, will. However, this can be a transactional process on product order.

The last category – RF Amplifier – may be less obvious. Over time, single amplifiers may have been added to a complex HFC network cascade to capture, for example, a new development at the end of a block. If these were not properly documented and put into the spatial database, then they could be missed in the upgrade process when existing amplifiers are swapped out for Mid-Split capable devices. The result is a Mid-Split blocker in a sea of newly upgraded Mid-Split capable actives. Fortunately, manifests itself as a cluster of localized iHAT failures that indicate more of a systemic issue than a home-by-home situation. Additional tools are then able to pinpoint the most likely ancestor device – which will be identified as a device close to the “missing” amplifier, a from which the actual location can quickly identified for remediation.

2.4. Evolution

A “Black Box” view of iHAT that includes the core functions of the initial design core is shown in Figure 4.

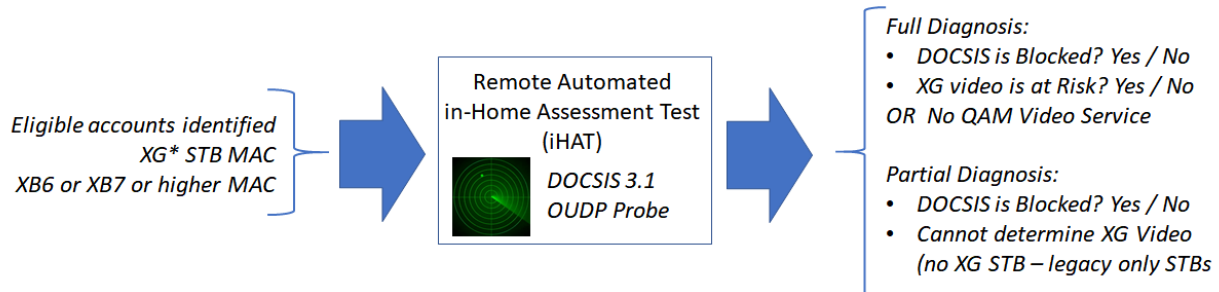


Figure 4 – The iHAT “Black Box” Method and I/O

2.4.1. iHATv1 – Proof of Concept and Functional Validation

With the introduction of the “N+0” architecture [Howald fiber frontier/2016, surf conditions/2018], Comcast began the upgrading of its HFC plant to an 85 MHz upstream. This has continued unabated, and is now on an accelerated path to cover more footprint, faster, by also upgrading in-place N+x HFC footprint to be Mid-Split capable. The system engineering and development of operational practices began at this time, and the mathematical basis for the potential for ACI due to Mid-Split quantified theoretically. Further lab testing was performed to assess the sensitivity of different STB models to ACI as a function of RF level and On/Off duty cycle of US traffic. The next step was to observe performance in the field against these projected impacts.

The first version of iHAT (iHATv1) was very much a trial-worthy, proof-of-concept (PoC) approach to facilitate learning about the drop amplifier and video interference deployment challenges. The approach used had high customer impact – multiple reboots to switch the CM diplexer between Low Split and Mid-Split, and a proprietary MIB managed. Fortunately today we have standardized MIBs via DOCSIS 3.1.

The method to evaluate readiness of a home for Mid-Split was to increase the upstream burst rate – i.e. change the speed tier to be higher to exercise the full Mid-Split spectrum and ensure the CM upstream signal would overlap with STB receivers, so the effect on the receiver could be observed when it did. Speed tests were then used to simulate utilization, while simultaneously measuring the effect on STB fidelity by observing modulation error ratio (MER) and uncorrected codeword error rate (UCER).

FEC-free range of operation for 256-QAM is $MER \geq 34$ dB and $BER \leq 1E-8$. Through lab measurements, we observed that video experience degrades – visible pixelization occurs – for 256-QAM when $ACI \leq -20$ dBc for the most sensitive of STB makes and models. By reading the MER and UCER during a speed test, a home could be identified as being at risk for video interference or not. Furthermore, the ability to demodulate the Mid-Split signals at the CMTS would identify that the bandwidth was not being blocked by a drop amplifier.

While this tool was a significant simplification of the in-home measurements that preceded it, where technician had to go to a home and take measurements and make a visual and RF diagnosis, it was still too limiting as a tool for use in scale:

- It could only be run during maintenance windows, at times when customers would be minimally impacted
- Test duration in a single home was long, due to multiple device reboots. For large node this would translate to several hours of MW down time for each node.
- Multiple APIs were required to orchestrate configuration, speed testing, and data collection

All the above led to a semi-automated engineering tool suitable for trial learning with 10s to 100s of customer assessments during a given maintenance window.

2.4.2. iHATv2 – Ecosystem Automation and Transition to D3.1 OUDP Test Probe

The iHATv1 tool validated the general premise that devices in the home and the tools available to activate and monitor them could be programmed for new functionality that would enable home-by-home diagnosis of Mid-Split activation readiness. The evolution to iHATv2 was driven by the need to become less disruptive to customers, less burdensome for the operations teams, more efficient in test duration and results availability, and more scalable for a production environment. Automation of the manual processes of pre-selecting customers and of launching iHAT itself. As shown in Figure 4, iHAT retrieves a list of devices, by account, on a particular Mid-Split-enabled RemotePHY Device (RPD) node, after the node is cutover, activated, and services restored. When an account is identified as having a Mid-Split-capable CM – for Comcast, this includes the DOCSIS3.1 gateway family of XB6, XB7, and XB8 – it is deemed eligible for an iHAT test. With one of those devices present, it will be possible to place the CM in Mid-Split mode to determine whether its upstream transmissions in the Mid-Split band are able to be seen and received by a Mid-Split enabled vCMTS and DAA node, or if they are blocked.

When an account also includes the “XG” class of QAM STB, the iHAT evaluation will look both for DOCSIS Mid-Split pass-through and the potential for video interference. This XG family, the majority of QAM STBs in the Comcast network, supports the proactive network maintenance (PNM) and SpectraCM functionality needed to capture RF measurements that are the basis for iHAT scoring of video interference potential. Older QAM STBs do not support this capability. In a home that includes an XG class STB, that measurement is a reasonable proxy for the expectation for other non-XG STBs with respect to their isolation from Mid-Split spectrum energy. If there is no XG-class STB present at all, but “legacy” QAM STBs are present, then no iHAT assessment can be made with respect to the potential for video degradation. At the outset, these homes default to Mid-Split activation. This policy will be revisited as production scale assessment can be made. This risk is low – single digit percentage [Howald MUSL 2021 paper]. By NOT defaulting to activating, the alternative is to roll a truck to each of these homes and take “iHAT” style isolation measurements manually.

2.4.2.1. OUDP Test Probe

The method iHAT uses to make its determination is based on the DOCSIS 3.1 OFDMA Upstream Data Profile (OUDP) feature, which allows a pre-defined “probe” signal to be scheduled by a CMTS and generated as a test signal. The probe can be defined by center frequency, bandwidth, and duration. When iHAT runs, it schedules this probe signal, home by home, to be burst into a portion of the Mid-Split spectrum.

Figure 5 shows the probe signal centered at about 80 MHz. It is 1.6 MHz wide (a common reference bandwidth for OFDMA bandwidth used in the DOCSIS 3.1 requirements), has a PSD at the ranged OFDMA power, and lasts 3-5 seconds. These are empirically-derived values through trial-and-error testing and optimization in the lab.



Figure 5 – Probe Signal Used in iHAT via DOCSIS 3.1 OUDP Feature

When the probe is fired, the time stamp is used to instruct the XG STB when to execute a Full Band Capture (FBC), and with that capture, samples are returned to that include levels of the OUDP probe and the first few downstream QAM channels. By determining the relative levels of these components and comparing them to an interference threshold value, making offset adjustments that account for the test probe not occupying the Mid-Split band completely, the home can be classified as to whether it needs remediation.

The OUDP method provides three major advantages:

- 1) It is part of the DOCSIS 3.1 specification, so a required feature to be compliant to the specification (when asked for!)
- 2) It can be a scheduled event within a system's normal operation, and therefore is very non-intrusive, happening without a customer's awareness or service interruption
- 3) As a scaled down (in total power) representation of an actual upstream signal, it does not actually create enough interference to impact video. Instead, it emulates what a small portion of the filled spectrum would look like and extrapolates mathematically to draw the proper pass/fail conclusion.

Iterative optimization of the parameters yielded a repeatable, reliable result that correlates well as a mathematical extrapolation with the video threshold testing that forms the foundation of ACI analysis.

The probe signal can also be used to evaluate blocking of the Mid-Split upstream by a drop amplifier, because if this is so, the CMTS will not be able to observe the probe. However, as part of the iHAT test, Mid-Split becomes active on a modem prior to an OUDP probe being launched, once the CMTS has a configuration that supports it. Ranging information of the OFDMA band (DOCSIS 3.1 ranging) is available to determine if the upstream was successfully sounded. If not, this is typically sufficient cause to identify a home with a drop amp issue, at which point the CM

can remain in partial service or reverted to Low Split mode. In either case, the home state is logged as “remediation required.”

An further benefit of the OUDP approach, looking to the future, is that it has a lot in common with the Sounding function that is part of DOCSIS 4.0 FDX.

2.4.3. iHATv3 – “Hitless” Implementation and Operation

While iHAT v2 provided significant improvements in the implementation of the test itself, in particular with respect to the execution of the core RF function itself, it still represented significant disruption because of the direct management of filters in the CM to be switched into one position or another. Modem resets that were required and device implementation differences for going back between filter settings, and onsite remediation challenges were among the obstacles. These ideas begat iHATv3, which was focused on “hitless” – or a substantially less intrusive iHAT operation sequence. A key change made to move in this direction was the use of the DOCSIS-standardized (and thus enforceable for all manufacturers) MAC Domain Descriptor (MDD) broadcasts to configure switchable duplex filters. Two major advantages to this approach are

- MAC re-initialization to Mid-Split is owned by configuration of the vCMTS and RPD, and not an iHAT function itself
- Dedicated boot files for the purpose of switching duplexers to assess Mid-Split serviceability is eliminated, removing a major disruptive step and minimizing boot file redundancy

An important element introduced by the integration of the vCMTS into iHAT’s functionality is the powerful role the vCMTS has in managing the broadband network. It puts much more capability at iHAT’s disposal, far above and beyond the initial simple iHAT functionality originally defined. iHAT is able to take advantage of this but moving away from duplex filter-based split configuration to manage whether a device is in Low Split or Mid-Split. With the vCMTS now an API away from iHAT, it is called upon instead to configure the bonding group of the modem to Low-Split or Mid-Split utilization, rather than move a switch in a CM. Referring to Figure 4, if iHAT finds that the device should remain in MS mode for either failure case it will dynamically adjust the bonding channel configuration (DBC) to Low Split so that there is no energy in the Mid-Split band, avoiding a drop amp blockage that could result in partial services flags at the vCMTS, or protecting a STB in the home.

2.4.4. Future Roadmap of iHAT

The iHAT tool has been launched, but as with any new launch and particularly in an agile SW-based environment, the road does not end at initial launch, it is just the beginning. Lessons learned will be rolled into future patches and releases, but also new use cases discovered and features needed to support those use cases developed, tested, and integrated into the production code.

Some of the use cases known at this time and part of the iHAT development backlog include;

- Low-Split amplifier, OFDMA blocker scenarios that extend beyond the home drop amp:
 - Use for outside plant amplifier cascades

- Cases of “missed” plant amplifier upgrades
- Hidden multiple dwelling units (MDUs) riser amplifiers
- In-line frequency-selective (and low-split) equalizers
- Enhanced orchestration between device APIs to minimize CM/STB communication failures
- Periodic iHAT results home health and maintenance checks
- Event-driven iHAT re-tests (for example, triggered by metrics in home or a new device in the home)
- Systemic failure results detection and automation
- Various iHAT dashboard filters and optimizations
- Extension to “nHAT” – detection of potential neighbor interference issues due to split changes, supporting expected phenomenon in both High Split and DOCSIS 4.0 FDX system

The features and capabilities will be addressed and rolled out in future software sprints.

3. From Lab to the Field

The mission statement for the original definition of iHAT in 2020 was to remotely go into a home with a test that would determine its readiness to be activate the home immediately with Mid-Split spectrum, and if it was not able to, diagnose why not. Now, two years hence, its fundamental technology has improved, it has been hardened, its capabilities and role expanded, and, more importantly, its APIs built out to turn it into a system that can be part of a large scale production operational step.

3.1. Completing the Ecosystem

Consider Figure 6 below, where iHAT is depicted at the heart of a broader upgrade workflow.

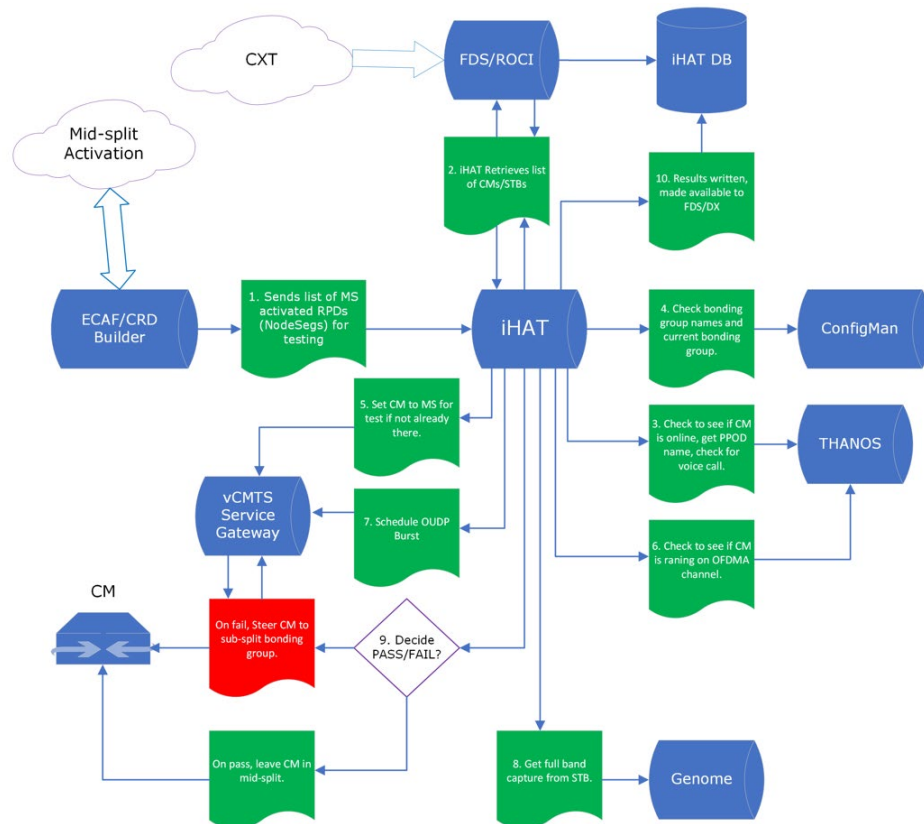


Figure 6 – iHAT as Part of a Production Upgrade Workflow

With a commitment to upgrade millions of HHP per year to mid-split capability, the tool cannot be operated in isolation, or be part of a workflow that has a “Stop” sign until human intervention to push a button. Automation is required to meet the pace required to cover the footprint target. Figure 6 contains a lot of homegrown names and acronyms associated with Comcast operations. The decoder ring is below, along with a description of their functionality as it pertains to Figure 6:

CXT = Customer Experience Technologies, the name of an organization within Comcast focused on systems and tools built for, and to maintain and support, the customer experience

FDS = Federated Data System. For iHAT purposes, it acts as an attribute and information sharing database for reading of essential information from the home and publishing out iHAT outcomes for subsequent disposition, processing and to other APIs. For example, what devices exist at this address? What type of device is it? What is the MAC address? What was the result of the iHAT test? The **iHAT DB** next too FDS/ROCI is where the output is pushed for other applications to use, in particular technician tools.

ROCI = An ML-based spatial tool that maps logical connectivity of devices to physical network architecture to the level of an individual Tap in the field. What modems are connected to which RF amplifier? Which node leg? Which RPD?

ECAF = A tool within the automated workflow managing construction projects. For iHAT purposes, ECAF signals when a Mid-Split RPD and HFC network upgrade project has been

completed in the field, triggering that the spectrum activation via iHAT can begin (there are some timing nuances to exactly when but the functional trigger is closure of ECAF ticket)

CRD = Configuration of an RPD is determined by the CRD it is programmed with. Its function in the iHAT workflow is to prepare the vCMTS to accept OFDMA upstream carriers for the iHAT testing and for data carriage itself after successfully passing the test.

ConfigMan = Configuration Manager. As described, this tool keeps track of the channel bonding configurations deployed and their state in the devices. Its role in iHAT is associated with the disposition of an iHAT result, whereby a “Pass” results in bonding if US SC-QAM and OFDMA carriers across the Mid-Split band for that home or business, and a “Fail” bonds only the 4 SC-QAM channels.

THANOS = A popular Marvel character, but for iHAT purposes it is an information store of vCMTS DOCSIS information of status, state, and logical connectivity.

Genome = A application that obtains metrics and data from devices and subsequently used to build dashboards and views in tools that maintain and support the RF network. Its role here in iHAT is to get the snapshot of the STB spectrum during the iHAT test to measure the isolation in the home.

vCMTS Service Gateway – An external application that in the iHAT application drives the vCMTS operation and the functions required to execute during the period that iHAT is running. It is apparent from Figure 6 that, while iHAT is at the heart of this flow, it is dependent upon a range of ecosystem software components to scale Mid-Split activation rapidly through the footprint. As part of the qualification of iHAT, a disciplined Integration and Test (I&T) process was developed to bring this full system to maturity and harden it for production. This process included phases of pairwise testing of relevant components to verify the APIs were communicating properly, exchanging the proper information, acting as expected on the information - reporting results, creating notifications or alarms, kicking off additional processes, etc., and executing cleanly and consistently before moving to the next phase of integration as the system was brought online component-by-component.

4. Integration with Production Ecosystems

While advancements in our network are necessary and important we have to be mindful about customer service interruptions. Minimizing the times we interrupt the customer’s service is crucial to limiting customer frustration, churn and sentiment. With the customer experience top of mind, we developed a roadmap for a customer’s journey and the touchpoints along their journey. We broke out the customer journey in three parts, as shown in Figure 7: Preparation, construction and appreciation.

PROCESSES: CUSTOMER COMMUNICATIONS

In an effort to convey reliability and to deliver a seamless customer experience, we're looking holistically across the customer lifecycle.

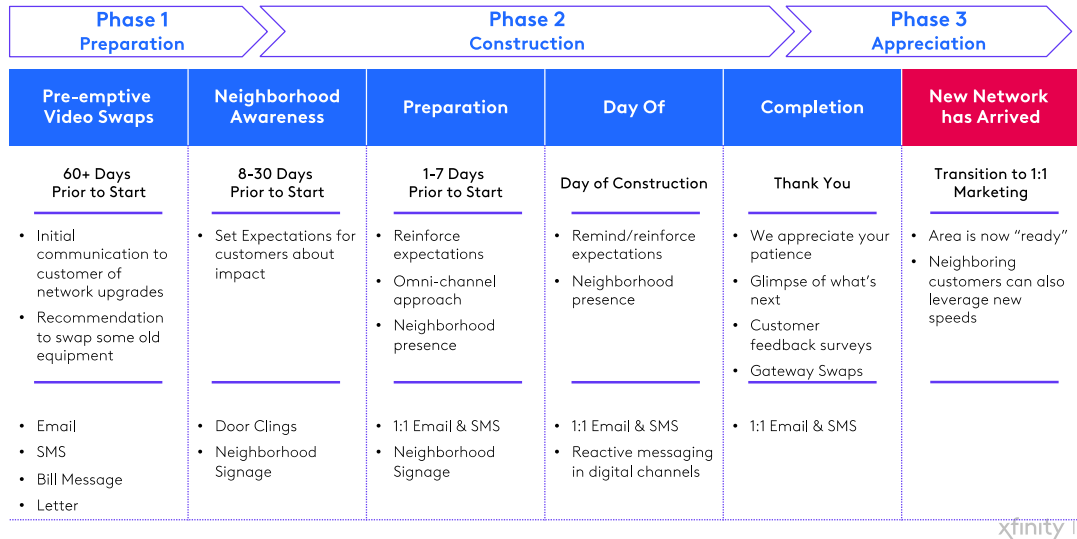
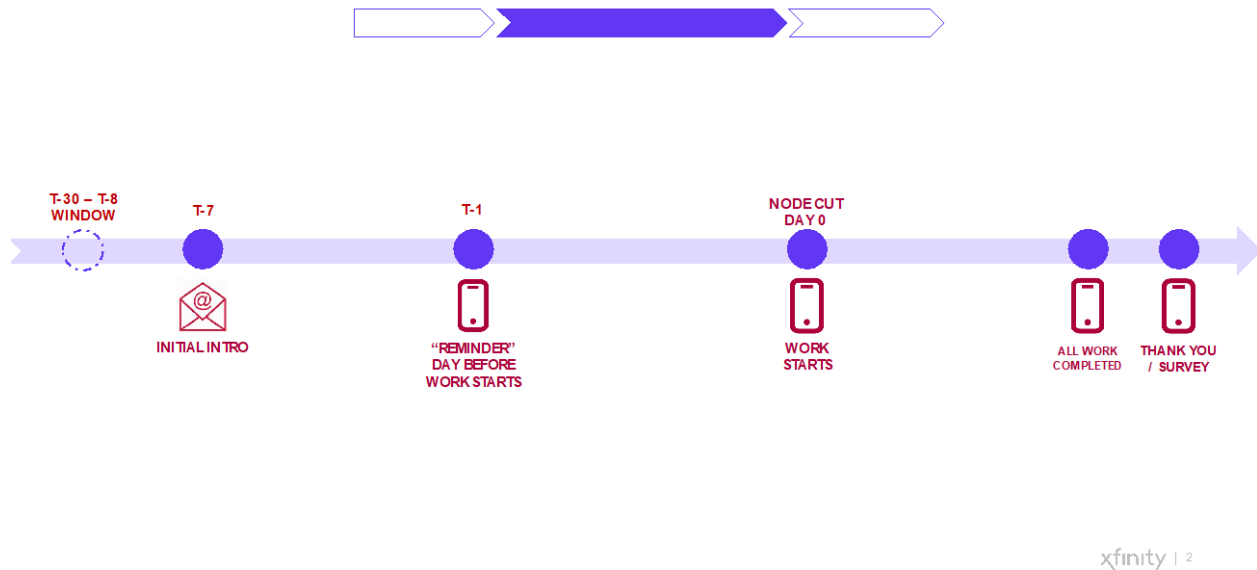


Figure 7 – Customer Journey During Upgrade

Along the customer journey we make it a point to notify them of upcoming service interruptions at least sixty days in advance. We reach out through multiple methods both physical and digital to ensure they are aware and have alternative options for connectivity during the times of interruption. During the construction phase our goal is to not have a customer hard down more than 180 minutes on average. Should we exceed 180 minutes hard down then credits and other forms of acknowledgment will be offered to the customer. When their journey is complete, we thank them and share our appreciation for their patience. We also share with them the new opportunities and capabilities offered to them with the upgraded network. This process is summarized in Figure 8.

PROCESSES: CUSTOMER COMMUNICATIONS

Proactive Communications journey – happy path of comms indicated by purple touchpoints



xfinity | 2

Figure 8 – Communications Strategy During Upgrade Cycle

From a CARE perspective we have also built-in the automatic detection of Mid-Split customers so that the moment they call us we know and meet their unique wants and needs. Furthermore, we integrate the intelligence of iHAT within our interactive troubleshooting guide to determine whether a truck roll is necessary to remediate an issue specifically caused by the enablement of the OFDMA spectrum. With both Mid-Split identification when a customer reaches out and quick diagnosis as to whether their issue is Mid-Split related will set us up for success in the future. This flow is shown in Figure 9.

PROCESSES

Example: Mid-Split Spectrum Activation – SIK+ Drop Amp Remediation (XB6/7 Only)

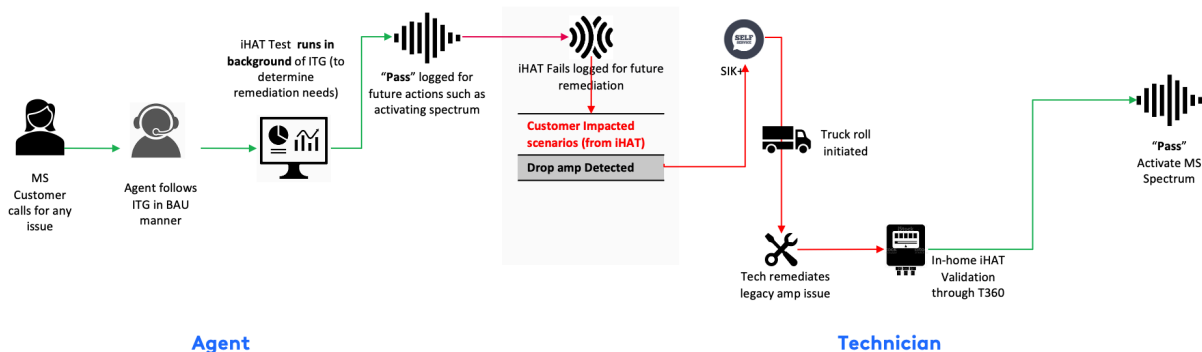


Figure 9 – Customer Care Workflow

5. Technical Operations Best Practices

Cross-functional teams across all Divisions and HQ are involved in executing the Mid-Split architecture at scale while caring for the end-to-end customer experience. A major advantage of our network architecture is that we can quickly and continuously evolve it with minimal disruption to customers — without digging up their yards and neighborhoods, as laying fiber to their homes would require. In most cases, customers should only be down for a few hours while our teams cut over to the enhanced technology on a neighborhood-by-neighborhood schedule.

Through neighborhood signage, SMS, email, and more, we will proactively and fully communicate to customers throughout the journey – well before work begins, during down time, and after the job is done.

In Field Operations, technicians will need to understand the importance of the work we’re doing to prepare our network for the future. They will also be responsible for:

- Preparing homes and businesses to use the new Mid-Split frequency range
- Identifying and removing hindrances preventing modems from leveraging the new frequencies

Due to the dynamic nature of Mid-Split deployment we’ve had to be very strategic about how, when and where we communicate. While technicians in a popular and competitive metro area would likely see activity in their market first it is equally important, we also communicate to their peers in more rural areas in a timely manner. **Table 1** is an example of the recommended communications timeline and tactics from both a national perspective to a small, localized trial.

Table 1 – Communications Timelines and Tactics

	Item	Audience	Key Message	Delivered By	Approximate Timing
	Final Comms Package and Templates	Division Communicators	Ensure the messaging is aligned and consistent across the enterprise with peer review	National Communications	
	Training Materials	Division Teams	Ensure the training is aligned and consistent across the enterprise with peer review	National Communications	
1	Leader Update	Supervisors+	Awareness to leaders that customers may start seeing advanced comms, program introduction & awareness	Launch Area Specific	Launch minus 3 weeks (or as determined by division)
2	Team Deck	Supervisors to use in team meetings	High level overview of Path to 10G	Launch Area Specific	Launch minus 2 weeks (or as determined by division)
3	Tech Awareness	All Technicians, Supervisors	High level overview of Path to 10G	Launch Area Specific	As Determined by local leadership
4	Tech Awareness	All Technicians, Supervisors	High level refresher of Path to 10G	Launch Area Specific	As Determined by local leadership

5.1. Hardware

Our network has evolved over the last few years to require less QAM set-top boxes in favor of QAM-less IP set-top boxes. The benefit of having less QAM set-top boxes in the premise means there is less of a need for in-home amplifiers in the home. For this reason we have been on a journey to bring field awareness about keeping the premise as passive as possible. Doing so sets us up for success and cost savings by not needing to return to the premise to remove a Mid-Split


amplifier when we move to High-Split or any unforeseen challenges when we go to DOCSIS 4.0.

An additional hardware consideration is the field meter necessary to support troubleshooting both OFDM and OFDMA in the field. Most of us are very familiar with how quickly technology can evolve so being prudent about which meter we deploy and when is crucial to our success. Furthermore, knowing the make and model technicians have with them is crucial to effectively route the ones with fully capable meters to a home that is leveraging the Mid-Split architecture. Lastly, the speed testing capabilities of the meter are very important for customer speed validation.

Training materials to onboard technicians for these challenges are shown in Figure 10 and Table 2.

TOOLS OF THE TRADE: HARDWARE

Keep homes as passive as possible and only use an amplifier when absolutely necessary



xfinity | 5

Figure 10 – Hardware Recommendations in the Home

Table 2 – Field Meters and Capabilities

TOOLS OF THE TRADE: HARDWARE

XM Capabilities Matrix

	XM1 - 2016	XM2 - 2019	XM2M - 2021
DOCSSDs Capability	D3.0 4x SC QAM	D3.132x SC QAM + 2 OFDM	D3.132x SC QAM + 2 OFDM
Ds Spectrum Visibility	50-1002 MHz	50-1212 MHz	50-1212 MHz
DOCSSUs Capability	4x SC QAM	8x SC QAM + 2 OFDMA	8x SC QAM + 2 OFDMA
Us Spectrum Diplex	5-65 MHz Fixed	5-42 / 5-85 MHz Diplex	5-42 / 5-85 / 5-204 Triplex
Ingress Widget Bandwidth *	5-125 MHz	5-125 MHz	5-204 MHz
Speedtest Bootfile	N/A requires XMTOdroid side car	d11_m_cgndp3_2g250m_c05.cm	d11_m_cgndp3_3g2g_c05.cm
Ethernet Capability	100Mbit	1000Mbit	2500Mbit

XM2 recommended for Mid-Split areas

*Noise and Ingress can be measured across the entire meter bandwidth using DSSpectrum widget

xfinity |

5.2. Software

Software is equally as important as the hardware because in many instances it can save a technician time by calling out impairments and opportunities ahead of their arrival. The software also serves as a be a second point of validation by confirming what the meter reads as opposed to a gateway. We've integrated tools like iHAT to help technicians be aware of previous conditions or impairments with a premise. The context enables them to know what they're likely going to a premise to remediate. Once they've remediated the issue we provide them the ability to force the gateway to start leveraging the OFDMA carrier.

Understanding enough about iHAT is important for technician to efficiently support the roll-out of Mid-Split and OFDMA. Figure 11 is a slide from the training material associated with this new software tool.

TOOLS OF THE TRADE: SOFTWARE

What is iHAT?

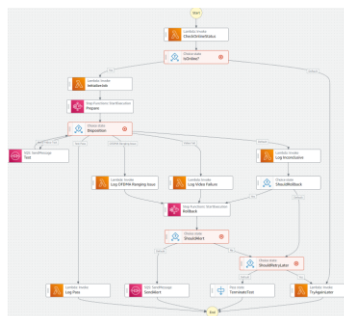
iHAT is a NGAN tool developed to help us determine whether a premise can support Mid-Split

How does it work?

When a node has been converted to Mid-Split, iHAT will be run by NGAN to determine which homes can activate MS spectrum without interruption and which homes may need remediation

What is iHATv3?

iHATv3 is an updated version of iHAT that will not require a device reset to perform a test



Identified Issue	Definition	Solution
OFDMA Ranging Issue/ Blocked Carrier	iHAT is unable to reach DOCSIS Device	Most likely cause is a drop amplifier in the way and will require a TC to remove or replace with Mid-Split Drop Amplifier
Video Interference	iHAT identified Video interference from one of the XG devices in the home	This is caused by low port-to-port isolation in the soft drop system TC will be required to ensure passive drop network meet Comcast specifications. (other solutions available depending on market availability)
No XG Devices		iHAT can perform a test but will not be able to provide video interference information
XB3/ BWG or older model		Customer does not meet minimum requirements to perform test

xfinity | [®]

Figure 11 – iHAT Definitions

6. The Terabyte Factory is Online

Comcast has been installing digital Mid-Split capable nodes for over 5 years as part of our “Fiber Deep” DAA network upgrade strategy using Remote PHY nodes (RPDs). Up until recently these deployments were exclusively in areas upgraded to “N+0,” which eliminates all RF amplifiers in the plant between the DAA node and customer homes. As we continue to build out our DAA network, deployments now include RPDs in typical “N+X” (X = number of RF amplifiers) using Mid-Split amplifiers swapped into existing amplifier locations to minimize plant rework. The entire DAA footprint is Mid-Split ready, and the process of activating the Mid-Split spectrum has begun, configured using both DOCSIS 3.0 SC-QAM channels and DOCSIS 3.1 OFDMA channels up to 85 MHz.

One of the advantages of using the vCMTS and DAA is the ability to perform extensive monitoring of the network and easily build custom tools and dashboards. These dashboards are critical during deployment to ensure we have visibility into the devices and network as the spectrum and upstream OFDMA are being activated.

We look at these aspects in more detail below.

6.1. Mid-Split Activation

Activating the Mid-Split spectrum is a multi-stage effort including planning, construction, spectrum activation, in-home health assessment, remediation and monitoring. Through June 2022, Comcast has over 30,000 DAA nodes built and activate, and the Mid-Split migration process has begun. The engineering effort, testing and planning for deployment is paying off and the monitoring tools and dashboards developed have been invaluable.

Spectrum activation consists of two stages:

- 1) Updating the system to use the OFDMA spectrum. This includes pushing an updated configuration to the RPD which has a 5 channel bonding group (4 SC-QAM channels + one OFDMA channel to 85 MHz). Figure 12 is illustrative of this configuration.




Figure 12 – Mid-Split Channel Configuration of 4x SC-QAM + OFDMA

2) Running iHAT

- a) Measures video interference on accounts with STB boxes and moves devices which are determined to have video interference to the standard 4 SC-QAM channel bonding group via Dynamic Bonding Change request.
- b) Moves devices which have OFDMA blocked and cannot use the OFDMA channel back to the 4 channel bonding group using Dynamic Bonding Change request.
- c) Steering the devices which cannot use OFDMA to the 4 channel SC-QAM bonding group prevents these from showing as in “partial” state in the monitoring tools.

During the spectrum activation and iHAT process, devices are monitored to ensure they are online and using the appropriate channel bonding group based on the device capability and iHAT analysis.

Upon initial spectrum activation, devices which have the spectrum between 42 and 85 MHz impaired will not bond to the Mid-Split OFDMA channel. These devices will initially show as upstream partial until iHAT is run. Once iHAT is run, these will be moved to the 4-channel bonding group and show as fully online.

As of June 2022, Comcast has over 35,000 DOCSIS 3.1 devices activated with Mid-Split and OFDMA across approximately 1,800 RPDs. The rate of activation will accelerate in the 2nd half of 2022 and through 2023 as new products enabled by Mid-Split are brought to market.

6.2. New Tools and Automation

Achieving the scale noted above would not be possible without monitoring dashboards. One of the main dashboards used during activation is the iHAT dashboard. This dashboard monitors the progress of checking for OFDMA partials and testing for video interference. In-home video interference, which caused by adjacent channel interference, is measured by having the cable modem send out a 1.6 MHz wide OUDP burst. This burst is sent out at the same spectral density as the OFDMA channel. The level of interference from this burst is measured at the adjacent STB in the home. If this burst exceeds a pre-determined threshold vs the downstream QAM signals at the STB, the cable modem is steered back to the 4 SC-QAM 4 channel bonding via DBC and will not use the OFDMA. This account is noted for remediation to be able to activate the OFDMA spectrum.

iHAT also checks the OFDMA channel and if it is registered and bonded. If the OFDMA channel is bonded and the account passes the video interference on the accounts with STBs, the CM remains in the 5 channel bonding group with OFDMA active.

If the OFDMA channel is not registered and bonded, iHAT steers the CM to the 4 SC-QAM bonding group via DBC. This removes the status of partial and shows the device online. The account is targeted for remediation to determine why the OFDMA channel is not bonded. The iHAT dashboard can look at the entire network, specific PPODs, and specific RPDs to see the number of devices with OFDMA blocked and video interference. Figure 13 is a sample of the iHAT dashboard after spectrum activation for a single PPOD which includes 154 RPDs.

This Dashboard shows the following:

- **RPD Count:** This is the total number of RPDs tested for the specific PPOD and timeframe selected.
- **Total number of tests:** This is the total number of tests performed which include testing for OFDMA Blocked, Video Interference
- **Pass:** Number of devices which can connect to the OFDMA channel and which pass the video interference test
- **OFDMA Blocked:** Number of devices which cannot connect to the OFDMA channel. This is caused by standard split in-home drop amps, or other in-home or plant issues affecting the Mid-Split OFDMA spectrum
- **Video Interference:** Number of devices which STB boxes on the same account, which connect to OFDMA, but where the OUDP burst received at the STB box is higher than the downstream QAM by the pre-determined threshold

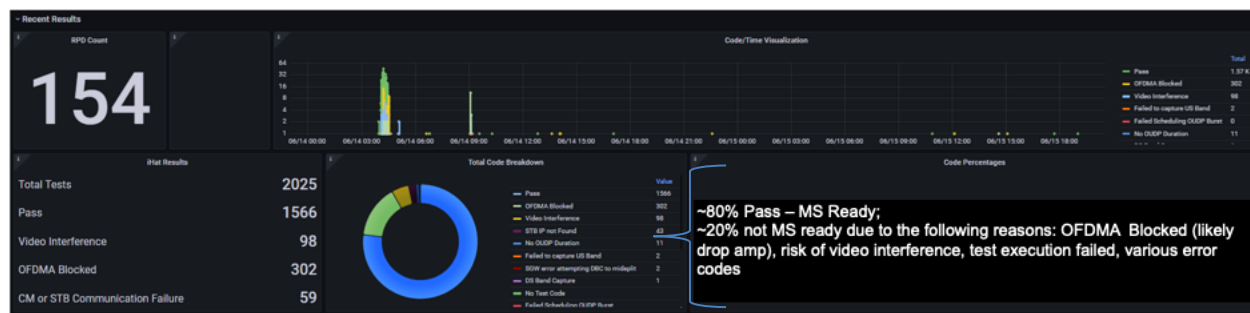


Figure 13 – Test Sample of an iHAT Dashboard

Additional tools and dashboards have been built to monitor the OFDMA health across the network. The dashboard in Figure 14 shows the OFDMA health for a specific PPOD. This dashboard can show the performance of critical OFDMA physical layer parameters of the whole network or down to a specific RPD.



Figure 14 – OFDMA Health Statistics

The dashboard in **Figure 15** shows the state of OFDMA cable modems and the percentage of OFDMA partials in a PPOD and the overall change in OFDMA partial count before and after running iHAT. Partial count is reduced significantly post iHAT.




Figure 15 – iHAT Execution Statistics

Standard existing tools are also used when activating the OFDMA spectrum. Yeti, which is an upstream spectrum analysis tool, has been updated to support Mid-Split can be used to verify the health of the upstream OFDMA spectrum. Figure 16 shows the upstream spectrum with the full 5-channel bonding group including OFDMA. In the Figure 17, the Yeti capture shows the 4 SC-QAM channels prior to adding OFDMA. Noise can be seen in the OFDMA band at 44-60 MHz and at 83 MHz. This is off air ingress from channels 2,3 and 5 which will have some effect on the total upstream capacity, but with PMA will be minimized.

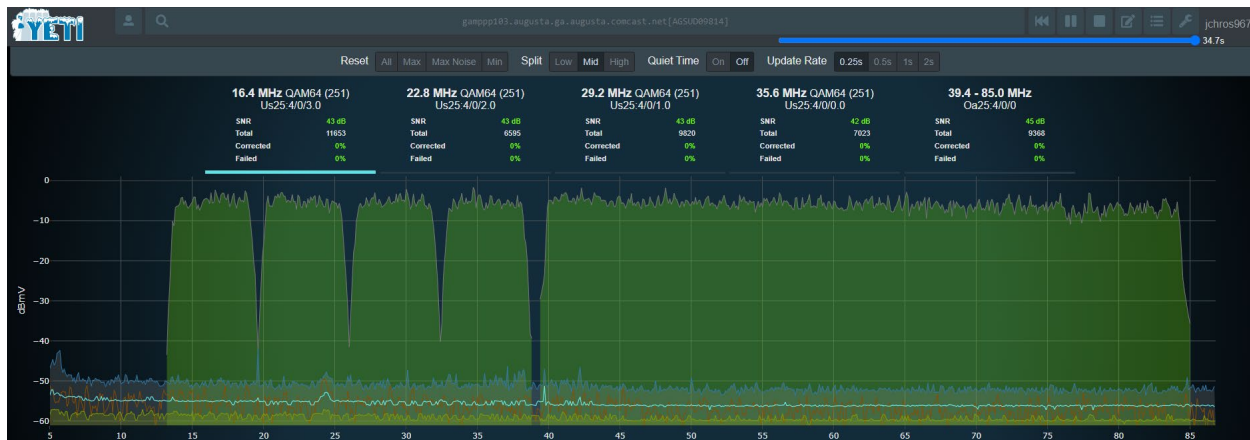


Figure 16 – Mid-Split Full Channel Configuration of SC-QAM + OFDMA

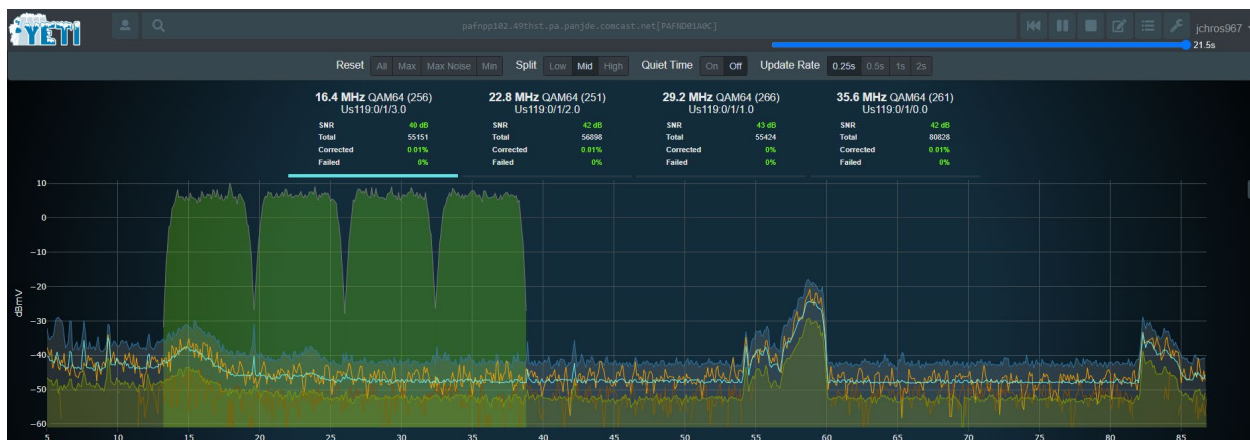


Figure 17 – Mid-Split Upstream Spectrum Capture Prior to Adding OFDMA

Numerous other tools and dashboards are available for scale deployment of mid-split and the ones above provide a look into the detail needed for this deployment.

Among the most important of these is the capacity impact of Mid-Split. Referring to Figure 16, a standard 4x SC-QAM plus an OFDMA channel from 39.4 MHz to 85 MHz offers a bonded total capacity of approximately 450 Mbps, with some variation around that depending mostly upon network fidelity variables. Of course, only the DOCSIS 3.1 devices can avail themselves to this bonded capacity. The DOCSIS 3.0 devices have their traffic contained completely within the SC-QAM allocation.

While only, roughly, doubling the upstream spectrum, we can calculate that the mid-split spectrum powered by DOCSIS 3.1 increases capacity > 4X compared to the standard 4-channel SC-QAM low-split upstream allocation. During the COVID outbreak, where an urgent need for capacity was felt everywhere and DOCSIS 3.1 had not yet been enabled (and still is not in the low-split spectrum), Comcast deployed 6 SC-QAM upstream channels to relieve congestion. This 6-channel capability represents effectively the “maximum” D3.0 only capacity in a low-split system by which to compare the Mid-Split allocation with the OFDMA augmentation.

Because of the interplay of device penetrations, capacity available to devices classes, and the variations of the network's ability to support QAM profiles, Comcast made an investment in capacity analytics to quantify how the introduction of OFDMA would impact network capacity in order to make good business decisions around device upgrades, network upgrades, and regional priorities.

For a complete analysis of the capacity impacts of Mid-Split with OFDMA, the reader is referred to [1].

6.3. Implications to Traffic Growth, Plant Upgrades, and Service Speeds

Mid-Split expansion takes the available upstream bandwidth from 37 MHz to a limit of 80 MHz. It was defined in DOCSIS 3.0, with the upper limit selected in part to fall just below the FM radio band in the US, while preserving the important downstream video out-of-band (OOB) signals widely used by legacy QAM set-top boxes (STBs). In recent history, it has been the upstream bandwidth limitation in the face of continued traffic growth that drove network upgrade activity.

The average per-user peak-busy-hour (pbh) upstream is still in the hundreds of kbps range – 400-500 kbps going into 2022. The downstream, by contrast, is about 10x that or more. As a general rule, the upstream payload has on average grown more slowly than downstream, although it has tended to more volatility year-over-year as different “killer apps” arose such as peer-to-peer file sharing (i.e. Napster) and home security video cameras. However, where DS was racing ahead at 50% per year for many years, driven by streaming video most recently, upstream averaged 20-30% - and less or flat in some years. An additional bonus of Mid-Split is that this new upstream spectrum is typically cleaner. As a result, when combined with use of much more spectrally efficient DOCSIS 3.1 OFDMA, the Mid-Split impact on network lifespan is extremely powerful. Roughly twice the total upstream spectrum converts to 3-4 times the available capacity.

Figure 18 shows the time runway generated by three options – node split, node split plus upgrade to Mid-Split, and finally N+0 with Mid-Split. While N+0, with smaller service group size, offers the longest runway of the three, an N+x migration tied to a node split is also a very effective way to extend HFC lifespan to nearly 7 years.

A key benefit of N+x with spectrum migration is its ability to add capacity quickly when compared to N+0. When the Covid-19 traffic spike eliminated months of CAGR lifespan, N+x upgrades brought more US bandwidth to the network quickly to reset the lifespan timeline. With a year of capacity growth runway erased by the pandemic – perhaps it flattens as the effect recedes and a new normal established – alternatives such as drop-in HFC upgrades that are both fast and effective make a sensible augmentation step.




Figure 18 – Mid-Split Upstream Spectrum Capture Prior to Adding OFDMA

The 7 years of lifespan offers a comfortable window of time to assess HSD trends (is CAGR changing), technology availability options (i.e. DOCSIS 4.0, R-OLT FTTH), and assess emerging applications and the speed, utilization, latency, and jitter requirements that they entail and adjust accordingly if necessary.

Lastly, while the “big story” for upgrading to Mid-Split is been its ability to deliver a long-term life span for the HFC plant without a major compromise of downstream bandwidth, there are significant upstream service speed opportunities as well. The Low-Split upstream payload data capacity using D3.0 SC-QAM with 64-QAM carriers is about 100 Mbps, and residential upstream service speeds limited to 35-50 Mbps today.

By contrast, the available capacity for a DOCSIS 3.1 CM in a Mid-Split network using the configuration shown previously in Figure 16, and as noted on the figure and observed in the prior section, is about 450 Mbps. This will enable upstream speeds of 100 Mbps, 200 Mbps, and up to 300 Mbps – under an empirical set of guidelines for penetration, utilization, service group size, etc. The field trial activity previously described validated speeds in scale up to 200 Mbps in preparation for market launches of these speeds in the future.

As DOCSIS 3.1 device penetration continues to increase, the DOCSIS 3.0 QAM carriers can be exchanged for more bandwidth efficient DOCSIS 3.1 OFDMA, which will make this bandwidth approximately 60-80% more efficient on average, resulting in the 450 Mbps of capacity increasing to over 600 Mbps if fully converted to DOCSIS 3.1, which would take practical speeds up to at least 400 Mbps.

So, the Mid-Split upgrade represents a strategy that addresses both effective capacity growth management and higher speed product offerings for residential and business customers.

7. Conclusion

Consumer usage trends continue to fluctuate as new applications and societal shifts take place. The year 2020, inclusive of the global pandemic, accelerated and altered some of the usage trends. Many companies sent employees to work from home at the onset of the pandemic while some schools went completely virtual. Online gaming has become less of a leisure hobby and adopted by many professional teams inspiring a young generation of one day having the opportunity to participate and potentially create and monetize for themselves. One thing is clear, all usage, including upstream, is likely to continue increasing. As we enter the last quarter of 2022 one would be hard pressed to believe we will ever see ourselves in a pre-pandemic state of usage trends. If anything, we should accept what we know today as the new normal and tactically plan to continue bolstering and augmenting our networks.

Mid-Split scratches the surface for the potential of our networks, but it is a realization of what was once only a vision not many years ago. We've built the tools, processes, training, technology and end-to-end systems to develop, launch, scale, support, and maintain a DOCSIS 3.1 OFDMA Mid-Split activate network, and Comcast and our customers are reaping the rewards of this effort through a healthier, smarter, more capable network that provides a better, high availability Internet experience for our customers with higher upstream speeds to deliver on these emerging trends and application

Looking ahead, the advent of High-Split and DOCSIS 4.0 FDX present their own new challenges that, much like Mid-Split, appear daunting at the outset. However, much like we've done in the past, we will use the collective wisdom of the industry to take on these new challenges. The learnings we have gained and will continue to gain building out DAA Mid-Split will make the path that much easier for the generations to follow.

Abbreviations

ACI	Adjacent Channel Interference
AGC	Automatic Gain Control
BAU	Buisness as Usual
BG	Bonding Group
CACIR	Carrier-to-Adjacent Channel Interference Ratio
CAGR	Compounded Annual Growth Rate
CDF	Cumulative Distribution function
DAA	Distributed Access Architecture
DBC	Dynamic Bonding Change
DSG	DOCSIS Settop Gateway
FDD	Frequency Division Duplex
FDX	Full Duplex
FTTH	Fiber-to-the-Home

HHP	Households Passes
iHAT	in-Home Assessment Test
LoQ	Line-of-Questioning
MER	Modulation Error Ratio
MTA	Media Terminal Adaptor
MUSL	Mid-Split Spectrum Upstream Launch
NI	Neighbor Interference
OFDMA	Orthogonal Frequency Division Multiple Access
OOB	Out-of-Band
ODUP	OFDMA Upstream Data Profile
OSP	Outside Plant
OTA	Over-the-Air
PHT	Performance Health Test
QAM	Quadrature Amplitude Modulation
SNR	Signal-to-Noise Ratio
STBs	Settop Boxes
TaFDM	Time and Frequency Division Multiplexing
TCP	Total Composite Power

Bibliography & References

- [1] Harb, Maher et al, Deploying PMA-Enabled OFDMA in Mid-Split and High Split, 2022 SCTE Expo, Sept 19-22, Philadelphia, PA
- [2] Howald, Robert, Execute the Upstream Makeover without Leaving Scars, 2021 SCTE Expo, Oct 11-14, Atlanta, GA.
- [3] Howald, Robert, Repair the Ides of March: COVID-19 Induced Adaption of Access Network Strategies, 2020 SCTE Expo, Oct 12-15 (virtual).
- [4] Thompson, Robert, and Rob Howald, Dan Rice, John Chrostowski, Rohini Vugumudi, Amarildo Vieira, and Zhen Lu, Rapid and Automated Production Scale Activation of Expanded Upstream Bandwidth, 2021 SCTE Expo, Oct 11-14, Atlanta, GA.

Broadband Capacity Growth Models

Will the end of Exponential Growth eliminate the need for DOCSIS 4.0?

A Technical Paper prepared for SCTE by

John Ulm

Engineering Fellow
CommScope
Moultonborough, NH
+1 (978) 609-6028
john.ulm@commscope.com

Zoran Maricevic, Ph.D.

Engineering Fellow
CommScope
Wallingford, CT
+1 203-303-6547
zoran.maricevic@commscope.com

Ram Ranganathan

Dir, Systems Engineering
CommScope
Toronto, ON
+1 905 568 731 7
ram.ranganathan@commscope.com

Table of Contents

Title	Page Number
1. Introduction.....	5
2. Network Capacity Planning Overview	6
2.1. The “Basic” Traffic Engineering Formula	6
2.2. The “Modified” Traffic Engineering Formula	7
2.3. Max Service Tiers in 10G Era	9
2.4. Determining Service Group QoE based on Probabilities	9
3. Broadband Subscriber Traffic Consumption – 2022 Update	11
3.1. Downstream Peak Period Average Bandwidth per Subscriber, DS Tavg	11
3.2. Upstream Peak Period Average Bandwidth per Subscriber, US Tavg	11
3.3. Downstream to Upstream Peak Period Average Bandwidth Ratio	12
4. Broadband Traffic Growth Trendlines	13
4.1. Broadband Traffic Growth in retrospect	13
4.1.1. DS Tavg Growth – 2010-22	13
4.1.2. US Tavg Growth – 2010-22	15
4.2. Correcting for the COVID Bump	16
4.3. Various Downstream Growth Trendlines	18
4.3.1. DS Exponential Growth Trendline	18
4.3.2. DS Linear Growth Trendline	19
4.3.3. Other DS Growth Trendlines	20
4.3.4. DS Adoption Curve (S-curve) Growth Trendline	21
4.4. Mapping DS Tavg Growth Projections over a 10-15 Year Span	28
4.5. Various Upstream Growth Trendlines	30
4.5.1. US Exponential Growth Trendline	30
4.5.2. US Linear Growth Trendline	31
4.5.3. Other US Growth Trendlines	31
4.5.4. US Adoption Curve (S-curve) Growth Trendline	33
4.6. Mapping US Tavg Growth Projections over a 10-15 Year Span	35
4.7. DS:US Tavg Ratio projections	36
5. Network Capacity Modeling for Low to High Growth Projections	37
5.1. Network Capacity Modeling Assumptions	37
5.2. Making the Most of 1218/204 MHz HFC Plant	39
5.3. Matching 10G PON on a 1794 / 396 MHz ESD Plant	43
5.4. Other 1794 MHz ESD and FDX Options	47
6. Conclusion	50
Abbreviations	52
Bibliography & References	53

List of Figures

Title	Page Number
Figure 1 – Example of Upstream Capacity Usage.....	6
Figure 2 – Mapping Traffic Eng Formula to US Capacity Usage Example.....	8
Figure 3 – Network Capacity Example for 1G Service Tier	8
Figure 4 – Mapping Traffic Eng formula unto SG Probabilities.....	10
Figure 5 – DS TavG, Average Subscriber Downstream Consumption	11
Figure 6 – US TavG, Average Subscriber Upstream Consumption	12
Figure 7 – TavG, Downstream to Upstream DS:US Ratio	12
Figure 8 – DS TavG YoY & 3-yr CAGR for 2010-22	13
Figure 9 – DS TavG Growth Trendline – 2010-2018.....	14
Figure 10 – US TavG YoY & 3-yr CAGR for 2010-22	15
Figure 11 – US TavG Growth Trendline – 2010-2017.....	16
Figure 12 – DS TavG – Estimating COVID bump in ‘21	17
Figure 13 – US TavG – Estimating COVID bump in ‘21	17
Figure 14 – DS TavG Exponential Trendline – 2018-2022.....	18
Figure 15 – DS TavG Exponential Trendline – 2016-2022.....	19
Figure 16 – DS TavG Linear Trendline – 2017-2022	19
Figure 17 – DS TavG Order of 2 Polynomial Trendline – 2016-2022	20
Figure 18 – DS TavG Power Trendline – 2017-2022	20
Figure 19 – DS TavG Logarithmic Trendline – 2017-2022.....	21
Figure 20 – The Technology S-Curve and Equation	22
Figure 21 – Prescriptive S-Curve Strategy	23
Figure 22 – Mapping a Single S-curve to 2010-2020 DS TavG	24
Figure 23 – Mapping a Single S-curve to 2010-2020 DS TavG thru 2040.....	24
Figure 24 – Mapping a Single S-curve to 2010-2022 DS TavG	25
Figure 25 – Mapping a Single S-curve to 2010-2022 DS TavG thru 2040.....	25
Figure 26 – Mapping two S-curves to 2010-2022 DS TavG.....	26
Figure 27 – Mapping two S-curves to 2010-2022 DS TavG thru 2040	26
Figure 28 – Mapping three S-curves to 2010-2022 DS TavG	27
Figure 29 – Mapping three S-curves to 2010-2022 DS TavG (Zoomed In)	27
Figure 30 – DS TavG Growth Projections for 2022 to 2032.....	29
Figure 31 – DS TavG Growth Projections for 2022 to 2037.....	29
Figure 32 – US TavG Exponential Trendline – 2016-2022.....	30
Figure 33 – US TavG Exponential Trendline – 2018-2022.....	30
Figure 34 – US TavG Linear Trendline – 2017-2022	31
Figure 35 – US TavG Order of 2 Polynomial Trendline – 2016-2022	32
Figure 36 – US TavG Order of 2 Polynomial Trendline – 2018-2022	32
Figure 37 – US TavG Power Trendline – 2017-2022	32
Figure 38 – Mapping a Single S-curve to 2010-2022 US TavG	33
Figure 39 – Mapping a Single S-curve to 2010-2022 US TavG (Zoomed Out to 2040).....	34
Figure 40 – Mapping a Double S-curve to 2010-2022 US TavG (Zoomed Out to 2040).....	34
Figure 41 – US TavG Growth Projections for 2022 to 2032.....	35

Figure 42 – US TavG Growth Projections for 2022 to 2037	36
Figure 43 – Max Subs per SG for Low, Moderate & High DS TavG growth, 1218/204 MHz.....	40
Figure 44 – 1218/204 MHz Max subs / SG vs. various DS TavG projection trendlines	41
Figure 45 – 1218/204 MHz System – Spectrum Utilization	41
Figure 46 – 1218/204 MHz System – DOCSIS DS Usage: Tmax, TavG, IP Video	42
Figure 47 – 1218/204 MHz System – DOCSIS US Usage: Tmax, TavG, IP Video	43
Figure 48 – Max Subs per SG for Low, Moderate & High DS TavG growth, 1794/396 MHz.....	44
Figure 49 – 1794/396 MHz Max subs / SG vs. various DS TavG projection trendlines	45
Figure 50 – 1794/396 MHz System – Spectrum Utilization	45
Figure 51 – 1794/396 MHz System – DOCSIS DS Usage: Tmax, TavG, IP Video	46
Figure 52 – 1794/396 MHz System – DOCSIS US Usage: Tmax, TavG, IP Video	47
Figure 53 – DS TavG Growth Projections for 2022 to 2037	51
Figure 54 – Max Subs per SG for Low, Moderate & High DS TavG growth, 1794/396 MHz.....	51

List of Tables

Title	Page Number
Table 1 – TavG Growth Projections: DS, US & DS:US	37
Table 2 – Network Capacity Model – Service Tier mix and 3.0/3.1 mix	39
Table 3 – DOCSIS 4.0 Maximum Service Tiers and Max Subs per SG	48

1. Introduction

Internet consumption grew exponentially last decade, doubling every other year. Several years ago, the growth rate began slowing, but COVID turned the world on its head with everyone living, working, and playing from home. This created a bandwidth (BW) bubble, especially upstream. Will this bubble continue, signaling a new paradigm, or return to our previous path?

ARRIS/CommScope has the most extensive broadband capacity monitoring history in the industry. Data collection started in 2010; done every year since; and covers 10's of millions of modems from numerous multiple system operators (MSOs). The 2022 data is in: this paper analyzes the new normal, quantifies the '21 COVID bump, and shows our recent consumption or average bandwidth per sub (Tavg) growth rates.

The real multi-billion-dollar question is what's the BW growth for coming decades? This drives our network investment strategies. Has Tavg growth slowed to a lower rate (e.g., doubling every 3-4 years) or is it no longer exponential? Some folks claim exponential growth is dead, and that BW growth is linear or following the Adoption S-curve. [S-curves have exponential growth in early years, linear growth during middle years, and flattens out in later years.] E.g., will Tavg reach a limit of 2½ ultra high-definition (UHD) streams per home (~25Mbps) and stay there? Or will another S-curve, potentially driven by virtual reality (VR) / augmented reality (AR), start a new era of high growth?

The paper highlights research on all alternatives; creates trendlines for each; and measures how accurately it matches last decade's data. These BW growth trajectories are mapped out for 5/10/15 years. The resultant spaghetti plots show a cone of uncertainty that grows over time, roughly doubling every 5 yrs.

The 2nd half of the paper plugs various Tavg growth trendlines into the CommScope network capacity modeling tool to analyze 1218/204 megahertz (MHz) & Data Over Cable Service Interface Specifications (DOCSIS[®]) 4.0 plant. It shows their useful lifetime using various growth models and explores whether low growth might eliminate the need for 4.0, or just delay it, giving precious time to transition.

Our wrap up recommends some migration strategies to minimize up front investments while maintaining flexibility to increase network capacity and manage uncertainty risks.

2. Network Capacity Planning Overview

Determining the capacity requirements for a service group (SG) is critical for providing customers with the appropriate quality of experience (QoE). Figure 1 shows an example of a SG upstream (US) capacity usage over a 16-minute window sampled at 1-second intervals. If an operator samples the BW usage once every 16-minutes, then it determines the average BW over that interval as shown by the purple line. This is a very useful datapoint but insufficient to determine the required capacity for the SG.

Sampling at 1-minute intervals would capture some of the variations, or ripples in the system but would still miss the many spikes which are individual modems bursting on top of the average BW. The gray line in the figure shows the high-water mark from 1-minute samples.

If an operator could sample at 1-second intervals or faster, then they could more closely determine the required SG capacity. The 1-second high-water mark is shown by the orange line. But sampling at these rates is not feasible in real systems, so another method is needed.




Figure 1 – Example of Upstream Capacity Usage

2.1. The “Basic” Traffic Engineering Formula

The CommScope (formerly ARRIS) team has provided industry leading traffic engineering research for over a decade. Originally, [CLO_2014] introduced QoE for broadband networks and developed a relatively simple SG traffic engineering formula that’s easy to understand and useful for demonstrating basic network capacity components. Over time, this evolved and a network capacity analysis [ULM_2019] gave an updated insight into calculating the SG capacity requirements. Some additional references of note include [EMM_2014], [ULM_2014], [CLO_2016], [ULM_2016], [CLO_2017], [ULM_2017] and [HOWALD_2022].

The “Basic” formula shown below is a simple two-term equation. The first term ($N_{sub} * T_{avg}$) allocates bandwidth capacity to ensure that the aggregate average bandwidth consumption generated by the number of subscribers (N_{sub}) can be adequately carried by the service group’s bandwidth capacity. The first term is viewed as the “Direct Current (DC) component” of traffic that tends to exist as a continuous flow of traffic during the peak busy period. The growth rate of T_{avg} has seen much research.

“COMMSCOPE/CLOONAN’S CAPACITY EQUATION” Traffic Engineering Formula:

$$C \geq (N_{sub} * T_{avg}) + (K * T_{max_max}) \quad (1)$$

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 period

K is the QoE constant (larger values of K yield higher QoE levels)...

where $0 \leq K \leq \text{infinity}$, but typically $1.0 \leq K \leq 1.2$

T_{max_max} is the highest Service Tier (i.e., T_{max}) offered by the MSO

Figure 1 shows there are obviously fluctuations that occur (i.e., the “alternating current (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 ($K * T_{max_max}$) is added to increase the probability that all subscribers experience good QoE levels for most of the fluctuations that go above the DC traffic level.

The second term in the formula ($K * T_{max_max}$) has an adjustable parameter defined by the K value. This parameter allows MSOs to increase the K value and add bandwidth capacity headroom that provides 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 service tier being offered to subscribers.

In previous papers [CLOONAN_2013, CLOONAN_2014, EMM_2014], found that a K value between 1.0 to 1.2 provides good QoE results for a 250 subscriber SG. Larger SGs need even larger values of K while very small SGs might use a K value ≤ 1.0 .

2.2. The “Modified” Traffic Engineering Formula

Over time, it was discovered that the optimum K value varies based on all the inputs: N_{sub} , T_{avg} and T_{max_max} . [ULM_2017] noted some of these limitations along with some refinements to the basic formula above. This resulted in the following which is still algebraically equivalent to the basic formula:

Modified “COMMSCOPE/CLOONAN’S CAPACITY EQUATION” Traffic Eng Formula:

$$C \geq (N_{sub} * T_{avg}) + (K-1) * T_{max_max} + T_{max_max} \quad (2)$$

The subtle change is that there are now three main components to the traffic engineering formula:

1. Peak Busy Period Average Consumption (i.e., $N_{sub} * T_{avg}$)
2. Peak Busy Period Ripple for managing QoE (i.e. $(K-1) * T_{max_max}$)
3. Headroom for maximum Service Tier Burst (i.e., $1 * T_{max_max}$)

Figure 2 shows how the modified formula maps to the US capacity usage example given in Figure 1. The basic formula might have used a value of $K=1.2$ in this example. This is now broken into a burst component equal to $Tmax_max$ plus a ripple component that is estimated by $20\% * Tmax_max$.




Figure 2 – Mapping Traffic Eng Formula to US Capacity Usage Example

Figure 3 shows a network capacity example for a 1 Gbps downstream (DS) service tier; 150 subs in the SG; and $Tavg = 3$ Mbps/sub. This requires SG capacity ≥ 1500 Mbps. This includes a ripple component of $5\% * Tmax_max = 50$ Mbps. Note that in this example, the ripple component as a function of $Tmax_max$ is much smaller than the initial $K=1.2$ values because $Tmax_max$ is significantly higher.




Figure 3 – Network Capacity Example for 1G Service Tier

2.3. Max Service Tiers in 10G Era

So, what kind of service tiers will subscribers enjoy in this new 10G bandwidth era? These tend to be marketing driven. The T_{max} value from the traffic engineering formula helps define the Service Level Agreement (SLA) that the operator can potentially offer to their customers.

First, 10G passive optical network (PON) provides a net downstream capacity of ~8.5 gigabits (Gbps) to the consumer. Using the traffic engineering formula, this capacity might support a downstream SLA of 7.5 Gbps. The SG utilization (i.e., $N_{sub} * T_{avg}$) for a 64 subscriber PON might grow to a 1+ Gbps over the next decade. That means a consumer with a 7.5 Gbps SLA will have a QoE coefficient of $K = \sim 1.0$ which is reasonable for this relatively small SG size.

The 7.5 Gbps SLA from 10G PON sets a bar that Hybrid Fiber Coax (HFC) systems must match with their DOCSIS 4.0 upgrades. Some network capacity modeling results are shown in a later section to show how many subscribers the HFC can support per SG at this SLA.

HFC does have an advantage because it can incrementally add capacity with additional spectrum to achieve a true 10 Gbps SLA that is equivalent to 10G Ethernet. Getting to a true 10 Gbps downstream SLA will mean providing greater than 10 Gbps network capacity. This will push the PON networks into next generation PON technology (e.g., 20+ Gbps). Some future technologies such as 1.8 and 3.0 GHz HFC plants are discussed further in [CLO_2019].

Choosing the upstream SLA is more complicated. As will be seen later in Figure 7 with the DS:US consumption ratio, there might be a 12:1 ratio between the two. However, in the new 10G era, a gigabit US SLA tier with high burst rates may be needed, even if the US consumption is much lower than downstream.

Looking at PON systems, they offer both symmetric and asymmetric data rates. A gigabit passive optical network (GPON) provides 2.5 Gbps downstream data rates with 1.2 Gbps upstream data rates for a 2:1 ratio. The Institute of Electrical and Electronics Engineers (IEEE) 10G ethernet passive optical network (EPON) downstream might be paired with either a 1G or 10G upstream for 10:1 or 1:1 ratio. In the International Telecommunication Union (ITU) world, XG-PON pairs 10 Gbps downstream with 2.5 Gbps upstream (i.e., 4:1 ratio) while XGS-PON provides a symmetric 10 Gbps in both directions for 1:1 ratio.

HFC systems have traditionally been extremely asymmetric, but these trends are changing. In the upcoming network capacity modeling section, a 2.5 Gbps upstream SLAs is paired with the 7.5 Gbps DS SLA on a 1794/396 MHz plant.

2.4. Determining Service Group QoE based on Probabilities

The Peak Busy Period Average Consumption and maximum Service Tier Burst components are well known and easily obtained. Much traffic engineering research has since focused on quantifying the Peak Busy Period Ripple component for QoE as it is impacted by all inputs, not just T_{max_max}.

It became apparent that quantifying the SG subscribers QoE needed to focus on the probabilities for SG capacity. And to predict behavior across different SG with different parameters, a network capacity transmit model for individual subscribers was necessary.

Our research found 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. 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 networks will morph and the QoE margins of the future. It is important to note that the operator can choose how much margin they would like to build in.

Once single subscriber BW probability distribution functions (PDF) were in place, it is then possible to create a SG BW PDF to analyze the behavior at the SG level. An example of the SG BW DS probabilities is shown in Figure 4. This is the output of a Monte Carlo simulation with 100K trials. It is a SG that consists of 128 subs, all of which have a 1G DS service tier. The DS Tavg = 15 Mbps which represents a time in the future.




Figure 4 – Mapping Traffic Eng formula unto SG Probabilities

Note the asymmetry in the curve – there is a much longer tail to the right. Various cumulative distribution function (CDF) probability thresholds are shown (i.e., 90%, 99%, 99.9%, 99.99%) to provide an insight as to the probability the SG BW reaches different capacities.

The blue candlestick in Figure 4 shows how the traffic engineering formula overlays the SG PDF. In this instance, the formula forecasts 3,120 Mbps is required. Note that this is sufficient for ~99.5% of the peak busy period time. For the remaining 0.5% of the time, the buffers may be temporarily filled, and some latencies introduced. Note that 0.5% represents ~50 seconds out of every evening. For most of those 50 seconds, the delays should be insignificant and not noticed by users. This is normal network behavior and why many users can share a single network pipe. Please keep in mind how low the probabilities are for the Tmax_max burst regions as shown above.

Some new revelations of CommScope's probability research are discussed further in [HOWALD_2022].

3. Broadband Subscriber Traffic Consumption – 2022 Update

ARRIS/CommScope has the most extensive broadband capacity usage monitoring history in the industry. Data collection started in 2010; done at the start of every year since; and covers 10's of millions of modems from multiple MSOs. Subscriber usage has been monitored from the same group of MSOs for consistency. This dataset has been compared and maps closely to many other MSOs globally.

3.1. Downstream Peak Period Average Bandwidth per Subscriber, DS Tavg

Figure 5 shows the average subscriber downstream consumption, DS Tavg, during peak busy period for several MSOs over a ten-year period. At the start of 2022, DS Tavg was 3.5 Mbps per sub averaged across all the MSOs.




Figure 5 – DS Tavg, Average Subscriber Downstream Consumption

It turns out that the DS Tavg growth rate was much higher last decade and continues to tail off in recent years. The DS Tavg is up only 12% from 2021. Looking across a 3-year window, the DS Tavg has grown roughly 21% per year. That is half the compounded annual growth rate (CAGR) from five years ago.

3.2. Upstream Peak Period Average Bandwidth per Subscriber, US Tavg

The 2022 US Tavg is almost 300 Kbps as shown in Figure 6. This is up 14.5% from the 2021 COVID inflated numbers. The 3-year CAGR for US Tavg is 28%, which is higher than DS Tavg CAGR for the first time and up from five years ago.

3.3. Downstream to Upstream Peak Period Average Bandwidth Ratio

The DS:US ratio as shown in Figure 7 peaked at 14.4 to 1 ratio in 2020. As of 2022, the average DS:US ratio in this post-COVID new normal seems to have stabilized around 12:1. There was significant variation between the different MSOs with a range of 8:1 to 18:1 ratio.




Figure 6 – US Tavg, Average Subscriber Upstream Consumption




Figure 7 – Tavg, Downstream to Upstream DS:US Ratio

4. Broadband Traffic Growth Trendlines

As operators plan their network migration for the next decade or two, Tavg growth becomes a critical component in that planning. This drives our network investment strategies. Growth rates appear to be changing dramatically of late and have cast a lot of uncertainty around future consumption growth rates.

The real multi-billion-dollar question is what will be the BW consumption growth for coming decades? Has Tavg growth slowed to a lower CAGR (e.g., doubling every 3-4 years) or is it no longer exponential? Some folks claim exponential growth is dead, and that BW growth is linear or following the Adoption S-curve. S-curves have exponential growth in early years, linear growth during middle years, and flattens out in later years. E.g., will Tavg reach a limit of 2½ UHD streams per home (~25Mbps) and stay there? Or will another S-curve (e.g., VR/AR driven) start a new era of high growth?

Before looking at each possible growth trendlines, it is useful to review last decade's growth patterns.

4.1. Broadband Traffic Growth in retrospect

4.1.1. DS Tavg Growth – 2010-22

The year-over-year (YoY) changes in DS Tavg are shown as the blue candlesticks in Figure 8. Through 2018, there were many years with ~30% YoY change and then a couple years with larger 50%-70% YoY change. The giant spurts in DS Tavg often coincided with new DOCSIS technology improvements (e.g., going from 8 to 16 or 16 to 32 bonded 3.0 single carrier quadrature amplitude modulation (SC-QAM) channels).

After 2018, things started to change. The 2019 and 2020 YoY changes were only ~20%. During the COVID lockdown of 2021, DS Tavg only went up ~32%. Then 2022 DS YoY clocked in at only 12%.




Figure 8 – DS Tavg YoY & 3-yr CAGR for 2010-22

Given the volatility of the YoY data, the CAGR for a sliding 3-year window was investigated. This is shown by the orange candlesticks in Figure 8. The 3-year CAGR hovered around 40% last decade, but then started dropping noticeably after 2018 as shown by the red arrow. The fact that the 3-year window is constantly dropping implies that the DS TavG growth rate is either dropping to a much lower percent or it is no longer exponential. The next section looks at these possibilities.

In any respect, 2018 appears to be an inflection point for DS TavG growth. There is no known obvious reason as to what caused this change. It started well before COVID and appears to be continuing in our new normal post-lockdown world.

Using the powers of Excel, multiple types of trendline were matched against DS TavG data from 2010 to 2018. For each trendline, Excel's R^2 (R-squared) metric is also shown, as an estimate of trendline's "goodness of fit". R-squared quantifies a trendline's percentage of variation compared to the actual subscriber consumption, TavG, over the given length of time [Investopedia R2], [Wiki R2], [Excelltip R2]. In cases where R-squared was not an available option on the graph, excel "RSQ" function was used to produce the R-squared value shown. (An alternative way to calculate it is to use the squared value of the correlation coefficient [CORREL], which is a covariance of the trendline with data, divided into the product of standard deviation of trendline times standard deviation of data).

The value of R-squared varies between 0 and 1, and closer to 1 it gets, better the trendline explains the data. A line fit with $R\text{-squared} = 0.9959$, for example, means that 99.59% of the variation is explained by the trendline, and this is valid in the time interval considered. However, projecting these trendlines outside of the time where data exists is best done with a "caveat emptor" (buyer beware!) warning – there are no guarantees the adherence to the trendline will continue in the future. This is analogous to the stock market adage of "past performance not indicative of future results"!




Figure 9 – DS TavG Growth Trendline – 2010-2018

Analyzing the 2010-2018 DS Tavgr data creates an exponential trendline that resulted in a 43.5% CAGR DS over that 8-year window. That equates to doubling in slightly less than every other year. This is shown in Figure 9. For the 2010-2018 DS data in Figure 9, R-squared (R^2) = 0.9959 which is an excellent match of better than two 9's of accuracy. So, in 2018, the DS Tavgr growth projection would have been:

- 2018 DS Growth projection => DS Tavgr = **100 Mbps/sub by 2030**

At this consumption growth rate, many people thought that operators would need Fiber to the Premise (FTTP) for all their subscribers by then.

4.1.2. US Tavgr Growth – 2010-22

The YoY changes in US Tavgr are shown as the purple candlesticks in Figure 10. The 3-year sliding window CAGR is shown with the pink candlesticks.

Up until 2018, the US YoY would bounce from 10% or less to the 20%-25% range. The 3-year sliding window smooths out some of that YoY variation and showed an US CAGR around 17%. The US data saw some large YoY spikes >40% in 2018 and then again during the 2021 COVID lockdown. However, the other years (i.e., 2019, 2020, 2022) showed less than 20% US growth. The US 3-year sliding window has shown a definite uptick and has been in the 20% to 30% range since 2018.




Figure 10 – US Tavgr YoY & 3-yr CAGR for 2010-22

The 2010-2017 US Tavg growth trendline was 17.3% CAGR as shown in Figure 11. There is much more volatility in the US data, so US $R^2 = 0.9805$ is not nearly well matched as the DS. If the 2018 spike is introduced into this data, then R-squared became significantly worse.

At this point in time, the US Tavg growth projection was:

- 2017 US Growth projection => US Tavg = <1 Mbps/sub by 2030

The US Tavg consumption was not on anyone's radar in 2017. However, this projection also implied that the DS:US ratio would exceed 100:1. [ULM_2017] speculated whether the DS:US ratio would "stabilize" around 15:1, but it wasn't clear if that would happen due to DS growth slowing or US growth increasing. It turns out that both happened.




Figure 11 – US Tavg Growth Trendline – 2010-2017

4.2. Correcting for the COVID Bump

The Coronavirus Bandwidth Surge wreaked havoc on our long-term consumption growth planning. Growth patterns started shifting pre-COVID in 2019 and early 2020, but then things were turned upside down during lockdown with our January 2021 Tavg numbers. [ULM_2021] spelled out this impact on cable networks.

The 2022 Tavg numbers represents the first data points being seen since a return to the "new normal". Some of the Tavg impacts from the COVID lockdown were anomalies that need to be disregarded while other impacts appear to be a permanent shift. To calculate our current growth trendlines requires the anomalies to be factored out.

Ignoring the 2021 data for a second, DS Tavg increased 48% from pre-COVID Jan '20 to post-lockdown new normal in Jan '22. That maps to a 21.8% growth each year if not for the COVID bump. The blue candlestick in Figure 12 shows the estimated 2021 DS Tavg without the COVID bump. The pink candlestick on top of the 2021 data shows the size of DS anomalies estimated during the lockdown. Note that this was less than 10% additional DS traffic. For our growth trendline analysis, the blue candlesticks are used, and the COVID bump anomalies shown in pink will be ignored.




Figure 12 – DS Tavg – Estimating COVID bump in '21




Figure 13 – US Tavg – Estimating COVID bump in '21

Figure 13 shows that the US Tavg grew 80.5% from pre-COVID Jan '20 to post-lockdown new normal in Jan '22. That maps to a 34.35% growth each year if not for the COVID bump. The blue candlestick in Figure 13 shows the estimated 2021 US Tavg without the COVID bump. The pink candlestick on top of the 2021 data shows the size of US anomalies estimated during the lockdown. Note that this was almost 20% additional US traffic, much higher than the DS. For our growth trendline analysis, the blue candlesticks are used, and the COVID bump anomalies shown in pink will be ignored.

4.3. Various Downstream Growth Trendlines

4.3.1. DS Exponential Growth Trendline

Conventional wisdom for the last couple decades was that broadband traffic grows exponentially. That is the first growth trendline considered. Since 2018 appears to be an inflection point, Figure 14 shows the exponential growth trendline from 2018 to 2022. This accurately maps to a 21.1% CAGR DS, which is Tavg doubling roughly every 3.75 years. The data was extremely well matched with DS $R^2 = 0.9998$, almost four 9's accuracy.




Figure 14 – DS Tavg Exponential Trendline – 2018-2022

By comparison, Figure 15 looks at the DS Tavg exponential for the 2016-22 span. Including the 2018 spike certainly skews the data. This results in a 27.6% CAGR but DS $R^2 = 0.9789$ shows a much poorer fit than the 2018-2022 trendline in Figure 14. Visual inspection of Figure 15 also shows how the trendline is obviously pulling away from the actual 2022 datapoint. Another reason why the 2018-2022 CAGR is the exponential trendline of choice for our DS Tavg growth projections.




Figure 15 – DS Tavg Exponential Trendline – 2016-2022

4.3.2. DS Linear Growth Trendline

Recent discussions have speculated whether the DS Tavg has entered a phase of linear growth rather than exponential growth. The 2017-2022 linear trendline analysis is shown in Figure 16. Various time windows were considered starting in 2016, 2017 and 2018. The 2017-2022 had the best linear match with the DS $R^2 = 0.9912$. While this is not as close as the exponential match above, it is still better than two 9's of accuracy. This is a viable contender and should be considered as one of the possible DS Tavg growth projections.




Figure 16 – DS Tavg Linear Trendline – 2017-2022

4.3.3. Other DS Growth Trendlines

Since the door was opened to look at trendlines other than exponential, our research also looked at polynomial, power and logarithmic trendlines to see how well those might match.

The 2016-22 order of 2 polynomial trendline is shown in Figure 17. It had a very good match with DS $R^2 = 0.997$. This was better than linear but not as good as exponential trendlines. For polynomial trendlines, starting in 2016 had a better match than 2017 or 2018. This is another option to consider for our long-term growth projections.




Figure 17 – DS Tavg Order of 2 Polynomial Trendline – 2016-2022




Figure 18 – DS Tavg Power Trendline – 2017-2022

The 2017-2022 power trendline is shown in Figure 18. It had a very good match with $DS R^2 = 0.9907$. This match tracked the linear trendlines and is another option for consideration. For power trendlines, a starting window of 2017 provided the best results.

The logarithmic trendline is shown in Figure 19. This was a relatively poor match with $DS R^2 = 0.9322$. This trendline is not a candidate for our long-term growth projections.




Figure 19 – DS Tavlg Logarithmic Trendline – 2017-2022

4.3.4. DS Adoption Curve (S-curve) Growth Trendline


One recent school of thought is that broadband consumption is following an Adoption Curve, a.k.a. S-curve, that other technologies have followed. The S-curve theory has been around for multiple decades. Some overviews are given in [CHR_1992] and [ARK_BLOG].

In S-curve theory, there is a period of exponential growth rate in the early years. This then levels off to linear growth rates in the middle years which is followed by slowing to zero growth in the latter years. Figure 20 shows this form in a chart from [CHR_1992] with the basic S-curve equation below it. We inserted the red arrow to show where the DS Tavlg might currently sit on a S-curve.

The rationale for broadband traffic following an S-curve goes like this:

1. From 2010 to 2018 – video usage over the internet grew rapidly due to applications like Netflix and YouTube. Video penetration and video bit rates increased extensively resulting in exponential growth year to year.
2. From 2018 to present – video use on the internet approaches 80%, so slowing increase in video streams results in slower growth rates. The net result is linear growth.
3. From present to late 2020's – video bit rates continue a slow shift towards 4K Ultra-HD resolution; operators migrate to IP video. Growth rates continue to decline.
4. From late 2020's into 2030's – video penetration reaches saturation with no more growth. Growth rates are very slow to zero.

As an example, the DS Tavg might peak once operators deliver 2½ unique 4K streams to every home, so every customer has their own private UHD video stream. This might lead to the DS Tavg point being ~25 Mbps/sub in 10-20 years.



$$f(x) = \frac{L}{1 + e^{-k(x-x_0)}}$$

$f(x)$ = output of the function

L = the curve's maximum value

k = logistic growth rate or steepness of the curve

x_0 = the x value of the sigmoid midpoint

x = real number

Figure 20 – The Technology S-Curve and Equation

In reality, one technology wave may be followed by another wave, and then another. [CHR_1992] showed this possibility in a chart that is copied in Figure 21. As an example, maybe Netflix + YouTube adoption was part of the 1st S-curve which has reached maturity. A second S-curve is underway that is a migration to managed IP video delivery with 4K UHD streams. A third S-curve based on Augmented Reality (AR) and Virtual Reality (VR) is in the early stages and goes mainstream in another decade.




Figure 21 – Prescriptive S-Curve Strategy

For our research, models were created for a single, double and triple S-curves. A best-fit algorithm was then run to match these curves to the 2010-2022 DS Tav_g datapoints. Figure 22 shows two possible single S-curve growth trendlines mapped to the 2010-2020 DS Tav_g (i.e., before COVID). The actual DS Tav_g is the blue curve. The green single S-curve was our original S-curve model that assumed an upper asymptote of 26 Mbps/sub. In 2020, it had a reasonably good match with DS $R^2 = 0.9903$.

The purple dotted curve in Figure 22 shows another single S-curve fit. This one did not make any initial assumptions about the upper limit and used a best-fit algorithm to optimize DS R-squared values. This resulted in an even better DS $R^2 = 0.9963$. But notice in Figure 23 how much both S-curves diverge even though both have pretty good matches. One hits a limit of 26 Mbps and the other 4 Mbps. This illustrates the sensitivity in making these projections.




Figure 22 – Mapping a Single S-curve to 2010-2020 DS Tav_g




Figure 23 – Mapping a Single S-curve to 2010-2020 DS Tav_g thru 2040

As the 2021 and 2022 datapoints are added to the blue curve, it plots a path that is right in between these two 2020 single S-curve projections. The best-fit single S-curve was run again with 2010-2022 data and the results are shown as the light blue dotted curve in Figure 24. S-curve '22 DS R^2 improved further to 0.9975, while the two '20 projections got worse. With this update, the upper limit increased 33% to ~5.4 Mbps/sub. This is still significantly lower than the original green single S-curve trendline in Figure 25.




Figure 24 – Mapping a Single S-curve to 2010-2022 DS TavG




Figure 25 – Mapping a Single S-curve to 2010-2022 DS TavG thru 2040

The next step makes a best-fit growth trendline combining two S-curves as shown in Figure 26. The midpoint of each S-curve is the year 2018 and 2028 respectively while the other parameters are optimized. There is an excellent fit with DS $R^2 = 0.9982$. The individual S-curves are dotted lines with the sum of the two represented by the dashed dark blue curve. The double S-curve reaches an upper limit of ~28 Mbps, very close to the original green single S-curve which is also in Figure 26 and Figure 27.




Figure 26 – Mapping two S-curves to 2010-2022 DS Tavg




Figure 27 – Mapping two S-curves to 2010-2022 DS Tavg thru 2040

Figure 28 and Figure 29 shows the best-fit results for a trendline using three S-curves. The mid-point is set for years '18, '28 and '38 respectively and the other parameters are optimized. There is a good fit again with DS $R^2 = 0.9982$. The individual S-curves are dotted lines with the sum being the dashed dark blue curve. Figure 28 zooms in on the triple S-curve scenario to give a better look at how it matches the 2010-22 data. It also clearly shows the contribution of each of the three individual S-curves over time.




Figure 28 – Mapping three S-curves to 2010-2022 DS Tavg




Figure 29 – Mapping three S-curves to 2010-2022 DS Tavg (Zoomed In)

The original green single S-curve is also included. The triple S-curve remains close to it through 2030 as seen in Figure 28, but then the two diverge significantly during that next decade as shown in Figure 29. This divergence is caused by the third S-curve ramping up. This might represent the adoption of VR/AR technology that takes off in the 2030's. The triple S-curve eventually reaches an upper limit of ~70 Mbps in 20+ years.

4.4. Mapping DS Tavg Growth Projections over a 10-15 Year Span

Now it is time to pull all these growth projections together to see how they compare. At this stage, there is not enough evidence to warrant one option over another. A handful of the trendlines at the extremes with a poor fit have been disregarded (i.e., 2016 exponential, power, log trendlines). All the DS Tavg growth trendlines for years 2017-2032 are plotted in Figure 30. The red circles show the growth trendlines of interest and represent the amount of uncertainty between the different projections.

In five years, 2027, DS Tavg is expected to be in 5-12 Mbps/sub range. A decade from now, 2032, that increases to a 7-25 Mbps/sub range. By considering these multiple potential growth trendlines, there is now a “cone of uncertainty” that grows over time. Notice that the window of uncertainty has almost doubled from five to ten-year window.

Figure 31 now extends the growth trendlines out 15 years to 2037. It resembles a spaghetti plot for a hurricane path! By 2037, the DS Tavg is 8-65 Mbps and the cone of uncertainty has doubled once again.

In an earlier section, it was noted that in 2018 with 43.5% CAGR, the expected DS Tavg would reach 100 Mbps by 2030. Looking at our current growth projections, the high growth projection (21% CAGR) hits 100 Mbps in 2040; while the low growth linear projection takes 200+ years to reach that!!!

From a network capacity planning perspective, our conclusions on considering multiple growth trendline options are:

- the 5-year window gives us reasonably high confidence for near-term planning
- the 10-year window gives us a range of high, moderate, and slow growth scenarios for longer term planning
- the 15-year window shows too much variance to be used for network capacity planning and is more of an academic exercise.

The authors view these growth trendline projections as an on-going process that gets updated and adjusted every year.




Figure 30 – DS TavG Growth Projections for 2022 to 2032




Figure 31 – DS TavG Growth Projections for 2022 to 2037

4.5. Various Upstream Growth Trendlines

Upstream consumption growth over the last 4-6 years has been very different from DS growth. US growth rates have increased while DS growth has declined considerably. Upstream consumption has also been more volatile than DS, with very large spikes in 2018 and again with the 2021 COVID lockdown.

4.5.1. US Exponential Growth Trendline

The 2016-22 US exponential trendline is shown in Figure 32 while the 2018-22 period is in Figure 33.




Figure 32 – US Tavg Exponential Trendline – 2016-2022




Figure 33 – US Tavg Exponential Trendline – 2018-2022

The 2016-22 exponential growth trendline has a 23.7% CAGR with US $R^2 = 0.9726$. The 2018-22 exponential growth trendline has a 22.1% CAGR with US $R^2 = 0.9244$. The 2016-22 trendline is a much better fit and the one that will be used as our high growth exponential scenario. Including the extra two years made for a better fit this time.

4.5.2. US Linear Growth Trendline

Figure 34 shows the US linear growth trendline. The US $R^2 = 0.9195$ which is a relatively poor fit. Linear growth assumes a slowing growth rate, but the upstream has seen an increased growth rate. As such, this trendline is not a candidate for consideration at this time.




Figure 34 – US Tavg Linear Trendline – 2017-2022

4.5.3. Other US Growth Trendlines

The order of 2 polynomial growth trendline does provide a good fit. The 2016-22 trendline is shown in Figure 35 with the 2018-22 trendline in Figure 36. The 2018-22 trendline has the better fit with US $R^2 = 0.9948$ compared to US $R^2 = 0.9737$. It is even a better fit than the 2016-22 exponential trendline.

The 2016-2022 US power trendline is shown in Figure 37. It had a very poor match with US $R^2 = 0.8984$. The logarithmic trendline was even worse and is not even shown. These trendlines are not candidates for our long-term growth projections.




Figure 35 – US Tavg Order of 2 Polynomial Trendline – 2016-2022




Figure 36 – US Tavg Order of 2 Polynomial Trendline – 2018-2022




Figure 37 – US Tavg Power Trendline – 2017-2022

4.5.4. US Adoption Cure (S-curve) Growth Trendline

As noted above, the upstream growth behavior is very different than downstream one. The 10G initiative will be opening up a world of new applications. Many of these could start driving upstream bandwidth consumption further. This may come from a plethora of IOT devices in the home, or maybe more affordable HD resolution video cameras pushing content to the cloud.

Figure 38 and Figure 39 show the mapping of a single S-curve to the 2010-22 US Tavg data. The first figure is a view up until 2028, while the next figure zooms out to 2040. The US Tavg data is shown both with and without the COVID bump. The trendline without the COVID bump provides the best results with $US R^2 = 0.9815$, a relatively good fit for the upstream.

The upper bound on the US S-curve is 3.71 Mbps/sub. Note that the single S-curve midpoint is the year 2033. From an S-curve perspective, the US consumption is still in its early years of exponential growth.

Our modeling then calculated a best fit for two S-curves shown in Figure 40. The second curve ended up with a midpoint that was 20 years into the future and has minimal impact on the present day (i.e. <0.1 Kbps). The result is that the single and double S-curve scenarios are nearly identical for the next 12 years. It is almost 15 years before a separation between the two becomes obvious. Again, this shows that mapping out to 15+ years is more of an academic exercise.




Figure 38 – Mapping a Single S-curve to 2010-2022 US Tavg




Figure 39 – Mapping a Single S-curve to 2010-2022 US Tavg (Zoomed Out to 2040)




Figure 40 – Mapping a Double S-curve to 2010-2022 US Tavg (Zoomed Out to 2040)

4.6. Mapping US Tavg Growth Projections over a 10-15 Year Span

The upstream traffic growth projections have fewer options since the slower growth trendlines like linear are not currently under consideration. This obviously could change in the future as more Tavg data is collected and there is a better understanding of the “new normal” in this post-COVID lockdown world.

Figure 41 shows the US Tavg growth projections for the next decade. The trendlines at the lower extremes with a poor fit have been disregarded (i.e., linear, power, log trendlines). The red circles show the growth trendlines of interest and represent the amount of uncertainty between the different projections.

In five years, 2027, US Tavg is expected to be in 0.7-1.1 Mbps/sub range. A decade from now, 2032, that increases to a 1.2-2.5 Mbps/sub range. Notice that the upstream “cone of uncertainty” is smaller than the downstream, but it still has almost doubled as it goes from the five to ten-year window.

Figure 42 now extends the growth trendlines out 15 years to 2037. By 2037, the US Tavg is now 2-7 Mbps/sub, and the cone of uncertainty has roughly doubled once again.




Figure 41 – US Tavg Growth Projections for 2022 to 2032




Figure 42 – US Tav Growth Projections for 2022 to 2037

4.7. DS:US Tav Ratio projections

Now let's put the DS and US Tav growth projections side-by-side. This is shown in Table 1. From here, an estimate for the DS:US ratio can be made over time. The authors believe that there is some interdependence between the DS and US Tav consumption growth, so it is highly unlikely that the DS and US will hit opposite ends of their low and high projections. The low and high ratio range in the table is slightly tempered from these extremes.

With the upstream showing a slightly higher growth trajectory, the mid-range of the DS:US ratio is expected to ease from the current 12:1 down to something closer to 10:1. Over the longer term, there is no evidence that the mid-range of the DS:US ratio will vary much from 10:1.

Table 1 – Tavg Growth Projections: DS, US & DS:US

Tavg Growth Projections	DS Tavg Range		US Tavg Range		DS:US Ratio Range		
Year	Low	High	Low	High	Low	Mid	High
2022	3.5		0.3		12:1		
2027	5	12	0.7	1.1	6:1	10:1	15:1
2032	7	25	1.2	2.5	5:1	10:1	18:1
2037	8	65	2	7	4:1	10:1	20:1

5. Network Capacity Modeling for Low to High Growth Projections

So, how do these BW growth projections drive our network investment strategies? Our next step plugs various Tavg growth trendlines into the CommScope network capacity modeling tool to analyze 1218/204 MHz and DOCSIS 4.0 plant. It shows their useful plant lifetime using various growth projections and explores whether low growth might eliminate the need for 4.0; or just delay it, giving precious time for the operator to transition.

Over recent years, there has been a slowing in the downstream usage growth rate (i.e., DS Tavg) compared to the service tier growth rate (i.e., DS Tmax_max). This has a number of consequences including the network becomes more “bursty”. It also means that the overall utilization of the network is lower. In this respect, it is important to try and maximize subscribers per service group (SG) to take advantage of statistical multiplexing and to get better economics. This analysis focuses on the number of subscribers that a service group can support over time for a given configuration.

5.1. Network Capacity Modeling Assumptions

The CommScope network capacity model contains 100’s of different inputs that can vary from year to year. The paper uses a relatively good to best case scenario that was taken from modeling work done for several operators based on their current network migration plans.

Many HFC plants today are still 870 MHz, 750 MHz or even lower. The model assumes that the HFC plant is upgraded to at least 1218/204 MHz plant by 2024 (and as needed to extended spectrum DOCSIS (ESD) by 2025). The modeling starts with 400 subs per SG as the maximum, then reduces max subs per SG as needed over time. Note that this is an extremely large SG that could represent a node in the 600 to 1,000 homes passed region.

From DOCSIS perspective, it assumes that up to 400 typical data subs per cable modem termination system (CMTS) Service Group (SG) as the starting point. And from a DOCSIS combining perspective – 1 node per CMTS SG for both DS & US to start (e.g., 1x1 Remote MAC-PHY device, RMD). If the upstream ever becomes the limiting factor, then the upstream will be split (e.g., 1x2 RMD). If an operator starts with a 1x1 RMD, then a migration to 2x2 RMD enables SG splits down the road.

Legacy video begins with 60 total quadrature amplitude modulation (QAM) electronic industries association (EIA) channels today. But a rapid migration to 100% IPTV in 2024 reduces that to zero QAM EIA channels. These scenarios assume there are no video on demand (VOD) QAM channels, no Analog video spectrum, and no switched digital video (SDV) QAM channels.

For DOCSIS Services, the beginning includes 32 DOCSIS 3.0 SC-QAM channels. DOCSIS 3.1 (D3.1) modems are used for all the top tiers. Older DOCSIS 2.0/3.0 modems are phased out over time until there is 100% D3.1/4.0 modems by 2025.

The model uses the broadband consumption numbers from this paper. DS Tav_g = 3.5 Mbps in 2022. Modeling runs were then done from 2022 to 2032 for high growth (i.e., 21% CAGR), moderate growth (i.e., 16% CAGR) and low growth (i.e., linear) scenarios. US Tav_g = 300 Kbps in 2022; with 23% CAGR growth from 2022 to 2032. Our findings show that the plant tends to be downstream BW limit so there was no need to consider low to moderate growth upstream scenarios. These result in:

- DS Tav_g High Growth: 21% DS CAGR = 24 Mbps by '32;
- 23% US CAGR = ~2.4 Mbps by '32
- Note that this is still a 10:1 DS:US ratio

The CommScope network capacity model has a lot of flexibility and allows us to adjust Tav_g per service tier (e.g., Top Billboard tier is 2X Flagship tier). The DS Tav_g and US Tav_g listed are the weighted average across all subscribers in the SG.

For the DOCSIS Physical (PHY) layer, the model assumes:

- D3.1 orthogonal frequency-division multiple access (OFDMA) US = 7.58 bps/Hz (1024 QAM, 4K fast fourier transform (FFT), 1.875us cyclic prefix (CP))
- D3.1 orthogonal frequency-division multiplex (OFDM) DS = 9.7 bps/Hz (4096-QAM, 8K FFT, 1.25us CP) up to 1218MHz
- 1794 MHz plant drops to 8.58 bps/Hz on average (2048-QAM, 4K FFT, 1.25us CP)
- 'Normal' amp spacing = ~9.0 bps/Hz; 'Stretch' amp spacing = ~8 bps/Hz

The service tier mix is shown in Table 2. The Top Billboard tier jumps to 5 Gbps DS with a 1 Gbps US in 2024 for the 1218/204 MHz plant. Note that a 5G DS Tier requires four bonded OFDM channels on a D3.1 plant. This can easily be handled with DOCSIS 4.0 modems operating in D3.1 environment. The 1794/396 MHz plant has a Top Billboard tier of 7.5G DS X 2.5G US. Any tiers 500/50 or higher are assumed to be immediately moved to D3.1 modems only.

Table 2 – Network Capacity Model – Service Tier mix and 3.0/3.1 mix

Service Tier	Mix %	2022	2024	2026	2028
Top Billboard	2%	2G/400	5G/1G or 7.5G/2.5G		
Performance	13%	1G/100	2G/200	5G/1G	5G/1G
Flagship	50%	200/20	500/50	1G/100	2G/200
Economy	35%	100/10	200/20	200/20	500/50
% D3.1 modems		50%	75%	100%	100%

Regarding video services, the managed IPTV bundle grows to 100% penetration by 2024 replacing all Legacy Video subs. It assumes 50% of high-speed data (HSD) subs take IPTV bundle (e.g. 150 IPTV subs out of 300 HSD subs). The IPTV service is 100% Unicast delivery with the following mix of video bit rates:

- HD streams: 5 Mbps per today; dropping to 3 Mbps by 2030
- 4K UHD streams: 5% @ 20 Mbps in '23 growing to 50% @ 8 Mbps in 2032

5.2. Making the Most of 1218/204 MHz HFC Plant

With so much interest in slowing DS Tavgr growth rates, the first case study considered standard D3.1 1218/204 MHz plant offering a 5G DS tier with a 1G US tier. What is the maximum sized SG that can be supported over the next decade?

The model was run for high (21% CAGR), moderate (16% CAGR) and low (linear) DS Tavgr growth rates. That leaves DS Tavgr a decade from now at 24, 15 and 8 Mbps/sub respectively. The US Tavgr with a 23% CAGR reaches 2.4 Mbps/sub while the lower growth projections are just above 1.5 Mbps in 2032.

Figure 43 shows the max number of subs per SG for low, moderate, and high DS Tavgr growth projections. All projections can support 400+ subs through 2025. In 2026, the high growth scenario needs to gradually reduce the max subs/SG each subsequent year. It is still supporting 300 subs/SG until 2028, 200 subs/SG through the end of the decade and 150 subs/SG in ten years. The moderate growth rate follows a similar curve only delayed about two years. The slow linear growth is still supporting 400 subs/SG at the end of the decade and 350 subs/SG in 2032, a decade from now.




Figure 43 – Max Subs per SG for Low, Moderate & High DS TavG growth, 1218/204 MHz

For comparison purposes, Figure 43 also shows the 2018 projections using the 43% CAGR uber growth. It had projected that max subs per SG would drop to 28 subs by 2030. This drove many people to think that FTTP would be required by then. The reality is that a 1218/204 MHz plant supporting 5G x 1G tiers can easily last into the next decade, maybe even further if the slower growth projections hold.

Figure 44 maps two max subs/SG scenarios onto the DS TavG trendline charts from Figure 28. When DS TavG = ~11 Mbps, the max subs/SG supported is ~300 subs. This is the lower horizontal red dashed line. When DS TavG = ~23 Mbps, the max subs/SG supported is ~150 subs. This is the upper horizontal red dashed line. For a given trendline, it can support that SG size until it crosses that dashed line. Looking at the 150 max subs/SG as an example, the highest growth projections show it lasting until the end of this decade while multiple slower projections don't even reach the mark in 15 years. The slow linear growth takes multiple decades. If an operator is at 150 subs/SG today, they will stay there for a very long time (provided the 5Gx1G tier is sufficient!). This might be true for markets where there is limited demand or need for multi-gig upload speeds.




Figure 44 – 1218/204 MHz Max subs / SG vs. various DS Tavg projection trendlines




Figure 45 – 1218/204 MHz System – Spectrum Utilization

Figure 45 shows the spectrum utilization for the 1218/204 MHz plant scenario. The 204 MHz high split is introduced in 2024. The US spectrum is on the bottom of the chart. In that same year, the operator has finished the IPTV migration and the legacy video spectrum (in red) is reduced to zero. The overall

spectrum requirements are reduced with these savings, even though the IPTV spectrum (in purple) had increased slightly. Note – there are many inputs that impact IPTV BW consumption, so savings may vary from operator to operator depending on their particular situation.

In 2025, the 5G DS tier is introduced which fills up most of the available 1218 MHz of spectrum. Note that the SG size is still at 400 subs. After 2025, the SG size is reduced as needed to keep within the allotted 1218 MHz. Note that the IPTV spectrum is shrinking over time. This is due to both fewer subs and the video bit rate reductions, despite higher UHD %.

Figure 46 shows the models outputs for DOCSIS usage in Mbps broken out into three components:

- QoE Delta (i.e. $K \cdot T_{max}$), $N_{sub} \cdot T_{avg}$ and IP Video

It shows these three components for both D3.1 modems and 3.0 modems (which are removed by 2025). The redline on the top is the total available system capacity. Note that the 1218/204 MHz plant is providing a total of >9 Gbps usable capacity. That is higher than a 10G PON at 8.6 Gbps.

Figure 46 gives a good visual of the proportion of capacity needed for each component. The sum of the three components needs to be below the total available capacity. Note that this scenario is reasonably balanced between the consumption and burst components. Cutting the SG size in half would only provide an extra 1.5 Gbps of capacity, not enough to drastically change the 5G DS tier.




Figure 46 – 1218/204 MHz System – DOCSIS DS Usage: Tmax, Tavg, IP Video




Figure 47 – 1218/204 MHz System – DOCSIS US Usage: Tmax, Tavg, IP Video

The US DOCSIS capacity usage is broken out in Figure 47. Supporting a 400 Mbps US tier in 85 MHz during 2022-23 is a bit tight, but there is plenty of capacity after the 204 MHz high split upgrade. As can be seen, the Tmax component dominates over time. Figure 47 assumes the SG size is being set based on a moderate DS Tavg growth projection of 16% CAGR while the US Tavg is on a high growth 23% CAGR. See that the US has sufficient capacity until 2032. At this point, the operator might consider switching from 1x1 to 1x2 RMD configuration. Also, the DS:US ratio has dropped all the way to 6:1.

Perhaps the key point of this 1218/204 MHz case study is that a node with 150+ subs can be upgraded to 1218/204 MHz and support a service tier of 5 Gbps x 1 Gbps for the next decade and beyond. There is no pressing near term need to push the HFC to very small (and inefficient!) SG sizes, that could be, for example, achieved in N+0 systems.

5.3. Matching 10G PON on a 1794 / 396 MHz ESD Plant

The next case study focused on a DOCSIS ESD 1794/396 MHz plant offering a 7.5G DS tier with a 2.5G US tier. The primary goal here is to match the 10G PON DS service tier level. As before, the model was run for high (21% CAGR), moderate (16% CAGR) and low (linear) DS Tavg growth rates. The DS Tavg and US Tavg both hit the same numbers in 2032 as the previous case study.

Figure 48 shows the max number of subs per SG for low, moderate, and high DS Tavg growth projections. All projections can support 400+ subs through 2024. In 2025, slightly earlier than the previous case study, the high growth scenario gradually reduces the max subs/SG each subsequent year. It

is still supporting 250 subs/SG until 2028 and ~130 subs/SG in ten years. The moderate growth rate follows a similar curve only delayed about two years. The slow linear growth supports 400 subs/SG through 2026 and 300 subs/SG in 2032, a decade from now. For comparison purposes, Figure 48 also shows the 2018 projections using the 43% CAGR uber growth.




Figure 48 – Max Subs per SG for Low, Moderate & High DS TavG growth, 1794/396 MHz

Figure 49 maps two max subs/SG scenarios unto the DS TavG trendline charts from Figure 28. When DS TavG = ~8 Mbps, the max subs/SG supported is ~300 subs. This is the lower horizontal red dashed line. When DS TavG = ~19 Mbps, the max subs/SG supported is ~150 subs. This is the upper horizontal red dashed line. For a given trendline, it can support that SG size until it crosses that dashed line. Looking at the 150 max subs/SG as an example, the highest growth projections show it lasting until 2027 while the slow linear growth will take multiple decades. For a 150 subs/SG today that is upgrade to 1794/396 MHz ESD plant, it can stay there until 2029 with high growth but for a very long time with slow linear growth. Thus, node splits on the ESD plant become very sensitive to which DS TavG growth trendline it tracks. A proactive operator might want to deploy a 2x2 RMD in that location even though 1x1 may be adequate for many years.




Figure 49 – 1794/396 MHz Max subs / SG vs. various DS Tavg projection trendlines




Figure 50 – 1794/396 MHz System – Spectrum Utilization

Figure 50 shows the spectrum utilization for the 1794/396 MHz plant scenario. The 204 MHz high split is introduced in 2024 at the same time the DS expands to 1794 MHz. Then the US split is reconfigured for 396 MHz in 2026. The legacy video and IPTV spectrum follow the same trajectory as before. The first big spectrum jump in 2024 is the 5G DS introduction. The 2nd jump in 2026 is the 7.5G DS tier introduction.

Note that the SG size is at 400+ subs until 2024, then the SG size is reduced as needed to keep within the allotted 1794 MHz. As before, the IPTV spectrum is shrinking over time due to both fewer subs and the video bit rate reductions, despite higher UHD %.

Figure 51 shows the 1794/396 MHz model outputs for DOCSIS usage in Mbps broken out into three components:

- QoE Delta (i.e. $K \cdot T_{max}$), $N_{sub} \cdot T_{avg}$ and IP Video

The redline on the top is the total available system capacity. Note that it is initially 1794/204 MHz plant with >13 Gbps usable capacity, but then settles back to ~11 Gbps for the 1794/396 MHz plant configuration. This is significantly higher than a 10G PON at 8.6 Gbps.

Figure 51 gives a good visual of the proportion of capacity needed for each component. The sum of the three components needs to be below the total available capacity. Note the T_{max} burst component is starting to dominate the usage. Cutting the SG size in half provides minimal benefit, ~ 1 Gbps of additional capacity. If an operator would prefer additional plant life, they might consider backing down the DS T_{max} to 5 or 6 Gbps.




Figure 51 – 1794/396 MHz System – DOCSIS DS Usage: Tmax, Tavg, IP Video




Figure 52 – 1794/396 MHz System – DOCSIS US Usage: Tmax, Tavg, IP Video

The US DOCSIS capacity usage is broken out in Figure 52. As can be seen, the Tmax component dominates over time. Figure 52 assumes the SG size is being set based on a moderate DS Tavg growth projection of 16% CAGR while the US Tavg is on a high growth 23% CAGR. See that the US has sufficient capacity through 2032. It remains a fairly balanced 1x1 RMD configuration despite the DS:US ratio dropping to 6:1.

Perhaps the key point of this 1794/396 MHz case study is that a node with 150+ subs can be upgraded to 1794/396 MHz, but the timing of additional node splits on the ESD plant becomes very sensitive to which DS Tavg growth trendline it tracks. A proactive operator might want to deploy a 2x2 RMD in that location even though 1x1 may be adequate for many years. As with the previous case study, there still seems to be no pressing near term need to push the HFC to very small (but inefficient!) SG sizes such as those found in N+0 systems.

5.4. Other 1794 MHz ESD and FDX Options

The D4.0 specification [DOCSIS_4.0_PHY] provides operators with lots of different choices and options. From the highest level, D4.0 offers both Extended Spectrum DOCSIS (ESD) and Full Duplex DOCSIS (FDX). ESD DS spectrum goes up to 1794 MHz, a.k.a. 1.8 GHz. The specification offers five different upstream split options going from 204 MHz up to 684 MHz. The DS spectrum starts after a guard band above the US. The size of the guard band increases proportionately with the US split.

FDX shares US + DS spectrum in the FDX Band. This is variable width. The FDX Band starts at 108 MHz and can go as high as 684 MHz. While the top of the DS spectrum is nominally 1218 MHz, many operators considering FDX do not want to replace their taps, so they may be limited to 1002 MHz of DS spectrum, at least initially.

Table 1 shows the maximum DS + US service tiers supported by the various ESD and FDX options. The table also provides the max number of subs per SG for two different DS Tav_g – 7 and 15 Mbps/sub. The 7 Mbps DS Tav_g might occur 3-8 years from now depending on high, medium or low growth rates; while the 15 Mbps DS Tav_g happens in a much longer window that is 8-20+ years from now.

The max tier calculation uses mostly best-case assumptions including:

- 100% IPTV, no Legacy Video spectrum
- 100% DOCSIS 3.1/4.0 modems
- DS Tav_g = 7 or 15 Mbps/sub
- 4096-QAM for dedicated DS below 1218 MHz and 1024-QAM US for dedicated US
- A range from 1024-QAM to 4096-QAM for ESD DS > 1218 MHz and FDX DS in FDX Band
 - Previous network capacity modeling used an average of 2048-QAM

Most of the max tiers are shown in Gbps except for ‘2G’ = 2 x 940 Mbps and ‘5G’ = 5 x 940 Mbps.

As a reference point, the final row of the table provides the max tier and SG sizes for a 10G PON network. This includes the symmetric versions: 10G EPON, 10G XGS-PON and NG-PON2. It does not include XG-PON which only has 2.5G US optics and is limited to a 2G US tier. Note the 10G PON SG size limits as DS Tav_g increases.

Table 3 – DOCSIS 4.0 Maximum Service Tiers and Max Subs per SG

DOCSIS 4.0 Max Service Tiers		Max DS Tier	Max US Tier	Max Subs per DOCSIS/PON SG			
				DS Tav _g = 7Mbps		DS Tav _g = 15Mbps	
				1K-QAM	4K-QAM	1K-QAM	4K-QAM
ESD 1.8GHz	204 / 258	12G	1.25G	194	371	90	173
	300 / 372	11G	‘2G’	182	360	85	168
	396 / 492	10G	2.5G	162	340	76	159
	492 / 606	9G	3G	151	328	70	153
	684 / 834	7.5G	‘5G’	54	232	25	108
FDX	1002 / 108-684	6G	‘5G’	220	353	114	176
	1218 / 108-684	7.5G	‘5G’	305	400+	154	216
10G PON		7.5G	7.5G	128		64	

With the ESD options, an operator can choose how symmetric or asymmetric to make their system. For this analysis, the service tiers were maximized at the expense of SG size. This is yet another trade-off the operator can make. Note – by reducing max DS tiers, the ESD SG size can be significantly increased as our case study showed.

With the asymmetric 204 MHz split, the operator could offer a 12 Gbps DS tier, far eclipsing the 7.5G DS tier with 10G PON. This system is roughly 10:1 DS:US ratio as a 1.25 Gbps US tier pairs with it. Even with DS Tavg = 15 Mbps, the SG size is still a reasonable 173 subs at 4096-QAM.

The most symmetric ESD option is the 684 MHz split. This can achieve 7.5G DS with ‘5G’ US, but SG sizes start to get squeezed.

The middle of the road 1794/396 MHz ESD case study previously considered can be optimized a bit further. Going with smaller SG, improving HFC plant for even better QAM modulations and eliminating the managed IPTV service can buy enough capacity to push the DS tier to a true 10 Gbps. This is paired with a 2.5 Gbps US tier. If an operator can’t make these improvements, they could drop to a 300 MHz split to still get the true 10 Gbps, but now the US tier is reduced to ‘2G’ (i.e., 2 x 940 Mbps). In the FDX camp, a 1002 MHz system can support up to 6G DS with a ‘5G’ US with reasonably large SG sizes. If the operator pushes this plant up to a true 1218 MHz system, then the DS tier goes up to match 10G PON at 7.5 Gbps DS tier.

6. Conclusion

The CommScope (formerly ARRIS) team has led industry traffic engineering research for over a decade. [CLO_2014] introduced broadband QoE using a simple formula with basic network capacity components. This evolved and [ULM_2019] gave an updated insight into calculating the SG capacity requirements:

Modified “COMMScope/CLOONAN’S CAPACITY EQUATION” Traffic Eng Formula:

$$C \geq (N_{sub} * T_{avg}) + (K-1) * T_{max_max} + T_{max_max} \quad (2)$$

The subtle change is that there are now three main components to the traffic engineering formula:

1. Peak Busy Period Average Consumption (i.e., $N_{sub} * T_{avg}$)
2. Peak Busy Period Ripple for managing QoE (i.e. $(K-1) * T_{max_max}$)
3. Headroom for maximum Service Tier Burst (i.e., $1 * T_{max_max}$)

While burst and ripple components manage a subscriber’s QoE, the consumption component is key to SG sizing. The T_{avg} growth rate has seen much research and is the focus of this paper. ARRIS/CommScope has the most extensive broadband capacity monitoring history in the industry, collecting continuously since 2010 from the same MSOs. The 2022 data is in and DS T_{avg} growth continues to slow.

The paper looks at the consumption growth so operators can drive their network investment strategies for coming decades. Has T_{avg} growth slowed to a lower CAGR or is it no longer exponential? Several possible growth trendlines were investigated including exponential, linear, Adoption S-curve and others. Our research measures how accurately each trendline matches last decade’s data. These BW growth trajectories are mapped out for 5/10/15 years. The resultant spaghetti plots in Figure 53 show a cone of uncertainty that grows over time, roughly doubling every 5 yrs. To understand the impact of these slowing growth rates, consider the following comparison to projections from just four years ago:

- 2018 DS Growth (43% CAGR) projection => DS T_{avg} = 100 Mbps/sub by 2030
- 2022 DS High Growth (21% CAGR) projection => DS T_{avg} = 100 Mbps/sub by 2040
- 2022 DS Low Growth (Linear) projection => DS T_{avg} = 100 Mbps/sub in 200+ years

The implication is that *the need for FTTP to all subscribers may be pushed back multiple decades*. From a network capacity planning perspective, our conclusions on multiple growth trendline options are:

- the 5-year window provides a reasonably high confidence for near-term planning
- 10-yr window provides high, moderate, and slow growth ranges for longer term planning
- the 15-yr window shows too much variance and is more of an academic exercise.

The CommScope network capacity model studies raise several key points. The 1218 MHz case study shows a 500HP node with 2x2 RMD and 150+ subs/SG upgraded to 1218/204 MHz supports 5G x 1G service tier for the next decade and beyond. This may work for markets where there is limited demand or need for multi-gig upload speeds. The 1794/396 MHz case study in Figure 54 shows that a node with 150+ subs can offer 7.5G x 2.5G tiers; but the timing of additional node splits on the ESD plant is sensitive to which DS T_{avg} growth trendline it tracks. In either case, there is no pressing need to push the HFC to very small (but inefficient!) N+0 SG sizes.

Operators need to consider the low/medium/high growth scenarios when formulating their network migration strategy. A companion paper to this, [ZORAN_2022], looks at the economic impacts of these

different growth scenarios for different cable and FTTP architectures. The goal is to minimize up front investments while maintaining flexibility to increase network capacity and manage uncertainty risks.




Figure 53 – DS Tavg Growth Projections for 2022 to 2037




Figure 54 – Max Subs per SG for Low, Moderate & High DS Tavg growth, 1794/396 MHz

Abbreviations

AC	alternating current
BW	Bandwidth
CAGR	compounded annual growth rate
CCAP	Converged Cable Access Platform
CDF	cumulative distribution function
CM	cable modem
CMTS	Cable Modem Termination System
CP	cyclic prefix
CPE	consumer premises equipment
D3.1	Data Over Cable Service Interface Specification 3.1
D4.0	Data Over Cable Service Interface Specification 4.0
DAA	distributed access architecture
DC	direct current
DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
EIA	electronic industries association
EOL	end of line
EPON	Ethernet Passive Optical Network (aka GE-PON)
ESD	extended spectrum DOCSIS
FDX	full duplex (i.e. DOCSIS)
FFT	fast fourier transformation
FTTH	fiber to the home
FTTP	fiber to the premise
Gbps	gigabit per second
GHz	gigahertz
HFC	hybrid fiber-coax
HSD	high speed data
HP	homes passed
HW	hardware
IEEE	Institute of Electrical and Electronics Engineers
ITU	International Telecommunication Union
K	QoE constant
MAC	media access control
MB	multi-port bridger
Mbps	megabit per second
MHz	megahertz
MSO	multiple system operator
N+0	node+0 actives
NCTA	The Internet & Television Association
Nsub	number of subscribers
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
PDF	probability distribution functions
PHY	physical interface
PON	passive optical network

QAM	quadrature amplitude modulation
QoE	quality of experience
SC-QAM	single carrier QAM
SG	service group
SCTE	Society of Cable Telecommunications Engineers
SLA	service level agreement
Tavg	average bandwidth per subscriber
Tmax	maximum sustained traffic rate – DOCSIS Service Flow parameter
TX	transmit
UHD	ultra high definition
US	upstream
VOD	video on demand
VR/AR	virtual reality / augmented reality
YoY	year over year

Bibliography & References

[ARK_BLOG] <https://blog.arkieva.com/basics-on-s-curves/>

[CHR_1992] “Exploring the limits of Technology S-Cure. Part I: Component Technologies”, Clayton M. Christensen, Harvard University Graduate School of Business Administration, Production and Operations Management Society, Vol. 1, No. 4, Fall 1992

[CLO_2019] T. J. Cloonan et. al., “Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years,” SCTE Cable-Tec 2019, SCTE

[CLO_2017] T. J. Cloonan et. al., “The Big Network Changes Coming with 1+ Gbps Service Environments of the Future,” SCTE Cable-Tec 2017, SCTE

[CLO_2016] T. J. Cloonan et. al., “Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond,” NCTA Spring Technical Forum 2016, NCTA

[CLO_2014] T. J. Cloonan et. al., “Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements,” SCTE Cable-Tec 2014, SCTE

[CLO_2013] “Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS,” T. J. Cloonan et. al., NCTA Spring Technical Forum 2013, NCTA

[CORREL] “How to find Correlation Coefficient in Excel” <https://www.exceltip.com/statistical-formulas/how-to-find-correlation-coefficient-in-excel.html>

[DOCSIS_4.0_PHY] “DOCSIS 4.0 Physical Layer Specification”, CM-SP-PHYv4.0-D01-190628, CableLabs 2019

[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

[Exceltip R2] “How to find R-Squared in Excel?” <https://www.exceltip.com/excel-functions/how-to-find-r-squared-in-excel-use-rsq-function.html>

[HOWALD_2022] “Collision-Free Hyper-Speeds on the Bi-Directional FDX Highway”, Dr. Robert Howald, John Ulm, Saif Rahman, Dr. Zoran Maricevic; SCTE Cable-Tec 2022, SCTE

[Investopedia R2] “What is R-Squared?” <https://www.investopedia.com/terms/r/r-squared.asp>

[ULM_2021] “Managing the Coronavirus Bandwidth Surge – How to Cope with the Spikes and Long-term Growth”, John Ulm, Dr. Tom Cloonan, SCTE Cable-Tec 2021, SCTE

[ULM_2019] “The Broadband Network Evolution continues – How do we get to Cable 10G?”, John Ulm, Dr. Tom Cloonan, SCTE Cable-Tec 2019, SCTE

[ULM_2017] “Traffic Engineering in a Fiber Deep Gigabit World”, John Ulm, Dr. Tom Cloonan, SCTE Cable-Tec 2019, SCTE

[ULM_2016] “Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies”, John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo

[ULM_2014] “Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning”, John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo

[Wiki R2] “Coefficient of determination” https://en.wikipedia.org/wiki/Coefficient_of_determination

[ZORAN_2022] “Network Migration to 1.8 GHz – Operational “Spectral Analysis” measured in nano-Hertz, a 30-year perspective”, Zoran Maricevic, Craig Coogan, John Ulm; 2022 SCTE Cable-Tec Expo

Broadening the Reach of Broadband, Powered by Distributed Access Architecture

An Operational Practice prepared for SCTE by

Katherine Aiello

Director, Program Management
Comcast

1800 Arch Street, Philadelphia PA 19104
856.296.2954

Katherine_Aiello@cable.comcast.com

Robert Howald, Comcast

Frank Eichenlaub, Comcast

Jason Combs, Comcast

Table of Contents

Title	Page Number
Introduction	3
1.1. Comcast History and Community Initiatives	3
1.2. Technology Path Alignment	3
2. The Foundation of Distributed Access Architecture	6
2.1. Benefits of DAA	6
2.2. Virtualization	8
3. DAA Extensibility and the Access-Agnostic Last Mile	9
3.1. Integration of EPON FTTH into DAA Infrastructure	10
3.2. The Road to RBB Is Through vBNG/R-OLT	12
3.3. Distance Matters	13
4. Operational Benefits	15
4.1. Software Over Hardware	15
4.2. Network Telemetry – Breadth and Depth Across the Ecosystem	17
5. Operational Impacts	18
5.1. Sustainability	18
5.1.1. Fiber	19
5.1.2. Coaxial Cable	19
5.1.3. HFC	19
5.2. Workforce	20
6. Conclusion	20
Abbreviations	21
Bibliography & References	22

List of Figures

Title	Page Number
Figure 1 – Diagram of DAA Basic Infrastructure from Headend to Home	5
Figure 2 – DAA Basic Infrastructure Demonstrating Converged Access	5
Figure 3 – Virtualization of Edge Access Platforms	8
Figure 4 – Comcast DAA Architecture based on Remote PHY Node and vCMTS	9
Figure 5 – DAA Enables Access-Agnostic IP Network Architecture Convergence	10
Figure 6 – Cable and Telco PON Network Models	11
Figure 7 – vBNG-Based PON: Service Components and Traffic Flow	12
Figure 8 – Abstract Layers of DAA System	13
Figure 9 – PON OLT, OTL and ODN Architecture	14
Figure 10 – Remote OLT (R-OLT) Hardware Components	15
Figure 11 – DAA-Based PON – Optical Connectivity	15
Figure 12 – Real-Time Network Stability	16
Figure 13 – Real-Time Customer Cable Modem Telemetry	18

Introduction

1.1. Comcast History and Community Initiatives

For nearly 60 years Comcast has held to a consistent truth: The customers and the communities that we support are at the focal point of our decision making. In recent years, Comcast has created a multitude of programs to best support the diverse communities we serve. In 2011 Comcast introduced Internet Essentials. This program has delivered high-speed Internet to over 10 million Americans, bringing the technology required in today's world to be successful in education and work, and which now has also become social gathering space and center of entertainment. In 2020 the Lift Zone program was created. Additional sites are created every year, with over 1000 sites now available. This past year Comcast started to participate in the Affordable Connectivity Program, a government run program which has been able to help over 11.5 million households obtain reliable, quality Internet services through its partnerships with Internet providers¹. These programs show that Comcast's commitment is more than simply providing broadband connectivity. The company prioritizes being a strong community partner as well. We have commitment to bridge the digital divide by partnering for these programs and creating additional grants which have supplied more than \$16 million in grants and \$75 million in-kind support to thousands of small businesses².

Comcast's commitment also spans into being a responsible steward of the environment, with various initiatives to hold ourselves accountable to specific environment-friendly target goals. A first of its kind recycling program for coaxial cables was established this year. With this program, approximately 70% of our coaxial cable waste will be recycled every year. We stand by our commitment to be carbon neutral by 2035 and have initiatives in place to drive to that outcome.

In this paper, we discuss how the landscape of rural broadband (RBB) is evolving quickly, and strategic network investments being made to serve these areas and help to close the digital divide in the years ahead. Today, over tens of millions of Americans do not have reliable high-speed Internet³, and many are in these rural areas. With a foundation of virtualization and distributed access architecture (DAA) serving our coaxial footprint, Comcast is building on this to bring to life the DAA vision of a converged access infrastructure, and its introduction is timely and effective for RBB initiatives.

1.2. Technology Path Alignment

In 2020, the world was faced with a challenging situation. Having pioneered the deployment of a DAA beginning in 2017 combined with business-as-usual node segmentation that pushes fiber deeper into our network, Comcast was in an enviable position to handle the instantaneous pandemic-related capacity demands that 2020 presented^{4,5}. The pandemic shifted daily norms of work, education, and social interaction.

Comcast's DAA is based on the remote PHY specification⁶, as shown in Figure 1. The introduction of DAA five years ago began the preparation for new processes, practices, tools and technology, starting with our most familiar and powerful workhorse – DOCSIS[®]. DAA was a fast follower to our prioritization of DOCSIS 3.1, which was deployed as soon as it was available. The launch of DOCSIS 3.1 represented yet another instance of recognizing the importance of building new capacity well ahead of demand. In 2020, this turned out to be very prescient, enabling Comcast and other operators to mostly absorb the COVID-induced traffic spike with loss of capacity margin as the primary impact – not customer issues as seen in **Figure 0**⁴. These investments not only allowed Comcast to meet the extreme capacity demands of the pandemic, but it validated the years of decisions based upon untapped potential

of the HFC network and continuing to leverage that network and the coaxial connections to millions of homes.




Figure 0 - Highly Utilized CMTS SG Trending⁴

As we continue to deliver on continuous capacity growth and demand for increased speeds, RBB highlights a renewed focus on the digital divide in rural America. We have evaluated the approaches to delivering on RBB requirements, which typically come in the form of request for proposals (RFPs). We have considered the underlay of the environment, needs and expectations of customers, and the growth projected for these areas in years ahead.

There is new technology being developed that is well-suited to deliver on RBB requirements in these RFPs. In this discussion, we highlight how Comcast's DAA and vCMTS foundation will enable us to drive the latest technology and capability into RBB by introducing a new remote OLT (R-OLT) that is particularly effective for RBB, as shown in Figure 2. The R-OLT and accompanying virtualized platform SW will enable efficient and effective delivery of fiber-to-the-home (FTTH) broadband services to the RBB customer base. Historically, to deploy passive optical network (PON) technology, companies used centralized optical line terminals (OLT). These are classic "big iron" chassis consuming significant power, space, and cooling capacity in headend facilities. Furthermore, with their proprietary hardware and software implementations and the large scale of users subscribed to a single chassis, they are relatively costly to deploy, operate, and upgrade.

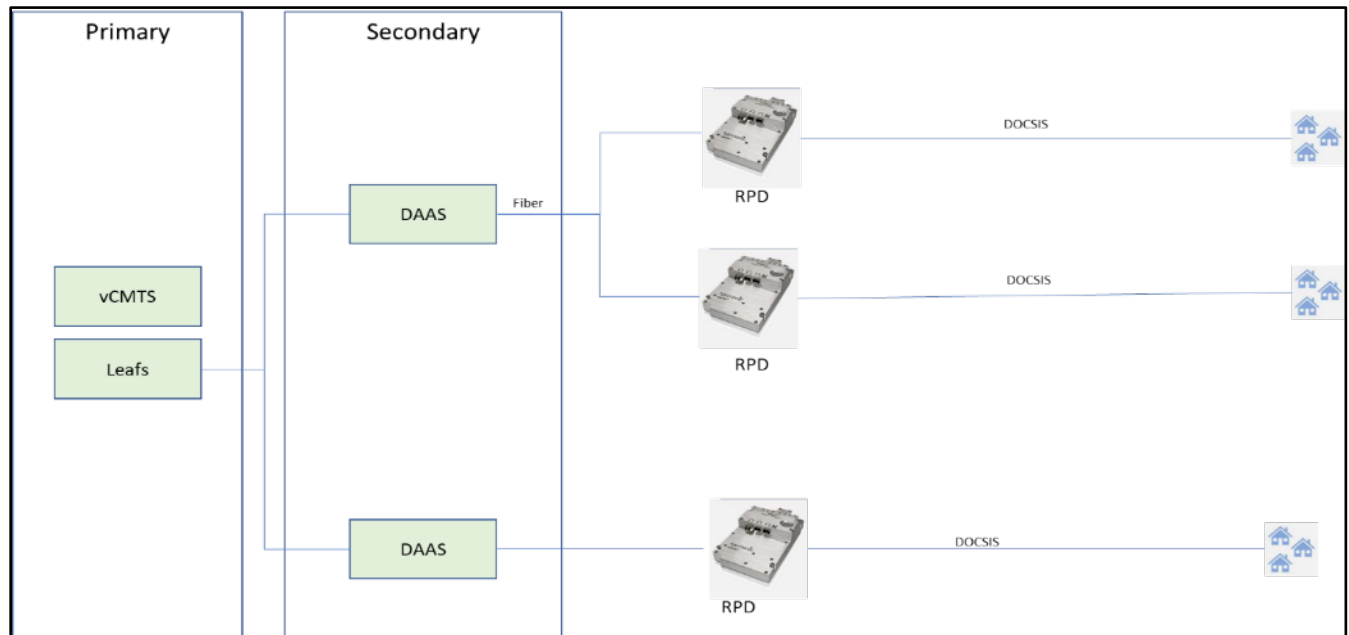


Figure 1 – Diagram of DAA Basic Infrastructure from Headend to Home

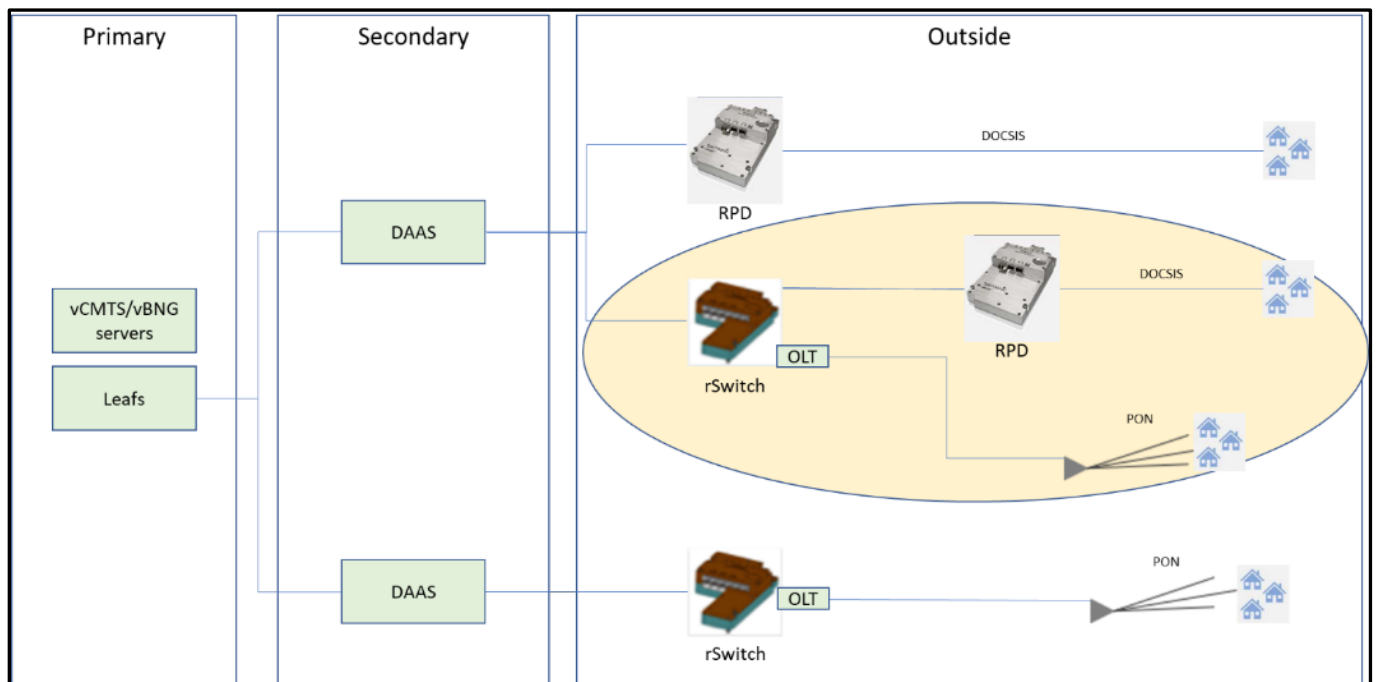


Figure 2 – DAA Basic Infrastructure Demonstrating Converged Access

The above description for PON chassis can also apply to DOCSIS, where “OLT” can be replaced with “I-CMTS” or “I-CCAP,” where “I” is for integrated. However, as noted, Comcast began moving to vCMTS

five years ago and continues to aggressively upgrade to vCMTS nationally. Nonetheless, most operators still have an integrated CMTS dominant footprint.

Working within the limitations of these integrated chassis platforms, major parts of budget strategies over the years have been focused on periodic SW upgrades and timing of procurements for growth and new features. In addition, while large scale chassis solutions do have efficiency advantages for supporting high density areas, they are not ideal for low density areas.

Other solutions have served and/or been proposed for these rural communities. Copper pair systems are simply inadequate – too slow – by today’s broadband standards. Wireless-based systems can be impacted by poor weather and other factors, and not be able to deliver the consistent, reliable performance customers expect and deserve. Leveraging DAA investments with vCMTS and digital nodes is a pathway to efficiently deploy a R-OLT-based FTTH solution into RBB areas. This architecture provides the flexibility to implement PON service very granularly and on demand. It will enable reliable services to penetrate anywhere deep into the existing footprint, while also expanding the radius of reliable Comcast broadband service to rural customers.

2. The Foundation of Distributed Access Architecture

Because of persistent year-on-year compounded annual growth rate (CAGR) of downstream and upstream traffic, operators found it increasingly difficult to keep up with upgrades to the network and spectrum using business as usual (BAU) processes such as node splits and incremental adjustments to spectrum allocation. New processes and tools were needed to replace BAU approaches before the scale made it too challenging or costly to execute upgrades.

Among the most powerful tools developed to address this were the distributed access architecture specifications developed through CableLabs. These specifications were a collaboration among operators, vendors, and CableLabs subject matter experts (SMEs) to address many of the anticipated challenges for operators as data traffic and speed demands continued to grow year over year. Among the first to be published was the “Remote PHY” (R-PHY) specification in 2015⁶. Technology development had begun well before the completion of the specification, and Comcast embraced the transition to DAA using R-PHY as part of its network evolution plan in 2015.

Note that the CableLabs work continued to develop another version of DAA based on the flexible MAC architecture, or FMA. As part of its strategic planning for DAA migration, Comcast deeply evaluated both options to determine which would better position the Comcast network for the long term. In the final analysis, the R-PHY approach was selected as it better aligned to key network objectives such as centralization of high SW complexity, lightweight distributed compute for low-touch, high-reliability outside plant locations, and simplifying interoperability of ecosystem components.

Comcast began deploying R-PHY based DAA in 2017 and has over 30,000 remote PHY device (RPD) equipped nodes from multiple vendors in production today, and that number is rapidly increasing as a major network upgrade plan to digital nodes and vCMTS takes place. There is no large-scale deployment of the FMA architecture to date. The Comcast DAA network now exceeds the availability performance of our mature, standard HFC deployments, as was anticipated with the migration to DAA.

2.1. Benefits of DAA

Distributed access architecture delivers multiple powerful advantages. It represents one of the most powerful technology upgrades to cable, nearly as powerful as the introduction of fiber itself to all-coaxial RF systems decades ago.

Among the benefits that can be attributed to the move to digital optics, and specifically Ethernet-based digital optics, in place of cable-specific analog optical (AM) technology are:

1) Optical wavelength efficiency

With increasing node splits and deeper fiber migration, significant numbers of new nodes are installed every year. The use of wavelength division multiplexing (WDM), where multiple links to RPDs can share a fiber, maximizes the use of the existing fiber infrastructure. The use of digital optics means up to 80 dense wavelength division multiplexing (DWDM) wavelengths/nodes can be aggregated on a single fiber (or more). Standard HFC AM optics are restricted typically to 16 wavelengths to manage the nonlinear effects on modulation error ratio (MER) for typical HFC optical links lengths. This means a more cost-effective deployment and less fiber needed to support the nodes.

2) Reach of digital optics vs AM optics

As noted above, longer optical links from headend transmitter to node historically meant that fewer wavelengths can be used on that fiber to meet a given end-of-line MER. Digital optics eliminate this dependency in practice for common HFC optical links. In fact, a DAA-based CMTS core, outfitted with digital optical outputs, could be moved closer to the core if headend consolidation is an objective.

3) SNR and MER performance improvement of eliminating the DS and US AM HFC optics

DOCSIS 3.1 and DOCSIS 4.0 enable higher order modulation profiles, increasing the bandwidth efficiency by up to 50% over DOCSIS 3.0 in the downstream and up to 100% in the upstream. In standard HFC, the end-of-line (EOL) MER is dominated by the performance of the AM optics. The CMTS RF ports provide an extremely high MER of 48 dB (set by the DOCSIS Downstream RF Interface Specification, or DRFI) for DOCSIS 3.1, which gets degraded by the AM optical link to the node in the field to below 40 dB, depending on optical link length and other variables and that is *before* the RF amplifier cascade acts to degrade it further.

With DAA, the digital link to the node eliminates the AM optics, and the MER degradation caused by it. Instead, the DRFI requirement is met at the node RPD port, gaining back the lost fidelity of the AM optics nearly completely. This maximizes the capacity possible for DOCSIS 3.1 DS output and US input signals.

4) Space, power, cooling efficiencies in the hubs and headends

Moving some of the CMTS functionality to the node leaves less behind in the hub or headend to power, cool, and consume space. The density of RF connectors and isolation requirements on a typical CMTS tend to set the density of these chassis. With the RF ports of the CMTS distributed into the plant, the density of the core can now be redefined since the density of optical connectors is much higher. Also, integration of digital infrastructure for video and data services allows for the elimination of large RF combining networks.

5) Alignment with virtualization of the CMTS, and more broadly towards convergence into a common virtualized platform serving multiple last-mile coaxial and fiber access technologies

This is discussed in the next section.

2.2. Virtualization

Cable operators' inside plant equipment and networks historically have been based on a collection of purpose-built video, voice, and data platforms of integrated hardware and software. This has been a successful formula for over 30 years of service introduction and service growth including digital video, high-definition (HD) video, voice, data, and video on demand (VOD) and FTTH implementations of the same services.

Unfortunately, with these monolithic platforms, in particular DOCSIS CMTS platforms, it is difficult to keep pace with exponential traffic increases. Also, considering the construction aspect of splitting nodes, typical network augments – node splits – cannot accelerate to match the trajectory of traffic growth.

Network function virtualization (NFV) and software defined networking (SDN) are enablers of cost effective, efficient, exponential network and service change velocity. Historically, the continued growth in traffic and mounting node splits to support that growth meant the addition of CMTS RF ports and line cards.

Bringing this concept down to operational practice, in a DAA implementation, these RF ports are distributed into the field. However, supporting line cards of the CCAP core would still be necessary if the digital node is connected to an existing I-CCAP.

In a virtualized implementation, however, this purpose-built hardware core, designed to be tightly coupled to the output RF interfaces, is instead implemented in commercial off-the-shelf (COTS) server hardware. This is made possible simply through Moore's Law. The compute power and resources needed are available in standard processors today, allowing CMTS functions to be executed in such platforms. This approach is shown in Figure 3.




Figure 3 – Virtualization of Edge Access Platforms

The potential to create SW-based CMTS on standard server platforms has enormous implications for cost, space and power savings in facilities, and service velocity. As with DAA nodes, the advantages of a migration to virtualized edge platforms are so compelling that figuring out this transition from integrated

architectures to distributed architectures, and to virtualized *and* distributed is essential. At Comcast, every RPD node is connected to a vCMTS core. Traffic growth and new product speeds now revolve increasingly around compute power, which scales with Moore's law, and spectrum re-allocation, which no longer involves massive re-wiring inside plant because it is empowered by automated RPD configuration for channel line-up changes.

In summary, the Comcast DAA implementation is based on an Ethernet switch fabric feeding R-PHY-based DAA nodes. The previously purpose-built CMTS hardware core is instead virtualized in COTS platforms. This architecture is shown in Figure 4, envisioned here as migration of existing HFC architecture – I-CMTS and analog fiber node to vCMTS, digital infrastructure, and digital nodes (RPDs). Not pictured here for simplicity is the substantial RF and analog optical headend equipment in the secondary that exists between the I-CMTS and analog node.

A significant item to note is that with DAA the vCMTS location is consolidated back to the primary – a benefit of the digital optics and scalability, with the switching infrastructure deployed into the secondaries rather than the CMTS itself.

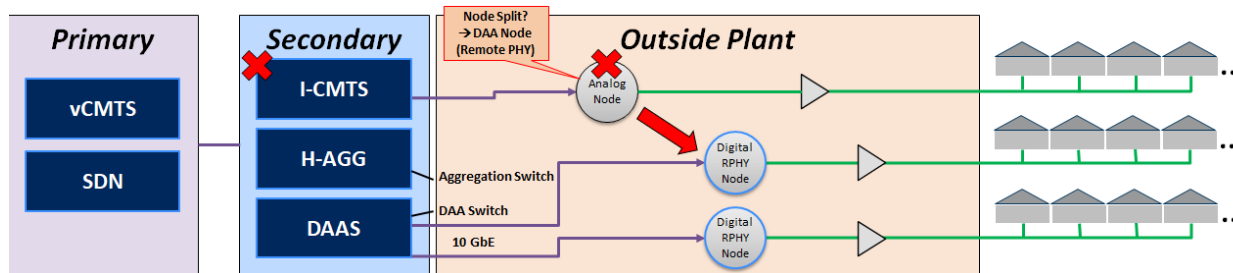


Figure 4 – Comcast DAA Architecture based on Remote PHY Node and vCMTS

3. DAA Extensibility and the Access-Agnostic Last Mile

DAA benefits extend beyond the considerable advantages it provides to cable systems. Most operators within their strategy also build and operate FTTH systems, consider the role of FTTH as HFC network upgrades take place, and think about how to efficiently transition to it where and when it makes sense to do so.

As the HFC architecture takes fiber deeper and builds out DAA based on vCMTS, the power of deep network Ethernet connectivity comes to the forefront. Ethernet is a Layer 2 foundation of global scale, and the workhorse of modern fiber and copper WAN and LAN communications. Virtually any last mile access technology will interface to an Ethernet-based system. Furthermore, by terminating these links in OSP in a fiber node, the node platform becomes a multi-purpose platform of plug-in modules supporting various last-mile options. For cable systems today, the most obvious next scenario and fast-follower of DOCSIS into DAA is EPON FTTH. Unfortunately, HFC architectures are not built to the physical standards of optical reach, homes passed per fiber, and wavelength plan of a “classic” EPON architecture. Because of these differences, a key component of the existing EPON solution is a “PON extender” module in a node housing, which solves both fiber utilization and distance constraints of classic PON system overlaying HFC. This node location now serves as the perfect place to instead utilize an R-OLT. As part of a node housing, the node platform itself now becomes a multi-access last mile, architected for dual use in this case as shown in Figure 5.

The basis for both last miles is similar – deliver 10 gigabit Ethernet connectivity to an access module inside of the node. Current nodes support RPD modules already. A node of the future for FTTH last miles now includes the addition of an EPON R-OLT. The PON portion of network (where the PON protocol lives) then extends from the R-OLT port to business and residential customers, based on the same PON optical budget standards that exist today. These PONs are standard-compliant for optical link budget, as the PON is being served from that deeper physical location closer to FTTH customers, just as with the PON extender.

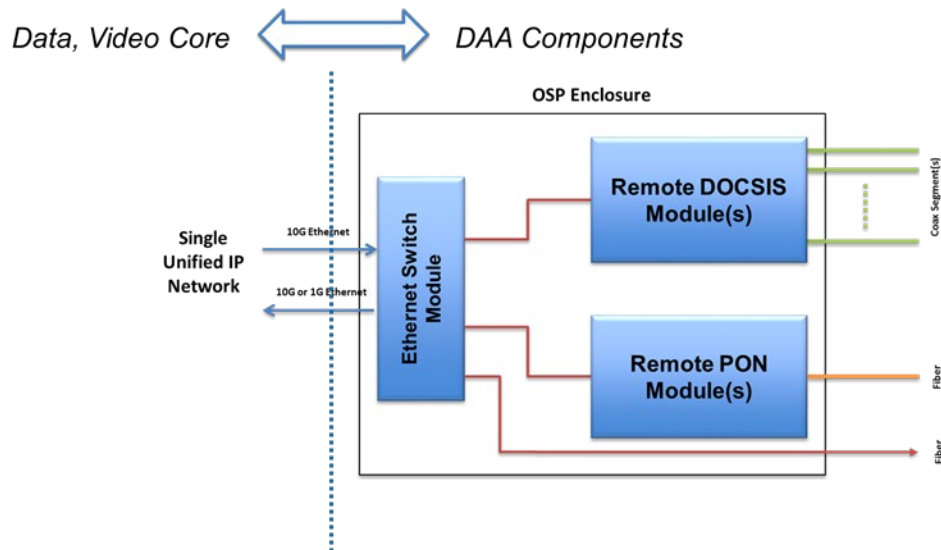


Figure 5 – DAA Enables Access-Agnostic IP Network Architecture Convergence

3.1. Integration of EPON FTTH into DAA Infrastructure

Last-mile access convergence does not end with an R-OLT HW module, however. Today's R-OLTs are simply smaller, hardened versions of centralized OLT platforms. Those were designed to support thousands of customers in large FTTH build-outs, such as the telcos have been doing in targeted ways for about 20 years. Cable operators generally deploy FTTH more tactically. Specific properties and developments built as FTTH are sprinkled about a ubiquitously deployed coaxial network. As mentioned, the R-OLT approach allows for ease of granularity of adding a single PON port at a time to manage growth most efficiently in the fashion it is deployed by most cable operators.

Aside from the properly sized HW platform, and as indicated previously, the virtualized infrastructure built and matured for DOCSIS DAA can be leveraged now for EPON, rather than operating DOCSIS and PON as two independent, parallel systems that happen to share the same optical infrastructure and switches. This introduces the virtual broadband network gateway (vBNG) platform SW component to the ecosystem

Consider Figure 6, which compares cable and telco implementation models for deployment of PON services, Open Systems Interconnection (OSI) alignment, and networking configuration. The telco PON architecture was developed to drop into existing DSL architectures, with a Layer 2 OLT taking over the role of the digital subscriber line access multiplexer (DSLAM).

Similarly, cable operators developed PON to drop in-line with existing DOCSIS deployment models, which has a more intricate integration of Layer 2 and Layer 3 functionality at the edge router (CMTS) for historical reasons. The more flexible EPON standard was used, and an SW abstraction layer called DOCSIS provisioning of EPON (DPoE) developed to accommodate the technology within a DOCSIS end-to-end system. DPoE assigns each ONU as a virtual cable modem (vCM) inside of the OLT, allowing operation and management of it via DOCSIS tools and back-office systems. This is what is commonly used in most EPON systems today in the cable operator community.

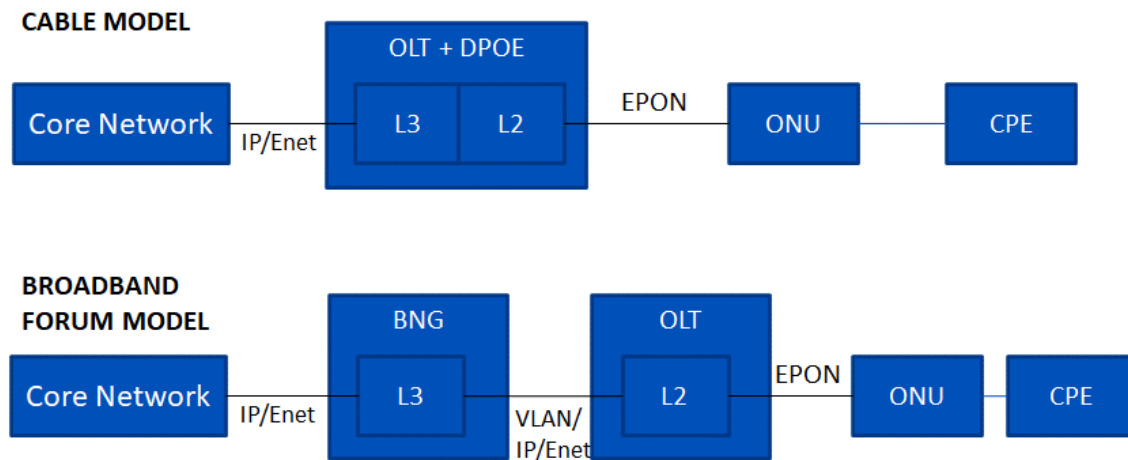


Figure 6 – Cable and Telco PON Network Models

For vBNG, however, this emulation of DOCSIS functionality in the R-OLT is eliminated. Rather than remotely deploying complex OLT SW with DPoE, the R-OLT is simplified, becoming more of a wavelength, scheduling/multiplexing, and protocol converter that integrates directly to subscriber management SW in the vCMTS platforms alongside DOCSIS. The vBNG data plane can be run in the network processing unit (NPU) of today's COTS server platforms, while the control plane can be run in the cloud.

Developing the architecture for incorporating FTTH into the production DAA system brought to light opportunities for synergy between DOCSIS and PON. The first obvious area of synergy is simply that the remote switch (rSwitch) and R-OLT are connected to already existing DAA infrastructure. This produces a very low barrier to entry to deployment time and saves on construction costs. It also allows the reuse of field processes that technicians have become accustomed to when installing RPDs. With this experience combined with experience from also having deployed EPON in scale in production, the learning curve is significantly reduced, and less new training needs to be developed.

Other major benefits to the sharing of the physical infrastructure include:

1. Transit fiber feeding DAA components in the field is efficiently used and shared between the access networks either through high-speed switch networks or via multiplexer combining of the Ethernet wavelengths to the RPD and rSwitch devices.
2. Routing infrastructure is combined, so the IP infrastructure costs are made to be as efficient as possible. This means that PON customers and DOCSIS customers will utilize the same IPv4 and IPv6 scopes and all the same security risk management practices.

3.2. The Road to RBB Is Through vBNG/R-OLT

One of the most powerful advantages of the DAA-based EPON implementation is the ability to take advantage of over five years of production deployments of DAA with vCMTS for DOCSIS customers, and leverage that knowledge for PON customers over that same DAA infrastructure with vBNG (instead of vCMTS) and R-OLT (instead of RPD). In the vCMTS, the core of the software is all built into a single Kubernetes “pod,” which is a tight knit group of containers. For PON, the software is built in three primary pods, identified in Figure 7: the vBNG itself, the access controller, and the service activation pod. Splitting these functions allows us to separate the complexities of service provisioning and device control away from the compute-intensive but relatively simple function of the vBNG to pass the customer’s traffic.

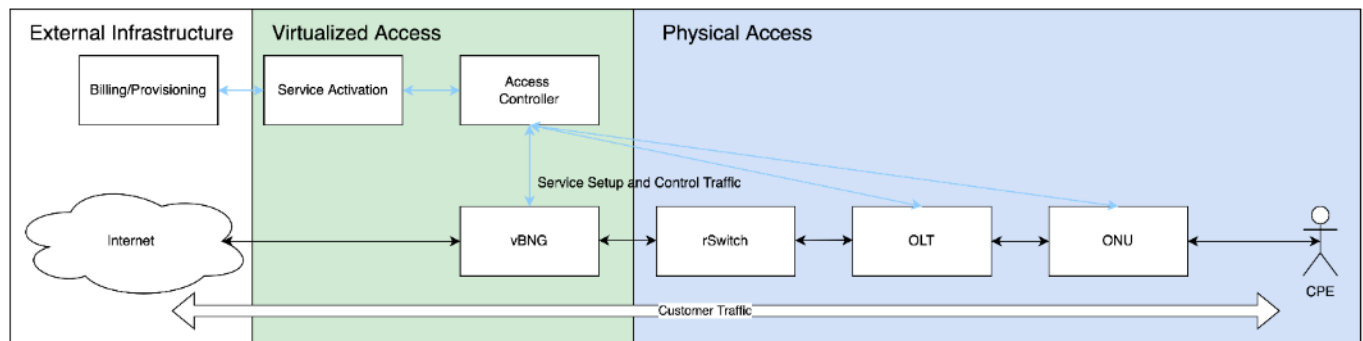


Figure 7 – vBNG-Based PON: Service Components and Traffic Flow

Let us detail further the functionality of these components:

Service Components – vBNG – This is the most familiar component in the system. The job of the vBNG is to guarantee the quality of service (QoS) parameters associated with the customer’s services, enforce security policies, provide basic IP services such as Address Resolution Protocol (ARP) and neighbor discovery, and to consistently pass the customer’s Ethernet frames to and from the Internet. This component has a high CPU resource requirement and must be highly available with a maximum downtime, counted in a handful of seconds. Simplicity and consistency are the primary operational considerations.

Access Controller – This component is primarily responsible for the configuration and monitoring of the R-OLT and ONU. It also provides service level configuration to the vBNG for its consumption. The PON interfaces and the service definition can become quite complex necessitating close adherence to both standards documents and actual implementation details within the R-OLT and ONU. This complexity does not impact the availability of customer services. However, for customer services, the access controller is only needed during initial provisioning in that the vBNG, R-OLT, and ONU will continue to provide the services they were configured with even if the access controller is down. It is still preferred that the controller be available if only for monitoring purposes, but the acceptable downtime can be measured in minutes. While the service activation and vBNG components will largely be common to all Ethernet based access technologies, the access controller must be replaced or redesigned significantly to support other technologies.

Service Activation – This component is responsible for all per subscriber service level configuration. Today for most cable operators that means that it will create a vCM to interface with a DOCSIS compliant back office. However, unlike a DOCSIS cable modem, information about the ONU devices can

be provided through telemetry rather than being polled via Simple Network Management Protocol (SNMP). This greatly simplifies the vCM itself, essentially becoming a DHCP emulator with a means of downloading configuration files and translating them into a common API to communicate with the access controller. The superpower of the service activation component, however, is that it is not tied to the DOCSIS back office. Having this component as part of the system allows the operator to add to, change, or replace this mechanism with the newest and/or most convenient APIs for back-office integration. This is an ability that is difficult to take advantage of in the short term but encourages movement in that direction just by existing. Much like the access controller, this component is only necessary at initial provisioning and the acceptable downtime can be measured in minutes.

The final critical and strategic attribute that should be highlighted about the DAA-powered PON architecture, founded upon the rSwitch, R-OLT, and vBNG components is not just the value it brings today, but also its value for future evolution. This system is constructed of six independent sets of services, as shown in Figure 8: management, provisioning, service fulfillment, routing, backplane, and access technology. These services can and are being improved upon without impacting the other services. The ideal future state is one where we have the best in breed services in each of these areas which can then be continually optimized within their own technology domains. By doing this, we will also be continuously improving network capability, services offered, and the customer experience while increasing network reliability for all commercial and residential customers. We can even consider third and fourth last mile access technologies to integrate in the future.

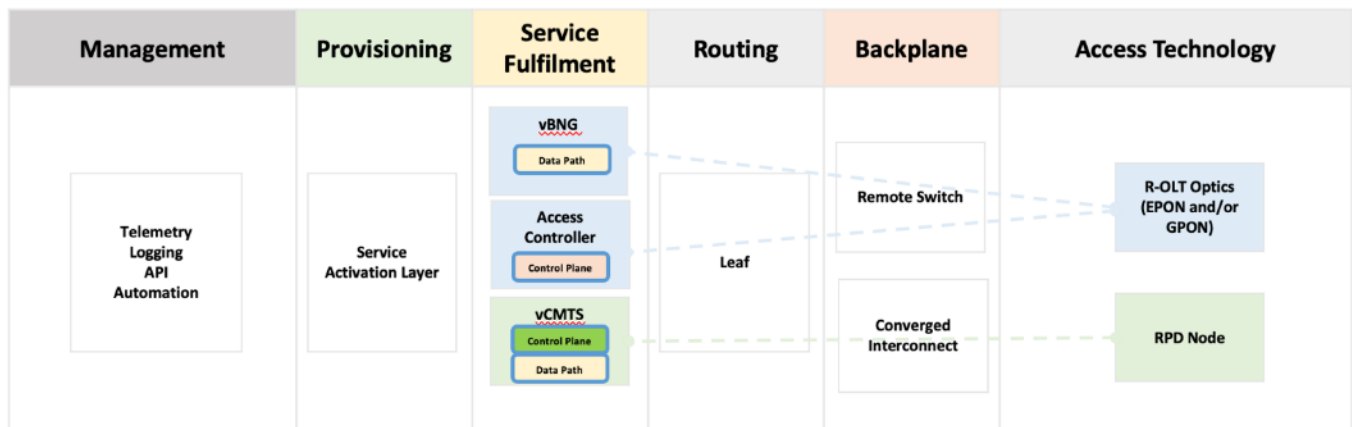


Figure 8 – Abstract Layers of DAA System

3.3. Distance Matters

Deployment of traditional HFC plant in very low-density has been limited historically. The cost of large-scale network construction has to be balanced against the number of customers that will be reached. Signal attenuation of the coaxial cable is a major factor in low density due to longer runs of cable, requiring long RF amplifier cascades and more power to run them.

Fiber optic cable has a much lower attenuation per mile compared to coax, so is better suited to cover very long distances efficiently. As such, the R-OLT and vBNG bring an opportunity to support rural deployments very effectively, as well as streamline the deployment of Comcast's EPON FTTH network by converging it with the existing DAA infrastructure.

In addition to the last-mile distance to homes in low-density areas, another challenge to serving these rural communities is the distance from Comcast's core network and primary facilities to these areas. A conventional chassis OLT requires deployment in a secure and environmentally controlled environment and access to the core network. From that centralized location the standard PON deployment can reach commercial and residential customers as far as 20 km of fiber away, assuming the 128-split ratio of the PON standard. Unfortunately, as noted previously, these parameters are not well-aligned to an implementation to cover a distant rural footprint. A proprietary PON extender solution developed by a technology partner enables the PON network to extend up to 80 km from the conventional OLT. In this implementation, the EPON network is transported from the OLT using conventional 10G DWDM transport to a node-type housing. Up to eight fiber link modules (FLMs) are contained in the node housing, making up to eight 10G EPON networks possible, supporting up to 1024 passings. This technology breaks up the optical transportation from OLT to ONU network into two parts, the OTL (optical transport link) and the ODN (optical distribution network) as shown in **Figure 9**.

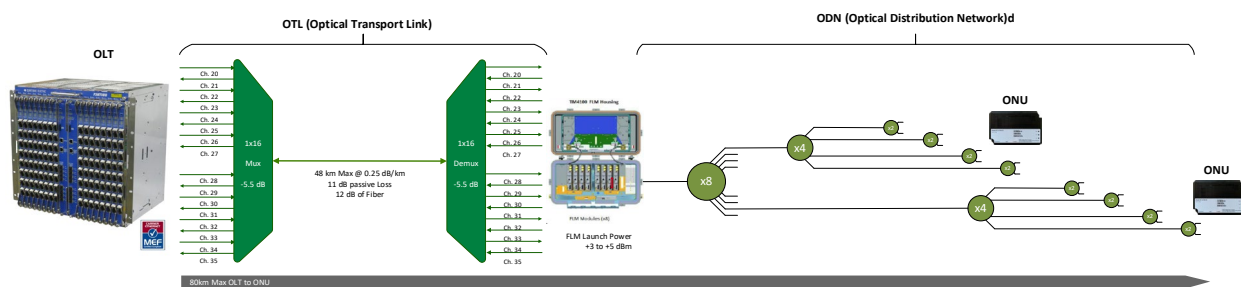


Figure 9 – PON OLT, OTL and ODN Architecture

As effective as PON extenders are to FTTH deployments, they represent a bandage to a fundamental architecture incompatibility that exists between telco-based PON standards and cable network principles. The R-OLT, by contrast, aligns PON technology to how HFC technology has evolved, in particular with DAA – use of Ethernet-based optical infrastructure feeding flexible, modular, multi-use platforms.

An example R-OLT is shown in Figure 10, in this case developed by Harmonic, one of Comcast's DAA business partners. The deployment is made up of three parts, the housing ("Ripple"), rSwitch "Jetty-1", and the single-port R-OLT itself, the "FIN-1".

This node housing for the R-OLT can also support an RPD, simultaneously. The Jetty-1 is a remote switch that connects to the distributed access architecture switch (DAAS) through a standard 10G DWDM link. The switch could support additional application, but for this paper we will focus on EPON. The Jetty-1 contains six SFP ports, two designated as uplink ports and four multipurpose ports. Up to two rSwitches can be mounted in this housing.

The FIN-1 enables simple plug-and-play 10G PON capability from rSwitch. Up to four FIN-1's per Jetty-1 are expected to typically support EPON deployments.

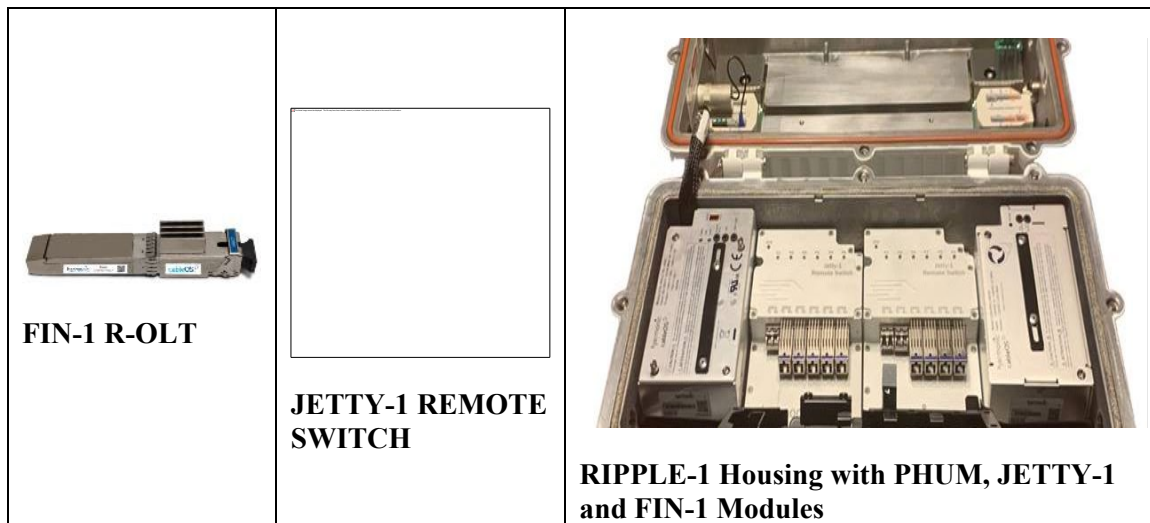


Figure 10 – Remote OLT (R-OLT) Hardware Components

Standard fixed or tunable 10G SFP+ are deployed to provide connectivity from the DAAS port and the rSwitch port. 12, 24 or 48 port DWDM 4.0 optical filters provide multiwavelength connectivity on a single fiber from the DAAS facility to the R-OLT housing. This is shown in **Figure 11**.

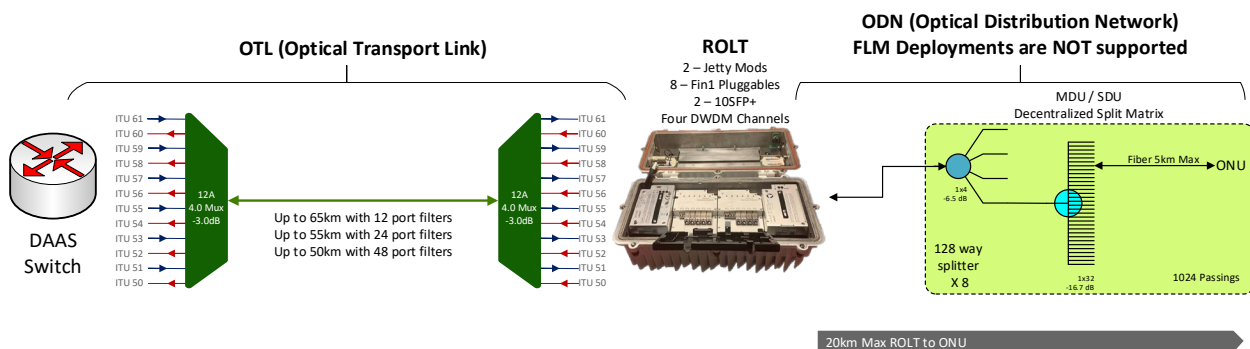


Figure 11 – DAA-Based PON – Optical Connectivity

The engineering and design of the ODN does not change with the deployment of the R-OLT when replacing an FLM-based EPON deployment. Of course, the OTL becomes standardized across the DAA deployment footprint including access to production tools and provisioning developed during the vCMTS rollout. The overall distance from the DAAS port to the ONU increases by 5 km reaching a maximum of 85 km from the DAAS to the ONU, enabling a solution that enables the deployment of EPON into the rural markets that is much simpler and faster – and where most of the end-to-end mileage is already built!

4. Operational Benefits

4.1. Software Over Hardware

As noted, Comcast has been investing in DAA for years. As part of the architecture development, decisions were made to drive to a more software-centric solution wherever possible. What does this mean

when talking about the access network? In many cases it means having the ability to upgrade software features and capabilities instead of requiring physical replacement of deployed elements in the network.

The large-scale deployment of vCMTS / RPDs demonstrates this, and the same vision, capability, and intelligence in the network is part of the vBNG / R-OLT design. Through DAA, we have never-before available visibility into the network, in real-time. We are leveraging the cloud and machine learning to create intelligent automation and configuration of components error-free and in-scale, all in SW. As an example, the process of increasing speeds and change spectrum allocations on the plant is a few clicks that can ultimately touch hundreds or thousands of production nodes.

Maximizing use of software in the architecture enables flexibility to, for example, push new speed tier offerings throughout the network in months rather than years. By designing a more intelligent outside plant, we are able to push configuration changes via software, obviating the need for a technician to go into the facility to perform re-wiring and/or into the field to make configuration changes. It opens the doors to automation that drives rapid scalability. With that scalability comes the need for improved network resiliency and visibility. Figure 12 is an example of a dashboard showing real-time information related to activation of OFDMA as part of a mid-split network upgrade. Currently, approximately 90% of updates and configuration changes introduced into the DAA footprint are delivered via automated tools, with a goal to increase this to 99% by the end of 2022.

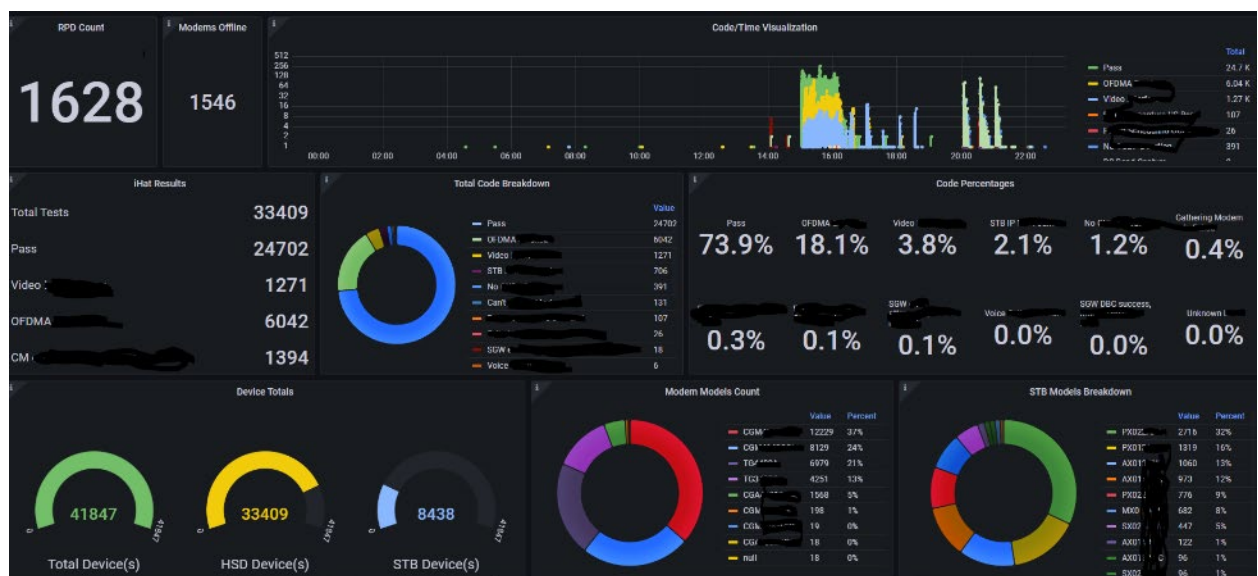


Figure 12 – Real-Time Network Stability

Increased intelligence means substantially more telemetry available to observe and manage the network. Data can be processed and filtered into dashboards focusing on performance, availability, and capacity management in real-time. These data sets will be available in R-OLT through the DAA investments, just as they are in DOCSIS RPD deployments. With metrics and real-time traffic data, new RBB deployments can be optimized very efficiently as the understanding of network characteristics and traffic dynamics associated with these new environments is learned. Available information includes operational data at each component of the network. This is an unprecedented amount of data compared to traditional HFC systems. The challenge then becomes how best to use this deep, real-time information. How should this data be managed? Who can use it effectively? Are there new tools that we must develop to convert this data into actionable information?

4.2. Network Telemetry – Breadth and Depth Across the Ecosystem

We break down the information and the benefits of this capability across equipment and serviceability elements below:

- **Platform**
This includes the base level hardware and software of the virtualized system that support the microservices, and the infrastructure to instantiate and maintain them. Invaluable information that speaks to the health, efficiency and workload within this view includes memory utilization, CPU load, Kubernetes state and host state.
- **Software deployment**
Deploying software in a non-customer impacting way – in-service upgradeability – is extremely important when there are hundreds of independent microservices that all rely on each other. This allows us to upgrade code and augment features with maximum flexibility and without traditional operational practices focused on customer care management. Current versioning, downtime, error notifications, and level of impact in case of a failure are all important metrics in this view.
- **Service Orchestration**
Bringing up the right software instances at the right time to the right customers is vital for success. How do we coordinate data from the field with information such as location, network devices, and physical connectivity of construction and network maintenance with the delivery of vBNG or vCMTS-based services? Automation and scripting are vital to doing this effectively at scale and without human error. This enables the efficient, widescale activation of services quickly as network builds such as RBB take place in parallel across the footprint.
- **Outside Plant**
With the distribution of smart components in the field, we have access to new and valuable data closer to the customer pertaining to their experience. This data can be analyzed via machine learning algorithms to trigger preemptive repair before a customer notices a problem. For DOCSIS, this includes RF data from the RPD and future smart amplifiers. For either DOCSIS or FTTH, real-time information on fiber impairments, fiber cuts, and power supply health and status provides important telemetry to preempt network outages or improve MTTR.
- **DAA Network**
As we push Ethernet out into the field via connectivity to the rSwitch and RPDs, we create increased connectivity complexity compared to traditional HFC, but also a much better opportunity to manage and monitor capacity, utilization, and scalability. Traditional network monitoring must be augmented to work effectively in a cable operator's physical network where an out of band connection cannot always be relied upon.
- **Customer Services**
This view of the network most directly impacts the customers. Who is impacted by a physical outage? Are the devices working in a degraded state? What are the network services that the customer is entitled to? This is the most traditional and vital view of the network for cable operators. Being able to observe real-time ONUs and CM behavior via streaming telemetry of the virtualized platform is an opportunity to address device issues before a customer notices service issues, which has a huge positive customer experience impact.

There are many other great ways that the system can be viewed, but this gives us a reasonable cross section to make some important observations.

The virtualized core and DAA infrastructure are very similar, no matter which last mile access technology is being used. Beyond the nuances of PON and DOCSIS, the system is run the same way using the same hardware with very similar, if not the same, software. Much of information gathered is similar, with some

obvious differences in the specific last mile RF data vs FTTH data. Most importantly, the information coming from this data is operationally transformative, creating significant efficiencies of scale while simultaneously adding deeper diagnostic value. They lend themselves to leveraging cloud services and infinite possibilities of dashboard views that can be spun up quickly and iterated upon to maximum effectiveness by users, as seen in Figure 13, showing a dashboard view of CM online status.

General / vCMTS Summary Landing Better View ☆ 🔊

Site All ▾

CM Status Per Site ▾

site_name	Online ▾	Unregistered	Offline	Partial
cc	156,629	54	749	9.25 K
ca	133,416	42	421	5.48 K
wa	126,933	24	491	12.90 K
nh	115,936	26	339	7.04 K
fl	108,666	66	1,521	8.00 K
co	105,931	42	803	3.61 K
ca	103,996	31	1,356	4.37 K
nm	103,614	51	1,013	6.74 K
fl	96,197	50	522	6.62 K
fs	85,945	33	498	5.30 K
fl	84,415	27	450	6.01 K
il	83,041	48	1,515	3.70 K
fl	79,430	23	386	5.43 K
tx	78,923	36	434	4.76 K
ma	78,826	21	292	5.28 K
az	76,643	27	264	3.78 K
mn	76,526	17	312	4.23 K
cal	76,355	34	317	4.68 K
al	74,483	119	458	4.83 K
nj80cc100	72,071	34	323	5.61 K
gampec100	71,632	18	416	4.53 K

1 2

Figure 13 – Real-Time Customer Cable Modem Telemetry

With the amount of new and possibly overlapping data, algorithms using machine learning are needed to transform this data into actionable intelligence. For example, with new data comes new thresholds for alarms and triggers. If there are multiple alarms going off, how do we ensure they are correlated with one another? Do they have the same root cause? Can the system help direct the people managing the network to the source of the issue faster? Can the system learn enough to be able to correct itself when it is possible? This is a key part of the operational ecosystem under development.

5. Operational Impacts

5.1. Sustainability

Sustainability has been at the forefront of Comcast practices for many years. Whether it is proactive decommissioning for a network upgrade project, reducing energy use or managing customer network traffic most efficiently, Comcast is considering the impact on our environment. As we build out our network for RBB and continue to upgrade our HFC footprint, we are committed to being environmentally responsible.

5.1.1. Fiber

For RBB, and more generally for all EPON FTTH deployments going forward, integration of PON into existing production DAA system – re-using already built infrastructure for FTTH services – provides construction-related savings, and outside plant power savings.

In addition, because of the significant difference in the signal attenuation of fiber compared to coaxial distribution the PON distribution network is approximately 90% more power efficient than the typical HFC network that would be needed to cover the same low-density footprint. Thus, for RBB, the vBNG with R-OLT represents a significant advantage in terms of sustainability.

5.1.2. Coaxial Cable

As we continue to upgrade the HFC footprint, there are sustainability opportunities available also. Comcast uses thousands of miles of coaxial cable every year. New partnerships are strengthening our recycling programs for coax. Earlier this year Comcast launched an enhanced recycle program for coaxial cables that have reached their end of life. In partnership with Echo Environmental, we can now break down the cable to create new raw materials that can be reintroduced, resold, and reused.

Here is how it works: Coaxial cables are multi-layered “cords” that consist of 27 different polymers, all of which need to be separated to make them usable in new products. Traditional recycling efforts can recover the metals contained within the wires, but not the insulation and jacketing around them. Echo Environmental created a solution to address this issue, all without the use of hazardous chemicals or incineration – and now, thanks to their technology, 70% of our coaxial waste can be recycled for reuse each year, significantly reducing landfill waste. The remaining 30% of our cable waste will still be recycled, as it is now, with plans to recycle 100% for the purpose of reuse in the future.

“It is incredibly gratifying to develop sustainable solutions for underserved market products,” says Brian Hays, Echo Environmental Project Developer. “Our unique process combines traditional recycling systems with methods from other industries, resulting in the ability to extract target polymers cleanly. This outside-the-box thinking not only sets us apart in our industry, but also enables us to deliver meaningful, real-world solutions for our clients and our environment.”

5.1.3. HFC

Comcast has hundreds of thousands of nodes and millions of amplifiers in our HFC network, and all of these devices require power to operate. To support this infrastructure, Comcast has one of the largest power distribution networks in the United States. Operation of this power network with a high level of efficiency is a high priority.

As Comcast upgrades HFC plant, there is major focus on the power efficiency of the network and the health of the grid to make sure it is operating efficiently. The simple act of upgrading 20-year-old amplifiers introduces significantly more efficient amplifier technology – so much so that we can extend the bandwidth to 1 GHz, launch that additional power downstream with a tilt, and still come out approximately neutral on total power consumed, or even slightly ahead. The bps/Hz/W has improved dramatically as families of power amplifiers – silicon-based, gallium arsenide (GaAs), gallium nitride (GaN), and iterations of each, have evolved over the last 20+ years.

New technologies aimed specifically at power savings are now also available, such as digital pre-distortion (DPD), envelope tracking, and automatic bias adjustment for RF load. These technologies are

reflected in our specifications, as are recommendations for power-factor corrected (PFC) supplies inside of our actives. PFC enables the AC supplies feeding the network to operate efficiently.

With historical efficiency gains over time and new technologies aimed at power savings we believe that today's HFC power grid is well-suited to support migration to 10G within the bounds of today's power grid. The perspective here is less about the absolute savings of power, but doing much more for the customer with the available existing power.

5.2. Workforce

As the vBNG and R-OLT solution moves into the field, we will be able to assess the actual implications to the field and to sustainability as identified previously. Another important aspect is how the solution will impact internal teams and field technicians. In 2017, Comcast launched the industry's first large-scale initiative to deploy DAA using RPDs and vCMTS. We did this expecting a bumpy road at the outset as the new technology and practices slowly matured. With vBNG and R-OLT, the advantage of lessons learned from that experience is with us and embedded into practices, documents, and the knowledge base. The deployment plan accounts for the change impacts to these downstream teams. Back-office and XOC support teams are working with this new technology in their labs to determine how it fits into existing processes and what potential impacts it would have to current staffing trends to ensure seamless deployments.

Previous PON solutions were so different than DOCSIS that the thought of having common staff maintain and manage both was a challenge, and at a minimum required technicians supporting both to augment existing knowledge with the "new" skills needed to support PON. Some of this is unavoidable in the last mile, of course. However, by aligning the architecture, platform, tools, and processes, and pushing the fiber and intelligence of the network further into the field, we've significantly converged the access network and correspondingly reduced the need for excessive additional training.

Our field technicians have become DAA experts with over five years of learned experience behind them. Training was developed internally and strong partnerships built between headquarters technologists and the practitioners in the field who would be called upon to operate the network. The smarter components themselves of DAA have played a role also, providing our technicians better insight into the network than they have ever had before once they have learned to use the tools. This commonality benefits the vBNG/R-OLT operational model, reducing the need for substantial new training. At the same time, the new DAA tools, processes, and converged technology offers growth opportunity for these technicians, much like technicians of prior years had in the vCMTS roll-out, further back in time, during the analog-to-digital video conversion.

6. Conclusion

Comcast has been building and operating DAA systems that virtualize CMTS platforms for over five years. We have gotten really, really good at it and have the network performance metrics to prove it. We are in the midst of rapidly scaling this foundational network upgrade across the vast majority of the footprint. The benefits have been enumerated herein – performance, scalability, availability, automation, deep telemetry, flexibility, sustainability... the list goes on. As we make our way along the path to 10G, DAA and virtualization are the most important pieces of the puzzle. Layering on DOCSIS 4.0 is the next step for the HFC network, a high priority strategy being executed to ready the network for multi-gigabit symmetrical speeds.

Exploding onto the scene in the midst of this very active cycle of HFC initiatives is the RBB initiative. Fortunately, having embraced the capability and vision of DAA, leading us to invest heavily in the baseline architecture, Comcast is in an excellent position to deliver best-in-class services to RBB areas using efficient architectures, align with operational simplification and sustainability objectives, and continue growth of our EPON footprint in areas where it makes the most sense to deploy FTTH.

We are not just bridging the digital divide as we turn up RBB – we are installing a 10G-enabling highway that will be maintained and managed with the power of proven DAA tools, back-office monitoring and automation, and deliver that best-in-class experience we strive for all of our customers, from the most dense high rise to the now-farther reaches of the Comcast plant.

Abbreviations

AC	alternating current
AM	amplitude modulation
API	application programming interface
ARP	Address Resolution Protocol
BAU	business as usual
bps	bits per second
bps/Hz/W	bits per second per hertz per watt
CAGR	compounded annual growth rate
CCAP	converged cable access platform
CM	cable modem
CMTS	cable modem termination system
COTS	commercial off-the-shelf
CPU	central processing unit
DAA	distributed access architecture
DAAS	distributed access architecture switch
dB	decibel
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data-Over-Cable Service Interface Specifications
DPD	digital pre-distortion
DPoE	DOCSIS provisioning of EPON
DRFI	[DOCSIS] Downstream RF Interface [Specification]
DS	downstream
DSLAM	digital subscriber line access multiplexer
DWDM	dense wavelength division multiplexing
EOL	end-of-line
FLM	fiber link module
FMA	flexible MAC architecture
FTTH	fiber-to-the-home
EPON	Ethernet passive optical network
GaAS	gallium arsenide
GaN	gallium nitride
GHz	gigahertz

HD	high-definition
HFC	hybrid fiber/coax
HW	hardware
I-CCAP	integrated converged cable access platform
iCMTS	integrated cable modem termination system
IP	Internet Protocol
km	kilometer
LAN	local area network
MER	modulation error ratio
NFV	network function virtualization
NPU	network processing unit
ODN	optical distribution network
OLT	optical line terminal
ONU	optical network unit
OSI	Open Systems Interconnection (the standard's full name is "ISO/IEC 7498-1 Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model – Part 1")
OSP	outside plant
OTL	optical transport link
PFC	power-factor corrected
PHY	physical layer
PON	passive optical network
QoS	quality of service
RBB	rural broadband
RF	radio frequency
RFP	request for proposal
RPD	remote PHY device
R-PHY	remote physical layer
R-OLT	remote optical line terminal
rSwitch	remote switch
SDN	software defined network(ing)
SFP	small form-factor pluggable
SNMP	Simple Network Management Protocol
SW	software
US	upstream
vBNG	virtual broadband network gateway
vCM	virtual cable modem
vCMTS	virtual cable modem termination system
VOD	video on demand
WAN	wide area network
WDM	wavelength division multiplexing

Bibliography & References

[1] "Remarks by President Biden on the Affordable Connectivity Program" President Joe Biden, Briefing Room May 9, 2022

[2] “Comcast RISE, National Initiative to Support Small Businesses, Awards 100 Allegheny County Businesses with \$10,000 Grants,” Comcast, Jul 26, 2022

[3] “42 million Americans don’t have high-speed internet. Local providers may be the key,” Adrian Florido, David Condos, Matt Stout, NPR May 11, 2022

[4] “Access Capacity Planning: Staying Well Ahead of Customer Demand Helped Ensure Stability During COVID-19,” Bruce Barker Jr., Claude Bou Abboud, Erik Neeld, SCTE 2020

[5] “Repair the Ides of March: COVID-19 Induced Adaption of Access Network Strategies” Dr. Robert Howald, SCTE 2020

[6] DOCSIS DCA-MHAv2, Remote PHY Specification, CM-SP-R-PHY-I12-190307, CableLabs

CableLabs[®] Custom Connectivity

An Architecture To Bridge The Digital Divide

A Technical Paper prepared for SCTE by

Darshak Thakore

Principal Architect

CableLabs

858 Coal Creek Circle, Louisville, CO 80023

+1-303-661-3456

d.thakore@cablelabs.com

Craig Pratt

Lead Software Architect, Security & Privacy Technologies

CableLabs

858 Coal Creek Circle, Louisville, CO 80027

+1 303.661.3408

c.pratt@cablelabs.com

Mohan Gundu

SVP, Engineering

Veeva Inc.

164 E 83rd Street

New York, NY 10028

+1 212 535 6050

mohan.gundu@veeva.com

Roger Lucas

SVP, Systems Engineering

Veeva Inc.

Veeva Systems Ltd., Cambridge House, Henry Street, Bath, BA1 1JS, England

+44 7776 234297

roger.lucas@veeasystems.com

Jose Quintero

Senior Director, Innovation Labs

Liberty Latin America

Boulevard Costa del Este, Panama City, Panama

+507 208-5197

jose.quintero@lla.com

Table of Contents

Title	Page Number
1 Introduction.....	3
2 Background	3
3 The Digital Divide	4
4 CableLabs [®] Custom Connectivity Architecture	4
4.1 Architectural overview	4
4.2 The Custom Connectivity Controller	6
4.3 The Custom Connectivity Portal.....	6
4.4 The Custom Connectivity Gateway Agent	6
4.5 The Custom Connectivity Administrative Interface	7
4.6 The Custom Connectivity APIs	9
5 Panama Trial.....	10
5.1 Service Requirements	10
5.2 User Experience Considerations.....	11
6 Implementation Details.....	12
6.1 Network segmentation on shared Access Points.....	12
6.2 Per-device credential management	13
6.3 Roaming within the Multi-AP mesh	14
6.4 Enforcement of Service policies	14
7 Conclusion.....	15
8 Abbreviations.....	15
9 Bibliography & References.....	16

List of Figures

Title	Page Number
Figure 1: Custom Connectivity high-level architecture	5
Figure 2: Custom Connectivity AP Components.....	7
Figure 3: Custom Connectivity Reference UI - Overview Page.....	8
Figure 4: Custom Connectivity Reference UI - Service Detail Page.....	9
Figure 5: Custom Connectivity Component Interoperation	9
Figure 6: Custom Connectivity User Experience	12
Figure 7: Adding a device using a Custom Connectivity-enabled mobile app.....	13
Figure 8: Inter-AP VXLAN, Extender and Access Network Interconnects.....	14

1 Introduction

In today's world, reliable internet connectivity is a necessity. The COVID-19 pandemic has accentuated the need to stay connected while also highlighting the digital divide that exists across the globe. When we talk about the "digital divide", we typically think about serving remote and rural areas. But a lesser-recognized digital divide exists in dense urban areas. Socio-economic factors in developing countries have provided broadband internet access to upper income, educated communities while economically disadvantaged neighborhoods are left unserved or, at best, underserved.

These unserved and underserved areas pose multiple challenges for providing affordable service using traditional deployment models where end-user service is provided as a post-paid recurring subscription via a Consumer Premise Equipment (CPE) installed to a household with a billable home address. These models focus on postpaid subscription for the address where, over time, the ROI model recoups the cost of installation, the CPE, and recurring costs to bring connectivity to the household.

Overcoming the challenges of installation cost, installation logistics, device / CPE security while unleashing the purchasing power of the unbanked in these economically disadvantaged neighborhoods is key to bridging the digital divide and allows for deployment models across high-density, shared residential and MDU communities.

In this paper we present a novel approach based on the CableLabs[®] Custom Connectivity (CC) architecture to delivering internet service, directly to a device and/or to a group of devices without the need for an in-home deployed CPE or requiring a fixed household address for broadband service delivery. This Custom Connectivity architecture, implemented in collaboration with Liberty Latin America and Veeva Inc, enables alternate deployment models utilizing edge compute based shared Wi-Fi access points that provide on-demand virtualized home gateways to subscribers. The rest of the paper is organized as follows

Section 2 (Background) talks about traditional CPE based deployment architecture and some of the shortcomings of that model.

Section 3 (The Digital Divide) dives into the unique constraints in some of the unserved and underserved communities that hinders broadband services delivery using the traditional CPE based model.

Section 4 (Custom Connectivity Architecture) provides an overview of the Custom Connectivity architecture and some of the unique capabilities provided by that architecture.

Section 5 (Panama Trial) describes the technical and business requirements that shaped the trial implementation of the Custom Connectivity architecture to provide broadband services to the underserved and unserved communities.

Section 6 (Implementation) explains the functional components that collectively allows delivery of broadband services directly to subscriber's devices through an on-demand subscriber specific virtual home gateway service that is hosted using the shared-CPE design.

2 Background

Traditionally broadband services are delivered to a physical home address with a consumer premise equipment (CPE) like a cable modem or an ONU terminating the access network. The subscriber's billing relationship is also tied to the address where the CPE is installed. This model has served well over the past few decades where the focus was on delivering cable video as well as broadband services to a household with a limited number of consumer devices connecting to the network. Over time the number of consumer devices connecting to the network have been increasing and continues to increase exponentially. Consumers are also connecting an increasingly wider variety of devices to the network and these devices may have specific connectivity requirements. For example, security cameras require

increased upstream bandwidth compared to a Smart TV and/or a work computer. These trends have highlighted several constraints with the existing deployment model:

- The need for a CPE to be installed and activated for each subscription adds to the cost of servicing a subscriber. This cost goes beyond just the cost of the CPE and includes the costs associated with inventory management, installation logistics, CPE security and maintenance etc.
- Operators typically do not have visibility into the devices behind the CPE that are connecting to the network. This makes it difficult to troubleshoot problems caused by individual devices and makes it harder to deliver custom service(s) to a device or group of devices.
- The billing for the subscription is typically tied to the home address where the CPE is installed instead of the devices to which the services are being delivered to. This prevents the operators from offering device-centric value-added services that can optionally be tied to different payers.
- It prevents sharing of CPE across subscribers which, as described below is one of the factors that contributes towards the difficulty in offering broadband services in certain unserved and underserved communities.

3 The Digital Divide

The global COVID-19 pandemic brought into focus the importance of reliable internet connectivity in our lives and the devastating consequences to people and communities that did not have access to reliable internet connectivity. While a majority of the population in developed countries were able to sustain themselves and their families by depending on reliable broadband services, there was a significant population in low-income communities in Latin America and other developing countries across the globe that were either unserved or underserved in terms of broadband connectivity. The impact to these communities was substantial. Children were unable to access educational content and adults were largely unable to participate in today's digital economy. A lot of these communities were composed of daily wage earners and the gig-economy participants that primarily lived in high-density neighborhoods. These neighborhoods had remained largely unserved or underserved because the traditional CPE based service delivery model did not scale economically in these neighborhoods. A number of factors contributed to it being non-viable for the operators. First and foremost, the costs associated with the CPE (see [Section 2](#)) resulted in a negative ROI in most cases due to the low-ARPU in these communities. A significant population in these communities were unbanked which made it difficult to establish a BSS relationship with the subscribers. In certain instances, there wasn't a fixed home address to deliver the service to. The subscribers preferred a subscription model that was more on-demand, where they could stop/disable their service on certain days and not have to pay for it. Such on-the-fly service activation/deactivation is hard if not impossible to support with the existing home-address based billing model. These factors contributed to creating a digital divide for these communities and put them at a significant disadvantage during the pandemic.

4 CableLabs[®] Custom Connectivity Architecture

4.1 Architectural overview

The Custom Connectivity architecture encompasses a set of technologies that collectively enable the delivery of broadband services using an alternate device-centric service delivery model. In particular, Custom Connectivity provides per-device credentials and per-device policy for consumer-grade wireless

devices – essentially providing enterprise-grade device access and management capabilities to non-enterprise devices.

The core of the Custom Connectivity system is the Custom Connectivity Controller (CC Controller) – which coordinates wireless access points to provide access for the provisioned network services and associated devices. The Controller interfaces with the operator OSS/BSS systems and customers via the Custom Connectivity Portal (CC Portal) – which is adapted to the operator's infrastructure. Service and device configuration and monitoring, as well as AP provisioning and deployment, can be performed using the Custom Connectivity Controller API – either via a stand-alone operations interface or integrated with existing network operations interfaces and tools. Similarly, the Custom Connectivity telemetry interface can be used to monitor the health of the network(s) and, along with the network operations interface, provide customer support.

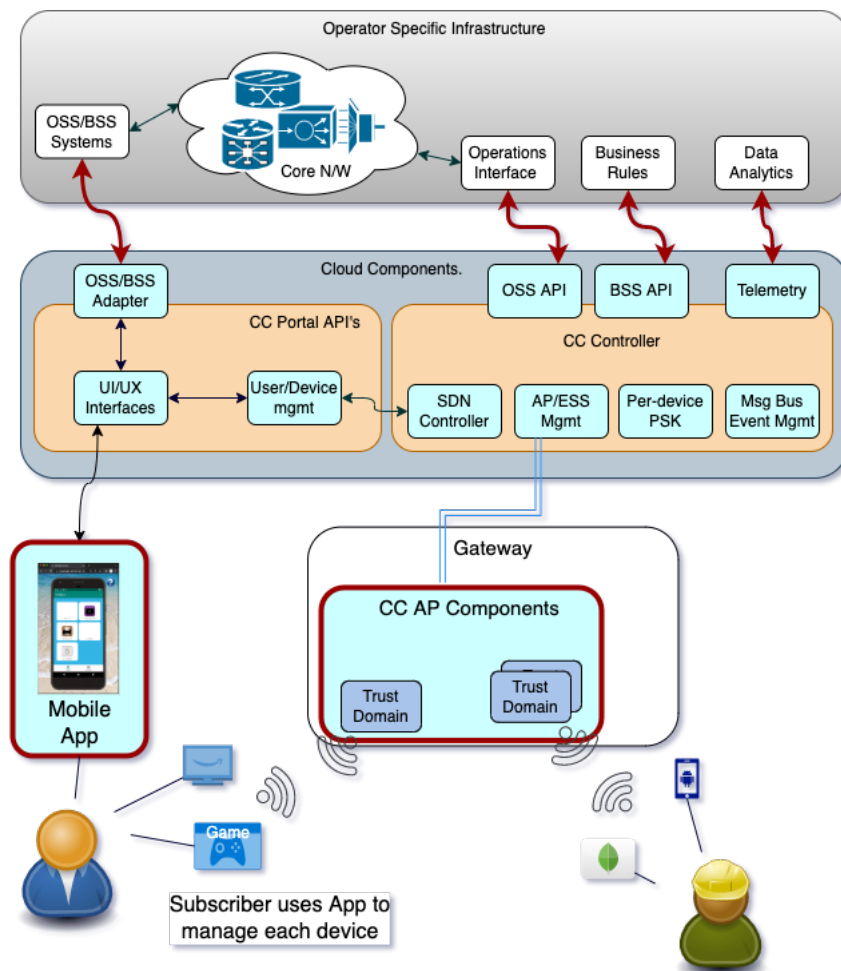


Figure 1: Custom Connectivity high-level architecture

4.2 The Custom Connectivity Controller

The Custom Connectivity Controller contains the information regarding ... and operational state information for the Custom Connectivity network and has the responsibility of keeping all parties in a Custom Connectivity network synchronized so that all elements are operating with a common network/device model. In particular, the Controller provides:

- REST API's to manipulate the Custom Connectivity Service model
- The Custom Connectivity PSK Engine – for identifying valid PSKs and determining device associations
- A reference operator interface for provisioning Custom Connectivity-enabled Wi-Fi access points
- An MQTT Broker to communicate model changes to provisioned access points
- REST API's to for access points to communicate Wi-Fi telemetry and device connectivity status

4.3 The Custom Connectivity Portal

The Custom Connectivity Portal is the communication bridge between the Controller, operator OSS/BSS infrastructure, and mobile application communication. The Portal is intended to be highly adaptable to the operator network and infrastructure. The Portal is responsible for:

- Interfacing with the provider's customer-facing systems/applications – which can be used by the customer to setup their Custom Connectivity Service, determine funding, establish per-device access rules, onboard Wi-Fi devices, confirm device additions, and manage their services/devices.
- Communicating with the Custom Connectivity Controller to manipulate the service/device objects on behalf of the customer.
- Communicating customer-initiated service/device changes with OSS/BSS infrastructure.
- Associating Custom Connectivity Services with customer subscription plans and maintaining service according to the subscribed plan and payment status – including handling expiration times/dates.
- Process notifications from the Controller regarding service/device object status changes (e.g. when a device connects/disconnects) and communicating the changes to the customer and OSS/BSS infrastructure.

4.4 The Custom Connectivity Gateway Agent

Every wireless access point that supports Custom Connectivity must include a Gateway Agent provisioned with credentials to connect and communicate with the Custom Connectivity Controller. The Gateway Agent is responsible for:

- Establishing and maintaining a connection with the Controller using credentials provisioned by the vendor and/or operator
- Performing initial setup of the AP by interrogating the Controller and setting up Wi-Fi credentials, VLANs, bridges, and inter-AP VXLAN connections to support the Custom Connectivity services and devices provisioned for the AP the agent is running on
- Handling DHCP requests for Custom Connectivity devices and providing IP addresses according to the Custom Connectivity service model provided by the Controller
- Delegating Wi-Fi 4-way handshake authentication requests for devices which are being onboarded to the Custom Connectivity Controller – which in-turn performs the necessary decryption steps to determine if/which device is associated with the presented password

- Handling MQTT messages from the Controller to update the access point to reflect changes to Custom Connectivity services/devices associated with the access point
- Setting up and maintaining inter-AP VXLAN tunnels – to support the CC inter-AP mesh
- Providing per-device telemetry to the Controller
- Routing Internet inbound/outbound traffic to the access network
- Routing intra-Service device-to-device traffic across inter-AP tunnels when devices within the same Service are connected to different APs (mesh routing)

The Custom Connectivity Reference Implementation includes a reference gateway implementation with the components outlined in Figure 2: Custom Connectivity AP Component

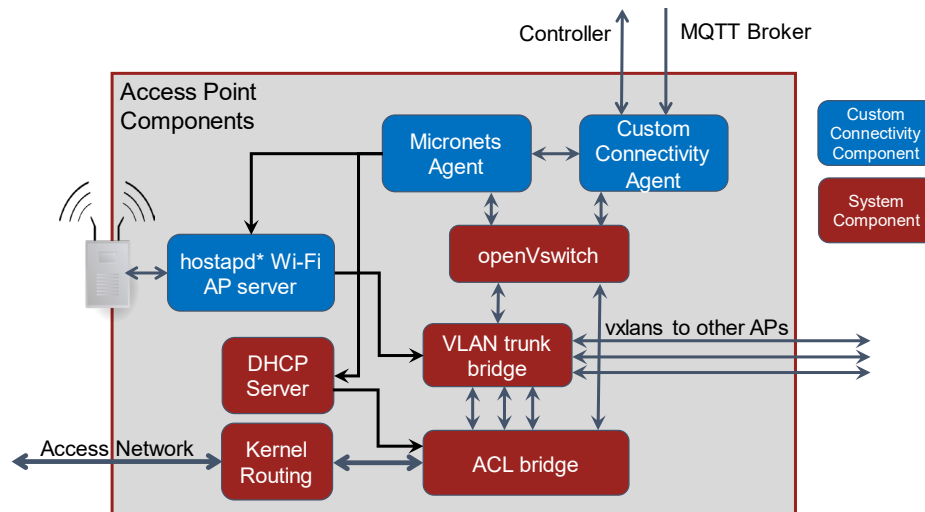


Figure 2: Custom Connectivity AP Components

4.5 The Custom Connectivity Administrative Interface

For monitoring and managing a Custom Connectivity system by network operators and customer support the Custom Connectivity Controller provides a REST Operations API and notification support for enabling web- and/or smart device-based management interfaces. This Operations API can enable existing operation center user interfaces to incorporate Custom Connectivity administrative/support functions and/or enable the implementation of stand-alone management interfaces.

The Custom Connectivity Reference Implementation includes a reference web interface which allows for:

- The enumeration of Custom Connectivity services, devices, AP groups, and APs and their various association
- The provisioning of APs using public keys and registration tokens
- The monitoring of devices with live updates on connected devices and associated APs
- The creation of APs and AP groups, as well as associating/disassociation of APs with AP groups

Some illustrations of the Custom Connectivity management reference web interface can be found in Figure 3 and Figure 4.

NRC A

AP Groups Access Points Services Devices

netreach-bethany AP Group Enabled ACTIONS

Details EDIT

Name netreach-bethany
 UUID cf2bcf41-f5f1-46fe-9f2c-5e7f7b3148de
 SSID netreach-bethany-01

Access Points + ASSOCIATE ANOTHER AP

Name	Status	Serial	Management Address	Registration Token	VXLAN Endpoint
CraigTestAP-2	ONLINE	CRAIG-TEST-AP-02	10.10.1.163		
CraigTestAP-3	ONLINE	CRAIG-TEST-AP-03	10.10.1.149		
CraigTestAP-4	ONLINE	CRAIG-TEST-AP-04	10.10.1.162		
CraigTestAP-5	ONLINE	CRAIG-TEST-AP-05	10.10.1.218		

Services

Name	Subscriber ID	Plan ID	Micronet ID	Micronet Subnet
craig@netreach-bethany	359c71ce-d82f-558d-99fa-5dd61a3d300a	5x5	CRAIG@NETREACH-BETHANY_0c93e631	10.0.5.0/24
craig2@netreach-bethany	fd61b632-6299-576e-ad07-02109b0247f0	5x5	CRAIG2@NETREACH-BETHANY_d4628143	10.0.6.0/24
craig3@netreach-bethany	d6f9d712-19fa-59ea-a202-6cc3058af738	5x5	CRAIG3@NETREACH-BETHANY_88bceeb1	10.0.7.0/24
craig4@netreach-bethany	4fe63585-3923-5526-8e71-bd950bfbd5c1	0x0	CRAIG4@NETREACH-BETHANY_eb8c6bb3	10.0.8.0/24

Figure 3: Custom Connectivity Reference UI - Overview Page

NRC A

AP Groups Access Points **Services** Devices

craig2@netreach-bethany Service Enabled ACTIONS

Details EDIT

Name craig2@netreach-bethany
 UUID d4628143-0246-4ef0-a6b7-5facd18bab39
 SSID netreach-bethany-01
 AP Group [netreach-bethany](#)
 Subscriber ID fd61b632-6299-576e-ad07-02109b0247f0
 Plan ID 5x5
 Micronet ID CRAIG2@NETREACH-BETHANY_d4628143
 Micronet Subnet 10.0.6.0/24
 Micronet Gateway 10.0.6.1
 VLAN 6
 Number of Devices Allowed 5
 Max Upstream Kbps
 Max Downstream Kbps
 VXLAN Endpoint
 Expires Jan 9, 2023, 8:19 AM

Devices

Name	Device Type	MAC Address	IP Address	Passphrase	Connected	Associated AP
CraigPi-4	Raspberry Pi	b8:27:eb:6e:3a:6f	10.0.6.6	*****	Connected	CraigTestAP-4
My SmartPhone	SmartPhone		10.0.6.10	*****	Not Connected	CraigTestAP-5
CraigPi-5	Computer	00:12:7b:21:6c:b6	10.0.6.11	*****	Connected	CraigTestAP-5
CraigPi-2	Computer	b8:27:eb:f2:7b:c0	10.0.6.12	*****	Connected	CraigTestAP-3

Figure 4: Custom Connectivity Reference UI - Service Detail Page

4.6 The Custom Connectivity APIs

The Custom Connectivity APIs are used to manage a common Custom Connectivity entity model. The ability to modify different elements of the model is grouped into four functional areas that have group-level access controls.

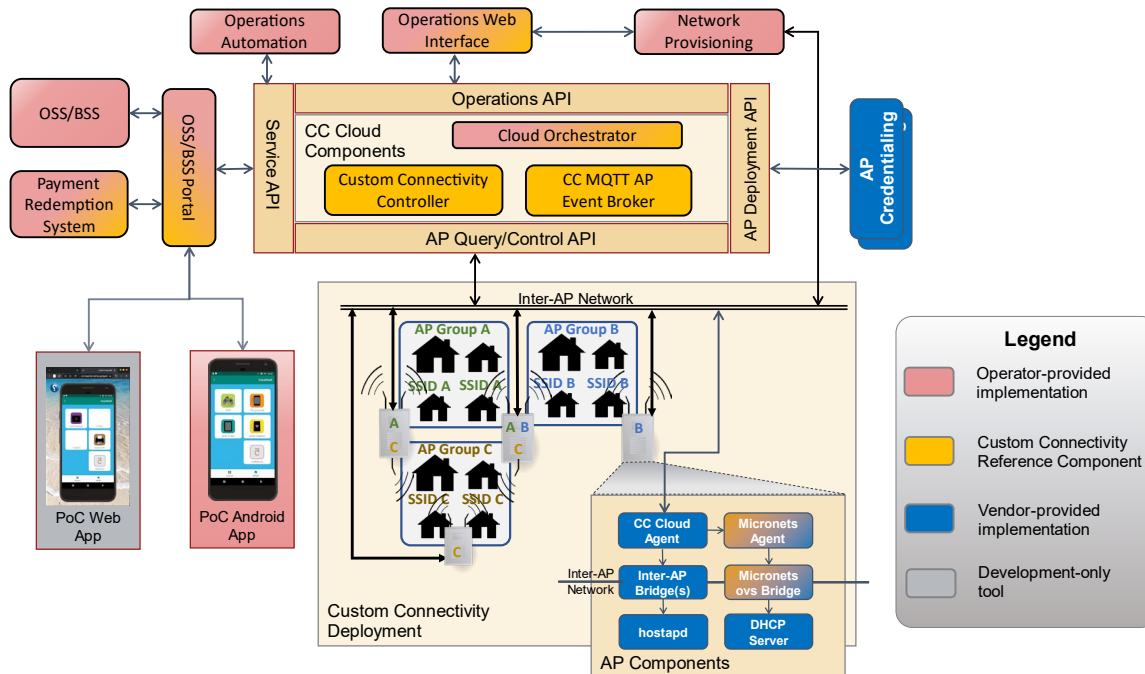


Figure 5: Custom Connectivity Component Interoperation

The Custom Connectivity Controller APIs help ensure interoperability between vendors components, operator components, and the Controller. The intent of the Custom Connectivity architecture – and having well-defined APIs - is to enable either the operator or a service provider to host and manage the Controller on behalf of the operator. This can enable different degrees of regional colocation based on desired cost, reliability, and scaling factors.

The Custom Connectivity APIs are divided into the following functional groups:

- **Service API:** While the Custom Connectivity Portal primarily communicates via interfaces which are OSS/BSS-specific, Portal implementations expose a single REST endpoint that the Custom Connectivity Controller uses to notify the Portal of changes to Custom Connectivity Services and associated Devices. The Portal in-turn uses the Controller APIs to enact changes and enable/disable services and associated devices based on customer subscription/payment changes, service expirations, and/or customer preference.
- **AP Deployment API:** Supports the creation of authentication credentials for APs and putting APs into/out of service. This interface enables credential provisioning in a variety of ways – either at time of AP manufacture/configuration or at initial power-on the AP using provisioning tokens.

APs can also be put into/out of service in a controlled fashion to facilitate replacement, relocation, debugging, and/or network upgrades.

- **AP Query/Control API:** Custom Connectivity-enabled APs are informed of changes in the Custom Connectivity data model via MQTT and update their internal state by performing queries to the Controller regarding affected model elements. APs also use these APIs to update the state of Device elements based on physical device state (e.g. “authenticated” or “connected”) and connection quality (e.g. the station’s RSSI and error rate).
- **Operations API:** These APIs enable operators to administrate and automate the Custom Connectivity system. Network operators and automated tools can monitor the state of the APs, Services, and Devices to help identify potential issues in concert with network backhaul monitoring. APs can be organized/reorganized into different AP groups to optimize the access and backhaul network and ensure sufficient wireless coverage. This API also enables customer service to diagnose device and network connectivity issues and take corrective actions.

5 Panama Trial

5.1 Service Requirements

The Custom Connectivity Architecture described in Section 4 provides an alternate service delivery model that is based on a shared-CPE paradigm and rather than linking a subscribers’ broadband services to the CPE, it focuses on providing those services directly to the devices. This paradigm made the architecture suitable for bridging the digital divide and providing broadband services to the unserved and underserved communities (see [Section 3](#)). Our goal was to validate the architecture by running a limited trial in a representative unserved community. Consequently, a small neighborhood in Panama was selected for the trial where the houses were grouped along a street with utility poles along the street. The service would be delivered through external AP’s mounted to the utility poles and the subscriber’s devices would connect directly to these AP’s. The following technical and operational requirements were addressed for the trial.

- Service to each subscriber was to be provided directly over Wi-Fi without the need to run any wiring to the subscriber’s home
- The subscriber should be able to connect their devices anywhere within the service area and not have to select any specific AP to connect to (seamless to them)
- The subscriber’s connectivity experience should resemble a traditional home Wi-Fi experience where they can connect any device including Smart TV’s, printers, casting devices, tablets, Chromebooks etc. and be able to discover and communicate amongst them.
- The subscriber should be able to connect any standard Wi-Fi enabled device without having to perform any special procedure on those devices (i.e no captive portal interactions, no mac address registration, no custom software). The only interaction required on the device to be connected is to select the service set identifier (SSID) and enter the Wi-Fi passphrase.
- The service provider has the capability to apply rate limits (packet data rates) by subscriber independently, with option to provide different service tiers per subscriber.
- Capability to add extenders (future capability) for improved indoor coverage (with Wi-Fi as a backhaul from the extender)
- Be able to use any backhaul technology including standard FTTH/xPON from the AP’s/shared-CPE’s.

5.2 User Experience Considerations

In addition to the technical and operational requirements described above, another goal for the trial was to ensure that the entire service experience, from service activation to service and device management to billing and payment – was to be kept simple and familiar to the subscriber. In this neighborhood, most subscribers were accustomed to activating and managing their pre-paid cellular plans using pre-paid vouchers that they purchased from the local stores. The following requirements were taken into account while developing the overall user experience

- Since the subscribers were accustomed to using pre-paid vouchers, we leveraged the existing voucher-based BSS system and integrated it with the BSS interfaces on the Custom Connectivity Portal. This allowed the users to use the same pre-paid vouchers to either renew their cellular plans or activate their Wi-Fi based home broadband service plan.
- The only requirement imposed on the user was for them to install and/or use a mobile app that allowed them to self-manage their services - creation of a new account, redeeming vouchers, service activation, adding/managing devices and their credentials etc.
- The subscriber should be able to self-manage their subscription, i.e be able to choose how long they would like their service enabled and be able to suspend their service if they did not need it on a given day.
- The subscriber should be able to renew their subscription as and when needed.

The sequence diagram in Figure 6 shows the overall user experience that was developed and provided in the trial.

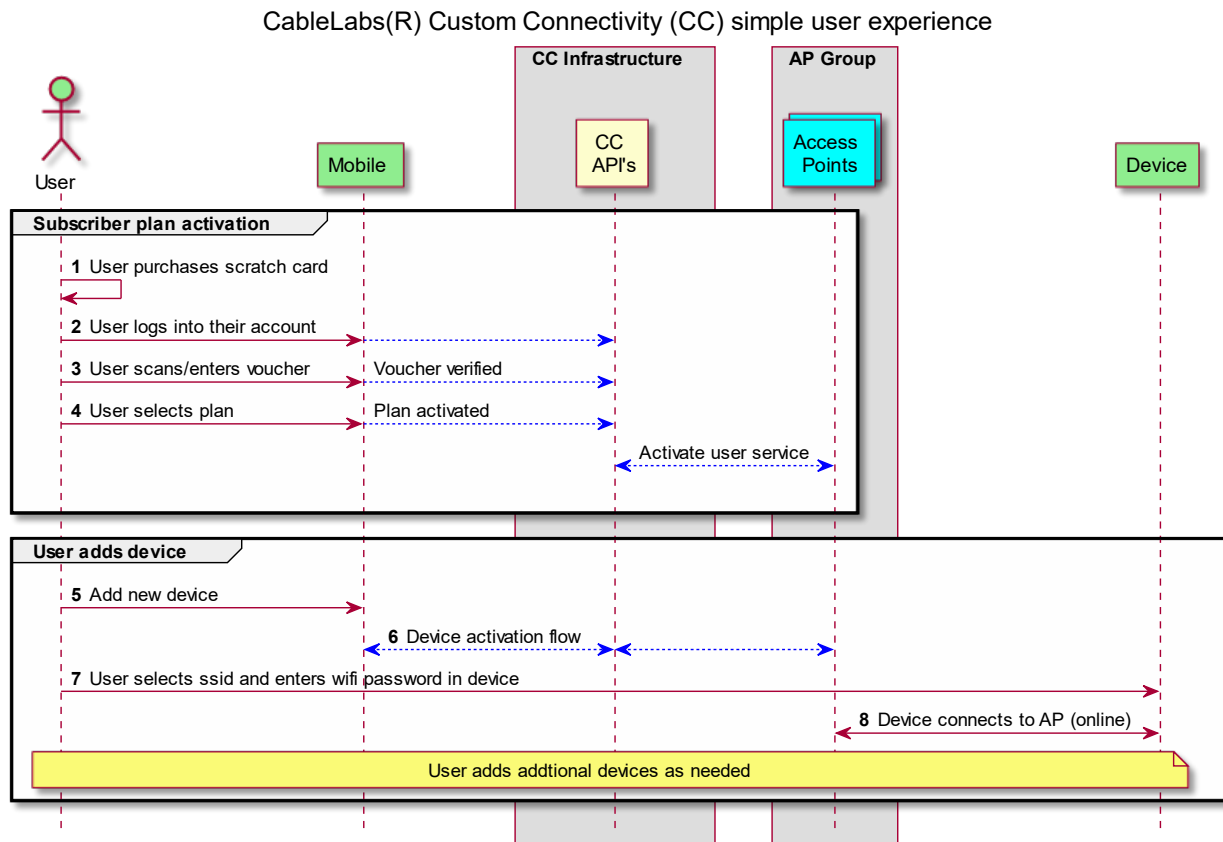


Figure 6: Custom Connectivity User Experience

6 Implementation Details

CableLabs[®] Custom Connectivity offers to deploy Internet access without a drop, a CPE, or an in-home wireless access point. And unlike hotspot or captive portal solutions, Custom Connectivity offers to also provide a complete home network that works with all Wi-Fi devices - including smart TVs, IoT, and home assistant devices. In short Custom Connectivity offers a full CPE+AP experience without a CPE.

We attempt to provide some details below on exactly how Custom Connectivity provides this functionality – and how it can be done in an interoperable multi-vendor way.

6.1 Network segmentation on shared Access Points

The logical segmentation of physical ethernet networks using VLANs has been a reality for decades. When Wi-Fi technology (see [2][3]) was adapted for enterprise use, capabilities were added to enable devices to be authenticated via the use of a user-associated identity/credential. Devices could all join the same SSID with each device being optionally designated for a particular VLAN. While these enterprise features require special Wi-Fi device support (that most home devices don't have) and a complicated AAA server to create/authenticate credentials, this fundamental capability for per-device credentials and VLAN association was and is present in the core Wi-Fi technology today. (see [1])

Using its unique device onboarding system, Custom Connectivity leverages the core capability present in Wi-Fi to associate non-enterprise Wi-Fi devices with an identity, encryption key, and VLAN all on a shared SSID. This means that devices cannot discover, communicate, or observe each other's traffic without explicit policy enabling them to do so. Custom Connectivity utilizes this separation capability to group devices by household (and optionally subzones within a household) into discreet segments – each associated with a VLAN – providing the same separation for a household as one would get with a discreet in-home Wi-Fi AP and switch.

6.2 Per-device credential management

As part of the onboarding process, the Custom Connectivity system provides each device a unique passphrase which is used to provide the device access and, additionally, as a means to robustly identify the device. The process for the user is straight-forward:

1. Using a mobile app (logged into the account provided by the operator), the user selects an option to add a new device to their service (see Figure 7),
2. The user enters a couple details about the device (e.g. a name and device type),
3. The app provides an SSID and passphrase for the new device,
4. The user enters the SSID and passphrase into the device,
5. When the new device attempts to authenticate with the AP, the AP – in conjunction with the Custom Connectivity Controller – performs the necessary cryptographic logic to determine which passphrase was provided by the device,
6. The device completes authentication, and the AP associates the Wi-Fi session with the device – setting up necessary interconnects for the device and signaling the Controller,
7. The controller signals all other APs of the new device to enable interconnects and roaming.




Figure 7: Adding a device using a Custom Connectivity-enabled mobile app

The details of how robust association of devices with unique credentials can be achieved is discussed in a previous SCTE paper, *Wi-Fi Passwords: The Evolving Battle Between Usability and Security* [1].

6.3 Roaming within the Multi-AP mesh

A single AP can only cover a limited area, so Custom Connectivity enables a large number of APs to be grouped together to form an AP group and hence cover a larger area such as a neighborhood or campus. All APs in the mesh will offer identical service, so a client may connect to any AP in the group and even move between different APs within the group at any time. Irrespective of the AP within the group that the client connects to, the client will see their own home network and all their devices connected to it. From a client's perspective, the Custom Connectivity AP mesh acts as a single large AP covering the entire area with a single SSID.

When devices which are part of the same household service are connected to different APs in a Custom Connectivity AP group, inter-device packets must be exchanged between the APs to facilitate the inter-device communication. In Custom Connectivity, this inter-AP communication is accomplished with VXLANs. When Custom Connectivity recognizes that there are devices in the same household/segment connected to multiple APs, the APs establish bi-directional VXLAN channels between the APs to facilitate the device-to-device communication over a shared AP backhaul network.

For device-to-Internet traffic the Custom Connectivity architecture enables each AP to independently egress/ingress Internet access for connected devices without any intermediate network components. See Figure 8 for an illustration of how APs setup VXLAN, inter-AP, and extender connections.

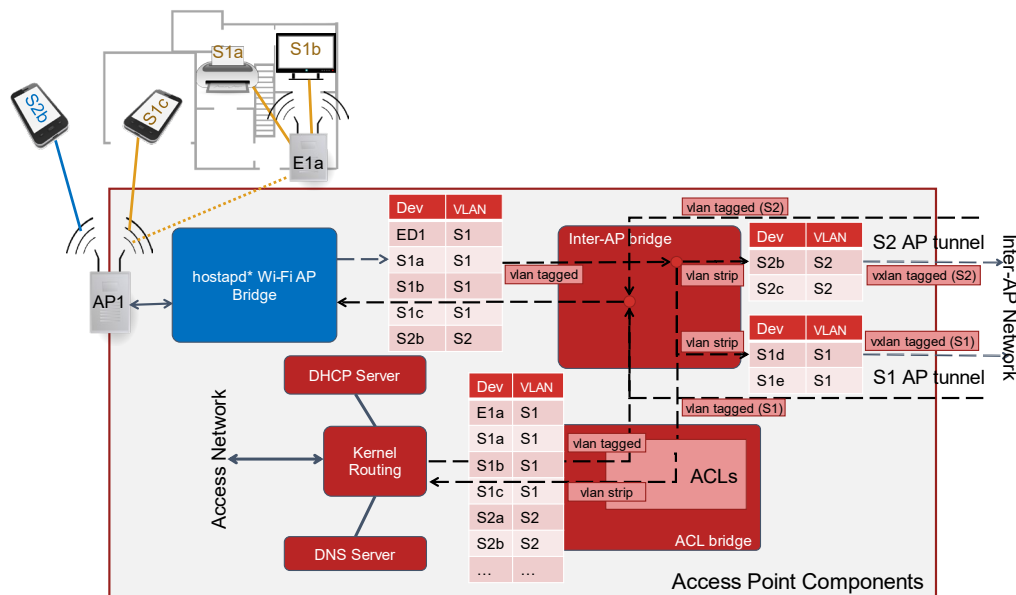


Figure 8: Inter-AP VXLAN, Extender and Access Network Interconnects

6.4 Enforcement of Service policies

Custom Connectivity's device-level identity and visibility allow for a wide variety of user-defined and operator-defined policies to be applied at both the device level and the service level.

For instance, operators likely will want to offer particular service offerings capped particular upload and download speeds. Custom Connectivity allows the maximum downlink/uplink BPS to be specified on a per-service basis – which can be specified by the operator via the Portal. As APs can independently egress traffic for a service, they share their respective egress/ingress rates to collectively enforce the limits using a distributed traffic shaping algorithm. This ensures that the subscriber’s service policy isn’t exceeded, irrespective of the distribution of their devices across the group.

Another example is the ability for users to assign per-device scheduling policy. Custom Connectivity allows for access hours limits to be defined and enforced more robustly than similar policy systems on other offerings. Since the policy is associated with both the identity and credentials of the device, a simple trick like changing the device’s MAC address will not result in policy circumvention.

Many other policies are easily defined for Custom Connectivity – with the same ability to be tightly bound to the service and/or device identity and associated credentials.

7 Conclusion

The success of the trial allowed us to validate the capabilities of the Custom Connectivity architecture and helped us in taking one step closer to addressing the digital divide discussed in [Section 3](#). We also learned a number of lessons from the trial that helped us refine our architecture as well some of the usability requirements.

This proof of concept helped us assert that as part of the overall evolution of cable broadband services and the build out of the 10G platform, the CableLabs[®] Custom Connectivity architecture provides another tool in the overall toolkit to the operators in the form of an alternate device-centric service delivery model.

8 Abbreviations

AAA	authentication, authorization and accounting
AP	access point
API	application programming interface
BSS	business support system
VLAN	virtual local area network
VXLAN	virtual extensible local area network
MQTT	MQ Telemetry Transport
OSS	operational support system
REST	Representational State Transfer
RSSI	received signal strength indicator
SCTE	Society of Cable Telecommunications Engineers
SSID	service set identifier

9 Bibliography & References

- [1] Wi-Fi Passwords: The Evolving Battle Between Usability and Security; Society of Cable Television Engineers; https://scte.org/documents/3070/1742_Pratt_3216_paper.pdf
- [2] WPA3[™] Specification Version 2.0; Wi-Fi Alliance
- [3] IEEE 802.11i-2004: *IEEE Standard for information technology-Telecommunications and information exchange between systems-Local and metropolitan area network - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*; Institute of Electrical and Electronics Engineers
- [4] CableLabs[®] Micronets; <https://www.cablelabs.com/technologies/micronets>
- [5] Device Provisioning Protocol Version 1.2; Wi-Fi Alliance
- [6] RFC 7348; Virtual eXtensible Local Area Network (VXLAN)

Challenges, Considerations, and Best Practices for Secure SD-WAN Operationalization for Business Services

A Technical Paper prepared for SCTE by

Xin Huang

Sr. Principal Engineer, Product Development Engineering
Comcast Cable
1800 Bishops Gate Boulevard
Mt. Laurel, NJ 08054
Xin_huang@comcast.com

Joshua Horton

Director, Product Development Engineering
Comcast Cable
1800 Bishops Gate Boulevard
Mt. Laurel, NJ 08054
joshua_horton@comcast.com

Hung Le

Sr. Principal Engineer, Product Development Engineering
Comcast Cable
11951 Freedom Dr., STE 900
Reston, VA 20190
hung_le@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Challenges and Solutions	3
3. Keep-It-Simple.....	4
4. Automation-First Lifecycle and Test-Driven Development.....	7
5. Data-Driven Proactive Monitoring	9
6. Conclusions.....	12
Abbreviations	12
Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 – SD-WAN Platform Architecture Design	5
Figure 2 – SD-WAN Platform Architecture Optimization	6
Figure 3 – Test Automation Framework.....	8
Figure 4 – SD-WAN Platform Data-Driven Monitoring.....	10
Figure 5 – Benefits of Observability	11
Figure 6 – Microservices Growth vs Platform Capacity	12

List of Tables

Title	Page Number
Table 1 – Per Platform Cloud Footprint Reduction for Different Scenarios	6
Table 2 – Platform Lifecycle Automation Benefits	9

1. Introduction

Network connectivity products such as Software-Defined Wide Area Network (SD-WAN) or cybersecurity are becoming critical enablers of needed connectivity as businesses of all sizes re-configure and consolidate network services and solutions using software-driven and virtualization technologies. Large service providers such as multiple systems operators (MSOs) who want to provide innovative solutions in this space have been developing expertise that can drive customer success. This paper will use insights from real projects to detail the ways in which those wishing to deploy these technologies can be guided by simple principles and industry best practices to kickstart successful networking platform initiatives.

In the past, connectivity providers have been very successful deploying networking gear, operating it at scale, and delivering value by executing well. They may not have created their business engines around a core of software technologies or on large-scale virtualization in quite the same way the largest internet platforms have been driving their businesses. Hardware-based technologies provide very high performance in a reliably fixed and predictable architecture, with key differences being variations in speeds, feeds, protocols, or connectors. The technologies and expertise needed to launch software products, by contrast, can often feature dynamic architectures having unpredictable variations, needing data-driven insights to manage.

Organizations with different strengths/expertise who now wish to adopt technologies that have grown up in the era of large internet platforms must become skilled in techniques tied to the software-based infrastructure which brought those platforms to life. Comcast's launch of software-driven networking services could be considered as one such case study. Luckily, many of the lessons that were learned were related to a few fundamental software best practices, which are very well-documented and to which all modern practitioners should already have access.

In this paper, we will share a few key challenges and lessons learnt through real projects and detail the ways in which those wishing to move ahead in deploying networking software at scale can be guided to successfully kickstart similar products and platform initiatives.

2. Challenges and Solutions

Software-based network services such as software-defined networking/network function virtualization (SDN/NFV) connectivity approaches involve abstracting network functions and services out of silicon and into software, separating the connectivity service into administration plane, control plane, and data plane vectors, each of which is coordinated and controlled via software components operated as a platform.

The data plane is transported physically over white-box devices with all the logic for routing and services instantiated across the platform. Control plane services allow update to the configurations and changing the behaviors of the data plane, including adding software-based network functions, such as firewall, traffic steering, or anti-virus in line with packet processing (in a process called "service chaining"). Customers and operators manage the control plane services via the admin plane, exposed via application programming interface (API) or graphical user interface (GUI) into complex orchestration software.

In hardware-based network services, all aspects of the service are well-defined and fixed into the design of the devices being deployed. The ways in which the devices can be configured is therefore more or less pre-determined by the vendors. If more capacity is needed, then new hardware is purchased and deployed.

In software-based network services, all aspects of the service might be distributed across multiple software components, each of which could be instantiated on a variety of different hardware options. Each choice in architecture and configuration results in a potentially wide range of capabilities and trade-offs that must be evaluated and carefully calibrated. Different layers of software abstraction, of operating systems and virtualization layers, and of interoperability between the layers, creates combinatorial numbers of variations which could affect the behaviors of the customer service.

Embracing this complexity and developing techniques to make the problems tractable and in line with the strategies for being handled within hardware was a core part of the challenge in successfully deploying a software-based networking product.

Below is a list of high-level principles we embrace in design and operations in order to resolve the above challenges:

- Keep it simple: Leverage **cloud-native architecture** and **standard technologies** like edge routers, border gateway protocol (BGP), generic routing encapsulation (GRE) tunneling, proxies, and load balancers (LB) for system integration.
- Emphasis on standardization of configuration in version-control, combined w/ logical inventory in change management database (CMDB) plus strict change control policies to facilitate **automation-first deployment, move/add/change/delete (MACD), & disaster recovery (DR)** for operations to reduce unforced errors
- Embrace **test-driven development** using fully-automated unit and integration tests to ensure version-after-version quality consistency
- Forwarding to data lake, aggregation of time-series data combined with intelligent machine learning, to achieve **observability and data-driven capacity planning**

3. Keep-It-Simple

When introducing the SD-WAN product, our main goal is to integrate vendor solutions seamlessly with our existing eco-systems and business strategies.

Nowadays cloud infrastructure virtualization has become a dominating technology because of the set of benefits it brings. These include a wide range of hardware selections, improved economies of scale, reduced costs to resource efficiencies, operational flexibility, and faster time-to-market, etc. To keep our product competitive in the market, we embraced the “keep-it-simple” design principle and make use of industry standard technologies and best practices. For example, we adopted cloud-native architecture design, deployed our platform services in geographical-redundant data centers (DCs), and utilized standard networking technologies to facilitate communications between DCs and to the Internet. Figure 1 below illustrates the high-level architecture design of our cloud-based SD-WAN platform.

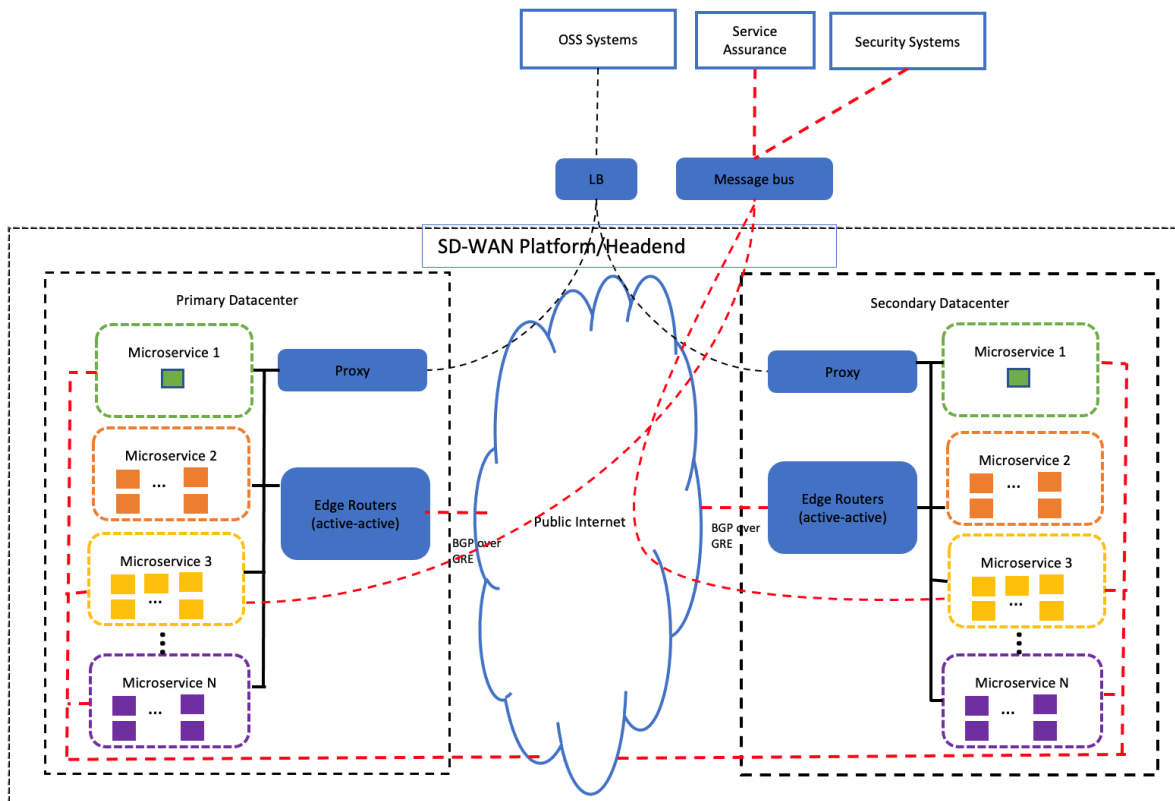


Figure 1 – SD-WAN Platform Architecture Design

Microservices are the core of cloud-native architecture design. The complex SD-WAN control plane and management plane functionality is broken down into multiple microservices, each of which serves a specific function and could scale in/out independently based on its workload. We take advantage of microservices because they support DevOps, and improve scalability, while also allowing flexibility with respect to infrastructure growth. Within the same data center, microservices are interconnected with one another via traditional technologies, e.g., application programming interfaces (APIs), load balancers (LBs), etc.

To achieve SD-WAN platform high-availability (HA) and guarantee business continuity, we deploy our platform (including microservices and data) across geographically diverse DCs (i.e., located in different regions of the country). This geographical redundancy approach is an industry standard best practice that provides business resiliency against natural disasters and catastrophic events which might bring a DC down for certain period of time. Even when disaster happens and one of the DCs is down, our platform remains available since services are still running in the other DC. Once the impacted DC is recovered, everything returns to normal. Different microservices in Figure 1 have different HA designs (e.g., active-backup, active-active, or cluster-based) depending on the nature of the functionality and requirements.

To keep the design simple but efficient, we also adopted standard network technologies to facilitate the inter-connectivity between data centers and the communication between the SD-WAN platform and applications/devices from the Internet. As shown in Figure 1, edge routers and Border Gateway Protocol (BGP) over Generic Routing Encapsulation (GRE) are used to provide HA and dynamic traffic steering for components to communicate with each other between DCs. Standard proxy and load balancer

technologies are adopted to facilitate the communication between upstream systems and components in our platform. Message buses are used to distribute platform telemetry data to service assurance systems and security monitoring systems.

With these high-level design principles in mind, we keep refining and optimizing our platform architecture to make it more scalable. For example, we observed that breaking down big microservices into smaller microservices is an effective way to reduce per-platform cloud footprint. Figure 2 and Table 1 show that our optimization could successfully reduce the per-platform cloud footprint by 9%, 9%, and 62%, respectively, for the best case scenario, the average case scenario, and the worst case scenario. This benefit grows with the platform capacity. When the platform capacity doubles, per-platform cloud footprint could be further reduced by 19%, 19%, and 82%, respectively, for the best case scenario, the average case scenario, and the worst case scenario.




Figure 2 – SD-WAN Platform Architecture Optimization

Table 1 – Per Platform Cloud Footprint Reduction for Different Scenarios

Scenarios	Current Capacity	1.5 X Current Capacity	2 X Current Capacity
Best case	9%	21%	19%
Average case	9%	15%	19%
Worst case	62%	82%	82%

4. Automation-First Lifecycle and Test-Driven Development

When silicon-based network capabilities get implemented as software distributed across multiple components, whether or not on a cloud platform, the pattern is still microservices, and each software component provides a subset of overall system functionality. An orchestration function is then frequently placed between the components to coordinate the components into more advanced processes, such as enabling a new network feature by updating the customer configuration. In some cases, each microservices component might be a separate software product in itself, with its own behaviors and release schedule. This distributed structure provides maximum flexibility and reuse, and can allow for simplification and different optimizations for the operator, at the possible cost of complexity of implementation.

In our case, service reliability was of paramount concern. Given the wide range of network features and functionality being launched, the tight integration between the administration plane and the control/data plane behaviors demanded a comprehensive validation of all capabilities for backwards compatibility, to prove the proper working of all services before moving any new software code to production. This was a non-negotiable requirement, in order to preserve the confidence of customers and operations that software changes would not be disruptive to their experience. But we soon found that traditional approaches to bench testing would not alone be enough to capture sufficient details regarding the individual health and wellness of each component independently, let alone to build a comprehensive picture of overall service reliability.

While we understood that minimizing any risk of disruption would necessitate comprehensive regression before every significant change, we also knew that the tedious manual testing exercises of the early development phase would not suit the needs of our customers. Our approach shifted towards development of a custom, reconfigurable testing platform, integrated with our continuous integration/continuous deployment (CI/CD) pipeline. This high-level framework is depicted in Figure 3. It resulted in reliably repeatable validation cycles, covering an ever-growing set of test cases across all components. This switch to automated testing added new development in the sense of coding test cases, but eventually test design and coding merged into the same practice. Overall, it cut our testing cycles by multiple orders of magnitude, allowing us the flexibility to increase velocity of deploying the latest code, resulting in faster improvement of reliability and features releases to our customers.

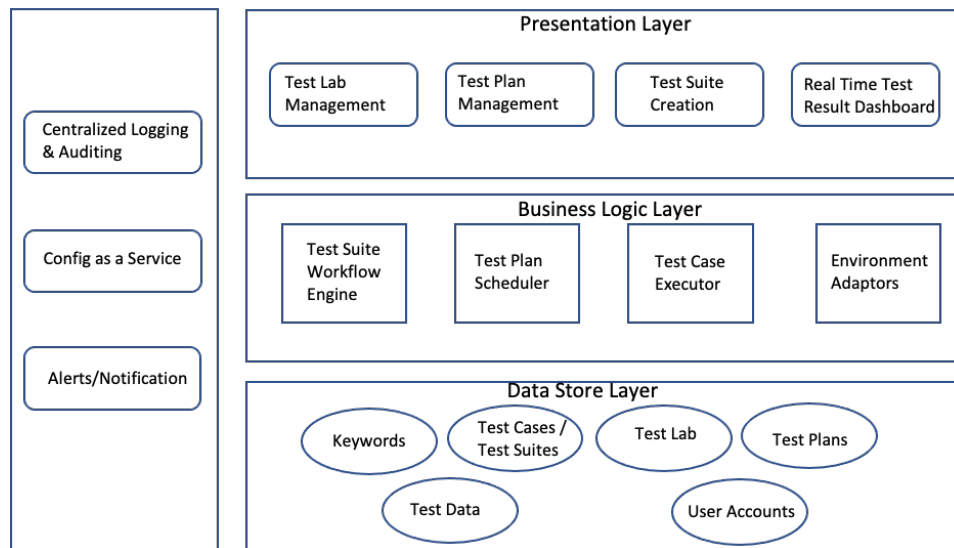


Figure 3 – Test Automation Framework

We also addressed several important challenges. For one, we could not always rely on a fixed set of software versions in production; that is, in the field, due to many different teams operating over time and different field requirements or business realities, there could be different versions of each component running, other than what got initially deployed. All these different variations would need to be supported uniformly from a feature perspective. Further, operations teams maintaining microservices-based systems, having numerous components with different behaviors distributed geographically, would face the tedious exercise of having to manage extreme amounts of detail during maintenance windows, requiring large teams of highly-skilled engineers maintaining superhuman focus for hours at a stretch, attending to every detail when performing upgrades.

The most crucial aspect of successfully managing these details turned out to be perhaps one of the most difficult to achieve in practice: configuration standardization, such that the configurations being tested and deployed have known behaviors that can be used as baselines when our teams are trying to resolve something that isn't behaving as expected in the wild. This is something done very well in software, but very difficult to achieve manually. These system complexities made clear that manual administration of even a small number of environments would be untenable over time. As has been discovered by other software-driven organizations, we resolved that an automation-first strategy was required to make even simple administration tenable.

Another early indicator that tipped the scales towards platform automation was the realization that there would be numerous instantiations of the fundamental datacenter software stacks which powered our service – so many, in fact, as to make manual administration of all those system instances impossible in practice. There could never be enough skilled engineers to manually log in and take care of all the many traditional Day 2 activities which invariably would arise when running complex software systems – password changes, template updates, patches, even disaster recovery. Neither could these systems reliably and repeatably be deployed, day after day, week after week, retaining the same level of quality with the 15th as with the first; nor could they reliably be restored after a disaster using a manual checklist alone.

We thus adopted a platform strategy for our systems lifecycle applications, a technique that is also widespread in the software industry. Starting with a common framework of basic services, such as data, API, GUI, communications, and logging, we ensured that all applications participating in this platform would also enjoy common performance and availability optimizations, such as blue-green deployment for live upgrades, and site-diversity for fault tolerance. On top of this framework was developed a portal, into which bits of functionality could be dropped. Initially just a wrapper for some crude management utilities, it has become the one-stop-shop for platform operations teams, who leverage automation at every step in the lifecycle of our production systems. The portal's extensibility enables it to be used not only for large milestones such as deployment, upgrades, or disaster recovery, but also to perform more routine tasks such as license management, password rotation, and security patches.

The most important benefits from the repeatability and reliability of this standardized approach to operations are clear and have proven value from the start; many serious issues that could typically have resulted from hand-crafted configurations, varying from environment to environment, have been completely avoided. Taken as a whole, the benefits due to the automation are irrefutable, with time-in-motion improvements typically measured in (sometimes multiple) orders of magnitude, as illustrated in Table 2. To paraphrase computing legend Larry Wall, it "makes hard tasks easy, and impossible tasks possible."

Table 2 – Platform Lifecycle Automation Benefits

Platform Lifecycle Category	Execution Timeline Improvement with Automation
Regression Testing	>99%
Production VM Build / Software Deployment	92%
Disaster Recovery / High Availability Testing	>70%

5. Data-Driven Proactive Monitoring

Modern operational visibility has expanded beyond sysadmins and ITOps analysts. It is required not only to monitor the status and performance of running applications/services/systems and to detect issues in real-time but also to understand why, project the trends, and provide feedback to DevOps teams and customers. Additionally, the nature of cloud technology and SD-WAN technology, namely the separation of data plane, control plane, management plane, virtual resource, and physical resources, adds more complexity to the platform monitoring and data analysis.

Following the industry standard, we embrace data-drive proactive monitoring approaches and cross-layer correlation to achieve observability at all layers. It is a straightforward architecture pattern which allows us great flexibility in adding or changing features and service. Figure 4 shows our high-level data-driven monitoring architecture design.

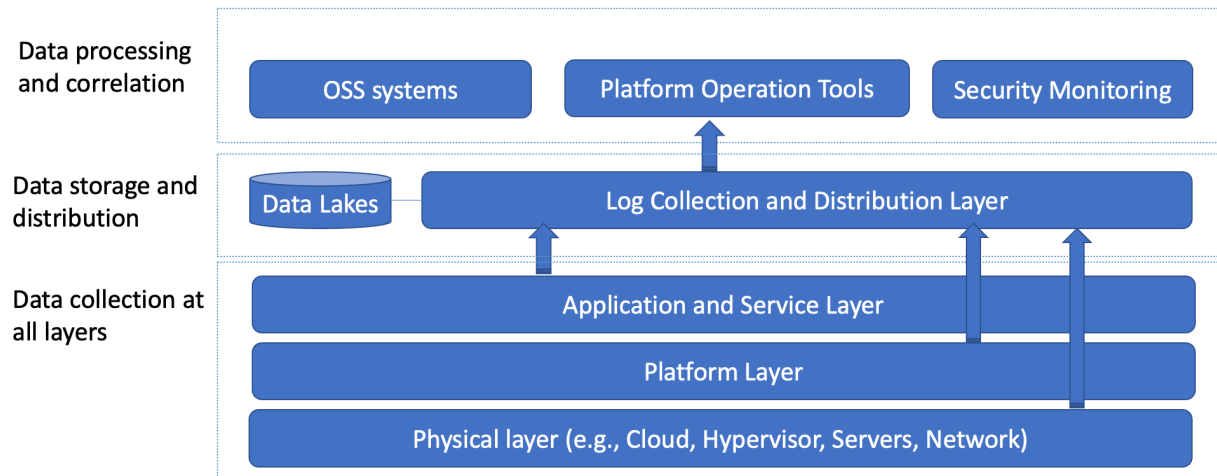


Figure 4 – SD-WAN Platform Data-Driven Monitoring

The data collection layer utilizes industry standard tools or vendor supported features to collect data and provide observability from all layers, including the physical layer (e.g., cloud, hypervisor, servers, hosts, network, etc.), the platform layer, and the application/service layer. The data collected includes all three pillars – logs, metrics, and traces – that are needed for observability.

The data storage and distribution layer uses industry standard technologies and shared platforms to store and distribute telemetry data to the upstream systems. The volume and velocity of the data needed for observability is huge. Thus, our design requirement on systems and platforms used at this layer mainly focuses on scalability, performance, and HA.

The data processing and correlation layer consists of multiple systems that are designed and developed to provide visibility from different perspectives. For example,

- Customer portal: provides the overall health status at the customer service level.
- Operation team tools and portal: provides in-depth health status of all layers from an engineering perspective.
- Security monitoring portal: provides in-depth telemetry data from a security perspective.

Our data-driven monitoring infrastructure has become a critical piece in the entire product ecosystem. It provides insightful information and feedback from many product perspectives, as in Figure 5 (a) below.

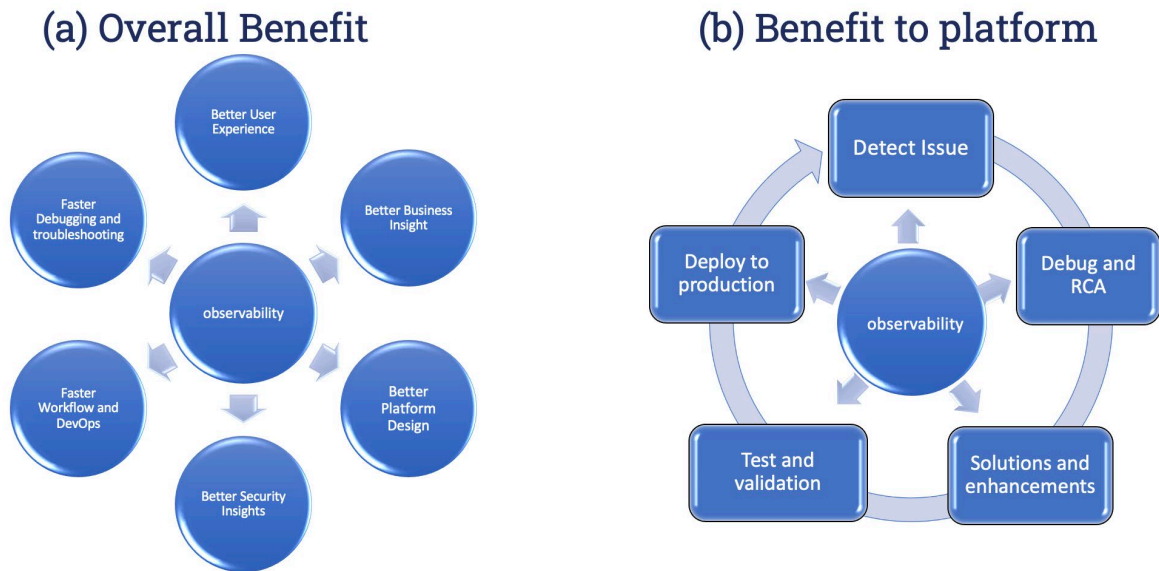


Figure 5 – Benefits of Observability

From a platform perspective the main benefits include but not limited to:

- Fast operation reaction to issues: we are collecting health metrics from all layers and triggering notifications to our operation teams to react. The correlation of collected data from all layers also assists in troubleshooting and debugging process, helping operation teams and platform architecture teams to identify the root causes and fix issues even before customers report them (as illustrated in Figure 5 (b)).
- Health metrics collected from the physical layer and the platform layer help us to look for signs that indicate resources may soon run out of capacity, enable us to predict the growth and trend, perform capacity planning, and trigger operation teams to scale out platform components. Figure 6 illustrates the microservices growth projection with increasing platform capacity. The calculation is based on the observability data collected in production. As shown in the diagram, with the increasing platform capacity, different microservices need to be scaled out differently depending on the projected workload. Some microservices do not require to be scaled out even when we plan to double the platform capacity.




Figure 6 – Microservices Growth vs Platform Capacity

In addition, using standard technologies in design and developing this data-driven proactive monitoring infrastructure helps to reduce development cost, to achieve required scalability and reliability, and to hire talent to maintain and operate the platforms.

6. Conclusions

It was the intention of this paper to detail the ways in which we have found that simple software industry best practices could be implemented to great effect as part of the operationalization of networking services. These include leveraging standard networking protocols for implementing core availability behaviors, standardizing configurations in order to apply an automation-first approach to change management, embracing test-driven development to validate changes as quickly as needed by the business, and employing insights collected using modern data management approaches to forecast growth and anticipate changes. Although common among many industries, these techniques differ from hardware-based approaches due to their inherent flexibility in relation to dynamic virtual and distributed software-based systems, allowing greater reliability and availability to be offered.

Abbreviations

API	application programming interface
BGP	Border Gateway Protocol
CI/CD	continuous integration/continuous deployment
CMDB	change management database
DC	data center
DR	disaster recovery

GRE	Generic Routing Encapsulation
GUI	graphical user interface
HA	high availability
MACD	move/add/change/delete
MSO	multiple system operator
SDN/NFV	software-defined networking/network function virtualization
SD-WAN	software-defined wide area network

Bibliography & References

Programming Perl, 2nd Edition (1996), Tom Christiansen, Randal L. Schwartz, and L. Wall; **ISBN-13:** 978-1565921498, **ISBN-10:** 1565921496; O'Reilly Media.

M. Casado, N. McKeown, and S. Shenker, "From ethane to SDN and beyond", in ACM SIGCOMM Computer Communication Review, vol. 49, issue 5, Oct. 2019, pp 92-95.

L. L. Peterson, C. Cascone, and B.S. Davie, "Software-Defined Networks: A System Approach"; System Approach, LLC.

"Network Functions Virtualisation – Introductory White Paper"; ETSI. Oct. 2012. Retrieved June 2013.
G. Fellows, "High-Performance Client/Server: A Guide to Building and Managing Robust Distributed Systems", in Internet Research, vol. 8, issue 5, Dec. 1998.

Observability Engineering, C. Majors, L. Fong-Jones, and G. Miranda ; ISBN: 9781492076445 ; Released May 2022 ; O'Reilly Media.

Observability in Google Cloud ; Google Cloud DevOps Research and Assessment (DORA) research.

N. Kumar, A. Leventer and A. Matatyaou, "Monitoring and Troubleshooting at Scale with Advanced Analytics", SCTE Cable-Tec Expo 2021.

A. Mohan and X. Huang, "Robust and Resilient Service Assurance System Design with Observability to Improve Enterprise Customer Experience", SCTE Cable-Tec Expo 2022.

Coherent PON Poised to Become Cable's Next Long Term Evolution Access Platform

A Technical Paper prepared for SCTE by

Zhensheng (Steve) Jia, Ph.D.

Distinguished Technologist and Director of Advanced Optical Technologies
CableLabs
858 Coal Creek Circle, Louisville CO, 80027
303.661.3364
s.jia@cablelabs.com

L. Alberto Campos, Ph.D.

Fellow
CableLabs
858 Coal Creek Circle, Louisville CO, 80027
303.661.3377
a.campos@cablelabs.com

Haipeng Zhang, Ph.D., CableLabs

Chris Stengrim, CableLabs

Curtis Knittle, Ph.D., CableLabs

Table of Contents

Title	Page Number
1. Introduction and Motivation	3
2. CPON Deployment Scenarios and Use Cases	6
3. CPON Key Technology Development	8
3.1. Optical Phase Modulation	9
3.2. Simplified Carrier Phase Recovery (CPR)	10
3.3. Preamble Design and Burst Signal Processing	13
4. Conclusion	15
Abbreviations	16
Bibliography & References	17

List of Figures

Title	Page Number
Figure 1 – Standardized PON Evolution	3
Figure 2 – Technolog Options in 100G PON	5
Figure 3 – Technology Comparison: Coherent Optics vs. IM-DD	6
Figure 4 – Split Ratio and Transmission Distances for 100G CPON link	7
Figure 5 – Extended CPON Use Cases	7
Figure 6 – Conventional IQ (a) and Simplified Optical Phase (b) Modulation Structures	9
Figure 7 – Signal in Different Stages of Receiver-side DSP Process. (a) optical signal phase after frequency-offset estimation; (b) estimated optical signal phase noise; (c) QPSK signal constellation without phase domain equalization; (d) corresponding phase of QPSK signal in (c) mapped back to PAM-4 signals; (e) equalized PAM-4 signals after phase domain post-equalization; (f) QPSK signal constellation with phase domain equalization	10
Figure 8 – Conventional CPR Process	11
Figure 9 – Simplified CPR Process	11
Figure 10 – Phase Estimation Results for Two Polarizations	12
Figure 11 – BER Performance Comparisons	12
Figure 12 – Preamble Design and Coherent Burst Signal Processing	13
Figure 13 – Experimental Verifications for Coherent Burst System	14

1. Introduction and Motivation

A passive optical network (PON) that has the benefits of point-to-multipoint passive topology for highly efficient fiber utilization has become the dominant optical access architecture for the operators in the current fiber to the premise (FTTP) deployments. The global PON equipment market is forecast to grow at a compound annual growth rate (CAGR) of 12.3% from 2020 to 2027, approaching \$16bn in 2027, from \$7bn in 2020. FTTP infrastructure builds are gaining momentum in most countries, with the global fiber access household penetration forecast exceeding 47% by 2026 [1], representing more than 1.1 billion subscriptions. Furthermore, 10G FTTP (and higher) is offered to residential subscribers by more than 55 communications service providers (CSPs). Ever-increasing bandwidth demand is the key driver for fiber access, along with demand for low-latency, reduced jitter, and improved quality of experience (QoE) by VR-based games and cloud applications [1].

In a typical PON system, an optical line terminal (OLT) at Hub or central office (CO) performs bidirectional communication with multiple optical network units (ONUs) at customer premises over a passive optical distribution network (ODN). A time-division-multiplexed PON (TDM PON) based on power-split ODN (PS-ODN) represents the most common PON architectures deployed so far. For a TDM PON, the OLT broadcasts the downstream (DS) signal to each ONU in continuous mode on a single wavelength channel. In the case of upstream (US) transmission, each ONU is assigned for a specific timeslot to transmit a burst of upstream data on another wavelength channel. The other type of PON architecture is wavelength-division-multiplexed (WDM) PON. WDM PON assigns dedicated wavelength for each ONU to realize virtual point-to-point optical connection. The corresponding ODN is wavelength-routed ODN (WR-ODN), which has an intrinsic wavelength routing capability through wavelength splitters. Time and wavelength division multiplexed (TWDM PON), in comparison, is the third type of multiplexing technique that combines TDM and WDM. Accordingly, the ODN is wavelength-selected ODN (WS-ODN), which relies on tunable optical filters to provide a wavelength selection capability in the ONUs. Among these different PON architectures, TDM PON offers much higher aggregated bandwidth for efficiently handling bursty traffic and with cost and complexity advantage, becoming the most popular optical access network. Unless otherwise stated, TDM PON is the focus of this paper.




Figure 1 – Standardized PON Evolution

The development of PON systems over the past 20 years has resulted in two families of standards: ITU-T Question 2/Study Group 15 (Q2/15) and IEEE 802.3 Ethernet working group, as shown in Figure 1. Guided by the telco-led Full Service Access Network (FSAN) Group, the ITU-T PON systems are defined primarily in ITU-T recommendations, which cover the system and architecture, physical medium dependent (PMD), transmission convergence (TC), and management layers. The IEEE PON systems are defined by a combination of PMD and data link protocols in IEEE 802.3 while the management/system layers are defined in IEEE 1904.

The G.983 asynchronous transfer mode PON (APON) and broadband PON (BPON) were the first standardized ITU-T PONs in late 1990s with a line rate of 155 Mb/s to 622.08 Mb/s symmetrically. The next PON was the most widely deployed gigabit PON (GPON), standardized as ITU-T G.984 in 2003. The GPON speed was extended to 2.5 Gb/s in the downstream. The maximal reach is 20 km, with a maximal split ratio of 1:64. The Ethernet PON (EPON) is the first type of PON standardized by the IEEE in 2004 as 802.3ah. Data are transformed by Ethernet frames with the line rate of 1.25 Gb/s symmetrically, maximal reach of 20 km, and a split ratio of 1:32. For 10G line rate per wavelength PON, the 10G EPON standard was released in 2009 as 802.3av with 10.3125 Gb/s symmetrical line rate. Then next-generation PON (XG PON) was standardized in 2010 as ITU-T G.987, and the speed was increased to 10 Gb/s downstream and 2.5 Gb/s upstream. The next-generation PON stage 2 (NG-PON2) was standardized in 2015 as ITU-T G.989 with TWDM architecture. The NG-PON2 supports 4–8 wavelengths with broadband speed of 10 Gb/s per wavelength and 40–80 Gb/s aggregated capacity over a single fiber. The next-generation symmetric PON (XGS PON) is a symmetric version of XG PON with 10 Gb/s line rate for both downstream and upstream.

Going beyond 10 G line rate per wavelength, the IEEE 802.3ca Nx25G EPON Task Force recently defined a system based on 25 Gb/s line rate [2]. The initial objective was to standardize a 100G-EPON by bonding four 25 Gb/s wavelength channels together and was scaled back to a 2×25 Gb/s system in November 2017 as neither the technological maturity nor the market needs were present. On the ITU-T side, in 2021, a suite of G.9804 Recommendations or G.hsp was developed in the ITU-T for a 50 Gb/s line rate PON system with a vision of laying the groundwork for future ITU-T PON systems. Common transmission convergence (ComTC) layer in G.9804.2 is defined in a line rate (with fundamental line rate of 12.4416 Gb/s and line rate factor of 1, 2, 4...N) to form the nominal line rate $12.4416 \times N$ agnostic of transmission rates, number of operating wavelength channels, and signal modulation and thus applicable to future TDM and TWDM PON systems [3].

For each generational PON development, operating over deployed ODNs and coexisting with legacy PON systems are two essential system requirements because ODNs are the biggest investment, and they have a lifetime much longer than technology cycle. These requirements can minimize the infrastructure cost and ease a smooth and non-disruptive migration towards new PON system. In comparison, the physical (PHY) layers in ITU-T and in 802.3 are largely equivalent because they both need to support similar operators' ODNs and leverage the common optoelectronic components and devices. The PHY layer generally handles the line rates, coding, forward error correction (FEC), wavelength plan, transceiver optical characteristics, ODN loss budget, reach, and coexisting systems.

All these standardized PONs employ simple intensity-modulation and direct-detection, IM-DD technology with non-return-to-zero (NRZ) format. Standardization of early generation PON was fairly straightforward. Between each adjacent generation, the main considerations were about increasing the data rate, higher launched power, or better FEC to fulfill a 29-dB or higher loss budget, which is the basic requirement for coexistence with legacy PON or reuse the deployed ODN. As shown in Figure 2, 25G EPON adopted low-density parity-check (LDPC) coding at the bit error rate (BER) threshold of 10^{-2} instead of keeping (Reed-Solomon) RS coding of XGS-PON at 10^{-3} . It is expected that 50G PON will

further improve the coding gain by using soft-decision LDPC. It is envisioned that 100 Gb/s per wavelength and beyond will be required in the future to meet continuously growing bandwidth demands. To reach 100 Gbps, IM-DD OLT and ONU transceivers are required to be more advanced, and a much higher degree of complexity compared to previous generational PONs. analog to digital converter (ADC) and digital to analog converter (DAC), and digital signal processing (DSP) unit are required to mitigate or equalize bandwidth limitation penalty, device linearity, and transmission impairments. A high launch power from a range of 8 dBm to 11 dBm is needed potentially with optical amplifiers on both OLT and ONUs for IM-DD 100G PON to stay within the required power budget. Achieving such a high launch power at an ONU side is still too challenging in the upstream. In addition, fiber dispersion causes significant penalties for high-baud-rate signals. Coexistence with traditional PON services makes it challenging to find a suitable transmission wavelength window in the O-band. Therefore, for increasing the data rate for a single wavelength, limited sensitivity will become too challenging for the 100G TDM-PON to meet the required power budget by using direct detection in the O-band, not to mention extended power budgets to cover the long-reach or high-density applications.




Figure 2 – Technolog Options in 100G PON

Coherent optics, on the other hand, more easily reaches 100Gbps because it divides the speed into 4 lanes, two polarizations each with amplitude and phase makes technology operate at 1/4 of the line rate [4]. Coherent optics, a game-changer optical fiber transport technology, has completely transformed optical transmission systems and enabled a widespread upgrade and new deployment of DWDM networks to speeds of 100 Gbps, 200 Gbps, and 400 Gbps per wavelength. Over the past decade, coherent optics has moved beyond its long-haul application origins to metro networks and now it has been introduced into access networks. Employing coherent optics in a PON offers many advantages. Coherent PON (CPON) is like traditional PON with point-to-multipoint topology over passive ODN. However, CPON uses multi-dimension modulation and coherent detection to provide longer reach and higher split ratio with improved optical power budget.

Compared to alternative direct-detect solutions that just use amplitude to represent the signal, only one dimension, coherent optical solutions (shown in Figure 3) use a high-power local laser source as a reference to achieve linear conversion of the optical field instead of optical power used in direct detection. This enables modulation and detection using four independent degrees of freedom, including amplitude and phase in two polarizations. With a local oscillator, significant coherent gain is provided along with wavelength selection without the need of an optical filter. Additionally, power fading induced by chromatic dispersion in direct-detection system is no longer an issue because of optical field recovery in

coherent detection. All these features coherent optics bring to optical transport systems, enable greater modulation efficiency and receiver sensitivity. Coherent detection for short-haul networks enables a superior receiver sensitivity that allows for extended power budget. Its high spectral efficiency enables dense WDM (DWDM) and lead to higher capacity channels and fewer optical ports providing operational simplicity that may lead to overall network savings. Moreover, the multi-dimensional signal recovered by coherent detection provides additional benefits to compensate linear transmission impairments such as chromatic dispersion (CD) and polarization mode dispersion (PMD) and its most efficient use of fiber spectral resources results in more optical spectrum available for future use and enables future network upgrades using multi-level advanced modulation formats [5].

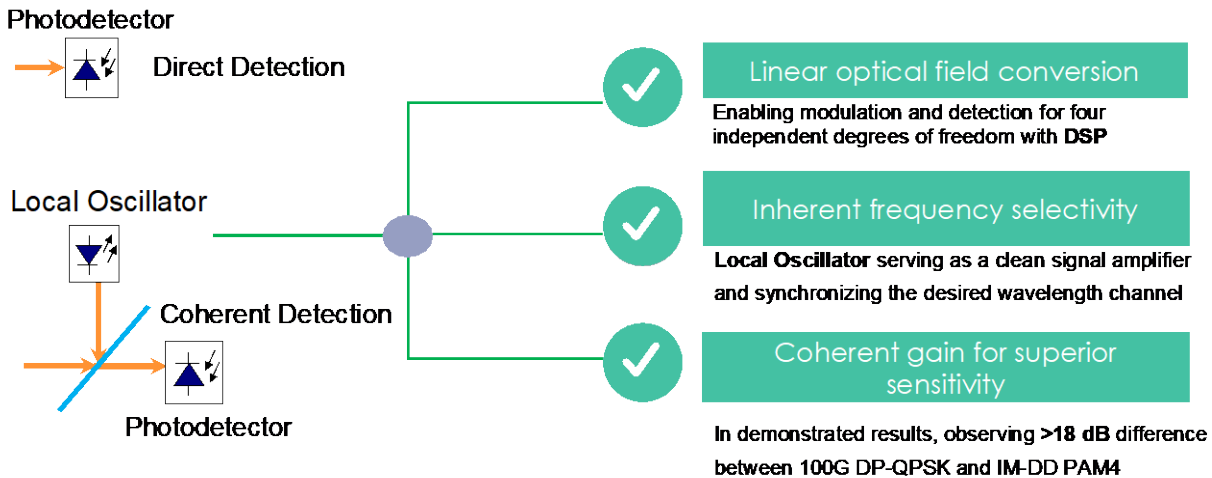


Figure 3 – Technology Comparison: Coherent Optics vs. IM-DD

2. CPON Deployment Scenarios and Use Cases

Leveraging its high sensitivity and powerful digital equalization of fiber transmission impairments, CPON offers unprecedentedly high degree of flexibility in terms of transmission distance and split ratio. Figure 4 shows measured 100G CPON split versus transmission distance topology options along with derived 100G target topology options after adding commercial implementation margin. CPON enables the service provider with significant flexibility on the network topology and reach [6]. Depending on different deployment scenarios, one can choose either very high split ratio at short reach, i.e., 512 split at 20 km transmission distance in Case 1, or low split ratio at longer reach, i.e., 16 split at 80 km transmission distance in Case 2 for rural applications. A hybrid mode in Case 3 that combining various split ratio at different transmission distances can also be achieved, i.e., a CPON port is first tapped at 20 km with a 90/10 passive splitter to support 64 splits, then tapped at 40 km with a 75/25 passive splitter to support 32 splits, finally at 80 km it supports another 8 splits. Case 3 represents the distributed CPON architecture to optimize optical power delivery over passive ODNs. The ample link margin enables flexible deployment of distributed CPON architectures where optical couplers with optimized coupling ratios can be used in a distributed fashion as the need for CPON connectivity is geographically dispersed. Each coupler diverts a fixed or adjustable portion of the CPON signal to a local splitter with the number of ports suitable for the demand of local CPON ONUs. As traffic and penetration evolves, the coupling ratios need to be adjusted. Incorporating adjustable and remotely controlled couplers enable automation and streamline operations.

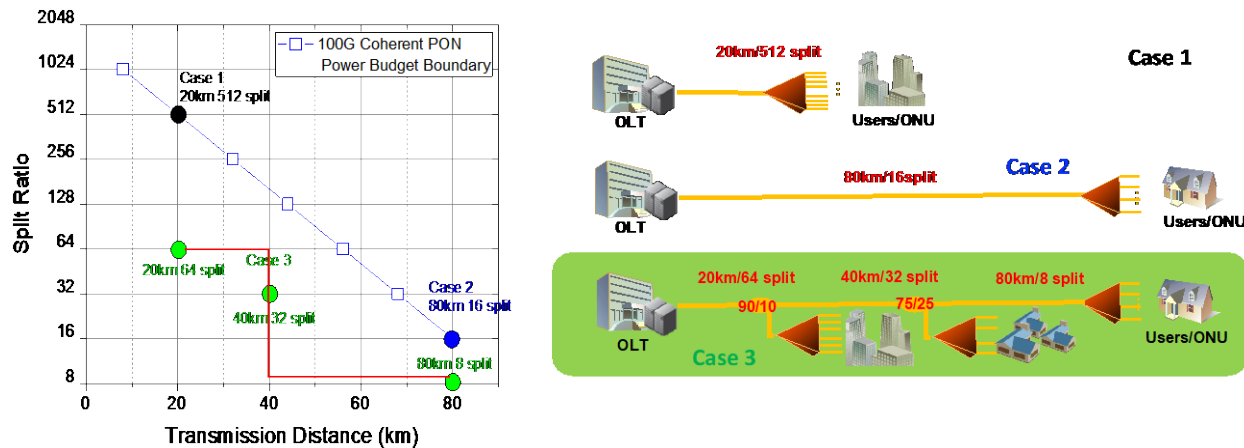


Figure 4 – Split Ratio and Transmission Distances for 100G CPON link

With the increase of CPON ability in higher capacity, longer reach, and higher split ratio, its applications are now extended to support many more optical connections. These include the new radio access network for the high-capacity x-haul transport of 5G and beyond 5G mobile networks, enterprise and business, access aggregation architecture, and even to data centers' aggregated connections. All these present new opportunities as well as challenges of a single platform for convergence of broadband wired and wireless access networks.

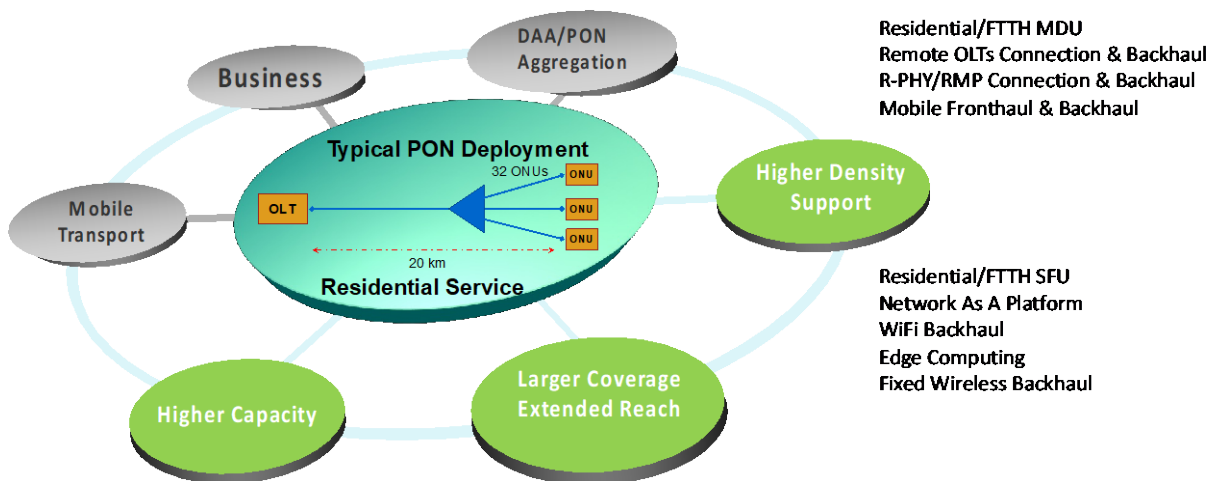


Figure 5 – Extended CPON Use Cases

CPON supports the use case of fiber-to-the-home (FTTH) connectivity. This includes direct fiber connection to either multiple dwelling units (MDUs) or single-family units (SFUs), as shown in Figure 5. CPON offers high density connectivity in densely populated areas, or long link distance connectivity in rural areas and other sparse/lower-penetration environments. In greenfield FTTH deployments without existing services, CPON deployment would ease operations in future years. In brownfield FTTH deployments with existing IM-DD PON, a CPON network could be overlayed using Coexistence modules (CEX), allowing the operator to gracefully migrate certain customers to increase revenue opportunities or lower congestion for other customers on the ODN. Network operators have wireline networks extending to multiple network edge devices that connect to end points such as residences, enterprises, or wireless

access points. Links connecting edge devices consolidate at an aggregation node and are transmitted on a single optical link to a Hub/CO. CPON can be used to support back-haul connectivity to the aggregation node. The aggregation node, for example, can contain multiple remote OLT (R-OLT) devices, which moves the OLT out of the Hub/CO to the aggregation node location that is located close to ONUs. This Remote OLT case could be an important bridge for applications where the distance exceeds IM-DD limits, but where the capacity does not yet exceed the offerings of an IM-DD PON service. The aggregation node can also contain Remote PHY devices (RPDs) and Remote MAC-PHY devices (RMDs), to support distributed Converged Cable Access Platform (CCAP) architectures which offer operators great flexibility and low-cost deployment benefits.

Another use case for CPON is to provide mobile x-haul services. For example, CPON can be utilized to provide base station connectivity and carry back-haul traffic in 5G RAN. CPON can also carry traffic from the mid-haul/front-haul segment of 5G RAN. In this use case, beside meeting peak capacity requirement, CPON will be designed with reduced system latency leveraging low latency MAC mechanism for example, to meet the low latency requirements of 5G RAN.

In addition to FTTH, aggregation connectivity and mobile x-haul, CPON also supports a wide range of existing or emerging applications. For example, CPON can provide access point connectivity to carry Wi-Fi traffic or support fixed wireless back-haul and provide radio connectivity to carry fixed wireless traffic. CPON can serve to connect the users to edge computing nodes/devices located at the Hub/CO sites or provide connectivity from Hub/CO sites to edge computing platforms such as edge nodes near end users/subscribers. Furthermore, a point-to-multipoint local network infrastructure such as Passive Optical Local Area Network (POLAN) can leverage CPON to deliver data to multiple end users within campuses or buildings. CPON can also be adopted in emerging applications such as IoT-based Smart City, or Industry 4.0 Smart Factory, providing connectivity to many sensors and smart devices.

3. CPON Key Technology Development

Despite the numerous advantages offered using coherent optics in next-generation PON systems, major engineering challenges associated with CPONs for access applications remain. The coherent technology in a long-haul optical system utilizes best-in-class discrete photonic and electronic components, such as state-of-the-art digital-to-analog converters (DACs)/analog-to-digital converters (ADCs) and DSP ASIC based on the most recent CMOS process. These solutions are over-engineered, too expensive, too big, or too power-hungry for the access PON. Cost is the first challenge that must be resolved to introduce coherent optics into access networks. A PON is eventually directly connected to end users and is therefore highly sensitive to the cost of the optical and electrical components used in the network because of the enormous market involved. Therefore, the lasers, modulators, coherent receivers, and DSPs in an access network must be optimized to lower the complexity, cost, and power consumption. Meantime, it is worth noting that the total cost of ownership (TCO) should be considered here instead of single device or equipment. Coherent optics fiber spectrum reclamation not only alleviates the pressure to retrench new fiber, but also the low transmit power of coherent optics leads to the ability in placing a larger number of optical carriers on a single fiber strand. Looking at this from an overall system cost perspective, the usage optimization of fiber resources leads to a lowest cost per bit transport [7].

The cost of coherent detection can be broken down into several parts: the laser source, transmitter, receiver, and DSP. The first three parts are mainly optical and electrical components that can be simplified in many ways. In long-haul transmission, laser sources on both the transmitter and receiver sides are generally expensive, high-quality, narrow-linewidth (typically less than 300 kHz) external cavity

lasers (ECLs). Basically, two laser sources are needed in one transceiver, including one signal carrier and one LO. One solution for cost reduction is to use fewer lasers in the ONU. Using a remote laser source from the centralized OLT for each ONU would reduce the overall cost by distributing the laser cost over the whole network. Lower cost distributed feedback (DFB)/distributed Bragg reflector (DBR) lasers, with linewidths of higher than 1 MHz, can also be introduced in coherent systems with minimal impact in performance for access applications [8-9].

Another major simplification is related to the transceiver, especially on the receiver side. A full field optical coherent receiver is typically used in long-haul transmission systems and consists of four pairs of balanced PDs (8 PDs) with polarization and phase-diversity homodyne coherent detection. Consequently, four ADCs are required for signal processing. This setup could be simplified by using fewer PDs. The use of heterodyne coherent detection can halve the number of required balanced PDs, ADCs, and associated TIAs. The system could be further simplified by removing the polarization-diversity receiving system and using polarization scrambling or a special polarization-time block coding method [7].

Therefore, coherent systems traditionally have come at a greater cost, but through access specific re-design, significant simplification and reduction in power consumption can be achieved. In this section, a few potential key technologies developed for CPON application are discussed.

3.1. Optical Phase Modulation

As modulators with low V_{pi} are developed, a simpler modulator design becomes an attractive alternative to traditional IQ modulators. An optical QAM coherent modulator could be implemented with a single-phase-modulator by directly modulating the laser to vary the amplitude, thus trading modulator driver signal complexity with a lower optical loss present in nested IQ modulators. Figure 6 compares a traditional nested IQ modulator structure with a simplified modulator using single-phase-modulator per polarization.

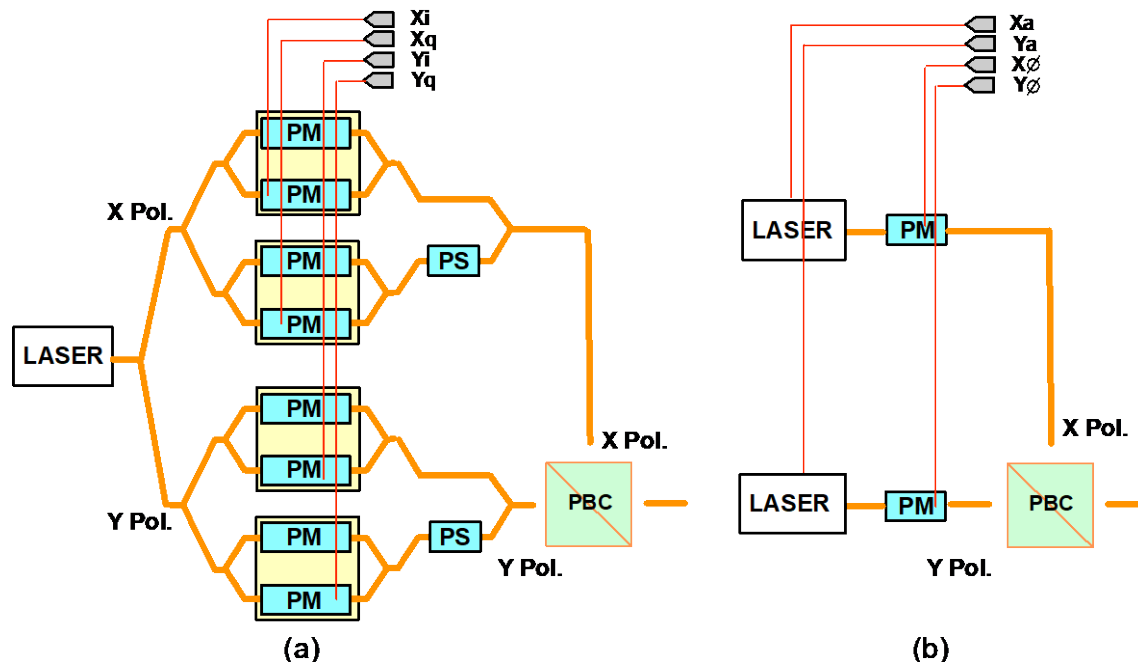


Figure 6 – Conventional IQ (a) and Simplified Optical Phase (b) Modulation Structures

This modulator alternative is simplified for a quadrature phase shift keying (QPSK) signal as no amplitude modulation drivers (X_a , Y_a) are needed. We verify the results by testing the performance of 12.5 GBaud polarization multiplexed PAM-4 signal, which is equivalent to 12.5 GBaud polarization multiplexed QPSK after phase domain modulation. Phase domain multilevel signal modulation and post-equalization is verified here with coherent detection. The results of a signal in different stages of receiver-side DSP process are shown in Figure 7. Clearly, obvious phase domain inter-symbol-phase interference (ISPI) exists, with large phase fluctuations on four QPSK phases. By comparing Figure 7 (d) and (e) or comparing Figure 7 (c) and (f), we can see significant improvements.

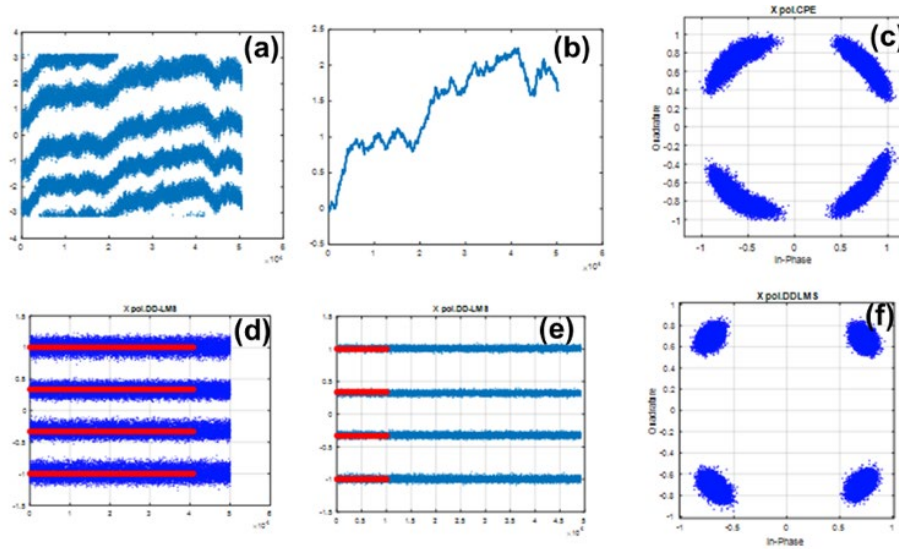


Figure 7 – Signal in Different Stages of Receiver-side DSP Process. (a) optical signal phase after frequency-offset estimation; (b) estimated optical signal phase noise; (c) QPSK signal constellation without phase domain equalization; (d) corresponding phase of QPSK signal in (c) mapped back to PAM-4 signals; (e) equalized PAM-4 signals after phase domain post-equalization; (f) QPSK signal constellation with phase domain equalization.

3.2. Simplified Carrier Phase Recovery (CPR)

In conventional CPR process, the independent phase estimation is performed for both polarization signals respectively. L-tap symbols are used for the center Symbol $S_{n+L/2}$ phase estimation based on 4th power Viterbi-Viterbi (VV) CPR or Blind phase search (BPS) algorithm as shown in Figure 8, where S_n is the n^{th} received symbol. As an example, for a QPSK signal with four phase states, the received complex symbols are first raised to the 4th power to remove modulation and make sure that only the phase noise is present. The $S_{n+L/2}$ is then added to N predecessors and successors to average the estimated phase. Because the phase varies over the range of 2π , the estimated phase must, therefore, be “unwrapped” to provide a continuous and unambiguous estimation of phase. After the phase “unwrapping”, the compensation of an estimated phase error is performed with respect to the received complex symbols.

In the simplified CPR process for the hardware-efficient DSP flow, two steps are proposed. phase noise estimation is firstly performed at only single polarization direction and then shared with the second polarization signals as shown in Figure 9. Data-aided or blind estimation methods is then used for the fixed phase rotation estimation & recovery of the second polarization signal. In this way, only one dynamic phase noise estimation is performed, which is time varying with high computation complexity.

Fixed phase rotation estimation is a one-time process and the computation complexity is negligible when compared to dynamic phase noise estimation.

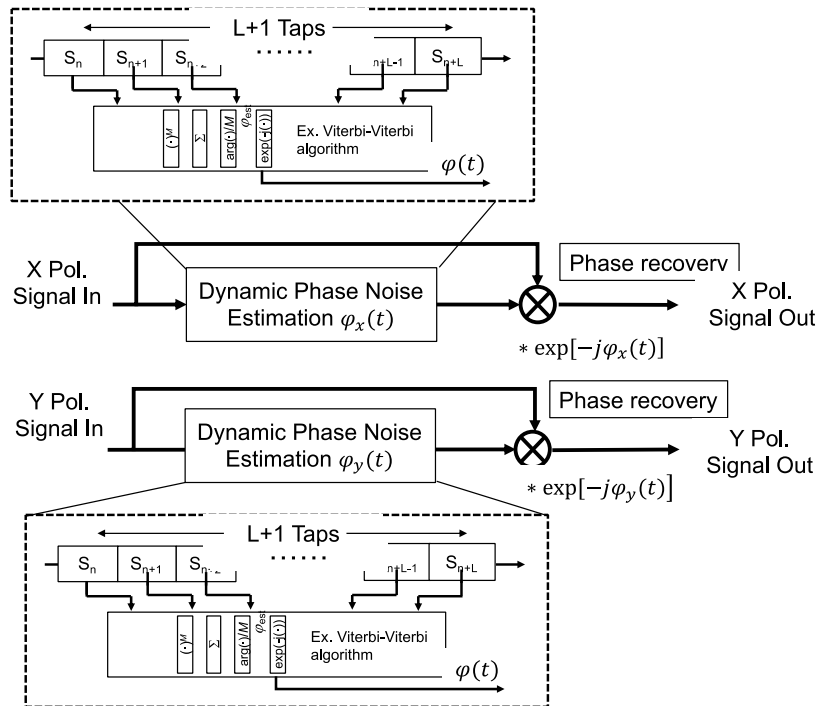


Figure 8 – Conventional CPR Process

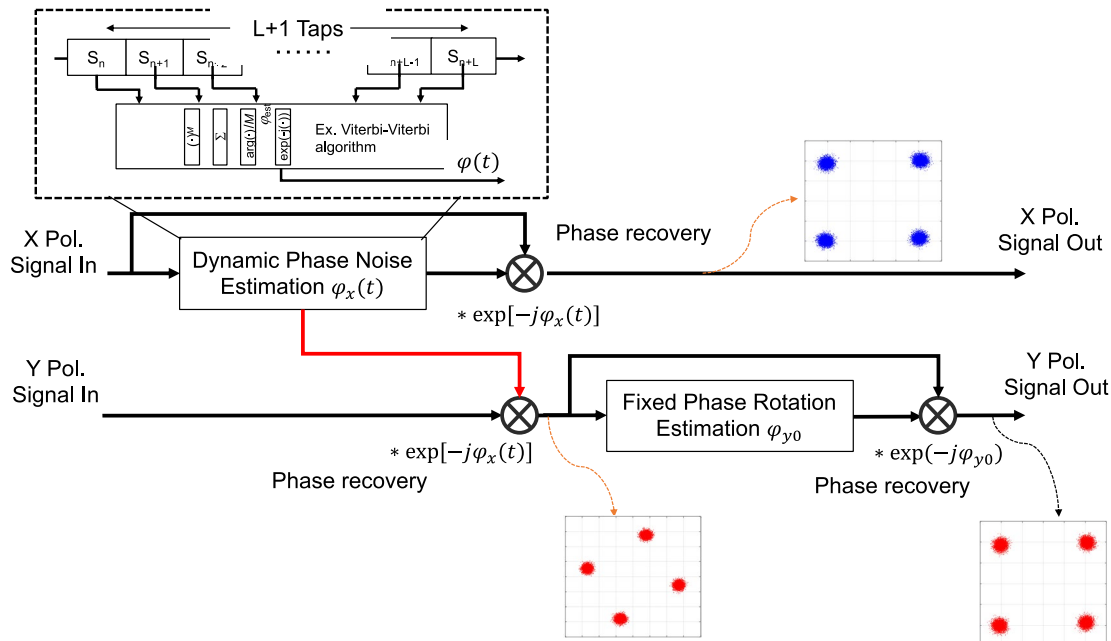


Figure 9 – Simplified CPR Process

Figure 10 shows the conventional and simplified CPR results. In the conventional algorithm (a), the independently estimated dynamic phase noise for X and Y polarization has the same phase evolution with a fixed phase offset. Since the independent phase noise from fiber nonlinearity is rather small at access distance, the phase noise in the two polarizations shows the same behavior except the fixed phase rotation. Figure 10 (b) shows the effectiveness of the simplified CRP process in the simplified DSP flow with the use of one polarization phase estimation result and fixed phase rotation for the other polarization.




Figure 10 – Phase Estimation Results for Two Polarizations

Figure 11 (a) shows BER performance with different fixed phase rotation estimation methods. It is clearly seen that training sequence (TS) and blind estimation have very similar performance in the condition of fixed receiver power at -38.3 dBm. It also shows that 64 symbols of TS or average window size shows the converged results for fixed phase rotation estimation. Figure 11(b) demonstrates the optical power sensitivity comparisons with conventional and proposed CPR methods for both QPSK and 16QAM signals. The results show the simplified method can apply for different modulation formats and no obvious performance degradation is observed by using fixed phase rotation estimation compared to conventional methods.

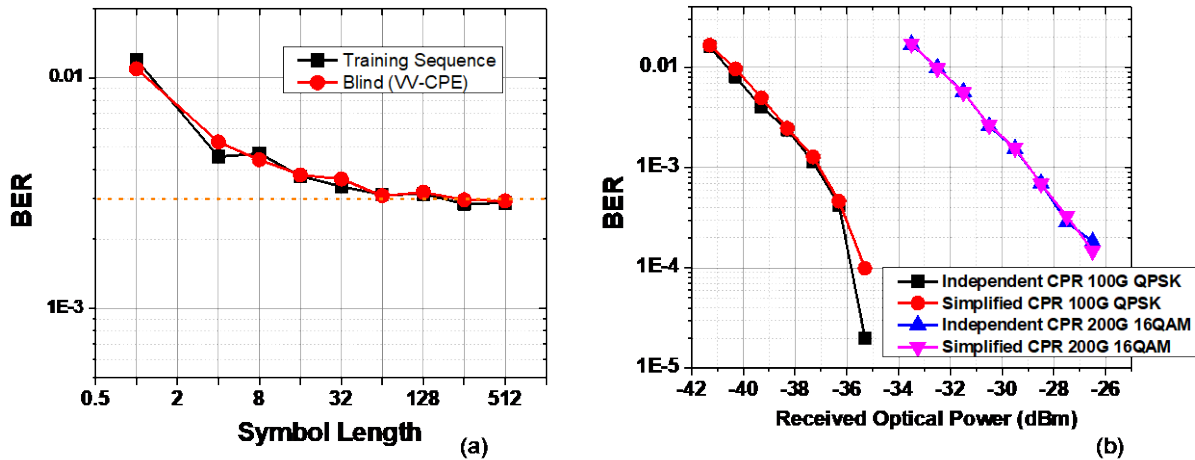


Figure 11 – BER Performance Comparisons

3.3. Preamble Design and Burst Signal Processing

One of the key challenges for CPON is the realization of upstream coherent burst-mode detection when coherent optics is brought to point-to-multipoint connection with a TDM fashion. Burst-mode coherent detection generally includes burst-mode linear amplification and burst mode DSP. Burst-mode amplification is required to deal with different signal power levels and enlarge the dynamic range. As considerable challenges are still associated with coherent receivers with burst-mode linear TIAs, optical burst-mode preamplification is an effective method of leveling burst signals before these signals are sent to coherent receivers via electrical, optical, or power control of local oscillator. Meanwhile, scheduler knowledge of which ONU is transmitting can also allow receiver to anticipate power level adjustment needed and ranging process can also be used to have the ONU pre-amplify Tx-burst so that it arrives at the OLT receiver at the desired level to alleviate some of the burden on the receiver.

Burst-mode DSP and an efficient preamble design at the receiver side are needed in addition to optical power control. The acquisition time for signal recovery at the coherent burst-mode receiver must be sufficiently short to improve the upstream efficiency. Specially designed preambles and fast DSPs have been reported. In pilot-aided time-domain estimation, a low-complexity widely linear compensation method is used to compensate for the IQ imbalance and burst-mode DSP based on precalculated finite impulse response (FIR) filter coefficients have also been demonstrated to shorten the acquisition time.

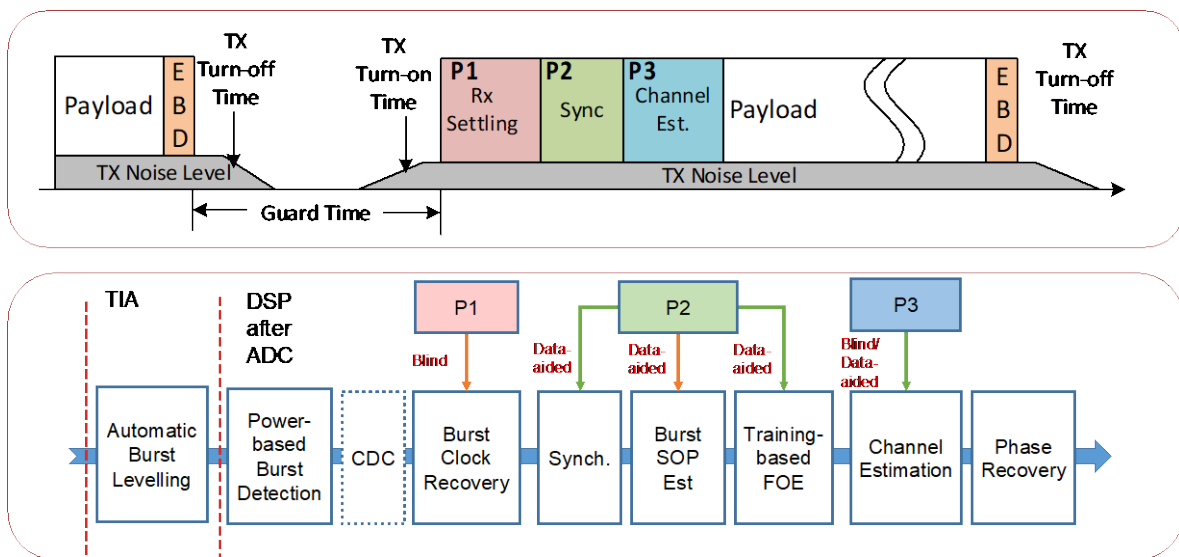


Figure 12 – Preamble Design and Coherent Burst Signal Processing

Figure 12 is a schematic showing the principle of the burst frame structure and upstream burst-mode signal recovery functions for a TDM CPON [10] with different synchronization patterns. Although the general concept is taken from the existing PON standards, most of the preamble patterns and related subfunctions are specially designed for dual polarization coherent bursts. Here, we consider the most complex case in which the optical signals are modulated and multiplexed on the phase, polarization, and amplitude. The proposed preamble consists of three sync patterns, P1, P2, and P3. Based on nearly equally distributed QPSK symbols, P1 is designed for receiver settling with burst-mode amplifications. P1 is also designed for burst clock recovery. P2 is designed for data-aided DSP synchronization functions, including frame synchronization, synchronization, and frequency synchronization (frequency offset estimation). In the design, 2 x N conjugate symmetric symbols are used to demonstrate the feasibility of efficient preamble design and signal processing, where three DSP functions share the same preamble to

reduce the overall preamble length. Based on low-order training symbols (e.g., QPSK symbols), P4 is designed for adaptive channel equalization. Based on the information of the state of polarizations (SOP) estimated by P2, the inverse of the Jones matrix can be utilized to reduce the convergence time of adaptive channel equalization. Finally, the payload process can be simplified using the information from the proposed preamble. High-order modulation formats can be used in the payload block. Finally, the order of these preamble patterns and burst mode DSP functions may differ from that shown in Figure 12, depending on the algorithms used.

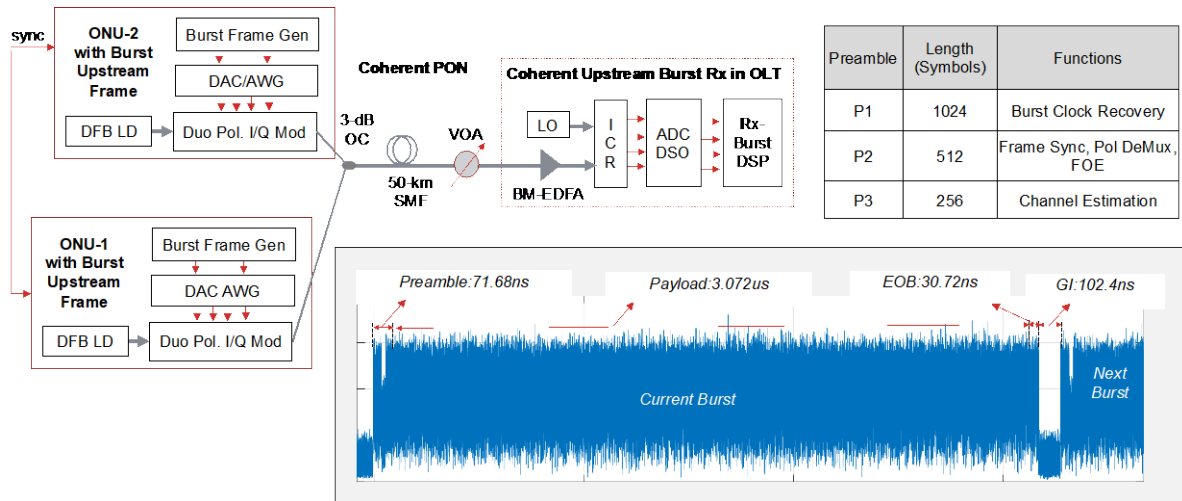


Figure 13 – Experimental Verifications for Coherent Burst System

Based on above preamble design and burst-mode DSP, we set up the experimental demonstration of coherent upstream burst detection in 100-Gb/s/λ TDM CPON, as shown in Figure 13. To evaluate the performance of upstream burst signal detection, two synchronized ONUs are running separately. At the ONU-side, the burst frames of 25-GBaud PDM-QPSK with the proposed preamble and structure as described in Figure 12, are generated by the 80-GSa/s arbitrary waveform generators (AWGs). Then the generated burst frames are fed into the dual-polarization I/Q modulators with four drivers for optical signal modulation. Here, we use a tunable DFB laser at 1550-nm wavelength with a linewidth of ~1-MHz as the laser source in each ONU. After modulation, the burst signals from ONU-1 are combined with a dummy signal from ONU-2 by a 3-dB optical coupler (OC). To avoid collision, the burst frames from two ONUs with same setup are staggered. Through the automatic bias-control and synchronization between two AWGs, the burst signal from one ONU is only coupled with the null signal from the other ONU.

The combined burst signals are then transmitted over 50-km single-mode fiber. The received optical power is controlled by a variable optical attenuator (VOA) for BER test. At the OLT-side, a burst-mode EDFA is used for signal pre-amplification. The pre-amplified signal is mixed with local oscillator (LO) in an integrated coherent receiver (ICR) for coherent detection. A tunable ECL at 1550m-nm is used as LO in OLT, and its linewidth is < 100-kHz. After coherent detection, the received signals are sampled by an 80-GSa/s scope and then processed via the offline burst-mode DSPs described in Figure 12. The detected power waveform of burst frame is shown in Figure 13 with the detailed configuration, and the summary of preamble unit length and functions is that P1, P2 and P3 have 1024, 512 and 256 symbols, respectively. Therefore, each burst frame contains a total preamble length of 71.68 ns (1792 symbols), a payload of 3.072μs, an end of burst (EOB) of 30.72 ns. A guard interval time of 102.4 ns is used to separate bursts. It is worth mentioning that the coherent preamble obtained is lower or comparable with preambles used by the different flavors of IM-DD PON, meaning that efficient preambles were achieved despite the dual polarization and IQ constellation complexity characterizing coherent transport.

4. Conclusion

This paper presents ITU-T and IEEE PON evolutions, the progress and limits of current intensity modulation and IM-DD PON technologies, and the major drivers and use cases for longer reach and higher split CPON. The high-level architectural requirements are discussed, such as the distributive optical taps for efficient power delivery, coexistence of generational PONs. This paper also provides the technology options to simplify the coherent OLT and ONU for cost reduction, and future PON upgrade strategies for different cable operators.

CableLabs[®] announced the launch of the CPON project in May 2021 toward the goal of future proofing cable's access architecture. The objective of the project is to initially develop specifications for 100G passive optical networks and devices that are multi-vendor interoperable, that coexist with existing infrastructure, and that can cost effectively be deployed at scale.

As reviewed in this paper, CPON demonstrates significant performance improvement, it can offer 10 times capacity compared to 10G PON and with the increase of 16 times in split ratio or 4 times in transmission distance, from 20km to 80km. Leveraging such capabilities and CPON synergies with existing HFC architectures, will allow CPON to significantly expand its capabilities beyond traditional residential deployment to support convergence needs at the network edge, from DAA aggregation, mobile x-haul, optical LAN, all the way to future fiber to the MDU and home.

It is worth mentioning the significant advantages that CPON brings beyond speed. Transitioning to coherent transport from IM-DD or analog optics results in a more efficient use of the optical wavelength spectrum. One can use one CPON channel instead of using ten 10 Gbps PON channels, which allow the operator to reclaim optical fiber spectrum just like they did years ago in the coaxial domain when transitioning from analog to digital video. Port reduction at the hub location is another key advantage when transitioning from 10 Gbps systems to 100 Gbps system as it reduces real estate in the router/switch and simplifies operation. The reach that CPON enables leads to pervasive all passive networks, forgoing the need in placing cabinets with repeaters to when link distances exceed 20 km and the operational overhead associated in maintaining and powering repeaters. The split and reach combinations possible with CPON results in great flexibility and homogeneity in the use of technology, meaning one technology that can be used for all services. CPON is a technology that is in a nascent stage with a lot of room to grow.

It is believed that CPON technologies can carry the cable industry for 20, 30, or 40 years into the future, just like DOCSIS has done since 1995!

Abbreviations

AN	aggregation node
b/s	bit per second
CEx	coexistence element
CO	coherent optics
CPR	carrier phase recovery
DSP	digital signal processing
FEC	forward error correction
HFC	hybrid fiber coaxial
IM-DD	intensity modulation and direct detection
IQ	in-phase and quadrature
MUX	multiplexer/demultiplexer
ODN	optical distribution network
OLT	optical line terminal
ONU	optical network unit
PHY	physical layer
PM	phase modulation
PON	passive optical network
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying

Bibliography & References

- [1] Omdia, *PON Vendor Landscape Report and Fiber and Copper Access Equipment Forecast: 2021–27 – 2022*.
- [2] IEEE 802.3ca, *Physical Layer Specifications and Management Parameters for 25 Gb/s and 50 Gb/s Passive Optical Networks*, 2020.
- [3] ITU-T G.9804.2, *Higher Speed Passive Optical Networks: Common Transmission Convergence Layer Specification*, Sept. 2021.
- [4] Z. Jia and L. A. Campos, “*Coherent Optics for Access Networks*”, CRC Press, Nov. 1, 2019. ISBN 9780367245764.
- [5] Z. Jia and L. A. Campos, “*Coherent Optics Ready for Prime Time in Short-Haul Networks*,” in IEEE Network, vol. 35, no. 2, pp. 8-14, March/April 2021, doi: 10.1109/MNET.011.2000612.
- [6] L. A. Campos, Z. Jia, M. Xu and H. Zhang, “*Coherent Optics for Access from P2P to P2MP*,” 2022 Optical Fiber Communications Conference and Exhibition (OFC), 2022, paper Th3E.1.
- [7] J. Zhang and Z. Jia, “*Coherent Passive Optical Networks for 100G/λ-and-Beyond Fiber Access: Recent Progress and Outlook*,” in IEEE Network, vol. 36, no. 2, pp. 116-123, March/April 2022, doi: 10.1109/MNET.005.2100604.
- [8] H. Zhang, M. Xu, J. Zhang, Z. Jia, L. A. Campos and C. Knittle, “Highly Efficient Full-Duplex Coherent Optical System Enabled by Combined Use of Optical Injection Locking and Frequency Comb,” in Journal of Lightwave Technology, vol. 39, no. 5, pp. 1271-1277, 1 March 2021, doi: 10.1109/JLT.2020.2998438.)
- [9] Z. Jia, L.A. Campos, M. Xu, H. Zhang, J. Zhang, C. Stengrim and C. Knittle, “Ultra Low-Cost Injection-locked FP Laser Source for Coherent Access Networks”, SCTE Cable-Tec Expo 2019.
- [10] J. Zhang, Z. Jia, M. Xu, H. Zhang, L.A. Campos and C. Knittle “*High Performance Preamble Design and Upstream Burst Mode Detection in 100-Gb/s/λ TDM Coherent PON*”, in Optical Fiber Communication Conference (2020), paper W1.E.1.

Collision-Free Hyper-Speeds on the Bi-Directional FDX Highway

A Technical Paper prepared for SCTE by

Dr. Robert Howald

Fellow

Comcast

+1 (267) 398-8104

robert_howald@comcast.com

John Ulm

Engineering Fellow

CommScope

+1 (978) 609 6028

John.ulm@commscope.com

Saif Rahman, Comcast

Dr. Zoran Maricevic, CommScope

Table of Contents

Title	Page Number
1. Introduction.....	4
2. DOCSIS 4.0 Full Duplex Overview	4
2.1. Key FDX Innovations	5
2.1.1. A Closer Look: Echo Cancellation	6
2.1.2. A Closer Look: Interference Groups and Transmission Groups	7
2.2. FDX-Capable Amplifiers.....	7
3. FDX Traffic Engineering: An Operational Perspective.....	9
4. Network Capacity Planning	10
4.1. The “Basic” Traffic Engineering Formula	11
4.2. The “Modified” Traffic Engineering Formula	12
4.3. Determining Service Group QoE based on Probabilities.....	14
4.3.1. Individual Subscriber Bandwidth Probabilities	14
4.3.2. Service Group Subscriber Bandwidth Probabilities	16
4.4. FDX Traffic Eng for Large SG with Overlapping Upstream + Downstream.....	17
5. FDX Network Capacity Modeling for Large SG with FDX Amps	20
5.1. FDX Use Case Overview	20
5.1.1. Traffic Growth vs Time Considerations.....	20
5.1.2. Spectrum Allocation Considerations.....	21
5.1.3. Network Capacity Modeling Assumptions	22
5.2. FDX Use Case 1 – Nearer Term Years 2023-25.....	23
5.2.1. Case 1 with 288 MHz Legacy Video, 396 MHz FDX Band.....	24
5.2.2. Case 1 with 288 MHz Legacy Video, 492 MHz FDX Band.....	27
5.2.3. Case 1 with IPTV, no Legacy Video, 684 MHz FDX Band	30
5.2.4. Probability of DS + US Overlapping Tails – Case 1	33
5.3. FDX Use Case 2 – Longer Term Years 2026-29.....	35
5.3.1. Case 2 with 144 MHz Legacy Video, 684 MHz FDX	35
5.3.2. Case 2 with IPTV, no Legacy Video, 684 MHz FDX Band	38
5.3.3. Probability of DS + US Overlapping Tails – Case 2	42
6. Results Summary	43
7. Conclusion.....	45
Abbreviations	46
Bibliography & References.....	47

List of Figures

Title	Page Number
Figure 1 – Massive New Upstream Bandwidth by Sharing Downstream and Upstream in the Same Spectrum using DOCSIS 4.0 Full Duplex (FDX).....	5
Figure 2 – Two Key New Innovations Power DOCSIS 4.0 FDX: Echo Cancellation and Interference Group / Transmission Group Formation.....	6
Figure 3 – Basic View of Echo Cancellation Concept.....	6
Figure 4 – FDX Resource Block Assignment and Sample Band Allocation “States”	7
Figure 5 – Topology of a DSP-Based FDX Capable Amplifier	8
Figure 6 – Potential Interference Group (IG) Elongation due to an FDX Amplifier.....	9
Figure 7 – Example of Upstream Capacity Usage.....	11

Figure 8 – Mapping Traffic Eng Formula to US Capacity Usage Example.....	13
Figure 9 – Network Capacity Example for 1G Service Tier	13
Figure 10 – Typical Subscriber Traffic Scenario	14
Figure 11 – Single Sub BW DS Transmit Probabilities, 1G @ 3 Mbps, 1-sec windows.....	15
Figure 12 – Single Sub BW DS Transmit Probabilities, 4G @ 15 Mbps, 1-sec windows.....	16
Figure 13 – Network Capacity Example for 1G Service Tier	17
Figure 14 – FDX Band as conceived, significant overlap, Multiple TG needed.....	18
Figure 15 – FDX Band as understood now, only burst overlap, Single TG needed	18
Figure 16 – FDX Band Size – SG Limits.....	18
Figure 17 – FDX Band Size – TG Limits.....	19
Figure 18 – Example Phases of Spectrum Migration vs Time.....	22
Figure 19 – Case 1 w Video, 396 MHz FDX – DS BW Capacity Limits	25
Figure 20 – Case 1 w Video, 396 MHz FDX – DS Spectrum Requirements.....	25
Figure 21 – Case 1 w Video, 396 MHz FDX – US BW Capacity Limits	26
Figure 22 – Case 1 w Video, 396 MHz FDX – US Spectrum Requirements.....	27
Figure 23 – Case 1 w Video, 492 MHz FDX – DS BW Capacity Limits	28
Figure 24 – Case 1 w Video, 492 MHz FDX – DS Spectrum Requirements.....	29
Figure 25 – Case 1 w Video, 492 MHz FDX – US BW Capacity Limits	29
Figure 26 – Case 1 w Video, 492 MHz FDX – US Spectrum Requirements.....	30
Figure 27 – Case 1 w IPTV, 684 MHz FDX – DS BW Capacity Limits.....	31
Figure 28 – Case 1 w IPTV, 684 MHz FDX – DS Spectrum Requirements	32
Figure 29 – Case 1 w IPTV, 684 MHz FDX – US BW Capacity Limits.....	32
Figure 30 – Case 1 w IPTV, 684 MHz FDX – US Spectrum Requirements	33
Figure 31 – Probability of DS + US Overlapping Tails – Case 1 IPTV, 200 subs	34
Figure 32 – Probability of DS + US Overlapping Tails – Case 1 IPTV, 200 subs (Log Scale).....	34
Figure 33 – Case 2 w Video, 684 MHz FDX – DS BW Capacity Limits	36
Figure 34 – Case 2 w Video, 684 MHz FDX – DS Spectrum Requirements.....	37
Figure 35 – Case 2 w Video, 684 MHz FDX – US BW Capacity Limits	37
Figure 36 – Case 2 w Video, 684 MHz FDX – US Spectrum Requirements.....	38
Figure 37 – Case 2 w IPTV, 684 MHz FDX – DS BW Capacity Limits.....	39
Figure 38 – Case 2 w IPTV, 684 MHz FDX – DS Spectrum Requirements	40
Figure 39 – Case 2 w IPTV, 684 MHz FDX – US BW Capacity Limits.....	40
Figure 40 – Case 2 w IPTV, 684 MHz FDX – US Spectrum Requirements	41
Figure 41 – Case 2 w IPTV, 684 MHz FDX – Various Service Tier Examples.....	41
Figure 42 – Probability of DS + US Overlapping Tails – Case 2 IPTV, 200 subs	42
Figure 43 – Probability of DS + US Overlapping Tails – Case 2 IPTV, 200 subs (Log Scale).....	43

List of Tables

Title	Page Number
Table 1 – CAGR Effects on per-user Peak Busy Hour (pbh) Average Utilization	21
Table 2 – Effective K values used for Case 1	24
Table 3 – Effective K values used for Case 2	35
Table 4 – TG Size Summary: Spectrum Scenarios and Speeds.....	44

1. Introduction

The promise of 10G is emerging, as Full Duplex DOCSIS (FDX) and Extended Spectrum DOCSIS (ESD) solutions make their way through laboratory testing to field trials towards first market launches. DOCSIS 4.0 consists of these two complementary technologies aimed at dramatically increasing upstream capacity and, correspondingly, upload speeds that customers will enjoy.

The FDX option for DOCSIS 4.0 is based on use of spectrum in the 108 MHz-684 MHz band (or subset of the band) for both downstream and upstream. Intuitively, signals overlapping in frequency and time interfere with one another. However, FDX has two key innovations that prevent this. Echo Cancellation (EC) technology removes potential interference where possible. Interference that cannot be removed by EC is *avoided* by controlling transmission timing within the DOCSIS scheduler. FDX sizes up the network and creates groups of users – “Interference Groups”(IG) – that should not access the same chunk of spectrum at the same time. This knowledge is incorporated into algorithms for allocating downstream and upstream resources.

Because the scheduler’s job is to fairly and efficiently allocate time and frequency resources, the effect on traffic engineering arises. In this paper, a deep dive into the mathematical modeling and analysis, based on empirical data from real DOCSIS HSD systems, will be described. The analysis will show:

- How large can an IG be before it effects performance
- What are the implications to IG size and FDX service group size
- What are the relationships among common traffic engineering variables – bandwidth, speed, penetration of services and tiers, service groups size – with the IG element introduced
- How do the results impact network migration strategy

This pioneering analysis breaks new ground on traffic engineering of FDX systems, speaks to key aspects of N+x FDX systems, and promises to be a foundation for field implementation guidelines of DOCSIS 4.0 FDX.

2. DOCSIS 4.0 Full Duplex Overview

Figure 1 illustrates the essential spectrum goal of FDX– enabling significantly more upstream. More interestingly, these new FDX upstream bands (Red “FDX” in **Figure 1**) are also available for downstream! This is quite different than typical Frequency Domain Duplex (FDD) operation, the approach taken in DOCSIS 4.0 FDD. How can both downstream and upstream data exist in the same spectrum? There are two essential innovations in DOCSIS 4.0 that enable this.

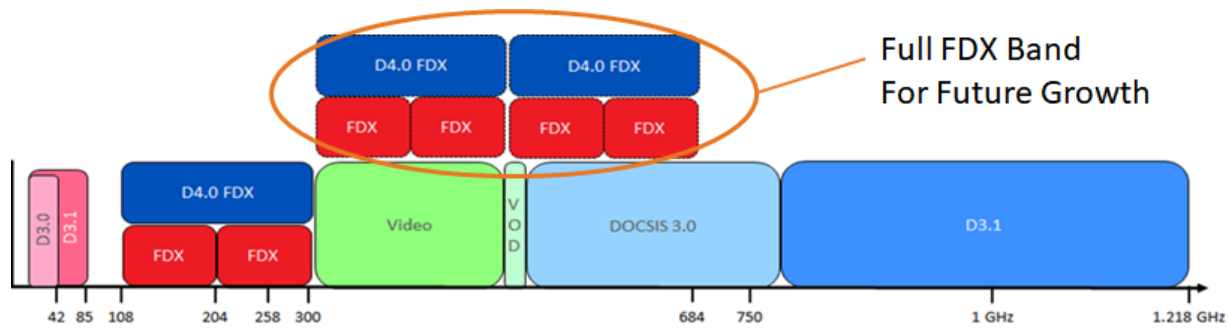


Figure 1 – Massive New Upstream Bandwidth by Sharing Downstream and Upstream in the Same Spectrum using DOCSIS 4.0 Full Duplex (FDX)

2.1. Key FDX Innovations

Although more upstream bandwidth is defined, it is the same 96 MHz OFDMA physical layer blocks that are defined in DOCSIS 3.1, except that the band over which they can operate is expanded. Thus, DOCSIS 4.0 leverages the power of the DOCSIS 3.1 PHY completely. Six additional 96 MHz blocks are added across the 108-684 MHz band, complementing an 85 MHz “mid-split” system.

Downstream and upstream can occupy the same band using a technology known as Echo Cancellation (EC). Echo Cancellation, in general, is a mature technology used in other telecom networks, such as xDSL and wireless. It has not yet been implemented in cable networks. The EC concept is very similar to those other applications, although the cable does introduce some new implementation challenges. EC is the first of the two critical innovations that power FDX.

The second key innovation is based on a fundamental architectural difference in cable systems compared to telco xDSL systems. Twisted pair telco networks are point-to-point connections from the DSL Access Multiplexer, or DSLAM, whereas HFC is a point-to-multipoint system. This logical architecture difference creates the need for another layer of innovation for FDX. This is the creation of Interference Groups (IGs) and Transmission Groups (TGs) for the scheduler to manage.

Figure 2 illustrates these innovations using a passive coaxial network (i.e. N+0) for simplicity. N+0 is NOT a requirement for FDX. In fact, one of the drivers for this paper is the evolution of FDX for N+x networks, whereas the DOCSIS 4.0 specification uses N+0 model types only as reference architectures. FDX-capable amplifiers will support FDX signals over N+x networks, allowing FDX to be implemented over a much wider swath of the footprint more quickly, and more of which is N+x than N+0. As we know from standard HFC networks, RF amplifier cascades impact end-of-line fidelity, and this is no different for FDX amplifiers. However, with amplifiers for FDX, an additional effect comes from the formation of IGs and TGs. As a result, an additional trade space created by N+x FDX networks are maximum speed tiers that can be supported for what amount of user penetration versus amplifier cascade depth. We will briefly discuss FDX amplifiers themselves in a subsequent section.

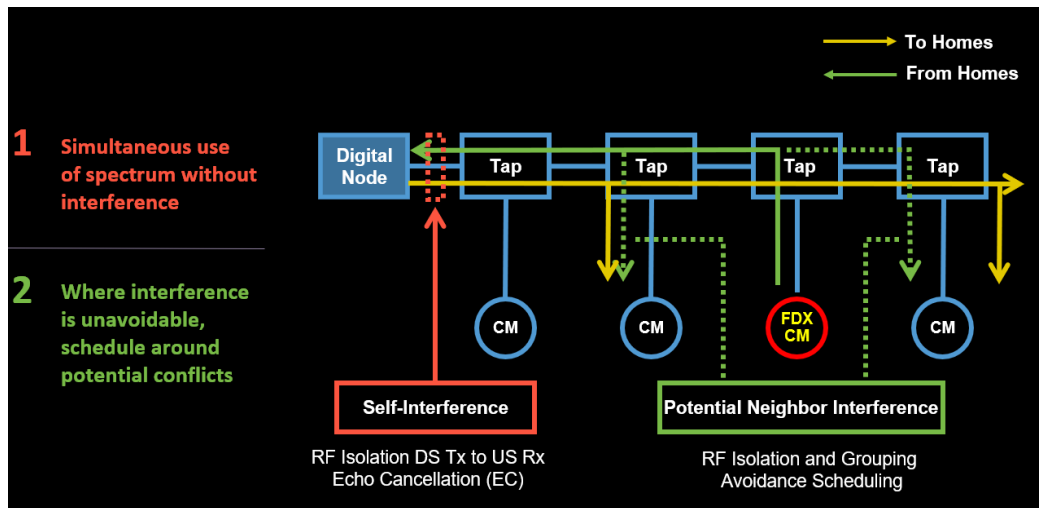


Figure 2 – Two Key New Innovations Power DOCSIS 4.0 FDX: Echo Cancellation and Interference Group / Transmission Group Formation

2.1.1. A Closer Look: Echo Cancellation

Referring to Figure 2, adding upstream signals at frequencies where downstream signals exist requires that the downstream signal be “subtracted” before the upstream (US) OFDMA receiver. This requires high RF isolation and strong EC of the much higher downstream signal that is reflected back into the US receiver. While the implementation details may be complex, the EC concept is a quite simple, and the digital signal processing (DSP) principles to build it are very mature. A simplified diagram illustrating the EC concept is shown in Figure 3.

The node downstream transmit signal will have some of its energy impose onto the US Rx simply through imperfect isolation characteristics of real hardware of an RPHY Node. It will also have some energy reflected back by imperfect RF interfaces as it travels down the coaxial plant, such as from the return loss of a tap, for example. These are the so-called “Echoes” that give the EC function its name. What is distinctive to EC for cable is the high cancellation required across a broad, multi-octave, bandwidth.




Figure 3 – Basic View of Echo Cancellation Concept

For more details describing Echo Cancellation and performance observed, please refer to [6][7][8][9][10].

2.1.2. A Closer Look: Interference Groups and Transmission Groups

As noted, the HFC network is a point-to-multipoint architecture. While an FDX modem knows its own upstream transmission, it cannot know that of his neighbor, and thereby has no way to “cancel” interference from a neighbor. It requires sufficient RF isolation among the homes sharing a coaxial RF segment to prevent FDX-band upstream users from interfering with a neighbor using that band for downstream. Unfortunately, RF isolation among homes cannot always be guaranteed to be high enough. Isolation relationships among homes are determined as part of FDX “sounding” process for DOCSIS 4.0 and DOCSIS 3.1 devices that support FDX-Light (aka FDX-L) functionality. A DOCSIS 3.1 CM that supports FDX-L becomes aware it is on an FDX-enabled network and can participate in the sounding process as a DS “measurer” device.

In situations without sufficient RF isolation, we instead call on the virtual CMTS (vCMTS) scheduler to avoid an interference scenario. Note that the vCMTS refers to the Comcast implementation of the CMTS function using commercial off-the-shelf servers hosting CMTS code and integrated with a DAA system based on a switched Ethernet architecture [7][8]. In an FDD system, the scheduler does not need to pay close attention to the relationship of downstream and upstream access to the coax. This changes in FDX. During FDX “sounding,” the FDX system determines these isolation relationships. Potentially interfering users are lumped into “Interference Groups,” or IGs. A logical set of IGs is called a Transmission Group (TG). Transmission Groups are created because not every IG needs to be treated independently – its overkill to do so, as we shall see later in this paper – and the vCMTS workload can be simplified by aggregating IGs into TGs. The scheduler assures that potentially interfering pairs are not accessing the same spectrum in the same time slot.

Because there are six OFDMA blocks, the vCMTS can service multiple IGs with uniform capacity and speeds by assigning different Resource Block Assignments (RBAs) to each IG. Figure 4 shows an example of how the 108-684 MHz FDX band might be allocated to simultaneously support a case with three TGs. These RBAs can adapt with time based on traffic and peak speed demand.

Note that the FDX band is not all of the DOCSIS spectrum available. Non-FDX DOCSIS 3.1 spectrum and DOCSIS 3.0 spectrum will also exist. Furthermore, not all of the FDX band needs to be assigned at all – more US bandwidth is allocated to FDX as peak upstream speed requirements increase.




Figure 4 – FDX Resource Block Assignment and Sample Band Allocation “States”

2.2. FDX-Capable Amplifiers

The FDX system specifications were written using the assumption of an N+0 network. However, as mentioned previously, FDX is not technically limited to an amplifier-free plant and the specification does not prevent it. In fact, shortly after the FDX specifications were completed, a CableLabs Study Group was formed to evaluate methods for implementing amplifiers that support FDX. The foundational EC

technology developed for the FDX RPD can be applied at any point in the network to manage overlapping spectrum, and this can include amplifiers. Of course, these are not traditional amplifiers, but a new class of device that includes digital signal processing (DSP).

An EC-based amplifier concept is shown in Figure 5. The nature of overlapping spectrum and gain in both directions creates a full-circle loop gain path in the FDX band. The EC must be capable of suppressing the FDX loop gain, such that the net gain around the path is < 0 dB across all frequencies to maintain a stable device. The EC must further be designed to act on the echo it is suppressing sufficiently that the aggregate residual echo noise, which becomes part of the amplifier's own noise floor, supports the US MER requirements effectively for DOCSIS 4.0, without introducing unacceptable MER degradation and subsequent loss of bandwidth efficiency.

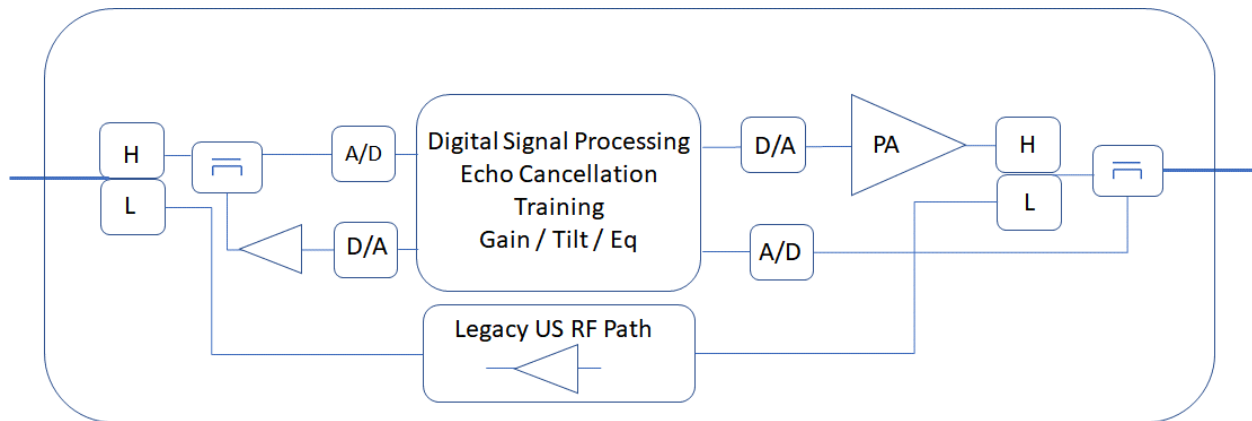


Figure 5 – Topology of a DSP-Based FDX Capable Amplifier

As noted previously, as an HFC amplifier cascade increases, noise contributions aggregate and the MER decreases. For the FDX amplifier, in the FDX band, there is an additional noise contributor in the form of residual echo. The amount of acceptable degradation due to the amplifier is a system engineering parameter that flows from performance specifications ultimately to the performance of the EC itself. A significant advantage for amplifier EC in comparison to an N+0 node output is that the levels on the DS Tx port are lower, and on the US Rx port higher. Thus, the DS to US level ratio is smaller, which is favorable for the EC.

While noise analysis is relatively straightforward for N+x with FDX, traffic engineering requires additional scrutiny due to the shared DS and US and the effect of amplifiers in expanding the size of an IG. Consider the N+1 system shown in Figure 6. When an amplifier is included, there is an expansion of Interference Group 4 (IG4) to the “south,” or home-facing, side of the amplifier. These users become part of the last IG of the tap string before the amplifier. This is because of the limited drop-to-output isolation characteristics of today's taps. This parameter can be optimized for high drop-to-output isolation, but until there was FDX to consider, there was no reason to drive more aggressive specifications for this parameter.




Figure 6 – Potential Interference Group (IG) Elongation due to an FDX Amplifier

3. FDX Traffic Engineering: An Operational Perspective

If we think about network segmentation triggers in current HFC networks, rules have been developed that link service group size, speeds, percent capacity utilization, and total capacity required. Then, with an awareness of device penetrations, this can be translated to DOCSIS 3.0 and DOCSIS 3.1 spectrum requirements and subsequently used to compare to thresholds that trigger a network augmentation.

Similar capacity-based analysis and empirical rule making will now apply to FDX, with one additional nuance. The FDX band, of course, can be allocated for both downstream and upstream. Infrequent bursts of peak speeds can therefore be called upon to service the downstream or the upstream, but they have a new shared resource dependency. The introduction of new terminology for FDX, in particular the terms “Interference Group” and “Transmission Group,” created an unfortunate aura of mystery around how the FDX band operates. This new terminology placed outsized significance on what turns out to be, from a traffic engineering point of view, a relatively benign, and certainly manageable phenomenon. What we have with FDX is a relatively modest “joint access” twist on HSD traffic engineering that we otherwise have been doing successfully for an independent downstream and independent upstream for decades.

Through these decades, operators have learned how to aggregate subscribers in a way that efficiently uses available capacity of the downstream and upstream while still delivering a high Quality of Experience (QoE) to the customer. Generally, operators have created empirically-based guidelines based on the key traffic engineering variables in play – capacity, service group size, HSD penetration, maximum speed, distribution of tiers, D3.1 and D3.0 mix, and also typically including some allowance baked in to account for a projected Compounded Annual Growth Rate (CAGR). Network augment rules have been developed over time to govern node splits, spectrum re-allocation, and device mix policies.

Instead of independent statistics of DS and of US, we now have a portion of the spectrum where resources are shared. However, we already have mature, reliable traffic models for downstream and upstream users! In either case – FDX or FDD – a set of users are sharing finite bandwidth and time resources, and the network design is deliberately NOT a non-blocking architecture, or it would be massively overdesigned. Users “compete” for the finite resources in the DS and in the US. The vCMTS makes decisions on packets sent DS and grants allocated US based on a set of QoE criteria determined empirically, as noted previously. Any individual user’s packets can be deferred for another’s *as part of*

normal CMTS operations – we just refer to it as *load balancing* and/or *fair scheduling*, and we developed rules and configuration settings for how to manage this effectively.

Again, for FDX, the only difference is there are users of DS and US both looking to access bandwidth and time resources, and the CMTS scheduler must accommodate them both. The mathematics and modeling of this joint access situation are in the sections to follow. They will underscore that an irrational “fear” of the Interference Group – seeking to minimize its size to a single Tap or two – was developed without any analysis or evidence to tell us just how big of a TG is too big. It just “felt” bad that US could interfere with DS unless we did something about it, that the two would both be looking to access resources, and one could be denied access in some mini-slot in deference to another. Of course, this happens ALL OF THE TIME in an independent DS or US!

A helpful *qualitative* way to think of the FDX band is to consider the DS users of the FDX band (D3.1+D4.0 CMs) as a service group (SG). Then, the US FDX band users can be considered a smaller set (smaller because D4.0 users < (D3.1 + D4.0) – and largely so for many years) of “new” users looking to, less frequently (because US utilization << DS utilization), access FDX band frequency and time resources. These FDX US “users” just happen to have the same MAC address and billing address of some of the DS users – which of course traffic engineering does not know nor care about. Because the upstream average and peak utilization is so much lower than downstream, and because the aggregate US traffic of a SG fits comfortably into the 85 MHz legacy band for many years of CAGR, except for brief, low-likelihood moments of peak speed bursts >300+Mbps, the FDX Band will be dominated by its use as a downstream channel. If the band was entirely dedicated to the upstream as it would be in an ultra-high split FDD system, it will be *idle* spectrum the vast majority of the time.

In summary, the DOCSIS 4.0 FDX Working Group (note – author’s included!) studied, in depth, the RF side of IGs and TGs, how to determine them, and incorporated it all into the specification. BUT – the working group did NOT do the math on practical implications of this phenomenon to the traffic engineering of the FDX band. We are now doing that Math! THE traffic engineering questions become:

- 1) How large can an IG be before there is an impact to the customer experience?
- 2) What service speed / IG size / spectrum rules exist when downstream and upstream traffic engineering become co-mingled in FDX?

These key questions are addressed in the traffic engineering analysis and modeling to follow. Before diving right into to answering these questions, however, we first take a step back to build up the foundation of burst traffic statistical behavior.

4. Network Capacity Planning

Determining the required amount of capacity needed for a service group (SG) is critical for providing customers with the appropriate quality of experience (QoE). Figure 7 shows an example of a SG upstream (US) capacity usage over a 16-minute window sampled at 1-second intervals. If an operator samples the bandwidth (BW) usage once every 16-minutes, then the purple line represents that average bandwidth for that measurement interval. This is a very useful datapoint but insufficient to determine the required capacity for the SG.

Sampling at 1-minute intervals would capture some of the variations, or ripples in the system but would still miss the many spikes which are individual modems bursting on top of the average BW. The gray line in the figure shows the high-water mark from 1-minute samples.




Figure 7 – Example of Upstream Capacity Usage

If an operator could sample at 1-second intervals or faster, then they could more closely determine the required SG capacity. The 1-second high-water mark is shown by the orange line. But sampling at these rates is not feasible in real systems, so another method is needed.

4.1. The “Basic” Traffic Engineering Formula

The CommScope (formerly ARRIS) team has been providing industry leading research in traffic engineering for many years which was most recently highlighted in [12]. The network capacity analysis in [12] provides an insight into how to calculate the SG capacity requirements. Previously, [4] introduced traffic engineering and QoE for broadband networks. From there, the paper develops a relatively simple traffic engineering formula for service groups which is easy to understand and useful for demonstrating basic network capacity components. Some additional references of note include [3],[5],[11],[13],[14], and [15].

The “Basic” formula shown below is a simple two-term equation. The first term ($N_{sub} * T_{avg}$) allocates bandwidth capacity to ensure that the aggregate average bandwidth generated by the N_{sub} subscribers can be adequately carried by the service group’s bandwidth capacity. The first term is viewed as the “DC component” of traffic that tends to exist as a continuous flow of traffic during the peak busy period. The growth rate of T_{avg} has seen much research. Some new revelations are discussed in [11].

“CLOONAN’S CLASSIC CAPACITY EQUATION” Traffic Engineering Formula:

$$C \geq (N_{sub} * T_{avg}) + (K * T_{max_max}) \quad (1)$$

where:

C is the required bandwidth capacity for the service group

N_{sub} is the total number of subscribers within the service group

T_{avg} is the average bandwidth consumed by a subscriber during the busy hour

K is the QoE constant (larger values of K yield higher QoE levels)...

where $0 \leq K \leq \text{infinity}$, but typically $1.0 \leq K \leq 1.2$

T_{max_max} is the highest Service Tier (i.e. T_{max}) offered by the MSO

There are obviously fluctuations that will occur (i.e. the “AC component” of traffic) which can force the instantaneous traffic levels to both fall below and rise above the DC traffic level. The second term ($K \cdot T_{\max_max}$) is added to increase the probability that all subscribers 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 service tier being offered to subscribers.

In previous papers [1], found that a K value of ~1.0 yields acceptable and adequate QoE results. [4] provides simulation results that showed a value between $K=1.0$ and 1.2 provides good QoE results for a service group of 250 subscribers. Larger service groups (SGs) would need even larger values of K while very small SGs might use a K value near or less than 1.0.

4.2. The “Modified” Traffic Engineering Formula

Over time, it was discovered that the optimum value for K would vary based on all the inputs: N_{sub} , T_{avg} and T_{\max_max} . Some of these limitations were noted in [13] along with some refinements to the basic formula above. This resulted in the following which is still algebraically equivalent to the basic formula:

Modified “CLOONAN’S CLASSIC CAPACITY EQUATION” Traffic Engineering Formula:

$$C \geq (N_{sub} \cdot T_{avg}) + (K-1) \cdot T_{\max_max} + T_{\max_max} \quad (2)$$

The subtle change is that there are now three main components to the traffic engineering formula:

1. Peak Busy Period Average Consumption (i.e. $N_{sub} \cdot T_{avg}$)
2. Peak Busy Period Ripple (i.e. $(K-1) \cdot T_{\max_max}$)
3. Headroom for maximum Service Tier Burst (i.e. $1 \cdot T_{\max_max}$)

Figure 8 shows how the modified formula maps to the US capacity usage example given above. The basic formula might have used a value of $K=1.2$ in this example. This is now broken into a burst component equal to T_{\max_max} plus a ripple component that is estimated by $20\% \cdot T_{\max_max}$.




Figure 8 – Mapping Traffic Eng Formula to US Capacity Usage Example




Figure 9 – Network Capacity Example for 1G Service Tier

Figure 9 shows a network capacity example for an operator who wants to support a 1 Gbps downstream (DS) service tier on a SG with 150 subs and an average peak consumption of 3 Mbps per subscriber. This requires at least 1500 Mbps of SG capacity. This includes a ripple component of $5\% * T_{max_max}$ which is 50 Mbps. Note that in this example, the ripple component as a function of T_{max_max} is much smaller than the initial $K=1.2$ values because T_{max_max} is significantly higher.

The Peak Busy Period Average Consumption and maximum Service Tier Burst components are well known and easily obtained. Traffic engineering research has since focused on quantifying the Peak Busy Period Ripple component as it is impacted by all inputs, not just T_{max_max} .

4.3. Determining Service Group QoE based on Probabilities

It became apparent that quantifying the SG subscribers QoE needed to focus on the probabilities for SG capacity. And to predict behavior across different SGs with different parameters, a network capacity transmit model for individual subscribers was necessary.

4.3.1. Individual Subscriber Bandwidth Probabilities

Data was collected across an entire CMTS for multiple years during prime-time evening hours. Every packet was captured; associated with a particular modem and service tier; and timestamped down to a millisecond. Big data analytics was then used to find the patterns across the different types of subscribers.

It was discovered that all subs had very similar patterns. Figure 10 shows a very typical sequence that occurs frequently. Around 70% to 80% of internet traffic these days are video which pre-dominantly uses adaptive bit rate (ABR) encoding. The ABR streaming (#1 in figure) is a “grassy” region where traffic is limited by the video resolution (e.g. 5-10 Mbps). A second ABR starts later (#5 in figure) and begins with a burst to pre-fill its video buffers. This burst size is relatively small, limited by the video rate and buffer depth.




Figure 10 – Typical Subscriber Traffic Scenario

Sandwiched between these ABR streams is a file download. For a short instance, it may burst above the modem’s Tmax rate before the DOCSIS token bucket algorithms kick in to limit the transfer to its Tmax value. In today’s world, a 12 MB file download might represent a couple high resolution pictures or a PDF/Word/Powerpoint document. A modem with a 100 Mbps tier might download this in ~1 second while a 1 Gbps tier modem only needs 0.1 seconds.

The Big Data Analytics then determined the BW transmit probabilities for each identified type of subscriber using 1-second windows. It was found that using a window too fast (e.g. milliseconds) made it difficult or impossible to see the patterns while too slow (e.g. minutes) caused the bursts as shown above to be missed. Figure 11 shows the probability mass function (pmf) for a subscriber with a 1 Gbps tier and an average peak consumption of 3 Mbps. [note – the pmf is just a discrete version of the probability density function (pdf).]

Using the 12 MB file download example above, the 1G sub transfers 96 megabits in 0.1 seconds. Assuming the remainder of the 1-sec window is idle, then that burst would correspond to 96 megabits in a 1-sec window below even though the instantaneous rate is 1 Gbps.




Figure 11 – Single Sub BW DS Transmit Probabilities, 1G @ 3 Mbps, 1-sec windows

Being a log scale shows how quickly the probabilities drop for large bursts. Looking at the area under the curve >900 Mbps gives us a probability of 0.00114% which equates to this occurrence of a sustained burst >900 Mbps happening for 1 sec for every 24 peak period busy hours. With ~3 peak period busy hours per day, this means 1 sec for every 8 days.

The Big Data Analytics also found patterns as the Tavg and Tmax are varied from subscriber group to subscriber group. This allows subscriber behavior to be predicted for future Tavg and Tmax values. Figure 12 shows the pmf for a hypothetical user in 2030 with a 4G tier consuming 15 Mbps during peak. Not only does the “Tmax bump” get pushed to the right, but the probability tail gets pulled lower too. The probability of a burst >3600 Mbps is now an order of magnitude smaller at 0.000117%. This maps to 1 sec every 240 peak period hours which might only be 1 sec every 80 days.




Figure 12 – Single Sub BW DS Transmit Probabilities, 4G @ 15 Mbps, 1-sec windows

In the FDX Amplifier Working Group previously described, an example was given for probabilities for subs with a 1 Gbps speed tier and a 2 Mbps Tavg:

- Probability (a single random sub creates a DS burst $> (0.9 \cdot T_{max})$ in a 1-sec sample window) = $3.4e-6$ = 1 second out of every 3.4 days
- Probability (50 subs sharing BW create a single DS burst $> (0.9 \cdot T_{max})$ in a sample window) = $(3.4e-6) \cdot 50$ = 1 second out of 1.6 hours
- Probability (2 subs simultaneously create a DS burst $> 2 \cdot (0.9 \cdot T_{max})$ in a sample window) = $(3.4e-6)^2$ = 1 second out of 2743 years

As can be appreciated, these bursts to $>90\%$ of T_{max} are relatively rare.

4.3.2. Service Group Subscriber Bandwidth Probabilities

Once single subscriber BW pdf's were in place, it is then possible to create a SG BW pdf to analyze the behavior at the SG level. An example of the SG BW DS probabilities is shown in Figure 13.

This is the output of a Monte-Carlo simulation with 100K trials. It is a SG that consists of 128 subs, all of which have a 1G DS service tier. The Tavg = 15 Mbps which represents a time later this decade.




Figure 13 – Network Capacity Example for 1G Service Tier

Note the asymmetry in the curve – there is a much longer tail to the right. Various cumulative distribution function (CDF) probability thresholds are shown (i.e. 90%, 99%, 99.9%, 99.99%) to provide an insight as to the probability the SG BW hits different capacities.

The blue candlesticks in the figure show how the traffic engineering formula overlays the SG PDF. In this instance, the formula said 3,120 Mbps is required. Note that this is sufficient for ~99.5% of the peak busy period time. For the remaining 0.05% of the time, the buffers may be temporarily filled, and some latencies introduced. Note that 0.05% represents ~50 seconds out of every evening. For most of those 50 seconds, the delays should be insignificant and not noticed by users. This is normal network behavior and is why many users can share a single network pipe.

The remainder of the document uses these candlesticks often. Please keep in mind how low the probabilities are for their burst regions as shown above.

4.4. FDX Traffic Eng for Large SG with Overlapping Upstream + Downstream

The FDX Amp use case is leveraging existing plant and may contain many 100's of subscribers. When FDX was conceived, this was not envisioned as a viable use case. In the early days, it was thought that there would be significant overlap between the US and DS in the FDX band. This is shown in Figure 14.

The more that the DS and US peak period average consumption overlaps, then the more important it becomes to have multiple FDX transmission groups (TG).




Figure 14 – FDX Band as conceived, significant overlap, Multiple TG needed

What has subsequently come to light is that the multi-gigabit burst region is dominating our traffic engineering formula. This means that the burst region (i.e. T_{max_max}) is much larger than the consumption component (i.e. $N_{sub} * T_{avg}$). Figure 15, drawn to scale from one of our upcoming scenarios, shows how only the DS + US burst regions overlap in the FDX band. Most of the US traffic fits below 85 MHz while most of the DS traffic stays in the dedicated DS spectrum above 684 MHz.




Figure 15 – FDX Band as understood now, only burst overlap, Single TG needed

In one of our examples, the US stay below 85 MHz for 99.8% of the time and only needs to access the FDX band for 0.2% of the time. Meanwhile, the DS stays above 684 MHz for 99.9% of the time and only needs the FDX band for 0.1% of the time. So, the FDX band is basically providing the burst bandwidth (i.e. 5G DS, 4G US in this case) and the probabilities of US + DS bursts overlapping is microscopic.

There are certain SG capacity requirements that must be met in any system, whether it is FDX or not. These include:

1. $DS\ SG\ Capacity \geq N_{sub}(SG) * DS\ T_{avg} + DS\ T_{max_max} + DS\ Ripple$
2. $US\ SG\ Capacity \geq N_{sub}(SG) * US\ T_{avg} + US\ T_{max_max} + US\ Ripple$

These two traffic engineering (TE) conditions are shown pictorially in Figure 16.




Figure 16 – FDX Band Size – SG Limits

In an FDX world, that means that the DS capacity above 684 MHz plus the DS capacity in the FDX band must be sufficient to meet TE condition #1. This must be met, totally independent of US requirements. Also, the US capacity below 85 MHz plus the US capacity in the FDX band must be sufficient to meet TE condition #2. This must be met, totally independent of DS requirements. In this respect, an FDX network's capacity requirements is no different than any other network.

19

5. FDX Network Capacity Modeling for Large SG with FDX Amps

Business decisions about the timing and pace of introducing multi-gig symmetric speed tiers will govern the pace of the migration of the network. There is uncertainty in the timing of launching these new tiers, the pace at which they roll-out nationally, and in the traffic growth to anticipate what will have occurred by the time they are launched. Therefore, a range of spectrum scenarios and target speeds were teed up for evaluation to cover what today appears to be realistic representative situations for FDX launch and years into the future of its use.

5.1. FDX Use Case Overview

5.1.1. *Traffic Growth vs Time Considerations*

Table 1 identifies two traffic utilizations that can apply for four different growth scenarios based on CAGR assumptions. If the CAGR is 30% in 2023 and persists,, then, as shown in the table, an average downstream user at peak busy hour will consume approximately 7-8 Mbps downstream and 500 kbps upstream. That same DS/US ratio will also be approximately correct in 2026 if the CAGR is instead 15%.

Similarly, if the 30% or 15% persists through 2026 or 2029, the DS and US will be approximately 15 Mbps and 1 Mbps, respectively. Of course, the further out the model is taken, the more uncertain it is, as CAGR typically wobbles year-to-year. However, in recent years, with the exception of the COVID-induced spike which is still settling, residential CAGRs have dropped from the 50% range of the 2010 time frame through the low 30% range pre-COVID, and it continues to decline. The upstream had been flattening out, decreasing to the <15% range pre-COVID, and appears to is returning to that range.

With these utilization boundaries, we have run models to draw out the acceptable TG size for speed tiers ranging from 1 Gbps to 5 Gbps for US and DS. We have also added some potential higher speed DS projections for 6 Gbps and 7 Gbps. Note that the speeds possible have a dependency on how efficient the bandwidth can be used. The models therefore also show the TG threshold range as a function of achievable QAM format for DS and US. Minor variations of QAM efficiency are expected, just as they are in HFC, as RF amplifiers incrementally reduce DS and US signal fidelity (MER) at the respective receivers. These fidelity losses translate to relatively minor variations in capacity. For example, if 4096-QAM is unachievable, but 2048-QAM is, then the capacity loss is about 8.3%.

Table 1 – CAGR Effects on per-user Peak Busy Hour (pbh) Average Utilization

Near term (2023) w aggressive CAGR (30%) or longer term (2025) with modest CAGR (15%): **7 Mbps DS / 500 kbps US**

Longer term (2026) w aggressive CAGR (30%) or even longer term (2029) with modest CAGR (15%): **15 Mbps / 1 Mbps**

		Network Migration Plan Year								
			1	2	3	4	5	6	7	8
		2022	2023	2024	2025	2026	2027	2028	2029	
	CAGR %	Tavg								
DS	15	4.5	5.18	5.95	6.84	7.87	9.05	10.41	11.97	13.77
	20		5.40	6.48	7.78	9.33	11.20	13.44	16.12	19.35
	25		5.63	7.03	8.79	10.99	13.73	17.17	21.46	26.82
	30		5.85	7.61	9.89	12.85	16.71	21.72	28.24	36.71
US	15	0.3	0.35	0.40	0.46	0.52	0.60	0.69	0.80	0.92
	20		0.36	0.43	0.52	0.62	0.75	0.90	1.07	1.29
	25		0.38	0.47	0.59	0.73	0.92	1.14	1.43	1.79
	30		0.39	0.51	0.66	0.86	1.11	1.45	1.88	2.45

5.1.2. Spectrum Allocation Considerations

There is a time element to spectrum allocation as CAGR and new speed tiers – forcing functions for new blocks of spectrum that must be large enough to deliver the highest speed tier, with margin – require continuous maintenance of the channel maps. This has been an HFC story for decades: managing spectrum as HDTV took hold, alongside VOD and niche/specialized video channels grew even as HSD CAGR was consistently high. This ultimately led to the elimination of inefficient analog carriage altogether, as well as technologies such as Switched Digital Video (SDV) to perform the spectrum balancing act.

The contemporary version of this balancing act is DOCSIS 3.1 vs less than DOCSIS 3.1 devices. Operators are looking to add more spectrum that supports DOCSIS 3.1, to take advantage of the increased bandwidth efficiency (reminder note: DOCSIS 4.0 uses the same DS and US OFDM and OFDMA technology), and to reduce video single carrier QAM as channels move to IP delivery. The IP video Nirvana is removing single carrier QAM video altogether, but more realistically a long term scenario is reducing it to a small block that is sufficient to support the segment of the customer base that relies on this service.

Referring to Figure 18, five spectrum allocation cases are shown. The five are broken into two categories, near term (scenarios 1-3) and longer term (scenarios 4-5). Scenarios 1-3 will use as their peak busy hour average usage for DS and US as 7 Mbps and 500 kbps, respectively per the above. Scenarios 4 and 5 will use the peak busy hour average utilization for DS and US as 15 Mbps and 1 Mbps, respectively.

Scenario 3 is based on near term CAGR values, but presumes that an All-IP conversion has taken place – i.e. no video single carrier QAMs. This scenario is to serve more as a boundary, “what’s possible” scenario under these particular average DS and US utilizations. Scenario 5 is also an All-IP conversion, in a longer term and potentially practical time frame, for comparison, using the increased utilization that would occur with the CAGR values described in the model descriptions above

The near term scenarios 1 and 2 show a spectrum still heavily weighted by video single carrier QAM spectrum – 48 slots to be exact. There is also significant bandwidth set aside for DOCSIS 3.0 (168 MHz) to support the still large volume of D3.0-only modems in the network. While the balance is rapidly shifting, as of this writing, there are still more D3.0 modems in the network than D3.1 modems at Comcast.

Scenario 1 can be viewed as an initial launch scenario, supporting 2 Gbps symmetric speed. Scenario 2 provides an additional OFDMA block that enables 3 Gbps symmetric services. And again, scenario 3 is a “what if” under these lighter DS and US traffic loads.

Scenario 4 and 5 presume that significant harvesting of video single carrier QAM and D3.0 modems has taken place, and all or most of the spectrum is DOCSIS 3.1 enabled DS and the FDX US band is maximized to 684 MHz. In these longer term scenarios, access to higher symmetric speeds – 4 Gbps and 5 Gbps – is evaluated, as well as DS speeds above 5 Gbps. In the DS, there is more total spectrum to allocate – out to 1218 MHz – than in the DOCSIS 4.0 upstream.

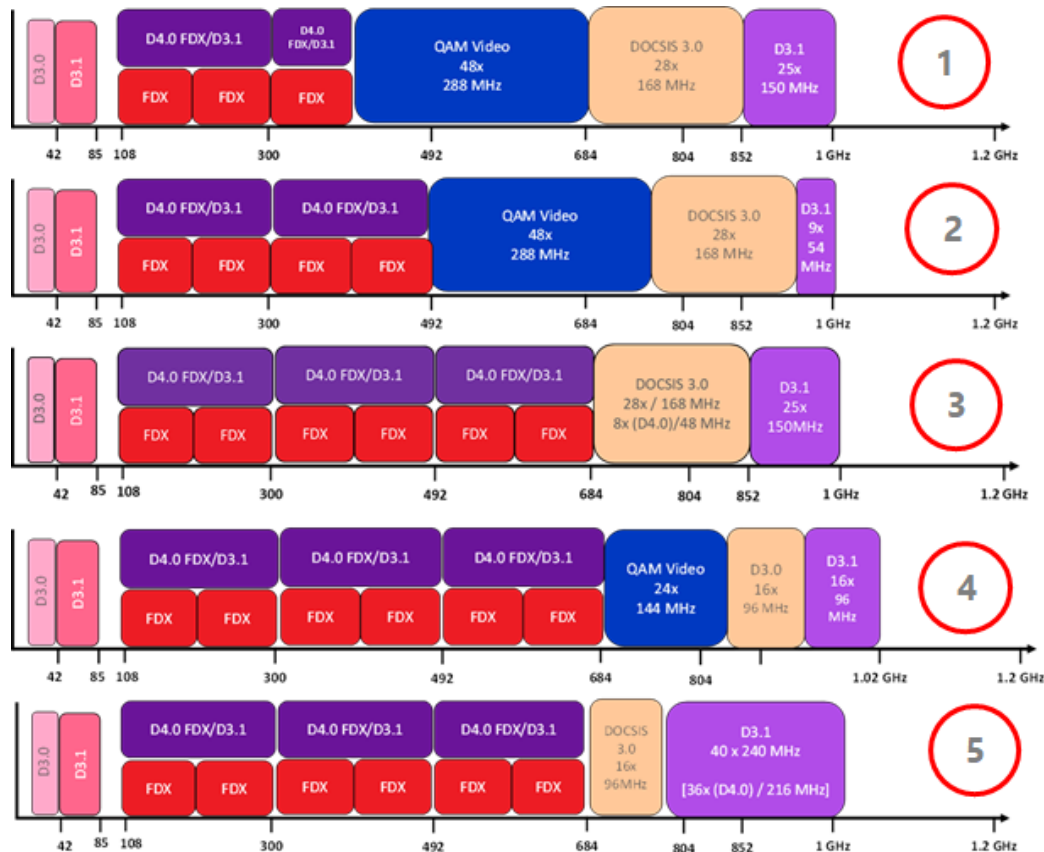


Figure 18 – Example Phases of Spectrum Migration vs Time

5.1.3. Network Capacity Modeling Assumptions

For our network capacity modeling, the following inputs were used in all scenarios:

- DS modulation = 4096-QAM for dedicated DS spectrum above FDX band
- Maximum DS frequency = 1002 MHz
 - Roll-off region not considered for this analysis, but could help
- DS modulation = 1024- to 4096-QAM inside the FDX band
- DS OFDM parameters: 1.25 usec Cyclic Prefix, 8K FFT
 - 4096-QAM => 9.70 bps/Hz

- 1024-QAM => 8.08 bps/Hz
- US fixed capacity below 85 MHz = 450 Mbps
 - Equates to 4 x 6.4 MHz SC-QAM + ~46 MHz OFDMA @ 1024-QAM
- US modulation = 64- to 1024-QAM inside the FDX band
- US OFDMA parameters: 1.25 usec Cyclic Prefix, 4K FFT
 - 1024-QAM => 7.70 bps/Hz
 - 256-QAM => 6.16 bps/Hz
 - 64-QAM => 4.62 bps/Hz
- SG sizes are varied from 32 to 384 subs
- Maximum service tier (Tmax_max) varied to fit up to 5G DS, 4G US
 - Every subscriber (i.e. 100%) takes the maximum service tier

5.2. FDX Use Case 1 – Nearer Term Years 2023-25

Case 1 assumes a DS Tavg = 7 Mbps with US Tavg = 0.5 Mbps. Three scenarios were considered: two with legacy video and one with 100% video over IP (IPTV) – meaning no QAM video carriers, all have been reclaimed for DOCSIS 3.1.

For our network capacity modeling, a SG CDF was created to allow us to determine the probability of various capacity thresholds. The 70th percentile point was used to effectively estimate the consumption (i.e. Nsub*Tavg) plus the ripple components.

Many folks are still familiar with the original traffic engineering formula using the K-value. To help relate our network capacity modeling results, Table 2 shows the effective K-value that would generate very similar results. As mentioned earlier, the optimum K-value varies with Nsub, Tavg and Tmax_max. The Tavg is fixed for Case 1, so the table shows the range of the other inputs.

For all US combinations, the service tier (Tmax_max) is >1,000 times larger than the Tavg. Thus, the effective K-value is extremely small, 1.025. For the DS, the Tmax_max to Tavg ratio isn't nearly as wide. For 4G & 5G tiers, the effective K-value is 1.05. For 3G tiers, the effective K-value is 1.05 until the SG size reaches 384 subs where it increases to 1.1. For 2G tiers, the effective K-value is 1.05 until the SG size reaches 256 subs where it increases to 1.1.

Table 2 – Effective K values used for Case 1

K-value Used	Original K		QoE modeling – Effective K			
Nsub/SG	DS	US	2G DS	3G DS	4G,5G DS	All US
32	1.2	1.2	1.05	1.05	1.05	1.025
64	1.2	1.2	1.05	1.05	1.05	1.025
128	1.2	1.2	1.05	1.05	1.05	1.025
256	-	-	1.1	1.05	1.05	1.025
384	-	-	1.1	1.1	1.05	1.025

5.2.1. Case 1 with 288 MHz Legacy Video, 396 MHz FDX Band

This scenario has 288 MHz of Legacy Video spectrum and the FDX band is limited to 108-396 MHz. This leaves DOCSIS with enough spectrum to fit 28 SC-QAM channels for D3.0/D3.1 modems plus 150 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~2.5 Gbps of fixed DS capacity.

Figure 19 shows the DS BW Capacity required for several different T_{max_max} tiers. The Y-axis is the capacity in Mbps while the X-axis is a log scale of the number of subs in a SG (i.e. Nsub). The yellow curve shows the DS peak period average consumption (i.e. Nsub*T_{avg}). The various green curves show the maximum capacity required for various T_{max_max}, from 2G to 5G. Note that the burst plus ripple components is added to the average consumption.

The dashed lines show various key capacity marks. The lowest one is the SC-QAM capacity of 1.05 Gbps. This is the upper capacity limit for 3.0 modems. A separate calculation must be done to ensure the 3.0 pool does not exceed their capacity limit. The next blue dashed line shows the Fixed DS capacity limit of ~2.5 Gbps. As long as the DS peak period average consumption is less than this, then the US TG limit (i.e. TE condition #4) never kicks in. The figure highlights where Nsub*T_{avg} crosses this point at 358 subs. Above this point, the DS average consumption starts taking away usable FDX BW from the US.

The top 3 blue dashed lines show the total DS capacity available for the various FDX DS QAM modulations (i.e. 1024-, 2048-, 4096-QAM). The point where a green curve crosses the blue dashed line is the maximum number of subs that can be supported for that combination, i.e. TE condition #1.

This scenario supports:

- 2G DS tier to 375+ subs
- 3G DS tier to 240 subs @ 1024-QAM; 306 subs @ 4096-QAM
- 4G DS tier to 90 subs @ 1024-QAM; 157 subs @ 4096-QAM
- 5G DS tier does NOT fit




Figure 19 – Case 1 w Video, 396 MHz FDX – DS BW Capacity Limits




Figure 20 – Case 1 w Video, 396 MHz FDX – DS Spectrum Requirements

Figure 20 shows the DS spectrum requirements for each of the above combinations. Note – this includes both the DOCSIS DS spectrum and 288 MHz Legacy Video. This analysis assumes 1002 MHz is the top spectrum. If the roll-off for the 1 GHz is usable or the taps are upgraded to 1.2 GHz, then the figure also shows what can be achieved up to 1200 MHz. Even at 384 subs where Nsub*Tavg has overflowed into the FDX band, it only needs >96MHz of FDX BW <0.25% of the time.

To put the DS in context, here are examples of FDX band probability usage. A SG with 256 subs and a 4G DS tier will stay within the fixed DS BW outside FDX band for 98% of the time and only need some additional FDX BW for just 2% of the time.

The US SG limits (i.e. TE condition #2) are examined in Figure 21. Like the DS chart, the yellow curve depicts the US $N_{sub} \cdot T_{avg}$ while the various orange curves show the total US capacity required for various US tiers (1.5G – 2.5G). The red dashed line on the bottom shows the fixed US capacity of 450 Mbps below 85 MHz. Notice that the $N_{sub} \cdot T_{avg}$ component never crosses this line. This means that the DS TG limit (TE condition #3) will never kick in.

The three upper dashed lines show the total US capacity available for the various FDX US QAM modulations (i.e. 64-, 256-, 1024-QAM). The point where an orange curve crosses the blue dashed line is the maximum number of subs that can be supported for that combination, i.e. US TE condition #2 from above. As can be seen, each jump in modulation effectively allows an additional 0.5 Gbps to be added to T_{max_max} .




Figure 21 – Case 1 w Video, 396 MHz FDX – US BW Capacity Limits




Figure 22 – Case 1 w Video, 396 MHz FDX – US Spectrum Requirements

Figure 22 shows the US spectrum requirements. The upper red dashed line is showing the top of the usable FDX spectrum. This is nominally 396 MHz when the DS peak period consumption does not overflow into the FDX band. However, Figure 22 shows that this point occurs at 358 subs. That is why this red dashed line starts to drop as Nsub increases past this point.

This scenario supports:

- 1.5G US tier to ~370 subs @ 64-QAM; 400+ subs @ 256- or 1024-QAM
- 2G US tier to ~350 subs @ 256-QAM; 400+ subs @ 1024-QAM
- 2.5G US tier to ~210 subs @ 1024-QAM
- 3G & 4G US tier does NOT fit

To put the US in context, a SG with 256 subs and a 2G US tier will stay within 85 MHz BW outside FDX band for 98% of the time and only need some additional FDX BW for just 2% of the time.

5.2.2. Case 1 with 288 MHz Legacy Video, 492 MHz FDX Band

This scenario increases the FDX band by 96 MHz to 108-492 MHz at the expense of the top OFDM channel. This leaves DOCSIS with enough fixed spectrum above the FDX band to fit 28 SC-QAM channels for D3.0/D3.1 modems plus 54 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~1.57 Gbps of fixed DS capacity. It also maintains 288 MHz of Legacy Video spectrum.

Figure 23 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now dropped ~1.57 Gbps. This causes the US TG limit (i.e. TE condition #4) to kick in at 225 subs where Nsub*Tavg crosses this point. The DS average consumption starts taking away usable FDX BW from the US at a much lower point in this scenario.

The top 3 blue dashed lines show the total DS capacity available for the various FDX DS QAM modulations (i.e. 1024-, 2048-, 4096-QAM). This scenario supports:

- 2G DS tier to ~354 subs @ 1024-QAM; 400+ subs @ 2048- & 4096-QAM
- 3G DS tier to ~218 subs @ 1024-QAM; 306 subs @ 4096-QAM
- 4G DS tier to ~68 subs @ 1024-QAM; 157 subs @ 4096-QAM
- 5G DS tier does NOT fit

The DS 4096-QAM results are the same as before, while the DS 1024-QAM results are reduced from the previous scenario.




Figure 23 – Case 1 w Video, 492 MHz FDX – DS BW Capacity Limits

Figure 24 shows the DS spectrum requirements for each of the above combinations. Note – this includes both the DOCSIS DS spectrum and 288 MHz Legacy Video.

To put the DS in context, a SG with 128 subs and a 4G DS tier will stay within the fixed DS BW outside FDX band for 99% of the time and only need some additional FDX BW for just 1% of the time. Even at 256 subs where Nsub*Tavg has overflowed into the FDX band, it only needs >96MHz of FDX BW <5% of the time.

The US SG limits (i.e. TE condition #2) are examined in Figure 25. The various orange curves show the total US capacity required for various US tiers (1.5G – 3G). Like the previous scenario, each jump in modulation effectively allows an additional 0.5 Gbps to be added to Tmax_max.




Figure 24 – Case 1 w Video, 492 MHz FDX – DS Spectrum Requirements




Figure 25 – Case 1 w Video, 492 MHz FDX – US BW Capacity Limits




Figure 26 – Case 1 w Video, 492 MHz FDX – US Spectrum Requirements

Figure 26 shows the US spectrum requirements. The upper red dashed line is nominally 492 MHz until the DS peak period consumption overflows into the FDX band at 225 subs. The drop in usable US BW above 225 subs is the result of TE condition #4. As Nsub increases, the impact becomes greater.

This scenario supports:

- 1.5G US tier to ~360 subs @ 64-QAM; 400+ subs @ 256- or 1024-QAM
- 2G US tier to ~235 subs @ 64-QAM; ~340 subs @ 256-QAM
- 2.5G US tier to ~245 subs @ 256-QAM; 330 subs @ 1024-QAM
- 3G US tier to ~250 subs @ 1024-QAM
- 4G US tier does NOT fit

To put the US in context, a SG with 256 subs and a 2G US tier will stay within 85 MHz BW outside FDX band for 98% of the time and only need some additional FDX BW for just 2% of the time.

5.2.3. Case 1 with IPTV, no Legacy Video, 684 MHz FDX Band

In this scenario, 100% IPTV migration is reached which frees up the 288 MHz Legacy Video to be added to the DOCSIS BW pool. The FDX band is now a full 108-684 MHz while the fixed DS capacity above 684 MHz returns to 28 SC-QAM channels for 3.0/3.1 modems plus 150 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~2.5 Gbps of fixed DS capacity.

Figure 27 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now back up to ~2.5 Gbps. The US TG limit (i.e. TE condition #4) kicks in at 358 subs like the first scenario.

Notice that the top 3 blue dashed lines showing the total DS capacity available has increased significantly to 7-8 Gbps. The Case 1 IPTV scenario supports:

- 2G, 3G & 4G DS tier to 400+ subs @ all modulations
- 5G DS tier to ~272 subs @ 1024-QAM; 400+ subs @ 4096-QAM

Figure 28 shows the DS spectrum requirements for each of the above combinations. The probability of using the FDX band is the same as scenario 1 above. The US SG limits (i.e. TE condition #2) are examined in Figure 29. The various orange curves show the total US capacity required for various US tiers (2G – 4G). Figure 30 shows the US spectrum requirements. The Case 1 IPTV scenario supports:

- 2G US tier to 400+ subs @ all modulations
- 3G US tier to 400+ subs @ 256- & 1024-QAM
- 4G US tier to 400+ subs @ 1024-QAM




Figure 27 – Case 1 w IPTV, 684 MHz FDX – DS BW Capacity Limits




Figure 28 – Case 1 w IPTV, 684 MHz FDX – DS Spectrum Requirements




Figure 29 – Case 1 w IPTV, 684 MHz FDX – US BW Capacity Limits




Figure 30 – Case 1 w IPTV, 684 MHz FDX – US Spectrum Requirements

5.2.4. Probability of DS + US Overlapping Tails – Case 1

The analysis until now has focused on capacity requirements and how many subscribers can receive adequate service for a given configuration. Now let's take a closer look at the probabilities of a DS tail overlapping with an US tail in the shared FDX band. Again, this assumes this is one single large TG, so the US and DS need to timeshare any FDX channels they need to use.

Figure 31 is a complex chart showing the probabilities of spectrum usage for both US + DS CDF with 200 subs. Our model assumes that the US capacity fills from the lowest channel in the spectrum and increases to the right until it reaches the top of the FDX band. The DS capacity is the opposite where it fills from the top channel in the spectrum down to the left. The fixed DS capacity is completely used before the DS uses the FDX band.

The yellow curve on the left is the US 'CDF'. Technically, it is the (1 – CDF) function. It tells us the probability that the US will need MORE than that amount of spectrum. The curve crosses the red line (85MHz) at 0.2%. This means that 99.8% of the time, the US capacity remains completely within 85 MHz, and only 0.2% of the time does it even need to request FDX BW. That equates to only 20 seconds every evening!

The blue curve on the right is the DS 'CDF'. It tells us the probability that the DS will need spectrum below that point. The curve crosses the top of the FDX band (684MHz) at 0.1%. This means that 99.9% of the time, the DS capacity remains completely above 684 MHz, and only 0.1% of the time does it even need to request FDX BW. That equates to only 10 seconds every evening!




Figure 31 – Probability of DS + US Overlapping Tails – Case 1 IPTV, 200 subs




Figure 32 – Probability of DS + US Overlapping Tails – Case 1 IPTV, 200 subs (Log Scale)

To take a closer look at the overlapping tails, Figure 32 shows the same data but on a log scale. The US datapoint for 0.01% is at 128 MHz. What this means is that the US needs more than 20 MHz worth of capacity from the FDX band 0.01% of the time. For 20 MHz worth of capacity, the scheduler could assign this as 25% of a single 96 MHz FDX channel or as 4% of the entire FDX band. So, the US can still burst to the top of the FDX band, it is just that it is very infrequent and short bursts that shouldn't congest the system. In the rare event that the US does have a sustained burst (e.g. speed test once a month), the yellow dashed line at 599 MHz shows the US upper capacity limit based on the traffic engineering formula. The US will not exceed this BW point over a sustained period.

Looking at the DS CDF, the 0.01% datapoint is at 533 MHz. This means that 99.99% of the time, the DS BW needs are above this point and only 0.01% of the time does the DS need to go below this point. 533 MHz represents ~80% use of a single 192 MHz FDX DS channel or ~26% of the entire FDX band. The blue dashed line at 257 MHz represents the lowest spectrum BW point that the DS needs based on the traffic engineering formula (e.g. a speed test on top of normal traffic).

During the normal course of events, the probability of the US + DS tails overlapping is on the order of 1 second every 100+ years! Case 1 is the easier one, let's see how things change with Case 2.

5.3. FDX Use Case 2 – Longer Term Years 2026-29

Case 2 assumes a DS Tavg = 15 Mbps with US Tavg = 1 Mbps. Two scenarios were considered: one with legacy video and one with 100% IPTV.

Table 3 shows the effective K-value for Case 2. As mentioned earlier, the optimum K-value varies with Nsub, Tavg and Tmax_max. This table is identical to Table 1 except for the 3G DS tier. For 3G tiers, the effective K-value is 1.05 until the SG size reaches 256 subs where it increases to 1.1.

Table 3 – Effective K values used for Case 2

K-value Used	Original K		QoE modeling – Effective K			
Nsub/SG	DS	US	2G DS	3G DS	4G,5G DS	All US
32	1.2	1.2	1.05	1.05	1.05	1.025
64	1.2	1.2	1.05	1.05	1.05	1.025
128	1.2	1.2	1.05	1.05	1.05	1.025
256	-	-	1.1	1.1	1.05	1.025
384	-	-	1.1	1.1	1.05	1.025

5.3.1. Case 2 with 144 MHz Legacy Video, 684 MHz FDX

This scenario increases the FDX band to 684 MHz while Legacy Video shrinks to 144 MHz. This leaves DOCSIS with enough fixed spectrum above the FDX band to fit 16 SC-QAM channels for D3.0/D3.1 modems plus 96 MHz OFDM channel for D3.1/D4.0 modems. These channels total ~1.53 Gbps of fixed DS capacity.

Figure 33 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now dropped ~1.53 Gbps. This causes the US TG limit (i.e. TE condition #4) to kick in at 102 subs where Nsub*Tavg crosses this point. The DS average consumption starts taking away usable FDX BW from the US at a much lower point in this scenario.

If the system is upgraded to 1218 MHz, then more than 200 MHz of OFDM capacity is added to the fixed DS capacity. See the lower purple dashed line in Figure 33. This moves the crossover point for TE condition #4 up to 242 subs.

The top 3 blue dashed lines show the total DS capacity available for the various FDX DS QAM modulations (i.e. 1024-, 2048-, 4096-QAM) up to 1002 MHz. This scenario supports:

- 2G DS tier to ~265 subs @ 1024-QAM; ~327 subs @ 4096-QAM
- 3G DS tier to ~200 subs @ 1024-QAM; ~260 subs @ 4096-QAM
- 4G DS tier to ~132 subs @ 1024-QAM; ~195 subs @ 4096-QAM
- 5G DS tier to ~62 subs @ 1024-QAM; ~125 subs @ 4096-QAM




Figure 33 – Case 2 w Video, 684 MHz FDX – DS BW Capacity Limits

Figure 34 shows the DS spectrum requirements for each of the above combinations. Note – this includes both the DOCSIS DS spectrum and 144 MHz Legacy Video.

To put the DS in context, a SG with 128 subs and a 4G DS tier will stay within the fixed DS BW plus one 192 MHz FDX band for 99.5% of the time and only need some additional FDX BW for just 0.5% of the time.

The US SG limits (i.e. TE condition #2) are examined in Figure 35. The various orange curves show the total US capacity required for various US tiers (2G – 4G). Each jump in modulation effectively allows an additional 1 Gbps to be added to Tmax_max. Note that the US Nsub*Tavg finally exceeds the 85 MHz BW beyond 450 subs which is much higher than the DS.




Figure 34 – Case 2 w Video, 684 MHz FDX – DS Spectrum Requirements




Figure 35 – Case 2 w Video, 684 MHz FDX – US BW Capacity Limits




Figure 36 – Case 2 w Video, 684 MHz FDX – US Spectrum Requirements

Figure 36 shows the US spectrum requirements. The upper red dashed line starts at 684 MHz until the DS peak period consumption overflows into the FDX band at 102 subs. The drop in usable US BW above 102 subs is the result of TE condition #4. As Nsub increases, the impact becomes greater. The slope is also steeper in Case 2 with the higher Tavq.

This scenario supports:

- 2G US tier to ~210 subs @ 64-QAM; ~265 subs @ 256-QAM; 310 subs @ 1024-QAM
- 3G US tier to ~175 subs @ 256-QAM; 225 subs @ 1024-QAM
- 4G US tier to ~150 subs @ 1024-QAM

TE condition #4 has now become the limiting factor in how many subs can be supported in a SG.

To put the US in context, a SG with 256 subs and a 2G US tier will stay within 85 MHz BW outside FDX band for 92% of the time and only need some additional FDX BW for just 8% of the time.

5.3.2. Case 2 with IPTV, no Legacy Video, 684 MHz FDX Band

Migrating to 100% IPTV frees up the 144 MHz Legacy Video which helps with the previous TE condition #4 limitations. The fixed DS capacity above 684 MHz is now 16 SC-QAM channels for 3.0/3.1 modems plus 240 MHz OFDM channel for 3.1/4.0 modems. These channels total ~2.93 Gbps of fixed DS capacity, almost double the previous scenario.

Figure 37 shows the DS BW Capacity required for several different Tmax_max tiers. The second blue dashed line shows the Fixed DS capacity limit has now back up to ~3 Gbps. The US TG limit (i.e. TE condition #4) now kicks in at 195 subs.

Notice that the top 3 blue dashed lines showing the total DS capacity available has increased significantly to 7.5-8.5 Gbps. The Case 2 IPTV scenario supports:

- 2G DS tier to ~353 subs @ 1024-QAM; 400+ subs @ 4096-QAM
- 3G DS tier to ~285 subs @ 1024-QAM; ~348 subs @ 4096-QAM
- 4G DS tier to ~225 subs @ 1024-QAM; ~287 subs @ 4096-QAM
- 5G DS tier to ~155 subs @ 1024-QAM; ~218 subs @ 4096-QAM

Figure 38 shows the DS spectrum requirements for each of the above combinations. The US SG limits (i.e. TE condition #2) are examined in Figure 39. The various orange curves show the total US capacity required for various US tiers (2G – 4G). Figure 40 shows the US spectrum requirements. The Case 2 IPTV scenario supports:

- 2G US tier to ~275 subs @ 64-QAM; ~325 subs @ 256-QAM; 350 subs @ 1024-QAM
- 3G US tier to ~250 subs @ 256-QAM; ~290 subs @ 1024-QAM
- 4G US tier to ~230 subs @ 1024-QAM




Figure 37 – Case 2 w IPTV, 684 MHz FDX – DS BW Capacity Limits




Figure 38 – Case 2 w IPTV, 684 MHz FDX – DS Spectrum Requirements




Figure 39 – Case 2 w IPTV, 684 MHz FDX – US BW Capacity Limits




Figure 40 – Case 2 w IPTV, 684 MHz FDX – US Spectrum Requirements

Figure 41 shows how three different service tier combinations map to spectrum using the candlesticks from the traffic engineering formula. The top example is a SG with 256 subs and 3G DS x 3G US tiers. Notice that the 3G US burst requires more spectrum than the DS 3G burst due to its lower modulation. This is a non-typical example of when US limits the DS from growing any further.

The middle example shows 256 subs with 4Gx2G tiers. The fully symmetric would not fit. As can be seen, there is still some headroom in both US + DS for additional Tav growth. The bottom example shows 128 subs with 5Gx4G tiers with headroom.




Figure 41 – Case 2 w IPTV, 684 MHz FDX – Various Service Tier Examples

5.3.3. Probability of DS + US Overlapping Tails – Case 2

As seen in the Case 2 scenarios with increased Tavg, the DS capacity needs more and more of the FDX BW. Figure 42 looks at the Case 2 IPTV probabilities of spectrum usage for both US + DS CDF for 200 subs.

As before, the yellow curve on the left is the US ‘CDF’. For Case 2, the curve crosses the red line (85MHz) at ~3%. This means that 97% of the time, the US capacity remains completely within 85 MHz, and ~3% of the time it needs to request FDX BW.

The blue curve on the right is the DS ‘CDF’. The DS curve crosses the top of the FDX band (684MHz) at ~50%. This means that 50% of the time, the DS capacity remains completely above 684 MHz, while 50% of the time it will need to request FDX BW. A much different situation than Case 1!




Figure 42 – Probability of DS + US Overlapping Tails – Case 2 IPTV, 200 subs

To take a closer look at the overlapping tails, Figure 43 shows the same data but on a log scale. The US datapoint for 0.01% is at 160 MHz. What this means is that the US needs more than 52 MHz worth of capacity from the FDX band just 0.01% of the time. For 52 MHz worth of capacity, the scheduler could assign this as 55% of a single 96 MHz FDX channel or 9% of the entire FDX band. Even in Case 2, the US can still burst to the top of the FDX band, it is just that it is very infrequent and short bursts that shouldn't congest the system. In the rare event that the US does have a sustained burst (e.g. speed test once a month), the yellow dashed line at 615 MHz shows the upper capacity limit based on the traffic engineering formula. The US will not exceed this BW point over a sustained period.

Looking at the DS CDF, it crosses 492 MHz just above 0.1%. This means that 99.9% of the time, the DS BW needs are above this point (i.e. DS Fixed spectrum + a single 192 MHz FDX DS channel) and only 0.1% of the time does the DS need to go below this point into a 2nd &/or 3rd FDX DS channel. The blue dashed line at 135 MHz represents the lowest spectrum BW point that the DS needs based on the traffic engineering formula (e.g. a speed test on top of normal traffic).




Figure 43 – Probability of DS + US Overlapping Tails – Case 2 IPTV, 200 subs (Log Scale)

While the probabilities of the US + DS tails overlapping becomes infinitesimal, the operator still needs to engineer the system for that instance when one modem does burst to its maximum (e.g. speed test). The worst case happens when the US bursts to its maximum, which is the yellow dashed line at 615 MHz. At this point, the US is completely consuming 108-492 MHz FDX band and needs 64% of the 492-684 MHz FDX channel (and 36% is available for DS BW). The critical question then becomes how much is the DS overlapping this?

To answer that, look to see where the blue DS curve intersects with US max at 615 MHz. The DS CDF intersects at the 5% point on its curve. This means that 95% of the time, the DS stays above this BW point, and everyone has sufficient capacity. However, 5% of the time there is a BW conflict with the DS overlapping the US Tmax_max burst. But how big is it? Breaking the 5% window into segments shows:

- US burst gets 4000+ Mbps for 95% of time
- US burst gets 3840-4000 Mbps for 2% of time
- US burst gets 3615-3840 Mbps for 2% of time
- US burst gets 3450-3615 Mbps for 0.5% of time
- US burst gets 3270-3450Mbps for 0.25% of time
- US burst gets <3270 Mbps for 0.25% of time

The impact of this contention is minimal. On average, it will reduce the US 4G burst by <0.5%. Just adding a small amount of over-provisioning into Tmax_max will easily compensate for this slight shortfall.

6. Results Summary

As we noted at the outset and have shown herein, traffic analysis of the FDX band is tractable using known statistical methods, tools, and empirical data. We have gone onto significant depth explaining the statistical foundation of burst traffic, the empirical basis of these statistics, describing modeling scenarios used to create operational guidelines, and creating parametric curves that quantify the relationships among bandwidth allocated, speeds, bandwidth efficiency, and transmission group sizes.

The foundational principle enabling large transmission group sizes is that, because of the statistical characteristics of real traffic, networks are able to be oversubscribed – meaning that many customers can effectively share capacity that is a number less than the sum of their individual peak usage needs. This has been the case for decades. As the peak speeds to average utilization skew to larger ratios in the Gigabit speed era, we see that the higher the peak burst, the increasingly infrequently it occurs. Years of evidence from broadband service delivery have been used to generate probability distributions that describe the traffic behavior, and which are leveraged and extrapolated to predict the expected performance in these models, and in particular how that applies to the dual-use FDX band.

The consolidated results are shown in Table 4. Results show that large TGs can be supported without loss of QoE. Furthermore, the TG sizes align well with what are current node sizes and service groups sizes that will result from network augments that are already part of the Comcast HFC upgrade plan.

As a point of context, the average size of a fiber node today is 350 hp and getting smaller. As nodes get split during DAA network upgrades, they are roughly halved – or become approximately 175 hhp. Adjusting then for residential penetration, these numbers would represent a service group of subscribers of about 200 and 100, respectively – very well aligned to the values in Table 4. *This suggests that, with proper spectrum management, and even under aggressive CAGR, N+x FDX – including with elongated TG – will comfortably multi-gig symmetrical speeds well within today's network architecture and expectations of future augments.*

Table 4 – TG Size Summary: Spectrum Scenarios and Speeds

	Scenario	2G / 2G	3G / 3G	4G / 4G	
Near Term	1	>350	306 / X	157 / X	
	2	>400	250	157 / X	
	3	> 400	>400	>400	
Longer Term	4	310	225	150	
	5	350	290	230	

*4k-QAM DS / 1k-QAM US

5 Gbps DS @ TG Size = 218
6 Gbps DS @ TG Size ~ 125
7 Gbps DS @ TG Size ~ 60

60

7. Conclusion

Traffic engineering of the HFC network has been at the root of our success deploying HSD services cost-effectively, and without over-engineering the network. The industry has decades of successful deployment experience. Over time, operators have developed mature empirically-based statistical models of DS and US. This has translated into robust processes for operating and augmenting the network in the face of DS and US CAGR and continually increasing speeds tiers. The essential relationships of capacity supply vs peak speed and user demand have led to predictable thresholds for configuring service groups. These same learnings can be applied to traffic engineering the FDX band.

However, we have even more tools, and more powerful ones, at our disposal today to manage capacity. Utilization data can be captured and delivered to the cloud on a service-group-by-service group basis. With algorithms to crunch the localized data and traffic trends, extremely granular, and thereby more efficient, cost-effective, and targeted migration planning can take place. On top of this, because there are now decades of aggregate traffic data history, the empirical data can be used confidently to create probability distributions using classical goodness-of-fit tools of statistical analysis and pattern recognition theory. Once these are crystallized into mathematical expressions, we have the foundation to quantifiably predict burst characteristics sliced and diced into durations and size, and use these to guide business decisions. The modeling exercise described in this paper is an outcome of this process.

From a traffic engineering perspective, an FDX US population represents a small number of less aggressive “users” looking to infrequently access frequency and time resources in a specific band that is managed jointly by the vCMTS. Similar to independent DS and US criteria, thresholds can be quantified for TG size versus peak speeds (for a fixed set of assumptions on CAGR, total capacity, penetration, BW efficiency). Infrequent (statistically) peak bursts and DS Avg BW >> US Avg BW means the FDX Band’s RBAs will be deployed as DS blocks the vast majority of the time.

By contrast, if a large chunk of the coaxial spectrum such as 5-204 MHz, 5-300 MHz, 5-396 MHz, etc., is dedicated to upstream traffic, it will be *idle* the vast majority of the time. Furthermore, for the increasingly high “ultra-high split” options, while this high-quality spectrum mostly idles as an outlet for a very occasional US burst, the approach will force more downstream into the least predictable and never before activated part of the coaxial spectrum above 1 GHz, and only *after* all of the taps and passives have been upgraded.

Traffic analysis of the FDX bandwidth has shown that

- 1) The initial intuitive instinct to minimize the size of an Interference Group turns out, in practice, to be evidentially unfounded.
- 2) Large TGs can be supported while maintaining customer QoE, similar to how oversubscription models have worked for operators for decades of broadband services.
- 3) The TG sizes determined for the multi-gigabit symmetric speed tiers of interest align well with current node sizing and the expected network augments in the Comcast upgrade plan in the years ahead.
- 4) Capacity, speed, penetration, and TG relationships can be used, as they are similarly used today, to provide guidance to network operators’ network augmentation and business (speed) planning

With these findings, and with the innovative development of FDX amplifiers already in the works and showing promise, the industry can now feel confident that multi-gigabit symmetrical speeds can quickly be enabled in their existing N+x HFC deployments, founded on a deep understanding of the practical realities of burst traffic engineering and how it applies in the FDX band.

Abbreviations

ABR	Adaptive Bit Rate
BW	Bandwidth
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CDF	Cumulative Distribution Function
CM	Cable Modem
CMTS	Cable Modem Termination System
CPE	Consumer Premises Equipment
D3.1	Data Over Cable Service Interface Specification 3.1
D4.0	Data Over Cable Service Interface Specification 4.0
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DSLAM	Digital Subscriber Line Access Multiplexer
DSP	Digital Signal Processing
EC	Echo Cancellation
EOL	End of Line
EPON	Ethernet Passive Optical Network (aka GE-PON)
ESD	Extended spectrum DOCSIS
FDD	Frequency Domain Duplex
FDX	Full Duplex DOCSIS
FDX-L	Full Duplex DOCSIS Light
FTTH	Fiber to the Home
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabit per second
GHz	Gigahertz
GOA	Grey Optics Aggregator
GOT	Grey Optics Terminator
HEO	Headend Optics
HFC	Hybrid Fiber-Coax
HP	Homes passed
HSD	High Speed Data
HW	Hardware
I-CCAP	Integrated Converged Cable Access Platform
IG	Interference Group
IP	Internet Protocol
IPTV	Internet Protocol Television
LDPC	Low Density Parity Check (FEC code)
LE	Line Extender
MAC	Media Access Control
MB	Multi-port Bridger
Mbps	Megabit per second
MHz	Megahertz
MSO	Multiple System Operator

N+0	Node+0 actives
N+ x	Node + x actives (amplifiers)
NCTA	The Internet & Television Association
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OPEX	Operating Expense
PDF	Probability Density Function
PHY	Physical interface
PMF	Probability Mass Function
PON	Passive Optical Network
PSD	Power Spectral Density
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
RBA	Resource Block Assignment
RF	Radio Frequency
RPHY	Remote PHY
SC-QAM	Single Carrier QAM
SDV	Switched Digital Video
SG	Service Group
SCTE	Society of Cable Telecommunications Engineers
SNR	Signal to Noise Ratio
STB	Set-top Box
Tavg	Average bandwidth per subscriber
TCP	Total Composite Power
TG	Transmission Group
Tmax	Maximum sustained traffic rate – DOCSIS Service Flow parameter
TX	Transmit
US	Upstream
xDSL	Digital Subscriber Line, unspecific type

Bibliography & References

Technical Papers

[1] Cloonan, Dr. Tom et. al., Advanced Quality of Experience Monitoring Techniques for a New Generation of Traffic Types Carried by DOCSIS, NCTA Spring Technical Forum 2013, NCTA

[2] Cloonan, Dr. Tom, John Ulm, Ayham Al-Banna, Frank O’Keefe, and Ruth Cloonan, Capacity Planning, Traffic Engineering, And HFC Plant Evolution For The Next 25 Years, SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.

[3] Cloonan, Dr. Tom et. al., The Big Network Changes Coming with 1+ Gbps Service Environments of the Future, SCTE Cable-Tec Expo 2017.

[4] Cloonan, Dr. Tom et. al., Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements, SCTE Cable-Tec Expo 2014.

[5] Cloonan, Dr. Tom et. al., Using DOCSIS to Meet the Larger BW Demand of the 2020 Decade and Beyond, NCTA Spring Technical Forum 2016, NCTA

[6] Howald, Dr. Robert and Jon Cave (Comcast), Olakunle Ekundare (Comcast), John Williams (Charter), Matt Petersen (Charter), Developing the DOCSIS 4.0 Playbook for the Season of 10G, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[7] Howald, Dr. Robert, Roaring into the 20's with 10G, 2020 SCTE Expo, Oct 13-16 (Virtual Event).

[8] Howald, Dr. Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, What is 10G – The Technology Foundation, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

[9] Prodan, Dr. Richard, 10G Full Duplex DOCSIS Implementation Exceeds Expectations, 2021 SCTE Expo Oct 11-14 (Virtual Event).

[10] Prodan, Dr. Richard, Optimizing the 10G Transition to Full Duplex DOCSIS 4.0, SCTE Expo, Oct 13-16, 2020 (Virtual Event).

[11] Ulm, John and Zoran Maricevic, Ram Ranganathan; Broadband Capacity Growth models – will the End of Exponential Growth eliminate the need for DOCSIS 4.0?, SCTE Cable-Tec 2022, Sept 19-22, Philadelphia, PA.

[12] Ulm, John and Dr. Tom Cloonan, The Broadband Network Evolution continues – How do we get to Cable 10G?, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

[13] Ulm, John and T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning, 2014 SCTE Cable-Tec Expo

[14] Ulm, John and Zoran Maricevic, 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

[15] Ulm, John and Dr. Tom Cloonan J. Ulm, T. J. Cloonan, Traffic Engineering in a Fiber Deep Gigabit World, 2017 SCTE Expo, Oct 17-20, Denver, CO.

[16] Wall, Dr. Bill, Practical Considerations For Full Duplex Deployments In N+x Environments, 2019 SCTE Expo, Sept 30-Oct 3, New Orleans, LA.

Press Releases

[17] *Comcast Demonstrates Fastest-Yet Speeds Over a Complete 10G Connection on a Live Network*, Comcast Corporate Press Release April 28, 2022.

[18] *Comcast Announces World-First Test of 10G Modem Technology Capable of Delivering Multigigabit Speeds to Homes*, Comcast Corporate Press Release Jan 13, 2022.

[19] *Announcing Another 10G Milestone Amidst a Flurry of Innovation*, Comcast Corporate Press Release Oct 14, 2021.

[20] *Comcast Conducts Groundbreaking Test of 'Full Duplex' Chip that Will Support Future Multigigabit Upload and Download Speeds*, Comcast Corporate Press Release Apr 22, 2021.

[21] *Comcast Reaches 10G Technical Milestone Delivering 1.25 Gig Symmetrical Speeds Over a Live, All-Digital HFC Network*, Comcast Corporate Press Release Oct 8, 2020.

Other

[22] *DOCSIS 4.0 Physical Layer Specification*, CM-SP-PHYv4.0-D01-190628, CableLabs 2019

Comcast Underground: Innovative Fiber Deployments Over Existing Underground Critical Infrastructure

A Technical Paper prepared for SCTE by

Venk Mutalik

Fellow

Comcast

1800 Arch Street, Philadelphia, PA 19103

+1 (860) 262-4479

Venk_Mutalik@Comcast.com

Pat Wike, Senior Director

Doug Combs, Consulting Engineer

Alan Gardiner, Director

Dan Rice, Vice President

Table of Contents

Title	Page Number
1. Abstract	3
2. Introduction.....	3
3. What is CiC?	4
3.1. The Basics.....	4
3.1. Pervasiveness of CiC.....	6
3.1. The Fiber Journey ... Underground and in the Air	7
4. CiC Parameters.....	8
4.1. Linear and Area Fill Ratios.....	8
4.1. Micro-Fiber Details	9
4.2. Fiber Strength and Micro-ducting.....	9
5. CiC in a Green Sandbox – The Control Trial	11
5.1. Basics of CiC Deployment	11
5.2. Identified Challenges.....	13
5.1. The Power Touch vs. the Human Touch	13
6. The CiC Trial out West.....	13
6.1. Locating the Trial.....	14
6.2. Dividing the Trial	14
6.1. Trial Highlights.....	17
6.2. Operational Considerations.....	18
6.3. Fiber Portfolio	18
7. Conclusions.....	18
Abbreviations	19
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1 – Comcast Network at a Glance	3
Figure 2 – Cable in Conduit and Direct Bury Cable [2] (Comcast)	5
Figure 3 – Illustrating pedestals in the neighborhood	6
Figure 4 – Illustrating the West Division Plant Composition	7
Figure 5 – Illustrating CiC Linear and Area Fill Ratios for various cables	9
Figure 6 – Controlled trial: Day 1, 2 and 3	11
Figure 7 – Illustrating the Rod-and-Rope process and blowing fiber into the installed micro-duct.....	12
Figure 8 – Finished micro-fiber and micro-duct installed	12
Figure 9 – Location #1 with “Left-Outside” and “Down-the-Middle”	14
Figure 10 - Location #2 Real World Comcast trial matching the Control Trial.....	15
Figure 11 – Location # 2 Left: Blowing of fiber, Right: Conduit spliced in ped for blowing continuity.....	16
Figure 12 – Location #1 down the middle direct pull of micro fiber with 600 pound pull strength	17

List of Tables

Title	Page Number
Table 1 – Table of Basic Attributes and Ratios [5,6,7].....	10

1. Abstract

London Underground or the ‘Tube’ began as a modest steam rail system over 150 years back and is now a sprawling transit system transporting over 5 million people daily. Similarly, Comcast began modestly but today has a large and growing optical footprint with fiber getting deeper into the network often leading to challenging underground fiber construction in the neighborhoods. Just as generations of Tube engineers innovated on their predecessors’ plans deep underground and grew their rail network, Comcasters innovate on critical infrastructure built by our cable predecessors and provide higher capacity to match current demands and future needs.

In this paper, we report on the use of innovative technology that enables us to use existing underground critical infrastructure and make fiber deployments in the neighborhoods simple, cost effective and minimally customer impacting. Cable companies have been laying underground RF cables inside conduits, in vast sections of cable builds since the mid 1990s. This process, called Cable in Conduit (CiC) has a fraction of the conduit occupied by the RF cable with a contiguous empty space in the conduit. Recently, fiber manufacturers have come out with ‘micro-fiber’ cable that bundles of up to 72 optical fibers occupying a diameter of only a few millimeters. This new fiber cable bundle is supple, affords tight bend radius and has good tensile strength. With existing rod-rope-pull equipment this micro-fiber can now be deployed within the existing conduit alongside the RF Cable cost-effectively, quickly and with minimal impact on the customer experience, all without the need for trenching or boring.

The paper describes details of trial activities on CiC in one of our divisions and the economics of this technology. Skillful use of this technology and innovative optical systems being developed bring fiber to the last active, simplify other architectures such as Full Duplex (FDX), improve performance and capacity overall and help propel new optical architectures such as Switch on a Pole/Pedestal (SOAP). Since this technology provides large fiber counts at RF tap locations very close to our customers, it provides great long-term opportunities to span the last few meters and reach customer homes (FTTH) when needed.

2. Introduction

As the largest broadband company in the US, Comcast serves millions of customers and businesses coast to coast. All of this is the result of a large optical network that spans core, metro and access layers as illustrated below [1].




Figure 1 – Comcast Network at a Glance

While the Core and Metro layers are all-optical circuits interconnected by ROADMs fed by large routers, the access layer comprises hybrid fiber and coaxial (HFC) cable network, some fiber to the home (FTTH) all optical networks and wavelength specific all optical very high-speed commercial connections. While the HFC and FTTH networks serve primarily residential customers and small-medium business customers, the commercial optical networks primarily serve 1, 10 or 100Gbps commercial enterprises with Metro Ethernet and cell tower back haul and 5G front-mid and backhaul links. Very often the HFC, FTTH and Commercial services are carried in the same fiber sheath, and over time could share the same access fiber. Over the years, Comcast has innovated on its access plant and through node splits and network expansion, driven fiber deeper into the network. In this context Cable in Conduit could be a valuable additional tool assisting in the fiber journey.

3. What is CiC?

At this point, Comcast has an almost equal share of aerial and underground plant overall. Due to varied plant practices many regions out West have a predominantly underground plant while in the Northeast a predominantly aerial plant exists, although there is an admixture of underground (UG) and aerial (AR) per node across the country. When new fiber is deployed, AR plant can more easily be converted to fiber since the infrastructure of overlaying fiber over coax cable is fairly well known and there is a high degree of infrastructure reuse. Such is today not the case however when underground fiber has to be deployed. Deploying UG fiber is a cumbersome process that requires digging up the streets or front and back lawns and is generally much more time consuming and much more expensive than of AR fiber deployment.

While both AR and UG plant require permitting and traffic management the nature of AR permits are of weight studies of additional cables being strung, but of underground are much more complicated and time consuming due to the need of non-interference between cable, telephone, power and natural gas infrastructure that is also buried below ground.

And so, it would seem that if a solution could be found for reuse of UG cable infrastructure, both the cost of fiber deployment as well as the time of fiber deployment would come down significantly. Furthermore, if such a solution were easy to deploy and in use in some way, it would immensely aid in pushing fiber deeper and accelerate our fiber journey by providing an additional high-speed lane. Such is the case with a technique that Comcast has started using called Cable in Conduit (CiC) also called sometimes internally called fiber override in the neighborhood.

3.1. The Basics

In a large portion of the Comcast UG cable plant West of the Mississippi built after around 1985~1990, the RF cable was laid inside of a conduit, and the RF cable then surfaces to pedestals containing amplifiers and nodes, hence our name for this approach “Cable in Conduit” (CiC). While the prevalence of CiC is high in the West due to the already high UG plant there, there are areas in Central and the Northeast divisions that are also CiC based. But a fair amount of Cable plant and particularly some of the older plant across the country is what is directly buried (DB) several feet under the ground and the RF Cable surfaces to taps and pedestals. Accordingly, we have three types of plant AR, CiC and DB that together describe the total Comcast RF plant today.




Figure 2 – Cable in Conduit and Direct Bury Cable [2] (Comcast)

The picture above illustrates CiC and DB underground plants. Here, we have shown DB as a larger hardline cable, while CiC a slimmer hardline cable contained within a conduit. With the above illustration, it is easy to see that the skinnier the cable and larger the conduit, the easier it might have been to sneak the cable into the conduit and surface it at required intervals to service pedestals that hold taps, amplifiers and passives.




Figure 3 – Illustrating pedestals in the neighborhood

The picture above illustrates various pedestals and enclosures that interconnect RF cables. These are roughly about 100ft to 250ft in distance. The pedestal on the left is the tap enclosure typically in the front lawn of homes, a San Pellegrino can is shown next to it for an indication of its size. A handhole from the local telephone utility is shown along-side it, sometimes Comcast itself may have handholes nearby as well. The middle pedestal is a bullet type enclosure that sometimes may contain RF amplifiers, a fire hydrant (in need of some paint work) is shown along-side of it for size comparison. The one on the right is a pedestal that might hold a node. In each case if CiC is the mode of deployment, the cable comes out of the conduit from each side and connected at two ends to appropriate devices.

3.1. Pervasiveness of CiC

In previous sections, how pervasive CiC could be in Comcast plant. To get a feel for it in real terms, we elected to check this out across the West division. We analyzed all the nodes in West and looked at their plant composition. This would include all hardline cabling in AR and UG plants, but NOT any of the drop cables. Drop cables connect homes to tap ports on the RF plant and are of varying lengths, but crucially these could be AR, DB or CiC.




Figure 4 – Illustrating the West Division Plant Composition

Comcast plant is a conglomeration of many acquisitions over the years and the relentless growth fueled by node splits and plant extensions. While documentation of hardline fiber plant necessarily follows the original plant, plant extensions generally follow modern guidelines. Presented above is the plant composition of the entire West division. One can see that the AR plant is only around 38% of the total while the remaining 62% of the plant is in UG. Around 25% of the total plant is clearly marked as CiC, but 37% of the plant is a mixture of CiC and Direct Bury, with no easy way of demarcating the builds. This is somewhat an artefact of the aforementioned acquisitions that results in a loss of clarity.

Still the results are quite revealing. In general, labor costs for fiber construction dominate over the cost of individual fiber costs. And UG fiber construction cost per foot is about 10 times the cost of AR construction. This is not surprising, as already mentioned, the labor cost of trenching and boring along with the more involved permitting and traffic control costs dominate over the more AR plant construction costs. So in that context, the total construction cost here could have been $0.38x + 0.62 \times 10x = 6.58x$. As will be shown later, the CiC process brings down the cost of construction in UG plant by around 7 to 10 times less than current UG deployment and bringing the CiC costs closer to AR deployments. In this context, the total cost of construction would have been $0.38x + 0.25x + 0.37 \times 10x = 4.33x$. This is a 34% reduction in construction cost overall! For construction budgets running into billions of dollars and spread out over years, these types of savings are quite impressive.

We would also stress again the importance of the time to upgrade types savings - using an existing duct infrastructure as compared to new trenching/directional boring construction in underground areas, and would also note that reentering the existing duct requires fewer permits, less traffic control and less restoration to the areas resulting in lower downtime and better customer experience.

3.1. The Fiber Journey ... Underground and in the Air

In previous papers [3,4] we have discussed that a move towards all fiber network across the country is a journey with multiple rest stops and not a single one-off event. And this is more so because of many

interesting developments in technology that enable the industry to serve customer bandwidth needs without an exclusive move to all-fiber solutions at once. Recent 10G industry moves help unlock RF cable potential via efficient and bi-directional use of RF spectrum and are of critical importance to our customers.

A point to note here is that each extra step in driving fiber deeper towards the customer has a force multiplying effect. Driving fiber to a node (which is N+x or N+0) is easier than driving fiber to the last active (FTLA), which in turn is easier than driving fiber to the curb (FTTC), which is easier than driving fiber to the home (FTTH).

By way of clarification, while N+0 and FTLA both have no actives save the nodes in the plant, but a crucial difference between the two is that N+0 reimagines the RF plant and with optimum node placement and RF modifications achieves the elimination of actives. In the FTLA however, the entire RF plant remains as before including the node and amplifier locations. By connecting up the nodes and each amplifier via optical fibers, each of the amplifiers is upgraded to a node (or a mini-node) and serves the existing homes attached to the said amplifier. FTTC entails running fiber to the current tap location and terminating it in micro-nodes and using existing RF drop cables to home, whereas FTTH requires an all-fiber circuit to the home and terminating it in analog or digital customer premise equipment (CPE).

As mentioned before, Comcast plant is today AR and UG, so a move towards deeper fiber should accommodate both plant types, else the end result cannot be accomplished. So, there is a need to optimize UG construction so that the entire process of fiber deployment become simple cost effective and predictable.

4. CiC Parameters

In this section, we define some important parameters that explain the CiC and discuss the ways in which CiC is implemented. In the picture below, the orange ring represents the conduit. The black circle on the left indicates the RF cable ensconced in the conduit when the UG plant was laid out originally. One can see here that there is an empty space in the conduit not utilized by the RF cable that might be big enough to accommodate an extra cable comprising optical fibers.

For many years, on short sections of conduit, folks sometimes SST fiber bundles. These SSTs were rectangular shaped stiff fiber build that held just 12 fibers. The peculiar geometry of the bundle and very limited fiber counts made it difficult and less attractive to consider this for a wider deployment. Recently, there has been a slew of development in so called micro-fibers that enable up to 72 count fibers in a diameter of just ~4.5mm! The middle picture above shows how such a fiber would look relative to the RF Cable. To compare these micro-fibers to what is generally used in AR plant the armored cable has a 48 count fiber with a total diameter of ~11.7mm. As can be seen, the wider the conduit and smaller the cables, the easier it is to accommodate within the conduit.

Of course, Comcast plant changes from place to place, but for purposes of discussion, we have considered here a 0.625in Coax in 1-1/4 Conduit for a good portion of our analysis and trial. For these conditions, the conduit inner diameter is ~1.4in, while the Coax outer diameter is ~0.85in.

4.1. Linear and Area Fill Ratios

In the left picture, the ratio of RF cable outer diameter to the conduit inner diameter is called the linear fill ratio (LFR) while the ratio of the respective cross-sectional areas is the area fill ratio (AFR). In the above example, the LFR is ~61%, while the AFR is ~37%. What this means is that with the cable in the conduit, the conduit still has an empty space that can accommodate an appropriately small fiber cable inside of

itself in addition to the existing RF cable. The smaller the LFR and AFR, the easier it is to get additional cables in. As the LFR and AFR grow, the accommodating space becomes smaller and friction and geometry start acting up and limit the addition of additional cables into the conduit.




Figure 5 – Illustrating CiC Linear and Area Fill Ratios for various cables

Consider the right picture above, where we are looking the typical conduit trying to accommodate an RF Cable along with a ~11.7mm armored 48 count cable. The LFR here which is the sum of the two cable diameters relative to the conduit inner diameter is 94%, while the AFR which here is the ration of the sum of the cross-sectional areas of the two cables relative to the inner cross-sectional area of the conduit is 48%. So, while there is a bunch of space (52%) in the conduit, it is not possible to pull or push the armored cable thru the conduit while also the RF cable rides the same conduit. The friction of the two cables is too much, some of it with the conduit walls, others thru the RF Cable itself. If on the other hand, one considers a ~5mm micro cable, the LFR is 76% and the AFR is 39%, which is considerably better

4.1. Micro-Fiber Details

Recent developments in micro-fiber technology have enables CiC. Previously only SST fibers accomodating just 12 fibers were in use, today these cables can accomodate 72 or mode fibers. Typically these fibers have a core strength memembr in the middle and 6 tubes arranged around it. Each tube contains 12 SMF fibers. Therefore a typical micro-fiber can have upto 72 fibers. There are options that may provide more or less fibers depending upon the geometry and population of fibers in the cable. Many of the micro-fibers have a rip cords on the side that enable easier peeling of the cable to expose required fibers. The entire 72 fibers with 200um buffer along with all cabling has an outer diamter of just 4.5mm and is an impressive achievement. In case of the standard 250um buffer with all cladding requires a 5.5mm buffer. Both these fibers have a 200lbs/ft pull and crush strength. Another type of 72 count fiber cable with tensile and pull strength that rivals traditional 48 count armored cable at 600lbs/ft is now available with 9.1mm diameter as compared to the 11.7mm of the standard cable.

4.2. Fiber Strength and Micro-ducting

It should be easy to see from the above section that lower fill ratios help in CiC deployments, but in broad terms, the fibers are pulled thru the conduit. Such pulling of micro fibers could result in fiber breakage in which case it will reset the entire CiC process on a bad path since it will end up needing a large amount of fiber splices in tight spots and significantly increase deployment time. So to prevent that we select lower

fill ratios, but also fibers which have a sufficient pull or tensile strength. A fiber of 600lbs/ft of pull strength should be able to handle most of the pulling encountered in the CiC deployment. In addition to the pull strength, we will also need to have good crush strength. This crush strength is super important for AR deployment generally so that the fiber can handle the periodic wiring holds on the fiber. But crush strength is important even in CiC in case where the CiC integrity is less than optimal or in case of tight bends and dents that might form over time. Typical crush strength ranges from 200 to 600 lbs/ft. The armored cable for example has a crush strength of 600lbs/ft, while the micro-fiber has a 200lbs/ft.

This is where the idea of micro-duct comes as an additional tool in simplifying CiC. Micro-ducts which are typically 8mm outer diameter and 6mm inner diameter are a bigger than the micro-fibers with LFR of 84% and AFR of 42% and with the same tensile strength as the micro-fiber. At first glance it appears to be not that great of a bargain in using micro-duct with its higher FRs without any consequent increase in tensile strength. But the main reason this is so useful is that the micro-duct does not have any fiber of its own, so pulling the micro-duct and having it break while inconvenient is not a show stopper. One could repair the micro-duct and continue one with micro duct deployment. For this reason, one can tolerate a higher FR in the case of a micro-duct. Once the duct is installed, the 4.5mm micro-fiber itself may be very easily blown in, this time without any extraordinary effort and risk of fiber breakage in installation. Incidentally, one could use the arrive at the FRs for a micro-fiber in a micro-duct which itself is inside a conduit with an RF cable in it already. These FRs for the micro-duct referred to above are an LFR of 75% and AFR of 56%.

Table 1 – Table of Basic Attributes and Ratios [5,6,7]

Construction Description	ID Conduit/ Duct (in)	OD 0.625 in RF Cable (in)	OD micro-fiber / micro-duct (in)	LFR micro-fiber / micro-duct	AFR micro-fiber / micro-duct	Pull Strength micro-fiber / micro-duct (lbs)	Crush Strength micro-fiber / micro-duct (lbs/in)	Bend Radius micro-fiber / micro-duct (in)	Weight micro-fiber / micro-duct (lbs/ft)
GC Basic with just RF Cable	1.4	0.85		61%	37%				
+ 11.7mm 48ct armored fiber (for comparison only)	1.4	0.85	0.46	94%	48%	600	125	4.6	80
+ Rodder (for initial install support only)	1.4	0.85	0.38	88%	44%				
+ 8.5mm/6.0mm micro-duct	1.4	0.85	0.33	84%	42%	100	125	3.3	18
+ mXT 4.5mm 72ct micro-fiber in micro-duct						200	30	3.3	30
+ mXT 5.4mm 72ct micro-fiber in micro-duct						200	30	3.3	35
+ LMHD 9.1mm 72ct micro-fiber Only	1.4	0.85	0.36	86%	43%	600	125	5.0	44

Above, we have summarized several options (we thank Duraline, Corning and AFL) and FRs for each of the options, again the smaller the FRs, the better is the outcome for CiC. It is however critical to remember that there are many more conduits of various diameters and different micro-ducts and micro-fibers available and a table including many more possible combinations might be needed as we proceed more into CiC. Note here that although all the optical fibers illustrated here are from Corning, the are micro-fiber cables and micro-ducts could be from multiple manufacturers.

Although standard techniques exist for AR construction, the use of micro-fibers and micro-ducts can still be extended there if needs be. For starters, micro-fibers weigh a lot less than traditional armored fibers, so with sufficient strength micro-fibers, either by themselves or in micro-ducts these fibers may find spots in AR construction. We see from the above table that traditional armored cable suitable for AR lash might weigh ~80lbs/kft but an equivalent micro-fiber or micro-duct and micro-fiber combination might be just half of the equivalent weight. This reduced weight will also be useful while seeking permits for AR construction.

As an aside, ‘squirrel chew’ the issue of pesky rodents determined to sharpen their teeth is a well known issue in construction. The ability to deter rodents will be a prime consideration ! Armored cable has long

been important for this effort, but so also the use of non-toxic bittering agents in the fiber cabling and micro-duct construction to deter squirrel chew would be helpful.

5. CiC in a Green Sandbox – The Control Trial

With the above understanding, Comcast and our partners decided to try CiC out in an outdoors trial location. Our microfiber partner Corning and micro-duct partner Duraline together elected to test out the CiC concept at the Corning Green Acres facility. Our many thanks to the partners and their dedication to see this work amongst the pandemic restrictions. This is an impressive outdoor underground plant laid out in a grid fashion over several acres. This facility has handholes with conduits connecting them up, and several of these conduits are the 1-1/4 sized. In addition, several of these conduits also have cables that are similar to the 0.625in RF Cable. As such a trial in this location could mimic Comcast plant.

Presented below is a grid diagram of the plant we had, on the left is the way we began with CiC on the first day, across 5 sections of plant of various lengths, the second day is represented by the middle picture where we installed CiC in 3 sections and on the third day we installed 2 sections with CiC. Each day we tested different concepts, and cumulatively we had tested our ability to install a single micro-fiber, two micro fibers and micro fiber installation in a micro-duct. At the same time important questions about slack, bend radius were answered, as we questions about the total thruput of CiC per day with a crew of 4 installers. All installations were manual installations, and towards the end of the trial, several mechanized versions of installations were discussed. These are described in detail next.

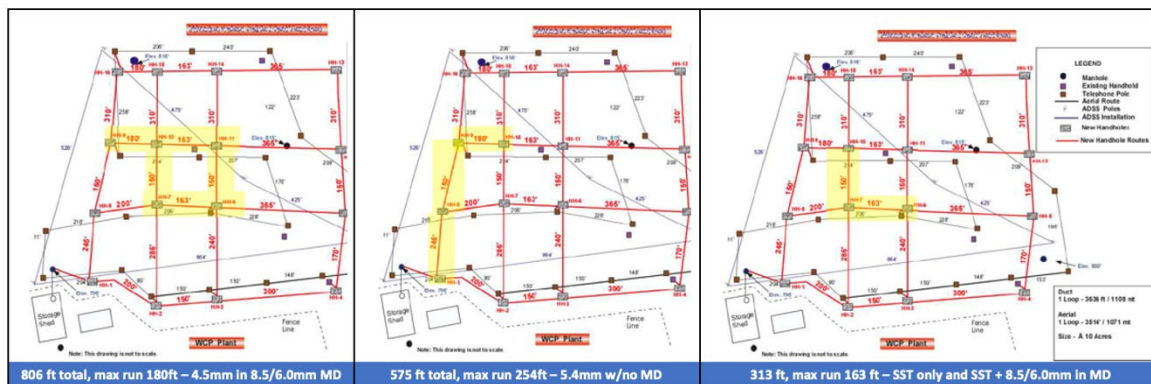


Figure 6 – Controlled trial: Day 1, 2 and 3

5.1. Basics of CiC Deployment

The basics of CiC deployment begin with selection of the RF Cable present and micro-fiber or micro-duct to be installed based on the FRs described earlier. Once that is done, a strong fiber glass “Rodder” of a diameter bigger than the micro-fiber or the micro-duct is inserted in the conduit. A nard bullet is affixed to the start of the rodder and is then pushed into the conduit until it surfaces at the other end of the conduit which may be 125 to 200 ft away. In doing so, the RF cable has been gently pushed aside and a continuous passageway has been opened inside the conduit.

At the surfaced end of the conduit, a cable pulling sock is attached to the rodder and the other end of the sock grabs on to the micro-duct or micro-fiber. The rodder is then pulled back the same way it went in, but at the other end, the micro-fiber or the micro-duct is then surfaced and thus the first phase of installation completed. This installation process is called the rod-and-rope method. During the many installations of the CiC, the Corning Green Acres trial clearly showed that re-entering the CIC using the rod-and-rope method was a viable solution.



Figure 7 – Illustrating the Rod-and-Rope process and blowing fiber into the installed micro-duct

The fiber or the duct is slacked and the next section is then begun. In practice one could go rather long distance fairly quickly if regular opportunity for surfacing the rodder are available. Notice that in this whole process, we never really had any reason to dig up the ground. Once the requisite length is reached, and if a micro-duct is used the fiber is simply blown in for the whole length, thus completing the process.



Figure 8 – Finished micro-fiber and micro-duct installed

The figure above shows the installed micro-fiber on the left and the installed micro-duct on the right.

5.2. Identified Challenges

In general, factors to consider for successful CiC include distance of the conduit system and the spacing between the handholes, the age of the system, condition of the conduit and blocks or damages, condition of the pedestal and space available around, terrain of the build and its proneness to rocks, elevation and ice. While none of these on their own are showstoppers, it is wise to prepare countermeasures ahead of time.

Most of the rodders also have a tonal strip, one that enables the exact path of the rodder to be known from above ground using tone detection equipment. So if the rodder is stuck or unable to proceed further, the exact spot of may be dug up and the conduit unblocked and the procedure continued. This is a way to precisely dig up only a small spot and limit impact. In our trials and tests we did not encounter this specific obstacle.

5.1. The Power Touch vs. the Human Touch

With a crew of 4 we spanned around 1700 ft of CiC spread over 10 sections, yielding an average of 170ft/section although, there was a section that was 245ft long. Based on these we estimate that a 200ft of CiC could be spanned by a crew of 4 within 20-65 minutes depending upon specific challenges. Rodding could be accomplished between a minimum of 10 to 40 minutes depending upon the conduit, pulling back the fiber our duct could be between 5 to 20 minutes and creating a fiber slack before proceeding to the next rod-and-rope could be about 5 minutes and if a micro-duct was used, blowing the fiber into it would be less than 5 minutes, but this last process is done after the micro-duct is installed thru all the sections.

We note here that if fiber is to be taken out of the micro-fiber cable, appropriate splice enclosures that can handle splices and the associated slack must be considered. In addition a good fiber management strategy should be adopted, one where tube colors and individual fiber colors in the tubes must be used as identifiers. The time associated with that process is part of node/network commissioning and not included here.

With this in mind a 4700-5000 ft of CiC could potentially be installed in one day with an 8 member crew. This is an important observation as the regular plant for an average sized node is about a mile (or 5200ft), and if it were in a planned development then that node could be fiberized to FTTC within a day in ideal conditions. Note here that there is no disruption of services as the fiber installation procedure has no impact on RF and power signals running on the cables.

There are power tools that could help the rodding and pulling process and all such equipment along with the fiber blower could be accommodated on a standard pickup truck, thus improving mobility and alleviating traffic concerns.

6. The CiC Trial out West

A decision was made to take the learnings gained from controlled testing that already had occurred and deploy it in an existing network. The market in the Denver area and the home of Comcast West Division office was selected as a suitable location. This market was selected as it has had a planned community with continual growth each year over the last thirty plus years. An important factor in performing this trial in the real world is determining where CiC has been deployed if mapping information did not capture conduit usage.

6.1. Locating the Trial

We were able to check neighborhoods by age of houses and using local knowledge of when CiC began widespread usage. Although CiC had been introduced in the 1980's the adoption amongst the multiple MSOs that existed at that time varied. There were variations of adoptions within the geographic areas of MSOs as well that affect what areas can be targeted for fiber override of existing conduit. After a couple attempts, we found consistent CiC usage in neighborhoods built after the year 2000. We did not go further in this area to narrow down what year CiC became prevalent but that information can be useful for particular geographic areas.

6.2. Dividing the Trial

For the field trial we decided to expand slightly over the controlled trial and push the application to gain additional data. Three different approaches were taken that are shown in the two pictures below that would give us additional data.

Location #1 what we called the outside left was an express run which would push the distance between each pedestal. Our plan here was to use the microduct being placed by rod and rope technique to establish the path for a final blow in of fiber.

Location #1 down the middle met the spacing consistent with what had been done in the controlled trial but here we would use a micro fiber with traditional six hundred pound pull strength. For deployment we would use rod and rope technique with no conduit, directly pulling the fiber in over the existing conduit.

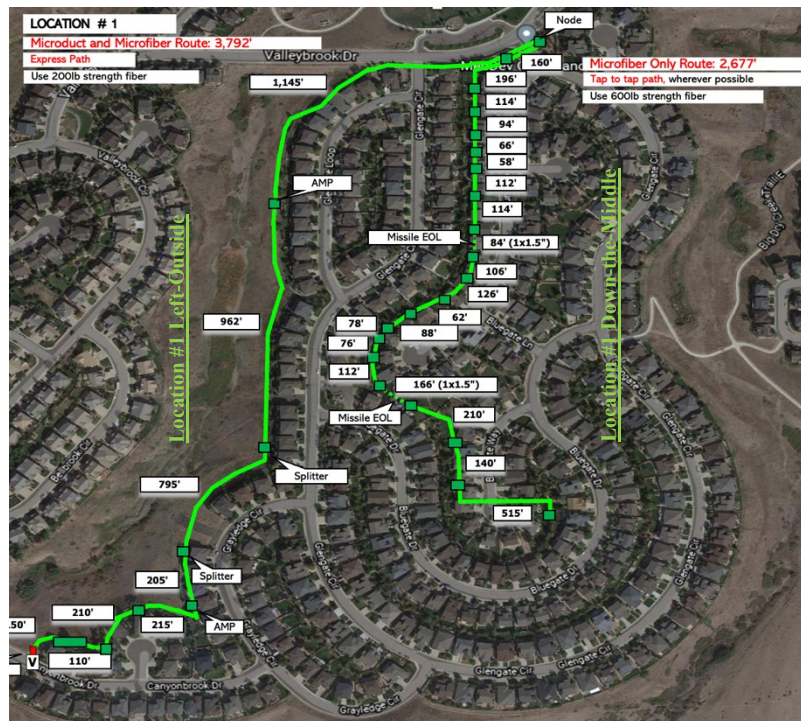


Figure 9 – Location #1 with “Left-Outside” and “Down-the-Middle”

Location #2 most fit what had been done in the controlled area and this was used to validate the lessons from the controlled trial. Pedestal distances were what we considered normal, the rod and rope technique was used to pull in micro duct to prepare for blowing in of fiber.



Figure 10 - Location #2 Real World Comcast trial matching the Control Trial

In each case, appropriate amounts of slack were rolled in into as part of construction. An important part of the trial was also the ability to identify fibers apart from RF cables. While this might look like a trivial part of the trial, getting this right is important to prevent needless fiber cuts by well-meaning techs out and about as they troubleshoot the RF plant. Some identification techniques work better than others and were incorporated in the builds.



Figure 11 – Location # 2 Left: Blowing of fiber, Right: Conduit spliced in ped for blowing continuity




Figure 12 – Location #1 down the middle direct pull of micro fiber with 600 pound pull strength

6.1. Trial Highlights

- No damaged conduit encountered in 7,924ft of CiC deployment
- Max length of conduit that can functionally be overridden is around 250ft (before rising above ground) this was determined in location #1 outside left. For distances longer than 250ft Rodder either became stuck or the micro-duct broke with too much pull force exerted on it
- 600-pound pull strength fiber is handled like BAU fiber today does not require any micro duct placement. We did experience scrapping of jacket making footage readings difficult from the jacket. Note that the 600-pound pull strength can be lashed in aerial plant like existing fiber (but this cable does not have armor)
- 200-pound pull strength fiber will require micro duct placement. Business Partners require additional skills and installation equipment not common to CATV construction. 200-pound pull strength cannot be placed aerially without micro duct, requires additional processes
- Both types of fiber cables can be successfully placed but must be operationalized for proper deployment
- Preliminary indications confirmed the significant reduction in cost indicated earlier relative to regular underground trench and bore in this location. Therefore cost savings can be substantial and time to deployment can be substantially decreased as well

6.2. Operational Considerations

- A new Statement of Work would likely be required to operationalize CiC across the footprint. This is because the amount of fiber that can be deployed and the speed of deployment are both different than for standard UG construction. From a strategic point of view, a new SOW focusing on CiC would also reduce conflict of interests within the builder community and help focus on CiC when that is better or regular construction when that is the only option available
- If smaller pedestals are in deployment, the CiC construction process and slack preparation would be longer and potentially impact construction costs. At this trial we encountered a number of smaller pedestals and considered upgrading them along the way
- Storage length and placement in each pedestal needs to be specified and documented. Identification of fiber vs. Coax needs to be vividly documented. Appropriate training material for maintenance and fulfillment teams would need to be developed

The technique of CiC fiber override has been introduced within Comcast as an option for local construction groups to use. This does require them to work with their business partner (contractor) to ensure they are properly prepared to execute this technique with adequately trained staff. Fiber override is not a one size fits all but the trial has proven that this is a viable technique which becomes another arrow in the quiver for the construction crews to use.

6.3. Fiber Portfolio

In other papers in this conference, we have presented on Hollow Core Fibers [8], this is a new type of fiber that enables light to be guided in a hollow core as opposed to the standard light being guided in solid core fibers. As light travels much faster (300,000km/s) in air than in glass (200,000km/s), there is a fundamentally large reduction in latency. This helps in important latency sensitive applications such as high frequency trading and 5G. In this context, CiC could play a major role in helping roll out fiber in UG plant. Recall that most micro-fibers have 6 tubes within with 12 fibers each in each tube. In this case, one or more tubes could be dedicated to hollow core fibers while others are for standard fiber and a portfolio of fibers may be installed when the conduit is re-entered. This type of innovative deployments may help overall to bring fiber to the neighborhood where high-capacity latency sensitive endpoints may be located.

7. Conclusions

In this paper, we reported on the use of innovative technology that enables us to use existing underground critical infrastructure and make fiber deployments in the neighborhoods simple, cost effective and minimally customer impacting. With innovations in fiber cabling and availability of higher count fibers, the ability to reuse existing conduit infrastructure opens up quicker ways of deploying fiber, all without the need for trenching or boring. As our own trials show, skillful use of this technology when combined with innovative optical systems could bring fiber to the curb (FTTC) and support the industry's 10G efforts and provide great long-term opportunities to span the last few meters and reach customer homes (FTTH) when needed.

Abbreviations

AFR	Area Fill Ratio
AR	Aerial
bps	bits per second
CiC	Cable in Conduit
DB	Direct Bury
FDX	Full Duplex
FEC	forward error correction
FTLA	Fiber to the Last Active
FTTC	Fiber to the Curb
FTTH	Fiber to the Home
HCF	Hollow core fiber
Hz	hertz
K	kelvin
LFR	Linear Fill Ratio
SCTE	Society of Cable Telecommunications Engineers
SOAP	Switch on a Pole
UG	Underground

Bibliography & References

- [1] Photon Avatars in the Comcast Cosmos: An End-to-End View of Comcast Core, Metro and Access Networks, Venk Mutalik, Steve Rupp, Fred Bartholf, Bob Gaydos, Steve Surdam, Amarildo Vieira, Dan Rice
- [2] CommScope Construction Manual – Figure used with Permission
- [3] The Yin and the Yang of a Move to All Fiber: Transforming HFC to an All Fiber Network While Leveraging the Deployed HFC Assets, Venk Mutalik, Marcel Schemmann, Zoran Maricevic, John Ulm; INTX 2015 Spring Technical Forum
- [4] Cable's Success is in its DNA: Designing Next Generation Fiber Deep Networks with Distributed Node Architecture, Venk Mutalik and Zoran Maricevic, SCTE/ISBE 2016
- [5] <https://www.duraline.com/micro-technology/MicroDucts%20HDPE-462>
- [6] <https://www.corning.com/optical-communications/worldwide/en/home/products/minixtend.html>
- [7] <https://www.aflglobal.com/Products/Fiber-Optic-Cable/Structured-Cabling/Outside-Plant/LMHD-Series-OSP-Heavy-Duty-MicroCore-Cable>
- [8] Approaching the Universal Speed Limit: Introducing Hollow Core Fibers for Low Latency and High Capacity, Venk Mutalik, Amarildo Vieira, Bob Gaydos, Elad Nafshi

Commercial Network Services Lifecycle Management

A Technical Paper prepared for SCTE by

Pattabi Ayyasami

Sr Principal Engineer
Comcast

1050 Enterprise Way #100, Sunnyvale, CA
+1 (408) 212-4383
pattabi_ayyasami@comcast.com

Tirumalesh Ramaiah Reddy

Director, Software Development & Engineering
Comcast

1800 Arch St, Philadelphia, PA
+1 (215) 286-5073
tirumalesh_ramaiahreddy@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Overview	3
3. Service Model (TOSCA).....	3
3.1. Why TOSCA.....	4
3.2. TOSCA Features.....	5
3.3. Use of TOSCA and Yang	5
3.4. TOSCA Meta Model Phases	6
3.5. TOSCA Object Model.....	6
4. Service Orchestration.....	7
4.1. Challenges with Service Life Cycle Management.....	7
4.2. Platform Goals.....	7
4.3. Service and Function Model Definition (Design).....	8
4.4. Service and Function Onboarding	8
4.5. Orchestration (Runtime).....	8
5. UseCases.....	9
5.1. SD-WAN.....	9
5.2. Firewall.....	10
6. Conclusion.....	11
Abbreviations	12
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 - What is TOSCA	4
Figure 2 - TOSCA Meta Model Phases	6
Figure 3 - TOSCA Object Model	6
Figure 4 – Functional Architecture of Service Orchestrator.....	8
Figure 5 – SDWAN VPN Site Service (Sample)	9
Figure 6 - Firewall Vendor Abstraction.....	11
Figure 7 - Firewall Vendor Substitution.....	11

List of Tables

Title	Page Number
Table 1 – Declarative versus Imperative Architecture	4

1. Introduction

As Software Defined Networking (SDN) and Network Functions Virtualization (NFV) have become modern approaches to quickly deliver and offer network services such as SD-WAN, Security and other value-added services using service chains. Declarative & model driven orchestration technologies like TOSCA has become the fabric to develop Network Automation platform to manage lifecycle of these network service applications and build an ecosystem for multiple vendors to participate & integrate with their VNF solutions.

Comcast architecture leverages several information models and orchestration tools including TOSCA and YANG to enable complete life cycle management of SDN and NFV functionalities and to configure physical and virtual devices that comprise an end-to-end service.

The paper describes how TOSCA is being used to describe orchestration between resources across complex end to end service(s) such as SD-WAN and Cloud Security. Demonstrate the capabilities of TOSCA like substitution/service decomposition for vendor abstraction to realize some of the VNF's, interface and life cycle operations used for the implementation aspects, requirement, capability types, node referencing features, define and visualize end-to-end service as a topology graph of inter-connected nodes and to navigate across connected services.

YANG is a data modeling language used to describe the device configuration. Device/network configuration properties are modeled as YANG. The syntax and semantic constraints and capabilities offered by YANG are made use of during service validation and activation.

2. Overview

At a high level, for Comcast use cases, Service life cycle management comprise of one or more of the following.

- 1) Staging/provisioning the base configuration for the network devices
- 2) Amend changes to the configuration.
- 3) View service configurations
- 4) Updates service configuration.
- 5) Add/Remove features to the services
- 6) Upgrade/Downgrade of services
- 7) Monitoring the health of services and corrective actions in case of service degrade and/or disruption.

This paper will focus on the model-based service management platform. The platform provides APIs for some of the service life cycle management functionalities. The platform is a middleware layer that interacts with OSS/BSS systems northbound and Network Layer, VNF Managers, Vendor specific EMS/NMS systems southbound.

The paper highlights how TOSCA model and Service Orchestrator enables to support service life cycle management functionalities for different domains such as SD-WAN, Security, Wireless etc.

3. Service Model (TOSCA)

OASIS TOSCA is an orchestration language for automating Service Lifecycle Management. Service designers use TOSCA to create Service Templates that contain, model-based descriptions of services, platforms, infrastructure, and data components, along with their relationships, requirements, capabilities,

and configurations. The TOSCA model optionally supports service-specific orchestration and lifecycle management directives and operational policies that operate on these models. TOSCA model assumes a runtime environment (an “orchestrator”) in which service models are processed with the goal of orchestrating the service and managing service lifecycles.

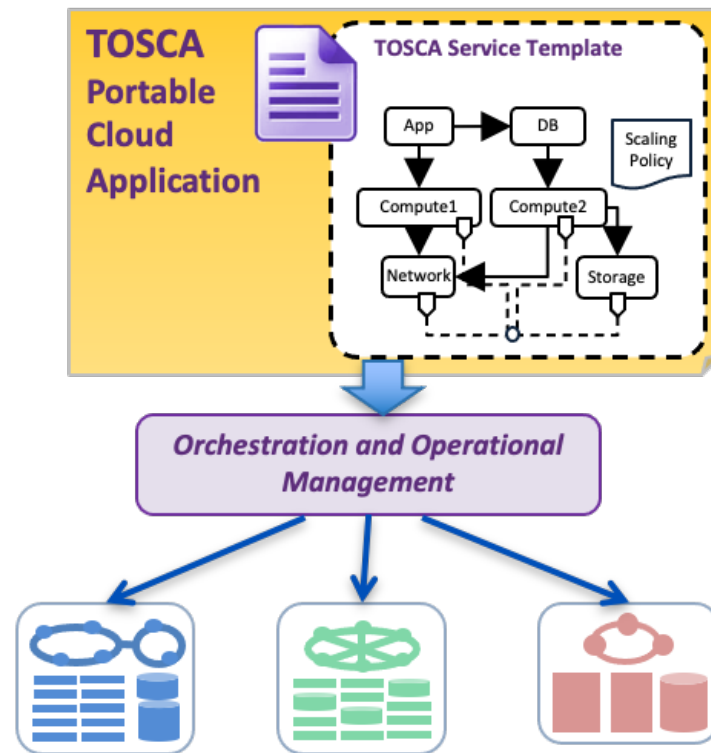


Figure 1 - What is TOSCA

3.1. Why TOSCA

TOSCA provides several features for the automation of service life cycle management via declarative directives. The following tables provide a high-level comparison of declarative versus imperative architecture.

Table 1 – Declarative versus Imperative Architecture

Declarative Architecture	Imperative Architecture
Arbitrary number of levels of recursion	Fixed number of layers
Federation built-in: North-south and east-west interfaces are the same	Federation must be added-on: North-south and east-west interfaces are different
Flexible resource layer: no architectural distinction between resources and services; resources are	Inflexible resources layer: distinction between resource layer and services layer is baked into the architecture

Declarative Architecture	Imperative Architecture
accessed as services, and services can be exposed as resources	
Identical DSL at each level in the recursion	Layer-specific APIs
Identical orchestration functionality at each level in the recursion	Layer-specific orchestration functionality
Interface implementations based on Domain-Specific Language (DSL)	Interface implementations based on APIs
Interface paradigm based on negotiation (request/response) and delegation	Interface paradigm based on management (higher layers control lower layers)
Organizing construct: recursive decomposition	Organizing construct: static layering

3.2. TOSCA Features

The following are some of the high-level features offered by TOSCA. These features are used as part of the service life cycle management.

- TOSCA Specification and TOSCA profiles. Profiles come with pre-defined type definitions.
- Ability to define type definitions (Node Types, Datatypes, Requirements, Interfaces, Capabilities, Artifacts etc.)
- Standard Interface & Lifecycle Operations - create / configure / start / stop / delete
- Custom Interfaces and Operations (Modify, Upgrade)
- Declarative workflows based on node relationships and capabilities
- Service Topology Template (Service inputs and outputs, Service components/nodes and their relationships)
- Substitution Mapping / Service Decomposition
- Node reference via node filter
- Package TOSCA models as a CSAR

3.3. Use of TOSCA and Yang

In certain use cases, there is a need to use multiple modeling languages for example, YANG. In such scenarios, the properties needed for the orchestration functionalities are modeled as TOSCA type definitions. The core feature configurations are modeled in YANG. The device configuration YANG model is parsed and represented as JSON schema and packaged as part of the CSAR. This allows the service orchestrator and implementation to validate input JSON data for the feature configurations against the schema at run time.

3.4. TOSCA Meta Model Phases

TOSCA Meta-Modeling Constructs support all phases of the service lifecycle

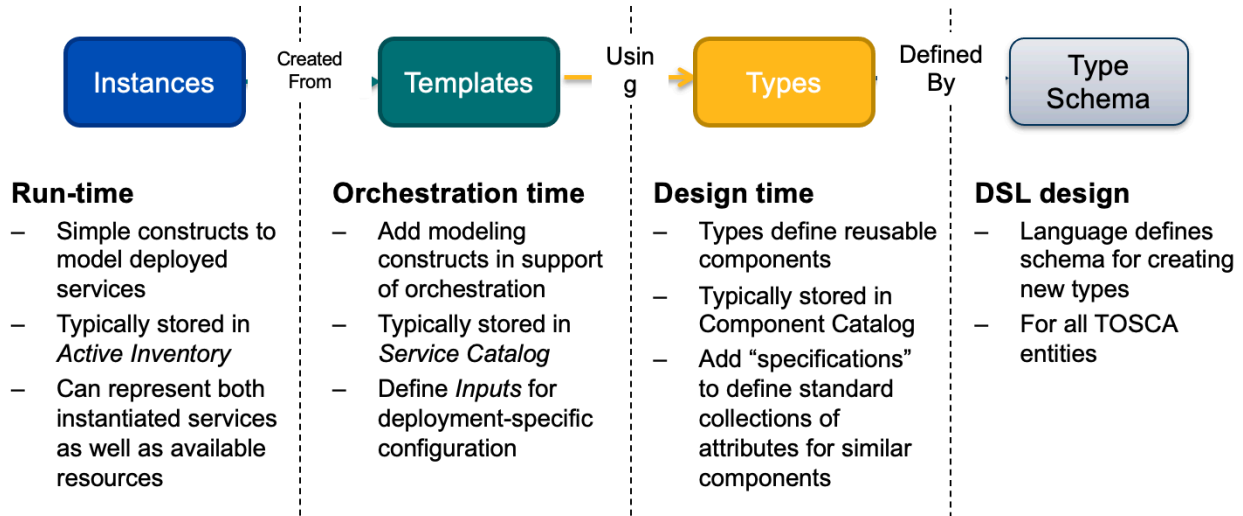


Figure 2 - TOSCA Meta Model Phases

3.5. TOSCA Object Model

Following is the TOSCA Object Model UML

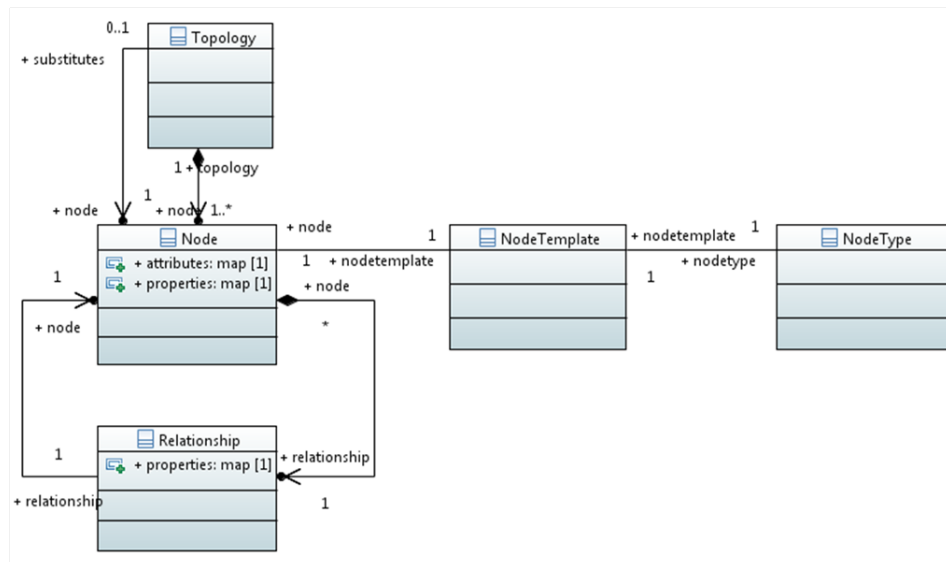


Figure 3 - TOSCA Object Model

4. Service Orchestration

4.1. Challenges with Service Life Cycle Management

The following lists some of the high-level challenges with regards to managing services.

- Service Model – How much to model/when not to model, Vendor specific versus Vendor agnostic, mix/match of model technologies (for example, TOSCA, YANG, SNMP, CLI etc.)
- Handing out of automation changes (Manage out of band changes and synchronization)
- Source of Truth – Network versus Automation System. This is critical in scenarios where out of band changes happen and the automation system does not have the capability to sync with the network.
- Rollbacks in case of failures – Automated versus Manual
- Managing service dependencies – Impact of a service change to another dependent service (Ripple effect)
- Service Upgrades – Model upgrade resulting in service instance(s) upgrades, Defect fixes, new feature rollouts etc.

4.2. Platform Goals

Software platform to unify and automate service lifecycle management functionalities via model driven (TOSCA for orchestration and YANG for configuration) architecture. The platform.

The platform offers capabilities to onboard models, provision/deploy/update services for different domains such as SD-WAN, Security, WIFI etc.

- A software platform to unify and automate complete service lifecycle
- Leverages Model driven Service Automation Language: TOSCA & YANG for configuration.
- Onboarding, Provisioning, Deployment, Monitoring
- Orchestrate multiple services (e.g., SD-WAN, Security, VNFs, other)
- Orchestrate hybrid environments – physical, multiple clouds and on premise
- Modeling and Service design tools
- Dynamic traffic steering and service chaining
- Active Inventory and Service Topology
- Southbound vendor specific implementations
- Enable us to define and implement business logic and functional logic independently




Figure 4 – Functional Architecture of Service Orchestrator

4.3. Service and Function Model Definition (Design)

Network Services and Network Functions are designed by creating TOSCA topology templates from which services can be instantiated and managed.

Service design involves the following steps:

- Defining set of type definitions (Data types, Node types, Relationship types and Capability types etc.). Reuse and/or enhance existing type definitions and add new type definitions.
- Define the topology template that represents the service as a topology of interconnected service components, relationships between the service components
- Define inputs that need to be provided at deployment time
- Define outputs that are expected to be returned upon successful service operation
- Provide implementation details (Plans/Workflows/Scripts references) for the life cycle operations (create, update, delete, upgrade, etc.) of the service components.
- Package as a cloud service archive (CSAR).

4.4. Service and Function Onboarding

The service packaged as a CSAR should be onboarded to the Service Orchestrator catalog repository. Service orchestrator validates (syntax and semantics) the CSAR and store in the catalog repository. Multiple versions of the CSAR can be onboarded.

4.5. Orchestration (Runtime)

At runtime, service(s) can be instantiated using the onboarded service CSAR. The service inputs as defined in the CSAR should be provided as part of service instantiation.

Service orchestrator performs validation of the inputs, invoke the implementations, and create service instances in the instance repository. Internally, TOSCA service orchestrator does the required substitutions, resolves references (matching requirements and capabilities) to the service components and resources. The service instance is a graph of service components nodes connected via relationships.

5. UseCases

5.1. SD-WAN

SD-WAN is an overlay service configured over the underlay infrastructure. Comcast offers physical/virtual CPEs in which SD-WAN is configured. CPEs are configured via VNF Manager. VNF Manager offers REST APIs. Configurations are staged at the VNF Manager and later committed to the CPEs (2 step process). SD-WAN logically modeled as 4 services. Services are linked. Information shared across services to enable configuration. For example, WAN Interface information, CPE details exposed by uCPE Service to VPN Site service. VPN information exposed by VPN Service to VPN Site Service.

Vendor VNF Manager modeled as a node type and node instance in the service (Not as a separate service/resource). Service Referencing (Orchestrator implementation extension, like the node referencing in TOSCA) feature used to link services and share information across services.

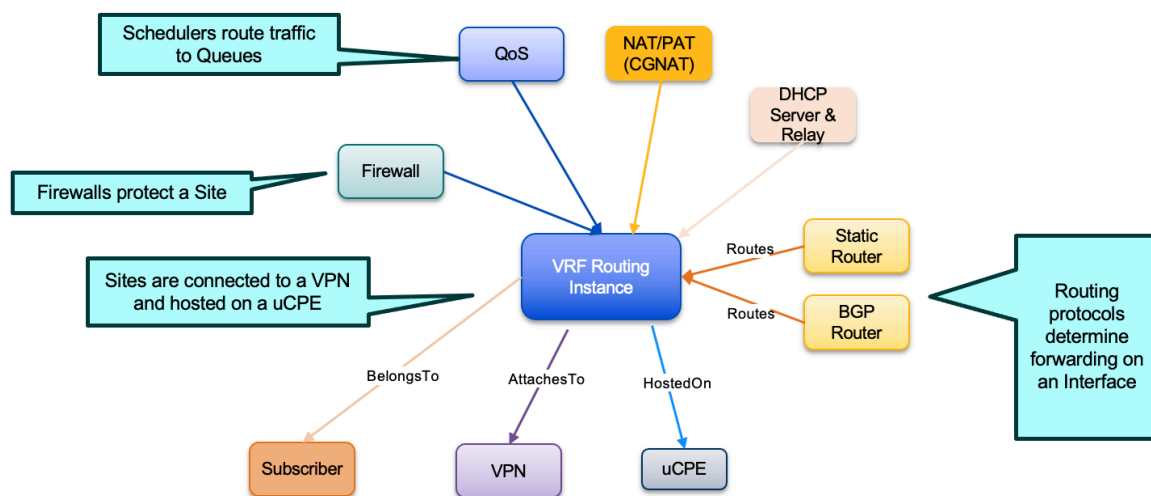


Figure 5 – SDWAN VPN Site Service (Sample)

The following are some high-level configurations represented as JSON type.

- Interface (LAN & WAN Interfaces) Configuration
- Routing Configuration (BGP, Static Routes etc.)
- DHCP Server & relay
- Class of Service
- CGNAT
- Firewall
- Traffic Steering Rules/Policies
- Custom Services & URL Filter Definitions

SD-WAN VPN Site Service Topology Template (Snippet only)

Following topology template snippet for a service component/node in a topology template. The sample highlights the node requirements and the life cycle operations of the service component (create, delete, upgrade) and the implementation references.

```
firewall:
  type: comcast.nodes.sdwan.fw.v1.FirewallConfiguration

  properties:
    vnfManager: { get_input: vnfManager }
    customerId: { get_input: customerId }
    firewallConfiguration: { get_input: firewallConfiguration }}
    zones: { get_property: [ vrfRoutingInstance, zones ] }
    interfaces: { get_attribute: [ SELF, ucpeService, interfaces ] }
    ## Properties omitted for clarity

  requirements:
    - protectable:
        node: vrfRoutingInstance
    - device:
        node_filter:
          properties:
            deviceId: { get_input: deviceId }

  interfaces:
    Standard:
      inputs:
        ## Omitted for clarity
      create:
        implementation:
          primary: Artifacts/Deployment/WORKFLOW/CreateFirewall
      delete:
        implementation:
          primary: Artifacts/Deployment/WORKFLOW/DeleteFirewall

  Upgrade:
    create:
      implementation:
        primary: Artifacts/Deployment/WORKFLOW/UpgradeFirewall
      inputs:
        ## Omitted for clarity
```

5.2. Firewall

Vendor agnostic models are defined as abstract models. The goal of abstraction is to avoid making technology/product decisions at design time. The abstract models will be decomposed (model decomposition feature of TOSCA/Service orchestrator) at runtime to vendor specific implementations.




Figure 6 - Firewall Vendor Abstraction




Figure 7 - Firewall Vendor Substitution

Substitution and Service decomposition - One Node instance substituted with a single node / set of nodes defined in the substituting topology template. In the above example, Firewall node substituted with Vendor A Firewall Template containing 4 nodes.

6. Conclusion

This paper highlighted how TOSCA as a modeling language and the TOSCA features combined with a service orchestrator simplifies the deployment of services for different domains.

The type definitions in TOSCA model simplifies the application to perform the validations. The service orchestrator combines all the node dependencies and generates the declarative workflow for the service minimizing the need to deal with dependencies in the application.

The service orchestrator enables to provide implementations for the supported life cycle interfaces and operations and is abstracted from the domain TOSCA models. In addition, it offers the flexibility to provide multiple implementations for a given model (Plug & Play).

Abbreviations

AP	Access Point
BGP	Border Gateway Protocol
CGNAT	Carrier-grade Network address translation
CPE	Customer Premises Equipment
CLI	Command Line Interface
CSAR	Cloud Service Archive
DHCP	Dynamic Host Configuration Protocol
JSON	JavaScript Object Notation
LAN	Local Area Network
REST	REpresentational State Transfer
SCTE	Society of Cable Telecommunications Engineers
SD-WAN	Software-Defined Wide-Area Networking
SDN	Software-Defined Networking
SNMP	Simple Network Management Protocol
TOSCA	Topology and Orchestration Specification for Cloud Applications
uCPE	Universal Customer Premises Equipment
VNF	Virtual Network Function
VPN	Virtual Private Network
WAN	Wide Area Network
YANG	Yet Another Next Generation

Bibliography & References

TOSCA Modeling Overview - OASIS, <https://www.oasis-open.org/committees/download.php/62645/TOSCA%20Overview%202018%2002%2026.pdf>.

Comparative Technical Analysis for 5G Fixed Wireless Access Rural Networks (2.6, 3.7 and 6.4 GHz)

A Technical Paper prepared for SCTE by

Dorin Viorel

Distinguished Technologist
CableLabs
858 Coal Creek Circle, Louisville, CO
303-661-3357
d.viorel@cablelabs.com

Ruoyu Sun

Principal Architect
CableLabs
858 Coal Creek Circle, Louisville, CO
303-661-6789
r.sun@cablelabs.com

Sanjay Patel,

Distinguished Strategist
CableLabs
858 Coal Creek Circle, Louisville, CO
303-661-3488
s.patel@cablelabs.com

George Hart,

Principal Architect
Rogers Communications
8200 Dixie Road, Brampton, ON
George.hart@rci.rogers.com

Table of Contents

Title	Page Number
1 Introduction.....	5
1.1 Executive Summary	5
1.2 Wireless Spectrum Intended for Rural FWA Applications	6
2 Analysis Tools and Assumptions	7
2.1 Analysis Methodology	7
2.1.1 Simulations Engine Block Diagram.....	7
2.1.2 CPE/BS Antenna Array Patterns.....	7
2.1.3 System Level Simulator	8
2.1.4 Link Level Simulator.....	10
2.1.4.1 LLS Methodology	10
2.1.5 Economics Performance Simulator.....	12
2.2 Assumptions.....	12
3 Key Results	13
3.1 Sensitivity Analysis for Key Inputs	13
3.1.1 BS Antenna Down Tilt.....	13
3.1.2 Network Load	15
3.1.3 BS Antenna Height.....	16
3.1.4 Small-Scale, Large-Scale Fading and O2I Loss	17
3.1.5 DL SINR and UL SNR.....	19
3.2 Cell Coverage.....	20
3.2.1 Cell and User (Data) Throughput.....	21
3.2.2 50, 100 and 300Mbps Coverage.....	23
4 Conclusions.....	26
4.1 The Impact of Coverage Limiting Factors	27
5 Abbreviations.....	29
6 Bibliography & References.....	30
7 Appendix	31
7.1 Simulations Assumptions	31
7.2 BS and CPE Antenna Array Performance Plots	33
7.2.1 BS Array	33
7.2.2 CPE Array	33

List of Figures

Title	Page Number
Figure 1. Block Diagram of the FWA NR Techno-Economics Performance Simulator.	7
Figure 2. Example of the 19 Cells Topology Used by SLS.....	9
Figure 3. BS Antenna Tilt Impact upon EIRP Distribution vs. Elevation Angle (Reference Boresight Horizontal Direction), 3.7GHz ($h_{BS}=60m$).....	14
Figure 4 Vtilt impact upon system interference and the related user throughput degradation as a function of Elevation Vtilt=0°/-15°/-21° ($h_{BS}=60m$, outdoor scenario, NetLoad=50%, 95% Service Availability, 3.7GHz).....	14
Figure 5. Comparative network interference impact upon (a) DL User Data Throughput and (b) Network Interference (expressed as CDF), impact for 25%, 50% and 75% network load (Outdoor scenario, 95% Service Availability , $h_{BS}=60m$, 3.7GHz)	16

Figure 6. Comparative network interference impact for outdoor and O2I scenarios (network load 50%, 95% service availability, $V_{\text{tilt}}=-15^\circ$, 2.6, 3.7 and 6.4GHz), $h_{BS}=30\text{m}$ and 60m .	17
Figure 7. CDFs of Large-Scale, Small-Scale and O2I fading, for an O2I Propagation Scenario ($h_{BS}=30\text{m}$ and $h_{BS}=60\text{m}$)	18
Figure 8. Comparative DL and UL SNR for Outdoor and O2I Scenarios (3.7GHz), $h_{BS}=60\text{m}$ vs. $h_{BS}=30\text{m}$	19
Figure 9. Comparative DL and UL SNR Degradation 2.6GHz vs. 3.7GHz vs. 6.4GHz, for Outdoor and O2I Scenarios (95% service availability, $h_{BS}=60\text{m}$).	20
Figure 10. Comparative Cell (User Data) and Throughput per User (User Data Only), 2.6, 3.7 and 6.4 GHz, for $h_{BS}=60\text{m}$ Outdoor and O2I Scenarios.	21
Figure 11. Comparative LOS, NLOS and Composite SINR and Tput, 2.6 3.7GHz	23
Figure 12. Outdoor (30 and 60m), 3.7GHz, 50, 100 and 300Mbps Coverage	24
Figure 13. Comparative O2I ($h_{BS}=30$ and 60m) Coverage, for 3.7 and 6.4GHz, When Subject to UL Coverage Limitations.	25
Figure 14. Coverage reduction between target 99% and 95% Service Availability, for 3.7GHz and 6.4GHz.	26
Figure 15. Sample of BS array parameters (3.7GHz). (a) Array geometry [2 2] subarrays, (b) Cross-Dipole antenna element geometry, (c) 3D radiation of a subarray (4x4x2) and (d) multi-beam azimuth radiation pattern.	33
Figure 16. Sample of Outdoor CPE Array Parameters (3.7GHz). (a) Array Geometry, (b) 3D Pattern of the Cross-Dipole Antenna Element, (c) 3D Radiation Pattern of the CPE Array and (d) Azimuth Radiation Pattern (Rectangular Coordinates).	34

List of Tables

Title	Page Number
Table 1 – Bands Available at 2.6, 3.7 and 6.4 GHz for FWA Services.	6
Table 2. CPE/BS arrays. Summary of the Main Configuration Parameters	8
Table 3. Summary of the CPE Array Performance Parameters	8
Table 4. Summary of the BS Array Performance Parameters.	8
Table 5. Assumed cell radius (MobEdge) and ISD for $h_{BS}=30\text{m}$ and 60m	9
Table 6. Cell Coverage Limitations	13
Table 7. V_{tilt} Impact upon Network Interference and DL Cell Throughput, referenced to the horizontal direction, (3.7GHz, Mobile Cell Edge=4000m, NetLoad=50%, 95% Service Availability)	15
Table 8. NetLoad impact upon overall cell radius (user data only), for $h_{BS}=30\text{m}$ (MobEdge=2000m) and $h_{BS}=60\text{m}$ (MobEdge=4000m).	16
Table 9. Comparative BS antenna impact upon network interference power impact upon victim cell (2.6, 3.7 and 6.4 GHz).	17
Table 10. The Link Budget Penalty Caused by Different Types of Fading Encountered by O2I Propagation ($h_{BS}=30\text{m}$ and 60m)	18
Table 11. Coverage and coverage reduction due to asymmetrical link budget (3.7GHz, $h_{BS}=60\text{m}$, 95% service availability, NetLoad=50%)	20
Table 12. Cell Coverage $h_{BS}=30\text{m}$ (95% Service Availability, NetLoad=50%, $V_{\text{tilt}}=-15^\circ$)	22
Table 13. Cell Coverage $h_{BS}=60\text{m}$ (95% Service Availability, NetLoad=50%, $V_{\text{tilt}}=-15^\circ$)	22
Table 14. Coverage Reduction (50, 100 and 300Mbps), for Outdoor 3.7GHz	24
Table 15. Coverage Reduction (50, 100 and 300Mbps), for O2I (3.7GHz)	25
Table 16. System And Cell Simulations Assumptions	31

Table 17. BS AND CPE Array Simulations 32
 Table 18 PHY/RF Assumptions 32
 Table 19. Atmospheric/Environment Conditions Assumptions 32

1 Introduction

1.1 Executive Summary

CableLabs developed a simulations engine intended to estimate a set of Key Performance Indicators (KPIs), concerning the user experience for Fixed Wireless Access (FWA), when estimated across a cluster of cells. This simulations engine is further used to provide input data for the FWA economics performance analysis tool. The Economics Analysis Methodology is presented in a companion paper.

The Technical Performance Analysis discusses the impact of different limitations upon the cell coverage:

- The support of frequency reuse 1, when smart arrays are employed. More specifically how could the network load impact upon victim cell coverage, could be controlled by dynamically controlling the base station (BS) array vertical down-tilt angle.
- The interfering cell network load impact upon the victim cell user throughput (T_{put}), indicating a network load in the 50% range or less, may be optimal to avoid significant cell throughput degradation.
- BS antenna height (h_{BS}) of 30m and 60m were analyzed, pointing to: i) a higher network interference impact when the frequency is decreased; ii) a lower network interference impact for the outdoor-to-indoor (O2I) case vs. the outdoor scenario; and iii) a very low or close to the noise floor outdoor 6.4GHz network interference.
- The impact of fading and loss mechanisms upon outdoor and O2I propagation is further analyzed. The results show that the main contributor to outdoors increased path loss is the small-scale fading, as long as the O2I fading is kept in check (CPE positioned close to the outer wall closest to the victim BS).
- The uplink (UL) link Budget limitations, pointing to i) downlink (DL) and UL coverage are in the same range for $h_{BS}=30m$ (3.7GHz) due the increased DL interference (MobEdge=2000m), ii) 2.6GHz coverage is slightly larger than the 3.7GHz, however the difference between the two is minimal due to the larger network interference impact on the 2.6GHz system, iii) the O2I UL coverage is more reduced vs. DL, due to reduced O2I DL interference and the reduced indoor CPE antenna gain vs. outdoor CPE gain,

The key results concerning cell coverage and user throughput are discussed from different angles:

- The outdoor coverage may extend beyond the serving cell edge, due to the highly directive outdoor CPE antennas.
- 2.6GHz line-of-sight (LOS) coverage is smaller than the 3.7GHz one ($h_{BS}=30m$ and 60m), due to the higher network interference impact.
 - Non-LOS (NLOS) becomes the dominant propagation mechanism, since the network interference is greatly reduced in NLOS conditions. Overall, the 2.6GHz composite coverage is NLOS driven being higher than the similar 3.7GHz coverage.
- The network interference ($h_{BS}=30m$) is higher than corresponding to $h_{BS}=60m$ for the 2.6GHz, due to the smaller cell radius selected for this case (MobEdge =2000m).
- The outdoor 6.4GHz cell edge is limited by the lower BS effective isotropic radiated power (EIRP) of 36dBm.
- Providing NR in Unlicensed Spectrum (NR-U) services in unlicensed 6 GHz spectrum may require a densified network.

- The O2I cell edge is UL limited, due to the additional UL path losses, caused by the O2I loss, the indoor CPE being required to either i) have a higher conducted RF power and/or ii) use a directive antenna array with a higher gain vs. an omni antenna.
- For 3.7GHz, $h_{BS}=60\text{m}$, outdoor case, all scenarios exceeding a service availability >97% are impacted by an UL asymmetric link budget.
- As the probability of achieving a particular radio link throughput (i.e. service availability) is increased to 99%, cell coverage is reduced.
- A FWA network planning targeting 95% service availability backed by 50% network load may be an optimal trade-off.

The paper is organized as follows:

- Section 2 introduces the analysis methodology and the simulations assumptions.
- Section 3 presents the main results.
- Section 3.1 discusses the limiting factors driving to a sub-optimal coverage.
- Section 3.2 summarizes the key results characterizing the user coverage and overall cell coverage, the coverage for FWA throughput thresholds.
- Section 4 summarizes the main findings of this analysis.

1.2 Wireless Spectrum Intended for Rural FWA Applications

FWA provides broadband service in areas where wired solutions are not prevalent or as a competitive alternative to wired broadband. One key performance factor is the optimization of wireless coverage for high throughput services. While earlier FWA implementations were hampered by spectrum and other technology implementations, 5G maximizes the potential FWA performance by employing channel bandwidths up to 100MHz (sub-7GHz) and up to 400MHz (24 – 52GHz spectra).

We identified the following spectra with channel bandwidth allocations in excess of 60MHz, which could be suitable for FWA in North America and Europe, following related spectrum auctions.

Table 1 – Bands Available at 2.6, 3.7 and 6.4 GHz for FWA Services

Frequency [MHz]	Band	Common Name	Max Channel BW [MHz] [4]	Market	Comments
2500 – 2696	n41	BRS	100	USA	
3550 – 3700	n48	CBRS	100	USA	
3450 – 3650 3650 – 3980	n77 Canada		100	Canada	
3300 – 4200	n77 global		100	Global	except USA and Canada
3450 – 3550 3700 – 3890	n77 USA	C-band	100	USA	
5925 – 7125	n96	6GHz	100 80	Regional USA, Canada	
5925 – 6425	n102		100 80	EU	

2 Analysis Tools and Assumptions

Throughout the paper we use the generic term BS for 5G BS (gNB).

2.1 Analysis Methodology

Our FWA coverage analysis is based on a simulation engine developed by CableLabs. This engine was designed to support an economics analysis based on the technical performance of a large-scale cluster of 5G of cells.

2.1.1 Simulations Engine Block Diagram

The simulations engine block diagram is presented (Figure 1). Within the simulation engine the System Level Simulator (SLS) module evaluates radio performance for unmodulated signals over a large number of iterations (Monte Carlo simulation). The SLS generates statistical results of aggregated interference across a cluster of 19 cells arranged in 2 rings surrounding the cell of interest. The 5G New Radio (NR) Link Level Simulator (LLS) simulates 3GPP compliant waveforms targeting the 5G NR performance in a simulated network environment including the network interference predicted by the SLS. The CPE/BS antenna pattern array block generates suitable BS and CPE antenna array patterns. The antenna arrays are critical for supporting frequency reuse (FR) 1 across the radio network by optimizing the link budget component.

The economics analysis block estimates the economic feasibility of the 5G FWA O2I and service delivery network under consideration, based on a set of technical KPIs. This paper focuses on the technical simulation consists of SLS, LLS and antenna characteristics. A companion SCTE Cable-Tec Expo 2022 paper presents the economics analysis.

The functionality of the component blocks is explained in sections 2.1.2 to 2.1.5.

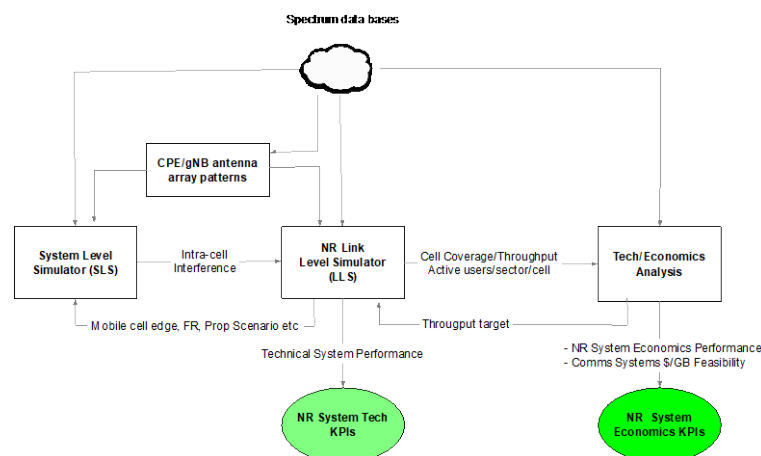


Figure 1. Block Diagram of the FWA NR Techno-Economics Performance Simulator.

2.1.2 CPE/BS Antenna Array Patterns

The CPE/BS antenna array patterns are critical for SLS and NR LLS to calculate the link budget in an interference rich environment, for both the victim path and interference paths in both DL/UL. The spectra used for this analysis are n41 (central frequency 2.6GHz), n77 (central frequency 3.7GHz) and n96

(central frequency 6.4GHz). For each of these bands, a number of antenna array parameters are estimated (Table 2, Table 3, Table 4). A selection of plots summarizing the performance of these arrays is presented in section 7.2.

An antenna array manufacturer may further optimize the electrical performance of the arrays described below. In this context, the performance of these array should be considered as readily achievable.

Table 2. CPE/BS arrays. Summary of the Main Configuration Parameters

	Array Type	Array of Subarrays	Subarray Size	Antenna Element Type	Vertical Electrical Tilt [°]	Estimated Mechanical Size L×W [mm]
CPE Indoor Antenna	UCA	4×1	1×4	Omni	0	
CPE Outdoor Antenna	URA	1×1	2×8×2	Cross-Dipole	0	400×150
BS	URA	4×1	4×4×2	Cross-Dipole	-15	

Table 3. Summary of the CPE Array Performance Parameters

	Max Gain [dBi]	Azimuth (Beam) HPBW [°]	Elevation HPBW [°]	Estimated Mechanical Size L×W [mm]
Indoor Antenna	8.2	360	10 ⁰	58x58x158 (cylinder)
Outdoor Antenna 2.6GHz	16.0	16	53	
Outdoor Antenna 3.7GHz	17	16	52	400×150
Outdoor Antenna 6.4GHz	17.1	16	52	400×80

Table 4. Summary of the BS Array Performance Parameters

	Max Gain [dBi]	Azimuth (Beam) HPBW [°]	Elevation HPBW [°]	Estimated Array Mechanical Size L×W [mm]
2.6GHz	15.6	29	32	
3.7GHz	16.4	29	30	800×200
6.4GHz	16.5	29	28	480×120

2.1.3 System Level Simulator

The SLS is 3GPP compliant in terms of cell topology [1] and propagation models [2]. As shown in Figure 2, the SLS's topology consists of 19 sites, with the serving BS (#1) in the center and two rings of interfering BS (#2 – #19). The assumed cell Radii are presented in Table 5, where inter-site distance (ISD) is calculated as:

$$ISD = \sqrt{3} \times Radius \quad \text{Equation 1}$$

For this analysis a three sectors per cell topology was assumed. All cells use a frequency reuse 1 (FR1). This topology allows each cell to employ the maximum channel bandwidth (BW), 100MHz for n41 and n77, 80MHz for n96, but triggering a significant intra-network interference for some scenarios. Adjacent

channel interference is ignored since is much weaker than co-channel interference. The BS and victim CPE employ the antenna patterns presented in section 7.2. For this paper, we conducted a comparative analysis on multiple scenarios, exercising $h_{BS}=30$ and 60m, outdoor and O2I scenarios for 2.6, 3.7 and 6.4GHz frequencies. While the analysis summarizes results for both $h_{BS}=60$ m and $h_{BS}=30$ m, for the sake of brevity, the authors chose to minimize the amount of simulations plots concerning $h_{BS}=30$ m.




Figure 2. Example of the 19 Cells Topology Used by SLS.

The related cell radius and ISD are presented in Table 5.

Table 5. Assumed cell radius (MobEdge) and ISD for $h_{BS}=30$ m and 60m

h_{BS} [m]	Cell radius (MobEdge) [m]	ISD [m]
30	2000	3464
60	4000	6928

SLS is a Monte Carlo simulation technique that runs a large number of iterations (e.g., 100,000 for each scenario analyzed) to output statistical results. For each iteration, a random CPE location is generated inside the serving cell. Distance, relative azimuth and elevation angles between the CPE and each of the 19 BS sites are calculated to estimate path loss, antenna gain, shadowing loss, small-scale fading, and O2I wall penetration loss for both the serving link and interference links. The distributions of these parameters are generated by the Monte Carlo simulation, based on which the service availability (e.g., 95% or 99%) is derived.

Network load is another variable given that not all radio resource blocks are fully allocated for all BSs at all times. For example, if the network load is set to 25%, each of the interfering cells and the victim cell have a 25% probability to transmit in each SLS iteration. All bands under consideration are time division duplex (TDD) bands. We assume TDD sync is enforced among BS s to avoid DL-to-UL and UL-to-DL interference. In this paper, we focus only on the aggressor-to-victim DL-DL interference from the aggressor BS to the victim CPE. Another internal study (not discussed here) indicated that UE to UE and UE to BS interference are negligible, even when TDD aligned.

The cell scenarios modeled in this paper are based on 3GPP channel models [2] as defined by the rural macro (RMA) environment. A random LOS probability is assigned for each SLS iteration, then the path loss, shadowing and fading are calculated for this probability of LOS or NLOS conditions. A Ricean K -

factor of 12 dB is used to generate small-scale fading in the LOS condition and small-scale fading Rayleigh distribution in NLOS conditions. The O2I wall loss is also based on the 3GPP model [2] following a Gaussian distribution for a residential home (wood outer wall, regular glass windows, with a glass/wall ratio of 0.3) centered on a mean value of 9.3 dB and standard deviation of 4.4 dB [2].

The calculation of interference power level and aggregated interference is presented in [12]. The aggregated network interference is further used by the NR LLS to estimate the signal-to-interference-plus-noise ratio (SINR) for the target victim CPE. Shadowing (large-scale fading), small-scale fading, and O2I loss variables for the signal link are also provided to the NR LLS to quantify service availability.

2.1.4 Link Level Simulator

The NR LLS is a Matlab-based simulator, developed in compliance with relevant recommendations from the 3GPP [2], [3], [4], [5], [6], [7], ITU [8], [9], [10] and TIA [11] targeting the behavior of a 5G waveform when subject to a target propagation environment and for a given geography, propagation model and system interference impact.

The LLS outputs the following technical performance KPIs.

- Dynamic DL/UL SINR, predicting the DL/UL link budget asymmetry per path length unit.
- User throughput versus path length for LOS, NLOS and composite propagation.
- DL user spectral efficiency (SE) versus path length as a function of NR Rel-16 link adaptation.
- DL received signal level (RSL) versus path length.
- BS array EIRP vs. path length and vertical tilt angle.
- Victim fade margin and system interference impact (based on the SLS inputs).
- 50, 100, 300 Mbps (configurable) and mobile cell edge coverage service availability.
- Coverage and user throughput vs. BS array vertical tilt angle.
- Coverage and user throughput vs. network load.
- Network interference power vs. BS EIRP or BS array tilt angle or network load etc.

2.1.4.1 LLS Methodology

The analysis methodology employs the following steps.

The PathLoss is calculated based on:

$$PathLoss(f, path, BS, CPE, environment) = PropagationLoss + prctile\{Shadowing + AtmosphericConditions, Outage\}$$

Equation 2

where:

PropagationLoss is a function of (distance, frequency), including BS Height, CPE height, CPE Outdoor/O2I scenario and clutter;

prctile function calculates the additional link fade margin for the target availability

Shadowing outdoor large-scale fading modeled by a Gaussian distribution (sigma scenario, mean path loss) and

AtmosphericConditions additional path loss caused by rain fading, water vapor fading and gaseous fading. For a path length <5 km below 10 GHz, the impact of atmospheric conditions upon the PathLoss may be negligible. More details about the environmental factors methodology upon the PathLoss could be found in [12].

The link SINR is calculated based on:

$$SINR = RSL - prctile\{LargeScaleFading + SmallScaleFading + ShadowingO2I + Interference(NetworkLoad) - WallLoss(sigmaO2I, MeanLoss)\} - NF + 10 \times \log_{10}(OS) \quad \text{Equation 3}$$

where

NF CPE noise figure (it includes the CPE modulation implementation losses);

OS Oversampling ratio (CPE PHY). Though the oversampling is effective on improving the SINR for noise-limited environments, it may have a limited capability on interference limited ones.

SmallScaleFading Modeled as the Rayleigh distribution for NLOS (Rice distribution with low K factor available is optional);

ShadowingO2I Modeled by a normal distribution {meanWallLoss, sigmaO2I}; Mean(WallLoss) is the mean outer wall penetration loss (O2I only), as a function of wall material, glass to wall area ratio, and glass material.

Interference(Outage, NetLoad) Network interference plus noise floor as a function of link outage probability and network load. The (System) Interference is calculated by the SLS for given outage, MobileCellEdge, and CPE/ BS antenna geometries.

The CodingRate function calculates the LDPC coding rate and the QAM modulation order as a function of SINR, based on the QAM256 modulation and coding scheme (MCS) table.

$$[MCS_{coding}, MCS_{qam}] = CodingRate(SINR) \quad \text{Equation 4}$$

where

MCS_{qam} QAM order for the target distance, subject to link adaptation,

MCS_{coding} MCS coding rate for the target distance, subject to link adaptation.

Spectral Efficiency is defined as:

$$SpectralEfficiency = MCS_{qam} \times MCS_{coding} \quad \text{Equation 5}$$

The User Throughput ($UserTput$) is calculated as follows:

$$UserTput(distance) = MCS_{coding} \times MCS_{qam} \times PRB \times (TDD_DL_sym - DMRS - ControlSym) \times MIMO \times Slots \times Subframes \times Frames \quad \text{Equation 6}$$

where

MCS_{coding} Emulates the link adaptation as a function of the SINR degradation by calculating the LDPC user coding rate,

MCS_{qam} Emulates the link adaptation as a function of the SINR degradation by calculating the QAM modulation order,

PRB The number of Physical Resource Blocks per slot,

TDD_DL_sym Chosen as 12 symbols/slot (TDD ratio 12:1:1), maximizing the DL output,

DMRS Selected as 1 symbol/slot,

ControlSym The number of control symbols per slot,

MIMO DL MIMO rank. MIMO performance in O2I propagation environments may be subject to degraded performance under severe multipath conditions due to the large amplitude imbalance between the different Rx air layers reaching the antenna receiver.

Slots 2^μ , where μ =numerology order: SCS= $\mu \times 15\text{kHz}$, where $\mu=1, 2, 3, 4$ (NR Rel 15-17)

Subframes Number of subframes/frame:10/frame

Frames Number of frames per second (10/s).

The overall BS User Throughput ($DLUserTput$) is based on:

$$DLUserTput = UserTput \times Beams \times Sectors \quad \text{Equation 7}$$

where

$UserTput$ User throughput data (it excludes the control and signaling data)

$Beams$ beams/sector and

$Sectors$ sectors/cell.

All the above steps are repeated three times to calculate above parameters for LOS, NLOS and composite (LOS/NLOS) for UMi, UMa, or RMa scenarios (only the RMa scenario considered for this analysis). The above functions are calculated for every meter of the path length, supporting high accuracy plots.

2.1.5 Economics Performance Simulator

This paper focuses on the technical performance parameters of mid-band spectrum. A companion strategy brief looks at the potential competitive implications of FWA services, based on main KPI of the Technical Performance Analysis (e.g., household (CPE) distribution per coverage as a function of path length to the BS). The analysis estimates how much capacity can be created by cell area under a range of assumptions. Using data from CableLabs' quarterly bandwidth usage report, the analysis forecasts future peak broadband demand per average household, factors household density for the areas of interest and estimates what percentage of FWA subscribers can be supported from a market penetration perspective.

2.2 Assumptions

There is a large set of assumptions backing the Key Performance Results, presented in this paper. All simulations results presented in this paper are based on this set of assumptions. These assumptions are grouped in the following categories:

- System and cell simulations assumptions (see Table 16)
- BS and CPE simulations assumptions (see Table 17)
- Atmospheric/environment conditions assumptions (see Table 18)
- CPE/ BS arrays summary of the main configuration parameters (see Table 2)
- Summary of the CPE array performance parameters (see Table 3)
- Summary of the BS array performance parameters (see Table 4)

Samples of BS and CPE arrays performance plots (3.7GHz) are presented in Figure 15 and Figure 16.

3 Key Results

The key performance parameters (coverage, throughput) are impacted by a series of other parameters. The following sub-sections analyze the impact of these parameters upon coverage and user/cell throughput. All throughput values presented hereby, represent user throughput only, the signaling and control data being de-embedded.

3.1 Sensitivity Analysis for Key Inputs

We analyze different cell coverage limitations, due to different factors.

Table 6. Cell Coverage Limitations

Sub-section	Parameter	Frequency (GHz)	BS antenna tilt	Network load	h_{BS} (m)	Cell radius (m)	O2I	Service availability
3.1.1	BS antenna tilt	3.7	0° vs. -15° vs. -21°	50%	60	4000	Outdoor CPE	95%
3.1.2	Network load	3.7	-15°	25% vs. 50% vs. 75%	60	4000	Outdoor CPE	95%
3.1.3	BS antenna height	2.6, 3.7 and 6.4	-15°	50%	30 vs. 60	2000/4000	Indoor/outdoor CPE	95%
3.1.4	Fading and O2I Loss	3.7	-15°	50%	30/60	2000/4000	Indoor CPE	95% and 99%
3.1.5	DL/UL	3.7	-15°	50%	30/60	2000/4000	Indoor/outdoor CPE	95%

3.1.1 BS Antenna Down Tilt

BS antenna array's performance is critical for controlling the radiated interference across the neighboring cells. Dependent on BS array performance, the network operator may enable or not frequency reuse 1. Frequency reuse is the number of times the same RF channel (frequency) is reused throughout the network. The most efficient spectrum utilization occurs when the same frequency is reused across the entire network, frequency reuse 1 being the most efficient one.

These simulations highlight the impact of vertical tilt upon frequency reuse 1 in a 5G network. We chose 3 vertical tilts: 0° (along the horizon), -15° (-3dB along the horizon) and -21° (equivalent to -6dB in the H direction), for a 3.7GHz 4×4×2 sub-array, grouped in 4×1 subarrays (2×2 subarray configuration).

Figure 3 is based on BS antenna array's ability to steer the beams vertically by electrical means, rather than mechanical ones (e.g., LTE case).

For the particular scenario simulated, the effective radiated power vs. the horizon is max EIRP ($V_{tilt}=0^\circ$), -2.5dB ($V_{tilt}=-15^\circ$) and -5 dB ($V_{tilt}=-21^\circ$). This EIRP reduction is further compounded, during the

simulations process with the path loss, system interference and different fading mechanisms for the propagation scenarios under consideration. Unlike a mechanically controlled vertical tilt, such an antenna array could accommodate dynamic tilt factors, dependent on the network load, potentially alleviating the impact of interference.

The related impact upon the system interference (assuming all BS antennas are tilted by the same angle) and the cell user throughput/coverage are presented in Figure 4. The impact of horizontal beam steering upon the network interference is not discussed in this paper.




Figure 3. BS Antenna Tilt Impact upon EIRP Distribution vs. Elevation Angle (Reference Boresight Horizontal Direction), 3.7GHz ($h_{BS}=60m$).

The related DL SINR degradation, for Path Length=2000m and 4000m (Cell Mobile Edge), network load is presented in Figure 4.




Figure 4 Vtilt impact upon system interference and the related user throughput degradation as a function of Elevation Vtilt=0°/-15°/-21° ($h_{BS}=60m$, outdoor scenario, NetLoad=50%, **95% Service Availability, 3.7GHz).**

Table 7. Vtilt Impact upon Network Interference and DL Cell Throughput, referenced to the horizontal direction, (3.7GHz, Mobile Cell Edge=4000m, NetLoad=50%, 95% Service Availability)

Vtilt Variation		Network Interference Variation		DL Cell Throughput Variation	
Vtilt	Gain Variation [dB]	Interference Power [dBm]	Degradation [dB]	Tput [Gbps]	Degradation [%]
0	-6	-57.8	3.7	4.3	0
-15°	-3	-59.5	2	4.3	0
-21°	0	-61.5	0	3.6	15.2

Observations:

- Due to the higher directivity of the outdoor CPE array, one beam's coverage could extend into the adjacent cell, though exceeding the target MobEdge (e.g., 4000m for the modeled case).
- No significant cell throughput degradation on the MobEdge (4000m, $h_{BS}=60m$), when the BS array is tilted from 0° down to -15°.
- The effective user throughput degradation when the BS array is tilted down to -21° from 0°, is 0.63% (24 users per cell) or 15.2% for the entire cell.
- Frequency Reuse 1 could be effectively implemented.
- The DL SINR gets improved by ~2dB when the BS antenna tilt gets tilted by 21°. A sharper vertical beam would reduce even more the network interference power radiated towards a victim cell, thus improving the victim user's SINR and related throughput. As a consequence, a larger subarray vertical size should be used in order to optimize even more FR1 coverage (e.g., 4 subarrays each subarray being 8×4×2 antenna elements, amounting to a 16×8×2 array).
- If the BS antenna's tilt is controlled dynamically, as a function of the network load, the system interference and subsequently the cell coverage and user throughput could be optimized.

3.1.2 Network Load

The impact of network load upon the cell performance in terms of cell (user data) throughput is analyzed, while holding BS Vtilt constant (-15°). We compare the cell throughput and the CDF of the network interference for NetLoad=25%, 50% and 75%, see Figure 5. The quantitative results are summarized in Table 8.

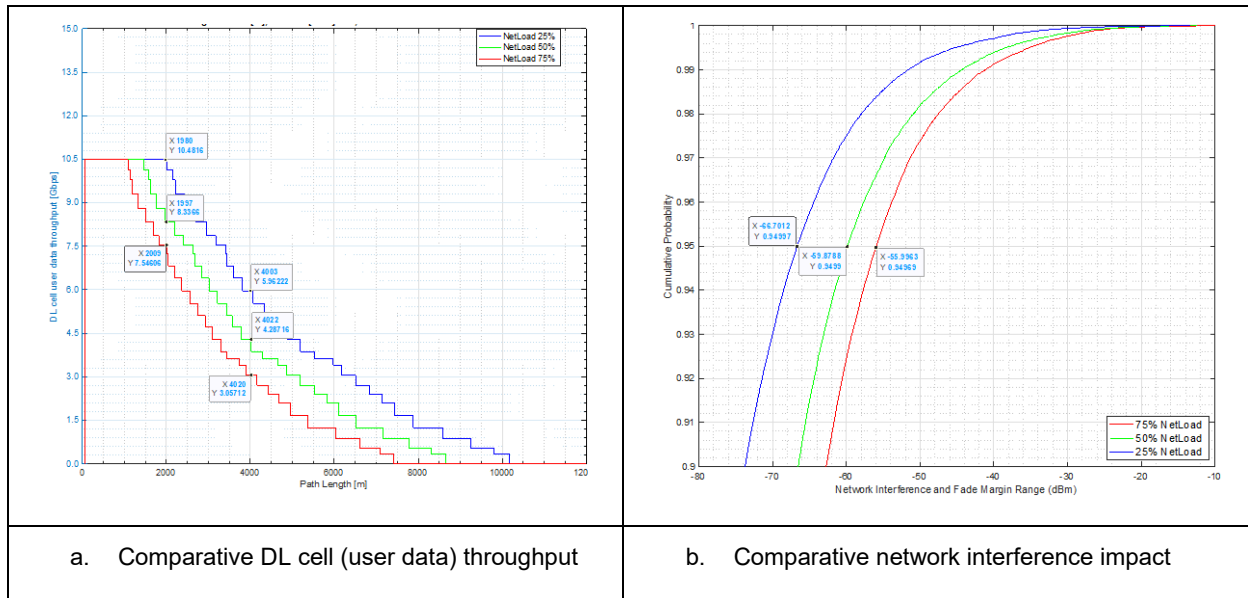


Figure 5. Comparative network interference impact upon (a) DL User Data Throughput and (b) Network Interference (expressed as CDF), impact for 25%, 50% and 75% network load (Outdoor scenario, 95% Service Availability, $h_{BS}=60m$, 3.7GHz)

Table 8. NetLoad impact upon overall cell radius (user data only), for $h_{BS}=30m$ (MobEdge=2000m) and $h_{BS}=60m$ (MobEdge=4000m).

	Cell Radius=2000m		Cell Radius=4000m (DL edge)	
NetLoad	Cell Tput [Mbps]	Variation	Cell Tput [Mbps]	Variation
25%	10481	0%	5962	0%
50%	8336	-21.5%	4278	-28.3%
75%	7540	-28.1%	3057	-48.8%

Observations:

- Cell throughput gets degraded by 48.8% ($h_{BS}=60m$, MobEdge=4000m) and by 28.1% ($h_{BS}=30m$, MobEdge=2000m), when NetLoad is increased from 25% up to 75% on the cell radius.
- The higher is the netload, the lower is the cell coverage and user/cell throughput.
- The coverage and throughput degradation caused by the increased network load highlights the significance of steering 5G antenna arrays, dynamically updating Vtilt as a function of network load, across a cluster of cells accordingly.

3.1.3 BS Antenna Height

The network interference impact caused by low and high BS antenna heights ($h_{BS}=30m$ and $h_{BS}=60m$), outdoor and O2I scenarios, for the 3 frequencies of interest (2.6, 3.7 and 6.4GHz) is analyzed. Summary results are presented in Figure 6. The quantitative results are summarized in Table 9.

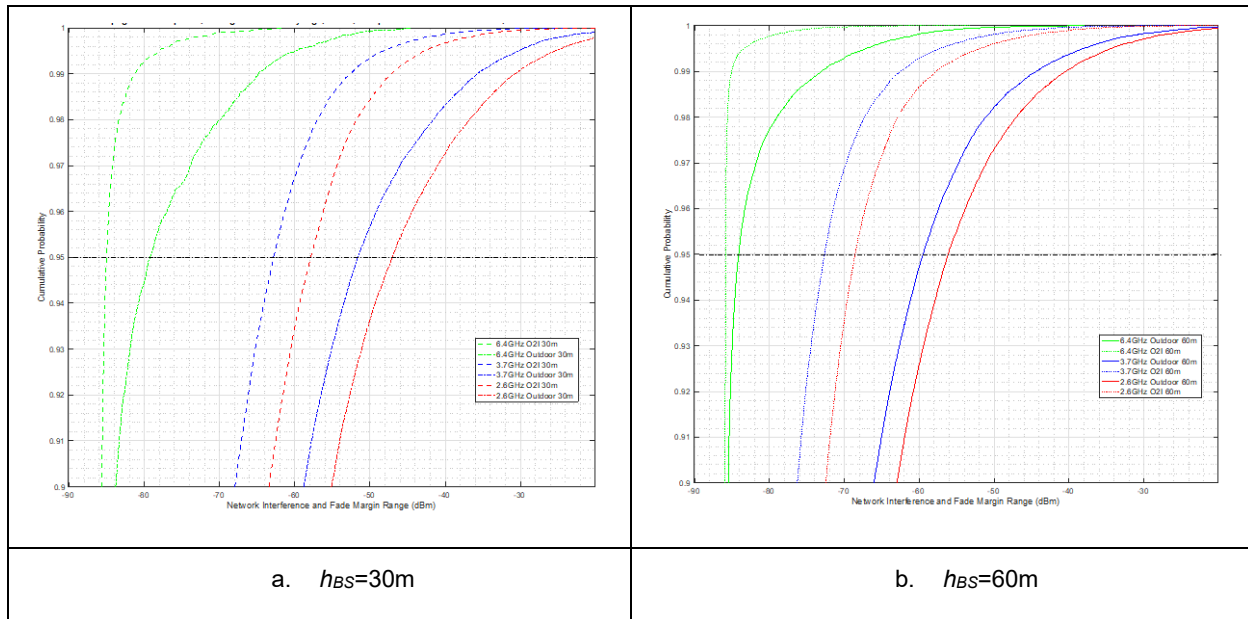


Figure 6. Comparative network interference impact for outdoor and O2I scenarios (network load 50%, 95% service availability, $V_{tilt}=-15^\circ$, 2.6, 3.7 and 6.4GHz), $h_{BS}=30m$ and 60m.

Table 9. Comparative BS antenna impact upon network interference power impact upon victim cell (2.6, 3.7 and 6.4 GHz).

	$h_{BS}=30m$		$h_{BS}=60m$		Degradation	
Central Frequency	Outdoor	O2I	Outdoor	O2I	Outdoor	O2I
2.6 GHz (ChBW=100MHz)	-47.2	-58	-56.3	-68.6	9.1	10.6
3.7 GHz (ChBW=100MHz)	-51.8	-62.9	-60	-72.8	8.3	9.9
6.4 GHz (ChBW=80MHz)	-79.5	-85	-84.2	-85.8	4.7	0.8

Observations:

- Mobile Cell Edge of 2000m is chosen for $h_{BS}=30m$, which triggers a higher network interference impact upon the victim cell than $h_{BS}=60m$ (MobEdge 4000m). DL SINR is degraded as follows, when remote head (RH) height is reduced from $h_{BS}=60m$ down to $h_{BS}=30m$:
 - Outdoor scenario: 9dB (2.6GHz), 8.3dB (3.7GHz) and 4.7dB (6.4GHz).
 - O2I scenario: 10.6dB (2.6GHz), 9.9 (3.7GHz) and 0.8dB (6.4GHz).
- The 2.6GHz system is subject to the highest network interference, among all considered scenarios, due to the lowest propagation losses vs. 3.7 and 6.4GHz cases.
- The outdoor 6.4 GHz system operates as a quasi interference free system, due to the lower EIRP density: 27dBm/10MHz vs. 50dBm/10MHz (2.6 and 3.7GHz).
- The network interference for the O2I case is lower than for the outdoor one, due to the additional O2I propagation fading impact.

3.1.4 Small-Scale, Large-Scale Fading and O2I Loss

The O2I propagation is subject to 3 different types of fading mechanisms:

- Large-scale fading, also known as shadowing

- The propagation delay is larger than the coherence time of the channel, hence the resulting received amplitude and phase are quasi constant. This type of fading is mainly caused by obstruction of the main path (e.g., shadowing, path loss).
- Small-scale fading
 - This is due to the multipath components in the propagation channel. The multipath arrived at the receiver can be constructive or destructive depending on the phase of each multipath, which cause the signal strength variation.
 - The small-scale Fading was modeled by a Rice distribution ($K=12$ dB) for LOS conditions and by a Rayleigh distribution for NLOS conditions.
- O2I loss is a Gaussian distribution centered on the mean value of the outer wall. It was assumed that the indoor CPE is positioned 1m behind the closest outer wall to the serving BS. The O2I fading does not apply for the outdoor scenario.

It should be noted that the Doppler spread fading, affecting mobile communications, doesn't impact FWA propagation.

All these three fading/loss mechanisms are modeled by three different distributions, which are summed up statistically, following 100,000 random victim CPE locations. The cumulative loss is further calculated by applying CDF function for the target service availability. The three distributions presented above are exemplified for O2I propagation ($h_{BS}=30m$ and $h_{BS}=60m$).

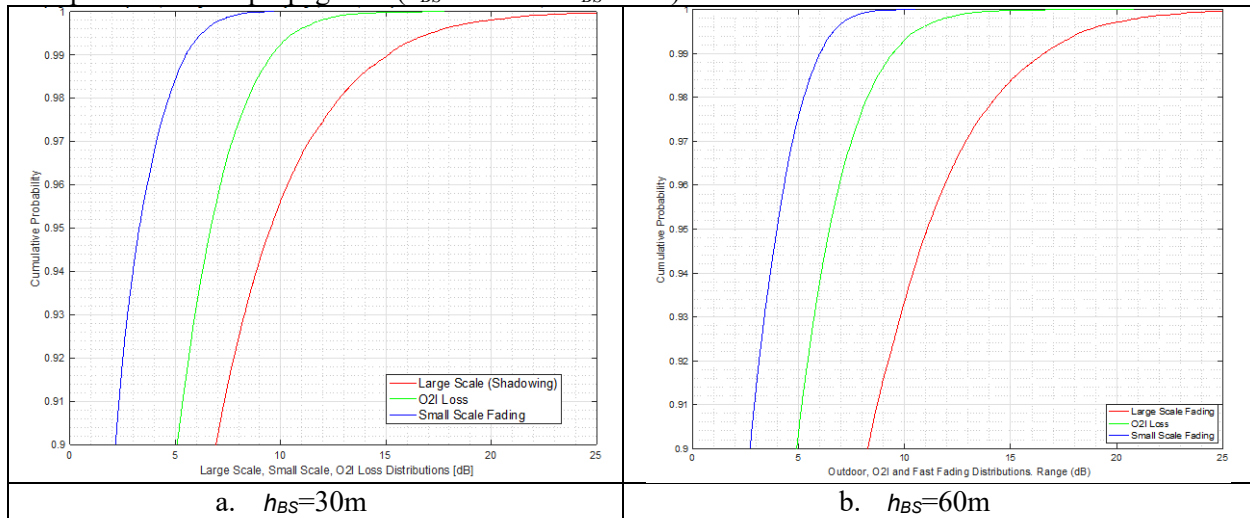


Figure 7. CDFs of Large-Scale, Small-Scale and O2I fading, for an O2I Propagation Scenario ($h_{BS}=30m$ and $h_{BS}=60m$)

The following table summarizes the quantitative analysis. The channel is more likely in LOS condition with the smaller cell radius with 30m BS height. Thus, the large-scale and small-scale fading at 30m h_{BS} are relative smaller than that at 60m h_{BS} .

Table 10. The Link Budget Penalty Caused by Different Types of Fading Encountered by O2I Propagation ($h_{BS}=30m$ and 60m)

	$h_{BS}=30m$			$h_{BS}=60m$		
Service Availability	Large-scale fading (Shading)	Small-scale Fading	O2I loss (dB)	Large-scale fading (Shading)	Small-scale Fading	O2I loss (dB)
95%	9.4	3.3	6.7	11.0	6.4	6.4

99%	15.3	5.5	9.6	14.3	5.2	9.4
-----	------	-----	-----	------	-----	-----

Observations:

- Large Scale (shadowing component) fading has the strongest impact among the three mechanisms analyzed, due to the NLOS propagation (Rayleigh type fading).
- The higher is the desired service availability, the higher is the composite fading impact upon the link budget.
- The O2I Loss could become the driving factor of the composite fading if the CPE is placed deep inside the house (vs. the outer wall facing the BS) and/or other construction materials used for the outer wall.
- FWA could use a lower target service availability (e.g., 95%):
 - Even for a higher service availability, the user experience impact may not be noticeable, as long as the user may not use the highest achievable allocated user data rate.

3.1.5 DL SINR and UL SINR

UL link budget is another coverage limiting factor. Usually the CPE/UE has a lower EIRP than the BS, driving to an asymmetrical link budget. This could cause coverage limitations, limited by the UL lowest MCS connection. In this section we examine the UL SNR impact upon the Cell Edge (MobEdge).

Firstly, we compare DL SINR and UL SNR for outdoor and O2I scenarios (3.7GHz, $h_{BS}=60m$), as presented in Figure 8. The max throughput is achieved for SINR=25.5dB (0.925 coding rate and QAM256), while the minimum throughput (cell edge conditions) is achieved for SINR=-4.5dB (0.117 coding rate and QPSK). It should be noted that the network planners may use a higher min MCS, for cell edge calculations, allowing cell overlapping in order to support seamless mobile handover between adjacent cells, however for FWA this may not be required.

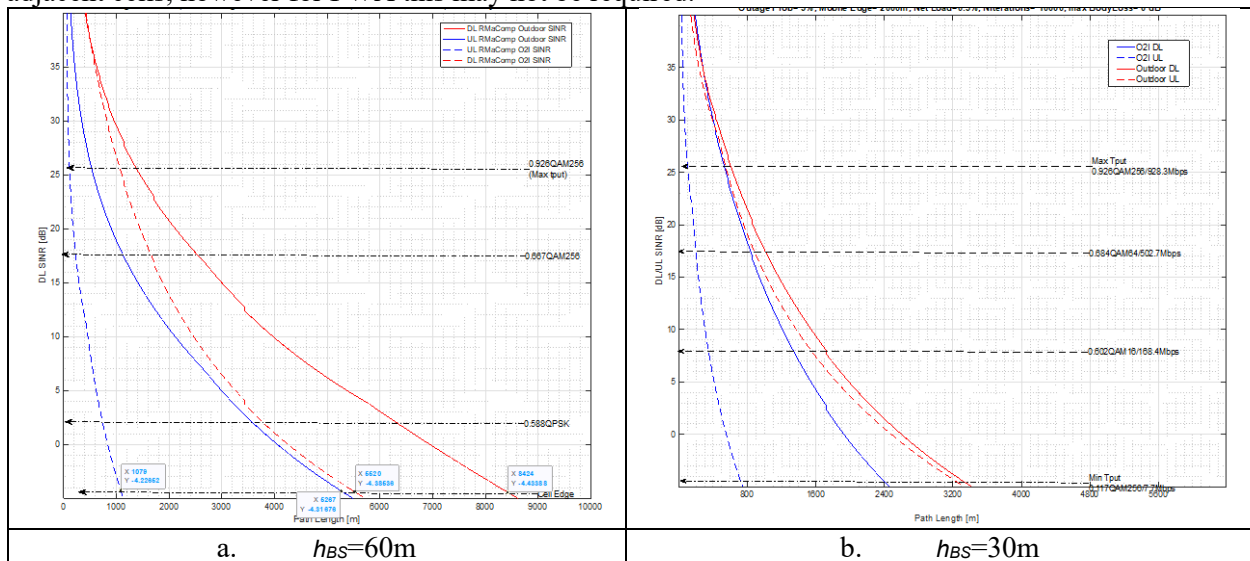


Figure 8. Comparative DL and UL SINR for Outdoor and O2I Scenarios (3.7GHz), $h_{BS}=60m$ vs. $h_{BS}=30m$

The coverage reduction (referenced to cell edge) is summarized in the following table. The cell coverage reduction is calculated against the respective MobEdge (for $h_{BS}=60m$ and $h_{BS}=30m$).

Table 11. Coverage and coverage reduction due to asymmetrical link budget (3.7GHz, $h_{BS}=60m$, 95% service availability, NetLoad=50%)

	$h_{BS}=30m$ (MobEdge=2000m)			$h_{BS}=60m$ (MobEdge=4000m)		
	Min(DL SINR)	Min(UL SINR)	Coverage reduction (%)	Min(DL SINR)	Min(UL SINR)	Coverage reduction (%)
Outdoor	3559m	3268m	16%	8424m	4316m	0%
O2I	2400m	720m	91%	5520m	1079m	92.8%

Observations:

- DL coverage is larger than the UL one for 3.7GHz, low and high RH height scenarios.
- Outdoor DL and UL coverage are in the same range for $h_{BS}=30m$ (3.7GHz) due the increased DL interference (MobEdge=2000m).
- The O2I UL coverage is more reduced vs. DL, due to reduced O2I DL interference and the reduced indoor CPE antenna gain vs. outdoor CPE gain.

We run the same analysis, comparing the cell coverage due to the UL link budget asymmetry, for different frequencies, as shown in Figure 9.

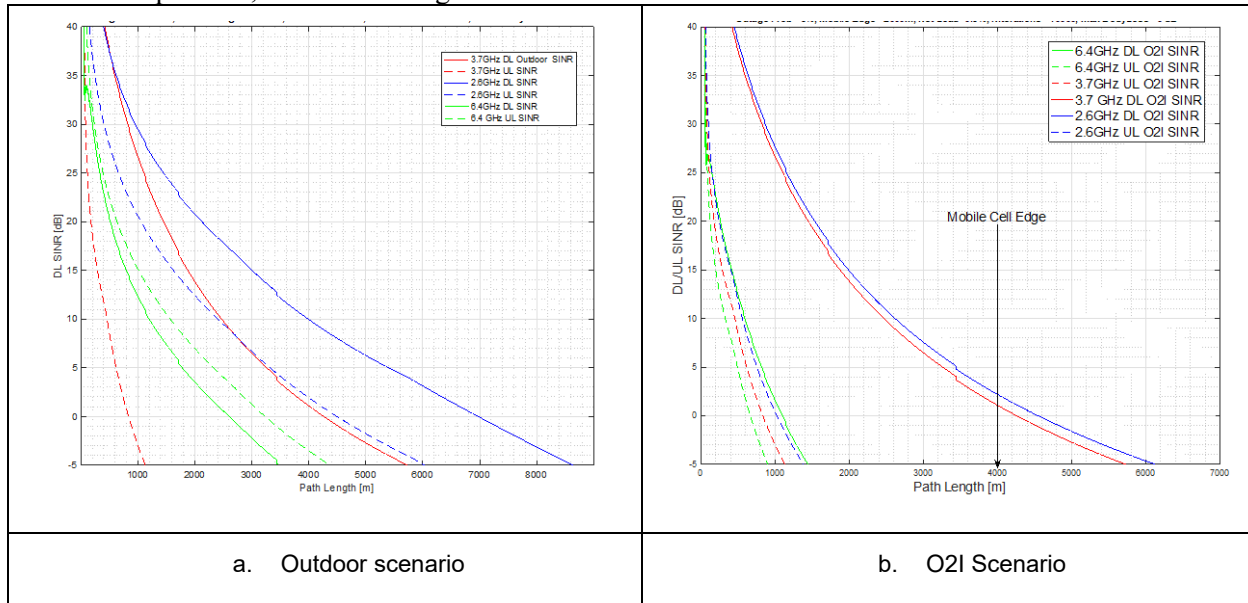


Figure 9. Comparative DL and UL SNR Degradation 2.6GHz vs. 3.7GHz vs. 6.4GHz, for Outdoor and O2I Scenarios (95% service availability, $h_{BS}=60m$).

Observations:

- The outdoor 2.6GHz coverage (both DL and UL) is slightly larger than the 3.7GHz.
- The DL interference limits the cell coverage for 6.4GHz (outdoor case), due to the close DL/UL EIRP difference (6dB) and increased UL coverage.
- The O2I coverage is severely limited by the UL coverage (poor UL link budget due to the additional O2I Loss).

3.2 Cell Coverage

The user data throughput is limited by the factors presented in section 3.1 and determined by the assumptions employed for this analysis:

- Subcarrier spacing=30kHz for all three frequencies
- 2000 slots/frame
- 4 beams/sector
- 3 sectors/cell
- 2 users/beam
- MIMO 2×2

The relationship between total user throughput and user throughput:

$$\text{CellThroughput} = \text{BeamsSector} * \text{Sector/Cell} * \text{User/Beam} * \text{UserThrupughput}.$$

Equation 8

where $\text{BeamsSector} * \frac{\text{Sector}}{\text{Cell}} * \frac{\text{User}}{\text{Beam}} = 24$. The DL comparative cell and user throughput for 2.6, 3.7 and 6.4GHz with 60m h_{BS} are analyzed.

3.2.1 Cell and User (Data) Throughput

The cell and user throughput plots are presented in Figure 10 (outdoor/O2I h_{BS} =60m). All throughput values represent user data (control and signaling data have been de-embedded).

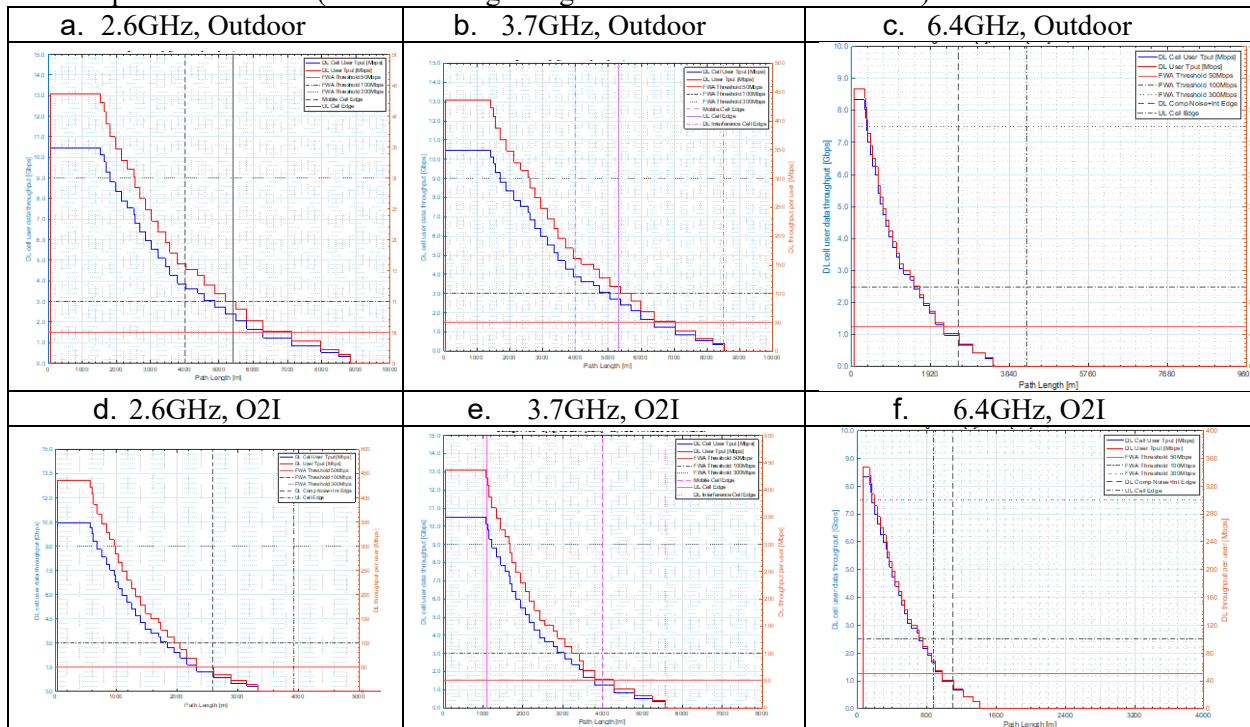


Figure 10. Comparative Cell (User Data) and Throughput per User (User Data Only), 2.6, 3.7 and 6.4 GHz, for h_{BS} =60m Outdoor and O2I Scenarios¹

It should be noted:

- DL cell edge accounts for the DL system interference
- The throughput jagged plots are generated by the link adaptation algorithm (QAM-256 MCS table).
- The UL System interference is minimal due to the lower CPE EIRP and lower CPE antenna height (increased obstruction probability).

¹ Network Load 50%, 95% service availability, VTilt=-15°

Similar plot shapes are obtained for all scenarios under consideration (30 and 60m, outdoor and O2I, 2.6, 3.7 and 6.4GHz). For the sake of brevity, we present the summarized data in Table 12 and Table 13, the limiting cell edge being highlighted. We analyze the limitations imposed on the cell edge by the following factors:

- DL cell edge, propagation and different fading and loss mechanisms limitations.
- Mobile cell edge (MobEdge) defined initially by the network planner (e.g., 2000m for $h_{BS}=30m$).
- UL cell edge limited by the asymmetrical link budget (DL driven).

There are 3 different cell edges. The relationship between them:

$$RealCellEdge = \min(DLCellEdge, MobEdge, ULCellEdge)$$

Equation 9

Table 12. Cell Coverage $h_{BS}=30m$ (95% Service Availability, NetLoad=50%, Vtilt=-15°)


Frequency [GHz]	Outdoor				O2I			
	DL Edge	MobCellEdge	UL Edge	Real Edge	DL Edge	MobCellEdge	UL Edge	Real Edge
2.6	3276	2000	3917	2000	2125	2000	1198	1198
3.7	3331	2000	3244	2000	2402	2000	723	723
6.4	1694	2000	2832	1694	1048	2000	560	560

Table 13. Cell Coverage $h_{BS}=60m$ (95% Service Availability, NetLoad=50%, Vtilt=-15°)

Frequency [GHz]	Outdoor				O2I			
	DL Edge	MobEdge	UL Edge	Real Edge	DL Edge	MobEdge	UL Edge	Real Edge
2.6	8846	4000	6398	4000	6002	4000	1344	1344
3.7	8461	4000	5339	4000	5572	4000	1097	1097
6.4	3436	4000	4249	3436	1395	4000	847	847

Observations:

- DL outdoor coverage is limited by the mobile cell edge (2.6 and 3.7GHz).
 - The network planner may allow a larger MobEdge than initially predicted
- DL outdoor 6.4GHz coverage is limited by the limited DL EIRP (36dBm)
- The O2I coverage is limited by UL cell edge, for all frequencies
 - The O2I coverage may require a higher CPE EIRP.
- While it may be expected to see a larger coverage for the 2.7GHz vs. the 3.7GHz case, the lower frequency advantage is partly offset by:
 - The higher interference impacting both the signal and network interference, but since $SINR=SNR-I$, the SINR impact may not be straightforward. A graphical explanation is provided in Figure 11 ($h_{BS}=60m$):






Figure 11. Comparative LOS, NLOS and Composite SINR and Tput, 2.6 3.7GHz²

We define

$$RealCellEdge = \min\{DLInterferenceEdge, ULEdge, MobCellEdge\}$$

Equation 2

Observations ($h_{BS}=60m$, $h_{BS}=30m^3$)

- 2.6GHz LOS coverage is smaller than the 3.7GHz one (both $h_{BS}=30m$ and $60m$), due to the higher network interference impact.
- NLOS becomes the dominant propagation mechanism because of the network interference is greatly reduced in NLOS conditions, 2.6GHz gets a larger RMa NLOS and RMa composite coverage. Overall, the 2.6GHz overall composite coverage is higher than the 3.7GHz one.
- The network interference ($h_{BS}=30m$) is higher for the 2.6GHz, due to the selected MobEdge (2000m).
- The outdoor 6.4GHz RealCellEdge is limited by the lower BS EIRP (36dBm).
 - The unlicensed 6GHz spectrum regulatory regulations restricts $\max(EIRP)=36dBm$.
- Providing NR-U services in unlicensed 6 GHz spectrum may require a densified network.
- The O2I cell edge is UL limited, due to the additional UL path losses, incurred due to the O2I loss. Under these assumptions, the indoor CPE is required either:
 - i) have a higher conducted RF power and/or
 - ii) use a directive antenna array with a higher gain vs. an omni antenna.

However, it should be noted that increasing the indoor CPE EIRP may also increase the indoor multipath, which may require a different analysis.

3.2.2 50, 100 and 300Mbps Coverage

We considered the 50, 100 and 300Mbps as three different traffic tiers for FWA services. We modeled the related coverage for these tiers, for 3.7 and 6.4GHz, outdoor and O2I scenarios. We also model the impact of different service availability rates upon coverage. The outdoor 3.7 and 6.4GHz coverage vs. service availability is plotted with and without UL limitation (Figure 12), for both $h_{BS}=30m$ and $60m$.

Observations:

² 95% service availability, 50% network load, 60 m h_{BS} .

³ Related plots available.

- Providing a higher target user rate availability (e.g., 99%) for outdoor 3.7GHz (same applies to 2.6GHz) a progressive coverage reduction may occur. The higher the target the service availability, the higher are the propagation losses and the shorter is the coverage.
- The higher the target user throughput (e.g., 300Mbps), the smaller is the related coverage.

The coverage reduction between 95% and 99% service availability is summarized in Table 14.

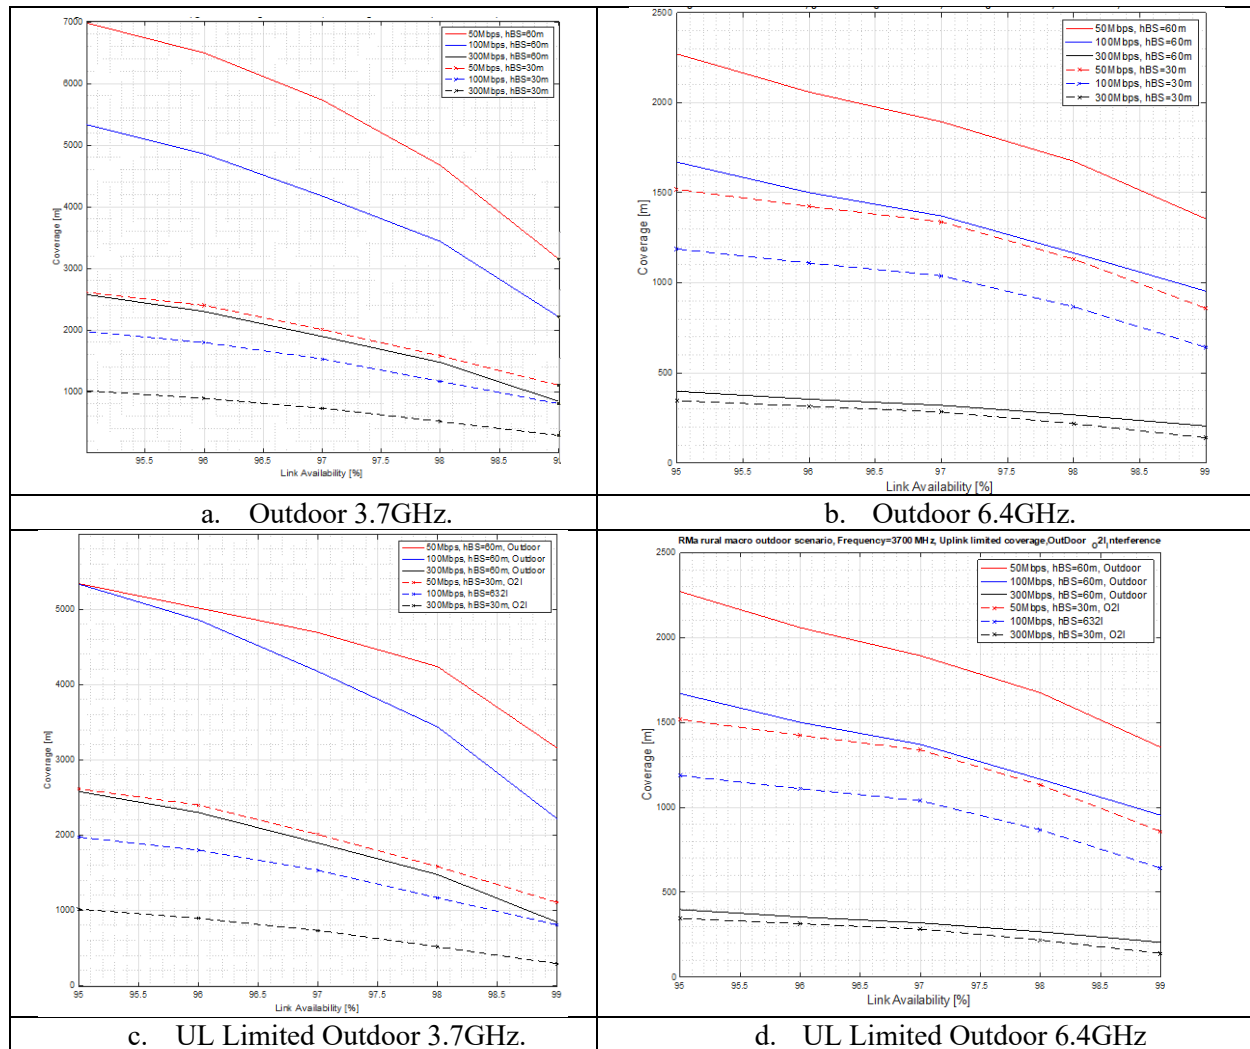


Figure 12. Outdoor (30 and 60m), 3.7GHz, 50, 100 and 300Mbps Coverage

Table 14. Coverage Reduction (50, 100 and 300Mbps), for Outdoor 3.7GHz

	$h_{BS}=60m$			$h_{BS}=30m$		
	50 Mbps	100 Mbps	300 Mbps	50 Mbps	100 Mbps	300 Mbps
95%	6984	5334	2577	2613	1972	1015
99%	3156	2216	806	1104	802	290
Coverage Reduction	80.6%	82.7%	90.2%	68.6%	83.4%	91.8%

Table 15. Coverage Reduction (50, 100 and 300Mbps), for O2I (3.7GHz)

	$h_{BS}=60m$			$h_{BS}=30m$		
	50 Mbps	100 Mbps	300 Mbps	50 Mbps	100 Mbps	300 Mbps
95%	4285	3235	1676	1937	1533	843
99%	2591	1944	1003	1152	880	437
Coverage reduction	63.4%	63.8%	64.1%	64.6%	67%	73.1%

The comparative O2I coverage reduction (3.7 vs. 6.4GHz), subject to UL link budget limitations, follows.

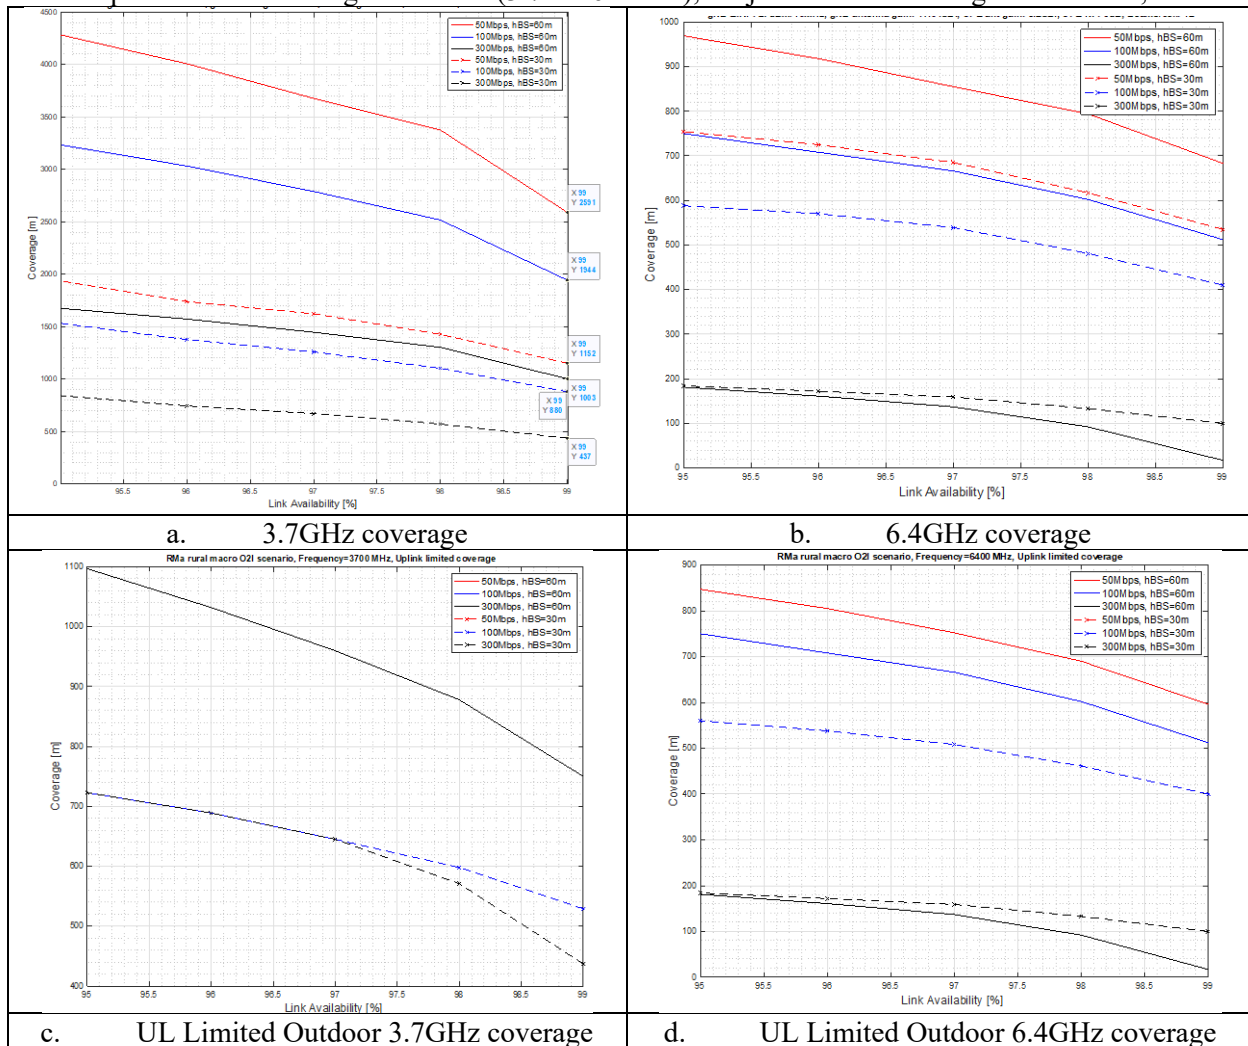


Figure 13. Comparative O2I ($h_{BS}=30$ and 60m) Coverage, for 3.7 and 6.4GHz, When Subject to UL Coverage Limitations

Observations:

- The O2I 3.7GHz coverage is severely UL limited, due the O2I loss.
- The O2I 6.4 GHz ($h_{BS}=60m$ and 30m) 50Mbps related coverage is UL limited, due the UL link budget impairment caused by the O2I fading.
- The UL indoor CPE may require a higher EIRP.

The cell coverage reduction between 95% and 99% service availability is calculated:

$$\text{CoverageReduction [\%]} = (\text{CellRadius}(99\%)/\text{CellRadius}(95\%))^2 \times 100$$

Equation 10

We analyze the comparative coverage reduction KPI, between 99% and 95% service availability.

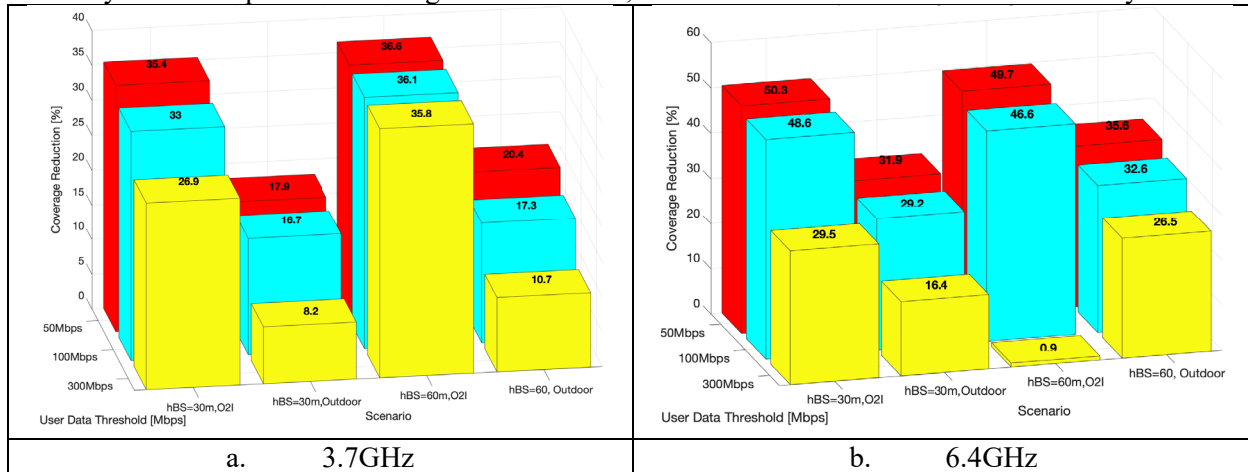


Figure 14. Coverage reduction between target 99% and 95% Service Availability, for 3.7GHz and 6.4GHz.

Observations:

- A severe coverage reduction occurs, when the service availability is increased from 95% up to 99%, for the outdoor, $h_{BS}=60\text{m}$ and 30m scenarios.
 - The outdoor propagation is impacted by large-scale and small-scale fading, resulting into a better coverage for service availability=95% and by a sharper reduction for 99%.
- The O2I propagation is impacted by three fading mechanisms (large-scale, small-scale fading and O2I loss), hence a relatively lower coverage (vs. outdoor coverage) for service availability=95%, and a milder reduction for 99%.
- For the 6.4GHz, O2I, $h_{BS}=60\text{m}$, 300Mbps case, the target service availability is practically non-existent (cell radius=17m), hence a severe coverage reduction degradation.

All of the above considerations highlight the challenges associated with high target service availability (e.g., 99%).

- However, due to the flexible 5G link adaptation algorithm, the user may not perceive the impact, since the user may not run consistently high data applications, close to the data threshold allocated for that user (e.g., most of the households may run applications demanding 30-60Mbps, while being subscribed for a 100Mbps plan).
- A FWA network planning targeting service availability=95% backed by NetLoad=50% may be a realistic target.

4 Conclusions

A variety of FWA scenarios was analyzed based on CableLabs' system simulations engine. The models used a comprehensive set of assumptions, presented in the Appendix. The simulations targeted a comparative technical performance of rural FWA networks, operating in 2.6, 3.7 and 6.4GHz, when operating with tall towers ($h_{BS}=60\text{m}$) and medium size towers ($h_{BS}=30\text{m}$).

4.1 The Impact of Coverage Limiting Factors

- A. The impact analysis of different limiting factors upon coverage indicates:
1. Vertical tilt of the BS antenna array, while supporting a Frequency Reuse 1 topology (3.7GHz)
 - Due to the higher directivity of the outdoor CPE array, A beam's coverage could extend into the adjacent cell, exceeding the target MobEdge (e.g., 4000m for the modeled case).
 - The effective user throughput degradation when the BS array is tilted down to -21deg (-6dB EIPR reduction vs. Horizontal Tilt alignment) from 0deg may not be significant (0.63% per user) but the network interference gets decreased by 3.7dB.
 - Frequency Reuse 1 could be effectively implemented.
 - A sharper vertical beam would reduce even more the Network Interference power radiated towards a victim cell, improving the victim user's SINR and related throughput. As a consequence, a larger subarray vertical size could be used in order to optimize even more FR1 coverage (e.g., 4 subarrays each subarray being 8×4×2 antenna elements).
 - If the BS antenna's tilt is controlled dynamically, as a function of the network load, the system interference, the cell coverage and user throughput could be further optimized.
 2. Comparative Network Load impact upon user throughput and network coverage:
 - The higher is the network load, the lower is the cell coverage and user/cell throughput.
 - The coverage and throughput degradation caused by the increased network load highlights the significance of controlling 5G antenna arrays, dynamically updating Vtilt as a function of network load, across a cluster of cells.
 3. BS Array height impact upon network interference.
 - The DL SINR is degraded, when RH height is reduced from $h_{BS}=60\text{m}$ (MobEdge=-4000m) down to $h_{BS}=30\text{m}$ (MobEdge=2000m):
 - Outdoor scenario: 9dB (2.6GHz), 8.3dB (3.7GHz) and 4.7dB (6.4GHz)
 - O2I scenario: 10.6dB (2.6GHz), 9.9 (3.7GHz) and 0.8dB (6.4GHz)
 - The 2.6GHz system is subject to the highest network interference, among all considered scenarios, due to the lowest propagation losses vs. higher frequencies.
 - The outdoor 6.4 GHz system operates as a quasi-interference free system, due to the lower EIRP density: 27dBm/10MHz vs. 50dBm/10MHz (2.6 and 3.7GHz)
 - The network interference is lower for the O2I case vs. outdoor, due to the O2I propagation.
 4. Small-scale, large-scale and O2I fading mechanisms' impact
 - Large Scale (shadowing component) fading has the strongest impact among the three mechanisms analyzed, due to the NLOS propagation (Rayleigh type fading).
 - The higher is the desired service availability, the higher is the composite fading impact upon the link budget.
 - The O2I Loss could become the driving factor of the composite fading if the CPE is placed deep inside the house (vs. the outer wall facing the BS) and/or other construction materials used for the outer wall.
 -
 - FWA could use a lower target service availability (e.g., 95%):
 - Even for a higher service availability, the user experience impact may not be noticeable (e.g., throughput), as long as the user may not use the highest achievable allocated user speed.

5. DL and UL SNR impact

- DL coverage is larger than the UL one for low and high RH height scenarios (3.7GHz).
- DL and UL coverage are in the same range for $h_{BS}=30m$ (3.7GHz) due the increased DL Interference (MobEdge=2000m)
- The Outdoor 2.6GHz coverage is slightly larger than the 3.7GHz.
- The DL interference limits the cell coverage for 6.4GHz (outdoor case), due to the close DL/UL EIRP difference (6dB/10MHz) and increased UL coverage.
- The O2I coverage is severely limited by the UL coverage (poor UL link budget due to the additional O2I Loss).

B. Cell and User coverage analysis

1. Cell and user throughput

- Outdoor 6.4GHz DL coverage is limited by DL EIRP (36dBm)

DL outdoor coverage is limited by the mobile cell edge (2.6 and 3.7GHz). In return, this supports a larger MobEdge.

- 2.6GHz LOS coverage is smaller than the 3.7GHz one ($h_{BS}=30m$ and 60m), due to the higher network interference impact.
 - NLOS becomes the dominant propagation mechanism; since the network interference is greatly reduced in NLOS conditions, 2.6GHz gets a larger Rma NLOS and Rma composite coverage. Overall, the 2.6GHz overall composite coverage is NLOS driven being higher than the similar 3.7GHz coverage.
- The network interference ($h_{BS}=30m$) is higher for the 2.6GHz, due to MobEdge (2000m).
- The outdoor 6.4GHz RealCellEdge is limited by the lower BS EIRP (36dBm).
- Providing NR-U services in unlicensed 6 GHz spectrum may require a densified network.
- The O2I cell edge is UL limited, due to the additional UL path losses, caused by the O2I loss. Under these assumptions, the indoor CPE is required either:
 - i) have a higher conducted RF power and/or
 - ii) use a directive antenna array with a higher gain vs. an omni antenna.

2. 50, 100 and 300Mbps coverage

- All service availability <99% ($h_{BS}=60m$, Outdoor, 3.7GHz) scenarios are impacted by the UL asymmetric link budget.
- The outdoor 6.4GHz case is subject to no impact by the UL link budget limitation, due to the lower DL EIRP (36dBm/80MHz).
- The higher the service availability target, the shorter is the coverage due to the higher path loss.
- The O2I 3.7GHz coverage is severely UL limited, due the O2I loss.
- The O2I 6.4 GHz ($h_{BS}=60m$ and 30m) 50Mbps related coverage is UL limited, due the UL link budget impairment caused by the O2I fading.

3. Service Availability

- Coverage is severely reduced, when service availability is increased to 99% (outdoor, $h_{BS}=60m$ and 30m).
 - The outdoor propagation is impacted by large scale and small-scale fading, resulting into a better coverage for service availability=95% and by a sharper reduction for 99%.
- The O2I propagation is impacted by O2I loss, hence a relatively lower coverage compared with the outdoor case.
- For the 6.4GHz, O2I, $h_{BS}=60m$, 300Mbps case, the target service availability is practically non-existent (cell radius=17m), hence a severe coverage reduction degradation.

- The user may not perceive the impact of high user data rate reduction, unless the user runs consistently high data applications (e.g., most of the households may run applications demanding 30-60Mbps, while being subscribed for a 100Mbps plan).
- A FWA network planning targeting service availability=95% backed by NetLoad=50% may be a realistic target.

5 Abbreviations

3GPP	3 rd Generation Partnership Project
BRS	<u>Broadband Radio service</u>
BS	<u>Base Station</u>
BW	Bandwidth
CBRS	Citizens Broadband Radio
CPE	Customer Premises Equipment
CDF	Cumulative Distribution Function
DL	Downlink
DMRS	Demodulation Reference Signal
EIRP	Effective Isotropic Radiated Power
FBR	Front-to-Back Ratio
FR1	Frequency Reuse 1
FWA	Fixed Wireless Access
gNB	5G base station
hBS	Height of the gNB antenna (remote head)
HPBW	Half Power Beamwidth
ISD	Inter Site Distance
KPI	Key Performance Indicator
LDPC	Low Density Parity Coding
LLS	Link Level Simulator
LOS	Line-of-sight
MCS	Modulation and coding scheme
MHz	MegaHertz
MobEdge	Network planner cell edge target
NLOS	Non-LOS
NR	New Radio
O2I	Outdoor to Indoor
PRB	Physical Resource Block
RH	Remote Head
SCS	Sub Carrier Spacing
SCTE	Society of Cable Telecommunications Engineers
SINR	Signal to Noise and Interference Ratio
SNR	Signal to Noise Ratio
SLL	Side Lobe Level (main side lobe level vs. main lobe boresight)
SLS	System Level Simulator
TDD	Time Division Duplexing
Tput	Throughput
UCA	Uniform Circular Array
UL	Uplink
URA	Uniform Rectangular Array

6 Bibliography & References

1. 3GPP TR36.942, “Radio Frequency (RF) system scenarios,” July 2020, Technical Report.
2. 3GPP TS38.901, “Study on channel model for frequencies from 0.5 to 100GHz,” December 2019, Technical Specification.
3. [38.101-1, NR; User Equipment \(UE\) Radio Transmission and Reception; Part 1: Range 1 Standalone](#), 3GPP Technical Specification, July 2021
4. [38.104, NR; Base Station \(BS\) Radio Transmission and Reception](#), 3GPP Technical Specification, July 2021
5. [38.214, NR; Physical Layer Procedures for Data](#), 3GPP Technical Specification, June 2021
6. [38.211, NR; Physical Channels and Modulation](#), 3GPP Technical Specification, June 2021
7. [38.901, Study on Channel Model for Frequencies from 0.5 to 100 GHz](#), 3GPP Technical Specification, January 2020
8. [P.676-11, Attenuation by Atmospheric Gases](#), ITU Radiocommunication Sector Recommendation, September 2016
9. [P.840-6, Attenuation Due to Clouds and Fog](#), ITU Radiocommunication Sector Recommendation, September 2013
10. [P.838-1, Specific Attenuation Model for Rain for Use in Prediction Methods](#), ITU Radiocommunication Sector Recommendation, October 1999
11. TIA-10, Interference Criteria for Microwave Systems, Telecommunications Industry Association, May 2019
12. Dorin Viorel, Ruoyu Sun, Sanjay Patel, “5G FWA Technical Performance Analysis for Mid-Band Small Cell Networks,” White Paper, CableLabs, Sept. 2021. [[Link](#)]

7 Appendix

7.1 Simulations Assumptions

Table 16. System And Cell Simulations Assumptions

SYSTEM	VALUE	CELL	VALUE
System interference	Per SLS feed	Service Availability (%)	95
Cluster of cells PLOS	As defined by [20]	Sector/Cell	3
Network traffic load (%)	25/50/75	Beam/Sector	4
O2I propagation scenario	O2I residential (TR38.901)	Carrier aggregation	1
Channel model	3GPP TR38.901	Cell edge SINR (AWGN driven) (dB)	-4.54
Number of SLS iterations	100,000	MIMO	2x2
Max body loss (dB)	0	Air layer (MIMO) EIRP reduction MIMO x2 [dB]	-3
NLOS small-scale fading	Rayleigh	O2I path length (behind outer wall) (m)	1
LOS small-scale fading	Rice, $K=12$ dB	O2I wall material	Wood
O2I large-scale fading	$N\{\text{mean } 9.35, \text{sigma } 4.4\}$	Glass/outer wall ratio	0.3
RF Waveform polarization angle	Cross-Polarized	Central frequency (MHz)	2596/3700/6400
NR band	n41, n77, n96	Link Adaptation	Enabled
Mobile cell edge (m)	2000 ($h_{BS}=30\text{m}$); 4000 ($h_{BS}=60\text{m}$)	Modulation implementation loss	3
ISD [m]	3640 ($h_{BS}=30\text{m}$) 6920 ($h_{BS}=60\text{m}$)		
Frequency Reuse	1		
Interference model	DL		

Table 17. BS AND CPE Array Simulations

BASE STATION	VALUE	CPE	VALUE
Antenna array		Indoor Antenna array	
Subarray	4x4x2	Outdoor Antenna array	URA 2x8x2
DL MIMO rank	2x2	Outdoor antenna element	Cross-Dipole
Antenna element	Cross-Dipole	Indoor CPE height [m]	2
SubArray structure	4x4x2	Outdoor CPE height [m]	4
Antenna height above clutter (m)	30 or 60	Outdoor array boresight Gain (3.7GHz)	16.4
Antenna array Tilt [°]	-15	Outdoor Azimuth HPBW [deg]	16.4
Subarray boresight gain 3.7GHz [dBi]	17.0	Indoor antenna gain (3.7GHz)	8.2

Table 18 PHY/RF Assumptions

BS	Value	CPE	VALUE
rmsEIRP/10MHz [dBm]	50	rmsEIRP (dBm)	30
Active users/beam	2	Noise figure (dB)	6
Sub Carrier Spacing [kHz]	30	PHY Oversampling ratio	X4
Slots/Subframe	2	DMRS symbols	1
Subframes/Frame	10	User symbols	1
TDD ratio	11:2:1		
DL Control (PCCh+DMRS) syms	2		

Table 19. Atmospheric/Environment Conditions Assumptions

ENVIRONMENT			
ITU rain region	Disabled	Average House Height [m]	8
Slanted path profiles	Disabled	Average Street Width [m]	20
Crane rain region	B2		
Atmospheric pressure	Sea level		

BS /CPE array assumptions could be found in section 2.1.2.

7.2 BS and CPE Antenna Array Performance Plots

7.2.1 BS Array

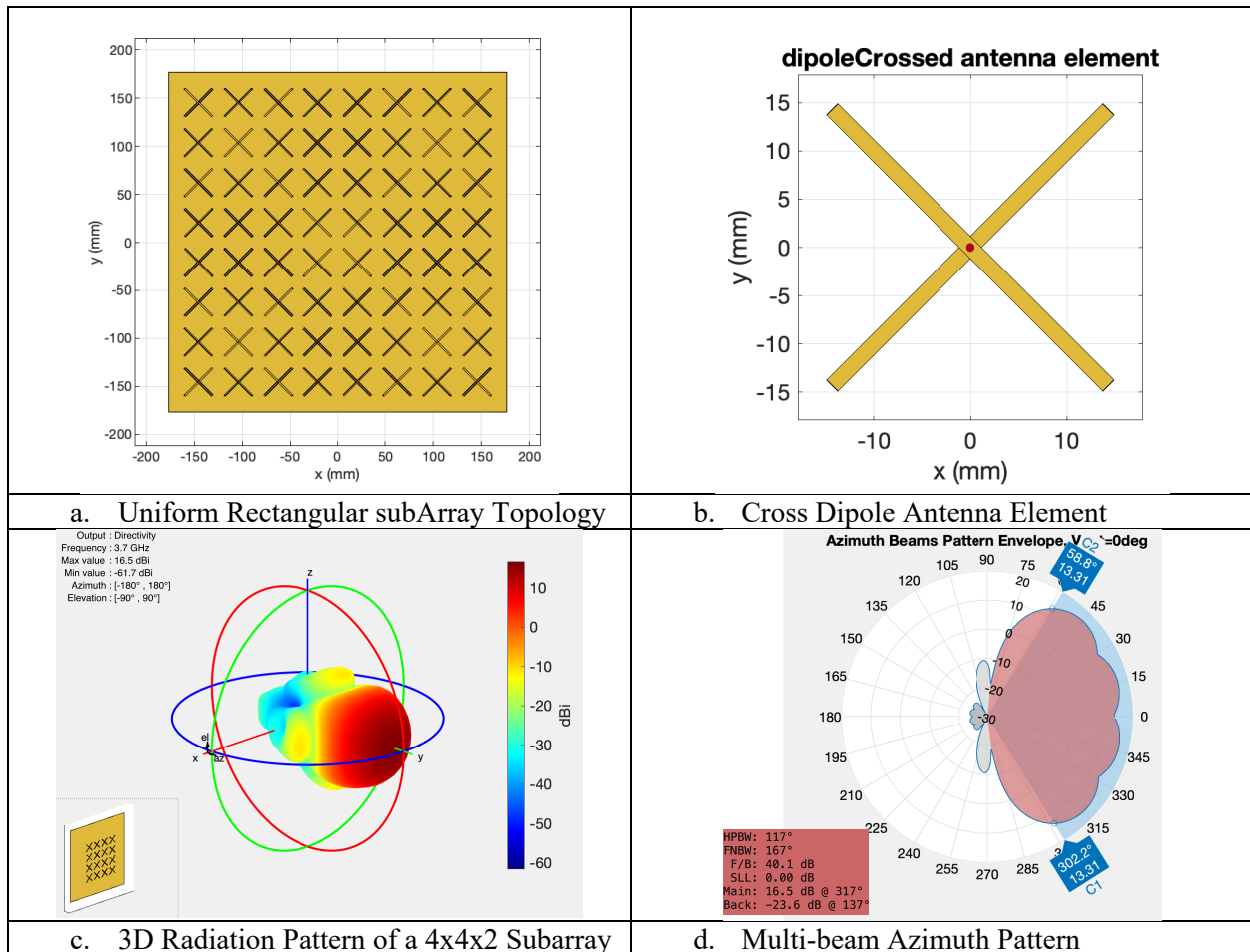



Figure 15. Sample of BS array parameters (3.7GHz). (a) Array geometry [2 2] subarrays, (b) Cross-Dipole antenna element geometry, (c) 3D radiation of a subarray (4x4x2) and (d) multi-beam azimuth radiation pattern.

7.2.2 CPE Array



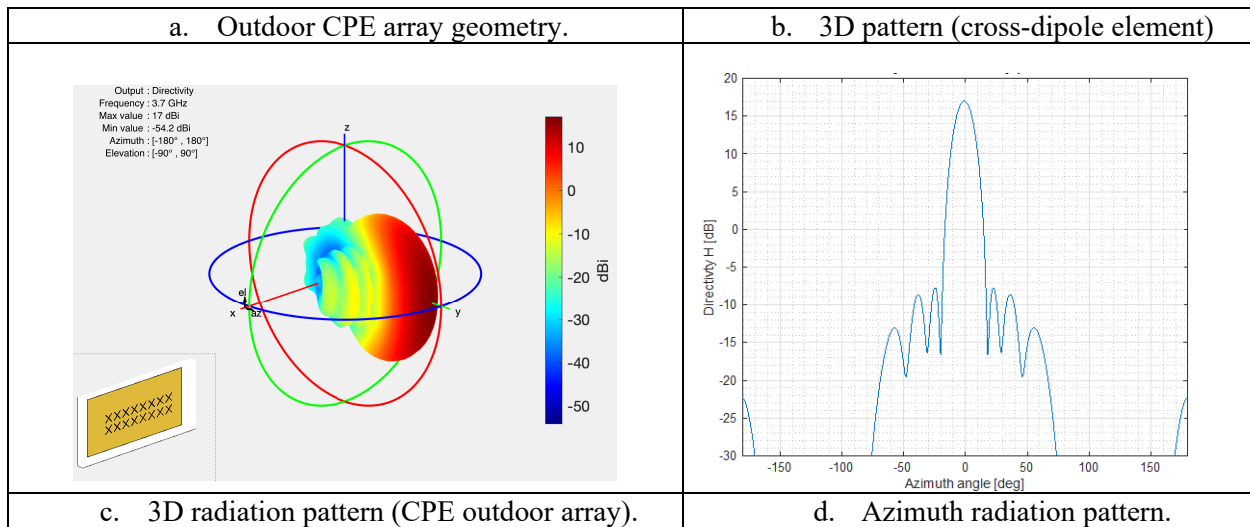


Figure 16. Sample of Outdoor CPE Array Parameters (3.7GHz). (a) Array Geometry, (b) 3D Pattern of the Cross-Dipole Antenna Element, (c) 3D Radiation Pattern of the CPE Array and (d) Azimuth Radiation Pattern (Rectangular Coordinates).

Composite Quality Metric KPIs

A General Framework for Interpretable Composite Metrics

A Technical Paper Prepared for SCTE by

Rohan Khatavkar

Principal Data Scientist I
Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
Rohan.Khatavkar@charter.com

Ryan Lewis

Manager Data Science
Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
Ryan.M.Lewis@charter.com

Veronica Bloom

Director Data Science
Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
+1 720-699-3798
Veronica.Bloom@charter.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Technical Approach	3
2.1. Methodology Approach Overview	3
2.2. Bucketing Metric Components	4
2.3. Metric Construction	5
2.4. Optimize Metric Weights and Model Selection	5
3. Conclusion.....	7
Abbreviations	8
Bibliography & References.....	Error! Bookmark not defined.

1. Introduction

We propose a novel method for constructing a composite metric for measuring product quality that is predictive/correlated with bottom-line key performance indicators (KPIs) such as customer churn rate or customer engagement. Traditionally, companies would track several individual metrics related to performance of a product. For example, to track customer experience across a hybrid-Mobile Virtual Network Operator (MVNO) network, metrics such as received signal strength, latency, throughput and link speed etc. are fundamental metrics that may be tracked in a dashboard. Subsequently, product owners can get a health check on the product by visualizing these metrics across different applications and across time. However, for deriving actionable insights, product owners would need to know if moving such individual KPIs (quality levers) would move the needle on a bottom-line metric such as customer churn.

AB testing would have to be conducted on a limited sample to derive such causal inferences. However, a bottom-line metric such as churn is a function of non-controlled exogenous variables such as macro-economic factors, competitor product, etc. Hence directly looking for difference in churn rate in test vs control sample would be difficult and hard to influence via a particular change in the product experience. A better approach would be to move the needle on a composite quality metric (CQM) constructed from individual KPIs which are essentially product levers that can be influenced by the product owners. The construction of the individual KPIs and their relative influence (weights) can be chosen such that the composite metric is predictive/correlated with bottom-line metrics. To achieve this, one could build a generalized linear model (e.g., a logistic regression model predicting churn) where the features can be the individual KPIs and the prediction equation will be the composite metric. However, such a metric will often be difficult to interpret for product owners unless they have prior training in statistical modeling.

In this paper, we propose a general framework for creating such an interpretable composite quality metric. The components or individual KPIs are discretized on a Fibonacci scale 8, 5, 3, 2, 1 such that a score of 8 can be interpreted as the best, 3 as bad and 1 as the worst experience. The weights for the components are selected via numerical methods such that the correlation with the bottom-line KPI is measured for each weight combination. The weight combination with the best correlation and non-skewed distribution can be selected as optimal. The result then is a single composite metric constructed from multiple components that is predictive of a bottom-line KPI.

2. Technical Approach

2.1. Methodology Approach Overview

Any customer-impacting initiative to improve product experience is a hypothesis that the product team would prefer to test on limited samples to answer key questions: 1) Did the initiative move the needle on the bottom-line KPI or not? 2) Did the initiative negatively impact customer experience as measured by safety metrics such as increased call volume? Conducting an A/B test on a random sample helps to ascertain the risk vs reward of rolling out a product change to the larger customer base. Since the engineering lift and resources involved in an A/B test are non-trivial, it is critical that the metrics for measuring success of an initiative are clearly defined and influenceable.

However, moving the needle on bottom-line KPIs can be difficult to do. On the other hand, moving the needle on individual KPIs that are quality levers begs the question: 1) What relationships do these individual KPIs have with the bottom-line KPI, if any? and 2) What is the relative importance of each individual KPIs? A composite quality metric gives us exactly this relationship between a bottom-line KPI such as churn, customer engagement, revenue, etc. and individual KPIs which are levers that a product

team can influence: received signal strength, latency, throughput, etc. The following three sub-sections provide a general framework for creating a composite quality metric: 1) Bucketing Metric Components, 2) Optimize Metric Weights, and 3) Model Selection.

We make an important call-out that for relatively mature and complex systems, unmitigated improvements are often impossible. That is, it is often impossible to improve one or more of the base KPIs without some kind of negative impact on others. This puts teams looking to improve the product in an ambiguous situation, which can be addressed by formulating a composite metric that makes explicit the effect of these tradeoffs on customer experience. Making explicit the tradeoffs also helps with institutional information. That is, as the organization's view on what is important to customers change, it is encoded in the CQM and can be tracked over time.

2.2. Bucketing Metric Components

For each individual KPI, we recommend partitioning the distribution of the values based on a large sample. For example, equal partitioning into fifths based on the 20th, 40th, 60th and 80th percentile can be used as thresholds to get roughly equal proportion of distribution in five buckets. This allows us to discretize the individual KPIs on a Fibonacci scale 8, 5, 3, 2, 1 such that:

1. Bucket 8: greater than 80th percentile (interpreted as the best experience)
2. Bucket 5: 80th – 60th percentile
3. Bucket 3: 60th – 40th percentile
4. Bucket 2: 40th – 20th percentile
5. Bucket 1: below 20th percentile (interpreted as the worst experience)

The advantages with the discretized set up are that: 1) product owners can get a sense of the gradation of the customer experience in a simple and consistent fashion across metrics on different scales, and 2) robustness to outliers. Even though we are proposing the Fibonacci scale 8, 5, 3, 2, 1, one could opt for a different scale such as 10, 8, 6, 4, 2, with the same strategy for thresholds. The primary goal is to pick thresholds based on a large sample data such that we avoid concentration of a point mass around one bucket.

In some cases, the discretization has to be specific to a product feature. For example, say we want to build a CQM to measure the effectiveness of the search feature in a particular streaming service. If one of the individual components is average number of characters typed in a search bar then it is important to consider differences in platforms that affect speed of input by the customer. On “1-foot” platforms such as mobile phones, tablets, and laptops it is relatively easy for most customers to enter more search characters compared to “10-foot” platforms such as Roku, Apple TV, Samsung TV, etc. In such a case, instead of creating different iterations of the same composite metric by platform type, it is efficient to simply customize the discretization by product feature. The goal of the product team would be to implement product enhancement that improve one or more components, preferably with high relative weight. As a result, certain components of the composite metric might change in distribution over time, shifting from lower to higher buckets. Therefore, it is worth checking the shift in distribution of the components for re-calibration of the thresholds, especially after a significant product change.

2.3. Metric Construction

Once the thresholds are computed to discretize the individual components, we can combine the components in one composite metric similar to a weighted average framework. We propose the following general formulation:

$$CQM = \left(\frac{\sum_{i=1}^n w_i C_i}{8 * \sum_{i=1}^n w_i} \right) * \prod_{j=1}^k A_j \quad (1)$$

where:

CQM stands for composite quality metric

C_i is the ith component of the composite metric with possible values (8,5,3,2,1)

w_i is the weight assigned to the ith component

$\sum_{i=1}^n w_i C_i$ *is the weighted average of the n components*

$8 * \sum_{i=1}^n w_i$ *is the normalization factor so that the range of CQM is between 0 and 1*

A_j is the jth binary filter with possible values (0,1)

$\prod_{j=1}^k A_j$ *is the product of all k binary filters*

The intuition behind incorporating binary filter terms is to allow a pass/no pass gate for the CQM depending on success of critical events for acceptable customer experience. For example, in a search quality metric, if the page fails to load or search API itself fails then we would want the score to be zero. Alternatively, if there is no page load failure and no search API failure then the metric is allowed to “pass” and would receive a score as per the first term in equation 1.

The normalization factor, which allows for the range of the CQM between 0 and 1 or alternatively between 0% and 100%, is intuitive for product owners to track overtime in a dashboard.

One caveat to note is that the CQM isn't going to be an absolute measure but rather a relative measure aimed at exposing areas for improvement. For a relative measure a low score doesn't necessarily mean the users had a bad experience, and a good score doesn't mean they had a good experience either. Since it is relative to the population base, it will necessitate updating the CQM as the experience of the population of users evolves.

2.4. Optimize Metric Weights and Model Selection

We introduced the construction of the CQM in the previous section. Given typical time constraints on feature engineering and productionizing such a metric, one could assume equal weights, for example 0.5,

for the components and start tracking a version 1 of the CQM. Alternatively, one could increase or decrease the weights based on prior product knowledge. However, the most optimal way to estimate weights would be benchmarking the CQM against a bottom-line KPI for several combinations of the component weights.

We recommend the following choices of weights for each of the n components: (0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0). If the metric has n component, that gives us a grid of 10^n combinations of weights to test for determining the optimal combination. The optimization steps are:

1. For each combination, calculate the CQM as shown in equation 1;
2. Determine correlation with the bottom-line KPI;
 - a. An interpretable way of determining correlation, which can also make it transparent to stakeholders, is bucketing the CQM score in 5 or 10 equal buckets.
3. Subsequently, calculate the average of the bottom-line KPI in each bucket;
4. A simple linear fit can give an idea about the strength of the linear relationship. We can also quantify the strength of the relationship by calculating the R^2 . In addition, the p-value of the slope can help weed out combinations that are not significant;
5. The combination with best R^2 and resembling a bell-shaped distribution of scores can be selected as optimal. The shape of the distribution of scores may not be bell shaped but as long as it is not skewed, we can select the combination as optimal.

As an example, if the optimal combination is found for the CQM, then as we move from a lower bucket of scores to a higher bucket, we should see higher customer engagement as measured by a bottom-line KPI such as customer churn. The advantage of keeping the selection process visual is giving ample transparency to stakeholders about the methodology. An output from all combinations tested can simply be visualized in a spreadsheet. The combination that has the best relationship with the bottom-line KPI while avoiding a skewed distribution of scores can be selected.

One caveat to this numerical approach is that the computational time is highly sensitive to the number of components. Since the number of combinations in the grid is 10^n beyond three components the compute time is unrealistic and unnecessary. In such a case, we advise reducing the combinations to be computed by taking the following approximate approach as shown for a metric with five components:

$$w_1 = (0.2, 0.4, 0.6, 0.8)$$

$$w_2 = (0.1, 0.3, 0.5, 0.7)$$

$$w_3 = (0.2, 0.4, 0.6, 0.8)$$

$$w_4 = (0.1, 0.3, 0.5, 0.7)$$

$$w_5 = (0.2, 0.4, 0.6, 0.8)$$

Even though we alternated through the even and odd sequence of weights, we cover most of the range of possible weights. By using this approximation, we reduced the number of combinations from 100K to 1,024. Also note that the relative combinations of weights are more influential than the absolute combination itself. In addition, we have observed that having a finer grid does not tend to lead to a drastically better metric.

We also recommend performing a basic regression analysis where we regress the individual components post-bucketing as regressors against the bottom-line KPI. This will help us understand which components are significant prior to running the optimization. In addition, we can possibly eliminate non-significant components or substitute for different components. One could argue using the prediction equation from the regression itself as the CQM. However, it is often difficult to explain the interpretation of such a metric to stakeholders who may be unfamiliar with statistical modeling. The advantage of the metric structure as laid out in equation 1 is that it is additive in nature and can be easily broken into components for further investigation by product owners.

3. Conclusion

In this paper, we have introduced a general framework for constructing a composite quality metric which is easy to interpret for stakeholders, effectively combines individual component metrics, and allows for benchmarking against a bottom-line KPI. The composite metric can be further used for testing product improvement initiatives via AB testing. A series of successful experiments that move the needle on the CQM helps make a case for improving the bottom-line KPI over time. The process of product improvements is iterative and the CQM is therefore appropriate as a relative metric. Depending on the industry, product maturity and changing consumption behavior, the threshold for what constitutes “good” quality will change. Hence, a relative measure is more useful than an absolute measure. The framework proposed in this paper is flexible so that the thresholds for the 8, 5, 3, 2, 1 buckets and subsequent re-weighting of the components to align with the bottom-line KPI can be conducted on a regular cadence. We also recommend including (or excluding) components as the range of quality levers at disposal changes to drive product improvement.

Abbreviations

API	Application Programming Interface
CQM	Composite Quality Metric
KPI	Key performance indicator
MVNO	Mobile Virtual Network Operator
OTT	Over the Top
R2	Coefficient of determination, or R^2

Considerations For the Delivery of Latency-Sensitive, Compute-Intensive Experiences Over a Communication Network

A Technical Paper prepared for SCTE by

Dhananjay Lal
Senior Director, Advanced R&D
Adeia
3025 Orchard Parkway, San Jose CA 95134
(513) 225 4948
Dj.lal@adeia.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. The Need for a Different Kind of Compute: Interactive Content	3
3. The Communication Service Provider Network Hierarchy: Architecture and Considerations	6
4. Delivering Latency-Sensitive Compute for AR/VR Experiences.....	8
4.1. A Priori Setup, CSP.....	8
4.2. A Priori Setup, Application/Media Service Provider	9
4.3. Connection Setup.....	9
5. Discussion of Compute Grade and Latency	12
6. Conclusion.....	13
Abbreviations	14
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1- Non-interactive content delivery	4
Figure 2 - Interactive content delivery.....	4
Figure 3- Sample Cloud Gaming System with latency budget	5
Figure 4- Sample Cable/Telecom Network.....	7
Figure 5 - Sample Session Setup	11
Figure 6 - Example System Specifications for Delivering Experiences	13

List of Tables

Title	Page Number
Table 1- Network Orchestrator "Available Compute Resource" Table	10
Table 2 - Application Service Provider Content Table	10
Table 3- Network Orchestrator "Network Hierarchy – Latency" Table.....	10

1. Introduction

For Augmented Reality (AR) and Virtual Reality (VR) devices to become truly immersive experiences, the industry is investing in R&D focused to address how they can achieve a goal of 20 millisecond (ms) motion to photon (MTP) latency [1] by modelling latency in concert with other systems.

In cloud gaming and other cloud content-streaming scenarios, 100ms “button-to-photon” latency is deemed adequate and is achieved through cloud rendering and bit rate encoding specified at the content source. While some latency variations occur in the transmission path, depending on network conditions, this 100ms “button-to-photon” is usually sufficient for these applications. But for AR/VR, prior art has made the case for 20ms MTP latency to avoid spatial disorientation, motion sickness and other adverse experiences for users.

In Cloud gaming and VR/AR, viewers are not passive—their interactivity (e.g., with objects and avatars) requires active modification of the content within the experience itself. In essence, this means that for non-interactive video experiences, content delivery is a “one-way” transfer of bits over the network (after transcoding) to the client (viewer), whereas interactive content delivery is two-way, where the viewer input from the client is used to determine subsequent content, which must then be rendered prior to transcoding and transmission. This imposes tighter latency requirements in AR/VR content rendering. Achieving 20ms MTP is practically impossible with current network architectures unless the cable service provider makes compute resources available deeper in its network.

This paper discusses a practical approach to cable delivery architectures where compute intensity and latency become the primary determinants of the end-user experience and are inherent properties of the content served to the subscriber. It describes a method and system for delivering latency-sensitive, compute-intensive experiences over a network that allows communications service providers (CSPs) to deliver latency-dependent compute to a subscriber, in concert with an application service provider (ASP), by employing seamless interactions between system components.

2. The Need for a Different Kind of Compute: Interactive Content

Current IP-based media delivery, for example, live and VOD (video on demand) TV content, has a latency sensitivity of ~500ms to a few seconds, dependent on various considerations, most importantly network conditions and buffer size at the client. While CSPs that offer programming can serve media experiences from within their own network, i.e., from a server in close proximity to their customers, there is little competitive advantage to be derived for the customer experience. This is evidenced by the rise of OTT (over-the-top) media service providers that provide programming via servers outside the CSP network. Other benefits, however, accrue with customer proximity to the CSPs such as reduced transit costs with shorter backhaul traffic, as well as access to consumer profile data.

In the context of current programming, we may logically separate the CSP that provides the packet route to the viewer from the Cloud, which we may view as a logical entity where content resides.

To explain an IP media delivery architecture, we may abstract out the network from the point of view of the media source in the Cloud and the consumer in the home. The original source file is encoded (compressed) at various bit rates, stored as a multi-bitrate asset (or chunks of the asset, each of which is multi-bitrate) and then delivered to a client based on requested parameters and network conditions, where it is then decoded and displayed.

Encoding, in this context called *transcoding*, remains the cornerstone of 2D media delivery on flat screens, big or small. At the server, transcoding may be performed by a combination of software and hardware, including CPUs, GPUs, FPGAs, etc. Depending on whether the programming is a live-stream broadcast or a VoD asset, the selection of transcoding platform and architecture as well as network protocols can vary significantly.

It is important to note that current media delivery architecture is non-interactive when it comes to content modification. While the quality with which the content is delivered may vary based on network conditions, and ads may be inserted through instantaneous decision-making in local zones, the actual programming content typically does not change based on viewer input.

Interactive content, on the other hand, can be altered by a viewer. In this sense, the viewer is not passive, but rather is immersed in the experience. Examples include games, Augmented Reality/Virtual Reality (AR/VR) experiences and holograms rendered on 3D displays. Interactivity may include content modification, for example, shooting and killing a formidable foe in a FPS (first person shooter) game that prolongs a player's ability to remain in the game (the content), but also content rendering in real-time encoded AR/VR based on the user's viewport, which controls what the user decides to focus on in their real (in case of AR) or make-believe (in case of VR) environment, ultimately changing what is rendered and displayed "on-the-fly". Streaming 360° video VR is a notable exception, where the viewport based on the user's pose is picked from within the pre-encoded video frame sent to the Head-mounted Display (HMD). It is non-interactive because the content is neither rendered nor encoded in real-time.

The key differences between non-interactive media like television and interactive media like gaming/AR/VR may be described as follows. Figure 1 and Figure 2 describe the abstracted delivery of media for non-interactive content versus interactive content.




Figure 1- Non-interactive content delivery

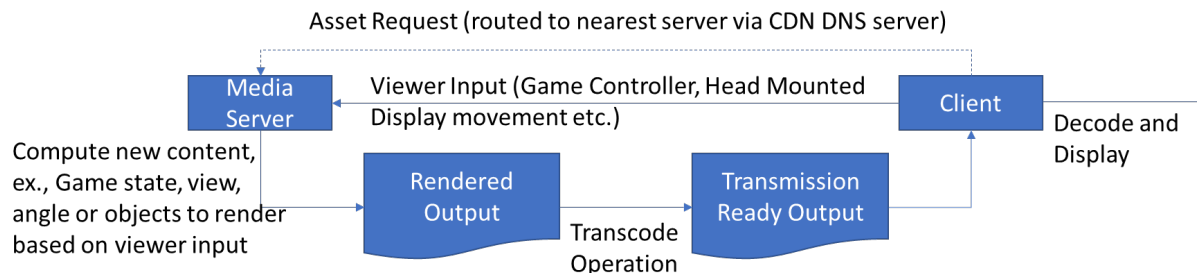


Figure 2 - Interactive content delivery

As mentioned, the key difference between non-interactive and interactive content is that non-interactive content delivery is a "one-way" transfer of bits over a network (after transcoding) to the client (viewer),

whereas interactive content delivery is a two-way street, where the viewer input from the client is used to determine subsequent content, which must be rendered prior to transcoding and transfer over the network.

In the case of non-interactive content, provided there are reasonable network conditions to deliver the bitstream to the client, and the client can buffer for intermittent interruptions in the network or processor use, the viewer gets a good experience. However, for interactive content, the client is typically expecting a higher bitrate (due to higher frames per second and resolution) while at the same time expecting the server to accept and process its current input to determine the subsequent state of the content.

This imposes much tighter latency requirements on interactive content rendering. Using gaming as an example, if the frame displayed at the client is more than three to six frames after the game input was issued, then the user experiences “lag”, i.e., feels that the game is non-responsive. Typical game experiences are served at 60 frames per second (fps) which gives the system 50-100ms to present a frame response back to the user from the time that the user issues their input.

There are many public discourses about optimizing latency for Cloud Gaming [2][3]. Figure 3, from [Nokia](#), is a sample representation of the system latency budget for Cloud gaming [4].

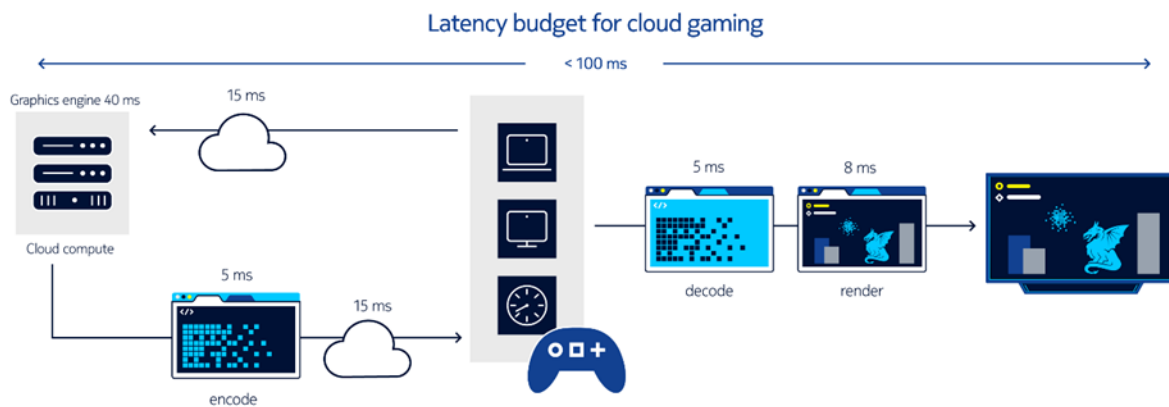


Figure 3- Sample Cloud Gaming System with latency budget

In the above example, the complete response time, i.e., the “button-to-photon” latency, is depicted as 88ms, with the network component of the latency being 15ms one way, or 30ms round-trip-time (RTT).

For a network, meeting 30-40 millisecond RTT latency is a tall order. For example, in a particular home, based on current network conditions, may take 5ms over WiFi and 25ms for DOCSIS, but it could take another 20ms to reach the Cloud due to routing over the backbone network, followed by egress to another carrier network or the Cloud gaming content delivery network (CDN), and ultimately, to the gaming server, i.e., the graphics engine – we assume that the gaming server is not in the same market or Metropolitan Statistical Area (MSA) as the home.

In this example, it becomes challenging to provide such interactive experiences from the Cloud. The unique position of CSPs to affect the quality of experience for services such as Cloud gaming by hosting the service from within their network, coupled with the financial incentives of higher ARPU, is driving some CSPs to bring these services to market.

However, Cloud gaming is only the beginning of immersive media delivery. AR and VR are on the horizon and the rumored launch of Apple Glasses promises to bring these services to mainstream adoption. AR and VR are, however, far less latency-tolerant than Cloud gaming. While public documentation proclaims that 100ms “button-to-photon” (system) latency is acceptable for Cloud gaming, the community of practice has concluded that VR will require 20ms “motion-to-photon” (system) latency. As seen from Figure 3, if a graphics engine takes 40ms to produce the next frame, or if the necessary pipeline steps of encoding and decoding take 10ms, then an exacting standard of 20ms “motion-to-photon” would be nearly impossible even if the network RTT latency is negligible.

In reality, there are several application-level techniques to “relax” the 20ms motion-to-photon latency requirement, such as re-projection, asynchronous time-warp, over-rendering, head movement prediction, etc. This has been validated by our empirical R&D. The type of content, “lean back” vs “lean forward”, also determines whether this “motion-to-photon” latency is a strict or relaxed requirement. For example, an experience such as “theBlu” [5], in which the user moves relatively slowly and interacts with elements in the content infrequently, is “lean back”, and the user will tolerate greater “motion to photon” latency for an acceptable experience. In contrast, VR “lean forward” content such as the popular “Beat Saber” [6], will tolerate less latency.

Overall, VR and AR will be much less tolerant to latency than Cloud gaming. Currently, the network RTT latency and jitter required to serve a particular AR/VR experience may be proprietary knowledge, dependent on the “motion to photon” latency as well as the technology stack of the Application Service Provider that manages rendering, encoding, decoding, etc. We expect that this will become public knowledge as the industry develops and users demand streaming options for delivery (similar to Cloud gaming today). It can be evaluated by ASPs under test conditions and may also be crowd-sourced through session evaluations of users.

In recent years, the Cable industry has embarked on a journey to explore the delivery of new media beyond video, such as gaming, AR/VR and light fields, over the network. For example, in the 2020 SCTE Cable Tec Expo Keynote demonstration, Charter Communications, together with partners, demonstrated a holographic transmission leveraging the “Power of 10G” [7][8][9][10].

Multiple System Operators (MSOs) have also set up the Immersive Digital Experiences Alliance (IDEA) [11] and are working with technology partners to advocate for standards and ecosystem development of immersive media. As MSOs become converged connectivity providers (wired and mobile), it is important that the Cable industry participate in the ecosystem and lead the definition of standards in Fog Computing and Edge Computing [12][13]. The methodology proposed in this paper crystallizes how MSOs can leverage their network to offer Edge Computing as a service to Application Service Providers (traditionally deployed in the Cloud).

3. The Communication Service Provider Network Hierarchy: Architecture and Considerations

When implemented, the methodology described here allows the CSP to provide latency-dependent compute to a subscriber, in concert with an ASP, for current and future interactive content experiences. We elaborate a method where system components interact seamlessly with each other for compute allocation.

We model compute grade and network latency as the most important factors that affect the end-user’s quality of experience (QoE). There may be other factors, such as specifications of an end-user device (ex., head mounted display/HMD); however, due to the nature of interactive immersive experiences (ex., VR

gaming or AR that factors in current context like detection of objects in view, SLAM [simultaneous localization and mapping], spatial anchoring of rendered 3D objects, etc., as opposed to watching a 360-video in VR), the intensity of compute and its latency to the end user become the primary determinants of whether the experience can be served at all. If it can be served, then optimizing delivery, i.e., encoding, streaming, buffering, etc., to match the end user device and pose is a secondary problem that must be solved tactically.




Figure 4- Sample Cable/Telecom Network

Figure 4 depicts the typical network hierarchy of a Cable/Telecom CSP. The access network, i.e., the “last mile” runs between a headend/CMTS (Cable Modem Termination System) to the customer premise for a Cable/broadband network and from a central office/base-station to the customer equipment (ex., mobile phone) for a telecom network. As more latency-sensitive compute is required for immersive media or other applications, it is expected that the central offices and headends may be retrofitted with more compute resources. There is public discourse on CORD (central office rearchitected as a data center) [14] and HERD (headend rearchitected as data center) [15][16] as the logical evolution of these CSP critical facilities.

The local headend/central office, however, is not the lowest level of the network hierarchy. Some of the new cable / fiber deep architectures such as Distributed Access Architecture (DAA) technology [17] [18] create opportunities for putting potential compute “closer” to the consumer along with Hybrid Fiber Coaxial cable (HFC) equipment. When combined with software “virtualization”, the downstream migration of DOCSIS functions may also free significant space and power in the headends, for retrofitting general purpose CPU/GPU compute into critical facilities. Similar cell site architecture proposals are also being considered in the 5G technology umbrella. The equipment on the customer premise (CPE, or customer premise equipment) comprises the lowest level in the network hierarchy. It includes cable modems, wireless routers, set top boxes, fixed wireless small cells, and user equipment such as TVs, holographic displays, AR/VR headsets, mobile devices, etc. As embedded compute becomes cheaper and more available, there is an opportunity to provide more capable devices that leverage built-in compute resources like GPUs.

Above the headends and central offices are the market and regional data centers, which may be considered at the same level or at different levels in the network hierarchy, depending on the specific topology of the network. Since these are already data centers, the CSP may augment these with new and more powerful compute like banks of GPUs.

Finally, the CSP network may have a national data center, or the traffic may be exchanged between a carrier network and the CSP network, to terminate in a Cloud location/CDN. This represents the highest level in the network hierarchy. These facilities are typically medium or large colocation data centers, and

providers of various services and applications such as immersive media may upgrade their equipment with adequate compute to serve new experiences.

It is important to note that today's IP media delivery connections are end-to-end, i.e., only have a source and a destination, meaning that, so far, there has been little need to specify a network hierarchy for media delivery. This also derives from the tolerance that current media applications have for network latency.

It is expected that a CSP will have a hierarchy of compute where, starting from the customer premise, each level of the network would typically have greater compute than the lower level, but less compute than the next level. Therefore, at the lowest level in the customer premise, the least amount of compute may be available. The headend/central office will have more compute than the customer premise, but less than the market or regional data centers, and so on. This is because each level of the network is serving more and more customers, requiring a larger serving radius. For example, while a CPE may serve only one home, a Remote-PHY (R-PHY) node [19] may serve 50 homes, a headend may serve 500 homes and a market data center may serve 10,000 homes. This is generally true for connectivity and packet routing today, however, as service providers augment their networks with compute for future applications, this will likely also be true of compute resources.

The other reason for this hierarchy is that it is easier to augment higher levels with compute as they are already data centers, and often have adequate space, power and cooling. Since the demand for compute at any point in time is statistical, adding servers with a bank of GPUs as a consolidated compute asset in a market or regional data center is logistically feasible. By comparison, a CSP may have to arrange for space, power and cooling to retrofit a central office or headend into a data center at the lower level. In a similar vein, a CSP may have to ship upgraded CPE, such as a router or a device to add compute in the customer premise. Therefore, it is relevant that the latency to the compute increases for each higher level in the network hierarchy.

4. Delivering Latency-Sensitive Compute for AR/VR Experiences

In this section, we describe how a CSP network may determine the intensity (or grade) of compute as well as the class of latency needed to serve an immersive AR/VR experience, as well as provide the specified compute to the customer from within its network.

4.1. A Priori Setup, CSP

- A. A CSP organizes its compute resources into units and compute grade/intensity. Each unit maps to a self-contained compute system capable of running a class of experiences, such as a Virtual Machine (VM) or a container. Each unit also has a compute grade associated with it, which refers to the quality of its resources, such as CPU cores, RAM, OS, GPU flops and vram. The compute grade may be a simple descriptor (say, a scale of one to ten) or a more complex descriptor that may be more descriptive on individual elements of the compute. Each unit is also associated with a hierarchy level in the network.
- B. The CSP sets up a global network compute orchestrator for management of all compute resources, as well as local compute orchestrators in each network compute element from data centers to CPE. A network compute orchestrator negotiates the compute resources on behalf of the subscriber from the network. The negotiation is based on the subscriber's service-level agreement (SLA) with the CSP – higher levels of service authorize higher compute grades for a subscriber, perhaps for longer time periods. Further, a network compute orchestrator may command a local compute orchestrator in a specific data center/headend/central office/CPE to reserve its resources.

- C. The network compute orchestrator has real-time visibility into the use of compute resources by receiving responses to queries, success/failure to commands or periodic messages from local compute orchestrators. It organizes this data using an efficient data structure that may be logically equivalent to an Available Compute Resource table wherein each entry represents Compute Resource Label - Compute Grade – Total Units – Available (Free) Units.
- D. The CSP maintains a data structure of network latencies required to reach its compute resources from each consumer's home (or a group of consumers' homes), at various levels in the network hierarchy. The CSP may deploy a latency measurement system between several probes at different points in the network. Some latencies may be directly measured and averaged. Other latencies may be deduced by addition or subtraction of aggregated, measured latencies based on knowledge of the topology of the network. Each measurement shall update a previously-averaged value using a weighted approach by allotting a higher weight to the most recently received value. A CPE may also periodically measure its latency to the central office/headend as part of this latency measurement system.

4.2. A Priori Setup, Application/Media Service Provider

- A. The ASP maintains a list of compute grade and the (worst case) latency required to serve a particular content/experience to a subscriber.

4.3. Connection Setup

- A. Subscriber requests an experience/content to be delivered to the premise on any of their devices. This request is routed to a network compute orchestrator.
- B. The network compute orchestrator queries the ASP for the compute grade and the latency required to serve the requested experience.
- C. The network compute orchestrator calculates the latencies from the subscriber to each of its compute units in an efficient data structure equivalent to a Network Hierarchy – Latency table.
- D. It then finds the lowest level of the network that has a latency equal to or better than the worst-case latency and a compute grade equal to or better than the compute grade for the requested experience using its Network Hierarchy – Latency (built from the point of view of the subscriber) and the Available Compute Resource (global compute resource information) tables, respectively. If the compute resource is not available, then the network compute orchestrator attempts to find the compute unit at the next higher level in the network hierarchy. This continues until either the compute unit is found, or the worst-case latency threshold is exceeded. If the compute unit is not found, the network compute orchestrator informs the subscriber that their requested experience cannot be served at this time.
- E. If a compute resource is available, the network compute orchestrator reserves it by issuing a command to the specific local compute orchestrator. It receives a token if the request is successful. It passes this token to the ASP and updates its Available Compute Resource table.
- F. If the Application Service Provider receives a token successfully, it proceeds to use the compute unit.

Table 1- Network Orchestrator "Available Compute Resource" Table

Compute Resource Label	Compute Grade	Total Units	Available (free) Units
.....
CPE – Router Subscriber S	3	1	1
CPE – Set Top Box Subscriber S	5	2	1
Headend / CMTS Domain D	9	10	4
Regional Data Center Region R	10	70	33
.....

Table 1 illustrates a network orchestrator Available Compute Resource table. In this example, we show compute grade (intensity) as a simple descriptor on a scale of one to ten. These descriptors can be mapped internally to specific attributes like GPU flops, vram, etc. The table shows entries for one subscriber, S. If the network makes compute available at the CPE level, then there would be entries for each subscriber. A tree data structure may be used for efficient storage and traversal, where the leaf nodes represent CPE at the subscriber level.

Table 2 - Application Service Provider Content Table

Content	Compute Grade / Intensity (descriptor 1-10)	RTT Latency, Customer to Compute (milliseconds)
.....
Content-X	8	30
Content-Y	6	10
.....

Table 2 illustrates an ASP content table. In this example, Content-X may be a “lean-back” experience where the user is immersed in a VR session but does not require fast-paced interaction with their environment. On the other hand, Content-Y may be a first-person shooter game in VR, where the targets are moving with high velocity. Thus, the RTT latency requirement for these experiences is significantly different.

Table 3- Network Orchestrator “Network Hierarchy – Latency” Table

Hierarchy	Compute Resource Label	RTT Latency, Customer to Level/Compute (milliseconds)
1	CPE – Router Subscriber S	3
2	CPE – Set Top Box Subscriber S	3
3	Headend / CMTS Domain D	9
4	Regional Data Center Region R	23
5	National “Cloud” Data Center NOC	55

Table 3 illustrates a “Network Hierarchy – Latency” table built by the network orchestrator for serving Subscriber S.

Consider that subscriber S wishes to experience Content-Y, which is both latency and compute sensitive, as observed from Table 2. From Table 3, it is evident that the content must be served from a headend/CMTS or a lower network hierarchy, due to latency constraints. Now, from Table 1, we observe that the CPE in Subscriber S's home does not have the compute intensity to drive this experience. In this case, the compute unit must be allocated at the headend/CMTS. In our example, four units of compute resources with grade nine are available, as seen in Table 1 (that is greater than the requirement of six for Content-Y, Table 2). Therefore, one of these units may be allocated to delivering Content-Y in Subscriber S's premise.

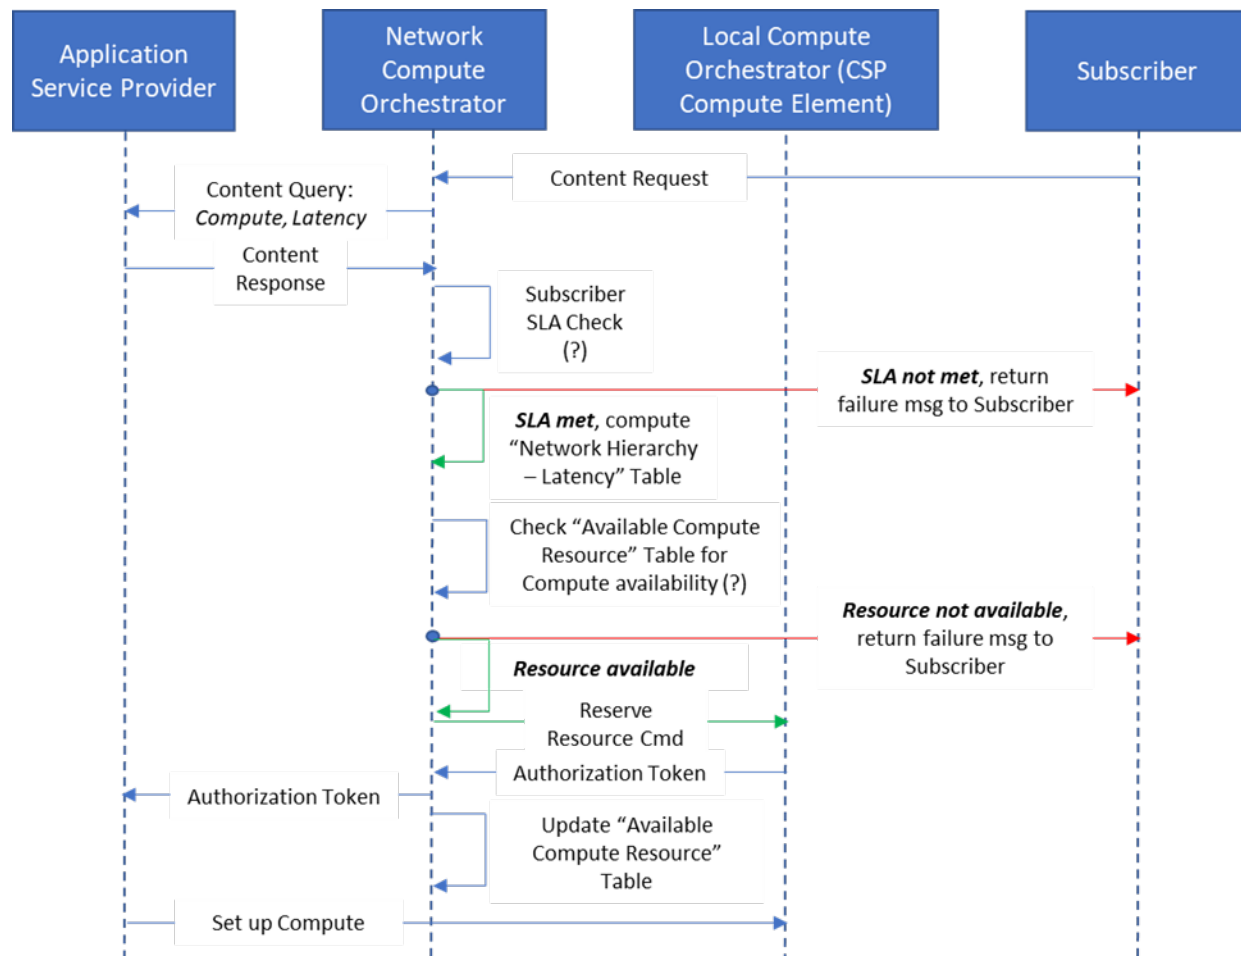


Figure 5 - Sample Session Setup

Figure 5 shows a sample session setup where a subscriber requests an AR/VR experience and the CSP allocates the compute resource for that content from within its network. The subscriber may request the experience either from the ASP that then contacts the CSP for low-latency compute, or directly from the CSP that may forward the request to the ASP. This may depend on the business relationship between the subscriber, the ASP and the CSP. The initial query to the ASP is used to determine the latency and compute intensity requirements. The CSP, after receiving these requirements and checking the subscriber SLA, proceeds to find the compute unit within its network. From the subscriber-centric "Network Hierarchy - Latency" table, the CSP determines the (one or more) levels in the network hierarchy from which the request may be served based on latency. Thereafter, a determination of the critical facility from which the request shall be served is made by the CSP's Network Orchestrator based on latency and

compute intensity. The Network Orchestrator updates its “Available Compute Resource” table accordingly. Once the compute is allocated, the ASP loads the application. It is worth noting that while the subscriber initiates the request, the delivery of content occurs when the ASP and the CSP work together.

5. Discussion of Compute Grade and Latency

In this paper, we used a simple descriptor on a scale of one to ten to denote Compute Grade or Compute Intensity. In reality, gaming/VR storefronts use detailed Minimum and Recommended Compute configurations to specify the grade. Figure 6 illustrates three examples from Steam [20], one of the largest distribution storefronts for gaming and VR:

Grand Theft Auto

SYSTEM REQUIREMENTS	
MINIMUM:	RECOMMENDED:
Requires a 64-bit processor and operating system	Requires a 64-bit processor and operating system
OS: Windows 10 64 Bit, Windows 8.1 64 Bit, Windows 8 64 Bit, Windows 7 64 Bit Service Pack 1	OS: Windows 10 64 Bit, Windows 8.1 64 Bit, Windows 8 64 Bit, Windows 7 64 Bit Service Pack 1
Processor: Intel Core 2 Quad CPU Q6600 @ 2.40GHz (4 CPUs) / AMD Phenom 9850 Quad-Core Processor (4 CPUs) @ 2.5GHz	Processor: Intel Core i5 3470 @ 3.2GHz (4 CPUs) / AMD X8 FX-8350 @ 4GHz (8 CPUs)
Memory: 4 GB RAM	Memory: 8 GB RAM
Graphics: NVIDIA 9800 GT 1GB / AMD HD 4870 1GB (DX 10, 10.1, 11)	Graphics: NVIDIA GTX 660 2GB / AMD HD 7870 2GB
Storage: 72 GB available space	Storage: 72 GB available space
Sound Card: 100% DirectX 10 compatible	Sound Card: 100% DirectX 10 compatible

Total War: WARHAMMER III

SYSTEM REQUIREMENTS	
Windows	macOS SteamOS + Linux
MINIMUM:	RECOMMENDED:
OS: Windows 7 64-bit	OS: Windows 10 64-bit
Processor: Intel i3/Ryzen 3 series	Processor: Intel i5/Ryzen 5 series
Memory: 6 GB RAM	Memory: 8 GB RAM
Graphics: Nvidia GTX 900/AMD RX 400 series Intel Iris Xe Graphics	Graphics: Nvidia GeForce GTX 1660 Ti/AMD RX 5600-XT
DirectX: Version 11	DirectX: Version 11
Storage: 120 GB available space	Storage: 120 GB available space
Additional Notes: 8GB Memory if using integrated GPU.	Additional Notes: TBA

Car Mechanic Simulator VR

SYSTEM REQUIREMENTS
MINIMUM:
OS: Windows 10
Processor: Intel Core i7 / AMD Ryzen 7
Memory: 16 GB RAM
Graphics: NVIDIA GeForce GTX 1660 Ti / AMD RX 5700
DirectX: Version 11
Storage: 13 GB available space

Figure 6 - Example System Specifications for Delivering Experiences

ASPs that provide Cloud compute, on the other hand, may use a single system configuration or a tiered system where a higher monthly rate entitles the subscriber to better system (compute) specifications.

Currently, content stores and ASPs may be separate or “all-in-one”. In this example, Steam is the content store that offers downloadable games, while an ASP would be a Cloud gaming/VR provider that runs game executables in its own data centers or in the public Cloud. In this paper, we treat ASPs as an “all-in-one”, wherein they have agreements with content publishers and storefronts to present content and, in concert with the CSP, compute for delivering the experience.

The immersive community has been exploring “split rendering”, wherein the compute required to serve immersive experiences is shared between a client device and a Cloud compute unit. Qualcomm’s “Boundless XR” [21] concept provides details on their implementation of split rendering.

The “split rendering” paradigm, an active area of R&D, divides compute between a client and the Cloud such that highly latency-sensitive rendering is delivered from the client (if the Compute Grade is available) while less latency-sensitive rendering is delivered from the Cloud Compute unit, and these renders are then composited to deliver a seamless experience to the user. Our methodology is easily modified to specify compute for split rendering. To enable this, the client device must send its available compute specifications, and the ASP, once aware of the locally-available compute, may provide a modified Compute Grade entry from its Content Table.

Finally, it is important to note that many latency measurement systems measure latency and jitter by sending a train of packets to a network device/Point of Presence (PoP), and monitoring RTT for each individual packet as well as the inter-packet delay when they are returned to sender. If a jitter value is available, the service provider may also make that a part of the Network Hierarchy – Latency table. If the jitter value from the subscriber to an element in the table exceeds a threshold, the network compute orchestrator shall reject that element as a potential render/compute candidate even if the latency is below the threshold required by the content/experience.

6. Conclusion

In this paper, we explored how a CSP such as a Cable and broadband service provider may work in conjunction with an ASP to commission on-demand compute for delivering immersive experiences to customers. We explained how it will be possible to serve AR/VR content from within the CSP network even if the “motion to photon” latency is published to be as low as 20ms. This has been validated through empirical observation.

Abbreviations

AR	Augmented reality
VR	Virtual reality
HMD	Head-mounted Display
ms	Milliseconds
MSA	Metropolitan Statistical Area
CSP	Communications Service Provider
ASP	Application Service Provider
RTT	Round Trip Time
CMTS	Cable Modem Termination System
CORD	Central Office Rearchitected as a Data center
HERD	HeadEnd Rearchitected as a Data center
CPE	Customer Premise Equipment
DAA	Distributed Access Architecture
R-PHY	Remote-PHY
CDN	Content Delivery Network
SLA	Service-Level Agreement
MSO	Multiple System Operator
SCTE	Society of Cable Telecommunications Engineers
PoP	Point of Prescence

Bibliography & References

- [1] What is Motion to Photon Latency, <http://www.chioka.in/what-is-motion-to-photon-latency/>
- [2] The Technology Behind a Low Latency Cloud Gaming Service, Parsec blog, <https://parsec.app/blog/description-of-parsec-technology-b2738dcc3842>
- [3] The Road to Success: How we are defeating Latency, Shadow Blog, <https://shadow.tech/en-gb/blog/news/roadmap-cloud-gaming-without-latency>
- [4] Game on! How broadband providers can monetize ultra-low latency services for gamers, Nokia Blog by Gino Dion, <https://www.nokia.com/blog/game-on-how-broadband-providers-can-monetize-ultra-low-latency-services-for-gamers/>
- [5] theBlu on Steam, <https://store.steampowered.com/app/451520/theBlu/>
- [6] Beat Saber – VR Rhythm Game, <https://beatsaber.com/>
- [7] Power of 10G, SCTE Cable Tec Expo 2020 Keynote Demonstration, <https://www.youtube.com/watch?v=I79WMpLrrGU>
- [8] Charter, partners stream 10G holographic demo at virtual Cable-Tec Expo, Broadband Technology Report, <https://www.broadbandtechreport.com/video/article/14185182/charter-partners-stream-10g-holographic-demo-at-virtual-cabletec-expo>

- [9] The Cable Network, Immersive Experiences and Lightfields, Ip, A. and Lal, D., Broadband Library, <https://broadbandlibrary.com/the-cable-network-immersive-experiences-and-lightfields/>
- [10] How we streamed a Light-field over a 10G Network, IDEA Workshop at 2020 SMPTE Annual Conference, <https://www.youtube.com/watch?v=cAg0A9gld5c&t=3s>
- [11] Immersive Digital Experiences Alliance website, <https://www.immersivealliance.org/>
- [12] Edge computing: current trends, research challenges and future directions, Carvalho, G., Cabral, B., Pereira, V. et al, Computing 103, 993–1023 (2021), <https://doi.org/10.1007/s00607-020-00896-5>
- [13] Fog Computing as an Enabler for Immersive Media: Service Scenarios and Research Opportunities, You, D. et al., IEEE Access, Volume 7, 2019, <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8716694>
- [14] Cabling considerations for CORD networks, <https://www.cablinginstall.com/data-center/article/14068510/cabling-considerations-for-cord-networks>
- [15] Have you Heard About HERD?, Chris Bastian, https://www.cablefax.com/cablefax_viewpoint/have-you-heard-about-herd
- [16] HERD for the Gigabit Era, Broadband Technology Report, <https://www.broadbandtechreport.com/docsis/article/16449156/herd-for-the-gigabit-era>
- [17] Evolving to Distributed Access Architectures, Chris Bastian, <https://www.cablefax.com/technology/evolving-to-distributed-access-architectures>
- [18] Distributed Access Architecture Is Now Widely Distributed – And Delivering On Its Promise, Howald, R., Eichenlaub, F., Peck, T., Bonen, A., SCTE Fall Technical Forum 2021, <https://www.nctatechnicalpapers.com/Paper/2021/2021-distributed-access-architecture-is-now-widely-distributed>
- [19] Remote PHY Distributed Access Architecture, Steven Harris, Broadband Library, <https://broadbandlibrary.com/remote-phy-distributed-access-architecture/>
- [20] Steam website, <https://store.steampowered.com/>
- [21] Boundless photorealistic mobile XR over 5G, https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/more_immersive_xr_through_split-rendering_-_web.pdf
- [22] QoE-aware dynamic service composition for immersive media-oriented services, Park, J., Lee, H., Yi, D., Kim, J. (2022), https://www.researchgate.net/figure/QoE-aware-dynamic-composition-framework-for-immersive-media-oriented-services_fig2_228930025
- [23] Implementation of a Media Aware Network Element for Content Aware Networks, Niculescu, D., Stanciu, M., Vochin, M., Borcoci, E., Zotos, N. (2011). CTRQ 2011 - 4th International Conference on Communication Theory, Reliability, and Quality of Service, https://www.researchgate.net/publication/228948173_Implementation_of_a_Media_Aware_Network_Element_for_Content_Aware_Networks

Converged Service Orchestration – Dynamic Cable Speed Boost

A Technical Paper prepared for SCTE by

Rahil Gandotra

Ph.D., Sr. Software Architect
CableLabs
858 Coal Creek Circle, Louisville, CO - 80027
(303) 661-3439
r.gandotra@cablelabs.com

Shafi Khan

Lead Software Engineer
CableLabs
858 Coal Creek Circle, Louisville, CO - 80027
(303) 661-3318
s.khan@cablelabs.com

Yunjung Yi

Ph.D., Principal Architect & Director of Wireless Standardization
CableLabs
858 Coal Creek Circle, Louisville, CO - 80027
(303) 661-3849
y.yi@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Converged Service Management Layer	5
3. Dynamic Cable Speed Boost	8
4. Conclusion.....	12
Abbreviations	13
Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 - Service orchestration as compared to workload or network orchestration.....	4
Figure 2 - NFV MANO framework [2].....	5
Figure 3 - Enhanced MANO framework using CSML for VNF and PNF management.....	6
Figure 4 - High-level CSML architecture	7
Figure 5 - Multi-cloud service orchestration using CSML	8
Figure 6 - Network topology of dynamic cable speed boost	9
Figure 7 - End-to-end interactions of CSML with different NFs	11
Figure 8 - Speed test results and 4K video connection speeds before and after speed boost	12

1. Introduction

Traditionally, networks serving different purposes operate in siloes with little-to-negligible overlap in their infrastructures or management processes. For example, DOCSIS networks, optical networks, and mobile networks, each employs specialized hardware functions managed by proprietary element management systems. Operators who offer both wireline and wireless services have dedicated teams to provision and manage each of their different services. This disjointed model of network management is challenged by operational economics. Consequently, designing, deploying, and operating end-to-end services are long and manual processes with long lead times (weeks to months) for effective service delivery.

On the contrary, the networks of tomorrow are envisioned to operate multiple different physical and cloud-native functions over a single flexible, programmable convergence platform whose hardware, software, and data storage resources are shared across multiple access technologies. Convergence is the process of unifying heterogeneous technologies to deliver seamless and ubiquitous connectivity. For operators, convergence creates opportunities for delivering novel, differentiated services, as well as allows for enhanced operational agility.

The umbrella term convergence consists of multiple building blocks such as access convergence, transport convergence, core convergence, platform convergence, operations convergence, and security convergence. Operations convergence entails creating a common operations framework for deploying, configuring, and managing network functions constituting a service. To make business models associated with converged networks and services increasingly attractive, in both wired and wireless domains, strategic research is required to devise the best integration architectures and appropriate accompanying integrated operations and management solutions.

When it comes to solving these challenges, technologies like software-defined networking (SDN) and network functions virtualization (NFV) have already addressed certain pieces of the puzzle. SDN enables decoupling the control plane (signaling/routing traffic) from the user plane (data/application traffic) to provide a global view of the network for efficient centralized control and allows for simpler forwarding devices. NFV allows for the separation of network functions from the underlying hardware and enables running them as virtual machines (VMs) or containers on commercial off-the-shelf (COTS) hardware. Additionally, the cloud computing paradigm provides efficient means to offer flexible resources and economies of scale. A converged service operator would need the capability to integrate all these disparate technologies into a single comprehensive framework to model end-to-end services and to abstract and automate the control and management of physical and virtual resources.

The Converged Service Management Layer (CSML) project — a key piece of the operations convergence puzzle — began in response to the rising need for a common automation platform for different network lifecycle processes. The goal for the project is to leverage and integrate existing next-generation network technologies to highlight the importance and novelty of converged service operations. We envision CSML acting like a master element management system (EMS) which communicates southbound to various domain-specific infrastructure layers, EMS, network functions (physical and virtual) in order to develop converged services. Typically, domain-specific management systems, such as for hybrid fiber-coaxial (HFC) or mobile networks, do not have the visibility or control over other domains for FCAPS tasks, and a central management entity such as CSML can enable that inter-domain communication to achieve multi-access convergence. By supporting management mechanisms of various domains, CSML can act like a single point of control that allows operators to employ virtual functions alongside physical functions and utilize their different network domains more coherently.

The term service orchestration is oftentimes used to describe distinct concepts in the industry. In addition, terms like workload orchestration or network orchestration are also employed in similar but different contexts. Workload orchestration typically refers to installing an application and its associated components over a single target virtual infrastructure, such as using OpenStack for virtual network functions (VNFs) or Kubernetes for cloud-native network functions (CNFs). It is predominantly limited to the management of a single cluster running its own control and user planes. Network orchestration refers to managing network configurations and policies for physical or virtual functions through a single pane of glass, typically employing a SDN controller for centralizing the control planes to manage multiple user planes. Service orchestration operates at a higher layer than workload and network orchestration and leverages them to deploy and manage services comprising of multiple applications to be installed on different target infrastructures, while integrating all control functionalities into a single framework. It uses workload orchestration to instantiate network functions as VNFs or CNFs over various target infrastructures and uses network orchestration to configure physical and virtual networks spanning multiple domains. Figure 1 highlights the difference between different layers of orchestration. Each domain, such as cable access and core, or mobile RAN and core, generally employ domain-specific orchestrators and controllers. For a converged service encompassing multiple domains and utilizing separate workload and network orchestrators, a service orchestrator such as CSML can help integrate all the different components to provide true operations convergence.

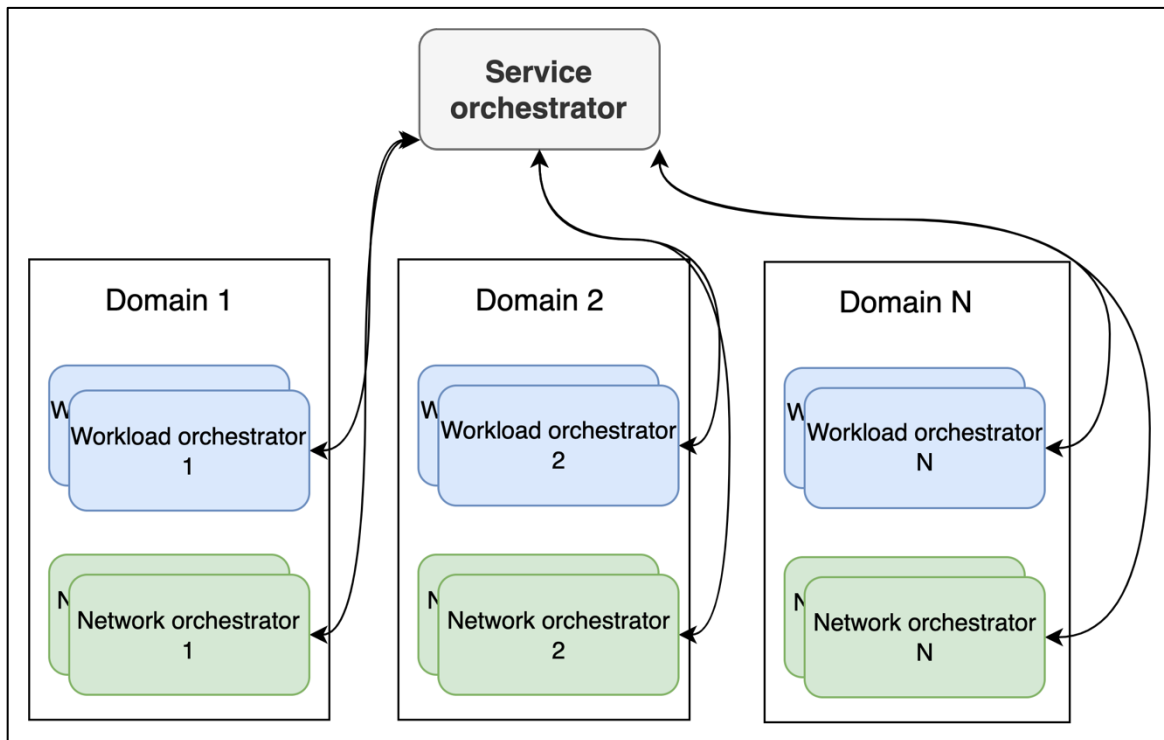


Figure 1 - Service orchestration as compared to workload or network orchestration

During the runtime of the project, several different use cases were developed and implemented targeting different network domains and lifecycles. The first use case targeted virtualized functions in the mobile and Wi-Fi domains, while the second use case focused on incorporating HFC-domain physical functions into the framework to demonstrate hybrid orchestration. This paper presents details on the second use case developed on HFC networks – dynamic speed boost - to demonstrate the value of operations convergence.

Typically, in cable networks, the process of requesting for and implementing a change in subscriber bandwidth or QoE is manual and requires a reboot of the cable modem (CM) for the change to take effect. The PacketCable Multimedia (PCMM) technology provides a mechanism to create dynamic quality of service QoS policies by leveraging functionalities defined in the DOCSIS and PacketCable DQOS specifications [1]. We onboard PCMM into CSML to enable dynamic, closed-loop, application-specific QoS control over the legacy CMTS to support existing and future QoS-enhanced services.

This paper presents details on the overall PoC architecture involving CSML, physical network functions (CMTS and CM) and a virtual PCMM domain controller, and discusses the end-to-end technicalities of the PoC. The rest of the paper is organized as: section II presents details on CSML, its different building blocks, and describes the first use case briefly, section III provides details of the dynamic cable speed boost PoC developed, and section IV presents the conclusion and directions for future work.

2. Converged Service Management Layer

The European Telecommunications Standards Institute (ETSI) Industry Specification Group for NFV (ETSI ISG NFV) developed a framework for NFV management and orchestration (MANO) of all resources in a virtualized data center [2]. Figure 2 shows the NFV MANO framework and the different components managing different network layers. The virtualized infrastructure manager (VIM) communicates with the NFV infrastructure (NFVI) layer to manage the underlying virtualized resources – virtual compute, virtual networking, and virtual storage (such as OpenStack or Kubernetes). The VNFM (VNF manager) interacts with both element managers (EM) and VNFs directly to manage the fault, configuration, accounting, performance, security (FCAPS) of individual VNFs. The NFV orchestrator (NFVO) interacts with OSS/BSS and the underlying management entities (VNFM and VIM) to manage the lifecycle of complex services comprising of multiple VNFs.

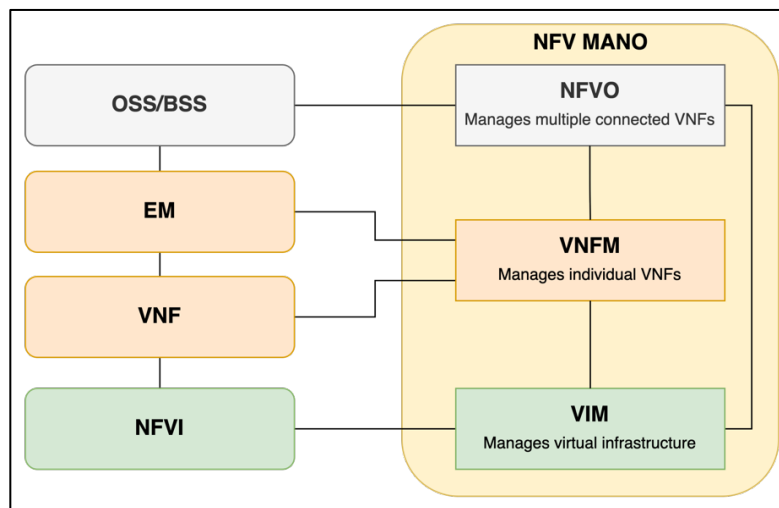


Figure 2 - NFV MANO framework [2]

In the context of different orchestration layers, the NFVO is a subset of the service orchestrator for NFV infrastructure, the VNFM maps to the workload orchestrator managing the lifecycle of individual functions, while the VIM compares with the network orchestrator providing underlying network connectivity. Since DOCSIS, optical, and mobile networks all include some functions implemented in specialized hardware, CSML extends the NFV MANO framework to include the management of physical network functions (PNFs) as well. Figure 3 illustrates the enhanced MANO framework employing CSML to manage PNFs as well. CSML acts as the NFVO and VNFM components from the NFV MANO framework. It includes the

capability to manage VNFs using either its in-built generic-VNFM (gVNFM) or by interacting with specific-VNFMs (sVNFM), including an interface to the VIM layer for virtual resource management over the NFVI. Similarly, CSML can manage PNFs either through specific EMs or by directly interacting with PNFs using its PNF manager (PNFM) over well-known management protocols.

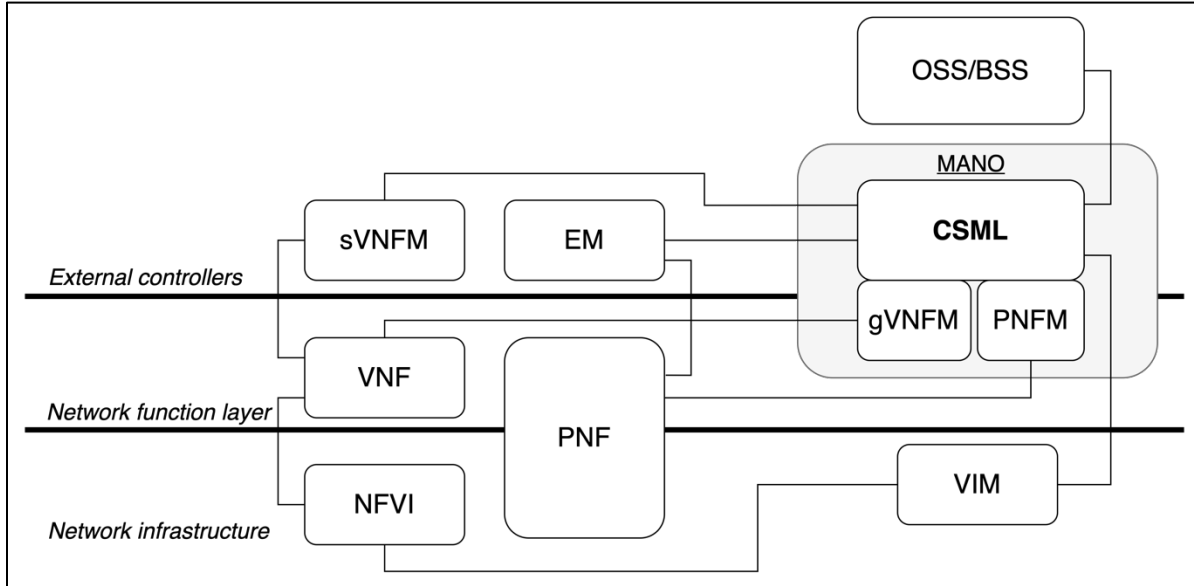


Figure 3 - Enhanced MANO framework using CSML for VNF and PNF management

On a high-level, as illustrated in Figure 4, CSML enables three types of lifecycle management (LCM) activities as a part of its two primary frameworks, namely the design-time framework and the run-time framework:

1. Service design: Involves composing services comprising of multiple xNFs, along with their Day-0 and Day-N configuration parameters and policy rules to enable elastic management of the service.
2. Service deployment: Involves instantiating the modeled services on to the target infrastructures (both physical and virtual) and supervising scale-in/out as needed.
3. Service assurance: Involves monitoring the deployed services and taking closed-loop actions using analytic tools to make the framework self-healing and self-optimizing.

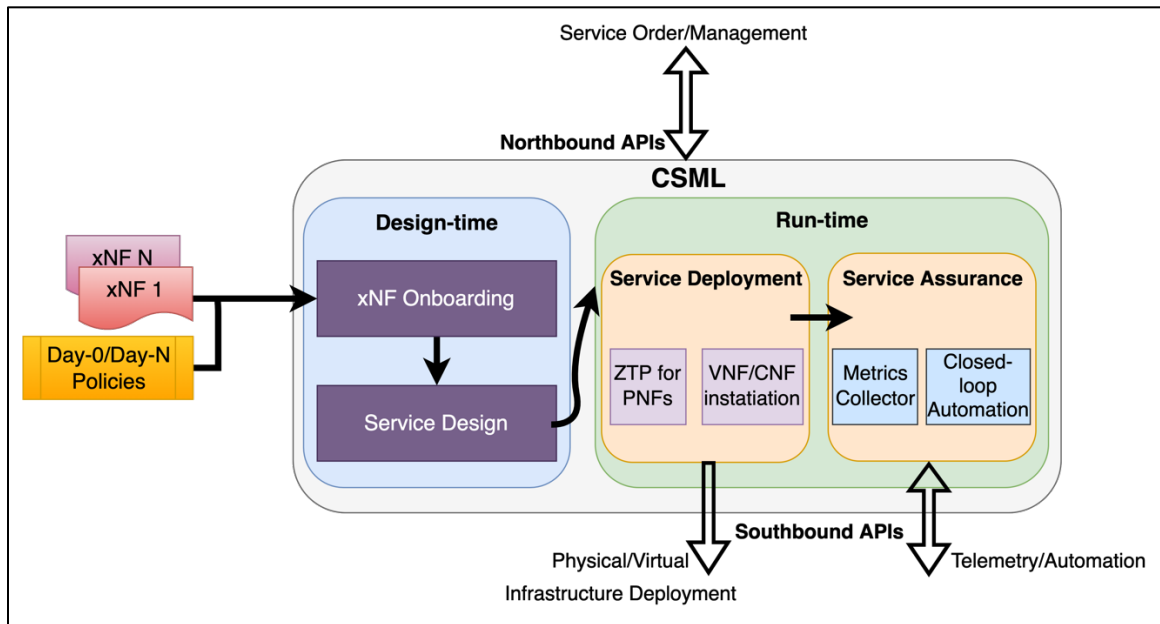


Figure 4 - High-level CSML architecture

The first use-case developed as part of the project activity was to orchestrate and assure a wireless steering application [3]. This application provides an over-the-top solution to steer user traffic over Wi-Fi or mobile network based on current network performance or operator requirements. It is a client-server-based application with the client running on user's mobile device and the server running in the cloud. The server service comprised of eight containerized network functions which needed to be running on different cloud locations simulating an operator's production network setup. This use case highlighted the value of using a service orchestrator such as CSML that can interact with multiple underlying workload orchestrators to deploy a distributed service over different cloud locations comprising of both public and private clouds.

Figure 5 illustrates the network topology developed for this use case. The different server components of the application are deployed on different cloud regions to showcase the capability of CSML to deploy network functions based on their network and compute requirements. For example, the control gateway is an edge component and needed to be geographically closer to the user, therefore it was deployed on Azure Central region, while the data gateway is the data path component through which all user traffic traverses, and therefore it is deployed on a private cloud. The VIM layer consisted of Kubernetes-based providers – Azure Kubernetes Service (AKS) for the public cloud locations and RedHat's OpenShift for the private cloud location. CSML itself was implemented to run in AKS, acting as the service orchestrator and was able to demonstrate multi-cloud service orchestration involving instantiating components over different locations, interconnecting them over virtual networks, and ensuring the end-to-end service operated as expected.

CSML was also responsible for assuring and healing the service by continuous monitoring and taking a specific closed-loop action when an issue was detected. The scenarios tested included - (i) Application-level healing – The control gateway component was responsible for leasing out IP addresses from a pre-defined IP pool to UEs connecting to the steering application, and as the IP pool would start to get exhausted, no new UEs would be able to connect to the application. CSML was able to identify this issue by continuously monitoring the control gateway's runtime via one of its REST endpoints and as 90% of the pool became exhausted, a new IP pool was allocated to by invoking another REST endpoint and modifying its Day-N configuration to ensure service continuity; (ii) Infrastructure-level healing – Since the control

gateway component is running on a shared cluster along with other services, there could be instances when the underlying compute resources become deficient. CSML was able to identify this issue by monitoring cluster resources via Kubernetes APIs, and resolving it by migrating the control gateway to another cluster and interconnecting all other components to the new cluster to ensure no service disruption.

Furthermore, additional utilities were developed in order to enable federated collection of telemetry data of both the infrastructure-level metrics as well as application-level metrics in order to enable centralized decision-making. This use case proved the value of service orchestration for virtualized environments spanning multiple locations and provided a common, abstracted framework for deploying and managing novel services.

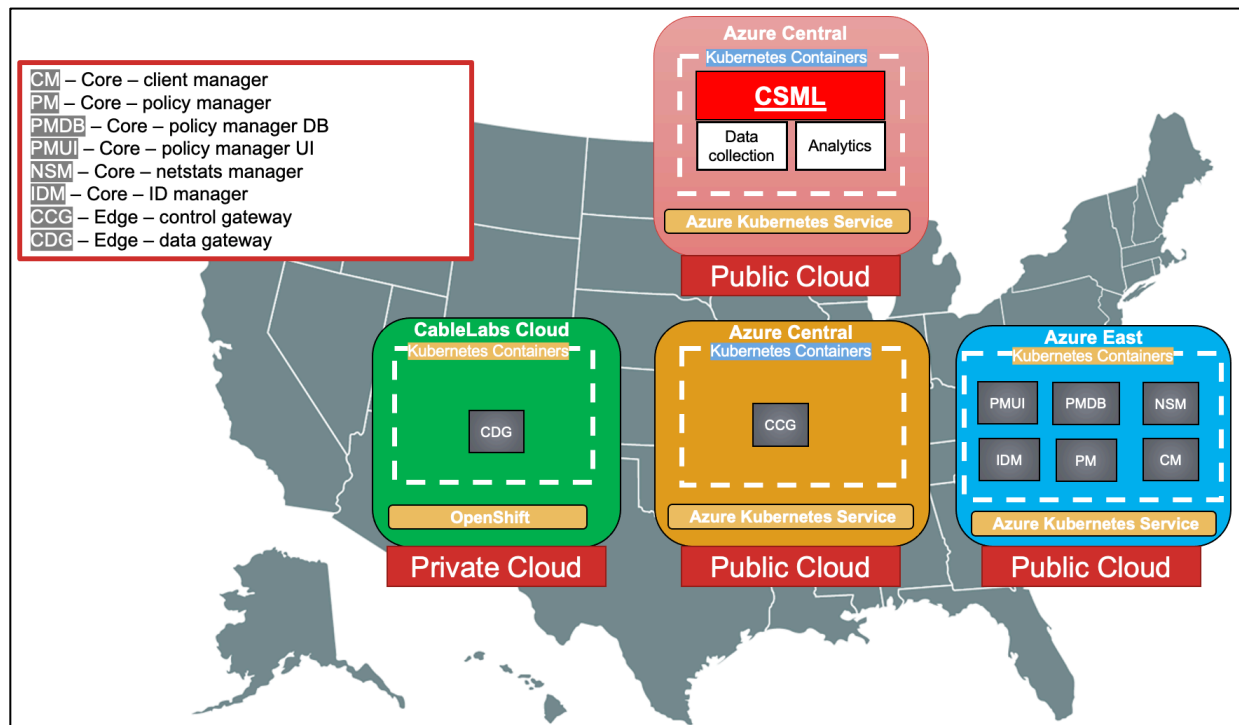


Figure 5 - Multi-cloud service orchestration using CSML

3. Dynamic Cable Speed Boost

The second use-case developed, to demonstrate hybrid physical-virtual orchestration, was to use CSML to enable dynamic speed boost in the cable network. DOCSIS provides two mechanisms to configure service flows (QoS policies) – (i) Static, which are established during the CM registration process, and (ii) Dynamic, which can be configured on an as-needed basis. Typically, operators employ the first mechanism which entails pushing a new configuration file to the CM manually and requires a reboot of the CM for it to take effect. This process is relatively lengthy and leads to disruption in traffic for the time it takes the CM to boot back up. The dynamic service flow mechanism requires an entity to interact with multiple controllers and network devices in order to translate the required QoS request into configurable DOCSIS service flows. CSML can act as the central orchestrator in this scenario to realize this use case by leveraging its capability to control both virtual and physical functions.

Figure 6 illustrates the high-level network topology of the use case implemented. The dynamic speed boost app is a web-based GUI developed using the Python-Flask framework, running in a VM in public cloud, which provides the user with the choice to boost or un-boost a specific application for a definite or indefinite time [4]. The application options included as part of this use case include YouTube (classified based on TCP port 443), Zoom audio (classified based on DSCP value 56), and Zoom video (classified based on DSCP value 40). Additional application options could be included based on requirements and differentiating parameters. CSML exposes northbound REST APIs which the dynamic speed boost app invokes to pass the user requested QoS parameters. In the southbound direction, interfaces were developed for CSML to connect to a PCMM controller and a CMTS in order to create, modify or delete service flows and verify if they have been pushed to the CM, respectively. OpenDaylight (ODL) was used as the PCMM controller in this use case as it is open-source and includes an implementation of the PCMM service as an additional module [5]. ODL was also running as a VM in public cloud. ODL exposes northbound REST APIs to provision a CMTS and service flows, and uses the Common Open Policy Service (COPS) protocol as transport to communicate with the CMTS. CMTS uses MAC-layer DOCSIS Dynamic Service Add/Change/Delete (DSA/DSC/DSD) messaging for service flow management for a specific CM. A Cisco cBR-8 was used as the CMTS and an Arris SURFboard model was used as the CM. A laptop was connected behind the CM for speed-testing purposes.

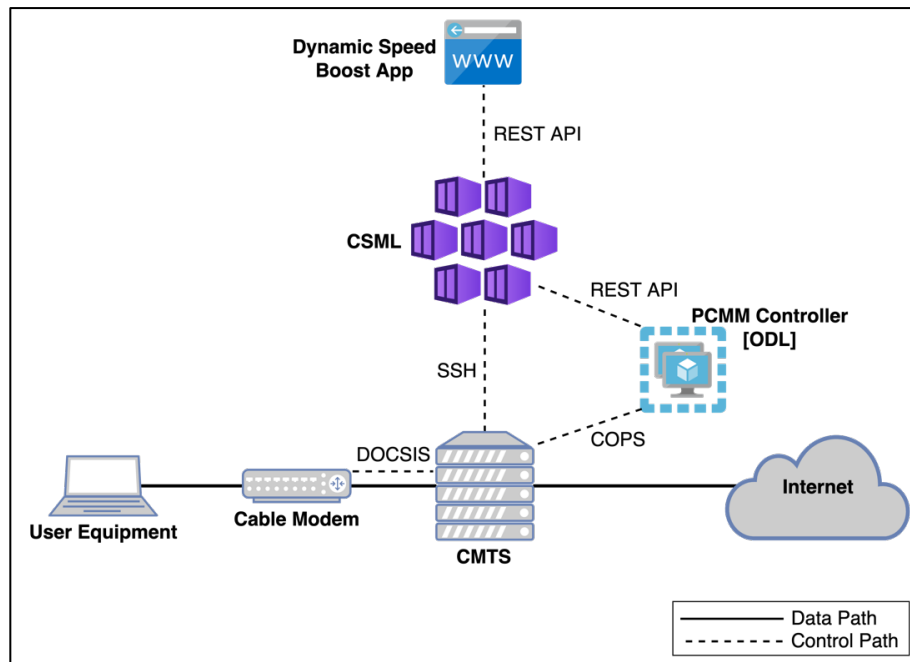


Figure 6 - Network topology of dynamic cable speed boost

The design-time work within CSML involved creating blueprint models for ODL and CMTS in order to enable their Day-N configuration. The network function (NF) model is essentially an archive consisting of TOSCA-based definitions represented in JSON, and specifications for run-time actions. The blueprint definitions include reusable data dictionaries, and component and node workflows to allow for intent-based configuration of the NFs. The run-time model defines the southbound interface and the associated actions that can be invoked to automate NF configuration. The southbound interface used to configure ODL was a REST-based environment that enables pipelining multiple API request-responses defined in a workflow, while the southbound interface used for the CMTS was a Python-based environment that allowed reusing existing scripts to SSH into the CMTS and change configuration. The Python-SSH script employed the

Netmiko library to set up connections with the CMTS, and send and receive commands related to service-flow configurations [6].

The run-time work within CSML involved developing northbound APIs and defining their associated payloads that could be used by an external application, the dynamic speed boost app in this scenario, in order to request for a boost for a particular application. CSML translated this request and invokes a defined workflow to transfer the requests to the correct run-time environments and take corresponding actions based on responses received from the NFs. This process enables an abstracted service framework which can be used by external northbound systems to interact with CSML, and in turn the underlying NFs, without knowing specific details about individual NFs.

Figure 7 depicts the multiple interactions of CSML with different NFs. The end-to-end flow works as:

1. User requests for a speed boost by logging into the dynamic speed boost app and providing details of the application requiring boost and an optional end time.
2. The request gets communicated to CSML over its northbound REST API including the CM IP of the requesting user.
3. CSML invokes a workflow defined for this purpose that includes the steps to fulfil this request.
4. CSML first checks if the COPS connection between the CMTS associated with the requesting CM and ODL is up. If it is not, CSML initiates the connection by sending an API request to ODL.
5. Next, CSML communicates the service flow request to ODL consisting of the CM IP, classifier to match specific application traffic, and either a pre-configured service class name on the CMTS or DOCSIS QoS parameters to define the speed boost requested.
6. ODL uses COPS protocol to communicate with CMTS the new service flow required.
7. CMTS uses DOCSIS to push the service flow to the specific CM.
8. Once ODL returns the confirmation to CSML that the flow has been pushed, CSML initiates a SSH connection directly to the CMTS and requests all service flows active for that CM IP.
9. CMTS returns the list of active service flows to CSML.
10. CSML verifies that the new service flow pushed matches the original request from the user.
11. After verification, CSML returns the boost confirmation message to the dynamic speed boost app over its northbound REST API.

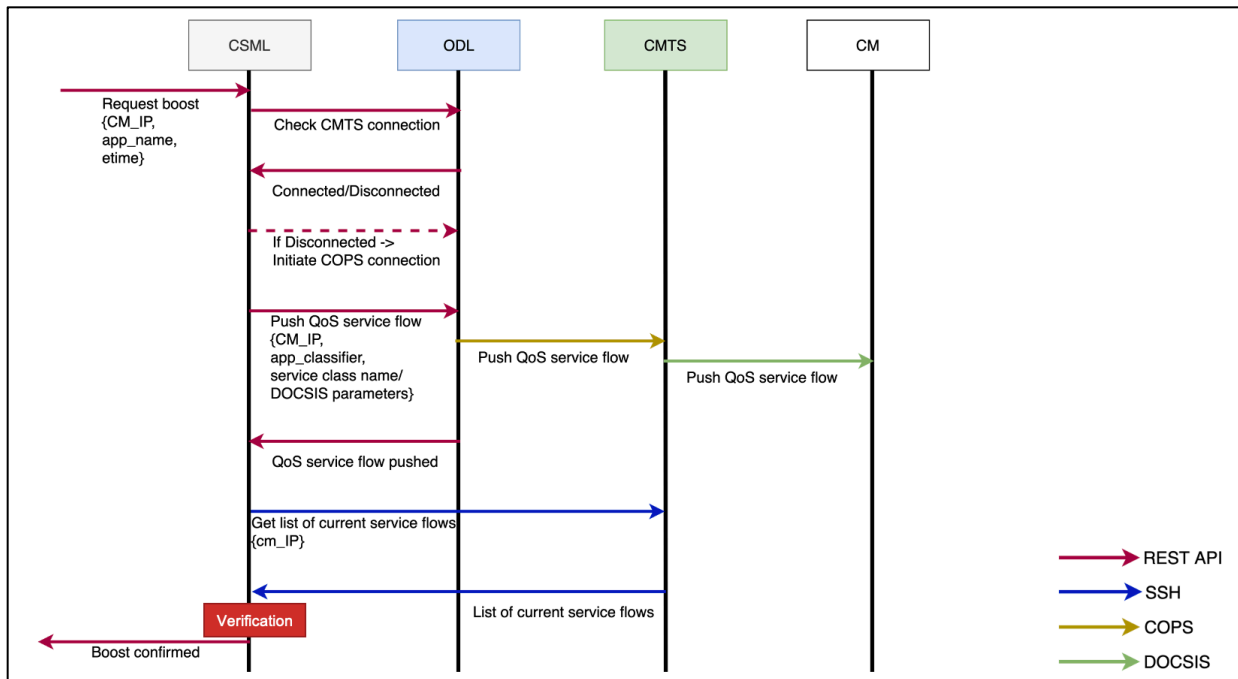


Figure 7 - End-to-end interactions of CSML with different NFs

The end-to-end flow described above illustrates the level of abstraction offered by using a converged service orchestrator. The OSS/BSS or external applications are only required to send a single request to CSML, and all the non-trivial interactions with different NFs, both physical and virtual, are implemented on the backend within CSML to fulfil the request. This enables integrations with multiple different NFs and controllers, and allows for the development of novel use cases with the orchestrator handling much of the operational complexities associated with them.

To test the overall functioning of this use case, the user laptop was used to test different application speeds before and after requesting for the boost. The user laptop was running a speed test application, YouTube application, and an iPerf client which was connected a public iPerf server on TCP port 5002. The default subscriber bandwidth was set to 10 Mbps for both upstream and downstream in the CM configuration file. Next, a 4K video was opened on the laptop and real-time video statistics were enabled to check connection speeds for the running stream. Initial speeds were capped at 10 Mbps, which resulted in excessive video buffering, and the iPerf speeds were also capped at 10 Mbps for both upload and download. To improve video performance, the user then requested for a speed boost of the YouTube application by logging into the dynamic speed boost app. Almost in real-time, CSML was able to fulfil this request and pushed a new downstream service flow (since YouTube video streaming is primarily downstream data-driven), and the video statistics indicated connection speeds going up to ~85-90 Mbps now, with the video no longer buffering anymore, as shown in Figure 8. At the same time, iPerf download speeds were still capped at 10 Mbps since it was running on a different port as that of the requesting application. This demonstrated the capability of CSML to achieve a dynamic cable speed boost for a specific application based on user requirements.

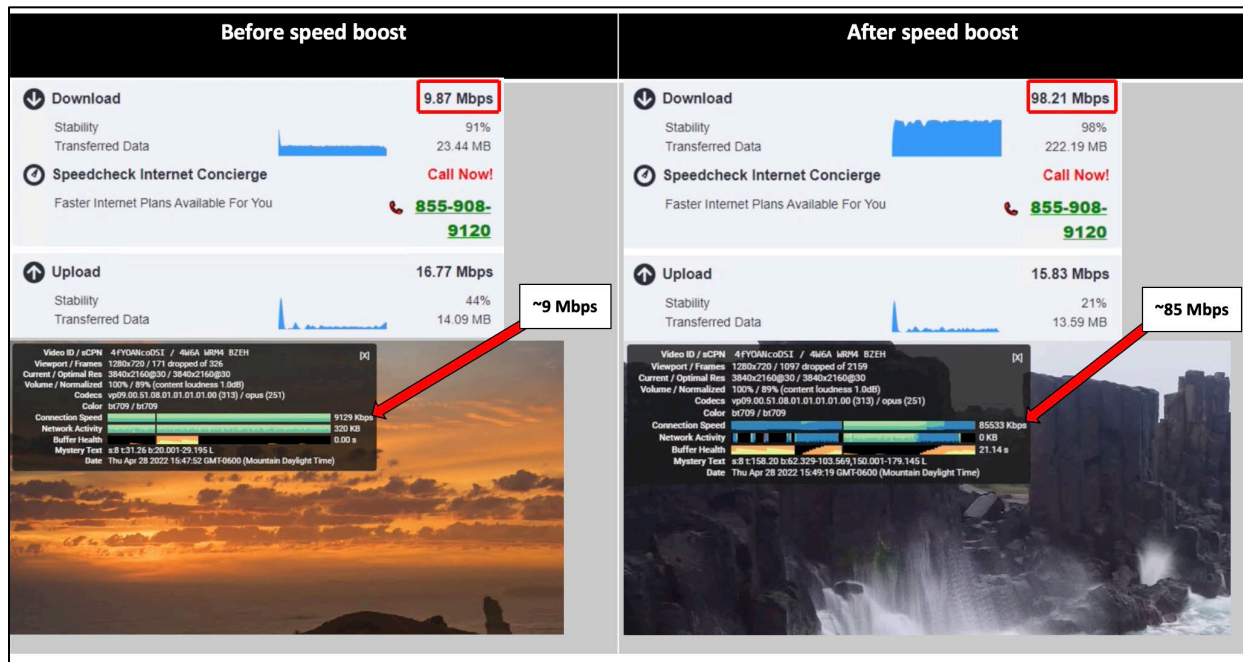


Figure 8 - Speed test results and 4K video connection speeds before and after speed boost

Although, individual scripts could also be used to achieve the same use case, however, orchestration provides additional value to the process. Firstly, developing scripts to communicate with different NFs is not trivial, CSML includes a run-time framework that removes much of the complexity to communicate with NFs and controllers, and simplifies creating complex workflows with multiple interactions. Secondly, once a NF or controller has been modeled within CSML, the same run-time environment can be extended to similar devices running the same management protocols. For example, while in this scenario a Python-SSH and a REST-API environments were developed to communicate with the CMTS and ODL controller respectively, the same could be leveraged to manage switches/routers over SSH and other controllers exposing REST APIs. Thirdly, once a NF or controller has been onboarded within the orchestrator, it can be reused for other use cases to create complex service chains employing different xNFs.

4. Conclusion

With the proliferation of virtualized network functions and the need to converge the management and operations of different network domains, a centralized service orchestrator is required for integrating different systems, thereby enabling novel use cases while reducing operational complexities. This paper introduced the CSML framework along with specific use cases developed to demonstrate the value of service orchestration and automation. The dynamic cable speed boost use case is described in detail which allows for changing the cable subscriber bandwidth on the fly. Both physical and virtual NFs and controllers were onboarded within CSML, and specific southbound run-time environments were leveraged in order to create a complex service chain including a PCMM controller and CMTS. Speed testing was performed to compare different application speeds before and after the user requests for a speed boost, and the results indicated that granular, real-time changes to subscribed bandwidth are achievable in DOCSIS networks using an intelligent, converged orchestrator framework such as CSML. The broader goals of the CSML project are to drive the adoption of network automation, virtualization, and operations convergence at scale. Also, as the transition to NFV is progressing, the project aims to demonstrate how physical network elements can be harmonized with virtual elements to preserve existing network investments.

While the use cases described in this paper demonstrated both physical and virtual orchestration, they were limited to domain-specific architectures. The next phase of the project is to develop a true converged use case that involves incorporating 5G core into the framework in order to provide additional value to operators that provide both cable and mobile services. For example, the cable speed boost use case could be extended by – (i) Identifying customer devices that are subscribed to both cable and mobile services from the same operator and, (ii) Boosting their Wi-Fi bandwidth to offer more value and help differentiating from other operators.

Abbreviations

AKS	Azure Kubernetes Service
CM	cable modem
CMTS	Cable Modem Termination System
CNF	cloud-native network function
COPS	Common Open Policy Service
COTS	commercial off-the-shelf
CSML	Converged Service Management Layer
DOCSIS	Data-Over-Cable Service Interface Specification
EMS	element management system
ETSI	European Telecommunications Standards Institute
FCAPS	fault, configuration, accounting, performance, security
HFC	hybrid fiber-coaxial
MANO	management and orchestration
NF	network function
NFV	network functions virtualization
NFVI	network functions virtualization infrastructure
NFVO	network functions virtualization orchestrator
ODL	OpenDaylight
PCMM	PacketCable Multimedia
PNF	physical network function
PNFM	physical network function manager
QoS	quality of service
SDN	software-defined networking
VIM	virtualized infrastructure manager
VM	virtual machine
VNF	virtual network function
VNFM	virtual network function manager

Bibliography & References

- [1] Cable Television Laboratories, Inc., “PacketCable™ Multimedia Specification”, PKT-SP-MM, November 2019. [\[link\]](#)

- [2] European Telecommunications Standards Institute, “Network Functions Virtualization (NFV); Management and Orchestration; Architectural Framework Specification”, ETSI GS NFV, June 2022. [\[link\]](#)
- [3] Cable Television Laboratories, Inc., “IWiNS—An Informed Approach to Mobile Traffic Steering”, January 2021. [\[link\]](#)
- [4] The Pallets Projects, “Flask - The Python micro framework for building web applications”. [\[link\]](#)
- [5] The Linux Foundation Projects, “OpenDaylight”. [\[link\]](#)
- [6] Kirk Byers, “Netmiko - Multi-vendor library to simplify CLI connections to network devices”. [\[link\]](#)

Data Channel Optimization

Managing Technology Borders

A Technical Paper prepared for SCTE by

Matthew Olfert

Principal Network Architect - Access Technology
Shaw Communications Inc.
2728 Hopewell Place NE Calgary AB T1Y 7J7
+1.403.538.5210
matthew.olfert@sjrb.ca

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Identifying Technology Boarders.....	4
2.1. HFC Borders	4
2.2. DOCSIS Version Borders.....	4
2.3. Cable Modem Borders	5
3. Key Data Elements	6
3.1. Cable Modem Distrabution.....	6
3.2. Cable Modem Consumption.....	6
3.3. Capacity By Service Group	7
3.4. Capacity By Capacity Type	7
3.5. Peak Traffic	7
3.6. Burst Capacity	8
3.7. Spectrum Efficiency	8
3.8. Cable Modem Upgrade Churn Rate	8
3.9. Current Product Offerings/Distribution	8
4. CMTS Capabilities.....	9
4.1. MAC Scheduler	9
4.1.1. Balanced Scheduler	9
4.1.2. Prioritized Scheduler.....	9
4.2. Dynamic Configuration Features.....	9
5. Optimization Drivers.....	10
5.1. Service Group Capacity	10
5.2. Customer Experience.....	10
5.3. Cost Reduction.....	10
5.3.1. License Reduction	10
5.3.2. Service Group Combining.....	10
6. Identifying Optimization Opportunities	10
6.1. Legacy Capacity Need Identification.....	11
6.2. New Capacity Need Identification	11
7. Optimization Cases Studies	12
7.1. Case Study 1 - Service Group Congestion	12
7.2. Case Study 2 – Spectrum Efficiency Congestion	13
7.3. Case Study 3 - Billboard Service with Low Burst Capacity.....	15
7.4. Case Study 4 - Low Utilization Service Group	17
7.5. Case Study 5 - High Utilization of Legacy Capacity.....	17
7.6. Case Study 6 – Spectrum Boundary.....	18
8. Looking Forward to DOCSIS 4.0.....	19
8.1. FDD Changes.....	19
8.2. FDX Changes.....	20
9. Conclusion.....	20
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 – Hourly Downstream Traffic Graph – Case Study 1	12
Figure 2 – Updated Hourly Downstream Traffic Graph – Case Study 1.....	13
Figure 3 – Hourly Downstream Traffic Graph – Case Study 2	14
Figure 4 – Spectrum Efficiency – Case Study 2	15
Figure 5 – Updated Hourly Downstream Traffic Graph – Case Study 2.....	15
Figure 6 – Hourly Downstream Traffic Graph – Case Study 3	16
Figure 7 – Updated Hourly Downstream Traffic Graph – Case Study 3.....	16
Figure 8 – Hourly Downstream Traffic Graph – Case Study 5	18
Figure 9 – Hourly Downstream Traffic Graph – Case Study 6	19
Figure 10 – DOCSIS 4.0 FDX N+1 Interference Groups	20

List of Tables

Title	Page Number
Table 1 – DOCSIS Version Borders	5
Table 2 – CM Bonding Capability	5
Table 3 – CM Duplex Filter Capability	5
Table 4 – Cable Modem Distribution Example	6
Table 5 – Cable Modem Consumption Example	6
Table 6 – Capacity By Service Group Example.....	7
Table 7 – Cable Modem Consumption Example	7
Table 8 – Peak Traffic by Capacity Type Example.....	7
Table 9 – Burst Capacity by Capacity Type Example.....	8
Table 10 – Example Data for Legacy Capacity.....	11
Table 11 – Example Data for New Capacity	11
Table 12 – Capacity – Case Study 1	12
Table 13 – Cable Modem Distribution – Case Study 1.....	13
Table 14 – Updated Capacity – Case Study 1.....	13
Table 15 – Capacity – Case Study 2	14
Table 16 – Cable Modem Distribution – Case Study 2.....	14
Table 17 – Updated Capacity – Case Study 2.....	14
Table 18 – Capacity – Case Study 3	15
Table 19 – Cable Modem Distribution – Case Study 3.....	16
Table 20 – Updated Capacity – Case Study 3.....	16
Table 21 – Capacity – Case Study 4	17
Table 22 – Capacity – Case Study 5	17
Table 23 – Cable Modem Distribution – Case Study 5.....	18
Table 24 – Capacity – Case Study 6	18
Table 25 – Cable Modem Distribution – Case Study 6.....	19

1. Introduction

As Data Over Cable Service Interface Specifications (DOCSIS[®]) access networks evolve through multiple versions technology Borders have been created. Evolving Hybrid Fiber Coax (HFC) plant upgrades, utilizing different Cable Modem Termination Systems (CMTS), and cable modems also create Borders as well. The evolution of technology within a coax access network has maintained backwards compatibility. However, we do not have forward compatibility through these evolutions. An optimized access networks must find the best balance between maintaining legacy services and building for future offerings.

This paper will define technology borders, and key data elements used in optimization. In addition, it will discuss the CMTS feature impacts, optimization drivers, and calculation methods to optimize an access network. Using case studies are also included which go through examples of utilizing this optimization method. Finally, it will discuss the impacts DOCSIS 4.0 will have in the future to this methodology.

2. Identifying Technology Borders

Technology borders can be summarized into two categories HFC and DOCSIS. These two categories are related but evolve independently from each other. Typically, the HFC plant is ready for technology evolution ahead of the needs of new DOCSIS version readiness.

2.1. HFC Borders

The HFC plant, utilized for data, is a bi-directional setup. There are two major data points for technology borders: total spectrum capability and diplex filter spectrum location. Total spectrum consists of the highest downstream frequency available, such as 750 MHz, 860 MHz, and 1 GHz. 1.2 GHz and 1.8 GHz. Builds of 1.8 GHz capable HFC plants are just beginning in preparation of DOCSIS 4.0 Frequency Division Duplexing (FDD).

The second data point of technology borders is the diplex filter which in North America is typically at 42, 85 or 204 MHz. With DOCSIS 4.0, operators will have options to either move the diplex filter higher, with FDD, or use software define spectrum division in Full-Duplex DOCSIS (FDX).

2.2. DOCSIS Version Borders

DOCSIS versions have modified spectrum range and offer multiple channel types. Every one of these changes creates a new border. The spectrum boundaries have changed each version on both the downstream and upstream, and the range of options have grown with each version. The range of support may differ between CMTS and cable modem devices. DOCSIS specifications state the following ranges as supported for each version:

Table 1 – DOCSIS Version Borders

DOCSIS Version	Downstream Plant (MHz)	Return Plant (MHz)	Full Duplex Plant (MHz)
2.0	54-864	5-42	-
3.0	54-1002 108-1002	5-42 5-85	-
3.1	54-1218 108-1218 258-1218	5-42 5-85 5-204	-
4.0 (FDD)	108-1794 258-1794 372-1794 492-1794 606-1794 834-1794	5-85 5-204 5-300 5-396 5-492 5-684	-
4.0 (FDX)	684-1218	5-85	108-684

2.3. Cable Modem Borders

Cable modems since DOCSIS 3.0 also have multiple variations of support within the same DOCSIS version. This variation usually revolves around bonding groups sizing and diplex filters. The diplex filters match the spectrum boundaries in the above table (Table 1). The below table (Table 2) shows common bonding group sizing for each DOCSIS version:

Table 2 – Cable Modem Bonding Capability

DOCSIS Version	Downstream Bonding Group Size	Upstream Bonding Group Size
2.0	1	1
3.0	4 8 16 24 32	4 8
3.1	32 (includes support of up to two OFDM channels)	12 (includes support of up to two OFDMA channels)

Table 3 shows the different diplex filter configurations that have typically been used in DOCSIS access networks:

Table 3 – Cable Modem Diplex Filter Capability

DOCSIS Version	Low-Split (42 MHz)	Mid-Split (85 MHz)	High-Split (204 MHz)
2.0	✓	x	x
3.0	✓	✓	x
3.1	✓	✓	✓

These tables do not include DOCSIS 4.0 cable modems yet, but we expect greater number of orthogonal frequency division multiplexing (OFDM) and orthogonal frequency division multiple access (OFDMA)

channel supported per cable modem, in addition to the many options of duplex settings available within the DOCSIS 4.0 specification.

3. Key Data Elements

As choices are made to modify the capability of the HFC plant or CMTS, data driven optimization decision making becomes necessary. The below data elements provide key information required to make these types of decisions. These data elements focus on each cable modem or DOCSIS channel. Note that service group data elements do not work for boundary optimizations.

3.1. Cable Modem Distribution

Looking at a service group's cable modem distribution can provide insight into traffic patterns seen at the service group level. If a certain border is operating below expectations, looking to see if the service group has enough cable modems that can access that capacity can provide a strong reason for this behavior. This can provide further insight into the potential for future capacity additions and can help predict the offloading of current capacity to the new capacity.

Table 4 – Cable Modem Distribution Example

Capacity Type	Count	% Of Distribution
DOCSIS 2.0	5	5%
DOCSIS 3.0	45	45%
DOCSIS 3.1	50	50%

3.2. Cable Modem Consumption

Consumption data from each cable modem within a service group can be very useful. This provides information on how much data was used during a given period. Adding the cable modem capabilities to this data allows for an enriched view to the service group's usage in that period, and by cable modem capability. Modifying the period can enable even more insights. For example, obtaining consumption data for certain peak hours and reviewing over a 30 to 90 day period can enable a better understanding of average consumption and overall usage for a particular service group. Using this type of consumption data ultimately enriches the cable modem distribution data set.

Table 5 – Cable Modem Consumption Example

Capacity Type	Count	% of Distribution	Total Consumption (Peak Hour) Gb
DOCSIS 2.0	5	5%	10
DOCSIS 3.0	45	45%	20
DOCSIS 3.1	50	50%	70

Table 5 illustrates the information gained from consumption data. In this example, the majority of traffic comes from DOCSIS 3.1 cable modems. Without this detailed view traffic would appear to be even between DOCSIS 2.0/3.0 versus DOCSIS 3.1 cable modems.

3.3. Capacity by Service Group

Understanding the capacity of a service group has long term been a key data element. Understanding the capacity you have made available on the downstream and upstream is very important. This is a very common data point historical for capacity planning an access network.

Table 6 – Capacity By Service Group Example

Capacity Items	Achievable Bit-rate (Mbps)
Downstream 32 SC-QAM (@256-QAM)	1216
192 MHz of OFDM (@1024-QAM)	1647
Total Downstream Capacity	2863
Upstream 4 SC-QAM (@64-QAM, 6.4 MHz)	104
42 MHz of OFDMA (@256-QAM)	285
Total Upstream Capacity	389

3.4. Capacity by Capacity Type

Understanding the capacity of the service group is very common data element. For this data element, it is important that the focus is on a *per* DOCSIS *channel* capacity level. Single carrier quadrature amplitude modulation (SC-QAM) channels are simple examples to understand; however, OFDM and OFDMA add complexity.

The spectrum location of DOCSIS channels is another important metric that can be used to enrich our data. The ability of the capacity type to be utilized becomes important when it is located within new boundary areas. For example, SC-QAM channel location within mid-split spectrum can only be utilized by mid-split capable cable modems.

Table 7 – Cable Modem Consumption Example

Capacity Type	Capacity (Mbps)
48 Downstream SC-QAM	1824
OFDM (192 MHz @1024-QAM)	1647
4 Upstream Low-split SC-QAM (@64-QAM)	104
2 Upstream Mid-split SC-QAM (@64-QAM)	52
OFDMA Mid-split (25.4 MHz @512-QAM)	195

3.5. Peak Traffic

Traffic usage at peak times provides insight into a congestion level of a service group. Breaking down this data into the channel level or channel type can provide insight on the performance of each capacity type.

Table 8 – Peak Traffic by Capacity Type Example

Capacity Type	Capacity (Mbps)	Peak Traffic (Mbps)
48 Downstream SC-QAM	1824	590
OFDM (192 MHz @1024-QAM)	1647	940
6 Upstream SC-QAM (@64-QAM)	156	46
OFDMA (25.4 MHz @512-QAM)	195	87

3.6. Burst Capacity

For service groups that do not have congestion, peak traffic is a powerful data element to understand burst capacity. Burst capacity becomes an output of capacity minus peak traffic. This data can also be organized to boundary focused elements as well. For example, DOCSIS 3.0 downstream burst capacity — which is based on only the SC-QAM capacity can be countered with DOCSIS 3.1 downstream burst capacity that includes both SC-QAM and OFDM capacity.

$$\text{Capacity} - \text{Peak Traffic} = \text{Burst Capacity}$$

The below table provides an example of utilizing this formula:

Table 9 – Burst Capacity by Capacity Type Example

Capacity Type	Capacity (Mbps)	Peak Traffic (Mbps)	Burst Capacity (Mbps)
48 Downstream SC-QAM	1824	590	1234
OFDM (192 MHz @1024-QAM)	1647	940	707
6 Upstream SC-QAM (@64-QAM)	156	46	110
OFDMA (25.4 MHz @512-QAM)	195	87	108

3.7. Spectrum Efficiency

This data point is a powerful indicator of operational issues. For SC-QAM capacity this indicator will show channel impairments and codeword error rates (CER), while for OFDM/OFDMA, this indicator will show low capacity profile use. This data is obtained by collecting information on each cable modems performance on each DOCSIS channel it utilizes and provides data on channel impairment, CER, and active profile/interval usage code (IUC) usage for each of these channels.

This data element is primarily an operational key performance indicator (KPI) but can be used to validate that an operational issue is not causing odd traffic patterns during optimization, avoiding a service group capacity change that impacts customers' services.

3.8. Cable Modem Upgrade Churn Rate

For long term planning (or forecasting), it is important to understand the churn rate of cable modems from older version (that have less capabilities) to newer versions that are more capable. Insights from churn rates can allow a more assertive decision making in transitions to new capacity methods. A good example of this is the rate of DOCSIS 3.0 modems upgrading to DOCSIS 3.1 modems. This can provide a better understanding of when OFDM capacity can be added (with a reduction to SC-QAM capacity) when spectrum is limited.

3.9. Current Product Offerings/Distribution

Understanding the products that are currently offered to customers can help determine the burst capacity needed per service group. In addition to the max speed, the limitations of the cable modems offered with each product is required. For example, a 1 Gbps product requires a OFDM capable cable modems, whereas a 100 Mbps does not require a OFDM capable cable modem, and DOCSIS 3.0 cable modems can still be used.

4. CMTS Capabilities

As each CMTS has differing capabilities, it is important to understand the possible options that can be utilized for managing technology borders, and if any capabilities can impact capacity and traffic patterns.

4.1. MAC Scheduler

The MAC scheduler is the controller of the DOCSIS access network and determines what channel a data packet is transmitted over. This functionality can assist or hinder the behavior the technology borders. There are two common schedulers: balanced scheduler and prioritized scheduler.

4.1.1. *Balanced Scheduler*

A balanced scheduler provides equally load-balanced traffic across all DOCSIS channels regardless of channel type and is usually based on percentage of utilization. This type of scheduler is effective where a large majority of cable modems have access to all DOCSIS channels. This is especially effective for DOCSIS 3.0 with low-split designed capacity.

As DOCSIS 3.1 was released and operators enabled several new technology borders with OFDM, OFDMA, and different upstream splits issues with type of scheduler started to become evident. Traffic offloading to new technology capacity was held back or legacy capacity was over utilized by more capable cable modems. This limitation forces an increased cable modem churn rate, or a different scheduler type is needed. However, if the new capacity is only dedicated to new product offerings, traffic offloading becomes manageable with this type of scheduler.

4.1.2. *Prioritized Scheduler*

While a balanced scheduler can be the most effective solution in certain situations, a prioritized scheduler—which is based on priority of certain channel types or user settings—is the preferred type of scheduler. By allowing priority to lower utilized capacity, the greatest amount of traffic offloaded from older capacity can be achieved. A simple example that most CMTS vendors have adopted is to prioritize OFDM channel traffic. OFDM capable cable modems need to utilize the full capacity of the OFDM channel before their bonded SC-QAM channels. This process offloads the OFDM capable cable modems' traffic from SC-QAM allowing more capacity for legacy services.

Though uncommon in the access network user controlled priority, would be the most powerful example of this type of scheduler. With DOCSIS 4.0 FDD, the expansion of OFDM to new spectrum would make the simple example above less effective. Ideally, a user would set the scheduler to prioritize the new spectrum OFDM channels over current OFDM channels.

Another example of the power of user controlled priority would be in low-split. With upstream channels that tend to have high forward error correction (FEC) error rates due to their location within the spectrum, these channels should be set to a low priority, which would mean that the channel is avoided until the capacity is required.

4.2. Dynamic Configuration Features

Dynamic configuration features like upstream agility and profile management application (PMA) can dramatically modify the capacity of a channel. Utilizing spectrum efficiency prevents these features from causing issues during planning.

5. Optimization Drivers

An understanding of the technology borders and data elements provides the background to start identifying optimization opportunities that can be applied to the access network. With numerous drivers for optimization, this paper will focus on service group capacity, customer experience, and cost reduction. The weighting of each of these drivers will be different for each network operator based on business goals.

5.1. Service Group Capacity

Historically the primary driver for access networks was congestion mitigation, achieved by adding additional service group capacity. This was typically done in tandem with segmentation to maintain enough capacity for IP services. However, as operators reach low levels of congestion more weight should be put behind offering higher burst capabilities for customers.

5.2. Customer Experience

Customer experience focuses on ensuring each customer in a service group achieves quality services. From the capacity management perspective, customer experience established through the burst capacity for that customer. Note the capacity planning is not focused on plant conditions or in-home issues.

5.3. Cost Reduction

When performing capacity management, cost reduction can be achieved through two methods. These methods include license reduction and service group combining.

5.3.1. License Reduction

Depending on the CMTS product, a license is likely utilized for each type of capacity. If the deployment of capacity is in excess, there is an opportunity to reduce the capacity. Each operator will have different agreements with their CMTS vendors, and as a result cost savings will differ from operator to operator.

5.3.2. Service Group Combining

Another option for cost reductions is the combining of two service groups that have low utilization. This reduces the license use by half between those two service groups. Furthermore, it may free up a service group resource for use for another HFC node, though this benefit would not reduce costs.

6. Identifying Optimization Opportunities

For long term forecasting our industry commonly utilizes the 2014 traffic engineering formula[1]. This formula is wonderful to forecast capacity needs for a service group going into the long term future.

$$C \geq (N_{sub} * T_{avg}) + (K * T_{max_max})$$

For access network optimization this formula is still interesting for forecasting but utilizing additional formulas to find opportunities. These opportunities are to maximize each serving group for products offered, and to continue to support legacy products. The starting point for optimization identification is burst capacity. The rest of the data elements discussed within this paper are supporting the decision making process around burst capacity.

6.1. Legacy Capacity Need Identification

As the access network moves towards new capacity methods, identification of legacy capacity need is beneficial in order to support legacy cable modems. Each access network will have service groups that behave outside of the norm and this process will identify them.

Utilizing burst capacity of legacy capacity, helps create a better understanding of the remaining capacity that is available during peak hours. By taking the highest offered service on a legacy cable modem and subtracting its value from burst capacity, you are left with remaining legacy capacity.

$$\text{Legacy Burst Capacity} - \text{Highest Legacy Service Tier} = \text{Remaining Legacy Capacity}$$

Table 10 – Example Data for Legacy Capacity

Data Element	Value (Mbps)
Legacy Burst Capacity (32 SC-QAM Channels)	342
Highest Legacy Service Tier	250

As an example, we can use the figures in Table 10 above to calculate the remaining legacy capacity:

$$342 \text{ Mbps} - 250 \text{ Mbps} = 92 \text{ Mbps}$$

There is 92 Mbps of remaining legacy capacity at peak burst capacity. We understand that the highest service on legacy services can achieve their burst speeds during peak hours. So, if the remaining legacy capacity was equal to or less than zero, this would reflect a greater need for legacy capacity.

6.2. New Capacity Need Identification

Typically, new technology capable cable modems have access to the new technology and legacy capacities. A simple example of this is a DOCSIS 3.1 cable modem would have access to both the SC-QAM and OFDM capacity. Using a similar formula as in the legacy example, you can calculate remaining capacity for new capacity.

$$\text{Burst Capacity} - \text{Highest Service Tier} = \text{Remaining Capacity}$$

Table 11 – Example Data for New Capacity

Data Element	Value (Mbps)
Burst Capacity (32 SC-QAM + 1 192 MHz of OFDM Channels)	848
Highest Legacy Service Tier	1000

We can apply the figures in Table 11 above to calculate the remaining capacity for new capacity:

$$848 \text{ Mbps} - 1000 \text{ Mbps} = -152 \text{ Mbps}$$

From this example we lack 152 Mbps of burst capacity for new capacity modems. A second OFDM channel could be utilized to gain further capacity for the DOCSIS 3.1 cable modems.

7. Optimization Cases Studies

In order to bring together the preceding information together thus far, this section will address several case studies. The scenarios presented below are commonly experienced by operators and can provide direction on how to optimize access networks. The cases studies are as follows:

1. Service group congestion
2. Spectrum efficiency congestion
3. Billboard service with low burst capacity
4. Low utilization service group
5. High utilization of legacy capacity
6. Spectrum boundary

7.1. Case Study 1 - Service Group Congestion

Congestion of a service group is typically handled by adding more capacity, and this additional capacity can be created by adding DOCSIS spectrum or a segmenting the HFC plant. Due to the state of congestion, using per DOCSIS channel utilization and burst data becomes challenging, but understanding each cable modem's usage and capability can provide insight into the type of capacity that will be required into the future.

This case study is based Node 14A, which is capable of 1 GHz with an 85 MHz return. The current DOCSIS configuration and utilization is:

Table 12 – Capacity – Case Study 1

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
32 Downstream SC-QAM	1216	1,140	76
OFDM (114 MHz @256-QAM)	774	644	130
6 Upstream SC-QAM (@64-QAM)	156	61	95
OFDMA (25.4 MHz @512-QAM)	195	77	118

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.




Figure 1 – Hourly Downstream Traffic Graph – Case Study 1

Because the upstream is not congested, the current configuration looks strong. However, on the downstream there is heavy usage on both channel types. At its current configuration, the downstream channel configuration only reaches 750 MHz, so spectrum expansion is recommended.

The Internet/IPTV cable modem distribution is:

Table 13 – Cable Modem Distribution – Case Study 1

Capacity Type	Count	% Of Distribution
DOCSIS 2.0	0	0%
DOCSIS 3.0	160	36.5%
DOCSIS 3.1	278	63.5%

This cable modem distribution and the OFDM channel flatline illustrates that adding OFDM capacity will also lower SC-QAM usage.

After expanding OFDM spectrum to 192 MHz and adding 1024-QAM flat profile, 873 Mbps of capacity was added to OFDM. The new DOCSIS configuration and utilization is as follows:

Table 14 – Updated Capacity – Case Study 1

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
32 Downstream SC-QAM	1216	694	522
OFDM (192 MHz @1024-QAM)	1647	1136	511
6 Upstream SC-QAM (@64-QAM)	156	53	103
OFDMA (25.4 MHz @512-QAM)	195	58	137

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.




Figure 2 – Updated Hourly Downstream Traffic Graph – Case Study 1

These updates have eliminated congestion; in addition, we can support the current product offerings. As seen in figure 2 a significant behavior change on the SC-QAM capacity can be observed.

7.2. Case Study 2 – Spectrum Efficiency Congestion

When dealing with multiple boundaries it is possible for legacy capacity to reach congestion without impact to the entire service group. This example shows how spectrum efficiency can cause congestion that appears like scenarios such as poor cable modem distribution. The root cause can be hidden and difficult to identify if the data elements being used are too few. In the example below, the problem presents as a legacy capacity issue but is an OFDMA channel performance issue.

This case study is based Node 3435B, which is capable of 1 GHz with an 85 MHz return. The current DOCSIS configuration and utilization is:

Table 15 – Capacity – Case Study 2

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
48 Downstream SC-QAM	1824	127	1697
OFDM (168 MHz @1024-QAM)	1438	134	1304
6 Upstream SC-QAM (@64-QAM)	156	153	3
OFDMA (25.4 MHz @512-QAM)	195	72	123

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.




Figure 3 – Hourly Downstream Traffic Graph – Case Study 2

The downstream capacity is strong for this service group, but upstream SC-QAM burst capacity is very close to zero. This appears to be a cable modem distribution issue. The Internet/IPTV cable modem distribution is:

Table 16 – Cable Modem Distribution – Case Study 2

Capacity Type	Count	% of Distribution
DOCSIS 2.0	0	0%
DOCSIS 3.0	68	54.4%
DOCSIS 3.1	57	45.6%

Upon closer examination, the majority of the codewords on OFDMA are passing at 16-QAM, a much lower modulation order than 512-QAM. This service group also has several cable modems that are currently impaired on OFDMA. At 16-QAM the OFDMA channel can only achieve 86 Mbps. Correcting this upstream performance issue could restore 109 Mbps of upstream capacity back to this service group. As seen in the information below, correction of this issue also initiated the correction of the traffic pattern.

Table 17 – Updated Capacity – Case Study 2

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
48 Downstream SC-QAM	1824	395	1495
OFDM (168 MHz @1024-QAM)	1438	991	447
6 Upstream SC-QAM (@64-QAM)	156	72	84
OFDMA (25.4 MHz @512-QAM)	195	108	87




Figure 4 – Spectrum Efficiency – Case Study 2




Figure 5 – Updated Hourly Downstream Traffic Graph – Case Study 2

7.3. Case Study 3 - Billboard Service with Low Burst Capacity

If service group congestion is not an issue, the focus moves to the next highest priority - supporting the highest service levels. Users of the highest services are generally heavier users, though in the majority of cases, the maximum service rate is rarely used. Ensuring efficient burst capacity for these services will enable the customer to achieve their max speeds during all hours. This will drive higher customer happiness as their service is capable during all hours.

This case study is based Node 402A, which is capable of 1 GHz with an 85 MHz return. The current DOCSIS configuration and utilization is:

Table 18 – Capacity – Case Study 3

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
32 Downstream SC-QAM	1216	899	317
OFDM (112 MHz @256-QAM)	774	624	150
6 Upstream SC-QAM (@64-QAM)	156	46	110
OFDMA (25.4 MHz @512-QAM)	195	87	108

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.




Figure 6 – Hourly Downstream Traffic Graph – Case Study 3

The upstream is not congested so the current configuration looks strong. On the downstream, despite heavy usage on OFDM, there is room on SC-QAM channels. The Internet/IPTV cable modem distribution is:

Table 19 – Cable Modem Distribution – Case Study 3

Capacity Type	Count	% Of Distribution
DOCSIS 2.0	3	1%
DOCSIS 3.0	92	32.1%
DOCSIS 3.1	278	66.9%

This service group has a high DOCSIS 3.1 distribution which explains the high OFDM usage and low SC-QAM usage. The top service tier for this service group today is 1 Gbps/100 Mbps, and the current burst capacity on the downstream is 803 Mbps, and most of it is on SC-QAM channels. Because the 1 Gbps service tier is only offered on DOCSIS 3.1 cable modems, increases to OFDM capacity would be the best path forward. In this example, reclaiming SC-QAM capacity for OFDM is not required since additional spectrum for OFDM is available.

After expanding OFDM to 192 MHz and adding 1024-QAM flat profile, 873 Mbps of capacity was added to OFDM. The new DOCSIS configuration and utilization is:

Table 20 – Updated Capacity – Case Study 3

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
48 Downstream SC-QAM	1824	590	1234
OFDM (192 MHz @1024-QAM)	1647	940	707
6 Upstream SC-QAM (@64-QAM)	156	46	110
OFDMA (25.4 MHz @512-QAM)	195	87	108

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.




Figure 7 – Updated Hourly Downstream Traffic Graph – Case Study 3

This configuration change has provided considerably more burst capacity, but due to the OFDM priority of the MAC scheduler the greatest usage change occurred to OFDM not SC-QAM capacity. Now, not only does this service group now have the burst capacity to support the 1 Gbps product, it can also support a higher service offering if offered.

7.4. Case Study 4 - Low Utilization Service Group

This case study applies to situations in which operators must reclaim capacity in order to reduce licensing costs. Licensing agreements with each multiple system operators (MSO) can vary from vendor to vendor. If reduction of DOCSIS spectrum can yield a license cost savings this case study provides a good example.

This case study is based Node 236A, which capable of 1 GHz with an 85 MHz return. The highest service tier offered is 1.5 Gbps/100 Mbps. The current DOCSIS configuration and utilization is:

Table 21 – Capacity – Case Study 4

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
48 Downstream SC-QAM	1824	100	1724
OFDM (192 MHz @256-QAM)	1647	154	1493
6 Upstream SC-QAM (@64-QAM)	156	72	84
OFDMA (25.4 MHz @512-QAM)	195	87	108

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.

This service group is efficient in its upstream capacities. However, the downstream is very much over built. Because this is an economic choice there is some difficulty in determining the “right” direction. Assuming cost savings were the same between SC-QAM and OFDM reductions, the most sensible option would be to reduce SC-QAM capacity by 50%. Pushing more capacity towards the ideal conditions for capacity per hertz.

7.5. Case Study 5 - High Utilization of Legacy Capacity

After the deployment of new capacity legacy capacity customers maybe impacted *if* their capacity was reduced during the process. Even with extensive planning, abnormal service groups can appear to have congestion or high utilization of legacy capacity.

This case study is based Node 1, which is capable of 1 GHz with an 85 MHz return. The highest service tier offered is 1.5 Gbps/100 Mbps. The current DOCSIS configuration and utilization is:

Table 22 – Capacity – Case Study 5

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
48 Downstream SC-QAM	1824	366	1458
OFDM (192 MHz @256-QAM)	1318	604	714
6 Upstream SC-QAM (@64-QAM)	156	116	40
OFDMA (25.4 MHz @512-QAM)	195	63	132

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.




Figure 8 – Hourly Downstream Traffic Graph – Case Study 5

The Internet/IPTV cable modem distribution is:

Table 23 – Cable Modem Distribution – Case Study 5

Capacity Type	Count	% of Distribution
DOCSIS 2.0	5	1.8%
DOCSIS 3.0	124	43.5%
DOCSIS 3.1	156	54.7%

The downstream is distributed very well and can support top tier packages, and the downstream SC-QAM has a significant amount of capacity to support legacy capacity customers. The problem shows up on the upstream, where there is low burst capacity on SC-QAM channels. While the highest upload package sold for this service group for legacy capacity is 30 Mbps, and there is sufficient burst capacity to support that package, this is playing very closely to the edge.

The legacy capacity could be increased by adding an SC-QAM channel and reducing OFDMA by 6.4 MHz. This would increase the legacy capacity by 26 Mbps to 66 Mbps, but reduce the OFDMA capacity by 36-37 Mbps.

7.6. Case Study 6 – Spectrum Boundary

As the industry moves towards DOCSIS 4.0 the possible net increase of new capacity becomes significant. Recently, high-split upgrades were completed that can provide insight into a scenario with a marked increase in poor cable modem distribution of capable modems. This is an interesting case study that showcases what can occur with an activation of a new spectrum boundary.

This case study is based Node 7518, which is capable of 1 GHz with an 85 MHz return. The current DOCSIS configuration and utilization is:

Table 24 – Capacity – Case Study 6

Capacity Type	Capacity (Mbps)	Peak Utilization (Mbps)	Peak Burst Capacity (Mbps)
32 Downstream SC-QAM	1824	68	1756
OFDM (192 MHz @256-QAM)	1318	140	1178
6 Upstream SC-QAM (@64-QAM)	156	23	133
OFDMA (25.4 MHz @512-QAM)	195	39	122
OFDMA (64 MHz @512-QAM) – High-split	476	52	458

The CMTS' MAC scheduler prioritizes OFDM traffic on the downstream and tries to balance traffic on the upstream.



Figure 9 – Hourly Downstream Traffic Graph – Case Study 6

The Internet/IPTV cable modem distribution is:

Table 25 – Cable Modem Distribution – Case Study 6

Capacity Type	Count	% Of Distribution
DOCSIS 2.0	0	0%
DOCSIS 3.0	7	24.1%
DOCSIS 3.1 (Mid-split)	21	72.4%
DOCSIS 3.1 (High-split)	1	3.5%

Even with this service group having very low utilization it still shows inefficient use of the high-split OFDMA channel. Given there is no congestion, immediate actions to rectify spectrum inefficacy from occurring is not required. Traffic distribution could normalize if more high-split capable modems were added. Another option to increase efficiency would be to prioritizing the high-split OFDMA channel over the remaining capacity, which would drive all high-split modem traffic to the high-split OFDMA channel until full. This would require fewer cable modems to be switched out to maintain capacity below the high-split spectrum addition.

8. Looking Forward to DOCSIS 4.0

As we move towards DOCSIS 4.0, what changes in the ways we manage technology borders to optimize access networks? The methods above can be applied, like an example reduction of legacy capacity in place of new capacity. DOCSIS 4.0 increases the importance of understanding the necessary planning in the reduction of legacy capacities, which has the potential to impact legacy services. DOCSIS 4.0 comes in two major designs: FDD and FDX. These designs will have different impacts on their technology borders from each other.

8.1. FDD Changes

For FDD DOCSIS 4.0 access network will largely have the same types of boundaries discussed in this paper. However, there is greater risk to legacy services due to the conversion of forward spectrum to return spectrum. Strong pre-planning is required to avoid poor legacy services post upgrade. After the upgrade is completed, the operator is able to return to modem upgrades and spectrum management as before. Given that is the largest spectrum upgrade that has occurred so far, operators will need to be aware of a number of important considerations - from spectrum efficiency for each customer to identified plant condition issues - to maintain capacity for each customer. Each OFDM channel representing a large chunk of capacity, cable modems that are unable to utilize most OFDM channels could have poor experiences, but a single OFDM impairment can be non-impacting. DOCSIS 4.0 FDD will require the management of the reduction of SC-QAM capacity. The continual forward progression towards OFDM/OFDMA only networks will require the application of the above optimization approaches in optimization to know when to take the next step forward.

8.2. FDX Changes

With FDX, the change for DOCSIS 4.0 access network is greater than that of FDD. Despite FDX having the advantage of software upgraded spectrum (up to 684 MHz), this capacity is *not* available to the whole service group, but is shared between interference groups (IG). This creates an additional technology boundary to manage. Due to the capacity gains achieved by this upgrade interference groups will not be an issue at the inception of this technology, as time goes forward it has the potential to restrict product offerings and capacity management options.




Figure 10 – DOCSIS 4.0 FDX N+1 Interference Groups[2]

In the diagram above, an N+1 FDX setup and corresponding interference groups are displayed. For capacity issues with IG 4, segmenting Amp 1 to its own Fiber Node is a logical solution. However, congestion on IG 1-3 how is this managed? Do you segment between these IG to management congestion? As is evident, the new technology boundary of interference groups will be challenging to overcome, but challenges can be overcome.

9. Conclusion

Operators will be required to build processes to manage capacity by channel or capacity type, in addition to the service group level moving forward. This is a key steppingstone as we move forwards towards mature DOCSIS 3.1 access networks, and to the path forward to DOCSIS 4.0. Other key steppingstones are distributed access architecture (DAA), PMA, and multiple OFDM/OFDMA channel configurations. The next few years will bring a lot of change to access networks, but the industry is ready to manage this. Enjoy the road towards 10G.

Abbreviations

CER	Codeword error rate
CM	Cable modem
CMTS	Cable management termination systems
DAA	Distributed access architecture
DOCSIS	Data over cable service interface specifications
FDD	Frequency division duplexing

FDX	Full-duplex DOCSIS
FEC	Forward error correction
Gbps	Gigabits per second
GHz	Gigahertz
HFC	Hybrid fiber coax
IG	Interference group
IPTV	Internet protocol television
KPI	Key performance indicator
Mbps	Megabits per second
MHz	Megahertz
MSO	Multiple System Operators
OFDM	Orthogonal frequency-division multiplexing
OFDMA	Orthogonal frequency-division multiple access
PMA	Profile management application
SC-QAM	Single carrier quadrature amplitude modulation

Bibliography & References

[1] SCTE 2014: *Traffic Engineering in a Fiber Deep Gigabit World*; John Ulm and Tom Cloonan; Society of Cable Telecommunications Engineers

[2] SCTE 2021: *Developing The DOCSIS 4.0 Playbook For The Season Of 10g*; Dr. Robert Howald and John Williams; Society of Cable Telecommunications Engineers

Delivering Network Agility and Automated Operations with GitOps

A Technical Paper prepared for SCTE by

David Bainbridge
Senior Director, Software Engineering
Ciena Corporation
dbainbri@ciena.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Kubernetes as an Intent Engine.....	3
2.1. Device Model.....	3
2.2. YANG to Kubernetes.....	4
2.3. Model to Target Association.....	4
2.4. Single Operator / Controller	4
3. Reconciling Model Controller	6
3.1. Order of Operation and Eventual Consistency	7
4. Intent	8
4.1. Abstract to Specific Decomposition	8
4.2. Intent Relationships.....	9
4.3. Intent Specification and Alternative Tooling.....	10
4.4. Unified Intents	11
5. Reconciliation	11
5.1. Drift.....	12
5.2. Drift Prediction.....	12
6. Multiple Sites.....	12
7. Source of Truth and the Pipeline.....	13
7.1. Test and Digital Twin Environments	14
7.2. Rollback.....	14
7.3. Kubernetes Drift from Git	14
7.4. Lockdown	15
7.5. Two Sets of Eyes	15
7.6. Cultural Change	15
8. Conclusions.....	16
8.1. Future Work.....	16
8.1.1. Abstract Intents	16
8.1.2. Troubleshooting	16
8.1.3. Multi-site Intents	17
9. Acknowledgements	17
Abbreviations	17
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - Example Model to Target Relationship.....	4
Figure 2 - Sample command to add YANG model into system	5
Figure 3 - YangMetadata Instance Examples.....	6
Figure 4 - Reconciliation Behavior	7
Figure 5 - Intent Decomposition Example.....	8
Figure 6 - Abstraction / Intent Decomposition Models.....	10
Figure 7 - Multi-Resource Domain Intent Decomposition	11
Figure 8 - Representation of GitOps Pipeline	14

1. Introduction

The term “DevOps” is derived from the combination of the terms “development” and “operations”. While there is no universal meaning of DevOps, it is generally accepted that it is a combination of specific practices, culture change, and tools [1] intended to reduce the time between committing a change to a system and the change being placed into normal production, while ensuring high quality [2]. Most commonly today DevOps is used in the operation of compute and storage resources, but increasingly these same concepts are being applied to connectivity (overlay and underlay network) resources.

The term “GitOps” describes additional practices, culture change, and tools that can be applied to DevOps to enable the state of a system to be version controlled through a persistent repository, such as the Git repository system [3]. As the state is updated in the versioned repository, automated processing pipelines are leveraged to verify, validate, and deploy the changes into normal production, while ensuring high quality. Like DevOps, GitOps is commonly applied to compute and storage resources today, but it is increasingly being applied to connectivity (overlay and underlay network) resources.

This paper describes how we integrated full network capabilities into a GitOps paradigm by first establishing a network model and integrating that model into Kubernetes to produce a capability that reconciles desired network state (intent) to a physical or virtual network. This paper then describes how the defined base network model can be leveraged to compose higher level or abstract intents changing the infrastructure from a “how” configuration model to a “what” configuration model. Finally, the network operator culture changes required to transition an organization to a controlled and stable GitOps infrastructure are discussed along with lessons learned and future work.

2. Kubernetes as an Intent Engine

Kubernetes is commonly used as an intent engine to orchestrate and reconcile the state of workloads onto a cluster of compute nodes. The implementation of Kubernetes provides many capabilities that intent systems require, including best practices for model definition, model control mechanisms, reconciliation capability, and model decomposition. Kubernetes also provides core context capabilities around the intent functions for security, scale, and tooling. This combination of capabilities makes Kubernetes a good platform on which to build a system that provides intent-based network operations. As Kubernetes already provides intent capabilities for non-network domains, building NetOps on Kubernetes means that a single intent can express a desired state across network, compute, and storage. Thus, Kubernetes can provide the backbone for a unified reconciling control plane.

2.1. Device Model

The most specific level of an intent system correlates to a specific target’s configuration, in that when a specific configuration parameter is set on a target the operator is essentially specifying their intent for the parameters value. When designing our intent system, we decided to start at this most specific level and considered the models that should be used for and directly map to the targets.

Evaluating the targets of which our network consisted as well as industry standards it was found that a common management paradigm of NETCONF and YANG models was applicable. We considered attempting to create a unique device model and then create adapters to various vendor targets, but instead opted to support the existing YANG models both to simplify the overall system and because historically, in our experience, least common denominator (LCD) or common models tend not to be complete or successful long term. Simply put, models already existed, and we saw no justification to not use them.

2.2. YANG to Kubernetes

Kubernetes operates on a set of resource definitions, or models, defined using open schema, known as the Kubernetes resource model (KRM). The first step in supporting the existing YANG models in Kubernetes was translating the YANG models into the KRM. Extending the KRM in this way is referred to as extending Kubernetes with custom resource definitions (CRD) and involves not only creating the schema model, but also developing a controller to provide the behavior or implementation behind the model.

KRM schema is defined using Open API Schema [5] specification and so we built a tool that produced Kubernetes compatible CRDs from YANG source models. These models could then be imported into Kubernetes and instances of these models could be created.

2.3. Model to Target Association

Leveraging the existing YANG models exposed a situation which was not originally considered, and is not common in existing Kubernetes models. Under a NETCONF/YANG configuration model, multiple, independent models are applied to a single target under control (TUC), where each associated model represents a unique configuration intent. In Kubernetes it is typical that a TUC and its configuration model are a single manifest or model definition. To account for this new associative model, we separated the definition of the TUC and the intent models into separate resource definitions with a reference from the models to the TUC to which they are applied, as depicted in Figure 1. By defining the association from the model to the TUC, the solution provides support for multiple models and allows associated models to be added or removed without having to update the TUC directly. Using Kubernetes's label selector pattern to associated models to TUCs was briefly considered, but discarded as this pattern tends to implement a 1:N relationship, while in a network environment a specific model instance is typically not applied to multiple TUCs, and thus only a 1:1 relationship was required.

It was determined that the model instance that represented a TUC would include the information required to provide connectivity to the target and the models to be applied to the TUC would reference the target via an "annotation" as described by the Kubernetes models [4].

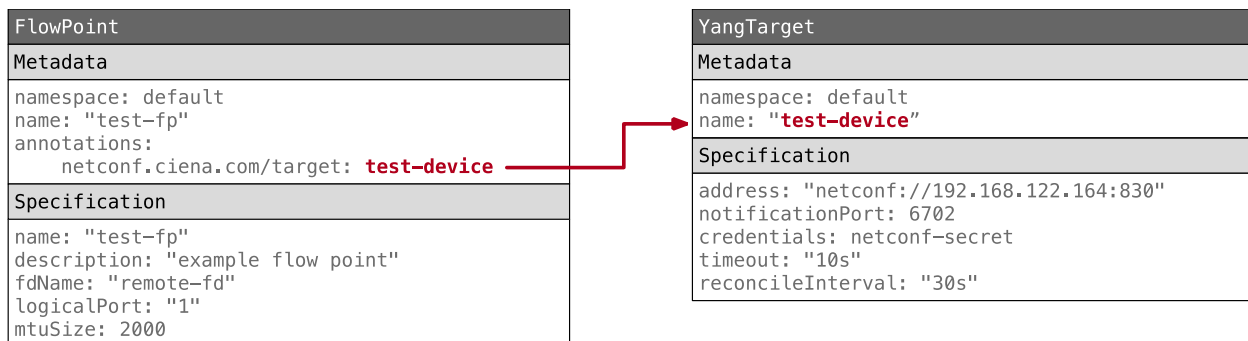


Figure 1 - Example Model to Target Relationship

2.4. Single Operator / Controller

In Kubernetes it is common that a single CRD is controlled by a unique controller that interprets the model changes and realizes them in the TUC, such as creating a Pod based on the Pod resource definition. When leveraging YANG as a CRD, if a unique controller was associated with each YANG model it would mean each model would require a controller and each of these controllers would have near 100% code overlap.

<working>

An additional goal was to allow YANG models to be added into the system without requiring the addition of more code, i.e., a new controller or the modification of an existing controller. This would allow individuals that were not “programmers” to more easily augment the system to support new models. This goal, coupled with the realization that most YANG controllers would have near complete code overlap, drove our decision to build a single controller that could provide the behavior for all YANG models. It was felt that this approach would be more sustainable as well as allow for dynamic model additions. Using this approach meant an individual could execute a tool that would generate a CRD schema from the YANG model and then add the new model into the Kubernetes environment with a simple command line action, as depicted in Figure 2. This capability also could be used in conjunction with the ability to query YANG schema from the TUCs to directly pull and add the models supported by a TUC into the system at runtime.

```
$ yang-to-crd -I /path/to/models ./new-model.yang | kubectl apply -f -  
customresourcedefinition.apiextensions.k8s.io/new-model.company.com created
```

Figure 2 - Sample command to add YANG model into system

During implementation, it was discovered that some YANG models had attributes that duplicate names already common to the KRM, such as name. We originally planned to eliminate this repetition to simplify usage of the model, but this led to the need to have custom mappings from existing KRM attributes, such as name, to YANG specific attributes, such as name or fdName. This quickly led to model specific behavior in the operator / controller or external metadata that represented that mapping that could not be automatically generated. Both these scenarios meant a more complex procedure to enable dynamic additions of YANG models at runtime. This path was abandoned, and it was determined that there should be a separation of the KRM object information, represented as Kubernetes metadata, i.e., Name, Namespace, and Labels, from the models which represented the TUC configuration, i.e., the YANG model.

Following this pattern of separation meant that we could provide a generic translation from the specification portion of the CRD to a NETCONF/YANG request, but also meant that there is some repetition between the Kubernetes metadata and the specification. It was determined that this tradeoff was worth being able to provide a generic translation capability, which additionally allows the system to dynamically add new YANG models without recompiling existing controller code, creating a customized controller per YANG model, or accommodating a manually generated mapping via metadata.

During the implementation of the portion of the controller that generates NETCONF/YANG requests from the CRD model, it was discovered that some level of YANG model specifics needed to be included and could not be avoided. To avoid embedding specifics into the operator / controller, patterns were determined and an additional CRD was developed to specify how these patterns applied to the CRD representation of a YANG model. This CRD, named YangMetadata, allows for the specifics such as key fields, YANG module information, and model XML namespace information to be specified externally and used to drive the translation of a CRD specified model to a NETCONF/YANG request. An example of a YangMetadata instance is depicted in Figure 3. It is important to note that the information contained in the YangMetadata resource relates only to the mechanics of translating one model to the other and does not involve mapping KRM attributes to YANG model attributes.

YangMetadata	YangMetadata
Metadata	Metadata
namespace: default name: "fp-metadata"	namespace: default name: "interface-metadata"
Specification	Specification
reference: kind: Fp group: netconf.ciena.com apiVersion: netconf.ciena.com/v1alpha1 wrapper: xmlRequestName: "fps" xmlName: "fp" keyField: "name" module: "ciena-mef-fp.yang" moduleSearchPath: "yang"	reference: kind: Interface group: netconf.ciena.com apiVersion: netconf.ciena.com/v1alpha1 wrapper: xmlRequestName: "interfaces" xmlName: "interface" keyField: "name" module: "openconfig-interfaces-ciena.yang" moduleSearchPath: "yang" augmentor: config.type: "xmlns=http://ciena.com/ns/yang/ciena-openconfig-interfaces" config.underlayBinding: "xmlns=http://ciena.com/ns/yang/ciena-underlay-binding" ipv4: "xmlns=http://ciena.com/ns/yang/ciena-openconfig-if-ip"

Figure 3 - YangMetadata Instance Examples

By externalizing the patterns required for special handling of the translations, the reconciling model controller maintains its independence from any specific model and allows models to be dynamically added to the system.

It is possible that there is overlap between the YANG models supported by a single TUC, i.e., two different YANG models can be used to manage a single TUC attribute. While this is acknowledged as a possible issue, in the system being described, the user controls which models are applied to which TUCs and, as such, this overlap can manually be avoided at present. As the implementation progresses some level of precedence may need to be introduced into the model structure to deterministically process model overlap.

3. Reconciling Model Controller

Kubernetes functions as a set of models that are reconciled by a set of controllers that perform operations when model instances or observed states change. The controllers can act on both virtual and physical resources. For example, when an instance of a standard Kubernetes [virtual] resource named Deployment is created, it creates another standard Kubernetes [virtual] resource named ReplicaSet. The ReplicaSet in turn creates Pod [virtual] resources, which are then physically realized on a compute capability as set of container instances. If an operator modifies the intent, as specified by the Deployment resource, the changes are propagated by the controllers through the ReplicaSet, Pods, and containers. If the status of the container changes, the status is propagated from the container, through the Pods, ReplicaSet, and Deployment.

As state and status are propagated the controllers can modify the resources that exist. For example, if the operator changes the value of the replicas property on the Deployment, this changes value replicas on the ReplicaSet, which in turn determines how many Pod instances are created. If a container fails, then this status is propagated to the Pod resource where the Pod's status is updated. This status is then propagated to the ReplicaSet controller, which in turn may delete the failed Pod and create a new Pod instance.

This behavioral model has been implemented as part of our solution. When the resource instances that represent a YANG model are applied to a TUC, the solution's controller converts the model from the CRD representation to a NETCONF/XML edit request to the TUC. The controller is also watching for asynchronous change notifications from the TUC so that if the TUC's configuration is changed the controller will receive a notification and reconcile the desired state, as represented as the Kubernetes resources, back down to the TUC. This reconciling control loop, depicted in Figure 4, ensures that the system continually maintains the operator's desired state.

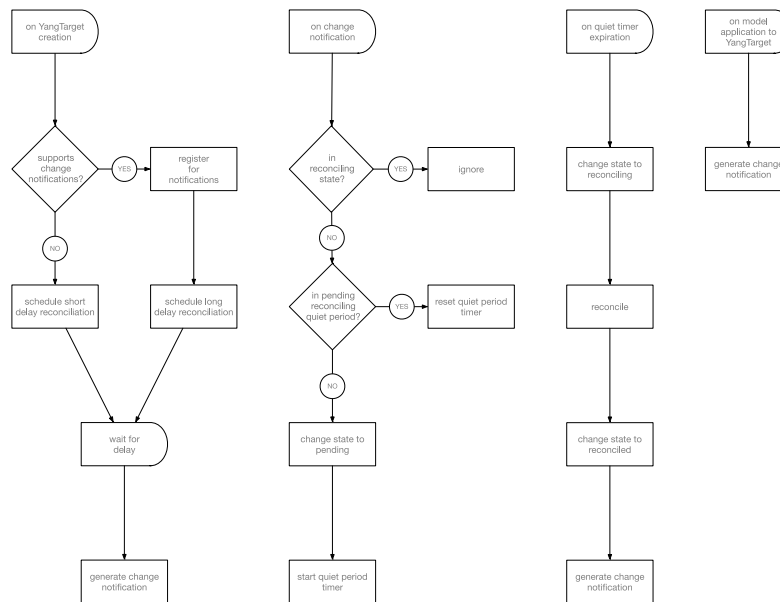


Figure 4 - Reconciliation Behavior

The difference of a TUC's observed or actual state from the desired state is known as drift. When a TUC supports asynchronous change notification the system utilizes that capability to optimize drift detection, using the change notification as an indication of drift. For TUCs that don't support asynchronous change notification, a periodic poll and drift determination is used.

A further optimization is the implementation of a quiet period. Experience demonstrated that changes to a device typically come in rapid succession and thus attempting a reconcile on every change is compute intensive and not efficient. To optimize for this situation, a quiet period is used so that reconciliation will not proceed until no change for a given TUC is detected within this quiet period. This mechanism has the effect of squelching changes to minimize reconciliation without compromising the time it takes to complete the reconciliation.

3.1. Order of Operation and Eventual Consistency

When configuring a TUC using traditional mechanisms or manual manipulation (CLI), order of operations can be important. This is due to dependencies between objects. If object A references B and then A cannot be created until B is created. Management systems, including orchestrators, often are designed, and implemented with knowledge of these order dependencies to ensure the higher-level goals are achieved.

In a model-based reconciliation system, order of operation does not have to be considered by the operator or the implementation. Using the previous example, if an instance representing A is created before an instance representing B, the reconciling of A to the TUC will fail, but will be continually retried. Eventually the reconciling of A to the TUC will succeed because the controller will also be working to reconcile B to the TUC. At some future time, B will exist on the TUC at the time A is reconciliation is retried, thus allowing A to be created.

4. Intent

An intent is the declarative specification of a desired state at some level of abstraction, the “what”, with limited or no direction to the implementation on “how” to achieve the desired state. There is a cultural change that an organization must make as they transition to intent-based specifications. This is because today, in general, the “how” is tightly bound to the “what” in many organizations, or at least the relationship between them is commonly known and utilized for daily operations including troubleshooting. The urge to include “guidance” as part of the “what” should be resisted so that the system that realizes the intent can fully optimize across multiple requests and the available resources.

4.1. Abstract to Specific Decomposition

Because an intent is specified at a level of abstraction, it can be decomposed to a set of more specific abstractions recursively until the level of abstraction is roughly the same as the specific configuration values of a TUC.

As an example, walking into a restaurant, asking the server to provide you food, and leaving the selection of the actual meal to the server can be viewed as an intent. The server might translate that level of abstraction to something more specific as the “daily special” when they submit the order to the kitchen. The kitchen staff decomposes the “daily special” into specific dishes and sides, and so on down to specific food elements and their preparation. At each level of intent specification, a set of declarative states can be defined that represents the translation of the more abstract request to one or more increasing specific intents. This same decomposition, in the context of network resources is depicted in Figure 5.

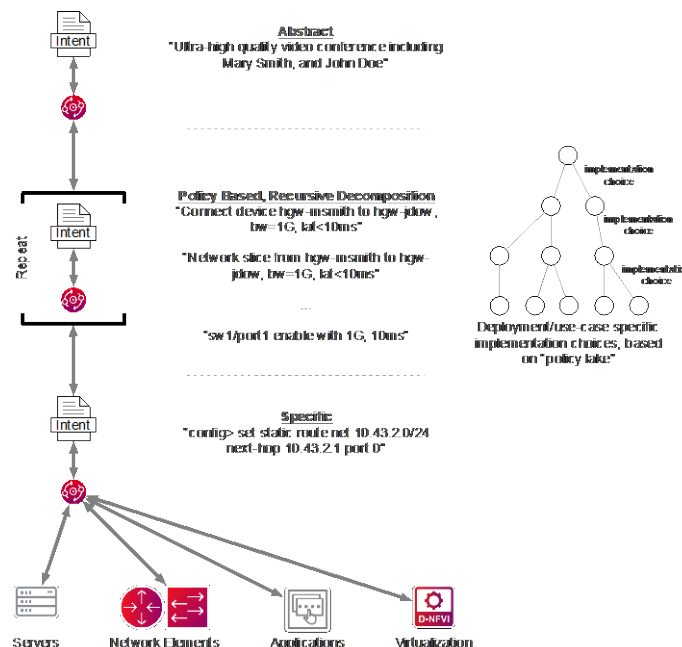


Figure 5 - Intent Decomposition Example

In the system being described, the most specific intent is represented by the YANG models that are applied to a TUC. On top of these intents increasing abstract sets of intents can be established, such as Fabric, Bridge, and Link. For each of these abstract intents there exists a controller that analyzes the intent and generates more specific intents that will realize the desired state.

As abstract intents are decomposed into a set of more specific intents decisions are made by the controller of the abstract intent with respect to how to implement the abstraction, as depicted in Figure 5. Decisions such as which more specific intents should be leveraged, i.e., VLAN or segment routing, as well as which TUCs should be considered.

How a more abstract intent is decomposed includes implementation logic that is outside the bounds of a single, “generic” controller. While a single controller manages all YANG models and their application to YangTargets, the larger system consists of multiple controllers to manage the abstract intents. This abstraction model, is common in Kubernetes for managing resources, as described in the Deployment example above and allows additional intents to be added to system at runtime as well as allows implementation up/down grades to existing abstract intents.

While we have experimented with more abstract intents, the bulk of the existing work has focused on the most specific intents, the YANG models. The abstractions we have experimented with, as a proof of concept (POC), express network fabric semantics to which end user hosts as well as intermediate network elements may be included as part of the fabric. While these POC abstractions successfully proved the point of intent decomposition, we have not formally defined a set of more abstract intents for our system

4.2. Intent Relationships

When considering decomposition, it is important to understand the relationship between the more abstract and the less abstract. In its simplest form this relationship could be considered a 1:N, creator to created relationship and is important during several lifecycle stages. For example, if resource A creates resource B, then when resource A is deleted, it is likely that resource B should be deleted. Additionally, if an attempt is made by an operator to directly delete a resource created by another resource (i.e., B), then this operation should fail because the resource that created B (i.e., A) should control the lifecycle of B. In short, when one resource creates another as part of a decomposition it has a responsibility towards the created resource’s lifecycle.

Complicating this relationship model is the fact that when dealing with network resources it is possible that more than one higher level abstraction may decompose and influence the configuration of more specific abstraction. So the relationship between an abstraction and its decomposition is not necessarily 1:N, nor a tree dependency relationship, as depicted in Figure 6(A), but could be N:N and a graph, as depicted in Figure 6(B). Thus, the relationship between intents (abstract to specific) might better be described as “interest” rather than ownership or creator to created.




Figure 6 - Abstraction / Intent Decomposition Models

With an interest relationship, the more specific intent cannot be removed until all more abstract intents are removed. At the same time, when a more abstract intent is removed, the more specific intent may be modified. Additionally, a more specific intent can be influenced by changes in more than a single more abstract intent, as such when an abstract intent changes, all “peer” abstract intents must be evaluated to determine the values of the shared, more specific intent.

While the Kubernetes framework does not prohibit the modeling of an acyclic dependency graph, it is not a common usage of the capabilities and may require special considerations when implementing such relationships. This pattern of reconciliation is not currently, to our knowledge, implemented as part of the Kubernetes framework. As this work continues, introducing this relationship will be avoided, if possible, as it significantly increases the complexity of the overall system.

One alternative to allowing an acyclic dependency graph may be to allow each more abstract intent to decompose into its own more specific intent and then predictably and consistently merge those intents when they are being applied to a single TUC.

4.3. Intent Specification and Alternative Tooling

The system being described utilizes YANG models to define the most specific intents as Kubernetes CRDs. Within this system, higher level intents can also be specified as CRDs, but can also be expressed using alternative tools including Helm [6] or Kustomize [7]. The result of these alternative methods is the generation of a set of KRM/CRD instances that represent some abstract intent. When using these alternative tools, there is no Kubernetes intrinsic relationship between the KRM/CRD resources deployed to Kubernetes and originating intent specification, i.e., Helm chart or Kustomize overlay.

As the Kubernetes operating model represents the state of the KRM, it is recommended that CRDs are used to define mature abstractions. We have found that tools such as Helm and Kustomize are valuable when prototyping abstractions, but as the abstraction matures being able to take full advantage of the KRM for implementation is advantageous and allows users to interact with the abstraction using the full ecosystem of Kubernetes based tooling.

It is important to note that while we do not recommend tools such as Helm or Kustomize for the modeling of abstract intents, other organizations may differ in opinion and nothing in the system being described prohibits from taking advantage of the full toolsets available through the Kubernetes communities.

Additionally, these tools can often be integrated into the GitOps processing pipelines, as described in section 7.

4.4. Unified Intents

Because the network intent model is built as CRD extensions to the KRM and it is instrumented via the Kubernetes controller framework, it is possible to create and deploy intents that include compute, network, and storage via a single model and toolchain. This allows infrastructure operators to define intents that are composed of these, and potentially other, resource types to ensure reconciliation across the resource domains, as depicted in Figure 7.

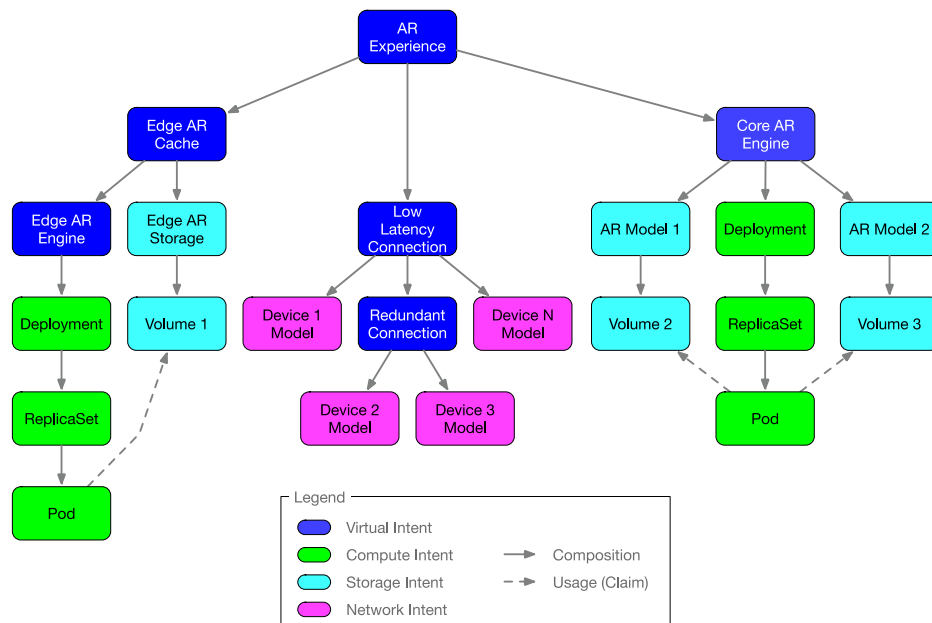


Figure 7 - Multi-Resource Domain Intent Decomposition

Again, while it is recommended that abstract intents be defined as extensions to the KRM as CRDs, this is not required, and they can be defined using many Kubernetes ecosystem tools including Helm and Kustomize. While GitOps was originally designed around the concept of directly deploying KRM resources into a cluster, many GitOps tools chains support tools like Helm and Kustomize directly so that, from the GitOps tool chain perspective, they are first class citizens. Integration of these tools into the GitOps tool chain extends the unification of the resource domains into the GitOps paradigm.

5. Reconciliation

An intent, as previously stated, declares the desired outcome, including capability and performance characteristics. To realize this outcome an intent is decomposed into more and more specific intents until the point where the intent is equivalent to a TUC configuration, including any configuration or modification to telemetry collection that is required to maintain performance characteristics of the intent. However, a TUC does not always maintain the state to which it was set. The difference between and intent's intended states and its actual or observed state is referred to as drift and can be caused by many factors, including errors, failures, or configuration change originating outside standard management practice, either intentional or malicious.

Drift in the system indicates that an intent or desired outcome is likely no longer being met, and thus, the drift must be addressed. When drift is encountered in a reconciling system, the system attempts to correct the drift in several ways, from the least disruptive to the most disruptive and from the most specific intent to the most abstract intent.

5.1. Drift

Within the system there are two types of drift which must be considered: static and constraint. Static drift (SD) is the difference between a desired configuration and the actual configuration. At the most specific intent level this can be the difference between the configuration set on a device and the device's actual configuration. In a "perfect" world SD is not encountered because all system change is driven through the control systems, but, as the world is not perfect, the control system must account for SD and when it is detected reconcile the desired state back to the actual. Reconciliation, in this context, consists of the control system resetting the TUC's configuration to the desired state.

Constraint drift (CD) is when the performance constraints specified by an intent are no longer being met even though there is no SD. CD is detected by telemetry collection and evaluation. When CD is detected the controller for the intent in which drift is detected re-evaluates how the intent can be realized and may create new more specific intents or modify/delete existing intents while attempting to eliminate CD. If a controller cannot eliminate CD, it will update the status of the intent. This status update serves as a trigger/indicator to more abstract intents that they need to be reevaluated to attempt to reconcile the CD. This process is recursive until a more abstract intent can eliminate the CD. If the CD cannot be eliminated then the status of the intents indicates the CD drift and it is up to the operator, or other controlling system to determine how to proceed.

It is important to recognize that CD may be a temporary situation, and that immediately attempting to correct for CD as soon as it is detected may not be the best action, as the effort and time to correct the CD may be more "painful" and take more time than the temporary CD might otherwise exist. For example, if the bandwidth for a connection is desired to be at 500 Mbps and CD is detected at 300 Mbps, the time it takes to reconfigure to a different path to maintain 500 Mbps may be longer than simply waiting for the CD or return to its normal operating state. This is not true in all situations, as in the case of a true failure, but can be true when the cause of the CD is a temporary "blip". Because of this, a waiting period should be observed before any correction of CD is attempted.

5.2. Drift Prediction

Drift prediction (DP) is the process of analyzing desired state, telemetry, and other sources of operation state and status to predict a future time where CD may occur. While our current implementation does not provide DP, it is believed that analysis, including AI/ML analysis, could be leveraged to calculate DP and use these predictions to re-evaluate intents before actual CD occurs, minimizing or eliminating potential issues.

6. Multiple Sites

Even in an organization where the desired state of the infrastructure, compute, network, and storage, is relatively static, the level of communication between the infrastructure TUCs and the reconciling controllers can be significant, particularly when you consider telemetry collection and CD detection. As such, it is prudent to separate an organization into domains each with their own control system. This practice can be seen in the evolution of Kubernetes with respect to workload control. In practice the

controller and the TUC should be “network close”, indicating sufficient bandwidth, connection stability, and reasonable latency.

When architecting multi-site control system, various patterns can be leveraged, such as hierarchical and peering. As our system is leveraging the Kubernetes platform, hierarchical control was a natural fit, where each site is a Kubernetes cluster and the clusters are aggregated via a multi-cluster console, such as provided by various open source and commercial products. A simple aggregation console does not address the issue of cross site intents but does allow administration of each site/cluster directly while providing a global view of a deployment.

Supporting cross site intents implies that an intent can be specified to a system such that the realization of the intent deploys and connects resources from multiple sites, while control of those resources are individually controlled by the site or Kubernetes cluster that “owns” them. To date, in our implementation we have not addressed the issue of cross site intents, but it is envisioned that a centralized Kubernetes cluster could house models and controllers that represent cross-site intents. These models, along with information about the remote clusters, could then be leveraged to implement intent-based configuration of our organization’s global infrastructure resources. As we believe our requirements and goals closely align to those of the larger Kubernetes ecosystem, we continue to monitor projects under organizations like the Cloud Native Computing Foundation (CNCF) and the Linux Foundation (LF) that consider cross-site or multi-site control architectures.

7. Source of Truth and the Pipeline

The operational source or truth (SOT) for the system being described is the model states as maintained by the Kubernetes reconciliation engine. These states can be made redundant, highly available, and persistently stored via standard Kubernetes deployment practices. These states represent the desired state toward which the system is working.

When extending the system beyond the operational state to utilize a GitOps environment, the SOT for the organization is maintained and versioned, in the git repository as a set of files that represent the desired states. In practice these states are represented as model instances formatted as YAML documents. These YAML documents, in Kubernetes parlance, are referred to as manifests. Events within the git SOT, such as the modification and commit of a manifest or the releasing (tagging) of a git repository that may contain many manifests, causes the state from the git repository to be applied to the operational system(s). At steady state, the state maintained in git and the operational state are equivalent. Figure 8 depicts a typical GitOps pipeline, including manual merge approval, automated unit testing, and multiple pre-production deployment environments.

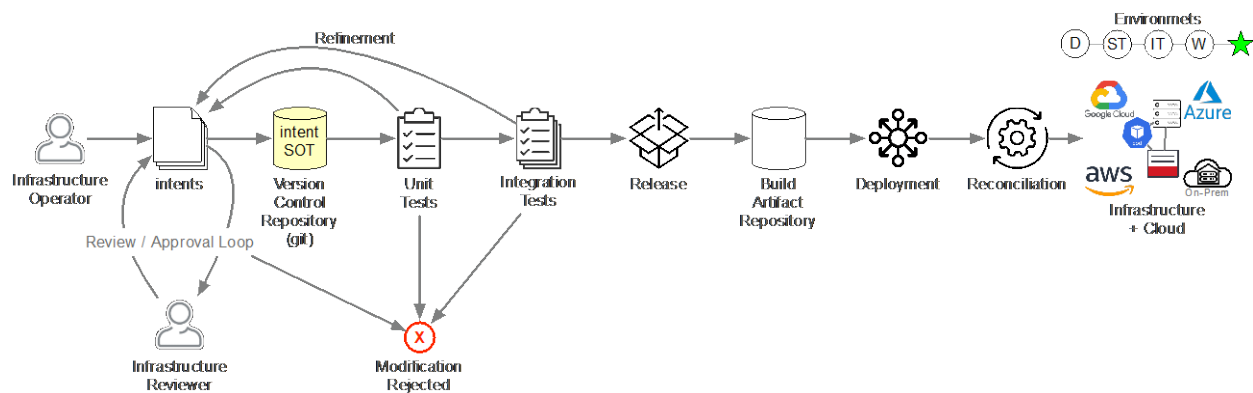


Figure 8 - Representation of GitOps Pipeline

7.1. Test and Digital Twin Environments

As part of GitOps, organizations can establish automated pipelines that verify a change in compute, storage, and network state before it is pushed into production. Minimally, this can include some basic static analysis and regression testing; but it can also include the deployment into a digital twin environment may include full system and failure tests.

The level of testing is up to the organization with the goal of not only protecting the production system from accidental configurations, but also improve the confidence that a new change will not adversely affect the capability or performance of the production system.

7.2. Rollback

Because the state of the environment (compute, storage, and networking) is maintained in a versioned repository and because the environment is controlled by a reconciling controller, it is possible to set the state of the environment back to any previous states by running the GitOps automated pipeline on a previous version of the desired state.

The pipeline will deploy the desired (previous) state to the Kubernetes reconciling control plane and the controllers within the Kubernetes control environment will reconcile the desired (previous) state to the actual state on the TUCs. Thus, rollback becomes the process of identifying the version of the desired state and executing the automated process pipeline, changing rollback from an anomaly in the environment to a normal process that is not to be feared.

7.3. Kubernetes Drift from Git

GitOps solutions do not typically monitor Kubernetes deployments for change or attempt to determine drift between the Git SOT and the operational SOT in Kubernetes. This is because it can be expensive and it is not required because of the Kubernetes reconciliation model.

When a state is applied to a Kubernetes control plane it compares values in the new state to values in the current operational state. If the values are equivalent, then no action is taken, and the application of the state is essentially a no-op. If the values are different, then the associated controllers are invoked, and the updated state is reconciled and propagated throughout the system.

This operational state allows GitOps solutions to simply reapply existing state multiple times with little computational consequence, rendering drift detection between git and Kubernetes unnecessary. Instead,

the GitOps systems re-apply the desired states at a [configurable] interval. Thus, git and Kubernetes are kept synchronized within the time error of this period.

7.4. Lockdown

A key aspect to any controlled environment is to minimize the ability of individuals to introduce unauthorized SD into the system. All changes to the TUCs **must** be handled through the GitOps pipeline, for normal operations as well as in times to failures.

This operational change is another instance of culture change an organization should be willing to accept before moving to GitOps management of their environment. While it may seem counter intuitive to force operators to drive changes through the GitOps pipeline in the presence of a customer affecting failure it is also, precisely at these times when extra care should be taken to remediate the problem to prevent a proposed remediation from exasperating the original issue and consequences.

While exceptions to this rule may be natural to consider, such as when the GitOps pipeline can no longer communicate to the Kubernetes control plane, it is recommended that this communication error between GitOps and Kubernetes be resolved first and then proceed through the GitOps process. One risk of not following this pattern and directly manipulating the configuration of the Kubernetes control plane or the TUCs is that when connectivity is restored, the GitOps pipeline will reconcile its version of the SOT, through the Kubernetes control plane, down to the TUCs, overwriting any direct manual change made and potentially reproduce the original problem.

7.5. Two Sets of Eyes

GitOps evolved from the software development practices related to continuous integration and continuous deployment (CI/CD). One of these practices is related to how code becomes accepted as part of the production code for a project. In this model, a developer develops some code and submits that code to be merged into the production code. Before the code is accepted, or merged, into production product, automated tests are run with the code and an individual with authority must review and approve of the code change. Within some projects, multiple people must review the code and the code is not accepted until a given number those individuals agree and approve of the code change.

With GitOps as part of infrastructure management, this process should be adopted. The equivalent of code when using GitOps for infrastructure management is an intent specification. When an operator makes a change to the desired intent state, this is submitted to the Git repository for automated testing, review, and approval. Only when the change is approved is it merged as part of the production SOT in Git and propagated to the Kubernetes control plane to be deployed in the production environment.

7.6. Cultural Change

GitOps, as previously defined is meant to provide a continuous, controlled, predictable, and stable change process while maintaining the highest level of correctness in the system. Before GitOps can successfully be deployed in an organization the wild west must be sacrificed as must the lone hero. Change process is controlled. Changes are automatically verified and reviewed regardless of the source. Even in the face of customer outages the process is followed. This can be a difficult pattern to follow and enforce when dealing with software development projects as many developers will attest. This can be even more challenging when attempting to apply this to infrastructure management.

8. Conclusions

A reconciling control plane is an effective control model for network elements and the service across those network elements. When the control plane used for network connectivity can be shared, as is the case with a Kubernetes control plane, the paradigm becomes even more powerful for infrastructure and infrastructure service control.

Adopting the culture and best practices from GitOps adds benefit on top of the core reconciling control plane and can lead to a more predictable and stable infrastructure change model. With the additions of manual review and automated checks in the GitOps pipeline, as well as system wide pre-deployment test in alternate environments, including digital twins, the incidence of misconfiguration can be minimized.

By introducing abstract intents into the system, the operators are allowed to specify the goals of the system while automated control loops optimize the implementation. It is important with the introduction of intents that the ability to troubleshoot in the context of explicit configuration be supported, including the ability to map from the explicit device configuration back to the intent or intents from which it was derived and why it was derived.

Kubernetes was a good choice for a system as described. It provided much of the “context” required for such a system, the ability to unify control across resource domains (compute, storage, and network), and provided opinionated processes for the definition of models and controllers. This opinionated process allowed most questions about the implementation of the system to be answered based on existing precedent and allowed the development to focus on those aspects that were unique to the solution.

When introducing a system as described into an enterprise it is important that the organization should feel empowered to extend the service intent model in ways that best meet their needs. This includes the development of new or organization specific CRDs that model required services. This is not an “out of the box solution” and an organization should expect and train their employees to be infrastructure model and control developers or rely on 3rd party product and support.

8.1. Future Work

8.1.1. *Abstract Intents*

As previous stated, the existing implementation has focused on the most specific intents. Work has started on developing more abstract intents. This work needs to continue in cooperation with the organization’s IT services groups so that the set of intents instrumented match the infrastructure services offered.

It is expected that as more abstract intents are developed issues may arise with decisions previously made and those decisions may have to be revisited. It is a guiding principle of this work to avoid complexity where possible, even at the cost of operational usability. It is believed that by keeping the underlying technology as simply as possible that the result will be a more stable system. It is also believed that operational usability can be managed through user interaction methods that implemented best practices or common usage patterns.

8.1.2. *Troubleshooting*

As this system evolves, the team will be investigating how to integrate better troubleshooting, performance monitoring, and predictive techniques. Being able to provide human consumable explanations as to why a system took a specific action or reevaluated an intent can be valuable when

understanding failures as well as it could provide a bases for an automated feedback system to the consumers of the services.

8.1.3. Multi-site Intents

Another area of exploration will be cross-site or multi-site intents and the comparison of a hierarchical solution vs. a peer-based solution. While we will heavily leverage the direction and learnings of the CNCF, there may be something unique with respect to network connectivity and decentralized environments that points us in a different direction. When controlling connectivity, the inter-site connections are as, if not more, important than the intra-site connections. Understanding that, for connectivity, the problem being addressed is not simply distribution of workloads across a decentralized environment. Site inter-connects are imported and must influence how workloads and storage are placed or replicated to optimize the connectivity to meet the desired performance and failure characteristics.

9. Acknowledgements

The author of this paper would like to acknowledge the efforts of Karthick Ramanarayanan and Himani Chawla for their contributions to both the design and implementation of the model-based reconciliation controller. Without their continued hard work this project might not have come to fruition.

The team would also like to show his appreciation to their manager, Marco Naveda, and their peers in the organization for their support, input, and suggestions. While they are not listed as authors on this paper, without their support it would not have been written.

Abbreviations

AI/ML	artificial intelligence/machine learning
CD	Constraint drift
CI/CD	continuous integration/continuous deployment
CLI	command line interface
CNCF	Cloud Native Computing Foundation
CRD	custom resource definition
DevOps	development operations
DP	drift prediction
GitOps	git operations
KRM	Kubernetes resource model
LCD	Least common denominator
LF	Linux Foundation
Mbps	megabits per second
NETCONF	network configuration
NetOps	network operations
POC	proof of concept
SD	static drift
SOT	source of truth
TUC	target under control
XML	extensible markup language
YAML	yet another markup language
YANG	yet another next generation

Bibliography & References

- [1] <https://en.wikipedia.org/wiki/DevOps>
- [2] Bass, Len; Weber, Ingo; Zhu, Liming (2015). DevOps: A Software Architect's Perspective. ISBN 978-0134049847.
- [3] <https://en.wikipedia.org/wiki/DevOps#GitOps>
- [4] <https://kubernetes.io/docs/concepts/overview/working-with-objects/annotations/>
- [5] <https://github.com/OAI/OpenAPI-Specification>
- [6] <https://helm.sh/docs/topics/charts/>
- [7] <https://kubernetes.io/docs/tasks/manage-kubernetes-objects/kustomization/>

Deploying PMA-Enabled OFDMA in Mid-Split and High-Split

A Technical Paper prepared for SCTE by

Maher Harb

Director, Data Science
Comcast
maher_harb@comcast.com

Dan Rice

VP, Access Architecture and Technology
Comcast
Daniel_rice4@comcast.com

Kevin Dugan

Data Scientist
Comcast
kevin_dugan2@comcast.com

Jude Ferreira

Principal Data Scientist
Comcast
jude_ferreira@comcast.com

Robert Lund

Principal Network Engineer
Comcast
robert_lund@comcast.com

Ramya Narayanaswamy

Director, Data Science
Comcast
ramya_narayanaswamy@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. The Upstream Split Change Evolution.....	4
3. OFDMA Capacity Analytics.....	5
4. Mid-Split Deployment Data	9
5. PMA for OFDMA	11
5.1. Why PMA for OFDMA.....	11
5.2. PMA Virtual Network Functions Development Starts and Ends in the Ingress Lab	13
5.3. OFDMA PMA Comes Together.....	13
5.4. State of the Upstream Spectrum.....	17
6. Mid-Split Ingress.....	19
6.1. VHF TV Ingress.....	19
6.2. Interesting Impairments to Keep in Mind for PMA	21
7. Conclusion.....	22
Abbreviations	22
Bibliography & References.....	23

List of Figures

Title	Page Number
Figure 1. Yesterday's low-split (sub-split) network utilizes mostly 750 MHz analog fiber coax network, with typically Node +6 Amps or less.....	5
Figure 2. Mid-Split expanded upstream using D3.1, increasing capacity by >4x and enabling speeds up to 300 Mbps; overall network spectrum is extended to 1 GHz enabling multi-Gbps downstream speeds.....	5
Figure 3. DOCSIS upstream capacities for the mid-split scenario. Total effective capacity is a function of PMA for OFDMA based on signal and noise quality, D3.1 penetration, and traffic, and consumption statistics.	6
Figure 4. Critical mid-split set of performance and capacity statistics.	7
Figure 5. Effective capacity correlated to OFDMA-to-D3.0 SC-QAM traffic ratio. Enabling PMA is expected to lift most nodes above the break-even line.....	8
Figure 6. Utilization and D3.1 cable modem penetration for ~20,000 to-be-converted analog nodes. The intersection of the 2 tails of the distributions reveals that only 1 node will be driven into alert level due to loss of the 5 th and 6 th SC-QAM channels to accommodate OFDMA in the mid-split.	9
Figure 7. SC-QAM and OFDMA 98 th utilization and traffic (Mbps). The vertical line in all panels marks the timestamp of the speed increases. The top row tracks the traffic on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines). The bottom row tracks the utilization on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines).	10
Figure 8. Per device (CM) total consumption for the 3 groups pre and post speed increase.	11
Figure 9: Lower VHF ingress into mid-split channel would have required 32-QAM or lower modulation without PMA solution.....	12

Figure 10: (right panel) Example high-split spectrum capture showing upper VHF and NOAA radio ingress. Mid-split channel has excellent signal quality. (left panel) PMA constructed profile for the high-split channel.	13
Figure 11. Lab testing flow. VHF Ingress is captured (left panel) and injected into the node using SDR platform (top right panel). PMA adjusts to the ingress by constructing the proper modulation profile (bottom right panel).	13
Figure 12. PMA architecture supports DS and US D3.1 and US D3.0 profile management.	14
Figure 13. Example PMA IUC/profile design across mid-split spectrum that improves the channel capacity compared to adopting a “flat modulation” IUC design.	15
Figure 14. Example PMA design for the high-split channel in Figure 10. Left panel shows the 7 constructed profiles. Since only 1 cable modem is bonded to the channel, 6 profiles were configured with flat modulation. The right panel shows the profile tailored to the noise detected on the channel.	15
Figure 15. Example MS PMA performance charts for a clean spectrum.	16
Figure 16. Example MS PMA performance charts for a noisy spectrum.	17
Figure 17. (Top panel) MER cumulative distribution from early trials on ~100 nodes shows that ~85% of mini-slots can support 2k-QAM. (Bottom panel) Actual traffic stats confirm this picture.	18
Figure 18. MER per minislot aggregated across tens of thousands of modems at 1-hour samples over 48 hours across varied network locations.	19
Figure 19. One primary source of ingress which impacts mid-split is the lower VHF OTA television broadcast. The channels of interest are 2-6 and FM broadcast	20
Figure 20. Variation in channel power vs distance for VHF transmitters for different device categories ...	20
Figure 21. Example video modulator in home transmitting into OFDMA spectrum.	21
Figure 22. Various Pre-EQ responses for OFDMA channel showing titlt.	22

1. Introduction

Comcast has deployed DOCSIS[®] 3.1 (D3.1) orthogonal frequency division multiple access (OFDMA) in the mid-split (MS) and high-split (HS) bands of the upstream (US) spectrum on our virtual cable modem termination system (vCMTS) platform. D3.1 OFDMA allows higher modulation levels with up to 2x efficiency increases, with 2048-QAM today and anticipating 4096-QAM in the near future, compared to single carrier-quadrature amplitude modulation (SC-QAM) of 64-QAM. More importantly, D3.1 US technology allows configuring the modulation per 400 KHz mini-slot, enabling adaptation to the ingress and distortions discovered in the network. Several ingress and impairment sources that have the potential to degrade capacity and customer experience have been identified; examples include off-air very high frequency (VHF) broadcast, linear distortions, in-home analog TV modulators, frequency modulation (FM) radio, NOAA weather radios, and pre-equalization stability. As we deployed OFDMA initially on the network we quickly realized that in order to take full advantage of the capabilities OFDMA has to offer, most notably delivery of hundreds of Mbps US product speeds, a profile management application (PMA) system is required similar in concept to the one deployed downstream for managing orthogonal frequency division multiplexing (OFDM) profiles, which we reported on in previous SCTE contributions [1-4]. This paper presents the initial results from adapting the OFDM PMA system for OFDMA and testing in both lab systems and production field deployments. We share some examples of spectrum impairments, characterization methods, our thinking around how profiles are constructed based on device-level MER measurements, and the validation of the profiles against the vCMTS internal PMA function. We also comment on some of the intricacies of expanding upstream spectrum that warrant revisiting the core algorithm in the future.

2. The Upstream Split Change Evolution

DOCSIS-based broadband access has historically been operated with lower upstream capacity capabilities due to the amount of spectrum available in traditional splits of the 42 MHz low-split (also known as sub-split), 85 MHz MS, 204 MHz HS, as well as the upcoming full-duplex (FDX) with adjustable upstream and downstream spectrum usage, including simultaneous spectrum overlap. Today many hybrid-fiber-coax (HFC) networks still operate with a low-split configuration as shown in Figure 1 and are limited to upstream speeds under 100 Mbps.

As operators migrate to an 85 MHz upstream MS spectrum band as shown in Figure 2, they can offer increased product speeds, for example 300 Mbps, and with a 204 MHz HS spectrum band, 1 Gbps upstream products are possible. While updating the network to one of these upstream split scenarios, downstream spectrum may also be expanded to include up to 1 or 1.2 GHz of spectrum. This spectrum update allows additional downstream data capacity. Managing this spectrum is the topic of another Comcast 2022 Technical Forum Paper focused on delivery of multi-Gbps DS services [5]. Ensuring the quality and speeds of this new US spectrum requires using D3.1 OFDMA technology, which is designed to adapt to network conditions to maximize capacity and robustness based on the same platform described in prior SCTE papers on the Comcast PMA solution [1-4].

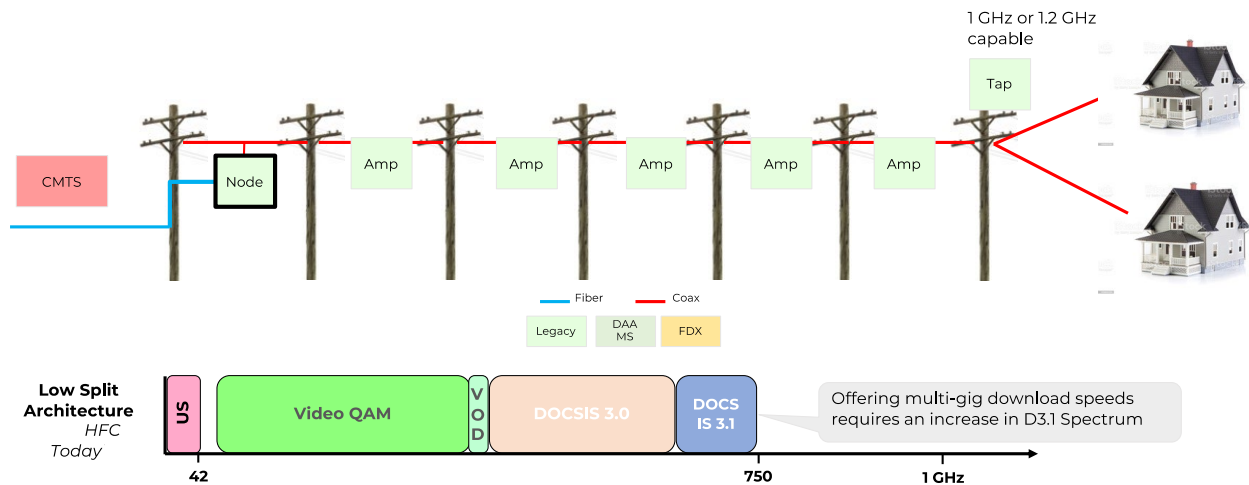


Figure 1. Yesterday's low-split (sub-split) network utilizes mostly 750 MHz analog fiber coax network, with typically Node +6 Amps or less.

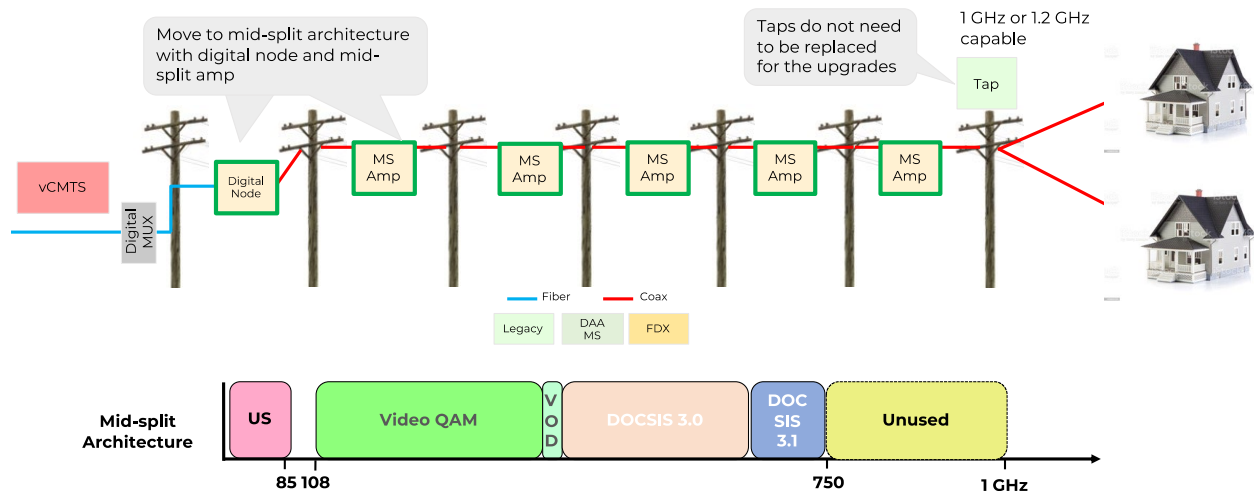


Figure 2. Mid-Split expanded upstream using D3.1, increasing capacity by >4x and enabling speeds up to 300 Mbps; overall network spectrum is extended to 1 GHz enabling multi-Gbps downstream speeds.

3. OFDMA Capacity Analytics

Deploying OFDMA in the mid-split spectrum increases capacity >4x compared to low-split pre-OFDMA US channel plans. During the COVID outbreak and the increase in bandwidth growth, particularly in the upstream direction, Comcast deployed six upstream channels and a PMA solution for D3.0 upstream channels to manage performance and rapidly add new capacity [2-3]. When deploying OFDMA spectrum on the nodes with high utilization and reducing the number of D3.0 channels, the resulting capacity needed to be equal to or greater bonded capacity than the six D3.0 channels in addition to enabling higher upstream product speeds. As a result, Comcast developed capacity analytics to understand in detail how the evolution to OFDMA would impact network performance from a capacity perspective.

It is instructive first to understand how this new capacity is provisioned to enable hundreds of Mbps product speeds. As shown in Figure 3 the channel plan was migrated from 6 US D3.0 channels to 4 D3.0 channels plus a single OFDMA channel. The blue line in the example node shows that the capacity of the 4 D3.0 channels is about 85 Mbps. This is less D3.0 capacity than the black line in the chart, which represents 100 to 120 Mbps of capacity available with the 6 US D3.0 channel lineup. As a result, Comcast needed to ensure that when we added the OFDMA capacity, enough traffic could be shifted into the OFDMA channel such that the net result would be equal or greater capacity available to the cable modem (CM) population on average. When the vCMTS bonds the D3.0 and D3.1 capacity for the D3.1 CMs in the service group (SG), it was possible – as depicted by the green line – to exceed 500 Mbps with the right PMA solution and network quality.

Depending on how the traffic is allocated across the channels, we are introducing a new metric known as “effective capacity” or “aggregate_speed” (yellow line in Fig. 3), which is the traffic-weighted capacity of the service group. This is a function of the amount of traffic enabled on the D3.1 OFDMA spectrum vs. the D3.0 spectrum along with the total capacity in the D3.0 and D3.1 channels based on the profile management solution optimization of the channels. Both the traffic allocation and the capacity of the channels change over time, based respectively on bandwidth demand and network quality. If the yellow line drops below the black line the goal of exceeding the D3.0 capacity when adding OFDMA is not met for that node and unit of time. This chart is one example node with 24% of the D3.1 modems able to access the OFDMA channel.




Figure 3. DOCSIS upstream capacities for the mid-split scenario. Total effective capacity is a function of PMA for OFDMA based on signal and noise quality, D3.1 penetration, and traffic, and consumption statistics.

The data of Figure 3 can be aggregated to the network level across a set of remote PHY device (RPD) Distributed Access Architecture (DAA) nodes. This aggregation for hundreds of nodes is illustrated in Figure 4 with the following key points:

- >1000 RPDs are mid-split activated (panel 1).
- ~20% of devices are OFDMA active with the potential to increase by > 10% as our in-home test solution [6-7] migrates each home to use of the OFDMA channel (panel 2).
- ~17% of traffic goes through the OFDMA channel on average; this is sufficient to increase the net effective capacity above the capacity of the full spectrum D3.0 (panel 3).

- Current traffic-weighted effective capacity (yellow) is higher than projected throughput if D3.0 spectrum was configured with 5th and 6th channels (i.e., the legacy six channel configuration shown as black line in panel 6).
- The consumption of individual customers on smaller nodes can influence the results such as a heavy consumer with a D3.1 vs. D3.0 CM.
- The abrupt increase in OFDM capacity (red line) represents the switch from 32-QAM to 256-QAM. As PMA is added and higher order modulations are included, as detailed in this paper, the net effective capacity will further increase (panel 6).
- 200 Mbps speed test passes > 98% of time with only 256 QAM.




Figure 4. Critical mid-split set of performance and capacity statistics.

While aggregating across all RPD nodes, the net effective capacity is positive at only 256-QAM flat modulation configured in the OFDMA channel. There may be specific nodes that fall below the break-even line due to challenges in CM's connection to the OFDMA channel, or due to the low penetration of

D3.1 modems along with the utilization of the D3.0 vs. the D3.1 spectrum.




Figure 5 illustrates how at the above OFDMA configuration, and without migrating the traffic reliably to OFDMA, the effective capacity could fall below the break-even line for D3.0 spectrum. Fortunately, for ~20k nodes in this analysis, only 1 is projected to have D3.0 utilization go above alert level without gaining benefit from OFDMA due to low D3.1 penetration as shown in Figure 6. Increasing the OFDMA capacity with the use of PMA, as illustrated in Figure 5 moves the modems all above the break-even line, as will the organic increase of the D3.1 relative CM population.




Figure 5. Effective capacity correlated to OFDMA-to-D3.0 SC-QAM traffic ratio. Enabling PMA is expected to lift most nodes above the break-even line.




Figure 6. Utilization and D3.1 cable modem penetration for ~20,000 to-be-converted analog nodes. The intersection of the 2 tails of the distributions reveals that only 1 node will be driven into alert level due to loss of the 5th and 6th SC-QAM channels to accommodate OFDMA in the mid-split.

4. Mid-Split Deployment Data

OFDMA deployment in mid-split spectrum results in increased capacity and higher upstream speeds for D3.1 customers. It is critical to track and manage consumption and traffic from a capacity analytics standpoint to ensure that there are no capacity constraints driven by change in customer behavior and that spectrum utilization is within the normal operational range. There are additional factors like operator-initiated speed tests to check OFDMA connectivity and available speeds that could be adding additional consumption that need to be accounted for in this analysis.

In our tests, speed increase analysis was based on capturing the 98th percentile traffic and utilization for OFDMA and SC-QAM interfaces at a digital node and PPOD level for the test group (OFDMA-enabled and speed tier increases) and the control group (OFDMA-enabled and no speed tier increases) over a period of 3 months. For calculating traffic (consumption), the speed test-driven traffic was subtracted from the total per subscriber traffic. Pre and post-speed tier upgrade traffic was compared to assess the impact of the speed tier increase between test and control group.

As noted in

Figure 7. SC-QAM and OFDMA 98th utilization and traffic (Mbps). The vertical line in all panels marks the timestamp of the speed increases. The top row tracks the traffic on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines). The bottom row tracks the utilization on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines).

analysis of OFDMA utilization and traffic post-speed increase in one deployment area, there were no notable increases in both utilization and traffic for the test group as noted by the blue trend line.

Additionally, SC-QAM traffic and utilization dropped for the test group (as noted by blue line) due to customers utilizing available OFDMA capacity.

Analysis of test and control groups showed that consumption per device post-speed increase saw a marginal increase of up to ~20%, as noted in

Figure 8. Per device (CM) total consumption for the 3 groups pre and post speed increase.

. Additional speed increases within the mid-split spectrum can absorb customers' increased consumption; overall traffic (SC-QAM and OFDMA) utilizations are within the current utilization trend.

With additional capacity gains due to OFDMA enablement, speed increases are not resulting in increased utilization.




Figure 7. SC-QAM and OFDMA 98th utilization and traffic (Mbps). The vertical line in all panels marks the timestamp of the speed increases. The top row tracks the traffic on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines). The bottom row tracks the utilization on OFDMA (left) and SC-QAM channels (right) for test group (blue lines) and control groups (green and red lines).




Figure 8. Per device (CM) total consumption for the 3 groups pre and post speed increase.

5. PMA for OFDMA

5.1. Why PMA for OFDMA

In addition to needing the capacity for bandwidth growth, and for higher product speeds, OFDMA managed through a PMA solution can significantly improve customer experiences, as we have discovered with OFDM in the downstream and D3.0 US profile management. During initial OFDMA deployments several challenging network impairments and distortions limited the capacity of the OFDMA channel without a PMA solution. In some cases, the ingress sources coupling into the network were so severe that modems were challenged even to connect to OFDMA and a more robust modulation was required for interval usage code (IUC) 13, which CMs use to try and perform initial ranging on the OFDMA channel. Several examples are shown here along with some analytics of different impairments discussed later in section 6 (mid-split ingress).

Figure 9 illustrates ingress from a very high-powered VHF transmitter that was impacting a variety of nodes. As can be seen, the lower VHF channels 2 and 6 are both coupled into the upstream at very high levels through loose connectors or damaged cables. During this initial deployment, a 256-QAM flat modulation profile was tested on the OFDMA channel. At the receiver, the spectrum analyzer showed noise at a level that would have required a flat 32-QAM, or even more robust configuration, to allow the modems to range on the OFDMA channel. As a result of this impressive ingress, Comcast did some of the network analytics detailed in section 6.1 (VHF TV Ingress) to identify the probabilities of this type of ingress impacting the wider network. Delivering the higher speed services cannot be done in the presence of this type of ingress without an OFDMA PMA solution.



Figure 9: Lower VHF ingress into mid-split channel would have required 32-QAM or lower modulation without PMA solution.

Figure 10 illustrates a second example of why PMA is required to achieve Gbps symmetrical speeds. This node had loose connections or cracks in the cable near three ingress sources:

- 1) An upper VHF spectrum transmitter with channels 9 & 10;
- 2) A NOAA weather radio transmitter, that although narrow, needed to be properly managed by PMA; and
- 3) FM radio ingress between 88 & 108 MHz.

The lower channel in the high-split node is stopped at 85 MHz primarily to enable mid-split capable CMs on this node to achieve the hundreds of Mbps speeds as these CMs cannot bond to a channel above their capabilities while the HS capable modems can bond to all of the spectrum. While not intended or preferred for capacity reasons, this does avoid the FM radio band which can impact HS node performance. Technology development is happening currently to allow a mid-split modem to use an OFDMA channel that extends above 85 MHz while only being scheduled in minislots below 85 MHz with appropriate upstream channel descriptors.

If a channel is deployed across the FM spectrum, it is clear that a PMA solution would be required while remediating the network. The right chart in Figure 11 shows the PMA response to this ingress. The modulation levels of the OFDMA profile are tailored around the spectrum and shown as yellow segments. The modulation in the minislot containing the NOAA radio needed to be modulated at QPSK to avoid using error correction capabilities to resolve it. Note that most of the spectrum is modulated at 2048-QAM with the future option of 4096-QAM (when the feature is available on the vCMTS platform). Without a PMA solution to manage the capacity, this node would not have been capable of delivering 1 Gbps upstream speeds.




Figure 10: (right panel) Example high-split spectrum capture showing upper VHF and NOAA radio ingress. Mid-split channel has excellent signal quality. (left panel) PMA constructed profile for the high-split channel.

5.2. PMA Virtual Network Functions Development Starts and Ends in the Ingress Lab

In addition to the analytics of field data, a critical component of developing a PMA solution for OFDMA is the ability to recreate ingress from the field within the lab as a test vector for the PMA algorithm. In the example shown in Figure 12 the VHF ingress in Denver is measured at a Comcast lab, injected into a software defined radio (SDR) platform, and played back into the lab node. The OFDMA IUC or modulation profile design of the algorithm is then evaluated to verify that it resolved the packet loss and automatically configured on the lab vCMTS. This SDR platform is also capable of converting MER per subcarrier or per minislot into a signal that simulates the ingress in the field that would have caused that MER profile. The OFDMA is run periodically in the continuous integration and continuous deployment (CICD) environment to ensure that all the test vectors are well managed as vCMTS, CM, or PMA software is evolved.




Figure 11. Lab testing flow. VHF Ingress is captured (left panel) and injected into the node using SDR platform (top right panel). PMA adjusts to the ingress by constructing the proper modulation profile (bottom right panel).

5.3. OFDMA PMA Comes Together

The PMA base platform has been described in several previous papers [2-5]. Figure 13 describes the base platform including 3 primary components: the configuration manager which automates the instantiation of

new IUC/profile configurations; the analytics engine which recommends changes to configuration; and the data engine that collects the real time streaming OFMDA telemetry and makes it available for the PMA solution along with other proactive network maintenance (PNM) tools used to direct fix agents that remediate the network. Building on this same existing platform for OFDM has dramatically accelerated the time-to-market for the OFDMA solution.




Figure 12. PMA architecture supports DS and US D3.1 and US D3.0 profile management.

Several examples of OFDMA support the platform descriptions in the previous papers are referenced in this paper. Figure 14 shows an example of a fairly consistent MER signature for a node in which most of the spectrum could support 2048-QAM except for a small region. Each modem MER per minislot is aggregated over several days of 5-min samples with statistics such as the 10th percentile calculated as representative of the minislot-aggregated MER value. In this example, a cable modem, labeled as “Device 9” was having some attenuation issues and could not operate with this IUC configuration, so a second configuration (row 2, column 1) was configured by PMA, along with a third to cover “Device 4” that was experiencing another drop in MER (row 1, column 4).

The PMA analytics engine compiles this data for each cable modem, clusters them into similar performance groups, and then sends a set of recommended configurations to the configuration manager. The configuration manager then schedules and automatically implements the new configuration along with pre- and post-checks and change management tickets (auto-open & close) with transactional integrity.




Figure 13. Example PMA IUC/profile design across mid-split spectrum that improves the channel capacity compared to adopting a “flat modulation” IUC design.

Figure 15 shows a HS channel with a single CM from one of the first HS trials. 7 IUC profiles were constructed by PMA, but only 1 was required for the HS modem. This is the same case shown at a different time for the spectrum analysis shown in Figure 10. This is a recommendation that was updated as the level of the ingress varied over time while the channel was being dynamically managed by PMA.

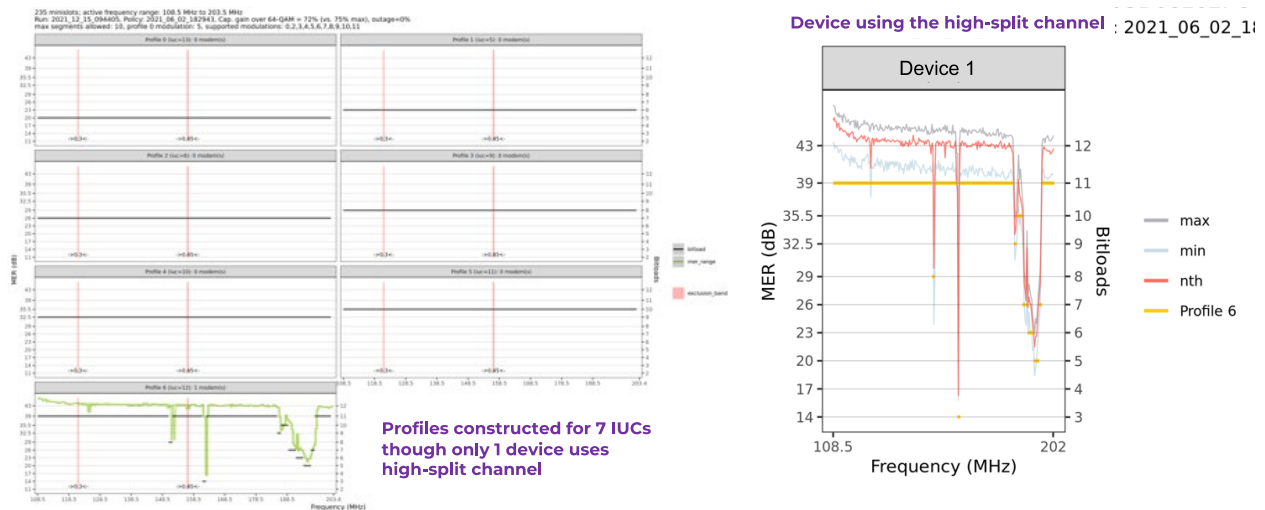


Figure 14. Example PMA design for the high-split channel in Figure 10. Left panel shows the 7 constructed profiles. Since only 1 cable modem is bonded to the channel, 6 profiles were configured with flat modulation. The right panel shows the profile tailored to the noise detected on the channel.

Figures 16 and 17 are dashboard type charts that track various metrics of interest relating to PMA performance. These include device counts, capacity, and traffic metrics both at the channel level (first row) and the IUC level (second row), as well as normalized traffic by IUC and MER statistics (last row). The first example shown in Figure 16 is for a field node with relatively clean spectrum. For this node, traffic is flowing mostly through IUC 12—the highest capacity profile at ~ 400 Mbps. MER statistics

show levels that are consistently above 41 dB and the uncorrectable codeword rate is extremely low. As a result, profile switches by PMA are minimal (indicated by vertical solid lines).




Figure 15. Example MS PMA performance charts for a clean spectrum.

In contrast, Figure 17 is for a node that exhibits variation in MER level within the 34-43 dB range. For this node, the internal PMA function responds by moving devices to suitable profiles. In this example, the dip in MER correlates with the shift of traffic from IUC 12 to lower capacity profiles. The PMA system is also responding by reconstructing the profiles in response to the dynamic nature of the noise. While the uncorrectable codeword rate is higher than the clean spectrum node, it is kept under 0.2%.

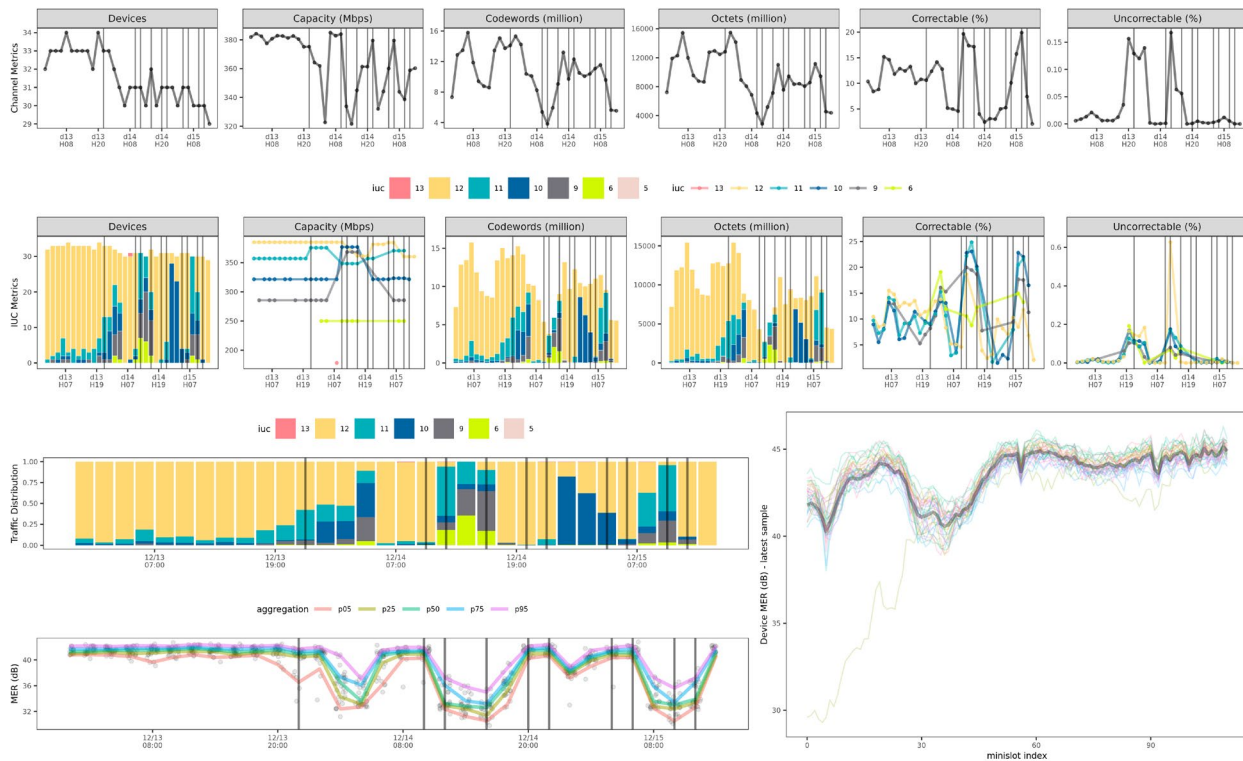


Figure 16. Example MS PMA performance charts for a noisy spectrum.

5.4. State of the Upstream Spectrum

While many of these ingress sources are problematic, the network is generally in great shape and easily supports the capacity and product speed expectations for the MS and OFDMA. Figure 18 describes the network from one of our very first OFDMA trials across 100 nodes in challenging RF environments. Even in these nodes it can be seen that:

- ~50% of network supports 4096-QAM US OFDMA based on DOCSIS PHY spec MER requirement.
- The IUC selection in the vCMTS is actually 3 dB better, implying that ~85% of OFDMA traffic can flow on 2048-QAM. This is consistent with lab testing that shows the OFDMA receiver performs at least 3 dB better than the CableLabs DOCSIS 3.1 PhHY specification requirements in the MER range of interest.
- The forecast is for ~70% of OFDMA traffic to use 4096-QAM based on this early data when that feature is available from the vCMTS.
- This performance exceeds our initial capacity models for both mid-split and high-split that were created to set goals for new product speeds.




Figure 17. (Top panel) MER cumulative distribution from early trials on ~100 nodes shows that ~85% of mini-slots can support 2k-QAM. (Bottom panel) Actual traffic stats confirm this picture.

While Figure 17 was from an initial small trial, the network performance trend is improving now that we are delivering OFDMA to hundreds of thousands of CMs with OFDMA activated. Figure 18 shows the distribution of MER per minislot over 2 days of time across tens of thousands of modems early in the writing of this paper. Several key points to note include:

- MER Performance statistics are very promising; against specification thresholds, ~90% of mini-slots show 4096-QAM speeds.
- Lab and field testing of OFDMA PMA indicates we have 3 dB more additional margin relative to specifican thresholds.
- Even at specification thresholds, this distribution represents an average of ~520 Mbps US capacity per MS SG.
- Work is ongoing to tune the OFDMA PMA, once deployment challenges are stabilized.




Figure 18. MER per minislot aggregated across tens of thousands of modems at 1-hour samples over 48 hours across varied network locations.

6. Mid-Split Ingress

6.1. VHF TV Ingress

VHF over the air (OTA) ingress is one of the more common ingress sources we discovered in our initial OFDMA deployments. As shown above this ingress can be problematic without a PMA solution. To understand how it might affect capacity, Comcast did an analysis of power levels expected at homes and within nodes. The analysis was based on a Federal Communications Commission (FCC) spectrum database for VHF transmitters and a propagation model informed by measuring actual ingress using our full band capture function. The bands evaluated are shown in Figure 19.

Channel Number	Frequency in MHz
Channel-2	54-60
Channel-3	60-66
Channel-4	66-72
Channel-5	76-82
Channel-6	82-88
FM Broadcast	88-108



Figure 19. One primary source of ingress which impacts mid-split is the lower VHF OTA television broadcast. The channels of interest are 2-6 and FM broadcast

A sample of VHF transmitters was selected and the overlapping channel power for devices at various distances from these transmitters was measured. Figure 21 shows the variation in channel power as a function of distance from the VHF transmitters. Key points to note include:

- Impact reduces as the distance from the VHF transmitter increases and it appears to follow a power trend line.
- The impact appears to vary across device types. A potential explanation for this variation may be due to shielding improvements in newer devices.
- Device Type 3 filters channels below 108 MHz and hence is minimally impacted.




Figure 20. Variation in channel power vs distance for VHF transmitters for different device categories

6.2. Interesting Impairments to Keep in Mind for PMA

The impairment in Figure 22 was seen in several nodes and diagnosed to see if it was distortions from DS signals. After investigating we believe these narrow-band ingressors are an old VCR, video game (i.e., Atari) or the wrong connector on an older set-top-box connected to an outlet in the home. This is the signature and power levels that would be generated from a channel 3 amplitude modulation (AM) modulator at 61.25 MHz.



Figure 21. Example video modulator in home transmitting into OFDMA spectrum

While stabilizing the OFDMA technology deployments we have observed several cases of high tilt compensation across the OFDMA channel. Some of these responses affected MER of the OFDMA channel resulting in some compensation by the PMA solution. Examples of these are shown in Figure 23. Some of these were due to network devices that were not replaced in the mid-split upgrade process, such as inline equalizers in the lower half of Figure 23 that were tuned for a low-split network. Other cases were from equalization optimization issues requiring a periodic unequalized probe to reset the equalizer and interaction with upstream transmit powers. Other cases were a result of incorrect use of upstream conditioning in new mid-split amplifiers.

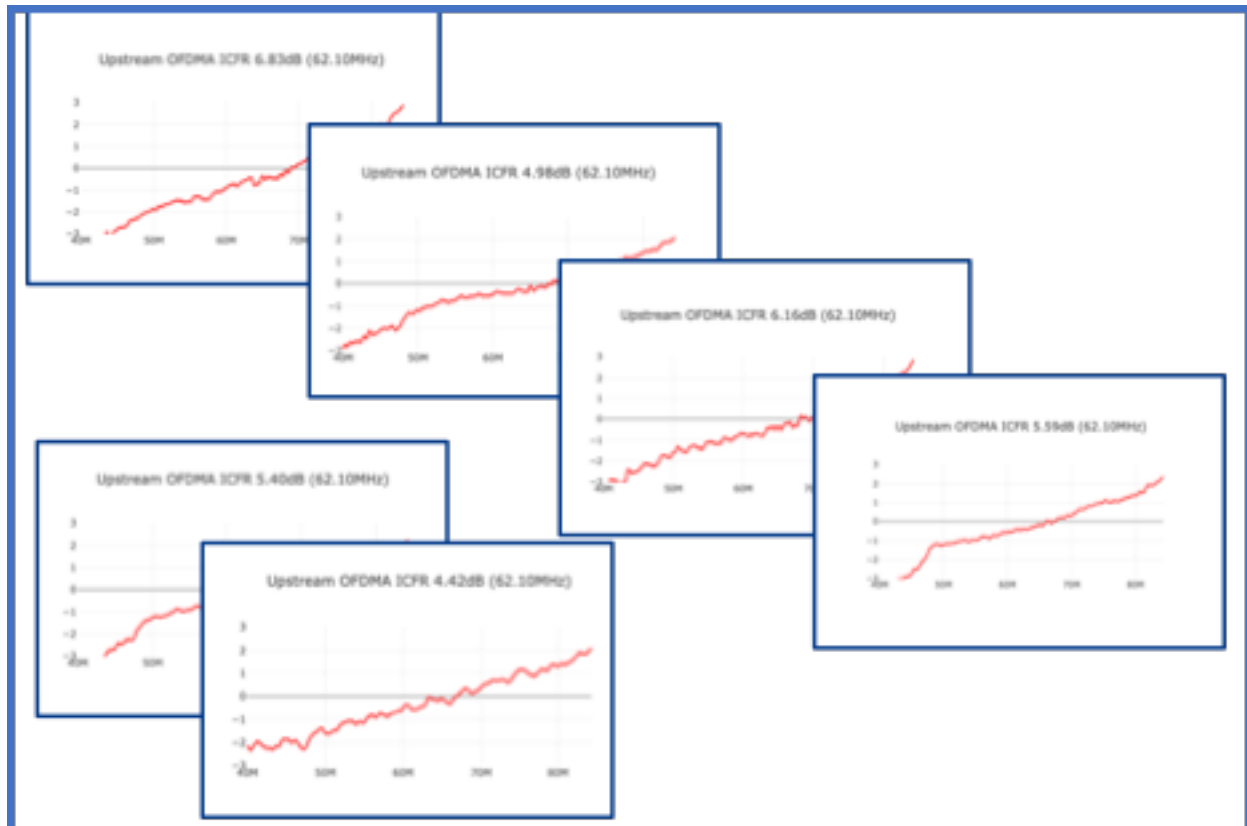


Figure 22. Various Pre-EQ responses for OFDMA channel showing tilt.

7. Conclusion

The paper details some insights into the challenges of deploying OFDMA technology. As part of the deployments and trials over the last 18 months we have shown the benefits and success of the OFDMA technology for delivering hundreds of Mbps of US speeds in MS networks and 1 Gbps US speeds in high-split networks with a positive customer experience. Comcast has determined that an OFDMA PMA solution is essential in delivering these new higher speed products. The PMA solution is essential because of the ingress and distortion challenges we have identified in a very small percentage of nodes that changes over time. Primary ingress sources include over-the-air broadcast transmission. These ingress sources are easily handled by a dynamic modulation profile management application that is an incremental evolution to the platform.

Abbreviations

AM	amplitude modulation
CM	cable modem
CMTS	cable modem termination system
D3.0	Data Over Cable Service Interface Specification 3.0
D3.1	Data Over Cable Service Interface Specification 3.1

DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
FCC	federal communications commission
FDX	full duplex
FM	frequency modulation
HFC	hybrid fiber coaxial
HS	high-split
IUC	interval usage code
MER	modulation error ratio
MS	mid-split
NOAA	National Oceanic and Atmospheric Administration
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OTA	over the air
PMA	profile management application
PNM	proactive network maintenance
PPOD	physical pod
RF	radio frequency
SCTE	The Society of Cable Telecommunications Engineers
vCMTS	virtual cable modem termination system
US	upstream
UVHF	Upper VHF band, Low VHF can be used to refer to channels 2-6, Upper VHF for channels above channel 6
VHF	very high frequency

Bibliography & References

1. “A Machine Learning Pipeline for D31 Profile Management”, Maher Harb, Jude Ferreira, Dan Rice, Bryan Santangelo, and Rich Spanbauer, NCTA Technical Paper, 2019.
2. “Full Scale Deployment of PMA”, Maher Harb, Jude Ferreira, Dan Rice, Bryan Santangelo, NCTA Technical Paper, 2020.
3. “Optimizing DOCSIS 3.0 Configuration in the Upstream through Applied Reinforcement Learning”, Kevin Dugan, Maher Harb, Dan Rice, Rob Lund, NCTA Technical Paper, 2021.
4. “PMA Improvements - Strategies employed for faster mitigation, increased capacity, and cost savings”, Leech et al, NCTA Technical Paper, 2022.
5. “Time is Ripe for D30 Farming: Achieving Optimal Spectral Efficiency by Allocating D30 Spectrum to OFDM”, Maher Harb, Chad Humble, Sebnem Ozer, Dan Rice, NCTA Technical Paper, 2022.
6. “Rapid and Automated Production Scale Activation of Expanded Upstream Bandwidth”, Rob Thompson, Rob Howald, John Chrostowski, Dan Rice, Amarildo Vieira, Rohini Vugumudi, Zhen Lu, NCTA Technical Paper, 2021.
7. “Bringing the Mid-Split Factory Online to Rapidly Produce Terabytes”, Howald et al, NCTA Technical Paper, 2022.

Design and Implementation of a Controls Framework to Secure a 10G Network

A Technical Paper prepared for SCTE by

Mike Gala

Executive Director
Comcast

1701 John F. Kennedy Blvd, Philadelphia, PA 19103
(215) 286-8937
mulchand_gala@comcast.com

Andrew Yun

Director
Comcast

1701 John F. Kennedy Blvd, Philadelphia, PA 19103
(215) 605-0722
andrew_yun@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Methodology.....	3
3. Control Design.....	6
3.1. Identify Risks.....	6
3.2. Identify Framework and Controls.....	8
3.3. Principles for Control Monitoring.....	9
3.4. 3 Lines of Defense Risk Management.....	10
4. Measurement and Reporting.....	11
4.1. Defining Control Metric and Thresholds.....	11
4.2. Identifying sources of reliable and scalable data.....	12
4.3. 1 st line and 2 nd line of defense dashboards.....	12
4.4. Accountability Summary and Responsibility Views.....	13
5. Onboarding and Continuous Monitoring.....	13
5.1. Knowledgebase.....	14
5.2. Control Trial Implementation.....	14
5.3. Stakeholder Alignment.....	16
5.4. Compliance Delegate Identification.....	17
5.5. Continuous Monitoring.....	17
5.6. Optional Attestation.....	18
6. Conclusion.....	18
Abbreviations.....	18
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 – Common Cybersecurity Control Frameworks.....	4
Figure 2 – DOCSIS 4.0 Distributed CMTS Reference Architecture [5].....	7
Figure 3 – Organizational Data Mapping to Controls [6].....	9
Figure 4 – Principles for Control Implementation and Monitoring.....	10
Figure 5 – 3 Lines of Defense.....	11
Figure 6 – Identification of Reliable Organizational Data and Insights Utilization.....	12
Figure 7 – Business Unit Onboarding Approach.....	14
Figure 8 – Trial Goals for Onboarding Readiness.....	15
Figure 9 10– Control Lifecycle.....	16
Figure 11 – Shift to Continuous Control Model.....	17

List of Tables

Title	Page Number
Table 1 – Example of NIST 800-53 Controls for a 10G Network.....	5

1. Introduction

Network security has been one of the top priorities for the industry since the transition from analog to digital. It is a combination of prevention, mitigation, remediation, and customer notification technologies that help reduce the risk of data loss, theft, and sabotage. [1] The landscape has also changed drastically recently due to the pandemic, geo-political instability, and evolving federal and state regulatory requirements and standards. Network security threats are constantly evolving as well with threat actors looking to exploit vulnerabilities to gain the initial access required to laterally move across the network, and ultimately exfiltrate sensitive information or impact the environment. The recent CISA alert AA22-158A [2] of state-sponsored cyber actors exploits of network providers and devices is an example of such threats potentially targeting modern networks. These actors exploited known vulnerabilities, primarily common vulnerabilities, and exposures (CVEs) associated with network devices to target and compromise major telecommunications companies and network service providers since 2020.

Often the response to widely impacting cyber threats is to deploy additional security safeguards across the network and its devices; initially at the perimeter and down through its many additional security layers (network, endpoint, application, and data) as well as externally to the subscriber premise devices. Patching and end-of-life infrastructure replacement can also be conducted, but these are usually reactive measures to exploits that are already utilized and potentially costly to remediate.

A security controls framework approach facilitates the proactive measurement of existing and implemented protections required to minimize risks across an organization's people, processes, and technologies safeguarding a 10G network. Network security threats such as DDoS attacks, fraud and phishing, and data breaches are top priorities for any organization requiring the ability to continuously measure control effectiveness and risks with an automated and scalable solution. The design and implementation of continuous monitoring for security controls is often a complex and challenging task requiring subject matter expertise, alignment across multiple organizations, and data aggregation and correlation to enable effective risk management. This whitepaper will elaborate on a unique strategy that showcases the best practices for design, implementation, and operation of a controls framework for a 10G network aligned with a risk management approach.

2. Methodology

To establish the continuous measurement and monitoring of security controls requires the implementation of a controls framework that is adaptable and customizable to the technologies, processes, policies and standards, and expertise of each individual organization. Many security frameworks exist where an assessment is required to determine the best fit in alignment to organizational cybersecurity goals and its risk management strategy.

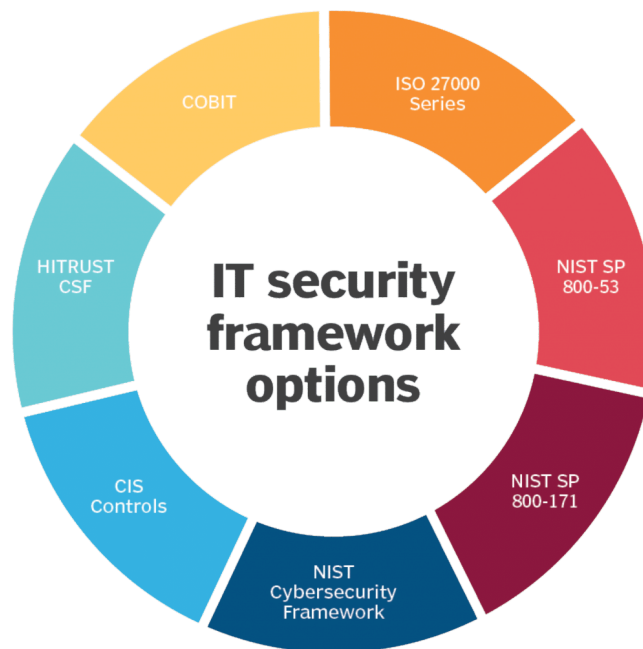


Figure 1 – Common Cybersecurity Control Frameworks

The cybersecurity framework utilized in this whitepaper to illustrate the design and implementation of continuous controls monitoring is NIST SP 800-53 [3]. The National Institute of Standards and Technology (NIST) is a non-regulatory agency of the United States Department of Commerce, where the NIST Cybersecurity framework was introduced in 2014 with revisions released to domains, controls, and enhancements (currently in Revision 5). The NIST Cybersecurity Framework is centered around five essential domains:

- **Identify (ID):** Determines an organization’s critical functions and the risks that could disrupt those functions. Identifying organizational systems and assets, information types, policies and procedures, risk management strategy, and roles and responsibilities are essential elements of this domain.
- **Protect (PR):** Defines the relevant safeguards required to deliver critical protective services by establishing priorities and cybersecurity efforts. Protection of organizational systems, assets, and information via access controls, awareness and training, data security, information protection processes and procedures, maintenance, and protective technology management are essential elements of this domain.
- **Detect (DE):** Implementation of an organization’s detective capabilities to promptly identify cyber risks and incidents. Event logging, anomalous activity detection, security continuous monitoring, and detection processes are essential elements of this domain.
- **Respond (RS):** Execution of measures related to a detected cybersecurity incident and an organization’s ability to manage its impact. Response planning, communication strategy, analysis procedures, mitigation, and application of improvements are essential elements for this domain.
- **Recover (RC):** Restorative capabilities to impacted systems and services as the consequence of a cybersecurity incident. Recovery processes and procedures, application of improvements, and cross organization collaboration are essential elements of this domain.

Taking into consideration an organization's goals, policies and standards, technologies, and processes are key in the successful implementation of a control monitoring framework. An organization's goals define its objectives, policies and standards provide governance and management, technologies illustrate its capabilities, and processes allow enablement. The domains in NIST 800-53 serve as the guidelines to be implemented concurrently to form an operational culture that addresses cyber risk. Within the domains are controls that provides the protective measures for systems, organizations, and individuals. The controls allow specification and customization to an organization's objectives allowing enablement of a continuous controls monitoring culture.

Table 1 illustrates how a series of controls can be customized to the specificity of a 10G network.

Table 1 – Example of NIST 800-53 Controls for a 10G Network

ID	10G Network Specific Controls
ID.AM	Network Inventory Management in defined source system, automated discovery, completeness of key attributes like Serial Numbers, Location for all key assets that encompass the 10G Network.
ID.BE	Alignment of the organization's mission, objectives, stakeholders, and activities with cybersecurity roles, responsibilities, and risk management of the 10G Network.
ID.GV	The policies, procedures, and processes alignment with regulatory, legal, risk, environmental, and operational requirements for the 10G Network.
ID.RA	Understanding of the 10G Network related cybersecurity risk to operations (including mission, functions, image, or reputation), organizational assets, and individuals.
ID.RM	The organization's priorities, constraints, risk tolerances, and assumptions are established and used to support operational risk decisions as they relate to the 10G Network.
PR.AC	Identity Management, Authentication and Access Control for physical and logical assets and associated facilities for authorized users managing and monitoring the 10G Network.
PR.AT	Providing personnel and partners the cybersecurity awareness education to perform their cybersecurity-related duties and responsibilities for the 10G Network.
PR.DS	Protecting Information and records (data) in transit and at rest in 10G Network against data leaks separating production and test environments.
PR.IP	Processes are in place for Software Development Lifecycle (SDLC), secure configuration, backups, vulnerability management and incident response for the 10G Network.
PR.MA	Maintenance and repairs of industrial control and information system components of the 10G Network are performed consistent with policies and procedures.
PR.PT	Audit Logging, media protection, control plane protection, load balancing are in place to ensure security and resiliency of the 10G Network.
DE.AE	Detecting Anomalous activity against a baseline and creating incidents for necessary events by correlating from multiple sources and sensors within the 10G Network.
DE.CM	Continuous monitoring and vulnerability scanning to identify cybersecurity events, malicious code, unauthorized access etc.
DE.DP	Detection processes and procedures are maintained and tested to ensure awareness of anomalous network events.
RS.RP	Response processes and procedures are executed and maintained, to ensure response to detected cybersecurity incidents.

RS.CO	Response activities are coordinated with internal and external stakeholders (e.g., external support from law enforcement agencies, vendors, etc.).
RS.AN	Analysis and forensics are conducted to ensure effective response and support the network recovery activities.
RS.MI	Activities are performed to prevent expansion of an event, mitigate its effects, and resolve the network incident.
RS.IM	Organizational response activities are improved by incorporating lessons learned from current and previous detection/response activities.
RC.RP	Recovery processes and procedures are executed and maintained for restoration of 10G network sub-systems or assets affected by cybersecurity incidents.
RC.IM	Recovery planning and processes are improved by incorporating lessons learned into future activities.
RC.CO	Restoration activities are coordinated with coordinating centers, Internet Service Providers, owners of attacking systems, victims, and vendors).

The rest of the whitepaper will discuss the process of designing the controls mentioned above, implementing related KPIs that can be matured over time, onboarding various teams running the parts of the 10G network, and setting up continuous monitoring of these controls aligned with a risk management approach.

3. Control Design

The number of devices connected to a subscribers' network continues to increase across variety of industries (e.g., healthcare, education, private and public sectors, etc.) requiring connectivity to reliable network services. Speed, capacity, reduced latency, enhanced reliable and security are all advancements that make 10G transformational in how people and industries operate in a multi-gigabit reality. "The 10G platform is a combination of technologies that will deliver symmetric multi-gigabit Internet speeds with a vision toward enabling symmetric 10 gigabits per second (Gbps) services. 10G will be significantly faster than what most consumers currently experience, and will offer lower latencies, enhanced security, and greater reliability" [4].

These advancements are examples of organizational 10G goals that are foundational to control design and the successful implementation of a security controls framework. The definition of these business objectives enables an organization to establish its scope of controls, prioritization of related security efforts, organizational adoption, and the measurability of control performance. With the organization objectives defined, an iterative approach to control design can be applied to establish a framework that is aligned to organizational goals, policies and standards, technologies, and processes.

3.1. Identify Risks

Development of the organization's holistic view of critical business functions starts with the identification of its assets, risks, policies, and owners. Identification of these aspects determines an organization's critical functions and the risks that could cause disruption to key network components or infrastructure. Identifying organizational systems and assets, information types, policies and procedures, risk management strategy, and roles and responsibilities are essential elements for a risk managed network.

The path to 10G and its network transformation to a distributed access architecture has required the controls framework to further adapt in the identification of risks in a virtualized and digital environment. From the virtualization of the cable modem termination system (CMTS) to the delivery of broadband signals across the network to the subscriber premise, the traditional approach to protect physical assets and key infrastructure has now expanded to the virtual environment.

Figure 2 illustrates an example of a virtualized CMTS network. [5] (Source: Cable Labs Data-Over-Cable Service Interface Specifications for DOCSIS[®] 4.0, MAC and Upper Layer Protocols Interface Specification CM-SP-MULPIv4.0-105-220328)

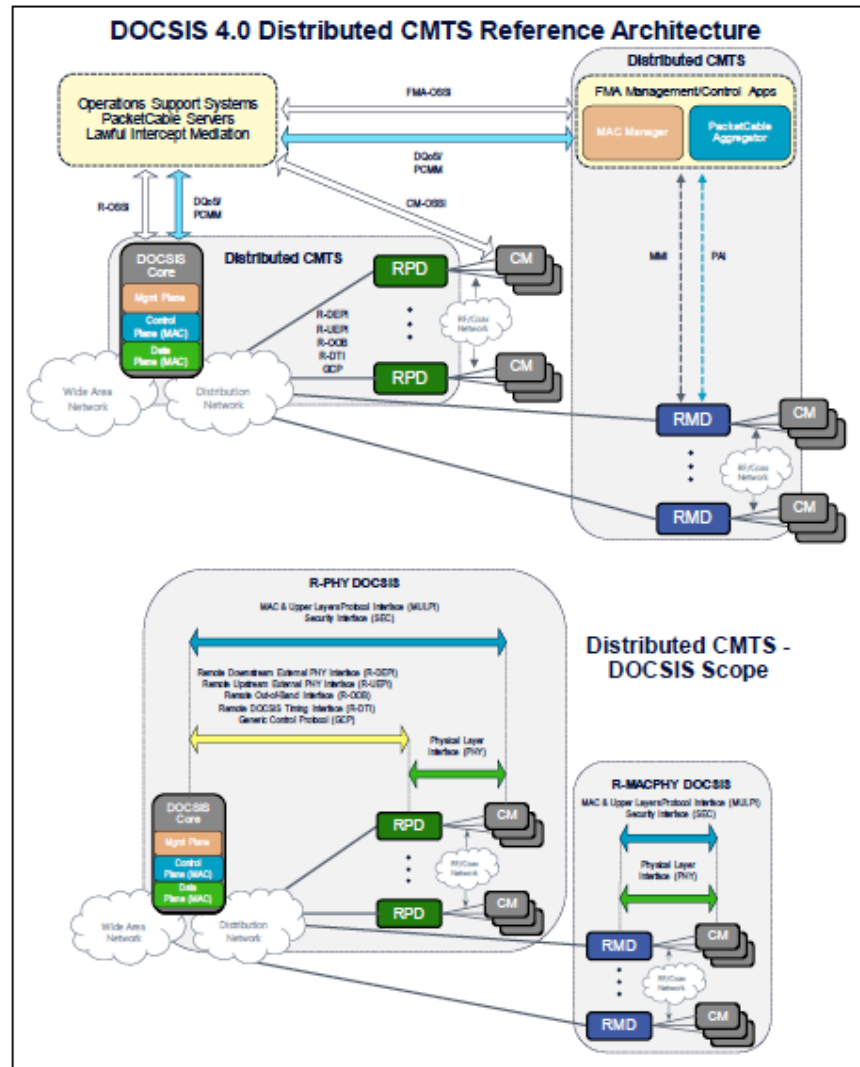


Figure 2 – DOCSIS 4.0 Distributed CMTS Reference Architecture [5]

Therefore, controls such as asset identification should be expanded to ensure coverage of the virtual environment, along with the design and implementation of policies and standards relevant to the security of a 10G network to ensure broader risk coverage. Similarly, with further reliance on automation to manage distributed and virtualized environments, traditional identity and access management risks should

be expanded to ensure appropriate coverage. Risks related to change management are also important to consider ensuring various teams involved in building and operating the 10G network are appropriately reviewing and approving deployments. Version control tools are appropriately configured and managed in alignment to organizational standards, along with the secure storage of secrets in organizational vaulting solutions and not present in source code.

Risk assessment and third-party management are additional critical aspects of risk identification in a 10G network. These areas allow a 10G network provider to understand the cybersecurity risks to organizational operations, priorities, and risk tolerance.

- **Risk Assessment:** Asset vulnerabilities to core 10G infrastructure are identified and documented; Cyber threat intelligence is received potentially impacting network providers (e.g., implementation of a Threat Intelligence Platform); Potential business impacts and likelihoods are identified; Threats, vulnerabilities, likelihoods, and impacts are used to determine risk; Risk responses are identified and prioritized
- **Third Party Management:** Cyber third party risk management processes are identified, established, assessed, managed, and agreed to by organizational stakeholders and network management teams; Suppliers and third party partners of information systems, components, and services are identified, prioritized, and assessed; Contracts with suppliers and third party partners are used to implement appropriate measures designed to meet the objectives of an organization's cybersecurity program; suppliers and third party partners are routinely audited to confirm they are meeting contractual obligations

3.2. Identify Framework and Controls

Identified internal and external risks require the implementation of controls within a framework that aligns best to organizational objectives. Frameworks such as NIST 800-53, NIST CSF, ISO 27001, etc. provide specific scope and functions where the selection may be based on various frameworks that align to the organization's goals and regulatory requirements (e.g., GDPR, CCPA, PCI DSS, etc.). Often, the legal team within the organization provides consultation on the framework selection to ensure appropriate coverage is provided to identified risks and specific regulatory requirements. Once the approach to framework selection is completed, the process to control identification can begin mapped to the identified risks.

Below details the best practices for control selection, design, and implementation:

- Mapping of controls to policies and standards, assets, risks, owners, and organizational data allows the ability to successfully design a control for implementation and measurement
- The management of controls is best accomplished utilizing domains (e.g., asset management, risk management, identity & access management, etc.) with mapping to controls for an organized implementation which avoids duplication of controls
- Each control should be paired with a specific process owner that has defined the scope of the control, designed and implemented the control, continuously supports the control and its system and users, and can accurately report on control performance
- Control design logic and requirements should be formally documented and continuously maintained with updates and changes. Additionally, a control maturity model should be considered and formally documented as the selection and implementation of a control may require multiple phases based on enhancements by the process owners

Another key aspect in the selection, design, and implementation of a control is the organization's data. The availability of control data allows the ability to measure control performance, identify risks, and enables the business to establish a continuously control monitoring culture. A data strategy should be designed based on organizational capabilities to ensure data ingestion, quality, transformation, analytics, sharing, and governance is appropriately and efficiently conducted.

Figure 3 illustrates both the upstream and downstream control mapping based on organizational data.

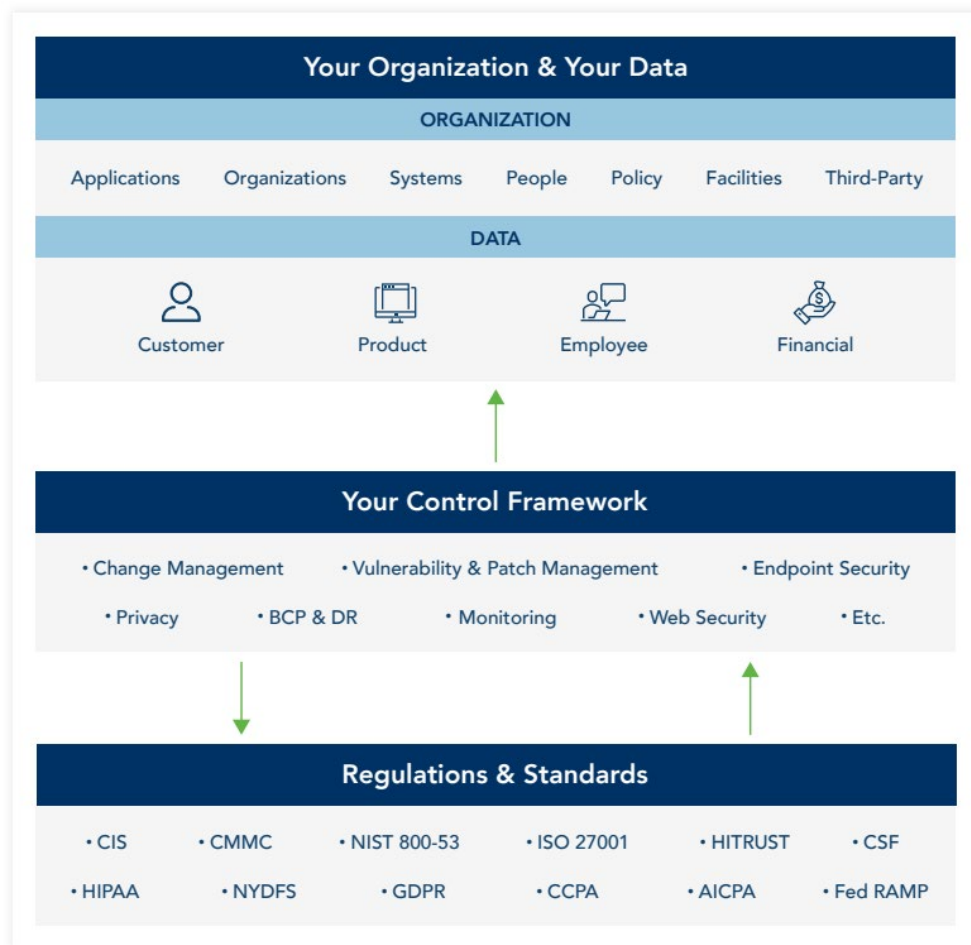


Figure 3 – Organizational Data Mapping to Controls [6]

3.3. Principles for Control Monitoring

Establishing principles for control monitoring is essential to ensure the implementation of controls meets the objectives for a 10G network. Following a set of defined and agreed upon principles allow the implementation, adoption, and support of the controls to be sustainable with roles and responsibilities clearly defined in the controls end-to-end process.

Figure 4 defines a set of principles that can be utilized to meet an organization's control monitoring objectives and requirements.

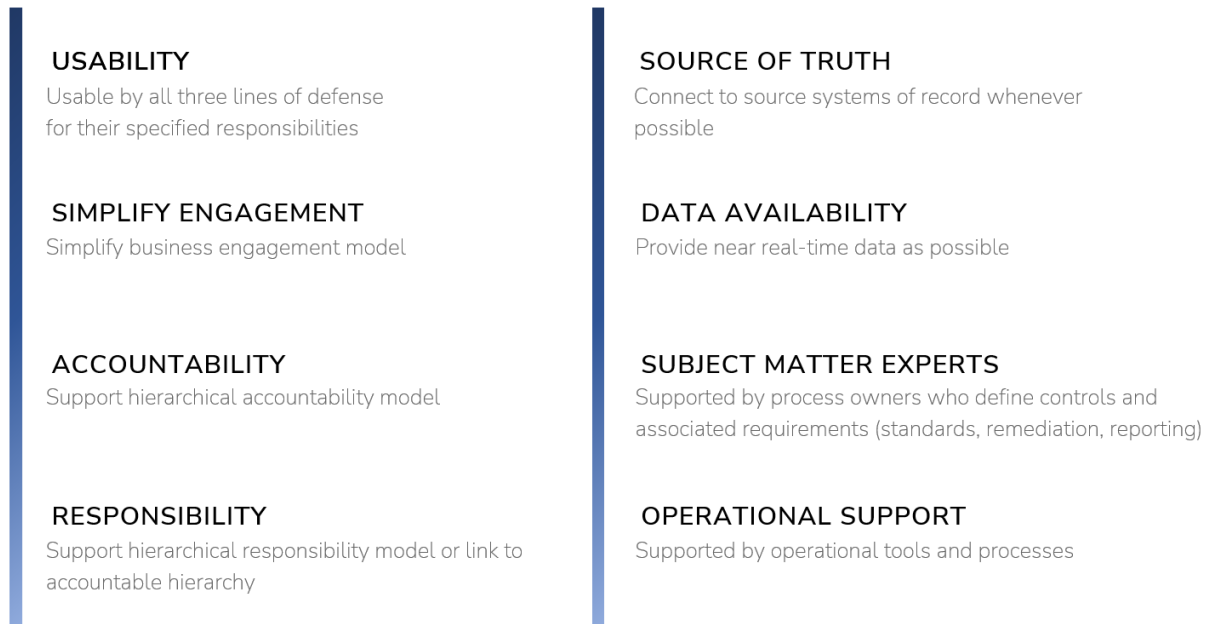


Figure 4 – Principles for Control Implementation and Monitoring

3.4. 3 Lines of Defense Risk Management

The concept of the 3 lines of defense in risk management is a term often utilized in the cybersecurity compliance and audit sector. It is essentially a model that provides guidance for effective risk management and governance for an organization with distinct roles and responsibilities across all 3 lines.

Figure 5 illustrates and details the 3 lines of defense and their responsibilities across an organization.



Figure 5 – 3 Lines of Defense

4. Measurement and Reporting

The representation of the control performance and analytics is key in the organizational adoption of the controls framework. The development of the control performance measurement is often accomplished utilizing a visualization tool (e.g., Tableau, Power BI, Splunk, etc.) that best fits organizational capabilities and skills. Prior to development, steps should be taken to gather and define development requirements which includes the determination of control metrics and thresholds, data source identification, and 1st and 2nd line defense views that demonstrate risk and control performance progress.

4.1. Defining Control Metric and Thresholds

Once the controls are defined, attention can be focused on the control metrics. Below are a few principles to consider in building the insights necessary for successful utilization:

- Understand your audience, identify their analytical capabilities and what best complements their visualization needs
- Define the appropriate control numerator and the denominator that provides the risk coverage and measures progress accurately
- Take a risk driven approach to prioritization of control measurements as opposed to easily available metrics in the data set and keep in mind the total risk posture (total attack or risk surface)
- Business Unit actions vs. Process Owner actions to maintain the accountability and responsibility for remediations
- Time-bound activities and their proper representation in the metric to measure SLAs and quarterly breakouts. Also consider the re-certification of a metric after an elapsed time
- “Keeping it simple but effective”. This adage still serves well for these metrics that can get very granular and detailed and could take a life of their own
- Ensuring that thresholds are defined in consultation with internal risk management teams and legal teams so internal and external compliance requirements are met

4.2. Identifying sources of reliable and scalable data

Following the control metrics definition, the next step is to look at various systems that source the data and build the data pipelines to integrate into the control monitoring dashboard. A few aspects to consider in this step:

- Secure interfaces and methods to integrate the data at the source
- Determine the frequency of the updates and apply automation where possible
- Ensure quality and accuracy of the data and its useability in the control metrics build out
- Deidentify the data to remove any Personal Information (PI)

Figure 6 illustrates a general inventory of an organization's data required to build the controls metrics dashboard and the lines of defense that utilizes its insights to perform risk identification and mitigation.

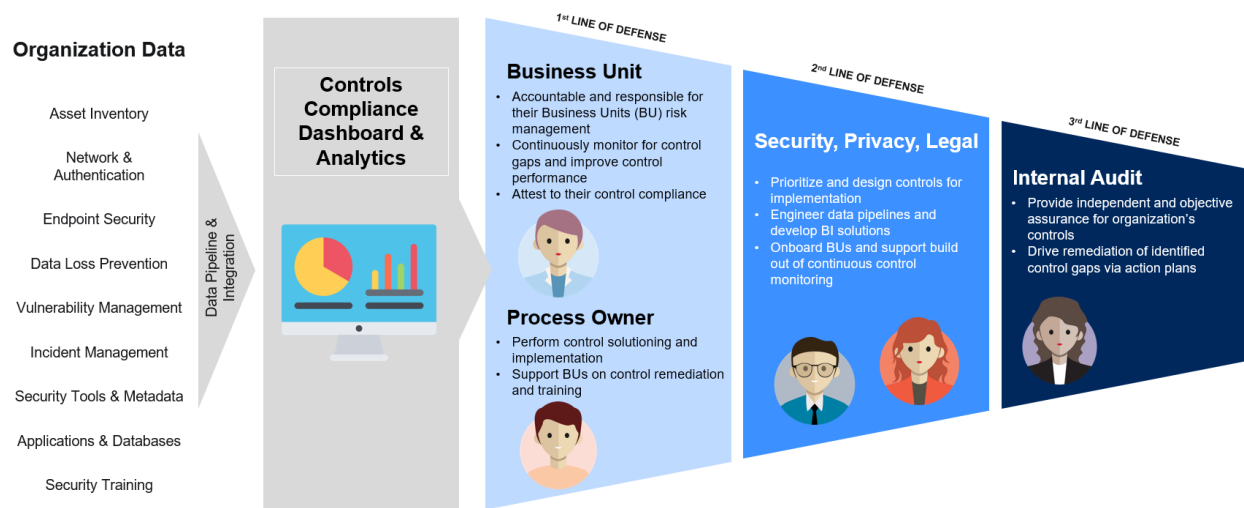


Figure 6 – Identification of Reliable Organizational Data and Insights Utilization

4.3. 1st line and 2nd line of defense dashboards

In line with the risk management framework, it is advisable to create dashboards that serve specific purposes to highlight risk for its intended use. The 1st line of defense dashboards provides the ability to highlight risk for the Business Unit to focus on control performance, gap remediation, and enablement to continuously monitor control performance. Ease of use, data accuracy, and alignment to the analytical capabilities of its audience is key in its adoption to drive enablement for the Business Unit.

The 2nd line of defense dashboards provides a risk view of the network across the entire organization with the ability to focus in on specific control risk areas requiring intervention or support. Also has the capabilities of highlighting the risk management maturity of a Business Unit with the added ability to assess progress via trending and identify additional risk areas. The 3rd line of defense is also able to utilize this view to support internal audits with the benefit of reducing the manual efforts associated to

identifying risks, driving remediation via action plans, and the validation required to confirm risk mitigation.

4.4. Accountability Summary and Responsibility Views

As the controls framework and its controls contain multiple and at times complex metrics, it is recommended to provide separate views that are related to accountability (the executive view) and responsibility (the operational view). The accountability view provides an executive the visibility required to understand control compliance at the organization level, identify escalation needs when control performance is lacking or use for attestation purposes (see section 5.6 for more details on attestations). Adding trending expands the control accountability to

The responsibility view provides the operational visibility required for the remediator to identify their control gaps and perform the corrections required to reduce risk. The responsibility view should provide the details necessary to perform the remediation; data points such as control gaps, insights on risk factors supporting prioritization, distribution by owner to forecast resource demand, and drilldown details are examples of the insights necessary to quickly and easily remediate the control gaps identified.

One of the challenges often experienced is the re-factoring of the dashboards due to organizational changes which may impact the hierarchical ownership of assigned controls and associated risks for the accountability and responsibility views. Associating the business unit roll up to a control metric that can be linked to immutable data objects such as applications and vendors provides a way to keep the dashboards agile and “self-correct” when these changes occur.

5. Onboarding and Continuous Monitoring

The onboarding of users onto the security controls framework and related dashboards is key in the organization’s overall risk management strategy. The goal in onboarding of users is to ensure a continuous control monitoring process and culture can be implemented, which improves the organization’s security posture and leads to a sustainable model for gap remediation. Sustainability is also central to onboarding as teams are often challenged with priorities and resource capacity, where consultation may be required to build in continuous control monitoring processes into existing workflows.

Onboarding onto the security controls framework can be methodically accomplished in defined phases to ensure those onboarded are equipped with the knowledge necessary to take ownership of their responsibilities across all lines of defense.

Figure 7 illustrates the onboarding approach to ensure alignment across key stakeholders and application of corrections or improvements based on feedback for Business Unit (BU) onboarding:

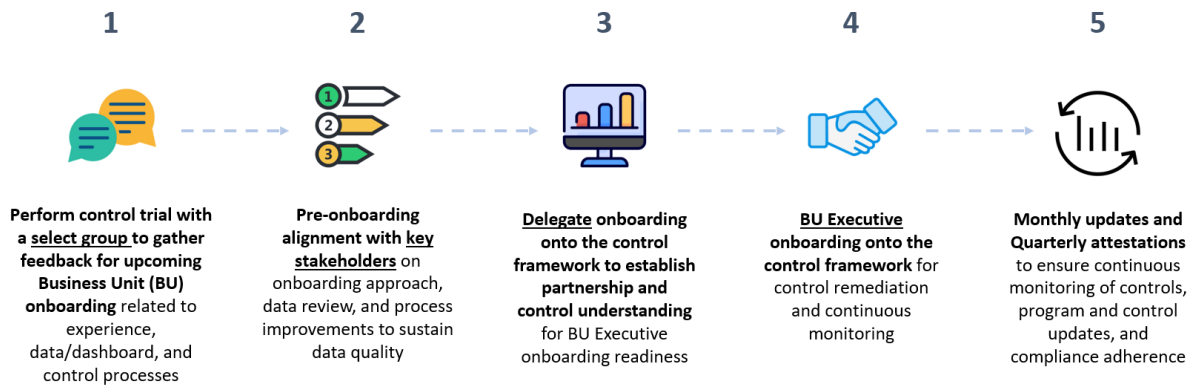


Figure 7 – Business Unit Onboarding Approach

5.1. Knowledgebase

Building a knowledgebase is key in the successful onboarding of users onto an organization's security controls framework. Depending on the size, implemented tools, and structure of the organization; control requirements and related processes may be complex. The ability for a user to understand a control, their responsibilities, and processes related to gap remediation allows the enablement necessary where accountability is taken in alignment to organizational cybersecurity goals and objectives.

A few guidelines can be followed in the development of an enablement focused knowledgebase:

- Accessibility should be considered to ensure it is usable by all users, but secured to ensure access is restricted to only those required in the organization
- Documentation is maintained periodically and expanded over time based on new control implementation, enhancements to existing controls or processes, and repeatedly asked questions
- Management of knowledgebase is simplified to ensure sustainability (e.g., links to policies and standards that can change over time)
- Control responsibilities and helpful tips for remediation are made easily available as it will drive improved adoption and understanding
- Avoid swivel where possible; if dashboards are developed measuring control performance, provide control details and remediation tips directly with the control performance metrics for users to easily address identified gaps

5.2. Control Trial Implementation

Similar to the market trials of a 10G network rollout, control implementation should be trialed with a small group of early adopters that can provide feedback on control effectiveness and design. As controls will evolve and mature over time, feedback is key in ensuring the effectiveness of controls and new features can be tested in alignment to its design and implementation. The trial goals should also be defined for the early adopters and can be targeted to specific goals during the testing period.

Figure 8 illustrates the trial control goals in preparations for a production implementation:

Control Trial Goals for Onboarding Readiness




Figure 8 – Trial Goals for Onboarding Readiness

An added benefit of implementing a trial model for future control implementation allows the ability to establish a release strategy following a control development lifecycle from start to end, while allowing the early adopters to address their control gaps early in the control implementation process.

Figure 9 illustrates the control lifecycle for new control implementation and existing control enhancements from intake to production release:

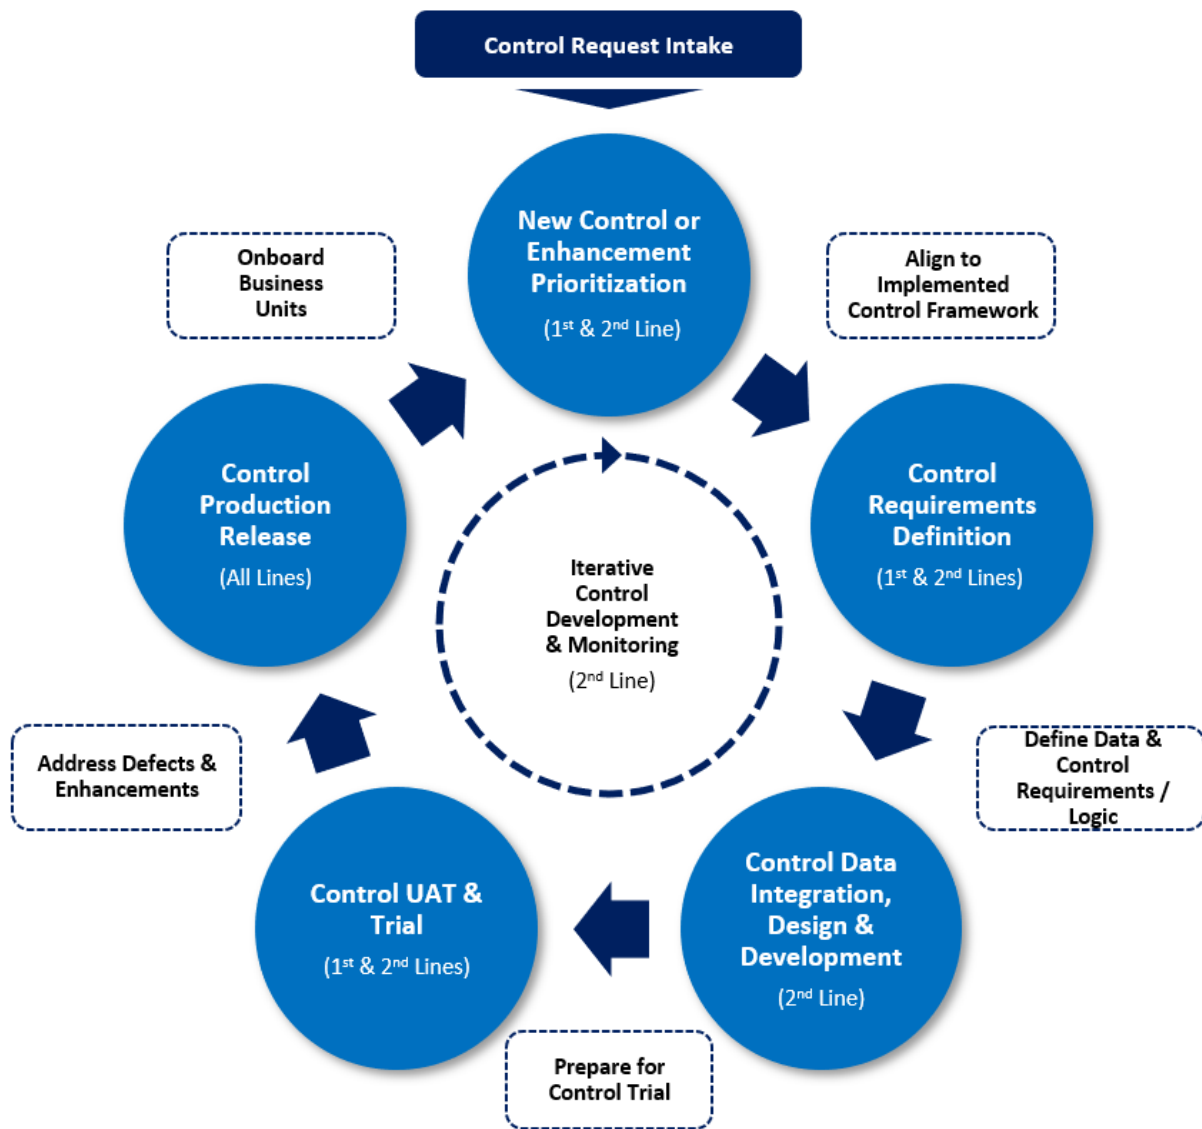


Figure 9 10– Control Lifecycle

5.3. Stakeholder Alignment

A key dependency in the success of a security controls framework implementation is the alignment established between key stakeholders. A 10G network implementation will require partnerships across many teams across engineering, field operations, to customer services. Similarly, a security controls framework requires partnerships with cybersecurity, privacy, and engineering counterparts to ensure the implementation of controls is in alignment to organizational goals and objectives. An example would be the partnership between the network engineering and cybersecurity teams in the establishment of access event logs to critical 10G network components. With access management policies and standards in place, the ability to measure control compliance via access event logs would be critical to proactively identify risks and gaps requiring remediation from both an engineering and process perspective.

5.4. Compliance Delegate Identification

Champions for a security controls framework implementation are instrumental in ensuring a continuous control monitoring culture can be adopted and implemented across the organization. These champions operate as delegates on behalf of the cybersecurity team, provide guidance on the implementation strategy for their respective teams, provide subject matter expertise to their teams related to control processes, and influence the change necessary to implement a continuous control monitoring mindset. As the delegate function in most situations are not always a part of their defined role, it is important to ensure that professional and personal growth related to this type of role is transparent for their active participation and partnership. This type of role often leads to increased visibility across the organization, with key leadership groups, and the ability to expand their knowledge as a subject matter expert in security.

5.5. Continuous Monitoring

The goal in the implementation of an effective security controls framework within an organization is to establish a continuous control monitoring process. This requires the ability to shift the cultural mindset from a traditional burndown approach to control gap remediation, to one that is proactive and continuous. Success should be measured not only in the control performance, but also with an organization's ability to shift security and privacy to the left early on in their internally processes. Only then will the security controls framework become sustainable, built in culturally across the organization, and a mechanism to ensure an initiative such as the design and implementation of a 10G network can be appropriately secured and risks mitigated.

Figure 10 illustrates the comparison between the reactive and proactive approach to control monitoring.

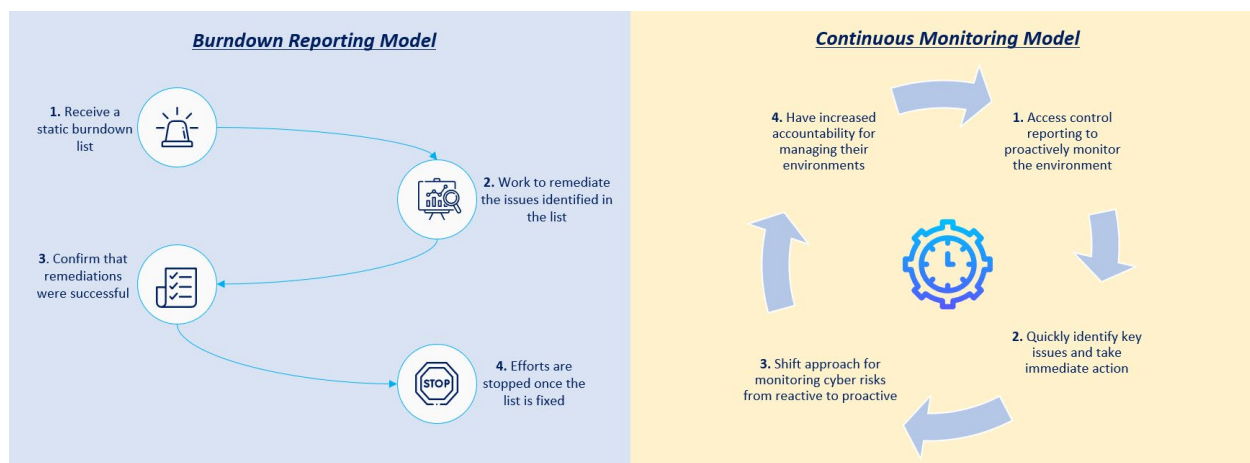


Figure 11 – Shift to Continuous Control Model

The benefits to proactive control monitoring are:

- Provides teams with the tools needed for holistic monitoring of control progress
- Promotes ongoing identification and remediation of cyber risks
- Enables teams to improve processes and go-forward strategy

5.6. Optional Attestation

In the past few years, attestations have grown in popularity to provide assurance over financial, security, and privacy related controls as required by compliance, regulators, or service clients. Adding a layer of formal attestations on quarterly basis provides the ability for companies to now gain confidence in the controls they have implemented in a method that is monitored by the financial, government, and regulatory agencies.

6. Conclusion

Establishing the foundations of a security controls framework for a 10G network provides the ability to proactively monitor implemented control effectiveness as well as identify risks in an environment that is continuously changing along with its attack surface. While it may seem challenging to implement a comprehensive framework such as NIST 800-53, it is possible to take a phased approach focused on risks with the highest impact to the network and organization. Engaging the Cybersecurity, Risk Management, and Compliance teams early on can ensure the appropriate controls are identified aligned to organizational goals and its risk strategy. The insights gained from a successful implementation can also start to showcase the security posture of the various organizational units managing the 10G network, establishing the accountability and responsibility necessary for continuous control monitoring. Securing a network is no small feat, making the right investments in security strategy and ensuring its measurable effectiveness provides the cyber protections necessary to deliver the transformative capabilities of a 10G network.

Abbreviations

GRC	Governance, Risk Management and Compliance
DoS	Denial Of Service
DevSecOps	A portmanteau of “Development” and “Operations”:
Botnet	Robot Network
IAM	Identity and Access Management
PI	Personal Information
NIST	National Institute of Standards and Technology

Bibliography & References

- [1] <https://www.cablelabs.com/10g/security>
- [2] <https://www.cisa.gov/uscert/ncas/alerts/aa22-158a>
- [3] <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf>
- [4] <https://www.cablelabs.com/10g>
- [5] <https://www.cablelabs.com/specifications/CM-SP-SECv4>
- [6] https://info.processunity.com/rs/638-QKL-150/images/PU_Sustainable_Cybersecurity_WP_Final.pdf

Designing Inclusive Learning Experiences Develops an Inclusive and Productive Workforce

A Technical Paper prepared for SCTE by

Erika Schoch

Sr. Program Manager, Technology Learning & Development
Comcast Cable
Erika_Schoch@comcast.com

Joel Moffatt

Principal Product Manager, Accessibility
Comcast Cable
Joel_Moffatt@comcast.com

Louis Rosario

Sr. Designer, Learning Solutions
Comcast Cable
Louis_Rosario@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background	3
2.1. Inaccessible Accessibility Training.....	3
2.2. An opportunity to create something new.....	3
3. Problem Statement.....	4
4. Accessibility and Inclusive Experiences (AIX) Principles.....	4
4.1. Why do we need AIX Principles? Mindset.	4
4.2. Why do we need AIX Principles? Skillset.	4
4.3. AIX Principles defined	5
5. Making training actionable and accessible	5
5.1. Understanding learners and learning strategy	6
5.1. Evaluating technology platforms	7
6. Determining value and impact.....	8
6.1. What demonstrates the value of training?	8
6.1.1. Learner feedback and ratings	8
6.1.2. What exists now that didn't before?	9
6.2. Value to business	9
7. Where do we go from here?.....	9
7.1. Inclusive tools and vendors.....	9
7.2. Learning design standards.....	10
7.3. A mindset established through the lens of the AIX Principles	10
8. Conclusion.....	11
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 – Audio Transcription Example.....	7
Figure 2 – Accessibility course star rating	8

List of Tables

Title	Page Number
Table 1 – Learner Feedback.....	8

1. Introduction

As the ADKAR Model - Awareness, Desire, Knowledge, Ability, Reinforcement - states “organizational change can only happen when individuals change” (Hiatt, 2006). Organizations within our industry are now challenged with creating in-person and virtual development opportunities for their employees that are both productive and inclusive. It is more important than ever that we create engaging, actionable, and accessible learning experiences for our entire workforce. This is especially important as the cable-tech industry focuses on creating a culture that celebrates inclusion in all aspects, from customer interactions to creating innovative new products.

Learning experiences designed with an inclusive mindset can increase the value of training through innovative uses of existing technology and drive the workforce’s adoption of accessibility practices and guidelines. This includes being intentional in framing disability within the Dimensions of Diversity (Loden, 1991) and making training actionable and accessible. In this paper we will explore how this can be done by using the lessons learned when creating new Accessibility training content and how this has made a larger impact on the business and workforce.

2. Background

2.1. Inaccessible Accessibility Training

The Comcast Accessibility team and Comcast’s learning and development team (L&D), called ULearn, originally started their partnership to create accessibility focused learning journeys in 2018. The intent was to create training for Comcast technologists that focused on accessibility awareness and outlined employee responsibilities in making products and experiences inclusive. While the two training modules developed were launched in 2019, they did not come without their own challenges. Not only were we asking a large portion of the workforce to change their mindset, accounting for accessibility in their work for the very first time, but we were using development tools to create this training that weren’t in line with Comcast accessibility standards. At the time, ULearn was only starting conversations on inclusive design, building in key components like alt text or video descriptions. We were also using tools that were not fully compatible with screen readers. In our original awareness training, learners were forced to toggle back and forth between browse mode and application mode, the second being a mode not widely used by many people who use screen readers. The reality was that we had unfortunately launched accessibility training that wasn’t accessible for all learners.

2.2. An opportunity to create something new

In 2020, ULearn and the Accessibility team at Comcast took the opportunity to improve on the work started in 2018, setting out to create role specific training for our technology department and solving for the accessibility issues discovered previously. Our goal was not only to create new Accessibility training content, but to use this content as a framework to make all future training content at Comcast accessible. This, in turn, will help create better training products overall while teaching our workforce how to create inclusive experiences at work and for our customers. As we launch the newly created Accessibility and Inclusive Experiences program (AIX) in 2022, we have influenced all designers and developers to begin thinking about accessibility at the beginning of product design, rather than try to retrofit solutions after a product launches. Using existing and past products as case studies, we have developed instructional templates that meet accessibility requirements and the needs of our users to shorten the development time and increase speed to market for our products.

3. Problem Statement

Our teams' reimagination of the accessibility content led us to some key questions: **How can we increase the value of learning experiences by designing them with an inclusive mindset? What change do we want to see because of these learning experiences?**

To solve for this and deliver content that makes an impact, we sought to achieve the following:

1. Be intentional in framing accessibility within the dimensions of diversity
2. Create actionable and accessible learning experiences using the tools currently available
3. Determine how we can measure the success of the learning experiences delivered

4. Accessibility and Inclusive Experiences (AIX) Principles

We start with framing accessibility within the dimensions of diversity. The Accessibility and Inclusive Experiences (AIX) Principles, developed by Comcast's Accessibility team and co-author Joel Moffatt (Principal Product Manager, Accessibility), are the pillars we used to create our learning products and the inspiration behind all of the content included.

4.1. Why do we need AIX Principles? Mindset.

Before we jump into the AIX Principles, we need to define some key terms. First, diversity is the set of characteristics — whether inherited or gained by experience — that make you uniquely you. This can include gender, sexual identity, race, religion, ability, age. At its core, inclusion is about allowing seats at the table for each of these dimensions of diversity while also hearing and valuing those voices because of, and not in spite of, those differences.

The nuances of diversity, equity, and inclusion can't be fully appreciated simply through repetition. Learning about and truly understanding cultural perspectives of others requires not just a learning posture, but an openness to the value of lived experiences other than one's own. For many learners, the idea of diving in and understanding the full cultural context of many different dimensions of diversity can be daunting. This need not be the case. Across those dimensions, there are many commonalities: in the way language is used, the value of empathy and understanding the experience of others, the function and role of allies, and so on. In other words, people don't need to start from scratch when embarking on a journey of learning about a particular dimension of diversity. An inclusive mindset is all anyone needs to get started.

4.2. Why do we need AIX Principles? Skillset.

Disability differs from other dimensions of diversity because there are some hard skills needed to succeed when delivering accessible employee and customer experiences. To be effective in gaining and employing those skills, learners need to understand the context of disability as a culture, the different methods and technologies people with disability may use to interact with content, spaces, and products. Once there's an understanding of why inclusion and accessibility are important, the workforce becomes more motivated to learn how they can create inclusive experiences and build accessibility into all they do.

In web and mobile development, developers adhere to W3C's Web Content Accessibility Guidelines (W3C, 2018), WCAG 2.1 Level AA. In learning and development settings, Universal Design for Learning principles (CAST, 2018) help educators meet students where they are by allowing for different means of expression and reception, like project based learning. For design, there are the Inclusive Design

principles. In the built world, the Americans with Disabilities Act (1990) codifies the minimal requirements to make a space accessible. Chances are, there are physical and digital aspects and platform and content considerations involved in everything you work on. And of course, you want to deliver the best experience possible for everyone who might engage with your product. And “product” here can be just about anything: email blasts, slide presentations, video entertainment, remote controls, theme park rides, apps, workspaces, and virtual learning.

4.3. AIX Principles defined

Instead of creating a curriculum that exists in a silo of disability and accessibility, we grounded the entire curriculum in Comcast’s Accessibility and Inclusive Experience principles. These AIX principles intentionally cut across all dimensions of diversity and take into consideration mindset and functional solutions. The principles are the initial lens you can look through to evaluate any project, product or other work right from the outset. They are:

INTEGRAL

“Inclusion and Accessibility are in our company’s DNA.”

Foster a culture of diverse teams, like-minded partners, and accountable leaders driven by inclusive policies.

RESONANT

“Diversity is represented in our content, products and workplace.”

Customers see themselves positively reflected in content and their needs met by our products. Use Inclusive language, cultural context, and accessible platforms.

EQUITABLE

“The experience accommodates the customer, not the reverse.”

Offer choice and control over settings, formats, and means of participation to create equitable experiences, whether physical or digital.

SITUATIONAL

“Combinations of where and how users do things are endless.”

Consider variables like venue, experience level, modes of interaction, and cultural context.

CONSISTENT

“Respect the user’s expectations and knowledge.”

Use familiar design and content conventions to make it easy to focus and understand the message and available actions.

VALUABLE

“Unnecessary features can break the experience.”

Add value rather than complexity and novelty, considering actual customer needs and preferences.

5. Making training actionable and accessible

Using the Accessibility and Inclusive Experiences (AIX) principles as the foundation to create inclusive training solutions, we began evaluating our tools. Developing learning solutions is divided into two groups: the learning strategy and the technology platform. The learning strategy is the core of

Instructional Design, outlining the who, what, when, where, why, and how. The technology is used to deliver the strategy to users. Combining the two creates an experience that should meet the user's needs.

5.1. Understanding learners and learning strategy

To fully understand how to meet the user's needs we need to understand who the users are, or the 'who' in the learning strategy, which is the most important part. The challenge we faced was that, while everyone can benefit from different aspects of inclusive design, how do we systematically incorporate elements of inclusive design into all of our content and processes? How do we create a product that meets the needs of an entire workforce?

This is where the AIX principles and WCAG guide our strategy when creating inclusive learning products. When designing the Accessibility and Inclusive Experiences program, we created a multi-part learning journey that introduces new users to the concepts of inclusion and accessibility. It also includes role specific training to help users put these principles into action. When asking our technologists to comply with certain standards it was instrumental to lead by example and ensure the design of our content was resonant, equitable, and situational.

To apply and demonstrate the use of the AIX principles and WCAG, we implemented the ULearn L&D Accessibility design standards. This included:

1. High contrast between text and background of at least 4.5:1
2. Closed captioning for video content
3. Do not use color alone to convey any information
4. All content is accessible when using keyboard only
5. Provide text alternatives for all non-text content

While these are the five core standards used on all ULearn L&D content, it was important to show our learners the impact of doing more than just the minimum required standards. Working within the capabilities of the eBook tools available to our design team, we set out to demonstrate how making a product accessible can be both simple and innovative. One solution was to create audio transcriptions for each page in a course eBook (Figure 1). By recording the text included on each page as an audio file, we used the following principles of Universal Design for Education referenced in the AIX principles section of this paper:

1. **Equitable use:** this feature is useful to people with diverse abilities
2. **Flexibility in use:** this feature accommodates individual preferences or abilities, for example, a learner may prefer to listen to learning content as opposed to reading it
3. **Simple and intuitive to use:** the feature is easy to understand with written directions on what the feature is used for and how to play the audio transcription



Figure 1 – Audio Transcription Example

While this solved for the content of the design, we needed to evaluate the technology used while creating an inclusive and accessible learning program to ensure it supports an inclusive experience for our learners.

5.1. Evaluating technology platforms

After discovering more about our learners and determining a learning strategy based on those findings, we needed to understand the tools and technology our learners use to interact with our training experiences. Partnering with the Comcast Accessibility Team, ULearn explored how to apply WCAG standards to develop equitable training experiences using industry instructional design tools that currently exist. Using those tools to help create speed to market, save resources, and save time, we learned that many instructional design tools used at Comcast had limited capabilities when it came to accessibility. The Comcast ULearn L&D team primarily uses two industry tools to create eBooks and interactive content. Those two tools have updated several times since 2019 to allow more control over the accessibility of the products we create for our users. However, these tools and most of the ones available in the industry, still don't have all the capabilities to create a fully accessible learning product. We have partnered with our vendors, especially providers of our development tools, to help them identify where they can improve their software.

Since these industry tools do not give us full control over the accessibility of the products we create with them, we often need to use features of these tools in non-traditional ways to meet the needs of our users. For example, we used an audio block within the eBook tool to allow users to listen to the page instead of reading it. While this isn't the expected use of the block, it fits the needs of our users. By using the tools in a different way we provide learners an equitable experience where they can choose to either read or listen.

In another example of how our tools restrict our capabilities, we can't access the code level of the products in a "What you see is what you get" WYSIWYG environment, since these tools are designed for people to create products without coding knowledge. However, since the tools don't provide the functionality to create a fully accessible experience for our learners, we now modify the output to include accessibility "fixes". This makes product maintenance more difficult. When a new version of the product is released, we implement the same fixes repeatedly. However, providing these fixes is still the better alternative for our users. As a short-term solution, we have templated our fixes to allow developers to copy and paste them throughout similar applications.

As we know this isn't a long-term solution, our ULearn Design team is adding to our current set of standards to help instructional developers interpret WCAG standards in order to create instructional experiences. Before any products go to market we include full quality assurance documentation.

6. Determining value and impact

6.1. What demonstrates the value of training?

When determining the factors that measure the success of our created learning journeys we had to answer a key question: **How do we determine the impact of building accessibility into our training programs when one of the primary goals is to create a mindset shift?**

As learners begin their accessibility training journey, the initial indicator of program success is learner feedback and course ratings. We also determined that success was measured in our own ability to create training products that follow the standards outlined in the content, allowing us to “lead by example.” Success can also be determined by recognizing new programs or opportunities that exist now, either as a direct result of the training or as a result of the company mindset shift that was influenced by developing inclusive accessibility training.

6.1.1. Learner feedback and ratings

Below are two examples of the initial feedback we gain from learners: anonymous quotes provided by the learner and a star rating, ranked on a scale between 1 (very dissatisfied) and 5 (very satisfied).

Table 1 shares just a sample of the quotes received, but they are a reflection of the actionable takeaways and mindset shift our learners gained as a result of the training content:

Table 1 – Learner Feedback

Anonymous quotes from learners		
"Very informative and helpful in terms of takeaways, specifically, terminology and mindful inclusion"	"Enlightening subject"	"This lesson really opened my eyes to how I can be more inclusive in ways that never occurred to me"

Figure 2, the Accessibility Course Star Rating, demonstrates that overall the content was well-received. Of the 13 courses launched as part of the accessibility programs, two of the 13 have received scores of 4 out of 5, while 11 of the courses have received 5 out of 5 stars, to date:




Figure 2 – Accessibility course star rating

The courses included in this data were released between November 2021 and May 2022 with a completion date of September 2022 for all learners assigned each course.

6.1.2. What exists now that didn't before?

One of the questions posed earlier in this paper was: **What change do we want to see because of these [inclusive] learning experiences?**

As we look towards the future, it's important to understand how the culture of our company has evolved since our original training content for Accessibility was launched in 2019. We have seen an increase in the number of programs and initiatives with a focus on accessibility and inclusion as more individuals have developed an inclusive mindset. They are also a direct result of the partnerships developed between the Comcast Accessibility Team and ULearn. This includes, but is not limited to:

1. The [Innovation for Inclusion of Diversity and Accessibility](#) conference, presented by Comcast Labs Connect and includes a partnership between ULearn and the Comcast Accessibility team. This conference focuses on the collaboration, culture and experiences of diverse teams driving the growth of accessible technology.
2. A Community of Practice, comprised of learning professionals and accessibility champions, this group comes together to discuss best practices in creating inclusive and accessible experiences for all
3. Focused effort between Accessibility, ULearn, leadership, and the business on implementing accessibility standards and practices
4. Peer learning and speaker series offerings with a focus on accessibility content

6.2. Value to business

By focusing on innovation with existing tools, we were able to create content using our own internal design and development resources. This gave our team control over the project timeline, cost, and to the best of our ability, the accessibility standards built in to the training created. We saved money by using existing platforms. Engaging a third party vendor to create the content can sometimes cost tens of thousands of dollars per eLearning module.

As is the case with most businesses, one of Comcast's most valuable resources is time. By using our own design teams we created accessible content without the extended lead times needed to secure contracts or onboard the right tools. As many know too well, the procurement process is one that can often take many month to complete.

7. Where do we go from here?

The experiences outlined in this paper have provided a better understanding of what we need to do in the future to create inclusive and valuable learning experiences. Our goal should be to create content that empowers our workforce without having to retrofit solutions using technology that isn't fully accessible.

7.1. Inclusive tools and vendors

The work we've already done to create inclusive training content provides a better understanding of what we need to do in the future to avoid retrofitting solutions.

This starts with the vendors we partner with and the tools we use to create content. Moving forward we need to ensure every vendor we engage with follows our own accessibility standards and policies,

including the Web Content Accessibility Guidelines. Over time, and with a commitment to changing our mindset and culture to one of inclusion, it should become apparent that we're working with a vendor that is compliant or inclusive. This standard should be met by any existing vendor or any new partnership we take on, starting immediately.

As discussed earlier, the tools and platforms we use are instrumental in creating accessible learning experiences for our workforce.

7.2. Learning design standards

Earlier in this paper we mentioned the five ULearn L&D design standards included in each of our learning products. As we move towards the future, it's important to note that our need for inclusion and standards are constantly evolving. We need to be open to that change, constantly re-evaluating our standards while providing our workforce the resources needed to feel confident in applying these standards. Our ULearn design team is currently working on updated design standards to align with newly implemented Accessibility Policies and Guidelines launched at Comcast this year. This includes PDF Accessibility standards, Office Documents (Word, Excel, PowerPoint, and Email) Accessibility Standards, and Instructor led guidelines for inclusive facilitation (both in-person and virtual), to name a few. ULearn L&D, in partnership with the Comcast Accessibility team, are launching these standards in 2022 with the goal of 100% compliance by 2024 for anyone creating content, including our instructional designers, developers and product managers.

7.3. A mindset established through the lens of the AIX Principles

Accessibility training at Comcast has two main objectives. First, teach individuals the skills they need to build accessibility into all they do. Second, the training must change the learner's mindset so that they are receptive to the content and understand why it's important. With the mindset established through the lens of the AIX Principles, individuals in any role can readily assess the inclusivity of their work and why building inclusive products or experiences benefits everyone.

The guidelines for living up to the AIX Principles are found in WCAG for web content, in the ADA for the built environment, in Universal Design for Learning Principles, and in the associated best practices for executing on each of those. Success looks like:

INTEGRAL

"Inclusion and Accessibility are in our company's DNA."

Example: Your company has established a policy on inclusive media, has tools and resources in place to make it easy to make content inclusive and accessible, and has training in place to enable employees to reliably adhere to that policy.

RESONANT

"Diversity is represented in our content, products and workplace."

Example: inclusive media guidelines call for diversity in casting and using inclusive language in scripting and copy. Customers and employees see their culture reflected in our content and messages.

EQUITABLE

"The experience accommodates the customer, not the reverse."

Media player controls include closed captions, audio description and alternate languages. Content includes those options across all asset versions at time of delivery.

SITUATIONAL

“Combinations of where and how users do things are endless.”

Experiences are evaluated across many display types and designed and built to work well on small as well as large screens.

CONSISTENT

“Respect the user’s expectations and knowledge.”

Media player controls use expected conventions. Player options are found in the same location and identified with the same iconography across platforms.

VALUABLE

“Unnecessary features can break the experience.”

Media player controls are not cluttered with niche options, e.g. too many caption font options, including illegible ones.

8. Conclusion

As we mentioned in the beginning, the ADKAR model states that “organizational change can only happen when individuals change” (Hiatt, 2006) This became apparent through the development and launch of the AIX learning journeys as we discovered more people than ever before are open to embracing an inclusive mindset. As we see more individual champions of accessibility within the cable-tech industry, there’s an opportunity to continue the conversations that will drive change within our industry. With the support of our leaders and peers, we can engage with the resources available and proactively create inclusive experiences. This includes a keen focus on how the Comcast Accessibility and ULearn Learning & Development teams can partner together to develop our workforce, how we can lead by example and build upon the accessibility standards we have already implemented. We can empower our learners and each other by being intentional in framing disability within the dimensions of diversity, basing our work on the Accessibility and Inclusive Experiences principles, and making training actionable and accessible for all.

Abbreviations

AIX	Accessibility and Inclusive Experiences
ULearn	Comcast Learning and Development Department
WCAG	Web Content Accessibility Guidelines
ADA	Americans with Disabilities Act
ADKAR	Awareness, Desire, Knowledge, Ability and Reinforcement
L&D	Learning and Development
WYSIWYG	What you see is what you get

Bibliography & References

CAST. (2018). *Universal Design for Learning Guidelines version 2.2*. Retrieved from UDL Guidelines: <https://udlguidelines.cast.org/>

Hiatt, J. (2006). *ADKAR: a model for change in business, government and our community*. Loveland: Prosci Learning Center Publications.

Loden, M. a. (1991). *Workforce America!: managing employee diversity as a vital resource*. Homewood: Business One Irwin.

W3C. (2018, June 5). *Web Content Accessibility Guidelines (WCAG) 2.1*. Retrieved from W3C: <https://www.w3.org/TR/WCAG21/>

Americans with Disabilities Act of 1990, 42 U.S.C. § 12101
(1990). <https://www.ada.gov/pubs/adastatute08.htm>

Detection of Passive Intermodulation in Drop Wiring by Burst Transmission Analysis

Diodes are common, but the network resists

A Technical Paper prepared for SCTE by

Tom Williams

Distinguished Technologist
Cable Television Laboratories Inc.
858 Coal Creek Cir. Louisville CO
303-661-3486
t.williams@cablelabs.com

Alberto Campos

Fellow
Cable Television Laboratories Inc.
858 Coal Creek Cir. Louisville CO
303 661 3377
a.campos@cablelabs.com

Larry Wolcott, Fellow, Comcast

Jason Rupe, Ph.D., CableLabs

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background	3
2.1. Background on Nonlinear Distortion	3
2.2. Background on CPD and PIM.....	7
3. Lab Experiments	8
3.1. FDX	13
3.2. PNM	13
4. Field tests	13
5. Detection Methods for Corrosion Diodes	14
6. Conclusion.....	15
Acknowledgements	15
Abbreviations	15
Bibliography & References.....	16
7. Appendix	17
7.1. Appendix A Diode Equation for a I-V Curve.....	17
7.2. Throbbing CPD Caused by Loose Seizure Screws.	17
7.3. Mathematics of nonlinear signals.....	19

List of Figures

Title	Page Number
Figure 1 –Sine wave input applied to nonlinear circuit produces a distorted sine wave.	4
Figure 2 – 2 nd and 3 rd order spectral components created by sine waves at f1 and f2.	4
Figure 3 – Plots of fundamental, 2 nd and 3 rd order distortions vs fundamental RF levels.....	5
Figure 4 – Non linear distortion	6
Figure 5 – Schottky Diode distortions	7
Figure 6 – CPD non linear distortions in downstream	8
Figure 7 – Circuit simulation of corrosion in a drop.	9
Figure 8 – OFDMA spectral plot	10
Figure 9 – OFDMA spectral plot (Shunting corrosion diode).....	11
Figure 10 – Cross -correlation plot.....	12
Figure 11 – MER of OFDMA signal	13
Figure 12 – 120 Hz CPD level vs. time.	18
Figure 13 – Shunt resistance-nonlinear distortion	19

1. Introduction

Home coaxial wiring and drop cable are expected to contain corrosion diodes created between dissimilar metals and metal oxides. Generally, corrosion diodes have not been a serious problem up to this point because continuous composite downstream radio frequency (RF) levels inside homes were not sufficient to force the corrosion diodes into hard conduction, which would create nonlinear distortion. DOCSIS[®] transmitted signals can be strong enough to force the corrosion diodes to conduct, but due to sub-split upstream frequencies, damage to downstream signals has been limited. It should be technically possible to automatically detect passive intermodulation (PIM) products at the headend due to their correlation with the main transmission and the spectral location of the nonlinear distortion. This paper discusses the mechanism for creating PIM in sub, mid and high split plant using orthogonal frequency division multiple access (OFDMA) transmissions. The paper also reports on PIM impairments created in our lab and analyzed using digital signal processing (DSP), discusses the potential impact of nonlinear noise on full duplex (FDX), and provides techniques to tackle this until now invisible plant impairment.

2. Background

2.1. Background on Nonlinear Distortion

Plant distortions can be separated into two categories: linear and nonlinear. Linear distortions consist of group delay, echoes, and channel ripple/tilt to name a few. Nonlinear distortions are created by mixing signals with themselves or other signals, and this mixing process creates energy at new frequencies. This distortion is also called intermodulation distortion (IMD). Cable plants have both types of distortions, and while linear distortion can generally be remedied with equalizers, nonlinear distortions are much harder to remove. Generally, it is best to avoid generating nonlinear distortions in the first place, if possible.

Mathematically, the input-output relationship of many devices can be described as a polynomial or Taylor series as shown below. More is explained in Appendix 8.3.

$$E_{out} = A \cdot E_{in} + B \cdot E_{in}^2 + C \cdot E_{in}^3 \dots \quad (1)$$

E_{out} is the composite output signal, E_{in} is a composite input signal, A is the gain, B is a coefficient for the second order distortion, and C is a third order coefficient. Additional higher-order coefficients typically exist and may not be ignored if input signal levels are large.

Figure 1 illustrates a transfer function modeled as a third order Taylor series. An input sine wave passes through a device with third order distortion, such as a push-pull amplifier. Most cable amplifiers are a push-pull design to cancel 2nd order distortions. In the figure, the resulting distorted output wave is illustrated as a red curve, while a desired undistorted amplifier waveform is black. The distance between the red curve and the black curve is the voltage of the distortion energy.




Figure 1 –Sine wave input applied to nonlinear circuit produces a distorted sine wave.

Figure 2 is a spectral plot of 2nd and 3rd order output products created by two sine waves, f_1 and f_2 . Consider an input signal E_{in} comprised of just two sine waves, f_1 and f_2 . They are mixed in a circuit with a transfer function with only 2nd and 3rd order distortions. Equation (1) may be expanded using trigonometric identities. Second order distortion products will be created at $2f_1$, $2f_2$, f_1-f_2 , and f_1+f_2 . Third order distortion products will be generated at $3f_1$, $3f_2$, $2f_1-f_2$, $2f_2-f_1$, $2f_1+f_2$, and $2f_2+f_1$.




Figure 2 – 2nd and 3rd order spectral components created by sine waves at f_1 and f_2 .

Generally, higher order distortions than the 3rd order should be considered at high composite RF signal levels.

System designers use an intercept point diagram, as illustrated in Figure 3, to model nonlinear components. A plot is made of input signal level vs. output signal level, and the fundamental slope of that plot is the amplifier's gain. A second plot is made of 2nd order distortion level vs. input signal level, and that plot has a slope twice as steep as the fundamental due to the squared term in the Taylor series. Likewise, a third order distortion plot has a slope three times as steep as the fundamental due to the cube term in the Taylor series. If the 3rd order distortion line is extrapolated (red lines), it crosses the fundamental line at a point called IP3, or the 3rd order intercept point. This key number is used to predict expected distortion levels, and usually is specified on amplifier and mixer data sheets. Another point of interest on data sheets is the 1dB compression point, where the fundamental signal's output falls 1dB below its ideal level.




Figure 3 – Plots of fundamental, 2nd and 3rd order distortions vs fundamental RF levels.

In Figure 3, fundamental and 3rd order plot can be projected to find a third order intercept point. This method can be used to predict distortion levels in amplifiers or other nonlinear components. This projection graphing method applies also for corrosion diodes and shows how higher RF fundamental levels produce higher levels of nonlinear distortions.

For test purposes, cable engineers and technicians sometimes raise (or lower) plant levels by a decibel to identify if resulting nonlinear distortion energy goes up by 2 or 3 dB, identifying the distortion as 2nd or 3rd order. To properly perform this test, pilot signals must be held constant because amplifier automatic gain controls (AGCs) will attempt to change gain.

Typically, digital cable signals being transported have a flat spectral energy with an approximately rectangular spectral shape. An analysis method for nonlinear distortion is to use a double or triple frequency-domain convolution of the input signal with itself. This method is used because a convolution in the frequency domain is mathematically equivalent to a multiplication (e.g. squaring or cubing) in the

time domain, and vice-versa. Figure 4 illustrates the creation of 2nd and 3rd order nonlinear distortion products. The convolution of a rectangular spectrum with itself produces a triangular shape on a linear vertical scale occupying twice the original rectangular spectrum, while a triple convolution of a rectangular spectrum with itself produces a “haystack” shaped spectrum occupying three times the original rectangular spectrum. Figure 4 illustrates a relatively narrowband input OFDMA signal 70 MHz to 100 MHz (red) which creates 2nd and 3rd order distortion products. Distortion products are not drawn to scale for the sake of clarity. The more convolutions a rectangular signal undergoes, the wider its resulting nonlinear distortion spectrum gets.




Figure 4 – Non linear distortion

The spectral plot in Fig. 4 shows how a rectangular-shaped block of spectrally-flat energy, such as an OFDMA or OFDM spectrum (in red) produces nonlinear distortion products (in black). The vertical scale is linear. The 2nd order products employ as double convolution, producing a triangular-shaped spectrum. 3rd order distortion products create haystack-shaped spectral energy. Occupied bandwidth increases linearly with each higher order distortion coefficient and the distortion takes on a relatively “flat” spectrum. This is particularly true for wide transmitted rectangular (bandwidth) blocks.

The spectral photo in Figure 5 explains why we generally have not been affected by nonlinear distortion products with a sub-split frequency plan inside home wiring up to this point. If, for example, an ATDMA upstream single carrier transmission at 30 MHz is 6.4 MHz wide, the second order difference term is 0 MHz to 6.4 MHz yet cable upstream amplifiers do not pass energy below 5 MHz. The second order sum products will be 12.8 MHz wide centered at 60 MHz, so 2nd order distortion can affect downstream signals located between 53.6 MHz and 66.4 MHz but received signal degradation occurs only for the duration of a burst upstream transmission.




Figure 5 – Schottky Diode distortions

Figure 5 is a lab experiment of a single carrier 6.4 MHz signal at 33.8 MHz being distorted by a Schottky diode producing 2nd and 3rd order distortions (lower plot). LeCroy DSO digital oscilloscope has a horizontal frequency range of 0 MHz to 200 MHz, and 10 dB/div is vertical scale. Upper plot is voltage vs. time, captured at 500 Msamples/second.

2.2. Background on CPD and PIM

CPD has long been a known impairment source for downstream cable plant. The distortion is created by corrosion diodes in cable plant, usually located in hard line plant where signal levels are large due to nearby line amplification. High level downstream signals produce nonlinear distortion in all bands, but the distortion is visible in the upstream band. [9]

PIM is a well-known impairment in wireless systems also caused by corrosion, and it sometimes occurs on RF feedlines. Because bandpass filtering is possible, the biggest concerns are third order distortions located just above, just below, and on the carrier. Both CPD and PIM are caused by corrosion diodes, and result in nonlinear distortions in RF. Therefore, CPD can be considered a subvariant of PIM, which is a generic term.

A question arises why CPD only seems to affect upstream signals, while it appears to have relatively flat spectral content. See Figure 6, which shows typical upstream and downstream carrier levels (per 6MHz) at the output of a downstream amplifier, which is also the input of an upstream amplifier. Assuming CPD is spectrally relatively flat, with a -40dBc distortion relative to downstream carriers each at +50 dBmV, upstream signals experience a 18dB carrier-to-noise ratio. This ratio is below the threshold for a 64-QAM upstream receiver.




Figure 6 – CPD non linear distortions in downstream

Figure 6 is a diagram explaining why CPD nonlinear distortion is seldom observed on downstream channels. As an example, in a 6 MHz band, if a downstream signal is at 50dBmV and creates a spectrally-flat CPD distortion -40dBc relative to the downstream carrier level at the output of a downstream amplifier, distortion power level is -10dBmV. So, it is relatively harmless to a downstream carrier. But an upstream carrier, attenuated by a flat loss of a 32dB tap, suffers a modulation error ratio (MER) degradation of only 18dB (28dB-10dB). This is below threshold and is a continuous impairment to all upstream signals in the node due to noise funneling.

One advantage that coaxial cable signals have over long term evolution (LTE) signals is cable composite signal levels are much lower, such as approximately 1-2 watts for cable versus 20 watts for cellular transmissions. However, cable plant has many more connectors relative to a cellular feedline and antenna. Cable downstream also employs uptilt, placing more energy into the higher frequencies. This design will cause some distortion components to land at higher frequencies, which hopefully are out of band.

3. Lab Experiments

It has been observed in lab experiments that a wide OFDMA transmission can produce relatively flat distortion energy over a wide band. If a strong signal, such as a wide OFDMA transmission, mixes with itself in a corrosion diode, a difference product will start at DC (0 HZ) and extend to the bandwidth of the OFDMA carrier with a triangular-shaped spectrum (viewed on a linear display). This is a 2nd order difference mixing product. However, if the OFDMA transmission mixes with another strong signal, such as a MoCA[®] transmission or a citizen band (CB) transmission ingress signal, a difference product will not be centered at direct current (DC), but at a difference frequency, as expected. For example, an OFDMA transmission centered at 200MHz can mix with a MoCA signal centered at 1150 MHz and produce interference centered at 950 MHz, which is in a downstream frequency band. The 2nd order interference will be a convolution of the MoCA signal's spectrum with an OFDMA signal's spectrum.

A first test was to pass an ATDMA single carrier 6.4 MHz signal through a corrosion diode simulator circuit, with a Schottky diode. The spectral plot was previously illustrated in Figure 5, showing the

expected 2nd and 3rd order distortion products. This plot can be compared to Figure 4, which illustrates frequencies at which distortion products are expected.

The circuit used is illustrated in Figure 7 was used to produce upstream-only FDD distortion products. This circuit simulates corrosion in the outer conductor of a coaxial connection, such as a ground block. A shunt resistor's value was varied to increase the amount of distortion caused by the nonlinear 1N5711 Schottky diode with a given RF drive level. Two duplex filters were used to prevent downstream nonlinear distortion, which can potentially cause interference with the cable modem's (CM's) reception. The optional inductor was used to prevent DC bias on the diode and has very large inductive reactance at test frequencies.




Figure 7 – Circuit simulation of corrosion in a drop.

Figure 7 is a circuit used in the lab to simulate corrosion in a drop. Two mid-split diplexers were used with a mid-split diplexer to limit the effects of distortion to only the upstream signal. The CM operates on a mid-split frequency plan. The shunt resistor (R=) is varied to increase or decrease MER caused by the nonlinear diode. The meter function (M), which may be located at the headend, samples upstream energy which is distorted.

Figure 8 is a pair of plots obtained from a LeCroy HDO 6104-MS digital oscilloscope. The top plot is voltage vs. time, and the lower plot is a fast Fourier transform (FFT) of the time samples. The sample rate was 500 M-samples per second. The corrosion diode is shunted with 0 ohms and thus had no effect. The OFDMA transmission was almost continuous due to heavy data loading, and the occupied bandwidth was 40 MHz to 80 MHz. Note that there is negligible nonlinear distortion.

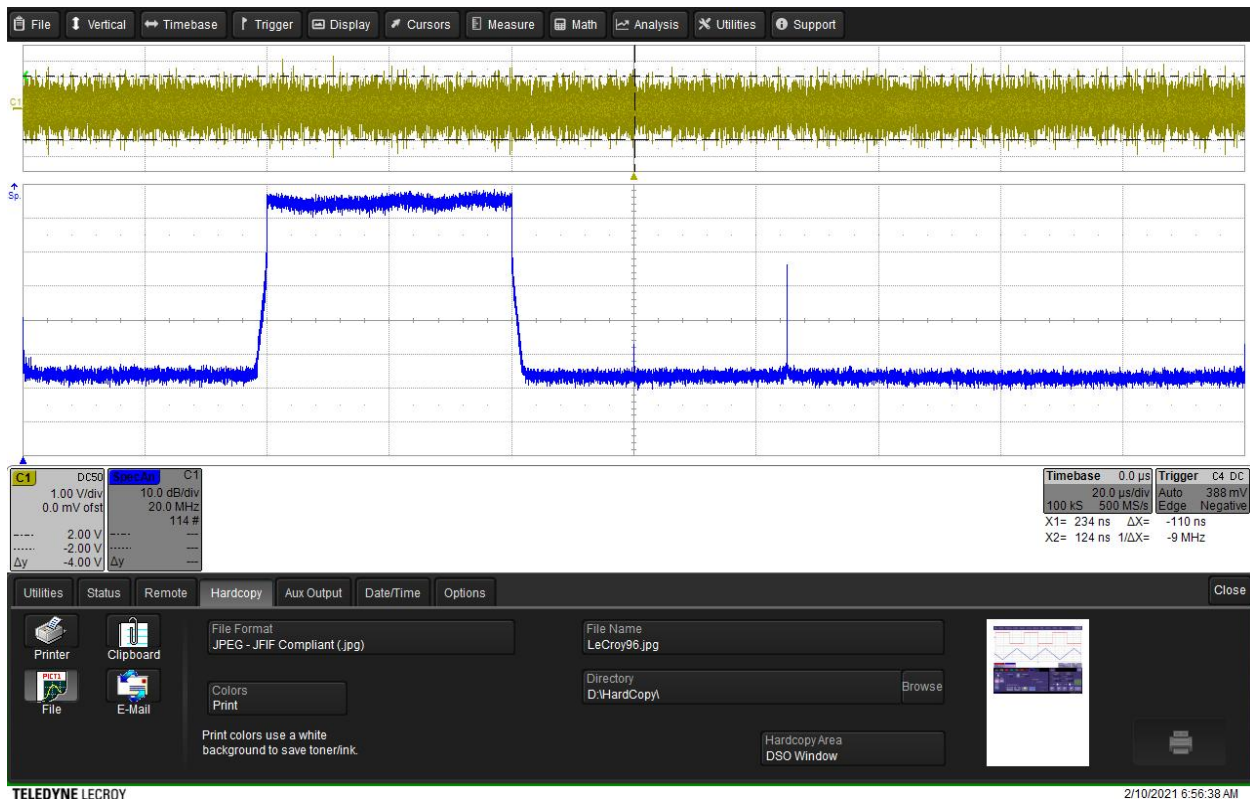


Figure 8 – OFDMA spectral plot

Figure 8 is a spectral plot of 40 MHz to 80 MHz OFDMA signal with $R = 0$ ohms (no corrosion diode effects). Top trace is voltage vs time, and bottom trace is dB magnitude vs. frequency. MER of 64-QAM OFDMA signal was an excellent 48.7 dB.

Figure 9 is a spectral plot with a shunt resistor value of 39 ohms, showing broadband distortion, notably sum and difference 2nd order.

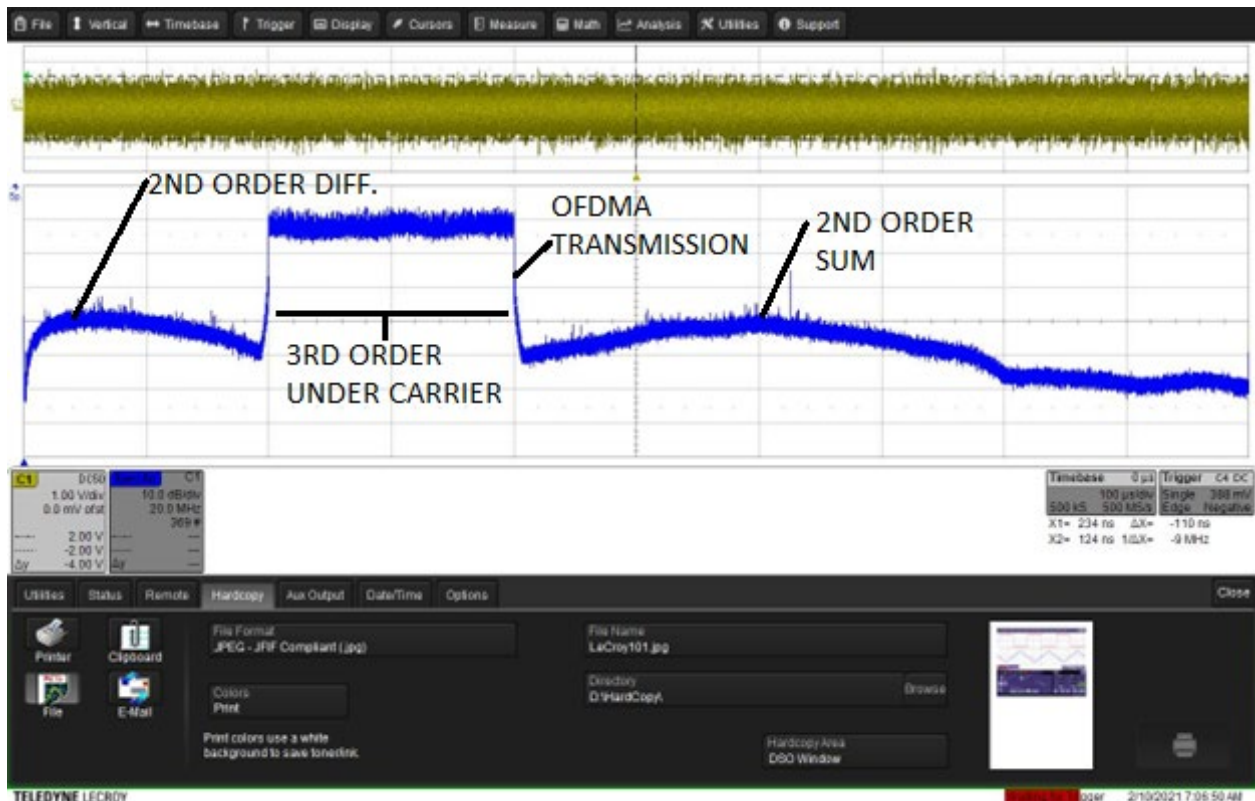


Figure 9 – OFDMA spectral plot (Shunting corrosion diode)

Figure 9 is a spectral plot of 40 MHz to 80 MHz OFDMA signal with $R = 39$ ohms (shunting corrosion diode). Transmit power was approximately 47 dBmV, and a MER of 64-QAM signal was 36.5dB. Observe wide spectrum of 2nd order distortion of difference products (on left) and sum products (on right). 2nd and 3rd order distortions overlap with a wideband transmission.

Figure 10 is a cross correlation plot of 2nd order measured lab distortion with ‘manufactured’ distortion. This method was published in 2013 [6]. It uses the received OFDMA carrier to manufacture a 2nd order distortion signal. That distortion energy is cross correlated with the filtered received energy at a frequency where 2nd order distortion is expected. This testing method is practical because upstream triggered upstream spectrum capture (UTSC) can capture a signal with nonlinear distortion. This cross-correlation method can be used for 3rd order distortion but works best where the cross-correlation capture is performed in a vacant band, such as with an active-quiet probe measurement.




Figure 10 – Cross -correlation plot

Figure 10 is a cross-correlation plot of the filtered signal time domain signal (measured) illustrated in Figure 8 with a mathematically-created signal (manufactured) created by squaring the filtered time domain signal illustrated in Figure 8, top trace. See 2013 technical paper [6]. An impulse at zero time offset indicates significant energy is 2nd order distortion.

Figure 11 is a plot of modulation error ratio (MER) on a 40 MHz to 80 MHz, 64QAM OFDMA transmission vs. shunt resistor value. Shunting the Schottky diode with a lower value resistor reduces distortion.




Figure 11 – MER of OFDMA signal

Figure 11 is a plot of MER of an OFDMA signal vs. shunt diode resistance, illustrated in Figure 6. When the shunt diode approaches zero, the corrosion diode is still present, but has no effect. A low value shunt resistance across a shunt diode is achieved by a high clamping force, achieved, for example, by a tight seizure screw or properly torqued housing bolts.

3.1. FDX

The process of PIM creation in FDX is going to be similar to FDD, but the lack of a duplex filter in the 108 MHz to 684 MHz FDX band means that distortion products can affect both upstream and downstream signals simultaneously.

Another concern is the adaptive equalizer in the fiber node won't be able to equalize non-linear distortion: it is designed to equalize linear distortion.

3.2. PNM

UTSC is essential in detecting anomalous US noise and distortion. Currently not all vendors implement this feature. Data should be in complex format for DSP. Furthermore, a UTSC capture needs to occur while a specific CM is transmitting at high power.

4. Field tests

CMs and cable modem termination systems (CMTSs) today are in some cases able to capture frequencies outside of their receiving ranges. Some CMs can capture low frequencies, and others can work in sub split mode on a mid-split or high split plant, etc. CMTSs can in some cases provide spectrum capture data (free run mode of upstream triggered spectrum capture or other means) for the downstream as well as upstream bands.

5. Detection Methods for Corrosion Diodes

Nonlinear distortion detection is a difficult problem because, in normal operation, upstream CM transmissions are brief RF bursts which are difficult for test equipment to capture (trigger on), and a source of which is generally difficult to identify.

A first detection method is to observe the MER for a particular CM's transmissions and compare them to the MER in transmissions from other CMs in the same node in approximately a same time period. Because of additive noise funneling, it is expected that all CMs in a fiber node should be experiencing a common noise background. However, a common noise background can vary with time. On the other hand, nonlinear distortion instantaneously follows the RF level of a transmission. A worse MER (relative to other subscribers in the node) could indicate nonlinear energy contamination in a signal path that a CM is using, but variable common background noise must be considered. An upstream MER test can be made on several CMs with nearly simultaneous upload speed tests at high RF power levels. This approach should yield a good indicator of homes with nonlinear distortion.

A second method is to look for distortion energy above the transmitter's frequency as a sum product or as harmonics. A challenge with this method occurs when the harmonics land in the downstream spectrum, which propagates away from the headend. Detection of nonlinear distortion in a CM on downstream caused by a CM's upstream transmissions may be considered a self-inflicted wound. However, using full band capture (FBC), if a CM can be forced into a speed test at a predictable transmit frequency, distortion energy may be observed either in a vacant band, or in the spectral null between two adjacent RF carriers. 2nd and 3rd harmonics can be observed with this method. A delta FBC plot can be made between CM transmitting and CM not transmitting.

A third method is to look for distortion energy below the transmitter's frequency, as a difference product. This method looks promising, but a UTSC needs to be done while a CM is at transmitting at high composite power levels. Testing needs to be initiated with a process such as a speed test, concurrent with a UTSC trace capture. It is expected that UTSC traces will contain 4096 in-phase and quadrature time samples.

A fourth method is to look for out-of-band energy that "follows" an RF transmission in time. Nonlinear distortion is created simultaneously with a transmission, so detection can be done in the time domain.

A fifth method is to plot a histogram of time domain voltage samples and look for non-symmetry, indicating signal clamping by a corrosion diode conducting. This method requires relatively severe clipping (distortion).

A sixth method is to analyze the MER error vectors in the OFDMA signal and correlate them with 3rd order distortion created from the error-free symbols. This requires the symbols in the OFDMA signal to not cross decision boundaries. [7]

Another method involving bonded upstream ATDMA signals is to get multiple upstream channels to transmit simultaneously from a single CM at high power. A block of three or four 6.4 MHz wide carriers can be approximately treated as one wide OFDMA-like block of random energy.

Distance ranging can be done to facilitate locating a corrosion diode in house. For example, the Optus (Australia) method using DSP can be employed using a test waveform such as a chirp signal. A display can show a technician TDR results for both linear and nonlinear distortions on the same screen. If timing matches, a linear discontinuity is the source of nonlinear distortion. This will speed troubleshooting.

The cable industry desperately needs functional, consistent, correct, upstream triggered spectrum capture capabilities. We don't yet generally have the ability to trigger an upstream spectrum capture for say a particular CM, which would be very helpful for finding PIM. Lacking that, we need alternate methods to identify and localize nonlinear impairments caused by upstream signals. Campos, Hamzeh, and Williams created a solution in 2013 [1]. Ultimately, capturing spectrum from a single CM transmitting would allow us to look for that PIM signature. If we are able to use UTSC to capture that signal, we could identify the distortion by the harmonics. If two narrow frequencies are transmitted, then a mixing product would reveal PIM if it appears. Note that two CMs transmitting at the same time, on two frequencies, could potentially create a mixing product in the plant, though that is less likely due to transmission loss leading to lower energy levels at potential points of corrosion diodes. Still, it may be possible, and may be found.

6. Conclusion

The question plant engineers have been asking is where are the corrosion diodes in my plant? A better question to ask is “where are the high value shunt resistors?” as very many junctions in the cable plant are candidates for a corrosion diode. A loose seizure screw is an example of how a shunt resistor may be created.

Calling a distortion-producing diode a PIM diode or a CPD diode is a relatively arbitrary distinction, as a single diode can be both a CPD diode producing upstream distortion, or a PIM diode producing downstream distortion. CPD is viewed as a subvariant or type of PIM.

For wide bandwidth transmissions, nonlinear distortion, particularly 3rd order, can be detected in adjacent frequency bands using UTSC, aided by a command to a CM to transmit a powerful signal, such as a speed test. An UTSC signal, captured as in-phase and quadrature component signals, can be processed to detect the presence of nonlinear distortions.

Clamping screws need to be correctly tightened to keep corrosion diodes shunted with low value resistors.

Acknowledgements

The authors of this paper would like to thank several people for their assistance with this work: Doug Jones, Sheldon Webster, James Lin, and Jay Zhu.

Abbreviations

AGC	Automatic Gain Control
ATDMA	Advanced Time Division Multiple Access
CB	Citizens Band
CMTS	Cable Modem Termination System
DC	Directional Coupler
DSP	Digital Signal Processing
FBC	Full Band Capture
FDD	Frequency Division Duplex
FDX	Full Duplex
FFT	Fast Fourier Transform

IMD	Intermodulation Distortion
LTE	Long Term Evolution
MER	Modulation Error Ratio
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
PIM	Passive Intermodulation
RF	Radio Frequency
TDR	Time Domain Reflectometer
UTSC	Upstream Triggered Spectrum Capture

Bibliography & References

- [1] A. Campos, B. Hamzeh, T. Williams, "Testing for nonlinear distortion in cable networks," CableLabs Whitepaper 2013.
- [2] R.Hranac, "Passive Device Intermod," Communications Technology, September 1998.
- [3] E. Sapienza "Analysis of Passive Intermodulation in a Junction: Physical Simulation Models for Non-Linear Electrical Contacts" Master Thesis 2018-2019- Dept. of Electronics and Telecommunications, Politecnico di Torino
- [4] H. Yang, H. Wen, Y. Qi and J. Fan, "An equivalent circuit model to analyze passive intermodulation of loose contact coaxial connectors", IEEE Transactions on Electromagnetic Compatibility, Vol.60, No 5, 2018.
- [5] International Standard IEC 62037 RF connectors, connector cable assemblies and cables – Intermodulation level measurement, 1st edition
- [6] Testing for Nonlinear Distortion in Cable Networks by A. Campos, B. Hamzeh, and T. Williams 2013 Cable Tech Expo.
- [7] US pat. 11,184,202
- [8] T. Williams, J. Rupe, "Common Path Distortion (CPD) in Cable Networks," Journal of Network Operations, Vol. 4, No. 1, December 2018, available at https://wagtail-prod-storage.s3.amazonaws.com/documents/SCTE-ISBE-NOS_Journal_V4N1.pdf
- [9] B. Patel, "Report on Common Path Distortions," Feb. 3, 1998, Dept.: Network Development.

7. Appendix

7.1. Appendix A Diode Equation for a I-V Curve

The I-V curve (diode characteristic curve) can be found by the following nonlinear equation. The equation is also known as ideal equation of a diode or diode law:

$$i = I_s(e^{\frac{qv}{kT}} - 1)$$

Where

i = Current flowing through the diode

I_s = Reverse or dark saturation current (Typical value for silicon is 10⁻¹² amperes)

e = Base of the natural logarithm (2.71828)

q = Charge on electron (1.602 x 10⁻¹⁹) in coulombs, which the absolute value of an electron charge.

v = applied voltage across the diode

k = Boltzmann's constant (1.380 x 10⁻²³ joules/kelvin)

T = Absolute temperature in degrees Kelvin (room temperature is about 300 kelvin)

7.2. Throbbing CPD Caused by Loose Seizure Screws.

In Williams and Rupe, "Common Path Distortion in Cable Networks," [8], CPD is studied.

Instead of using two fixed CW carriers to implement the SCTE-109 CPD lab test procedure, if one (or both) of the two carriers is made tunable, the frequency of a returning CPD CW difference beat will also vary. The two downstream CW carriers are combined at the headend (or node insertion point) to make a headend 2-CW test signal. CPD in plant is revealed by the presence of the difference signal in the return band.

The difference signal is a 2nd order impairment and was observed in 7 out of approximately 70 nodes. 3 of the nodes had a static difference product (CW), but 4 were dynamic.

In [8], dynamic upstream CPD on a sub-split plant was observed to go up and down in power level at a 120 Hz rate. This was observed in four nodes in the hub site containing approximately 70 nodes, so dynamic throbbing was about a 5.7% problem, and static CPD was also about a 4.3% problem. Two CW signals, one at 800 MHz and one at 840.5 MHz, were applied at analog video carrier level into the downstream, and a modulated difference product at 40.5 MHz was observed. Figure 12 is a plot of RF level in dB vs. time, and repetition rate was 120 Hz. The spectrum analyzer's span was set for 0 Hz.

Fortunately, currently available CPD location gear detects both throbbing and static CPD.

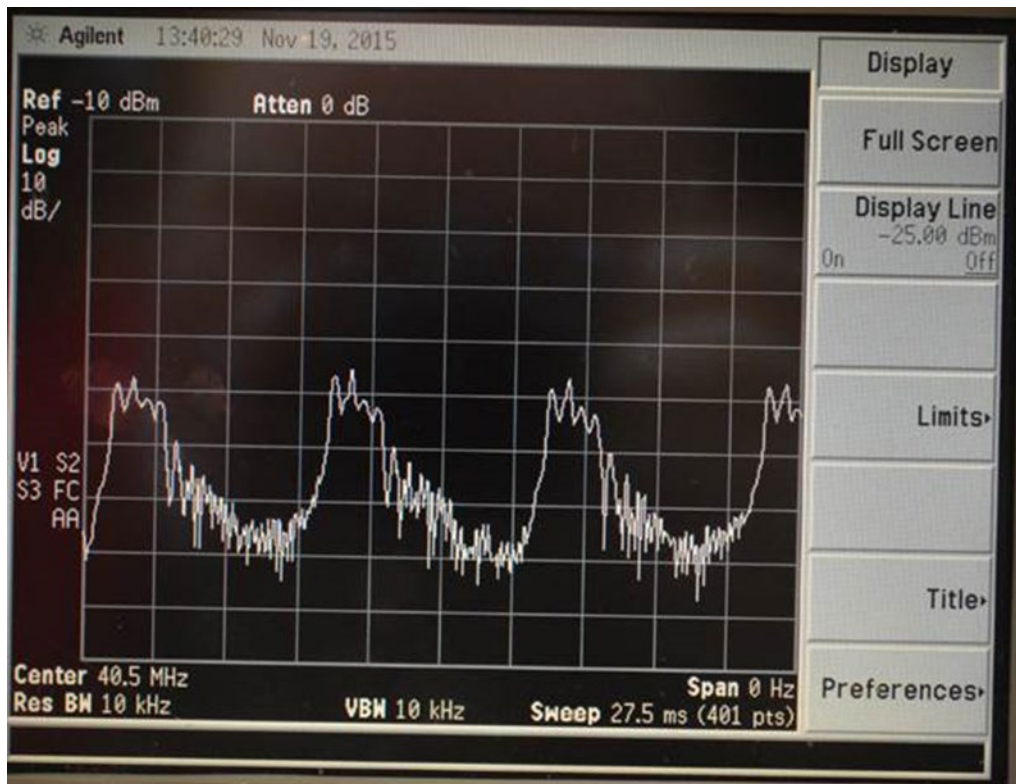


Figure 12 – 120 Hz CPD level vs. time.

In the field the level of a 40.5 MHz CW distortion (difference) signal at was observed to vary at a 120 Hz rate. The 40.5 MHz CPD signal is a distortion product created by mixing 800 MHz and 840.5 MHz CWs.

Figure 13 is a diagram of a seizure mechanism, which is one possible way that 120 Hz modulated CPD can be created. At point “A,” a resistive connection is made from a loose seizure screw, and several amperes of AC current are flowing to power amplifiers. This condition causes an AC voltage difference between the seizure screw and coax center conductor that is proportional to AC current. Nearby, at points “B,” corrosion diodes are present. The corrosion diodes can only actively mix downstream signals while the AC voltage difference is small. Increasing clamping force makes the distortion disappear.




Figure 13 – Shunt resistance-nonlinear distortion

A diagram showing how a loose seizure screw creates a high value shunt resistance, which enables nonlinear distortion to occur.

7.3. Mathematics of nonlinear signals

To evaluate non-linear behavior up to a third order we represent the composite input E_{in} as:

$$E_{in} = \sin(\omega_1 t) + \sin(\omega_2 t) = \frac{e^{i\omega_1 t} - e^{-i\omega_1 t}}{2} + \frac{e^{i\omega_2 t} - e^{-i\omega_2 t}}{2}$$

where $\omega_1 = 2\pi f_1$ and $\omega_2 = 2\pi f_2$.

The composite output E_{out} is defined in equation 1 as:

$$E_{out} = A \cdot E_{in} + B \cdot E_{in}^2 + C \cdot E_{in}^3$$

Expanding the composite input signal

$$A \cdot E_{in} = A(\sin(\omega_1 t) + \sin(\omega_2 t)) = \frac{A}{2}(e^{i\omega_1 t} - e^{-i\omega_1 t} + e^{i\omega_2 t} - e^{-i\omega_2 t})$$

Expanding the square of the composite input signal

$$B \cdot E_{in}^2 = B(\sin(\omega_1 t) + \sin(\omega_2 t))^2 = B((\sin(\omega_1 t))^2 + 2 \sin(\omega_1 t) \sin(\omega_2 t) + (\sin(\omega_2 t))^2)$$

$$4E_{in}^2 = (e^{i\omega_1 t} - e^{-i\omega_1 t})^2 + 2(e^{i\omega_1 t} - e^{-i\omega_1 t})(e^{i\omega_2 t} - e^{-i\omega_2 t}) + (e^{i\omega_2 t} - e^{-i\omega_2 t})^2$$

$$4E_{in}^2 = (e^{i2\omega_1 t} - 2 + e^{-i2\omega_1 t}) + 2(e^{i\omega_2 t} e^{i\omega_1 t} - e^{i\omega_2 t} e^{-i\omega_1 t} - e^{-i\omega_2 t} e^{i\omega_1 t} + e^{-i\omega_2 t} e^{-i\omega_1 t}) + (e^{i2\omega_2 t} - 2 + e^{-i2\omega_2 t})$$

$$4E_{in}^2 = (e^{i2\omega_1 t} - 2 + e^{-i2\omega_1 t}) + 2(e^{i(\omega_1+\omega_2)t} - e^{-i(\omega_1-\omega_2)t} - e^{i(\omega_1-\omega_2)t} + e^{-i(\omega_1+\omega_2)t}) \\ + (e^{i2\omega_2 t} - 2 + e^{-i2\omega_2 t})$$

Expanding the cube of the composite input signal

$$C \cdot E_{in}^3 = C(\sin(\omega_1 t) + \sin(\omega_2 t))^3$$

$$E_{in}^3 = (\sin(\omega_1 t))^3 + 3(\sin(\omega_1 t))^2(\sin(\omega_2 t)) + 3(\sin(\omega_1 t))(\sin(\omega_2 t))^2 + (\sin(\omega_2 t))^3$$

$$8 \cdot E_{in}^3 = (e^{i\omega_1 t} - e^{-i\omega_1 t})^3 + 3(e^{i\omega_1 t} - e^{-i\omega_1 t})^2(e^{i\omega_2 t} - e^{-i\omega_2 t}) \\ + 3(e^{i\omega_1 t} - e^{-i\omega_1 t})(e^{i\omega_2 t} - e^{-i\omega_2 t})^2 + (e^{i\omega_2 t} - e^{-i\omega_2 t})^3$$

$$8 \cdot E_{in}^3 = e^{i3\omega_1 t} - 3e^{i\omega_1 t} + 3e^{-i\omega_1 t} - e^{-i3\omega_1 t} + 3(e^{i2\omega_1 t} - 2 + e^{-i2\omega_1 t})(e^{i\omega_2 t} - e^{-i\omega_2 t}) \\ + 3(e^{i\omega_1 t} - e^{-i\omega_1 t})(e^{i2\omega_2 t} - 2 + e^{-i2\omega_2 t}) + e^{i3\omega_2 t} - 3e^{i\omega_2 t} + 3e^{-i\omega_2 t} - e^{-i3\omega_2 t}$$

$$8 \cdot E_{in}^3 = e^{i3\omega_1 t} - 3e^{i\omega_1 t} + 3e^{-i\omega_1 t} - e^{-i3\omega_1 t} \\ + 3(e^{i\omega_2 t}e^{i2\omega_1 t} - 2e^{i\omega_2 t} + e^{i\omega_2 t}e^{-i2\omega_1 t} - e^{-i\omega_2 t}e^{i2\omega_1 t} + 2e^{-i\omega_2 t} - e^{-i\omega_2 t}e^{-i2\omega_1 t}) \\ + 3(e^{i\omega_1 t}e^{i2\omega_2 t} - 2e^{i\omega_1 t} + e^{i\omega_1 t}e^{-i2\omega_2 t} - e^{-i\omega_1 t}e^{i2\omega_2 t} + 2e^{-i\omega_1 t} - e^{-i\omega_1 t}e^{-i2\omega_2 t}) \\ + e^{i3\omega_2 t} - 3e^{i\omega_2 t} + 3e^{-i\omega_2 t} - e^{-i3\omega_2 t}$$

$$8 \cdot E_{in}^3 = e^{i3\omega_1 t} - 3e^{i\omega_1 t} + 3e^{-i\omega_1 t} - e^{-i3\omega_1 t} \\ + (3e^{i(2\omega_1+\omega_2)t} - 6e^{i\omega_2 t} + 3e^{-i(2\omega_1-\omega_2)t} - 3e^{i(2\omega_1-\omega_2)t} + 6e^{-i\omega_2 t} - 3e^{-i(2\omega_1+\omega_2)t}) \\ + (3e^{i(2\omega_2+\omega_1)t} - 6e^{i\omega_1 t} + 3e^{-i(2\omega_2-\omega_1)t} - 3e^{i(2\omega_2-\omega_1)t} + 6e^{-i\omega_1 t} - 3e^{-i(2\omega_2+\omega_1)t}) \\ + e^{i3\omega_2 t} - 3e^{i\omega_2 t} + 3e^{-i\omega_2 t} - e^{-i3\omega_2 t}$$

Therefore, by expressing ω as $2\pi f$ and analyzing the first second and third order distortion, we have that the frequency components in E_{out} are:

$f_1, f_2, 2f_1, 2f_2, 3f_1, 3f_2, f_1+f_2, f_1-f_2, f_2-f_1, 2f_1-f_2, 2f_2-f_1, 2f_1+f_2, 2f_2+f_1$

DOCSIS 4.0 Security: A Comprehensive Guide to Successful Deployments

A Technical Paper prepared for SCTE by

Massimiliano Pala

Director, PKI Architectures
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
+1 303.661.3334
m.pala@cablelabs.com

Doug Jones

Principal Architect, Wired Technologies
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
+1 303.661.3326
d.jones@cablelabs.com

Yuan Tian

Security Engineer, PKI Architectures
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
+1 303.661.3330
y.tian@cablelabs.com

Craig Pratt

Lead Architect, Security & Privacy Technologies
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
+1 303.661.3408
c.pratt@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	3
1.1. DOCSIS 3.0 and The Enhanced Secure Provisioning (ESP).....	3
1.2. DOCSIS 3.1 and The Second Generation DOCSIS [®] PKI (The New PKI).....	3
1.3. Distributed Access Architectures and New Security Needs	6
2. DOCSIS 4.0 Security Principles.....	7
2.1. From DOCSIS 3.1 to DOCSIS 4.0 Security: What Does NOT Change?.....	9
2.2. From DOCSIS 3.1 to DOCSIS 4.0 Security: What Changes?	10
2.3. A New Authentication Mode (BPI+ V2).....	13
2.4. Certificate Revocations Updates.....	14
2.4.1. Enabling Mutual Authentication without Revocation Checking.....	16
2.4.2. Enabling Mutual Authentication with CM's Certificate (client-side) Revocation Checking	17
2.4.3. Enabling Mutual Authentication with CMTS' Certificate (server-side) Revocation Checking	17
2.4.4. Enabling Mutual Authentication with Mutual Revocation Checking.....	18
3. Deployment Examples	18
3.1. Preparing Your Networks for D4.0 CMs.....	19
3.2. Upgrading Speeds, Not Security.....	20
3.3. Enabling Advanced Security Features With BPI+ V2	20
3.4. Enabling Revocation Checking and DOCSIS 4.0	20
4. Acknowledgments	21
5. Conclusion.....	21
Abbreviations	21
Bibliography & References.....	23

List of Figures

Title	Page Number
Figure 1 - BPI+ Authentication Messages (V1 and V2).....	7
Figure 2 - BPI+ V2 authentication process	13
Figure 3 - Integrated BPI+ and Revocation Checking Flow.....	15
Figure 4 - Example Deployment with revocation checking support and local overrides	18

List of Tables

Title	Page Number
Table 1 - DOCSIS Security Evolution (1.0-3.1).....	5
Table 2 - Object Identifiers for ECU enabled functionalities in DOCSIS 4.0.....	11

1. Introduction

DOCSIS[®] 4.0 security introduces several important enhancements when compared to previous generations of the protocol [SECv4.0]. To better understand the impact and use of DOCSIS 4.0 new features and how they relate to today's deployments and practices, let's start from reviewing the history of DOCSIS security and its evolution.

The first version of the Data Over Cable Service Interface Specification or DOCSIS[®] was released in 1997. The document specified the first standard approach to providing Internet access to subscribers over a cable operator's shared-access Hybrid Fiber-Coaxial (HFC) network (i.e., cable network).

The initial DOCSIS security architecture supported two major schemes: the Baseline Privacy Interface (BPI) and the Full Security (FS), a Security System with a removable security module. These two schemes specified the requirements to implement DOCSIS' two main security goals of protecting users and operators from data privacy issues and theft-of-service. The DOCSIS 1.0 specification eventually dropped FS due to a lack of support from the community. DOCSIS 1.1 strengthened BPI with its implementation of BPI+, which later evolved into the DOCSIS Security Specification in DOCSIS 3.0 and 3.1.

1.1. DOCSIS 3.0 and The Enhanced Secure Provisioning (ESP)

The DOCSIS 3.0 Security Specification [SECv3.0] introduced several new components that built upon the BPI+ (V1) Specification such as the Enhanced Secure Provisioning (ESP), extended support for Revocation Status Checking, and PKI Updates.

The Enhanced Secure Provisioning, or ESP, refers to securing the cable operator's operational support systems (e.g., the CM provisioning process, including Dynamic Host Configuration Protocol (DHCP), Time of Day (ToD), and TFTP). Securing these processes played a critical role in protecting the CM and the cable network from unauthorized access and theft-of-service attacks that cable operators were experiencing. Specifically, it prevented hacked modems from requesting unauthorized services.

DOCSIS 3.0 Security Specification also introduced a very important tool: certificate revocation status checking. Specifically, DOCSIS 3.0 supported two standard methods of certificate revocation: Certificate Revocation Lists (CRLs) and Online Certificate Status Protocol (OCSP). DOCSIS 3.0-compliant CMTS-es support configuration of none, one, or both certificate revocation methods.

On the PKI side, DOCSIS 3.0 Security leverages the same infrastructure (PKI) and the same authentication protocol that was used in the previous version of DOCSIS, but it introduces a new CVC infrastructure that supports a three-tier certificate chain architecture (i.e., Root, Intermediate, End-Entity). The support for the new CVC SubCAs was actually employed later in DOCSIS 3.1 security specification as part of the 2nd Gen DOCSIS[®] PKI, namely the "New PKI".

1.2. DOCSIS 3.1 and The Second Generation DOCSIS[®] PKI (The New PKI)

The new DOCSIS 3.1 Security Specification [SECv3.1] was introduced to strengthen the cryptographic parameters used during authentication. While DOCSIS 3.1 retains all the DOCSIS 3.0 Security features, it upgrades the PKI to a new infrastructure (2nd Gen DOCSIS[®] PKI) and increased the allowed size of cryptographic keys for both authentication and encryption.

On the PKI side, for DOCSIS 3.1, CableLabs defined an entirely new PKI hierarchy. The legacy PKI was halfway through its 30-year lifecycle, and it was starting to show its age with key sizes that needed

increasing and Hash Algorithms (i.e., SHA-1) whose use was being deprecated by NIST due to discovered weaknesses in the algorithm. To address this problem, the new PKI introduced the use of SHA-256 for digital signatures.

The overview of the security features of DOCSIS (1.0-3.1) are summarized in Table 1.

Table 1 - DOCSIS Security Evolution (1.0-3.1)

DOCSIS 1.0-3.1 Security Overview					
		DOCSIS 1.0	DOCSIS 1.1 & 2.0	DOCSIS 3.0	DOCSIS 3.1
Baseline Privacy Key Management (BPKM)	CM Public Key	768-bit key, 1024-bit key	768-bit, 1024-bit	768-bit, 1024-bit	768-bit, 1024-bit, 2048-bit key
	Authorization Key (AK)	64-bit key	160-bit key	160-bit key	160-bit key
	AK Generation	Random key generated by the CMTS and sent to the CM			
	AK Encryption	RSAES-PKCS1-v1_5	RSAES-OAEP	RSAES-OAEP	RSAES-OAEP
	KEK	64-bit Key	128-bit key	128-bit key	128-bit key
	KEK Generation	Both CM and CMTS derive KEKs from a function using the Authorization Key and the SHA-1 Hash Algorithm.			
	TEK	64- bit key	64- bit key	64- bit key, 128-bit key	64- bit key, 128-bit key
	TEK Generation	Random key generated by the CMTS and sent to the CM			
	TEK Encryption	56-bit DES	56-bit DES, 112-bit DES	56-bit DES, 112-bit DES	56-bit DES, 112-bit DES
	Message Authentication Key (MAK)	160-bit key	160-bit key	160-bit key	160-bit key
	MAK Generation	Both CM and CMTS derive MAKs from a function using the Authorization Key and the SHA-1 Hash Algorithm.			
	Hash Algorithm	SHA-1	SHA-1	SHA-1	SHA-1
MIC	MIC	HMAC-MD5	HMAC-MD5	MMH-MAC	MMH-MAC
Data	Traffic Encryption	56-bit DES, 40-bit DES	56-bit DES, 40-bit DES	56-bit DES, 40-bit DES, 128-bit AES	56-bit DES, 40-bit DES, 128-bit AES
Auth	CM Authentication	MAC Address	X.509v3 RSA certificate	X.509v3 RSA certificate	X.509v3 RSA certificate
Code	SSD	Proprietary	1024-bit CVC, 1536-bit CVC, 2048-bit CVC	1024-bit CVC, 1536-bit CVC, 2048-bit CVC	1024-bit CVC, 1536-bit CVC, 2048-bit CVC
ESP	ESP	No	No	Yes	Yes

1.3. Distributed Access Architectures and New Security Needs

Since DOCSIS 4.0 relies on the use of distributed architectures to deliver increased speeds both downstream and upstream, it is important to assess and understand the boundaries of this new attack surface and the associated threats.

Distributed Access Architecture (DAA) is an evolved cable network architecture that decentralized the headend network functions by moving the PHY and/or MAC layer functions to the remote node of the access network, while other functions remain in the Converged Cable Access Platform (CCAP). Currently, there are three types of DAA networks: (1) the Remote PHY or R-PHY (2) Remote MACPHY architecture or R-MACPHY, and (3) Flexible MAC Architecture or FMA.

In R-PHY and R-MACPHY, the MSO network is split into “trusted” and “untrusted” domains. The RPD/RMD, coax access network and edge devices (e.g., CMs) are in the “untrusted” domain (i.e., either in the customer’s premises or deployed in the field). On the contrary, the CCAP core, authentication server and provisioning servers are located in the “trusted” domain that is usually under strict MSO physical control and where only authorized employees are allowed access (e.g., NetOps and NetSecOps).

With the introduction of FMA, virtualization and containerization techniques are used to ensure service availability and to optimize resource allocation for enhanced services capability. In FMA, the security boundary between trusted and untrusted domains is even more blurred, and, therefore, the need to verify the identity of all elements that are present in the network is critical and necessary. This approach to network security is referred to as Zero Trust Security. Operators may find themselves facing difficult choices when considering tradeoffs between security and service availability: strict revocation checking, frequent software updates, or unnecessarily short certificates’ validity periods can cause unexpected service availability issues, while loose security controls and exceptionally long certificate validity periods can put customer and organization assets under the potential risk of compromise.

The original threat model used to design the DOCSIS security protocol did not incorporate or anticipate the kinds of security issues introduced by distributed architectures and without specific hardware protections (e.g., secure key storage), it’s still possible to maliciously modify or replace device software and credentials, which might enable attacks such as Modem Cloning or Service Uncapping. For example, attackers could exfiltrate the device’s credentials or install malware (e.g., backdoors, bot agents, etc.) to perform active attacks such as ***Denial-of-service*** or ***Man-in-the-middle*** that can have very disruptive effects for the operator’s network.

To mitigate the possibility to carry out these new class of attacks, DOCSIS 4.0 introduced two new security controls that are meant to work together: device physical security and network identities.

On the physical security side, the new requirements are described in section 15 of [SECv4.0] and mandate for increased security of stored secrets (keys) (see Section 15.1 of [SECv4.0]) and the use of secure boot processes for CMs (see Section 15.2 of [SECv4.0]). The use of such measures is aimed at reducing the risk of unexpected/malicious changes in the software running on CMs.

When it comes to network identities, before DOCSIS 4.0, only the CMTS could verify the Cable Modem’s certificate, not the other way around. In fact, since no verifiable identity is used on the server side (CMTS) in BPI+ V1, an attacker may be able to intercept CMTS functionalities and redirect the messages to its own device or service by targeting, for example, fielded RPDs or RMDs. In this scenario, it is easy to show how an attacker could completely take over large amount of internet connections via malicious network configurations (i.e., DHCP, DNS, etc.) and services (i.e., Web, Mail, etc.). Similarly,

passive attacks are also possible where the customer's session is not actively modified but it is just monitored for "interesting" data.

As discussed in great detail in the next section, DOCSIS 4.0 supports a new version of the authentication protocol, namely BPI+ V2, that greatly reduces the possibility to carry out such attacks by requiring both the CM and the CMTS to verify each other identities before establishing a connection.

2. DOCSIS 4.0 Security Principles

One of the challenges faced in designing DOCSIS 4.0 was how to integrate the needed new security features such as Mutual Authentication or Perfect Forward Secrecy in a minimally disruptive fashion. The answer to this challenge was twofold: (1) provide support for the same BPI+ (V1) authentication protocol in use in previous version of DOCSIS (1.1-3.1) and, (2) introduce a new version of BPI+ (V2) that encapsulate the needed security enhancements.




Figure 1 - BPI+ Authentication Messages (V1 and V2)

Figure 1 provides an overview of the BPI+ authentication process (common across BPI+ V1 and V2) where the Auth Request and Auth Reply messages are used to establish a common encryption key (i.e., AES) between the CM and the CMTS.

To better understand the differences between the security features supported in the two versions of BPI+, let's compare their security features and how they relate to and solve the new threat model.

Message Authentication and Integrity. This security principle is related to protecting the integrity and providing verifiable origin information for exchanged messages during authentication. In BPI+ V1, neither the CM nor the CMTS authenticate (sign) the messages they generate. In BPI+ V2, the CMTS uses the public key of the CM's Device certificate to encrypt the authorization key for the destination CM. The lack of authentication is the source of many possible security vulnerabilities that can lead to very disruptive attacks (i.e., modifying exchanged messages, spoofing device identities, etc.). In BPI+ V2, both the CM and the CMTS are required to sign outgoing authentication messages sent to the other party and authenticating their peer's messages before processing them.

Perfect Forward Secrecy or PFS. This security principle is related to protecting data across sessions. Specifically, when PFS is supported, its use protects against the decryption of pre-recorded data even when the target Cable Modem (CM) private key has been compromised. In BPI+ V1, the session encryption key is directly encrypted under the public key of the CM certificate and an attacker can decrypt any prerecorded sessions between the CM and the CMTS by compromising the certificate's private key at some date in the future. In BPI+ V2 the encryption key is negotiated when establishing a new session by using a Diffie-Hellman key exchange over finite field or Elliptic curves. This prevents an attacker from decrypting previously recorded traffic that was protected under BPI+ V2 even if/when an attacker compromises the CM's private key.

Mutual Authentication or MA. This security principle is related to guaranteeing that no malicious entity is able to impersonate the other party when establishing/authenticating the communication session between the CM and the CMTS. This type of attack is usually referred to as a "Man-In-The-Middle" (MITM) attack and requires the ability to manipulate the traffic on the network. Because BPI+ V1 does not secure the CMTS (or network) identity (it does not provide certificate during connection initiation), BPI+ V1 is vulnerable to such attacks. Malicious actors could impersonate a CMTS without the CM being able to distinguish between the real CMTS and the attacker. In BPI+ V2, the CMTS provides a certificate to prove its identity and the CM can, in this case, easily identify the real CMTS by validating the CMTS' message signature and certificate. The incorporation of this CMTS certificate-based identity allows the CM to properly authenticate the CMTS via a digitally signed Auth Reply, thus preventing any non-authenticated or modified messages to be ignored.

Algorithm Agility. This security principle is related to the possibility to execute the authentication protocol independently from the underlying cryptographic algorithms selected for proving identities (e.g., RSA, Falcon, Dilithium, etc.). In BPI+ V1, as mentioned when describing PFS, the CMTS directly uses the RSA public key associated with the CM's certificate to secure (encrypt) the authorization. In BPI+ V2 the authorization exchange enables a variety of methods to determine an authorization key. This mechanism provides a path for enabling post-quantum-safe cryptography and classic/post-quantum hybrid identity certificates.

Increased size of encryption keys. This security principle is related to the normal evolution of cryptographic algorithms over time where it is understood that larger keys are needed to keep the same level of confidentiality. In BPI+ V1, although there are few options when it comes to negotiating line-rate symmetric ciphers (encryption algorithms), AES-128 is the only supported option that is considered secure according to today's best practices. DOCSIS 4.0 introduces support for negotiating larger key sizes (AES-256) to encrypt user-data traffic that aligns with current best practice and provides the same level of protection against quantum attacks that users enjoy today against classic ones.

Downgrade Protection. This security principle is related to protecting against malicious actors trying to negotiate a more vulnerable version of a protocol when multiple versions are supported. This is a very difficult problem to solve that generally affects Access Networks architectures such 3GPP networks. Since DOCSIS 4.0 supports two different versions of BPI+, without providing any protection, DOCSIS 4.0 could suffer from vulnerabilities similar to ones observed in mobile networks. To address this issue, DOCSIS 4.0 introduces the concept of Trust on First Use (TOFU) that requires a CM to store the minimum allowed version of the BPI+ protocol for subsequent authentications, as indicated by the CMTS, in a secure memory location on the CM after successful authentication. For example, a CM that authenticates with BPI+ V2 can be signaled by the CMTS to only communicate using BPI+ V2 for subsequent sessions – preventing an imposter CMTS from downgrading to BPI+ V1 on a subsequent connection.

In the rest of the section, to better understand the differences and similarities of DOCSIS 4.0 and 3.1, we provide a summary about what changes and what does not change when it comes to deployments.

2.1. From DOCSIS 3.1 to DOCSIS 4.0 Security: What Does NOT Change?

DOCSIS 4.0 supports two versions of the BPI+ authentication protocol: BPI+ V1 and BPI+ V2. While the latest version of the protocol (BPI+ V2) is only available in DOCSIS 4.0 mode, the first version of the protocol, i.e., BPI+ V1, is shared across almost all versions of DOCSIS, including DOCSIS 4.0 (1.1-4.0). A first advantage for retaining support for BPI+ V1 in DOCSIS 4.0 is the possibility to deploy new DOCSIS 4.0 devices without the need for updating existing procedures, thus reducing the extra overhead of deploying new technologies that might come with new requirements. A second advantage is related to the fact that since in BPI+ V1 the CMTS does not need a device certificate, deploying DOCSIS 4.0 with BPI+ V1 does not introduce new requirements (for security) when compared to previous versions of DOCSIS such as DOCSIS 3.0 or DOCSIS 3.1. However, as discussed previously, with the introduction of distributed nodes outside the operator's trusted domain, the attack surface has increased. When enhanced authentications are needed, BPI+ V2 can be enabled to address the new security risks (see Section 2.3 for more details).

From the *Secure Software Download (SSD) standpoint*, DOCSIS 4.0 design applies the same principle that was used for the authentication protocol: keep support for current procedures and provide the possibility for upgrading to a more efficient one when needed. In fact, in DOCSIS 4.0 *there are two different mechanisms that can be used to enable SSD on a device*. The first one is the same leveraged in DOCSIS 3.1 where the use of a Manufacturer's CVC and/or an Operator's co-signer CVC certificate(s), in the config file (or via SNMP SET), triggers SSD procedures. The second mechanism that directly use the Firmware Authentication Header (FWAH) is detailed in the next section.

Another important similarity between DOCSIS 3.1 and DOCSIS 4.0 is *the use of the same PKI*. Differently from the previous DOCSIS update (DOCSIS 3.0 → DOCSIS 3.1), DOCSIS 4.0 uses the same Root of Trust that is used in DOCSIS 3.1. This means that DOCSIS 4.0 can use the same procedures and Trust Anchor used in DOCSIS 3.1 to validate device certificates.

When looking at the *backward compatibility with previous versions of DOCSIS*, it is important to understand what type of certificates might be needed. Specifically, while DOCSIS 4.0 devices use a single certificate to connect to both DOCSIS 4.0 and DOCSIS 3.1 CMTS¹ by using a *Common Cable Modem* certificate profile, they will still need a DOCSIS 3.0 CM certificate to be able to connect to pre-3.1 CMTS. The reason for this is that DOCSIS 3.0 and DOCSIS 4.0 do not share the same Root of Trust and, therefore, separate certificates are still needed, exactly as for DOCSIS 3.1 devices.

Support for revocation also remains untouched from D3.1 and D3.0 specifications. In fact, besides the use of CMTS-related revocation information when BPI+ V2 is enabled, the validation of CM certificates is delegated to the CMTS. During this process the CMTS can use CRLs or OCSPs to check the status of the CM's certificate and take proper action about it (or just report it). Support for revocation status checking for the network (CMTS) certificate is discussed in the next section.

What about EAE? DOCSIS 4.0 support for EAE is backward compatible with previous versions of DOCSIS where the CMTS support for EAE is advertised via TLVs in downstream MDD messages. In DOCSIS 4.0, the usual TLV Type 6 that is used to advertise support for EAE for BPI+ V1 is joined by the

¹ DOCSIS 4.0 certificates are different from DOCSIS 3.1 ones, and a software upgrade might be needed for D4.0 devices to be supported by DOCSIS 3.1 CMTS.

new TLV (Type 23) that allows to specify additional options for BPI+ V1 and BPI+ V2 as described in the next section.

In summary, the security features and operations that remain common to both D3.1 and D4.0 are:

- DOCSIS 4.0 can use the same authentication protocol in use for DOCSIS 3.1 (BPI+ V1)
- DOCSIS 4.0 can use the same PKI in use for DOCSIS 3.1 (2nd Gen DOCSIS[®] PKI)
- DOCSIS 4.0 can use the same SSD procedures in use for DOCSIS 3.1
- DOCSIS 4.0 devices require, exactly as D3.1 devices, an additional certificate from the 1st Gen DOCSIS[®] PKI (“legacy” PKI) to authenticate in DOCSIS 3.0 networks (if the device supports D3.0 environments).
- DOCSIS 4.0 supports the same revocation options available in DOCSIS 3.1 and DOCSIS 3.0.

In the next section we look at the aspects that have changed and their impact on DOCSIS 4.0 deployments.

2.2. From DOCSIS 3.1 to DOCSIS 4.0 Security: What Changes?

While support for BPI+ V1 is shared across almost all versions of DOCSIS (1.1-4.0), DOCSIS 4.0 is the first version of DOCSIS to support BPI+ V2 that delivers new and enhanced security controls. Since BPI+ V2 introduces the most profound changes in DOCSIS security since the introduction of digital certificates in DOCSIS 1.1, a detailed description is provided in the next section while here we focus on the rest of the differences with DOCSIS 3.1.

On the certificate side, there are four important changes that need to be discussed.

First, although DOCSIS 4.0 uses the same PKI as DOCSIS 3.1, the contents of CM certificates for the two environments are different. In fact, while DOCSIS 4.0 uses the same algorithm (RSA) and key sizes (2048 bit) already in use in DOCSIS 3.1, certificates for D4.0 CMs are larger than their DOCSIS 3.1 counterpart because of the introduction of several standard extensions that deliver new security controls. The first change to notice is the use of the Authority Information Access (AIA) extension to carry the location of the authoritative OCSP server that the CMTS can use to check for the revocation status of certificates. This change fixes the status of revocation in D3.1 and D3.0 where the absence of such data makes revocation checking very hard in practice. Another important change in the certificate profile is related to the introduction of a new concept in DOCSIS: **roles or functions**. Indeed, D4.0 certificates use a well identified set of Object Identifiers or OIDs inside the Extended Key Usage (EKU) extension that allow the verifier to check, based on the presence of specific values in the extension, if the connecting device is authorized (or not) to provide specific services. DOCSIS 4.0 defines two OIDs to identify CMTS functionality (`svcCMTS`) and CM functionality (`svcCM`) respectively. For example, when a CM is validating the CMTS certificate, it will look for the `svcCM` value in the EKU extension of the certificate and, if not found, rejects the connection if the `svcCMTS` value is not present in the certificate (i.e., the device was not authorized to provide CMTS services). In other words, in DOCSIS 4.0 not all certificates are created equal to prevent attacks where a legit certificate (e.g., a CM certificate) would be used to impersonate different roles (e.g., a CMTS). Table 2 provides the list of the EKU values supported in DOCSIS 4.0.

Second, before the operator enables BPI+ V2, the CMTS needs to be provisioned with a DOCSIS certificate. As we discussed earlier, the availability of CMTS credentials enables the CMTS to authenticate its own messages and, therefore, turning on BPI+ V2 not only lowers the risks of network compromise, but it also enables the possibility for Trusted Services from the network enabled by the

possibility to validate, and subsequently trust, the CMTS identity. The provisioning and management of CMTS certificates (i.e., renewal and installation) is a very new process for the DOCSIS community and it is expected to require dedicated support and/or automation. CMTS certificates have a validity of up to five years.

Four, DOCSIS 4.0 CMs use a single certificate to authenticate in both DOCSIS 4.0 and DOCSIS 3.1 modes. This means that D4.0 devices connecting to existing D3.1 networks will use certificates that although compatible with the DOCSIS 3.1 environment (i.e., same algorithm and key sizes), they may be larger in size than the D3.1 ones because they contain new standard extensions that are not present in D3.1 certificates.

Table 2 - Object Identifiers for EKU enabled functionalities in DOCSIS 4.0

Short Name	Name	Value	Description
svcCMTS	id-cl-pki-eku-CMTS	1.3.6.1.4.1.4491.2021.2.1.1	CMTS functionalities
svcCM	id-cl-pki-eku-CM	1.3.6.1.4.1.4491.2021.2.1.2	CM functionalities

On the CVC side there are some important changes. Differently from the DOCSIS 3.1 environment, DOCSIS 4.0 does not support CVCs from the 1st Gen PKI (or “legacy PKI”) when it comes to Secure Software Download and Firmware signing. Because the validation of “legacy” CVCs does not require to check expiration times against the current time, retaining support for “legacy” CVCs would introduce the possibility to modify the code loaded onto D4.0 devices with CVCs that use outdated (and possibly weak) cryptographic parameters such as 1024-bit keys and the SHA-1 algorithm for signatures (any time in the future). To eliminate the security risk, differently from D3.1, DOCSIS 4.0 devices only support CVCs from the 2nd Gen DOCSIS[®] PKI (i.e., the “modern” PKI). Additionally, a new mechanism to initiate the SSD process has been introduced that optimized validations of the firmware by directly supporting the use of the Firmware Authentication Header (FWAH) in config files or via SNMP SET (see Section 14.2 of [SECv4.0]).

On the revocation side there are also some changes. Although when BPI+ V1 is used there are no changes in how revocation works since D3.0, when BPI+ V2 is enabled and revocation checking is desired there is a new certificate to be validated, the CMTS one. Indeed, in BPI+ V2, the CMTS MUST transmit the OCSP response related to the status of its own certificate in the Auth Reply or Auth Reject message. The CMTS’s OCSP response can be cached by the CM and the CMTS for its entire validity time, to minimize the load on revocation infrastructures.

In the previous section we mentioned **some changes on the EAE side**. Although support for TLV Type 6, i.e., the EAE Enabled/Disable TLV (see Section 6.4.28.1.6 of [MULPIv4.0]), is maintained for backward compatibility, DOCSIS 4.0 devices use the new TLV of Type 23, i.e. BPI+ Supported Version and Configuration (see Section 6.4.28.1.22 of [MULPIv4.0]), to discover which BPI+ versions are enabled on the CMTS and what services are available with each one. This new TLV is a compound TLV that uses two bytes to indicate (a) the enabled version of BPI (i.e., 1-byte integer value), and (b) the associated enabled features (i.e., bitmask where Bit 7, when set to one, indicates EAE support). Multiple TLV 23 can be used to announce the enabled features for each BPI+ version that is enabled on the CMTS.

On the BPKM layer side, DOCSIS 4.0 introduces an important update: enabling fragmentation support to extend the supported maximum size of BPKM payloads to ~28 Kb. This change is implemented by introducing two new MMM messages (i.e., the BPKM-REQ5 and the BPKM-RSP5) in [MULPIv4.0] that leverage MMM V5 (instead of V1 as used in BPI+ V1) to support large BPKM

messages that may span more than one frame (i.e., up to 16 fragments). This solution not only enables the use of extra cryptographic material in the Auth Request / Auth Reply process without the need to add new messages and states in the State Machine, but it also opens up future paths for the deployment of new cryptographic algorithms such as Kyber (for key exchange) and Dilithium (for public/private keys), or even hybrid approaches that combine RSA with new types of algorithms [Pala21].

In summary, the security features and operations that change between D3.1 and D4.0 are:

- DOCSIS 4.0 supports multiple versions of BPI+ (i.e., BPI+ V1 and BPI+ V2)
- DOCSIS 4.0 can use advanced authentication features when enabling BPI+ V2
- DOCSIS 4.0 devices use certificates that are larger in size than D3.1 certificates
- DOCSIS 4.0 can use updated SSD procedures that optimizes early error detection
- DOCSIS 4.0 introduces two new approaches for delivering SSH access to devices without static secrets on devices
- DOCSIS 4.0 can enable or disable, for each enabled BPI+ version, the use of EAE independently.
- DOCSIS 4.0 introduces fragmentation support for BPI+ V2 messages

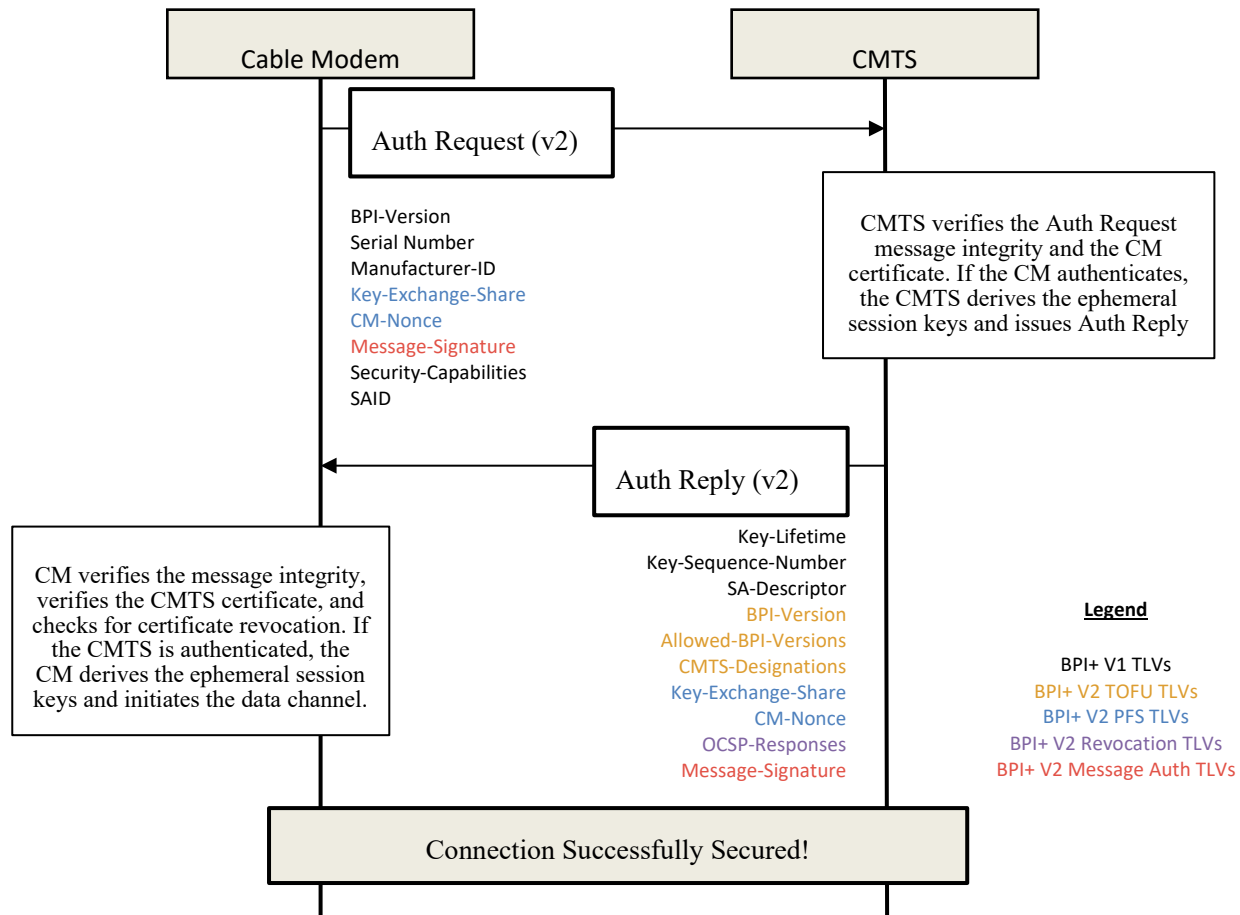


Figure 2 - BPI+ V2 authentication process

2.3. A New Authentication Mode (BPI+ V2)

The new version of the BPI+ authentication protocol supported in DOCSIS 4.0 is called BPI+ V2. Although this new version has the same structure of its predecessor (i.e., it still uses only two (2) messages to establish a secure connection between the CM and the CMTS and there is no change in the state machine), the security properties offered by BPI+ V2 are very different from the ones offered by BPI+ V1. Figure 2 depicts a BPI+ V2 message exchange where the newly defined TLVs are detailed.

A first fundamental difference between BPI+ V1 and V2 is the use of different versions of Mac Manager Messages (MMM) to encapsulate the protocol. In fact, while BPI+ V1 use MMM V1 messages that are limited in size to a single frame, BPI+ V2 defines two new messages, the BPKM-REQ5 and BPKM-RSP5, that leverage version 5 of the MMM headers. The new version of the messages supports payload fragmentation for up to 16 different fragments, thus pushing the maximum supported size for BPKM close to ~28Kb. This change allows for the deployment of larger cryptographic material that may require larger data structures during transport such as post-quantum keys, certificates, and signatures.

A second important feature of BPI+ V2 is the use of digital signatures to authenticate BPKM messages. The Message-Signature TLV, depicted in red color in the figure, carries a DER representation of a Cryptographic Message Syntax (CMS) data structure together with the signer's certificate and stores a

detached signature that is calculated over the entire BPKM message (i.e., the code, length, and payload up to, but excluding, the Message-Signature TLV itself). The use of this TLV implements the Message Authentication and Integrity principle. Moreover, the use of message authentication on both sides of the communication reflects a very important paradigm shift in DOCSIS security where, up to DOCSIS 3.1, the primary focus has been the authentication of Cable Modems only. As the deployment models for DOCSIS have become more distributed, the new BPI+ V2 introduces the support for Mutual Authentication principle via the presence of the Message-Signature TLV in both the Auth Request and Auth Reply (or Auth Reject) messages.

Another important enhancement introduced with BPI+ V2 is the use of the Diffie-Hellman (DH) key exchange mechanism to provide support for Perfect Forward Secrecy (PFS). Indeed, differently from BPI+ V1 where the authorization key is directly encrypted with the RSA public key of the CM's certificate, BPI+ V2 use the Key-Exchange-Share TLV to carry the cryptographic material to derive a common key across the CM and the CMTS. This new key exchange mechanism prevents the decryption of pre-recorded traffic for a CM even when the private key of the CM has been compromised, thus implementing the PFS principle.

The combination of Message Authentication and Mutual Authentication principles (implemented in BPI+ V2 via the Key-Exchange-Share and Message-Signature TLVs), enables a third one, i.e., Algorithm Agility. In fact, because of the separation of the algorithm used for the public key in the certificate and the algorithm used for key exchange, BPI+ V2 is algorithmically agile with respect to the device certificate, thus being able to support not only RSA-based certificates but other classic (e.g., ECDSA), post-quantum (e.g., Dilithium or Falcon), or hybrid (e.g., Composite-Crypto) algorithms.

With the introduction of multiple versions of the BPI+ protocol, a Downgrade Protection mechanism (TOFU) was introduced to provide protection against unauthorized downgrades. To support TOFU, BPI+ V2 uses two TLVs during the authentication process: the BPI-Version and the Allowed-BPI-Versions whose value can be used to manage which version of BPI+ should the device use after an initial successful connection. For example, during a BPI+ V2 authentication, the CMTS can use the Allowed-BPI-Versions TLV with the value of (1) to indicate that the CM can still use BPI+ V1, if needed, for subsequent authentications (i.e., legit downgrades).

In the rest of the section, we focus on the different options available for managing revocation status checking, an important aspect of successful deployments when it comes to security.

2.4. Certificate Revocations Updates

Support for checking the status of revocation for DOCSIS devices has been integrated into the specifications since DOCSIS 3.0 where both CLRs and OCSF checking were introduced to lower security risks associated with providing services to potentially compromised devices.

In DOCSIS 4.0, there are two main changes in the protocol that affect revocation checking procedures: the introduction of the CMTS (or network) identity, and the updating of certificate profiles.

The first change, the introduction of the network identity, required the definition of a new security control that would allow Cable Modems to know when (or not) to demand OCSF responses from the CMTS (i.e., in DOCSIS 4.0, CMs do not perform OCSF queries directly), and when they can ignore it. Since malicious attackers would try to remove checking of revocation information to make it easier to use compromised credentials, DOCSIS 4.0 does not leverage the usual control interface for configuring CM revocation checking requirements (i.e., config file TLVs or SNMP SETs), but the certificates itself.

Indeed, when CM revocation checking is desired, operators should use CMTS certificates that contain the URL of the OCSP responder (i.e., the Authority Info Access extension with the OCSP access method): when the URL is embedded in the certificate the CM understands that revocation status checking is required and will reject messages that do not carry the needed OCSP responses for the CMTS certificate. Vice versa, when revocation checking is not desired on the CM, operators should install CMTS certificates that do not contain the OCSP URL in them: when the URL of the OCSP responder is not embedded in the CMTS certificate, the CM understands that revocation checking is not required and, therefore, OCSP responses are not needed in the CMTS' messages. CMTS certificates that do not carry any OCSP revocation information are referred to as NRI certificates or No-Revocation Information certificates.

In other words, *the protected value, i.e., the extension, inside the certificate is the secure equivalent of configuring revocation checking on CMs via SNMP or configuration file options* since the presence of the extension in the certificate is protected by the CA signature on the device certificate itself (i.e., even the CMTS cannot lie about the requirement).

In this view, CMTS vendors should consider the possibility to support a dual-certificates configuration for their devices: one certificate for when CMTS certificate revocation checking is enabled and one certificate for when CMTS certificate revocation checking is disabled.




Figure 3 - Integrated BPI+ and Revocation Checking Flow

Figure 3 provides an overview of the generic workflow for the BPI+ protocol basic authentication messages when combined together with the revocation information checking process. We will use this figure throughout this section of the paper to discuss the details of the different deployment options available in DOCSIS 4.0.

As depicted in Figure 3, the retrieval of the revocation information can be triggered by new Authorization Requests coming into the CMTS, however it might be a good practice to keep a caching mechanism on the CMTS to allow for faster authentications (and re-authentications) and reduce the number of external

requests the CMTS issues (and rely upon) to complete the authentication. This is particularly true in the context of OCSP processing when a CMTS is rebooting and, in general, for CRL-based revocation checking where usually the validity period of CRLs is quite long (days or weeks). Since DOCSIS 4.0 explicitly allows for caching the revocation information for their entire validity period, it is important to leverage this option on the CMTS to deliver fast and efficient revocation status checking.

Another important consideration is related to the ability of the CMTS to directly connect to resources on the internet (not directly available on the operator's network). Usually, because of security concerns, CMTS-es are not allowed to access any resource on the Internet. When this is the case, the CMTS would not be able to query CRLs repositories or OCSP responders unless some "intermediary" is used. For example, operators can route all requests for revocation information via an HTTP proxy, thus allowing for easier monitoring (and restrictions) for which resources the proxy can access (e.g., only official OCSP repositories). When OCSP validity period and HTTP caching headers carry the same values, the HTTP caching mechanism can be used for both OCSP and CRLs.

However, when HTTP proxy services are not enough and operators would like to take control over the revocation information for the devices deployed in their own network, operators should consider the possibility to deploy local OCSP servers that can provide OCSP responses (locally) and, moreover, override revocation status (locally). The deployment of such services in the operator's network can enable the possibility to manage the revocation status of device certificates locally. For this option to be enabled, the operator needs an OCSP responder certificate from the DOCSIS infrastructure (for each D4.0 issuing CA) that can be used to setup the local service. Efficient open source implementations for in-line OCSP responders are already available (e.g., OpenSSL [OSS] or OpenCA OCSPD [OCAOD]).

In the rest of this section, we examine different deployment scenarios for enabling revocation checking on the server side, on the client side, or both.

2.4.1. Enabling Mutual Authentication without Revocation Checking

The simplest and most common deployment model is the one where there is no support for revocation status checking. This is the default deployment scenario in today's DOCSIS network where revocation checking is disabled or practically very difficult to implement because of the lack of OCSP URLs inside D3.0 and D3.1 certificates.

In this authentication mode, when BPI+ V1 or V2 are executed, the CMTS (in both BPI+ V1 and BPI+ V2) does not download any Certificate Revocation List (or CRL) nor any OCSP Responses to validate the revocation status of the device certificate. This means that authentications only rely on the information presented by the device (or the network) to decide if to allow the connection with the device (or the network). This means that the extra steps $\{1.a, \dots, 1.f\}$ from Figure 3 are not needed and will not be executed in both BPI+ V1 and BPI+ V2.

To achieve this configuration, the CMTS must be first configured to disable CRL and OCSP response (both). Changing the CMTS configuration is sufficient when only BPI+ V1 is enabled. However, to correctly handle the BPI+ V2 case, the CMTS must also be provisioned with a CMTS NRI certificate: the absence of the OCSP URL inside the certificate is used as the security control to communicate to the CM that no revocation checking is needed on this certificate when executing BPI+ V2 only (i.e., BPI+ V1 does not use any network identity and, therefore, there is no CMTS certificate to validate).

2.4.2. Enabling Mutual Authentication with CM's Certificate (client-side) Revocation Checking

Similarly, to D3.0 and D3.1, DOCSIS 4.0 supports checking the revocation status of connected devices when using either BPI+ V1 or BPI+ V2. Similarly, to the previous case, only when executing BPI+ V2 the CMTS NRI certificate is needed to be installed on the CMTS. In this configuration, OCSP responses are not sent inline from the CMTS to the CM during BPI+ authentications (i.e., because NRI certificates do not carry the URL of the OCSP responder), thus limiting the dependencies on the availability of revocation information to the CMTS only.

Since the CMTS is the ultimate controller that can allow or reject a CM during authentication, CMTS vendors have the possibility to implement different authorization policies can be enabled to better accommodate NetOps needs. For example, CMTS vendors could provide the possibility to have strict or permissive policies for allowing devices on the network only after passing revocation checking (strict policy) or allowing them even when revoked and, in that case, report it for monitoring or investigative purposes (permissive policy). Although a strict revocation is the most secure option (e.g., not allowing compromised devices to access any service), permissive policies might be implemented to monitor for potentially compromised or otherwise misbehaving devices. This approach allows for decoupling the decision to provision services to devices from the certificate revocation status (i.e., the revoked status becomes a factor in the decision, not the decision itself).

In this case, although both OCSP and CRL mechanisms can be used for checking the status of devices, it might be more efficient to enable CRL-based validation since a single CRL carries all relevant revocation information for all the certificates issued from the CA while OCSP responses are related to a single certificate entry. When using OCSP, individual request/response roundtrips have to be used for each CM the CMTS needs to validate the revocation status for, while a single CRL can be used to lookup the revocation status of all certificates issued from a CA (with the downside of being quite large if many certificates have been revoked).

2.4.3. Enabling Mutual Authentication with CMTS' Certificate (server-side) Revocation Checking

On the opposite side of the spectrum, this deployment model enables the revocation status checking for the CMTS certificate to authenticate the network. This use-case, because it involves the CMTS certificate, it is relevant only for BPI+ V2 authentications.

To achieve this configuration, operators must disable revocation checking on the CMTS, both CRL-based and OCSP-based mechanisms for the client side. By disabling revocation checking on the CMTS, the CMTS will not execute steps {1.a, ..., 1.f} from Figure 3 to validate the CM's certificate.

However, the CMTS still need to procure the OCSP response for its own certificate and send it to the CM in Auth Reply messages during BPI+ V2. Therefore, in this case, the CMTS still needs to execute steps {1.d, 1.e, 1.f} to be able to retrieve the OCSP response from the server (if not cached).

Differently from the two previous use-cases, the CMTS certificate must carry the URL of the OCSP responder in it to communicate to the CM that OCSP responses validation is required for this certificate. CMTS certificates that contains revocation information are referred to as "Full" or "Full CMTS" certificates.

As described in Section 2.3, that when validation checking is to be performed by the CM, the CMTS must provide the OCSP response inline during BPI+ authentications and that means that if a response cannot be fetched (or the cached version is expired), the CM will reject the CMTS certificate.

2.4.4. Enabling Mutual Authentication with Mutual Revocation Checking

The last use-case we want to explore is the scenario where both the client-side (CMs) and the server-side (CMTS-es) are required to check the validation status of the other party. This setting, as the previous one, is only relevant when BPI+ V2 is used.

To achieve this setup, the CMTS must be configured to enable the revocation checking of CMs' certificates (either via CRLs or OCSP) must be enabled on the CMTS. Moreover, the CMTS, exactly as the previous case, must be provisioned with a Full CMTS certificate to enable revocation checking on the CM.


In this configuration, the authentication message flow will actually use steps {1.a, 1.b, and 1.c} to validate the CM's certificate when the CRL mechanism is used or steps {1.d, 1.e, 1.f} to validate the CM's certificate when the OCSP mechanism is enabled. Additionally, steps {1.d, 1.e, 1.f} must be repeated to gather the revocation information of the CMTS certificate that needs to be sent inline when executing BPI+ V2.

3. Deployment Examples

As we have seen, DOCSIS 4.0 offers a series of new security features that can be independently enabled to lower the network's security risks. This section provides an overview of one of the possible paths to DOCSIS 4.0 deployment. We start by providing considerations on how to support D4.0 devices in existing networks (such as DOCSIS 3.1 or earlier) and then we focus on the impact of enabling the new security features when deploying new DOCSIS 4.0-enabled networks. The envisioned architecture is depicted in Figure 4 where it is assumed that the CMTS (i.e., or where the functionality is provided such as the CCAP Core or MAC-NE elements in distributed architectures) does not have direct access to external services. In this architecture, the CMTS can route all the HTTP request for CRLs and OCSP processing via the HTTP Proxy where strict access rules can be easily enforced.

Additionally, to improve network reliability, operators may deploy local OCSP Proxies and/or responders that can directly sign (or cache) valid responses locally. This solution can be used not only to improve the network reliability, but also to provide support for locally managed revocations that are specific for the operator's ecosystem such as tracking (permissive policy) or reject (strict policy) cloned or otherwise potentially compromised devices.

Figure 4 - Example Deployment with revocation checking support and local overrides



3.1. Preparing Your Networks for D4.0 CMs

To prepare the network for operating with DOCSIS 4.0 devices, there are important details that operators need to be aware of for successful deployments planning.

In Section 2.2, we discussed the additions made to DOCSIS 4.0 certificates when compared to DOCSIS 3.1 or earlier profiles (or configurations). It is therefore evident why DOCSIS 4.0 certificates are bigger in size than their DOCSIS 3.1 equivalent.

Because of the increased size of the certificates, even BPI+ V1 messages might hit the software limit imposed on the size of BPKM messages (1490 bytes)². The relaxation of the software limitation on BPKM messages is needed to enable the use of a common certificate when operating in DOCSIS 4.0 and DOCSIS 3.1 modes. Although this choice lowers devices cost because of the use of a single certificate, in some cases DOCSIS 3.1 CMTS-es might require a software update to enable processing BPKM messages that are larger than 1490 bytes. It is important to notice that the same issue does not affect pre-D3.1 backward compatibility: the common certificate cannot be used in this case, and, exactly as for DOCSIS 3.1 devices, a separate certificate issued from the 1st Gen DOCSIS PKI (i.e., the “legacy” PKI) is still needed

When operators are ready to offer higher speed tiers and additional services that come with it, networks can be upgraded by deploying new DOCSIS 4.0 CMTS (or their DAA equivalent), thus enabling new speeds and the possibility to upgrade the security features, if needed, at a later time. Let’s see how.

² The software limit stems from limitations in the frame size of DOCSIS 3.0 MAC and Phy layers that was wrongly directly imported in the DOCSIS 3.1 specifications without any update.

3.2. Upgrading Speeds, Not Security

When the time is right for a network upgrade, new DOCSIS 4.0 CMTS services can be enabled on the network. Besides the needed work to configure and upgrade the PHY layer to be able to deliver the new speeds, DOCSIS 4.0 does not require any changes from the currently used deployment model when it comes to security.

In previous Sections of the paper, we have shown how DOCSIS 4.0 networks can be configured to leverage the same authentication protocol used in DOCSIS 3.1 (and previous) networks (i.e., BPI+ V1) and how, additionally, it is possible to duplicate the existing networks configuration for revocation checking in the new environment. The combination of these two controls allows DOCSIS 4.0 deployments to leverage new speeds without requiring changes in NetOps because of security: a feature aimed at ease the transition, on the operator's time, to more secure options.

In other words, only when and if new security features are needed (i.e., PFS or MA), operators may decide to enable the use of BPI+ V2 according to their own deployment plans and schedule.

3.3. Enabling Advanced Security Features With BPI+ V2

During the development of DOCSIS 4.0, the need for providing trusted networks where the identity of the network is validated was quite evident: not only the use of a network identity enables more secure authentication and privacy options, but it also introduces the concept of authenticated and trusted networks that is the basis on top of which networks can offer trusted services. The need for trusted networks, combined with the need to mitigate new attack vectors related to new distributed deployment models, are some of the core reasons for enabling the new BPI+ V2.

When the new authentication protocol is enabled, a new set of controls is available to the operator via the configuration of the Persistent Security Attributes (PSAs). PSAs are stored in the secure memory of the CM and allow operators to further restrict accepted network/CMTS identities and protocol's versions during authentications (i.e., via the `Allowed-BPI-Versions` TLV) and are used to implement the downgrade protection mechanisms for BPI+. Additionally, PSA attributes are used to restrict what is considered valid with respect of network identities by requiring, in the CMTS certificate, the presence of specific values. For example, it is possible to configure a CM to only accept "Operator A" as the Organization field (O) in the certificate's subject name by using the `CMTS-Designations` TLVs in the Auth Reply message from the CMTS.

Enabling BPI+ V2, however, requires the provisioning and management of the CMTS certificate as discussed earlier. This means that managing CMTS (only when BPI+ V2 is enabled) will require supporting a new set of operations: from requesting the initial certificate (if not already installed by the Vendor) to regularly renewing it before expiration (i.e., once every 5 years). Although not part of DOCSIS 4.0 specifications, it is expected that automated certificate renewal protocols and tools will be developed and integrated with the increased enablement of BPI+ V2 across networks.

3.4. Enabling Revocation Checking and DOCSIS 4.0

As discussed throughout the paper and specifically in Section 2.4, DOCSIS 4.0 allows for very flexible configurations when it comes to revocation status checking. In fact, it is important to notice how enabling or disabling revocation status checking can be done independently from enabling or disabling the use of BPI+ V2. This capability is the key for empowering operators to choose the deployment path that is more relevant for their networks today and their upgrade path(s) tomorrow – even if they need to change it

during execution. In other words, because of the different options available in DOCSIS 4.0, revocation checking should not be considered a limiting factor for its deployment.

This said, because revocation checking has not been widespread enabled in the broadband community, enabling support for it (especially when requiring CMs to check revocation status of CMTS certificates) should be carefully planned and might require additional infrastructure services such as HTTP proxies or local OSCP responders as depicted in Figure 3.

4. Acknowledgments

The work described in this paper is the ultimate result of the contributions to DOCSIS 4.0 design and specifications from the whole community. We would like to recognize the many great contributions from both the Operators and the Vendors community without which DOCSIS 4.0 could not have happened. We would also like to recognize the continuous support and commitment to the DOCSIS 4.0 SEC WG activities from Ali Negahdar, Colin Dearborn, Dan Torbet, David Taylor, Margo Dolas, Jeff DeMent, Jeff Finkelstein, Owen Parsons, Onur Zengin, Pawel Sowinski, Philip Anderson, Ramy Elmoneiry, Sasha Medvinsky, and Satish Mudugere.

5. Conclusion

In this paper, we provide an overview of the many new security features in DOCSIS 4.0 with particular attention to the impact of its deployment on existing and new networks.

After a brief introduction where we describe the history of DOCSIS security together with considerations about new deployment models, the paper continues with a description of the security principles adopted in DOCSIS 4.0 and how they address new possible threats when considering distributed architectures. In particular, we have seen how DOCSIS 4.0 can be deployed by using the same authentication protocol and procedures that are in use in today's DOCSIS 3.1 networks and how operators can enable existing and new features by enabling BPI+ V2.

When it comes to revocation status checking, we also provided important considerations on how to support efficient revocation checking and described how to support different degrees of enforcement (i.e., strict vs. permissive policies).

Ultimately, DOCSIS 4.0 and BPI+ V2 open new future possibilities for the broadband industry and paves the road for practical solutions to address upcoming security issues or threats (such as post-quantum cryptography deployment for DOCSIS) while providing a cost-effective and efficient path to get there (algorithm agility).

Abbreviations

3GPP	3rd Generation Partnership Project
AES	Advanced Encryption Standard
AIA	Authority Information Access
AK	Authorization Key
BPKM	Baseline Privacy Key Management
BPI	Baseline Privacy Interface
BPI+	Baseline Privacy Interface Plus
CA	Certificate Authority

CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
CRL	Certificate Revocation Lists
CVC	Code Verification Certificate
DOCSIS	Data Over Cable Service Interface Specification
EAE	Early Authentication and Encryption
EKU	Extended Key Usage
ESP	Enhanced Secure Provisioning
FMA	Flexible MAC Architecture
FS	Full Security
FWAH	Firmware Authentication Header
HFC	Hybrid Fiber-Coaxial
HMAC	Keyed-Hash Message Authentication Code
HTTP	Hypertext Transfer Protocol
KEK	Key Encryption Key
MA	Mutual Authentication
MAC	Media Access Control
MAK	Message Authentication Key
MD5	Message-Digest algorithm 5
MDD	MAC Domain Descriptor
MITM	Man-In-The-Middle
MMM	MAC Management Message
MSO	Multiple Systems Operator
NetOps	Network Operations
NetSecOps	Network and Security Operations
NIST	National Institute of Standards and Technology
NRI	No Revocation Information
OCSP	Online Certificate Status Protocol
OID	Object Identifier
PFS	Perfect Forward Secrecy
PKI	Public Key Infrastructure
R-PHY	Remote PHY
R-MACPHY	Remote MACPHY
RMD	R-MACPHY Device
RPD	Remote PHY Device
RSA	Rivest, Shamir, Adleman (a public key cryptographic algorithm)
SA	Security Association
SCTE	Society of Cable Telecommunications Engineers
SHA-1	Secure Hash Algorithm 1
SHA-256	Secure Hash Algorithm 256-bit
SSD	Secure Software Download
SSH	Secure Shell
TEK	Traffic Encryption Key
TFTP	Trivial File Transfer Protocol
TLV	Type/Length/Value
ToD	Time of Day
TOFU	Trust on First Use

URL	Uniform Resource Locator
ZTN	Zero Trust Networks

Bibliography & References

[SP-SSI] MCNS Data Over Cable Security System Specification SP-SSI-I01-970506 (SP-SSI) May 6, 1997. Cable Television Laboratories, Inc.

[RSMI] MCNS Cable Modem Removable Security Module Interface Specification SPRSMI01-970425 April 25, 1997. Cable Television Laboratories, Inc.

[BPI] Data-Over-Cable Service Interface Specifications 1.0, Baseline Privacy Interface Specification, SP-BPI-C01-011119, November 19, 2001. Cable Television Laboratories, Inc.

[BPI+08] Data-Over-Cable Service Interface Specifications, Baseline Privacy Plus Interface Specification, CM-SP-BPI+-C01-081104, November 4, 2008.

[MULPIv4.0] Data-Over-Cable Service Interface Specifications, DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv4.0-I05-220328, March 28, 2022, Cable Television Laboratories, Inc.

[SECv3.0] Data-Over-Cable Service Interface Specifications 3.0, Security Specification, CM-SP-SEC3.0-I16-160602, June 2, 2016, Cable Television Laboratories, Inc.

[SECv3.1] Data-Over-Cable Service Interface Specifications 3.1, Security Specification, CM-SP-SECv3.1-I10-2111110, November 11, 2021, Cable Television Laboratories, Inc.

[SECv4.0] DOCSIS 4.0 Security Specification, CM-SP-SECv4.0-I04-220328, March 28, 2022, Cable Television Laboratories, Inc.

[Pala21] Massimiliano Pala. Enabling Encryption and Algorithm Revocation for Pos-Quantum DOCSIS Certificates. SCTE Cable-Tech Expo, September 2021.

[OSSL] The OpenSSL project. <https://www.openssl.org>

[OCAOD] The OpenCA OCSPD project. <https://www.openca.org>

DOCSIS[®] Time Protocol Proof of Concept Phase II Results

A Technical Paper prepared for SCTE by

Ruoyu (Roy) Sun
Ph.D., Principal Architect
CableLabs
858 Coal Creek Cir, Louisville CO, 80027
303-661-6789
r.sun@cablelabs.com

Aaron Quinto, Mark Poletti, CableLabs, Inc.

Charles Cook, Vikas Sarawat, Praveen Srivastava, Lei Zhou, Charter Communications, Inc.
Elias Chavarria Reyes, Ph.D., Hitron Technologies, Inc.

Table of Contents

Title	Page Number
1. Introduction.....	4
1.1 What Is DTP?	4
1.2 Where Is DTP Needed?	4
1.3 How Does DTP Work?	5
1.4 DTP History and Status.....	6
1.5 DTP PoC Testing	6
2. DTP PoC Phase II Test Plan and Configuration	6
2.1 Test Plan	6
2.2 Test Configuration	7
3. DTP PoC Phase II Results and Observations	8
3.1 Default Case.....	8
3.2 PTP Over-The-Top.....	9
3.3 Load Testing.....	9
3.4 Fiber and Coaxial Cable Length.....	10
3.5 Amplifiers.....	10
3.6 HFC Configurations.....	11
3.7 NCS Boundary Clock	12
3.8 Upstream OFDMA Channel Frequency	12
4. Conclusion.....	13
4.1 Suggestions for MSOs	14
Abbreviations	14
Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1 – DTP supports 5G NR/LTE in the field.....	4
Figure 2 – Where Is DTP Needed?.....	5
Figure 3 – How Does DTP Work?.....	5
Figure 4 – DTP PoC Testing Lab Configuration.	8
Figure 5 – VNA Measured Group Delay for a Five Diplexers Plant.....	13

List of Tables

Title	Page Number
Table 1. Mobile Network Timing Sources.	4
Table 2. DTP PoC Phase II Test Plan [8].	7
Table 3. Default Case Results.	8
Table 4. PTP Over-The-Top Results.	9
Table 5. Load Testing Results.	9
Table 6. Fiber and Coaxial Cable Length Results.	10
Table 7. Amplifiers: VNA Results.	10
Table 8. Amplifiers: DTP Results.	11
Table 9. HFC Configurations Results.	11
Table 10. NCS Boundary Clock Results.	12

Table 11. US OFDMA Channel Frequency Results. 13

1. Introduction

1.1 What Is DTP?

Mobile networks require high accuracy time and frequency synchronization. 3GPP mandates 1.5 μ s of timing accuracy from the base stations (BSs) to a common time reference [1]. Meeting this 1.5 μ s target avoids inter-base station interference in time-division duplex (TDD) networks. In addition, inter-network synchronization is required by the FCC for Citizens Broadband Radio Service (CBRS) and is recommended for other bands. In outdoor environments, Global Positioning System (GPS) signals are reliable and can be used to derive an accurate timing signal. However, when BSs are deployed indoors where GPS signals are unreliable, BSs need an accurate timing signal delivered over the backhaul link. DOCSIS Time Protocol (DTP) is designed to provide such a high-accuracy synchronization signal on the Hybrid Fiber Coaxial (HFC) network, serving as the Xhaul (Xhaul refers to backhaul, mid-haul, or fronthaul) for mobile networks.

Figure 1 shows a DOCSIS network as mobile backhaul and DTP as the timing source for a 5G New Radio (NR) or 4G Long-Term Evolution (LTE) base station. The CMTS gets timing signals from the Primary Reference Time Clock (PRTC) using the Precision Time Protocol (PTP), also known as the IEEE 1588 Standard. PTP cannot be directly used on an HFC network (see the reason in Section 3.2). The HFC network (CMTS, remote-PHY device (RPD)/remote-MAC-PHY device (RMD), amplifiers (A) and cable modem (CM)) uses DTP instead. The CM delivers PTP timestamps to BSs using PTP, which BSs widely support. Note that the terminology of PTP/DTP master/slave was used in the IEEE 1588 standards and the DOCSIS specifications. The IEEE 1588 working group is considering using more inclusive language: PTP timeTx and PTP timeRx to replace master and slave, respectively, and we use this new terminology in this paper.




Figure 1 – DTP supports 5G NR/LTE in the field.

1.2 Where Is DTP Needed?

Table 1. Mobile Network Timing Sources.

Timing sources	Accuracy	Applications
NTP	tens of ms	Not usable for BS
PTP	< 1.5 μ s	Widely supported by LTE/NR BS, PTP OTT on the HFC network performs poorly
GPS	a few ns	For outdoor BS
DTP	< 1.5 μ s	<i>Fills the gap!</i>

Table 1 lists potential timing sources for mobile networks. The Network Time Protocol (NTP) accuracy is in the order of tens of milliseconds, which does not meet the 3GPP requirements.

PTP over-the-top (OTT) on the HFC network performs poorly. GPS signal is unreliable in indoor or urban canyon environments. Multi-Service Operators (MSOs) are interested in deploying indoor small-cell networks using HFC as backhaul, for which DTP is the only option, as summarized in Figure 2.




Figure 2 – Where Is DTP Needed?

1.3 How Does DTP Work?

The CMTS, RPD/RMD, and CM all have their local timestamp. The task for DTP is to synchronize these local timestamps. As illustrated in Figure 3, the CMTS and RPD/RMD have a fiber connection, hence, they use PTP to update their local timestamp. The RPD maps its local timestamp into the DOCSIS 3.1 timestamp that is transferred to the CM via the coaxial plant. The DOCSIS 3.1 timestamp is delayed in the downstream path. The DOCSIS ranging procedure measures the round-trip delay between RPD and CM. This round-trip delay is defined as the true ranging offset (TRO). If the downstream and upstream delays in the cable plant are the same (ideally symmetrical), half of TRO is applied to correct the downstream delay for the CM local timestamp. The CM maps its local timestamp to a PTP timestamp output for mobile networks. The DTP timeTx in the CMTS and the DTP timeRx in the CM exchange messages that include parameters of the HFC plant. The DTP timeRx also reports the real-time TRO to the DTP timeTx. The CMTS uses half of TRO and other parameters to calculate the time adjustment, t_{adj} , that is applied in the CM.

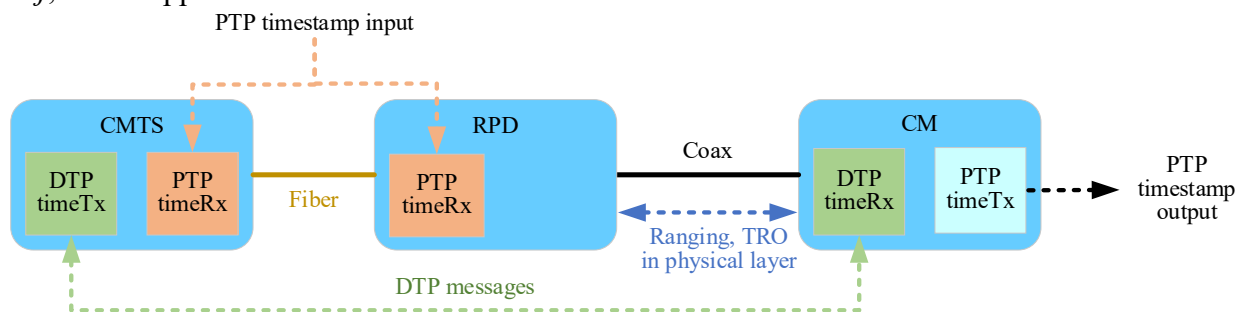


Figure 3 – How Does DTP Work?

DTP has many challenges, such as the asymmetrical delay in the HFC plant (the downstream and upstream delays in the cable plant are different, so half of the TRO cannot perfectly compensate the downstream delay), as well as some non-ideal effects in the CMTS, RPD/RMD and CM devices. A three-step DTP calibration procedure (see Section 7 in [8]) is needed to address these challenges.

1.4 DTP History and Status

DTP was invented in 2011 [2], standardized in the DOCSIS MULPI 3.1 Specification [3], and further incorporated in the DOCSIS SYNC specification [4] in 2020. DTP leverages existing DOCSIS hardware-based timestamp, accounts for path asymmetry, and provides timing performance independent of the traffic load on the DOCSIS network. DTP is being implemented in the industry. From September 2020, CableLabs, Charter Communications, Cisco, and Hitron started DTP proof-of-concept (PoC) tests. A DTP calibration procedure was designed to correct asymmetrical and other non-ideal delays in the HFC network. In December 2021, CableLabs/Kyrio developed a cloud database [5] that distributes the calibration data. The DOCSIS SYNC spec [4] was updated on July 15th, 2022, to describe the DTP calibration procedure and define the interface between the cloud database and CMTS.

1.5 DTP PoC Testing

The PoC testing was split into two phases. Phase I tests were conducted in a basic lab environment, and the results were published in an SCTE 2021 paper [5] and a CableLabs Technical Report [7]. This paper presents the phase II testing results. This second phase evaluates DTP performance in a complex environment similar to a field deployment with different downstream and upstream traffic load levels and fiber and coaxial cable lengths with multiple amplifiers. Phase II testing also considered HFC physical layer configurations such as modulation, interleaver, and cyclic prefix (CP) in the downstream orthogonal frequency-division multiplexing (OFDM) channel. In the upstream orthogonal frequency-division multiple access (OFDMA) channel, Phase II testing considered modulation, frame size, and CP. DTP PoC testing demonstrated that DTP could successfully deliver accurate timing information to end applications. TDD mobile networks successfully met the 3GPP timing accuracy specifications using DTP on the backhaul.

While working on the DTP PoC phase 1 and phase 2 testing, the group made many other contributions including designing the DTP cloud database [5], and updating the DOCSIS SYNC Specification [4].

2. DTP PoC Phase II Test Plan and Configuration

The DTP PoC phase 2 testing was conducted in Q3 and Q4 of 2021. This paper summarizes the key observations. More details are presented in [8].

2.1 Test Plan

The DTP PoC phase 2 testing aims to evaluate the DTP performance in a complex field environment. The cases listed in Table 2 cover most of the cases for MSOs' field deployment scenarios. We firstly confirmed that PTP over-the-top of HFC networks performs poorly. Then we tested DTP with different levels of traffic load in both downstream (DS) and upstream (US), fiber and cable lengths, number of amplifiers, and with different HFC network configurations at the CMTS.

The second column in Table 2 contains the baseline values that define the default test case, in which the traffic load, fiber and coax cable length and number of amplifiers are set the minimum value in our lab, and the CM and CMTS configurations capture the most commonly used values

by MSOs. Columns 3 and 4 contain comparative and extreme values that may or may not often be used in field deployments. Each case only changes one value from the default case. Several factors limit the phase 2 testing: (1) the CMTS serves as the DTP timeTx and the CM serves as timeRx (the SYNC spec also supports the CM serving as the DTP timeTx and the CMTS serving as the DTP timeRx); (2) the test bed uses the distributed access architecture (DAA) with RPD; (3) the primary upstream channel is the OFDMA channel instead of the single carrier quadrature amplitude modulation (SC-QAM) channel (SC-QAM channel has a different TRO, etc.); (4) assumes 4k fast Fourier transform (FFT) size in the US OFDMA channel and 8k FFT size in the DS OFDM channel; (5) assumes the DS OFDM channel has a flat profile; (6) PTP multicast was used between the CM and the Paragon-X and PTP unicast was used between the Paragon-X (or PRTC) and the CMTS & RPD.

Table 2. DTP PoC Phase II Test Plan [8].

Parameters		Baseline test value	Comparative test values	Extreme values (optional test)	Note
DS load impact on PTP over the top		0	25%, 50%	75%, 100%	
US load impact on PTP over the top		0	25%, 50%	75%, 100%	
DS load impact on DTP		0	25%, 50%	75%, 100%	
US load impact on DTP		0	25%, 50%	75%, 100%	
Fiber length (NCS to RPD)		90 m	5 m, 25 km		
Coax length (RPD to CM)		3 m	244 m (800 ft) 591 m (1938 ft) 835 m (2738 ft)		
Number of amplifiers		0	1, 2, 3, 4, 5		
CM configs (US)	Frame size	$K = 6$	$K = 18$, BW < 48 MHz		Assume 4k FFT size. K is number of symbols in a frame
	OFDMA modulation	256-QAM	64-QAM	1024-QAM	1024-QAM only in a clean environment with no noise
	CP	6: 256 samples	4: 192 samples		
CMTS configs (DS)	Interleaver	2	1	16	Assume 8k FFT size
	Modulation	4096-QAM	1024-QAM, 256-QAM		Assume flat profile
	CP	1 (1.25 μ s, 256 samples)	2 (2.5 μ s, 512 samples)	3.75 μ s (768 samples), 5 μ s (samples), and 0.94 μ s (192 samples)	

In addition to Table 2, we also tested three cases:

- Impact of the Cisco Network Convergence System (NCS) boundary clock (BC), see Section 3.7.
- The upstream OFDMA channel frequency range, which impacts the group delay, see Section 3.8.

2.2 Test Configuration

The lab test configuration is shown in Figure 4. Because there is no solution that can measure DTP performance directly, we used the Paragon-X to measure the PTP time error (TE) between the input of the CMTS/RPD and the output of the CM. The PTP timeTx in the Paragon-X is

connected to the Cisco NCS 55A1-24Q6H-SS that serves as a PTP BC, which provide the PTP timing source for both the integrated CMTS (I-CMTS) and RPD. The I-CMTS is the Cisco cBR-8 with software version 16.12.z1. The RPD is the Cisco SmartPHY 120 with software version v7.8.2. The CM is the Hitron ODIN1112 with software version ODIN-724GA-7.2.4.0.152 that uses MaxLinear's Puma 7 solution.

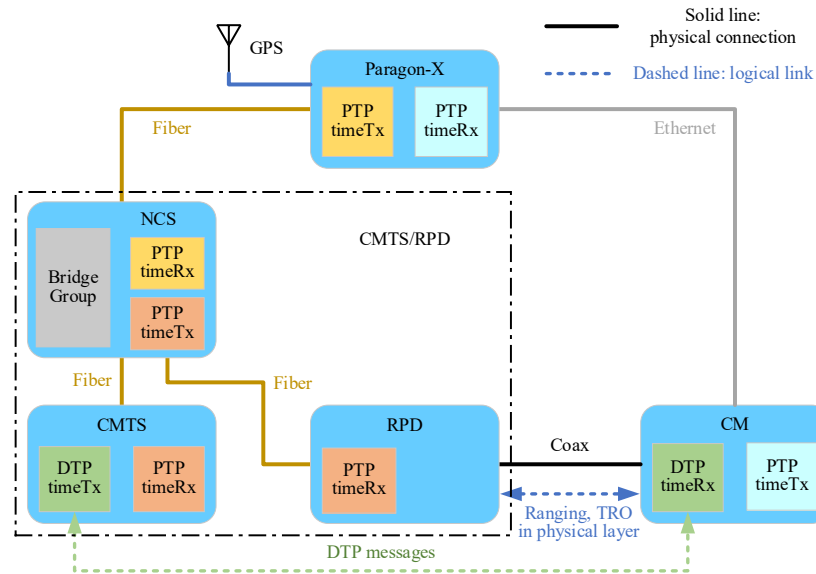


Figure 4 – DTP PoC Testing Lab Configuration.

3. DTP PoC Phase II Results and Observations

3.1 Default Case

The default case uses the baseline values listed in Table 2. Calibration was done before doing the PoC testing, as reported in Section 5.2 in [8]. Each of the phase 2 cases were run five times. Each run was one hour long. The TRO was captured on the CM before and after the run. The Paragon-X generated many statistics of the time error. In this paper we only focus on the constant TE and dynamic TE. The Paragon-X compares the PTP timestamp sent from its timeTx to the HFC network and the PTP timestamp received by its timeRx from the HFC network. The constant TE is the difference (time error) averaged over one hour and five runs. The dynamic TE is the maximum variation of the TE over a 1000 second moving window in each run, then averaged over five runs.

The default case results are provided in Table 3. The time error budget for this DAA scenario is 980 ns, see Table 4 in [7]. The time error has an average value of -35 ns and variation of up to 214 ns, which meet the 980 ns requirement.

Table 3. Default Case Results.

Case Index	Case	Average TRO (ns)	Constant TE (ns)	Dynamic TE (ns)
1	Default Setting	5,315,965	-35	214

Observation 1: DTP will work. The default test case with baseline configurations met the DTP time error budget.

3.2 PTP Over-The-Top

In this case, DTP is not running in the HFC network. PTP messages are over-the-top traffic in the HFC network. Thus, PTP messages are impacted by HFC queueing. Upstream PTP messages go through a DOCSIS upstream service flow configured with a best-effort scheduling service. Case 2 in Table 4 does not have any traffic load. The constant TE is over 3 ms. The TE varies over 4 ms. Both the constant TE and dynamic TE are multiple orders larger than the 3GPP requirement of 1.5 μ s.

Two other CMs and a tap were added after the RPD, see Figure 5 in [8]. The two CMs are controlled by a load tester (ByteBlower) to create a certain amount of traffic load either downstream or upstream. Cases 3 to 6 are with DS load from 25% to 100%, and cases 7-10 are with US load from 25%-100%. The constant TE and dynamic TE for all these cases do not meet the 3GPP requirement.

Table 4. PTP Over-The-Top Results.

Case Index	Case	TRO (ns)	Constant TE (ns)	Dynamic TE (ns)
2	No load	5,315,925	3,012,010	4,440,810
3	25% DS load	5,315,925	3,026,168	4,273,112
4	50% DS load	5,315,929	3,017,390	3,656,898
5	75% DS load	5,315,929	3,008,489	3,845,175
6	100% DS load	5,315,929	735,443	6,074,600
7	25% US load	5,315,929	3,173,979	5,383,784
8	50% US load	5,315,929	3,459,314	6,052,003
9	75% US load	5,315,917	3,807,350	5,559,820
10	100% US load	5,315,917	4,672,663	7,603,164

Observation 2: PTP over-the-top does not meet the 3GPP requirement. The dynamic TE is on the order of milliseconds with or without traffic load. The cTE changes with different levels of downstream and upstream load by multiple milliseconds.

3.3 Load Testing

Three CMs share the same channels. Two CMs generate 25% to 100% traffic load in either DS or US. The other CM runs DTP. The results are presented in Table 5. Because DTP messages are designed to be a control message, and the DOCSIS 3.1 timestamp is transferred on the physical layer link channel (PLC), neither the constant TE nor dynamic TE is impacted by HFC network traffic.

Table 5. Load Testing Results.

Case Index	Case	TRO (ns)	Constant TE (ns)	Dynamic TE (ns)
11	25% DS load	5,315,933	-58	216
12	50% DS load	5,315,933	-53	219
13	75% DS load	5,315,925	-64	215
14	100% DS load	5,315,988	-33	222
15	25% US load	5,315,941	-48	212
16	50% US load	5,315,941	-36	208
17	75% US load	5,315,996	-28	225
18	100% US load	5,315,996	-21	189

Observation 3: Neither the DS nor the US traffic load impacts the DTP performance.

3.4 Fiber and Coaxial Cable Length

The fiber length between the NCS and the RPD in the default case is approximately 90 m. Cases 19 and 20 change the fiber length to 25 km and 5 m. The TE is not impacted by the fiber length, see Table 6. Ranging and TRO between the RPD and the CM are also not impacted by the fiber length in such a DAA-RPD architecture.

Table 6. Fiber and Coaxial Cable Length Results.

Case Index	Case	TRO (ns)	Constant TE (ns)	Dynamic TE (ns)
19	25 km fiber	5,315,901	-42	223
20	5 m fiber	5,315,984	-15	215
21	244 m coaxial + 1 Amp	5,318,125	49	223
22	591 m coaxial + 3 Amps	5,320,750	185	237
23	835 m coaxial + 5 Amps	5,323,038	194	224

Observation 4: Fiber length in the DAA-RPD architecture does not impact DTP, nor does it impact TRO.

The coaxial cable length from the RPD to the CM in the default case is approximately 3 m. We replaced it with 244, 591, or 835 m long cables. To compensate high attenuation, multiple amplifiers need to be used. The TE changes slightly from the default case to cases 21-23, which is not due to the asymmetrical TE introduced by amplifiers but instead due to impact from the long cable. The TRO increases with cable length correspondingly to compensate for the additional delay introduced by the long cable.

The delay of the 244 m cable is also measured by a Vector Network Analyzer (VNA); see Table 7. The DS delay at 591 MHz is 1004 ns and the US delay at 32 MHz is 1005 ns, which are almost the same.

Observation 5: Cable length does not impact DTP. VNA measurements indicate that coaxial cable does not introduce any asymmetrical delay. The additional round-trip delay from cable length is symmetrical, and the corresponding TE is well compensated by the TRO, which is verified in the DTP measurements.

3.5 Amplifiers

The constant TE for cases 21-23 in Table 6 are slightly different. We designed cases 24-27 to further check if the additional constant TE is due to the cable length or amplifiers. Cases 24-27 use the same cable length of 3 m, but change the number of amplifiers (from QDAX) from one to four. Each of the amplifiers introduce additional TE from 38 to 83 (case 1 vs. case 24) ns with an average of 54 ns, see Table 8. The VNA measurements for the QDAX amplifiers are listed in Table 7. The QDAX amplifiers have a large asymmetrical delay between DS and US, which introduces an additional TE of 36-40 ns per amplifier. The additional TRO and TE are in the same range with the DTP results. The 591 m cable plant is built with three Arris amplifiers. The Arris amplifiers also have a large asymmetrical TE that introduce additional TE of approximately 27 ns per amplifier.

Table 7. Amplifiers: VNA Results.

VNA Measurements	QDAX Amplifier					591 m hardline cable + 3 Arris Amps	244 m RG-6 Cable
	#1	#2	#3	#4	#5		
DS delay $D_{DS'}$ (ns)	6.3	6.3	6.3	6.4	6.4	2278	1004
US delay $D_{US'}$ (ns)	81.5	79	80.9	81.3	86	2442	1005

Additional TE ($(D_{US}' - D_{DS}')/2$ (ns))	38	36	37	37	40	82	0.5
Additional TRO $D_{US}' + D_{DS}'$ (ns)	88	85	87	92	92	4720	2009

Table 8. Amplifiers: DTP Results.

Case Index	Case	TRO (ns)	Constant TE (ns)	Dynamic TE (ns)
24	3 m coaxial + 1 Amp	5,316,160	48	218
25	3 m coaxial + 2 Amps	5,316,250	101	215
26	3 m coaxial + 3 Amps	5,316,394	139	229
27	3 m coaxial + 4 Amps	5,316,480	182	217

The additional TE from amplifiers varies in the field. The MSOs need to consider this factor in the field DTP deployment. If a large number of amplifiers is used, the TE from other network elements will need to be reduced in order to meet the entire TE budget. For example, a higher class of CMTS or CM with better quality and smaller TE may be used. An alternative solution is to reduce number of cascading boundary clocks used in the DTP network. These solutions are suggested in Section 8.4.2.5 in the SYNC spec [4].

Observation 6: Amplifiers introduce an asymmetrical delay in the HFC plant and additional TE to DTP. The additional TE varies with amplifier make and model. The QDAX amplifier has an additional TE of 36–40 ns that is accurately characterized by the VNA. Such additional TE from each QDAX amplifier is, on average, 54 ns measured in the DTP testbed.

3.6 HFC Configurations

The baseline value for the US frame size is 6 symbols per frame, the US modulation is 256-QAM and the US CP is 256 samples. The US frame size is changed to 18 symbols per frame in case 28. The US modulation is changed to 64-QAM and 1024-QAM in cases 29 and 30. The US CP is changed to 192 samples in case 31. The results are provided in Table 9. The TE is not impacted by these US configurations.

Observation 7: HFC upstream network configurations of frame size, modulation scheme, and cyclic prefix do not impact DTP performance.

Table 9. HFC Configurations Results.

Case Index	Case	TRO (ns)	Constant TE (ns)	Dynamic TE (ns)
28	US frame size: 18	5,315,965	-32	212
29	US mod: 64-QAM	5,315,925	-50	221
30	US mod: 1024-QAM	5,315,933	-48	227
31	US CP: 192	5,315,941	-48	215
32	DS Interleaver: 1	5,315,968	-33	226
33	DS Interleaver: 16	5,315,992	-30	224
34	DS mod: 1024 for data 256 for control	5,315,965	-44	209
35	DS mod: 256 for data 64 for control	5,315,972	-29	225
36	DS CP: 192	5,276,152	19,898	211
37	DS CP: 512	5,475,554	-79,820	230
38	DS CP: 768	5,635,687	-159,914	228
39	DS CP: 1024	5,795,644	-239,959	221

The default case uses DS interleaver depth of 2 and DS modulation of 4096-QAM for data and 1024-QAM for the control channel. Cases 32 and 33 compare the DS interleaver depth of 1 and 16. Cases 34 and 35 compare DS modulation of 1024-QAM for data and 256-QAM for the control channel, and 256-

QAM for data and 64-QAM for the control channel, respectively. The TE is not impacted by the DS interleaver and modulation.

Observation 8: HFC downstream interleaver depth and modulation scheme do not impact DTP performance.

The baseline value of DS CP is 256 samples. Cases 36-39 compare DS CP values of 192, 512, 768 and 1024 samples. Both the constant TE and TRO are impacted significantly by the DS CP. This is likely due to a frame alignment issue at the CM. MaxLinear and Hitron are working on fixing the frame alignment issue. Before this issue is fixed, an alternative method is to have DTP devices calibrated for each individual DS CP value to compensate the impact.

Observation 9: The downstream cyclic prefix significantly impacts DTP. Every 1.25- μ s CP length reduces cTE by approximately 80 μ s. As of August 2022, this issue is being investigated by the CM vendor and chipset vendor.

3.7 NCS Boundary Clock

The PTP timeTx in the Paragon-X is connected to the PTP timeRx in the NCS, and the PTP timeTx in the NCS is connected to the PTP timeRx in the Paragon-X directly to evaluate the performance of the NCS as a boundary clock. The NCS employs a class B boundary clock with a theoretical TE of 20 ns. The measured TE is listed in Table 10. The constant TE ranges from -8 to 1 ns with an average of -3 ns. The dynamic TE is 15 ns.

Observation 10: The NCS class B boundary clock TE is between -8 and 1 ns, which is smaller than the 20-ns TE budget defined in the SYNC spec [4].

Table 10. NCS Boundary Clock Results.

Case Index	Case	Constant TE (ns)	Dynamic TE (ns)
40	NCS	-3	15

3.8 Upstream OFDMA Channel Frequency

Diplexers, amplifiers, and filters have frequency dependent group delay, which may impact DTP performance. A plant with five cascade diplexers is used to verify if the upstream OFDMA channel frequency may impact DTP. The diplexers have an upper cutoff frequency of 42 MHz. The group delay from 5 to 45 MHz for the five-cascade-diplexer plant is measured by a VNA. As shown in Figure 5, the blue curve is the group delay over frequency, which increases gradually from 115 ns at 5 MHz to 375 ns at 42 MHz, then increases dramatically after 42 MHz. Two US OFDMA channels are selected for the comparative analysis that are on the two edges of the diplexer frequency range: (1) 5-17 MHz (green box in Figure 5); and (2) 31-42 MHz (red box in Figure 5). The average group delay in the 5-17 MHz channel is 128 ns, and in the 31-42 MHz channel 241 ns. In comparison between these two channels, the impact on TRO is 113 ns, and the impact on TE is 56.5 ns (half of TRO).




Figure 5 – VNA Measured Group Delay for a Five Diplexers Plant.

This five-diplexer plant is plugged into the DTP test bed to replace the coaxial cable shown in Figure 4. The DTP measurement results are provided in Table 11. The TRO increased by 99 ns between 5-17 MHz and 31-42 MHz channels, and the constant TE increased 49 ns between the two channels. The DTP results confirmed that the group delay and US OFDMA channel frequency does impact DTP.

Observation 11: The US OFDMA channel frequency does impact DTP depending on the group delay variation over frequency.

Table 11. US OFDMA Channel Frequency Results.

	DTP Results		VNA Results	
	TRO (ns)	Constant TE (ns)	Impact on TRO (ns)	Impact on constant TE (ns)
5-17 MHz vs. 31-42 MHz	99	49	113	56.5

4. Conclusion

DTP provides an accurate timing source for mobile networks. DTP is particularly helpful for the scenarios where the GPS signal is unreliable (e.g., indoor) and the HFC network is used for mobile backhaul, where over-the-top PTP performs poorly. The PoC testing proved DTP meets the 3GPP requirement of 1.5 µs. A three-step DTP calibration procedure is required to correct non-ideal effects in the CMTS, RPD/RMD, and CM devices. Here are the observations from the PoC phase 2 testing:

Observation 1: DTP will work. The default test case with baseline configurations met the DTP time error budget.

Observation 2: PTP over-the-top does not meet the 3GPP requirement. The dynamic TE is on the order of milliseconds with or without traffic load. The cTE changes with different levels of downstream and upstream load by multiple milliseconds.

Observation 3: Neither the DS nor the US traffic load impacts the DTP performance.

Observation 4: Fiber length in the DAA-RPD architecture does not impact DTP, nor does it impact TRO.

Observation 5: Cable length does not impact DTP. VNA measurements indicate that coaxial cable does not introduce any asymmetrical delay. The additional round-trip delay from cable length is symmetrical, and the corresponding TE is well compensated by the TRO, which is verified in the DTP measurements.

Observation 6: Amplifiers introduce an asymmetrical delay in the HFC plant and additional TE to DTP. The additional TE varies with amplifier make and model. The QDAX amplifier has an additional TE of 36–40 ns that is accurately characterized by the VNA. Such additional TE from each QDAX amplifier is, on average, 54 ns measured in the DTP testbed.

Observation 7: HFC upstream network configurations of frame size, modulation scheme, and cyclic prefix do not impact DTP performance.

Observation 8: HFC downstream interleaver depth and modulation scheme do not impact DTP performance.

Observation 9: The downstream cyclic prefix significantly impacts DTP. Every 1.25- μ s CP length reduces cTE by approximately 80 μ s. As of August 2022, this issue is being investigated by the CM vendor and chipset vendor.

Observation 10: The NCS class B boundary clock TE is between -8 and 1 ns, which is smaller than the 20-ns TE budget defined in the SYNC spec [4].

Observation 11: The US OFDMA channel frequency does impact DTP depending on the group delay variation over frequency.

4.1 Suggestions for MSOs

A three-step DTP calibration procedure is required to guarantee DTP performance. The calibration test needs to be done for each pair of CMTS/RPD/RMD and CM devices. The calibration test needs to be repeated for each of the key software releases of these devices. The calibration data will be distributed by a cloud database. The CMTS will query the calibration data and apply them in the field.

DTP performance is impacted by amplifiers, diplexers, and filters, as well as the number of boundary clocks in the field plant. Each amplifier may introduce 36-40 ns additional TE to DTP due to the asymmetrical TE. When multiple amplifiers are used for a DTP CM, in order to meet the entire TE budget, a higher class (with smaller time error) of CMTS or CM may be needed, or a smaller number of boundary clocks may need to be considered [4].

The upstream channel frequency may impact DTP depending on the number of diplexers, amplifiers and filters, which change the group delay over frequency. In the case that the group delay in the upstream varies significantly over frequency, consider using the portion of the channel closest to the center of the upstream band in order to stay clear of the expected delay variation/increase at the band edges [4]. For example, the band edge frequency above 42 MHz should be avoided for the specific HFC network discussed in Section 3.8.

Abbreviations

3GPP	3rd Generation Partnership Project
A	Amplifier
BC	Boundary clock
BS	Base station
CBRS	Citizens Broadband Radio Service
CM	Cable modem

CP	Cyclic prefix
cTE	Constant time error
DAA	Distributed access architecture
DS	downstream
DTP	DOCSIS time protocol
FCC	Federal Communications Commission
FFT	Fast Fourier transform
GPS	Global Positioning System
HFC	Hybrid fiber coaxial
I-CMTS	Integrated cable modem termination system
LTE	Long-term evolution
MSO	Multi-Service Operator
NCS	Network Convergence System
NR	New radio
NTP	Network time protocol
OFDM	Orthogonal frequency-division multiplexing
OFDMA	Orthogonal frequency-division multiple access
PoC	Proof of concept
PTP	Precision time protocol, also known as IEEE 1588
PLC	Physical-layer link channel
PRTC	Primary Reference Time Clock
QAM	Quadrature amplitude modulation
RMD	Remote-MAC-PHY device
RPD	Remote-PHY device
SC-QAM	Single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
TDD	Time division duplexing
TE	Time error
TRO	True ranging offset
US	upstream
VNA	Vector Network Analyzer

Bibliography & References

- [1] 3GPP Technical Specification 36.133 v16.9.0, Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for Support of Radio Resource Management, June 2021. [\[link\]](#)
- [2] John T. Chapman, Rakesh Chopra, Laurent Montini., “The DOCSIS[®] Timing Protocol (DTP), Generating Precision Timing Services from a DOCSIS System,” *INTX/SCTE Spring Technical Forum*, 2011. [\[link\]](#)
- [3] Cable Television Laboratories, Inc., “DOCSIS[®] MAC and Upper Layer Protocols Interface Specification,” CM-SP-MULPI, December 2020. [\[link\]](#)
- [4] Cable Television Laboratories, Inc., “Synchronization Techniques for DOCSIS[®] Technology Specification,” CM-SP-SYNC, April 2021. [\[link\]](#)

- [5] Ruoyu Sun, Rahil Gandotra, Jennifer Andreoli-Fang, Elias Chavarria Reyes, John T. Chapman, Mark Poletti, “Designing a Cloud-Based DOCSIS Time Protocol Calibration Database,” in *SCTE-Expo 2021*, Atlanta, GA, October 11-14, 2021. [[link](#)]
- [6] Ruoyu Sun, Jennifer Andreoli-Fang, Elias Chavarria Reyes, John T. Chapman, et al., “DOCSIS Time Protocol Proof of Concept,” in *SCTE-Expo 2021*, Atlanta, GA, October 11-14, 2021. [[link](#)]
- [7] Ruoyu Sun, Charles Cook, Ryan Tucker, John T. Chapman, Elias Chavarria Reyes, “DOCSIS[®] Time Protocol Proof of Concept Phase I Technical Report CM-TR-DTP-V01-210915,” Technical Report, CableLabs, CM-TR-DTP-V01-210915, September, 2021. [[link](#)]
- [8] Ruoyu Sun, Elias Chavarria Reyes, Charles Cook, “DOCSIS Time Protocol Proof of Concept Phase 2 Technical Report,” Technical Report, CableLabs, CM-TR-DTP-Phase2-V01-220307, March 7, 2022. [[link](#)]

Don't Be Passive About Passives

A Technical Paper prepared for SCTE by

Kyle Hohman
Standards Specialist
Shaw Communications
751 Enterprise Crescent, Victoria, BC
403-538-5252
kyle.hohman@sjrb.ca

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Passive Replacement Programs.....	3
2.1. Passive Evaluation Criteria	5
3. Operationalization of Passive Upgrades.....	6
3.1. Blanket Method	7
3.2. BAU Method	8
4. Passive-Enabled Spectrum Expansion	10
5. Conclusion.....	11
Abbreviations	12
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 – Sample Plant Illustration	4
Figure 2 - Faceplate Upgrade Comparison.....	6
Figure 3 – Blanket Method Illustration	7
Figure 4 – Blanket Method Truck Stock / SKU	8
Figure 5 – BAU Method Sample Illustration.....	8
Figure 6 – BAU Method Truck Stock / SKU	9
Figure 7 – BAU Method with Overlay.....	9
Figure 8 – Passive Enabled 1.8 GHz Testing	10

1. Introduction

As operators look ahead to expand their speed tiers and offerings, they will also be looking to expand their plant capacity to 1.8 GHz Extended Spectrum DOCSIS (FDD) and DOCSIS 4.0. Passives are the highest quantity device in the outside plant that limits the usable bandwidth and will take the most amount of time and energy to replace.

This paper will discuss the methods that an operator can use to replace passive equipment in their plant and examine the operational and technical benefits of each method. We will explore the considerations for evaluating passives in the plant during field trials from both plant performance and maintenance savings perspectives.

For context, the paper will present some of the advantages of having passive upgrades completed before actives, specifically around the ability for 1.8 GHz signals to pass through the plant to measure for ingress, egress, and other test signals. We will also discuss the benefits of being able to set up plant actives by already having passives upgraded to 1.8 GHz, as opposed to a partial tilt setup, which would require the identification of where rolloff exists in order to do a drop-in style upgrade.

2. Passive Replacement Programs

Before we enter in depth into a discussion around passives and how they can be operationalized in the plant, we need to ensure a consistent language is being used around the devices themselves. Passives devices are a category of devices that includes taps, splitters, couplers and power inserters. A passive is also defined as a device that operates in the plant without consuming power directly. It acts as a gateway for the signals passing through it and is designed with specific RF characteristics in mind. Passives are the most common device in the plant, with an industry average of around six passives per active device. For the purpose of this discussion, we will be using an illustration of a sample plant for the various methods of passive upgrades to outline the benefits of each method, as in the figure below.




Figure 1 – Sample Plant Illustration

As cable operators begin their planning for DOCSIS 4.0, the work begins on selection of equipment for evaluation. When an operator is looking to bring a new passive into the plant, there are several factors that need to be explored and understood before proceeding with a field trial.

The first set of considerations are around higher-level plant design requirements for the radio frequency (RF) performance of the devices. The first thing on this list is the insertion loss, or the amount of signal that is lost as the signals pass through the device. As the spectrum is expanded, and to accommodate for higher spectrum transmissions, there is potential for tradeoffs related to the RF design of passive devices. These properties are actively under development by the vendor community and will not be discussed at length in this paper. It is sufficient to understand that the caveat of this performance exists, and how it will impact the operational methodologies examined throughout.

This leads to the signal from the passives to the customer and how it feeds into drop loss. As we look at these higher bandwidths, the total composite power (TCP) of the plant to the modem—and the modem back to the plant—become more of a balancing act than ever before. (Such analysis is outside the scope of this paper but something that will need to be considered as part of the larger design considerations.) After these levels are captured and understood, the secondary RF characteristics of the passive devices must be reviewed.

As spectrum is expanded, some operators are taking a step back to evaluate their network holistically to determine what levels they should be running out of their actives, as well as what levels they would like their modems to be transmitting at. The passive evaluation needs to be included in this discussion as the expanded spectrum greatly impacts the performance achievable through a passive device—along with the design and component usage to build the device—which in turn impacts power passing capability, low-end usable frequency, insertion loss, return loss and isolation between various ports.

Once the high-level design considerations and RF performance characteristics of the passive evaluation are understood, an operator can move to the operational efficiencies and savings that can be realized by selecting the correct passive for their network. The Society of Cable and Telecommunications Engineers (SCTE) has developed a standard for 3 GHz passives (known as SCTE 265 2021) that has completed the foundational work to ensure that the options available to operators have value and longevity in the plant.

At Shaw, we have spent the last few years performing a mid-split upgrade in our plant, coupled with a downstream spectrum expansion. Part of this effort was to replace all passives in the plant from legacy, defined as sub-1 GHz passing, to a minimum of 1 GHz passing. Over the course of this exercise, we learned several operational lessons that will significantly influence our next rounds of passive evaluations and selection.

2.1. Passive Evaluation Criteria

The most immediately impactful piece of selection criteria, aside from questions around advertised spectrum bandwidth, is what the insertion is and through loss of the new device. As touched on above, once the design of the passive is enhanced for almost double the spectrum, there will be tradeoffs with RF performance in other areas of the spectrum. These will vary from vendor to vendor, but they will add up in a cascaded environment. As evidenced in the formulae below, n is the number of passives and x is the loss of the existing passives, which leads you to the legacy loss. The additional loss from the ESD passive, that is the delta from the legacy passive to the new passive, is defined by y . Once a new passive has been added to the plant, every decibel of loss that this expanded spectrum passive could potentially contribute to the overall loss is added for every tap in the run.

$$\textit{Legacy Loss} = nx$$

$$\textit{ESD Loss} = n(x + y)$$

One of the often forgotten but incredibly important considerations is the pin length required for the passive device. As the spectrum that the passive needs to transport expands, the pin length becomes much more important than legacy passives. The accuracy of the pin length creates a much tighter window from a technician craft perspective. Hardline passive connector reusability is a key factor in saving time, and in some cases, there may not be excess cable available to be re-cored and have connectors added.

Another key consideration is the dimensions of the passive devices themselves. Upgrading passives in an existing environment means that the device needs to fit into existing pedestals, into housings and on existing strands. There may not be extra room in pedestals, or enough bend radius or cable length, to implement a large change in passive housing size.

Another often neglected factor in the analysis is the availability of an upgradeable faceplate. Upgradeable faceplates are key features when considering a new passive for the plant. To remove connectors and re-install a device into the plant, rather than unscrewing a faceplate and dropping in a new one, is a massive operational difference.




Figure 2 - Faceplate Upgrade Comparison

Something that we have found over years of maintaining cable plant through all imaginable geography is the importance of the seizure mechanism. The seizure mechanism is the way that the mainline connector pin makes contact for RF and power transmission within the passive device. Fundamentally, there are two types of seizer mechanisms: one that is controlled with a technician tightened set screw and one that has an effective and often proprietary mechanism that applies constant pressure to the pin of the connector. After the initial installation, passives tend to be in the plant until they fail, or the network is upgraded. Since the passives are such long-lived devices in the plant, even the smallest design choices can have a great impact on maintenance cost over the lifetime of the device.

One item that has stood out is the seizure screw for the mainline cable entering the passive. Over the life of a passive, our analysis has shown that the seizure screw may have been forgotten and not tightened or working loose over time, which occurs in a surprisingly large number of cases. Each operator will need to apply their own cost structure to how these issues translate into cost, but each one would be a customer impacting truck roll. The self-seizing mechanism has been shown to be a key success factor in our analysis, as it not only saves the costs associated with truck rolls for set screws that have some looseness, but also speeds up installation by removing some labor and tool requirements from plant technicians.

For any operator is looking to expand their plant to enhance capacity and increase service offerings to customers, the category of equipment that will require the most upgrades is the plant passive. The scale of upgrading every passive in the plant is enormous; as previously discussed, the number of passives overshadows the number of actives in the plant. There are two main methods to operationalizing the upgrades to the plant that we will discuss in the next session: the Blanket Method and the BAU Method.

3. Operationalizaton of Passive Upgrades

For any operator is looking to expand their plant to enhance capacity and increase service offerings to customers, the category of equipment that will require the most upgrades is the plant passive. The scale of upgrading every passive in the plant is enormous; as previously discussed, the number of passives overshadows the number of actives in the plant. There are two main methods to operationalizing the upgrades to the plant that we will discuss in the next session: the Blanket Method and the BAU Method.

3.1. Blanket Method

The Blanket Method refers to the complete, targeted upgrade of plant passives in an area or region, as shown in Figure 3.

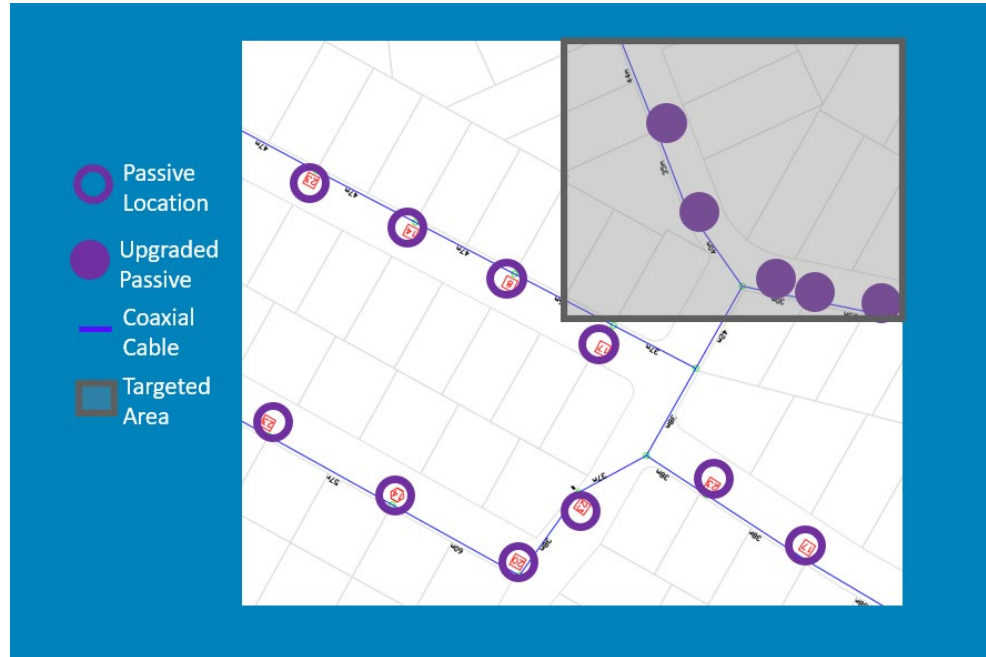


Figure 3 – Blanket Method Illustration

This method is advantageous if the upgrade is complex and requires more than a simple drop-in upgrade. As we will explore further in the paper, upgraded passives may introduce different loss characteristics that will require subsequent changes in the actives portion of the plant.

When using the blanket method, the number of SKUs (Stock Keeping Units) required to maintain the plant doubles. There will be two types of passive devices contained within a system, the existing 1.2 GHz and new 1.8 GHz. Operational teams will be required to stock and maintain equipment for both types of passives. This not only strains local technician resources in terms of truck stock and ensuring the swapping on passives, but it can also exacerbate existing supply chain issues, causing the like-for-like strategy to become unmanageable.




Figure 4 – Blanket Method Truck Stock / SKU

3.2. BAU Method

The second method to consider is the BAU (Business As Usual) Method that we leveraged at Shaw during our mid-split upgrade program. This method is based around making a wholesale switch for the business to a new passive device at a single point in time, across the entire network.

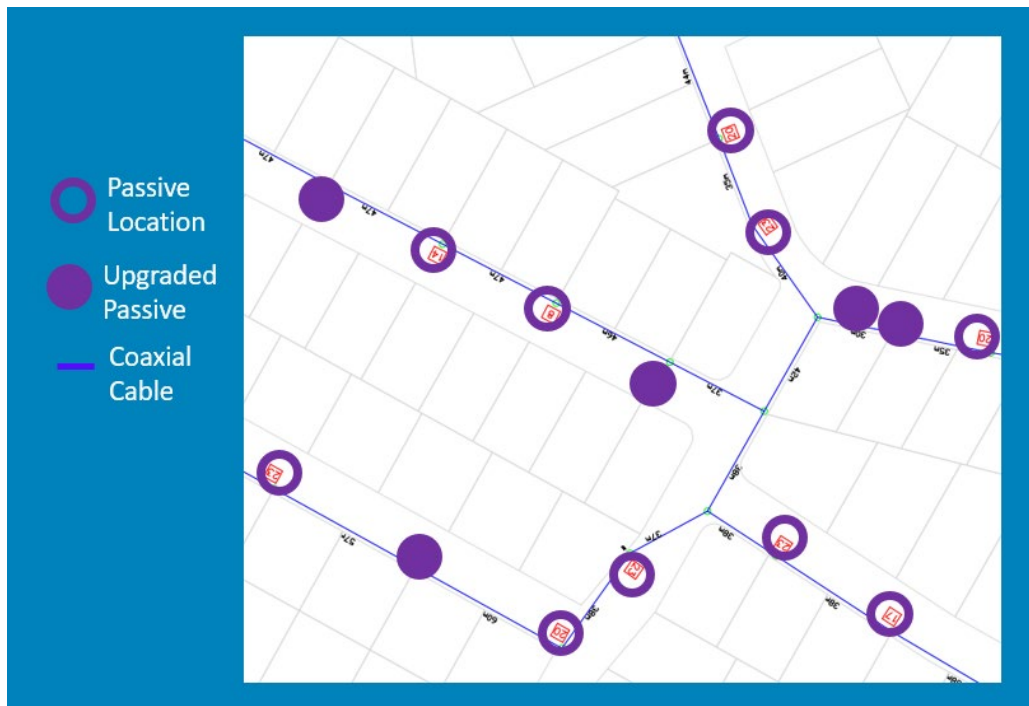


Figure 5 – BAU Method Sample Illustration

This mitigates one of the key operational issues that is seen with the Blanket Method, which is the doubling of the relevant SKUs. This is avoided by doing a straight part-for-part replacement within the inventory system.

Passive	Frequency		Passive	Frequency
Coupler	Legacy	➔	Coupler	1.8 GHz
Power Inserter	Legacy		Power Inserter	1.8 GHz
8 Tap	Legacy		8 Tap	1.8 GHz
12 Tap	Legacy		12 Tap	1.8 GHz

Figure 6 – BAU Method Truck Stock / SKU

The other key operational benefit of this method is that it does not preclude using the Blanket Method when targeting an area ready to offer enhanced service tiers or for an upgrade of actives to complete spectrum expansion.




Figure 7 – BAU Method with Overlay

However, it is important to consider that this method is built upon a drop-in style upgrade, where existing loss characteristics are matched by newly deployed passive devices. Since passives are being replaced on an as-needed basis, they can be dropped into any point in the plant and in any concentration. This creates a scenario in which there is a mix of 1.2 GHz and 1.8 GHz passives in the plant through maintenance and failure replacements. Care and attention will be required to ensure that these are properly documented in the network documentation, or at the very least, that they are easy to physically identify. Additionally, issues may occur if the loss characteristics change because of expanding the spectrum, which would require levels of actives to be reset.

4. Passive-Enabled Spectrum Expansion

When an operator has made the decision to expand their spectrum, there would inevitably be discussions around which equipment to upgrade first—actives or passives. When considering the options in this decision, it becomes clear that the answer is not straightforward—what are the benefits to upgrading the passives when there are no actives to generate the signal, and why upgrade actives without having any passives to pass the signal through?




Figure 8 – Passive Enabled 1.8 GHz Testing

To support the decision-making process, we will argue that there are considerable operational advantages to upgrading passives first that have been, thus far, under explored given that operators have not had a way to characterize or utilize the expanded spectrum.

When we expand spectrum through the use and upgrade of passives, we are able to capture both incoming ingress into the plant, as well as the signals that are already entering the plant on the higher frequencies. This has been noted as we move into the more heavily used LTE frequencies and will only become more notable as the spectrum expands. As operators are becoming increasingly aware, the higher spectral transmissions on the plant can cause egress and be impacted by ingress from smaller physical faults in the cable. It is seemingly no longer the pedestal that has turned into a squirrel nest causing the issue, but an improperly tightened back nut now acting as a slot radiator.

Another often-overlooked benefit is that having a head-start on the passives will minimize downtime and outages as the actives are upgraded in the area due to the larger number of devices in the plant already having been upgraded. This is becoming more and more of a customer experience piece in the work from home era, and with the rollout of more intelligent and power-hungry nodes and amplifiers that may have longer boot times, ensuring that customers' uptime is maximized is becoming increasingly important when making upgrades to the plant.

At the time of this writing, there are massive supply chain issues worldwide that show no sign of relenting. This also feeds into an operator's strategy and is worth including in any evaluation of timing. There has been a shift from an on-demand availability to an as-available environment for network equipment. The impact this has on the strategy for upgrading passives can be felt immediately in lead times for ordering equipment, but also in trying to create and upgrade new plant.

Having a flexible strategy and understanding for implementation of the equipment is a key component of any strategy, especially in the competitive landscape of telecommunications. As operators look to expand their downstream spectrum to reap the benefits of the spectrum available in coaxial cables, there are many options on the path forward, and multiple factors that can impact the decision for which avenue to proceed down first when updating the equipment in the plant to expand services.

5. Conclusion

Every cable operator must continue to upgrade their plant in order to remain competitive in the telecommunications landscape and be able to offer their customers best-in-class services. To do this, a well-considered and evaluated strategy for plant upgrades, and especially for passive upgrades, is vital. There is no longer a single path, or even one simple strategy, when it comes to opportunities to innovate within the telecommunications space, but with that freedom of choice there are often pitfalls, especially as the demands for speed and reliability grow and evolve exponentially. Passives are the foundation on which the rest of the plant is built and allow for our products and services to reach millions of customers around the world every day, and as such, a well-devised strategy can ultimately be the differentiator between competitors.

Abbreviations

BAU	Business As Usual
DOCSIS	Data Over Cable Service Interface Specification
FDD	Fixed Duplex DOCSIS / Extended Spectrum DOCSIS
GHz	Gigahertz
RF	Radio Frequency
TCP	Total Composite Power
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

ANSI/SCTE 265 2021: *Broadband Radio Frequency Hardline Passives for Cable Systems*; Society of Cable and Telecom Engineering

Emerging Trends in Continuous Integration and Continuous Deployment

CI/CD Across Data and Integration

A Technical Paper prepared for SCTE by

Rohit Kuruppath
Lead Architect
Cox Communications Inc
Atlanta
+1 (770) 878-0567
Rohit.K@cox.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Data at Scale.....	3
3. CI/CD Deployment Pipeline	4
4. Hybrid Data Lake and Network Data Ecosystem.....	5
5. Data and Integration CI/CD Lifecycle.....	6
6. Orchestration Architecture	7
7. Conclusion.....	8
Abbreviations	9
Bibliography & References.....	9

List of Figures

Title	Page Number
Figure 1 – Deployment Pipeline	4
Figure 2 – Network Data Ecosystem	6
Figure 3 – Reference Architecture	7

1. Introduction

Continuous Integration and Continuous Deployment (CI/CD) procedures have been extensively used across software development life cycles specifically around application development. The CI/CD implementation around data engineering tools was challenging within the UI centric tool stack. The delivery requirements were not agile without the need of paired programming and common code base. In most cases, the versioning capabilities of the data tools itself was utilized without utilizing the enterprise CI/CD capabilities.

This paper discusses how data engineers can take advantage of the new flexibility that comes with the deployment of software across different data and integration tools, more quickly and securely. It will provide an overview of CI/CD, with a specific focus on how it is being applied across on-premises, cloud and hybrid data environments.

Over the years, real-time data integration and automation have been highly critical for all organizations to meet their analytic needs and providing better insights and capabilities to customers and business leaders. With the huge amount of data to be collected and processed from different systems across the network, before building analytics models around it, there is a pressing need to adopt DevOps guidelines to deliver quality data outputs.

This paper will describe how Cox Communications leveraged the existing CI/CD tool capabilities and expanded upon the configuration tools to support Agile DevOps delivery for building a hybrid data lake. Data Engineers and analysts could leverage a stable pipeline flow and common set of automation tools for delivering the solutions without much onboarding and transitioning into these CI/CD processes and tools.

2. Data at Scale

The transformation of network devices and the possibility of different data points at higher frequency to derive analytical and business insights have been tremendously growing during the last few years and it will continue to evolve throughout decades to come. When exploring the possibilities of big data and cloud technologies, enterprise teams started collecting data points from all the different devices in and out of network. The understanding of 5v's of Big Data – velocity, veracity, volume, value, and variety, allowed data scientists to derive more value from their data while also allowing organizations to become more customer centric, while being proactive to changes in the network.

With the advancement of analytical capabilities in Cox, all the analytics teams across Cox's Technology and Operations organizations were centralized back in 2017. Analytic initiatives around network and customer service health is solely driven by the data across the network and is focused on four key areas.

1. Deep understanding of data that describes quality and use of service by our customers
2. Ability to predict future service impacts and mitigate or prevent them
3. Machines executing the right actions to take, at scale
4. Continuous learning capabilities to improve over time

Different flavors of data from the edge of the network and across the elements of the network, sampled at different intervals, must be accommodated. Formats must be collected, cleansed and processed for both real time and historical trending. Analytics are required to have more accuracy and lesser development

turnaround time. Hybrid data lake environments coupled with traditional reporting and analytical solutions have resulted in a diverse tool stack for the developers to work with.

The traditional extract, transform, and load (ETL) tools, along with the data warehousing technologies and relational databases, which used to be the primary tool kit for the data engineers, took second place with the new advancement in big data and cloud technologies. Data engineers ended up building and supporting solutions across different technologies and programming languages like the software programming application development. Data engineers are using traditional ETL tools like Informatica, Alteryx for the structural data transformations, along with big data technologies like Apache Spark and using Scala, Python to process unstructured data sets from sources and cloud services for their hybrid solutions.

3. CI/CD Deployment Pipeline

CI/CD are a series of steps that must be performed in order to deliver a new version of software. A pipeline is a process that drives software development through a path of building, testing, and deploying code. The objective is to reduce human error and maintain a steady process to release software through automation. Figure 1 below depicts the continuous integration and continuous deployment pipeline flow here at Cox for a traditional software development application life cycle.

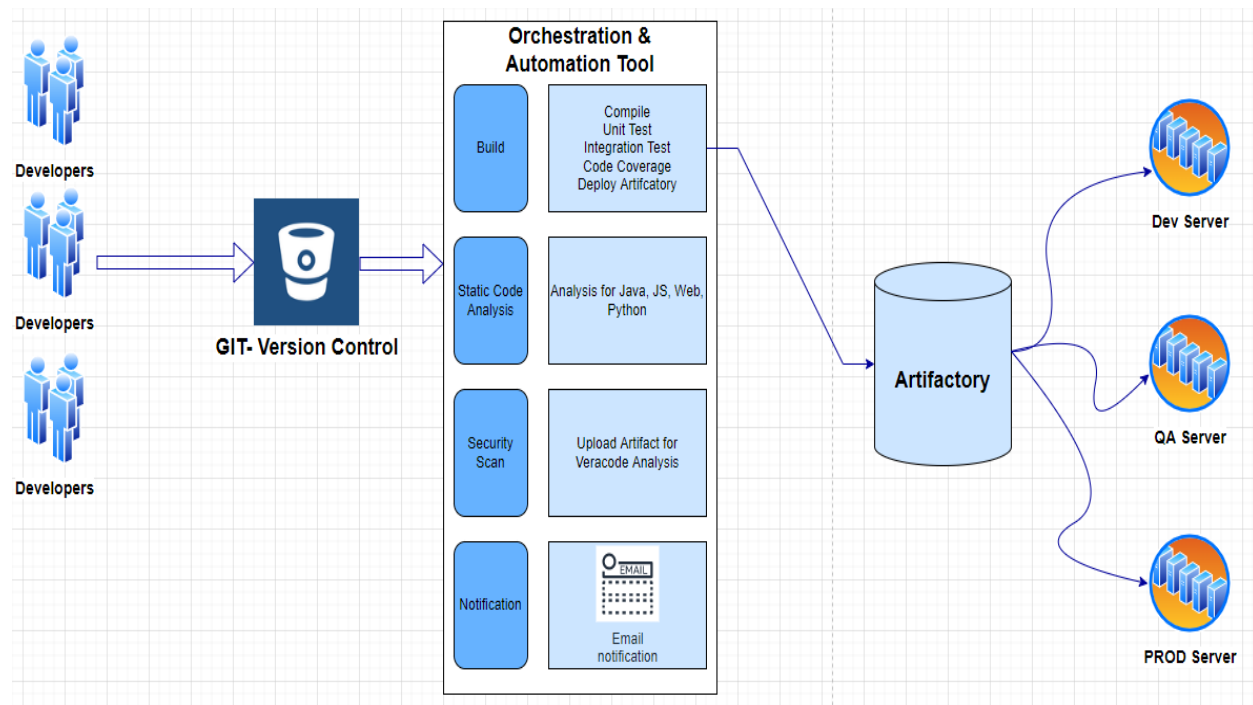


Figure 1 – Deployment Pipeline

In a GitOps-driven deployment process, all activities emanate from the version-controlled code repository. Here in Cox, we drive the deployment process directly within the code repository or version control service BitBucket. Storing the deployment instructions in a central version control repository provides a single source of truth for all deployment activities. Using the repository as a single source of

truth provides reliable change management and auditing capabilities. Version control and access security are also built into the service.

Orchestration tools play the role of managing the deployment process, which can range from a single task – such as running a unit test case against a source code – to a complex deployment process embodying many tasks. When an automation is triggered, it will refer to the playbook in the deployment repository stored within the version control and execute the tasks. Artifactory's build integration to orchestration tool allows our build jobs to deploy artifacts and resolve dependencies to and from Artifactory. Artifactory is a centralized binary management system providing management binary artifacts generated and used by the organization.

4. Hybrid Data Lake and Network Data Ecosystem

Data from different source systems gets loaded into the operational data store either through database (DB) replication or ETL loads built by data engineers over the years. These were built in different tool sets including Informatica, Oracle and loaded into relational databases. With the onset of big data technologies and possibilities of analyzing unstructured data assets, more data points at the least possible frequency started getting streamed (using streaming applications like Apache Kafka, Spark Streaming), pulled from sources on a specified interval (using tools like Apache Sqoop, Spark Batch) and even written into Hadoop Distributed File System (HDFS) for further processing and consumption.

Data engineers started working across multiple tools to build this ecosystem leveraging on-prem ETL processes to relational databases and big data lake for the unstructured data assets for analytical insights. There are different production solutions to meet both real-time and batch reporting needs in place now that are built across different tool stacks that consume data from systems like ticketing, polling platforms and so on. Before the expansion into on-prem big data lake and onto the cloud data lake, we realized the need to have a stable automated deployment process with a common CI/CD tool stack to support most of our needs and deliver in an agile fashion.

Traditional batch loads to meet the enterprise reporting through an enterprise data warehouse (EDW) which have been a stable process for more than a decade with 1000s of daily loads handling data assets from billing, network alarms, construction workflows, ticketing, customer calls, chats, geospatial assets, and many more. At the very beginning of defining the requirements for a more agile delivery and more frequent data needs, the team invested on setting up and configuring an automated CI/CD pipeline, as well for the traditional ETL tool and databases working with the enterprise software configuration management (ESCM) team.

Within the last four years, as Cox Communications started the data cloud journey, most of the on-prem ETL, big data deliveries are managed through CI/CD pipelines that provide more visibility into the source code. The CI/CD pipeline is auditable for security scanning and vulnerabilities and provides the developers the ability to follow a common deployment process for easy handoff for reviews and deployment across all the different tool stacks. End-to-end release cycles starting from the source to the very end at the data lake, are all getting managed through CI/CD pipelines that are configured and managed through individual application and data teams. This makes any changes seamless and controllable, even for an easy rollback in the case of faulty deployments.

For example, the polling data from the modems through a CMTS/CCAP devices goes through all different layers including streaming applications, an on-prem data lake, cloud services for analytics/machine learning models, and the traditional warehouse for enterprise reporting and integration. A change to a device element or the addition of new pollers, would have been a very intensive software

development life cycle (SDLC) process, if the end-to-end architecture, including all data spaces, are not able to support CI/CD processes. With the cloud journey, data is getting processed across multiple cloud accounts based on the department needs for different use cases. The codebases involving streaming applications to ingest raw data using Spark Streaming and landing them into file systems, ETL processing and analytical solutions using Python, PySpark and machine learning models are all stored in an on-prem code repository for orchestration for their respective environments. This provides the ability for development leads and architects across respective organizational units to have a better view of the solutions, and an ability to review and commit to the codebase for production deployments through a single CI/CD pipeline.

Because different servers are involved across the tool stack for ETL, big data processes also pushed the need of automation for the environment life cycle management to include upgrades and server additions. This reduced the need for downtime during the data tool stack upgrades and server additions. This is critical for the proactive service and health analytical solutions that are built on top of these data assets where data unavailability over a time period will have impacts to the model outputs and customer notifications. Figure 2, at a high level, depicts the network data assets and their data processing through a relational database, big data streaming applications, and ETL tools. Enterprise cloud data lake accounts for the hybrid data lake for enterprise consumption, storage and processing serves the downstream consumer accounts for their analytical models and outputs.

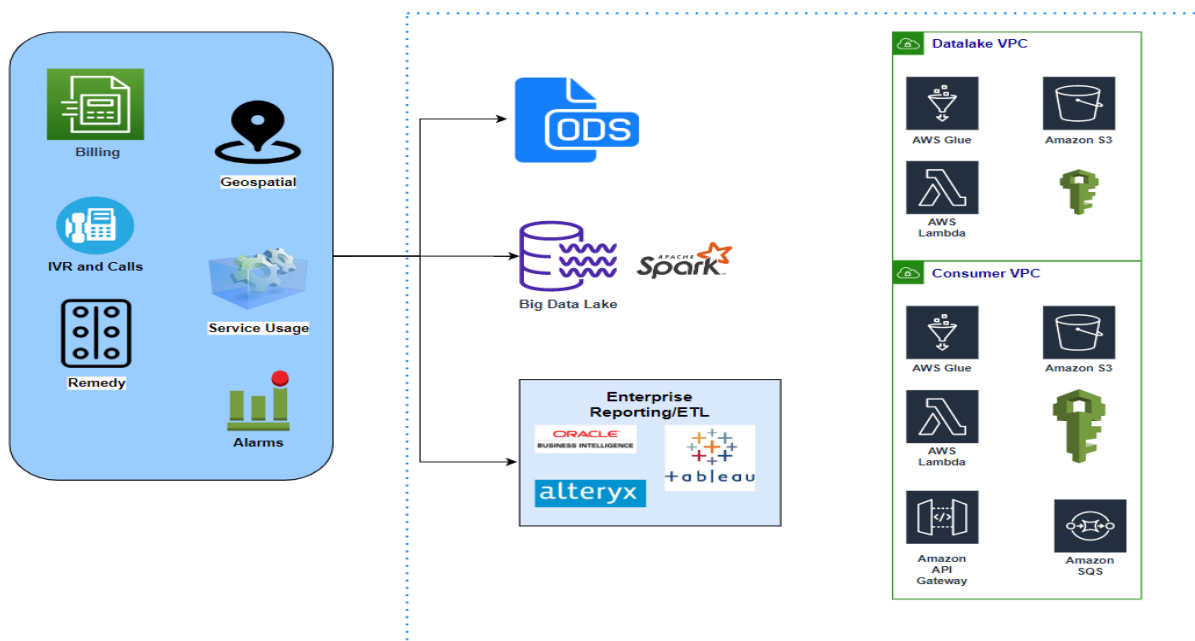


Figure 2 – Network Data Ecosystem

5. Data and Integration CI/CD Lifecycle

At Cox, we have an ECSM team to provide automation services that increase the velocity of development and operations teams across Cox to drive value to the business. DevOps teams across Cox collaborate with the ECSM team to define a deployment pipeline and support the right set of tools for delivering a CI/CD process to help the development teams to deliver the software or infrastructure as a code easily and

securely. Compared to the earlier days of infrastructure automation and traditional application development, the ESCM team plays a pivotal role in identifying, configuring and supporting different automation tools to meet the needs of engineers across data, automation, integration, application development, databases, and infrastructure.

Working with the ESCM team, the first step towards this journey was the identification of a unified platform and pipeline configuration for data engineers that provided minimal overhead for developers and did not impact the end users or customers. Starting with the ETL tool stack, where we had a built-in tool for deployment automations and validations, we tried reusing that stable process orchestrated through BitBucket and Jenkins with the goal of making it easy to use for data engineers new to the programming world and CI/CD tools. CI/CD tools in this case acted as a wrapper process to invoke ETL built in code validations, test scripts, and deployment processes, while also setting up a repository and pipeline flow for unified deployments.

The approach of building an initial prototype and establishing the pipeline configuration was followed during the big data and cloud journey to identify the right set of CI/CD tools and process that is supportable and extensible for different services within these cloud providers. The upfront design and setup time spent with the ESCM team helped the ongoing development teams to reuse the templates, Ansible scripts, and pipeline flows to build their own ETL solutions with minimal changes. User adoption is a key factor for a successful CI/CD model along with the DevOps enablement for agile program delivery. Individual teams will be responsible to maintain the Git projects along with their code base and delivery, while the ECSM team supports the tools and the environments to make this happen for DevOps teams. This allows enterprise teams to reuse the CI/CD tool stack, environments, and to follow standard guidelines along with the enterprise security checks on the codebase and processes.

6. Orchestration Architecture

Enterprise teams at Cox Communications have been using Bitbucket and Jenkins for code source control and automation/orchestration for the CI/CD flow as explained in Figure 3. Ansible is used for automating operations and building templates for repetitive and reusable functions. Figure 3 shows the high-level architecture and tools involved for an AWS Cloud deployment pipeline flow.

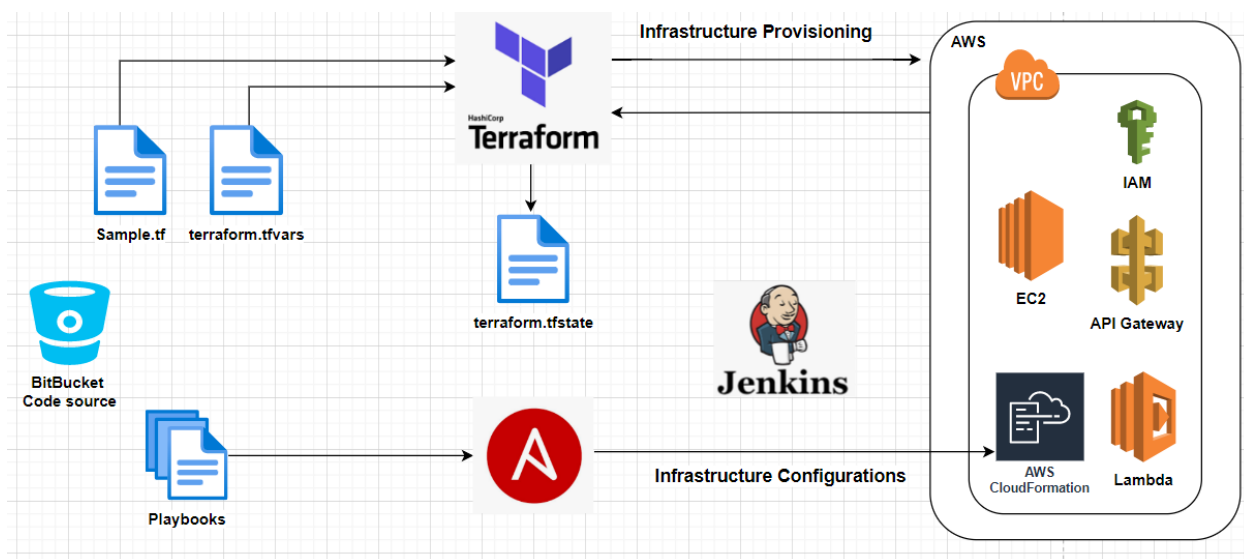


Figure 3 – Reference Architecture

Bitbucket is the source code controller for developers to commit their code. We are following the terraform deployments for infrastructure provisioning in the cloud, including the role provisioning, user accounts. Terraform variable files, modules, and the Ansible playbooks are all stored in the source control tool. Jenkins orchestrates the deployment automation process by invoking the cloud formation templates through Ansible scripts for serverless application deployments such as AWS Glue and AWS Lambda. These serverless processes are heavily used for the integration and ETL processing. Terraform records the state files and compares them against the requested resources while running the plan and deploy the infrastructure as a code to the requested AWS instance.

For the on-prem deployments, we have the built in Informatica power center command shells that manage the deployment through the enterprise supported CI/CD tools Bitbucket and Jenkins. This helps the developers to maintain a common codebase and repository while reusing the deployment pipelines wherever possible. Similarly, the big data applications across different services and tools like Spark Streaming, batch Spark, HBase, Python applications, are all deployed following the above architecture. The exception to this is the on-prem Hadoop file systems that are orchestrated through Ansible automations instead of AWS Terraform and AWS CloudFormation. This provides the developer a simplified and well-established deployment automation pipeline that is configured upfront and can be extended based on the specific use cases.

7. Conclusion

The journey into expanding a culture of Continuous Integration and Continuous Deployment into the Data Integration space has been one that mirrors the actual process of continuously iterating and deploying. Adoption of existing capabilities into the data space, and the partnership with enterprise configuration management teams to extend the capabilities for different software development tools, helped us to build a robust and useful CI/CD process.

People, process and technology are three pillars of a successful delivery model. Rolling out CI/CD into data space was done without jeopardizing the customer experience with frequent releases and lesser overhead on software developers to adopt the process across different tool stacks. Zero down time server upgrades, what would have once been a unique phenomenon, is now a common practice for data platforms with the introduction of standby clusters. Those lessons are now able to be transferred into other sides of engineering, to help drive improvements at a wider scale and without the need of constant manual intervention.

I would like to take the time to thank those who lent their time and expertise into the completion of this paper and presentation – Lisa Pinkey Coleman, Kesha Heard, Greg Garrison and Chandrasekhar Anne.

Abbreviations

AWS	Amazon web services
CI/CD	continuous integration and continuous deployment
DB	database
DevOps	development and operations
EDW	enterprise data warehouse
ESCM	enterprise software configuration management
ETL	extract, transform, and load
HDFS	Hadoop distributed file system
SDLC	software development life cycle
UI	user interface

Bibliography & References

1. Informational interview, Chandrasekhar Anne, Lead Systems Engineer, Cox Communications Inc
2. Informational interview, Kesha Heard, Lead Architect, Cox Communications Inc
3. “Multi Region Deployment in AWS with Terraform”, by Yadav Lamichhane
4. “Terraform: Beyond the Basics with AWS”, by Josh Campbell and Brandon Chavis

Empowering Smart Communities with a Digital Twin

Why detailed 3D city models are critical for enabling modern IoT solutions

A Technical Paper prepared for SCTE by

Scott Casey
VP of Sales
Cyclomedia Technology, Inc.
8215 Greenway Blvd., Middleton, WI
720-289-1399
scasey@cyclomedia.com

Table of Contents

Title	Page Number
1. The Case for a 3D Digital Twin	3
2. Key Components of a Smart Community Planning and Engineering Solution	3
2.1. 3D model of the real world	4
2.2. High precision and accuracy	4
2.3. Immersive online digital environment	4
2.4. Integration with GIS and design tools	5
2.4.1. Support for 3D data	6
2.4.1. Open API	6
2.5. Timeliness of information	6
3. Technical Considerations and Specifications	6
3.1. Street level imagery & LiDAR capture	7
3.2. Imagery & LiDAR integration and positional accuracy	8
3.3. Data analytics: asset and feature extraction	8
4. Conclusion	9
Abbreviations	11
Bibliography & References	11

List of Figures

Title	Page Number
Figure 1 – High Resolution Imagery & Extracted Assets	5
Figure 2 – Example of Vehicle-Mounted 5-Camera, LiDAR and GPS System	8

List of Tables

Title	Page Number
Table 1 – Commercial Data Capture Options	6
Table 2 – Asset and Feature Extraction Examples	9

1. The Case for a 3D Digital Twin

MSO's have two powerful and unique advantages towards delivering Smart City/Community programs – an expansive high speed, low latency network footprint and tremendous experience and expertise in the local markets they serve. However, delivering these expansive projects at scale, including Fiber-to-the-Premise/Home (FTTP/H), fixed wireless access, streetlight modernization, digital kiosks, parking, accessibility, and other intelligent Internet of Things (IoT) deployments presents unique challenges far different from the traditional cable business.

In the traditional business, a “best efforts” approach to data with uneven quality and currency, for example two-dimensional (2D) Geographic Information Systems (GIS) maps that are out of date, in many cases have been “good enough.” In contrast, connected community programs require up-front three-dimensional (3D) information models covering buildings, sites, surface features, existing infrastructure, and ROW (right of way). These models need to be current, hyper-detailed, and comprehensive, and they must offer pinpoint accuracy. As a result, the traditional approach to data gathering with manual field walkouts is time inefficient, costly, and doesn't meet the new requirements for smart community planning, IoT engineering, and project execution.

To effectively address the “Smart City” challenge, a new and innovative approach to capturing field data – infrastructure, assets, and site conditions – is required to support the more precise requirements of fixed and wireless IoT deployments. This includes high resolution 360-degree imagery, dense LiDAR (Light Detection and Ranging) point clouds, and machine learning. Working with Cox2M, a division of Cox Communications, and building on the latest technologies, R&D efforts, and project experience across several relevant use cases, a blueprint for a smart community enablement platform emerged, largely enabled by the power of a detailed and accurate 3D digital twin.

This technical paper will outline the business rationale and specific use cases related to new MSO initiatives in fiber-to-the-home or -premise (FTTH/P), wireless, and private-public partnerships, and then dive into the technical requirements, challenges, and implementation of a 3D digital twin to accelerate and meet the end-to-end deployment needs of leading edge, connected community projects. Key learnings and best practices will be highlighted, along with performance metrics from actual projects – such as cycle time reduction, data quality improvements, and cost savings – that quantify the true benefits and impact of developing and integrating a 3D digital twin for field data automation.

2. Key Components of a Smart Community Planning and Engineering Solution

There are several high level requirements that are unique to smart community IoT planning and engineering (P&E) projects. These include the following key capabilities and attributes:

- 3D model of the real world
- High precision and accuracy
- Immersive online digital environment
- Integration with GIS and design tools
- Timeliness of information

Here is a deeper dive into the five essential components of the digital twin solution for smart community/IoT planning and engineering.

2.1. 3D model of the real world

Any assessment of the existing field conditions, such as infrastructure, relevant assets, surface features, and for that matter everything within the right of way (ROW), has to be presented in a three-dimensional context. This means both imagery and LiDAR are needed in order to produce a digital 3D representation of the true field conditions. The primary reason for this is that optimal placement of IoT-enabled devices and supporting connectivity (fixed and wireless) requires comprehensive and detailed horizontal and vertical references to address factors such as line of sight, clutter, visibility, accessibility, form and fit, and aesthetics. Even something as simple as taking measurements need to be performed in 3D space using x, y and z coordinates. A traditional 2D GIS map falls flat!

2.2. High precision and accuracy

Engineering smart kiosks, lighting, sensors, fixed wireless access (FWA), Wi-Fi coverage, and pre-connected FTTH cables and drops, requires very accurate site information and high precision reference data. GIS maps are typically based on legacy information that has been converted and/or migrated over time with poor positional accuracy, at best 2 to 3 meters but often much worse. Taking measurements, which are dependent upon the precision or relative accuracy of the data, cannot be trusted for engineering tasks. This means that a field survey, often referred to as “boots on the ground” or “field walkout”, is required but this takes a lot of time and expense. The results have worked adequately for traditional cable and telecom engineering for many years, but this approach does not deliver the on the imagery and 3D model requirements, and it does not provide an immersive digital twin (see next section).

2.3. Immersive online digital environment

An extremely important part of the smart community solution is the notion of a “virtual field walkout”. This refers to the ability for project stakeholders to visualize the real world environment from an application on a connected device such as a PC, laptop, tablet, or smart phone – in essence a digital twin. There are many benefits to having this capability, but the primary ones are:

1. Eliminate or reduce time, cost, and risk of “boots on the ground” field work
2. Ensure high quality measurements and calculations based on true 3D data
3. Collaborate and solve problems quickly and confidently
4. Avoid remedial work and mistakes related to poor or missing field information



Figure 1 – High Resolution Imagery & Extracted Assets

A large international engineering firm refers to high resolution street level imagery and LiDAR as “the single source of truth” for planning, high level design (HLD), low level design (LLD), and permitting.

2.4. Integration with GIS and design tools

Planning and engineering work is typically performed using GIS, computer-aided design (CAD) or other geospatially-enabled software tools. Digital field data must be integrated with these tools for productivity, accuracy, and completeness. There are two core requirements for the integration of a 3D geospatial digital twin with GIS and CAD.

2.4.1. Support for 3D data

The GIS or CAD platform being used for smart community/IoT projects must support 3D data management and visualization. Platforms that are used extensively for managing and designing in three dimensions include Esri ArcGIS and ArcGIS Online, and AutoCAD. In addition, the platform must allow high density LiDAR point cloud files and extracted 3D asset/feature vector data layers (points, lines and polygons) to be automatically loaded into the database.

2.4.1. Open API

There are two application programming interface (API) capabilities that are needed for effective integration of imagery/LiDAR and extracted 3D asset and feature data. The first is a 360-degree viewer plugin that works interactively within the system's user interface (UI) so that designers can perform tasks in an immersive 3D workspace. These sorts of plugins are normally set up using Java Script or Practical examples of this include recording measurements and placing proposed design elements in three-dimensional space with x, y and z coordinates. The second API capability that's important to the overall usability of digital field data for design and permitting is live rendering of third party data such as LiDAR, imagery, and vectors in real-time. This would be the approach used for data that is stored on the cloud and made accessible using a standard file format (such as LiDAR .LAZ, or vector DGB or Shapefile), or as an online internet protocol like WMS (web map service) for map tiles or WFS (web feature service) for geospatial data records.

2.5. Timeliness of information

Having the ability to capture field data where and when needed is critical to the overall smart community solution process, sequence of events, and project timeline. Ideally, field data for planning and engineering should be captured within 3-9 months of HLD/LLD and permitting, and no more than 12 months old. Similar to mobilizing field workers to inspect, survey and inventory field information, this means that imagery and LiDAR data capture needs to be mobilized early in the project cycle since it's required as input to planning, engineering and permitting activities. In essence, the field data needed to build a comprehensive 3D digital twin model of the real world must be performed on demand for the target area of interest (AOI).

3. Technical Considerations and Specifications

The foundation of creating a 3D digital twin of the outside plant (OSP) environment is the recording system. There are several options when considering field data capture, such as the method of capture, the type, frequency, resolution, accuracy, and level of effort.

Here is a simple matrix showing the available commercial options for field data capture supporting smart community use cases:

Table 1 – Commercial Data Capture Options

Data Capture Options	Type(s)	Method	Frequency	Resolution and/or Accuracy	Level of Effort*
<i>Fixed Wing Aerial</i>	Ortho & oblique imagery, LiDAR	Low altitude at slow speed, grid pattern flight plan	On a set schedule	30- to 150-megapixel imagery, 50 to 100 points-per-square-meter LiDAR, 0.1 to	Low

				1.0 meter positional accuracy	
<i>Drone</i>	Oblique & 360° imagery, some LiDAR options	Flown with ground-based pilot, requires permit or authorization	On demand	30-megapixel imagery, 500+ points-per-square-meter LiDAR, 0.2 to 0.5 meter positional accuracy	High
<i>Street Level</i>	360° & ortho imagery, LiDAR	Vehicle mounted, follow public and private roads, parking lots, etc.	On demand	30- to 100-megapixel imagery, ~1,000 to 2,000 points-per-square-meter LiDAR, 0.1 meter positional accuracy	Medium
<i>Backpack/hand-held</i>	Conventional & 360° imagery, LiDAR options	Heavy backpack, trolley or hand-held device	On demand	10- to 20-megapixel imagery, 1,000+ points-per-square-meter LiDAR, 0.2 to 1.0 meter positional accuracy	High

* Level of Effort is based on the time and cost per unit of measure, such as square miles or linear miles.

Based on the above criteria along with testing, trials and production projects performed by companies like Cox2M (a division of Cox Communications), Verizon OneFiber and Wireless, Ledcor Technical Services, Byers Engineering, and CalComm Consulting, street level imagery & LiDAR capture is best suited to the specific requirements of IoT and communications planning and engineering projects. The factors that have contributed to this conclusion are imagery/LiDAR resolution and quality, data capture flexibility, speed of capture and post-processing, overall cost, and ease of integration.

3.1. Street level imagery & LiDAR capture

The on-going evolution and recent innovations in vehicle-based street level imaging has enabled the creation of a geospatially accurate 3D digital twin. A great example of this is a recording system comprised of five individual cameras that fire off in sequence to the front, right, up, left, and rear as the vehicle crosses theoretical recording locations, also known as recording points, that are taken 5 meters apart as the vehicle drives and captures the AOI. The source images are merged using a patented process, creating a “GeoCyclorama” – a seamless, parallax-free, spherical, high-resolution, panoramic image taken at street level. At 5-meter (16.4 feet) intervals, the GeoCyclorama” is generated covering the entire road

network for the target AOI. The resulting 100-megapixel resolution imagery, in which each pixel location is identified, provides users and systems with imagery that is both incredibly detailed and accurate.



Figure 2 – Example of Vehicle-Mounted 5-Camera, LiDAR and GPS System

The 100-megapixel resolution provides a full 14,400x7,200 resolution image. There is a 0.025 arc degree separation between pixels which translates to ~4.0mm at 10 meters, and ~6.5mm at 15 meters, between pixels. At 10 meters this is roughly a density of 62,500 pixels per square meter on a vertical plane such as a wall.

The LiDAR scanner is running the entire time so that the imagery and point cloud data capture is aligned and covering the exact same AOI. A high specification global positioning system (GPS) is integrated with the cameras and LiDAR sensor to ensure the best possible position identification and recording.

3.2. Imagery & LiDAR integration and positional accuracy

GeoCycloramas are high-resolution, parallax free images with a built in spatial component providing not only a precise location for each 360-degree image but an x, y & z coordinate for every pixel within the image. The patented capture technology allows for the precise pixel arc separation mentioned previously, which allows the system to accurately tie the imagery and LiDAR together. With the level of integration, the LiDAR point cloud is colorized based on the RGB (red, green, blue) values from the imagery. This is a very important part of the unique and innovative approach – any other process where imagery and LiDAR are collected separately will not deliver the accuracies and immersive capabilities of the GeoCyclorama, even if it's possible to integrate them after the fact (which is extremely uncommon). During this post-processing stage, other positional improvements are applied through steps such as geospatial positioning alignment with a reference datum (control points) or external model, and relative positioning improvement (RPI) that utilize overlapping datasets to improve the data even further.

On average, the absolute positional (x, y & z location) accuracy is +/-10cm at 1 standard deviation. The relative accuracy, for example for taking measurements like pole spans, height of attachments, pavement width, concrete surface areas, is on average +/-2cm at 1 standard deviation.

3.3. Data analytics: asset and feature extraction

Extracting assets and other features from the imagery and LiDAR data is generally called data analytics. Data analytics tools and processes include a combination of machine learning (ML), artificial intelligence (AI), automation, partial automation, and human review and input from skilled analysts and subject matter experts (SMEs). Over time the algorithms improve, and the level automation increases based on

lessons learned on production projects and as feature recognition libraries continue to be further developed.

Data analytics and extractions run through a proprietary and optimized platform that provides a host of automated, semi-automated, and manual tool sets to efficiently detect, extract and deliver a quality asset and feature-rich data analytics product. Over time the teams working on next-generation communications, IoT and smart city projects have built up an extensive library of ML models that can detect common infrastructure, roadway, roadside, architecture, and public works features from the 360-degree imagery and place the geometric representation accurately within the LiDAR point cloud data. The data analytics process has matured to become an entire workflow that includes project tracking, data management, automation, QA/QC, and client delivery to ensure that a quality product is maintained and that consistently meets the service level agreement (SLA) and required specifications.

To show the extent of the collective team's efforts, below are examples of some of the most widely used data analytics extraction assets and features, called data dictionaries:

Table 2 – Asset and Feature Extraction Examples

Pole Details	Overhead	Ground Features	Right of Way
Pole Base & Top	Span	Cabinet	Edge of Pavement
Material	Midspan	Manhole	Back of Curb
Cross Arm	Lowest Vertical Clearance	Handhole	Sign & Structure
Guy Wire POA	Power Vertical Clearance	Vault	Traffic Signal
Guy Wire Anchor	Primary Conductor	Pedestal	Bus Shelter
Equipment POA	Comms Cable	Wall	Sidewalk
Streetlight POA	Streetlight	Fence	ADA Ramp
Power POA		Building	Lane Markings
Comms POA			

4. Conclusion

The teams at Cyclomedia, Cox2M, and the other service providers and engineering firms – not to mention numerous local governments, cities, counties, and municipalities – have firsthand experience that collectively adds up to hundreds of projects and hundreds of thousands of miles of imagery and LiDAR data capture and processing. The use cases that have been successfully delivered include long-haul fiber, FTTH/P, 5G planning, pole inventory and engineering, joint use reconciliation, tax assessment, first floor elevations, sign condition and inventory, accessibility, and highway sign and roadway condition

assessment. Several best practices have emerged and more continue to be developed all the time. The most important findings are summarized as follows:

Remote sensing technologies have advanced so much in the last 3-4 years that imagery and LiDAR solutions have become a viable alternative to manual field data capture, including use cases like pre-engineering walkouts, visual inspections, infrastructure and inventory surveys, and condition assessment.

High quality and precise imagery, LiDAR and data analytics can replace 80-90% of traditional boots on the ground field data capture and accelerate project cycle times by 50% or more. Safety is improved while risk and liability are reduced.

Street level data capture is best suited to smart community, communications, and IoT planning and engineering due to the close proximity to assets, infrastructure and features, high data quality, immersive experience (virtual fielding), speed and ease of capture, and flexibility of the AOI.

Machine learning and automation are continuously able to handle more and more asset and feature extractions over time, but human subject matter experts are essential to training the routines, expanding the capabilities, and ensuring that quality and completeness goals are met.

Street level data is a foundational component of creating a 3D digital twin city model to support smart community, IoT and communications planning, engineering, permitting, and construction.

Cyclomedia and Cox2M are looking at short-term requirements to support additional assets and features along with closer integration with P&E workflows and tools. Longer term, the companies and other partners are looking to automate a greater number of assets and features, streamline the integration to the design and project management systems, and look to incorporate additional data sources into the analytics process. The Cyclomedia and Cox2M leadership teams welcome feedback from other SCTE, CableLabs, and NCTA members, and are open to development collaboration, lab testing and functional pilots to advance the 3D digital twin capabilities for smart community and IoT projects and evolving use cases.

Abbreviations

2D	Two-dimensional
3D	Three-dimensional
AI	Artificial Intelligence
AOI	Area of Interest
API	Application Programming Interface
FTTP	Fiber-to-the-Premise
FTTH	Fiber-to-the-Home
GIS	Geographic Information System
HLD	High Level Design
IoT	Internet of Things
LiDAR	Light Detection and Ranging
LLD	Low Level Design
ML	Machine Learning
OSP	Outside Plant
P&E	Planning and Engineering
ROW	Right of Way
RPI	Relative Positioning Improvement
SCTE	Society of Cable Telecommunications Engineers
SLA	Service Level Agreement
SME	Subject Matter Expert
UI	User Interface
WFS	Web Feature Service
WMS	Web Map Service

Bibliography & References

Wikipedia – Web Map Service. A standard protocol developed by the Open Geospatial Consortium in 1999 for serving georeferenced map images over the Internet. These images are typically produced by a map server from data provided by a GIS database. See https://en.wikipedia.org/wiki/Web_Map_Service.

Wikipedia – Web Feature Service. In computing, the Open Geospatial Consortium Web Feature Service (WFS) Interface Standard provides an interface allowing requests for geographical features across the web using platform-independent calls. One can think of geographical features as the "source code" behind a map, whereas the WMS interface or online tiled mapping portals like Google Maps return only an image, which end-users cannot edit or spatially analyze. See https://en.wikipedia.org/wiki/Web_Feature_Service.

Encourage EVERY Employee to Learn and Utilize Data, Analytics, and Machine Learning (DAML)

A Technical Paper Prepared for SCTE by

Robert Gray Wald, MS
Explorer Initiatives--Technology Innovation and Strategy
SCTE[®] a subsidiary of CableLabs[®]
BWald@SCTE.org
[LinkedIn.com/in/BobWald](https://www.linkedin.com/in/BobWald)

Table of Contents

Title	Page Number
1. Introduction.....	5
1.1. A tale of two CEOs.....	6
1.2. Paper approach.....	7
2. WHY DAML must become your organizational thinking and “vibe”.....	7
2.1. How the world’s leading companies are cultivating a new business culture—Case Study JLL.....	8
2.2. Motivations for making DAML an enterprise wide initiative.	9
2.2.1. DAML adoption is already exceedingly high.....	9
2.2.2. New DAML knowledge areas are exponentially increasing.....	9
2.2.3. Establishing a “data-driven culture” is the biggest contributor to DAML success.....	10
2.2.4. DAML must become EVERYONE’S strategy.	10
2.3. DAML value for <i>me</i> as an employee, and how will it make my job easier?	11
2.3.1. YOUR current job may cease to exist.....	12
2.4. How can my organization benefit?	13
2.4.1. How does DAML reduce and eliminate errors?	15
2.4.2. Greater customer insight and retention.	15
2.4.3. Retention and recruiting	15
2.4.4. Increased performance through DAML organizational obsession.....	17
2.4.5. Collaboration and working in teams.....	18
2.5. How can the cable industry benefit?	18
2.6. Ignoring DAML will doom existing companies.	18
3. WHAT DAML areas to initially address.....	19
3.1. A quick aside on “determinant” vs. “assisted” DAML processes.	20
3.2. Focus on data literacy.....	20
3.3. “There is No AI without IA (Information Architecture).”	22
3.3.1. Create data dictionaries that span the organization and align with industry.	22
3.3.2. Seeking Common Acronyms	23
3.3.3. Taxono Mies? Taxi please.....	23
3.3.4. Aim for INFORMATION CONTEXT.....	24
4. HOW DAML can be understood, embraced, and utilized by all employees.....	24
4.1. Start both “Fast and Furious” and “Slow and Steady”?	25
4.1.1. Create a grassroots-led effort.....	25
4.1.2. Dub DAML development director.....	25
4.1. Methods to infuse data driven competencies.....	27
4.1.1. Publicize DAML program and add participation into job evaluation criteria.	27
4.1.2. “Tiger Team” taming.	27
4.1.3. “DAML Academy”?.....	28
4.1.4. Form “DAML Translators” to bridge technology and business realms.	28
4.1.5. Develop communications and training to inform organization and external stakeholders.....	29
4.1.6. A little DAB’ll do ya!.....	30
4.2. Brown bag training.	30
4.2.1. Find mundane and repetitive tasks and prioritize.	31
4.2.2. Create a competition.....	31
4.2.3. Present chosen automation processes to DAB.	31
4.3. Constantly sharpen the hard and soft skills of technical personnel.....	32
4.3.1. Incentivize technical team participation at local DAML groups and events.....	32
4.3.2. HOST a DAML event.	32
4.3.3. Crossbow training?	32
4.4. Optimizing the learning experience with training and certification programs.....	33

4.4.1.	Use DAML to optimize DAML training.	34
4.5.	DAML big picture ideas.	35
4.5.1.	Showcase individual contributors or teams that are exceeding expectations.	35
4.5.2.	Publicize DAML effort to all customers, shareholders, on website, and in press releases.	36
4.5.3.	Alter all hiring descriptions to focus on measurable DAML skills and accomplishments.	36
4.5.4.	DAML Program integrates with organizational change management.	36
4.5.5.	Tie DAML program into innovation initiatives.	37
5.	Conclusion.	37
	Abbreviations	38
	Bibliography & References.	39


List of Figures

Title	Page Number
Figure 1: Wordcloud of DAML Technologies and Applications.	5
Figure 2: intensely deep programmers on dark screens in dimly lit cubicles,	6
Figure 3: neuroscientists to optimize all learning, behavior, and processes,	6
Figure 4: or automation executives that desire very few humans doing all complex work with their single push-button fingers;)	6
Figure 5: Hidden Figures' Dorothy Vaughn Learns New Computer Technologies Before It Replaces Her Team	8
Figure 6: DAML Production Deployments Increased from 35% in 2019 to 47% in 2021.	9
Figure 7: Gartner Hype Cycle for Data Science and ML July 2021	9
Figure 8: Areas Critical to DAML Success.	10
Figure 9: When an Organization Embraces DAML, Various Team Strategy Components Must Be Implemented.	10
Figure 10: Importance of DAML Skills by Department.	11
Figure 11: Gartner's Top Data and Analytics Trends for 2022 (Partial)	12
Figure 12: Funny Because it is True	12
Figure 13: Employment Changes from Automation 2017-2030.	13
Figure 14: Impacts of Employee Training Programs	14
Figure 15: Data Source: Gartner's TalentNeuron, and represents jobs posted between 4/22/2021 4/22/2022	16
Figure 16: Every Job Function has Adjacent Skills and Roles. Example Shown for Cloud Engineering Skill Adjacencies	17
Figure 17: What You Don't See Can Kill You (and Your Organization).	19
Figure 18: Operating Model of a DAML Organization.	19
Figure 19: Always start with assisted DAML processes	20
Figure 20: Despite Company Efforts Most Do Not Have Adequate Data Skills.	21
Figure 21: Forrester Approach to Building a Data Literacy Program.	21
Figure 22: Most DAML processes fail because of data preparation and access.	22
Figure 23: "Garbage In, Garbage Out" (GIGO).	22

Figure 24: Insert Some Humor When Content Appears Boring or Irrelevant.....	23
Figure 25: Dewey Decimal Library Classification System	24
Figure 26: Grass Fed Efforts Produce the Tastiest Results	25
Figure 27: Essential CDO Actions to Create DAML Success Across an Enterprise	26
Figure 28: Choose Your Tigers Carefully or One Might Dominate	28
Figure 29: Make Strange Uniforms and the DAML Salute an Academy Requirement.....	28
Figure 30: Best to Outsource for this Communications Approach.....	29
Figure 31: DAML! What the Heck is This Guy Doing?.....	30
Figure 32: Critically Observing People Managers “Change the Face” of Brown Bag Sessions.....	30
Figure 33: Frequent, Repetitive, and Error-Prone Processes are Where Initial Automation Excels.	31
Figure 34: Try Hard to Actively Network with Meeting Attendees.....	32
Figure 35: STOP!!! Cross Training NOT Crossbow Training!.....	32
Figure 36: SCTE's Unique Focus on Cable Workforce Development Will Be Key.....	33
Figure 37: Yes, shameless self-promotion here.	33
Figure 38: Fusion Teams Must Achieve 100 million Kelvin to Become a Self-Sustaining Reaction	34
Figure 39: Average value DAML brings to the company based on the organizational model.....	35
Figure 40: Gift cards blur tax reporting requirements;).....	36
Figure 41: YES, It's Real! Think We Just Made This Ourselves 🍌	36

1. Introduction

DAML is intended to include a list of terms that are ever-expanding in *both their breadth and depth*—far too many to list here. This is like looking at a wide-field telescope and a powerful microscope *at the same time*. Very few people can keep up with the changes or make sense of the interpretations between these views.



What this paper is *NOT* is an appeal for everyone to become:




Figure 2: intensely deep programmers on dark screens in dimly lit cubicles,




Figure 3: neuroscientists to optimize all learning, behavior, and processes,



Figure 4: or automation executives that desire very few humans doing all complex work with their single push-button fingers;)

Highly skilled experts and visionaries are essential to create *meaningfully unique* products and services. However, by having a much higher level of awareness and collaboration from the rest of the organization, communicating to the expert's "wavelength" is greatly improved. This enables achieving better results on the most important needs faster, more efficiently and with less resource cost.

1.1. A tale of two CEOs.

The author learned a great deal about the importance of an enterprise wide DAML initiative through two vastly different consulting engagements.

The first was entirely CEO-led at a one-hundred-year-old financial company. It began during a golf game where his competitor suggested he use a particular robotic process automation (RPA) tool and completely plan and implement it with an outside consulting company. Around the same time this CEO's Chairman requested that he implement a specific intelligent document processing (IDP) system used by two sister companies to eliminate manual efforts and errors. Five RPA processes were targeted, but only one was delivered and put into production; one was eliminated far too late in the process because of the high complexity; three others were put on hold when the consultant asked for additional compensation for the second time. The IDP software was eliminated after months of planning and implementation because the cost benefits for the applications considered were extremely poor compared to the sister company's high transaction volume. Sadly, there was another huge effort that targeted AI to transform the call center into

a digital contact center as well. The only people that knew this information in detail was a knowledgeable call center manager who was laid off, a fantastic technical executive that battled valiantly to reduce the weekly management changes, and a two technical contract project managers that were never made full-time employees.

The second CEO was at a software as a service (SaaS) company beginning to utilize DAML in a profound way to improve performance over time once the SW administrators approved the automations. This CEO wanted to completely change the mindset of all his employees by teaching them DAML; the author helped create a strategic organization plan and a potential product roadmap for new applications to be considered in this effort. The results were amazing. Not only did the employees learn these skills but began to suggest new product ideas and help bring them to reality while also looking for ways to use DAML methods in their daily work. This company has continued to evolve, filing several new patents, and being awarded at least two more. All of this helped attract attention from some of the largest multi-national companies in the world.

Much of the HOW information contained below was used from this approach.

1.2. Paper approach.

This paper will first discuss **WHY** DAML is so critical; next **WHAT** to *initially* teach company employees as the DAML bedrock; then finally **HOW** to deploy this strategy using simultaneous top-down and bottom-up approaches.

2. WHY DAML must become your organizational thinking and “vibe”.

In the 2016 motion picture *Hidden Figures*ⁱⁱ, Dorothy Vaughan is a mathematician who worked as a ‘human computer’ for NASA during the 1960s. When NASA installed their first IBM 7090 mainframe computer, Dorothy feared she and her team would soon become redundant, then she correctly predicted NASA’s urgent need for a programming team. Instead of worrying or doing nothing, she taught herself about the computer and the FORTRAN programming language, and then taught it to her team who became NASA’s very first programming team. Vaughan embraced their undeniable future with enthusiasm, helping not only herself but her colleagues, their families, NASA, and eventually society for their contribution.



Figure 5: Hidden Figures' Dorothy Vaughn Learns New Computer Technologies Before It Replaces Her Team

“No doubt about it, data is now a prized business asset, which means organizations of all kinds will want to employ data-literate individuals. In an average business context, data literacy means being able to access appropriate data; work with data confidently (creating/gathering data, keeping it up to date, etc.); extract meaning from data; and communicate those data-based insights to others. It is also important to question the integrity and validity of any data you are working with, rather than blindly following data”ⁱⁱⁱ.

2.1. How the world's leading companies are cultivating a new business culture—Case Study JLL.

JLL is one of the largest owners and operators of commercial real estate worldwide with an astounding 4.6B square feet in property and facilities management. Over the past three years, Paul Chapman, JLL’s global director of business intelligence and technology, has worked to develop a data culture across every aspect of the company. “Our facility managers can see the data for themselves, from showing them how old air conditioning units are and when they should be replaced, to how much each facility is costing per square foot.” Employees at every level of the company have access to this data through a dashboard, helping them to determine the root cause when problems arise and to figure out how to respond.^{iv}

2.2. Motivations for making DAML an enterprise wide initiative.

2.2.1. DAML adoption is already exceedingly high.




Figure 6: DAML Production Deployments Increased from 35% in 2019 to 47% in 2021^v.

Above figure shows high in-use DAML projects, and they are increasing over 15% CAGR.

2.2.2. New DAML knowledge areas are exponentially increasing.




Figure 7: Gartner Hype Cycle for Data Science and ML July 2021

Gartner publishes annually expectations for technologies in specific areas called Hype Cycles where the white and light blue circles depict those knowledge areas advancing the fastest. Organizations implementing these technologies typically must wait for these technologies to advance to at least the 4th

“Slope of Enlightenment” segment for outcomes to change business processes. Beginning a comprehensive DAML program across the entire organization better prepares everyone for the continual change of skillsets, mindsets, and toolsets that these innovative technologies require. One could easily claim these knowledge areas are not simply increasing but *exponentially* increasing.

The time to embrace DAML is *right now*.

2.2.3. Establishing a “data-driven culture” is the biggest contributor to DAML success.




Figure 8: Areas Critical to DAML Success^{vi}

2.2.4. DAML must become EVERYONE’S strategy.

As seen in the figure below, the impacts of enterprise wide DAML initiatives are far reaching.




Figure 9: When an Organization Embraces DAML, Various Team Strategy Components Must Be Implemented^{vii}

DAML skills are also seen as critical in every department of an organization:




Figure 10: Importance of DAML Skills by Department^{viii}

2.3. DAML value for *me* as an employee, and how will it make my job easier?

The initial outcomes of DAML processes starts by *augmenting* human capabilities, reducing mundane, repetitive, and even dangerous tasks to allow a greater concentration on the unique skills that define us as individuals.

These distinctively human talents are the most difficult to automate and therefore will become the most indispensable:

- Effective communication and collaboration through emotional intelligence,
- Empathy, compassion, and authenticity;,,
- Curiosity plus instigation,
- Strategic analysis and analytical thinking,
- Complex, multi-disciplinary judgement, and decision making,
- Adaptability and extensibility,
- Creative thinking leading to real innovation,
- Conflict resolution,
- Cultural intelligence and diversity consciousness,
- Ethical awareness,
- Negotiation and persuasion, and
- Leadership through genuine trust, transparency, inclusivity, and respect.

Can you envision an exemplary, engaged fellow employee NOT wanting to shed monotonous tasks to focus on these critical thinking and empathy skills?




Figure 11: Gartner's Top Data and Analytics Trends for 2022 (Partial)^{ix}

The above figure describes the need to help employees where they work currently and to aid in their efficiency and quantity/quality of work output. This means to augment people with a broad array of contextual data, to build up components optimized on how teams really work, to understand how decisions are made (aka “decision engineering”) to optimize people and process impact, as well as to build a higher level of DAML competency.

Curiosity and desire for continuous learning, growth, and improvement is cited at the one skill everyone must cultivate. Whatever your age, whatever your industry, if you can spark your desire to learn, you will be giving yourself the best chance of a successful, fulfilling life. Especially in a work context, curiosity and continual learning are fundamental to being able (and willing) to embrace change. It ensures your skills stay sharp, that you can keep up with the major transformations taking place, and that you stay relevant. DAML is a major transformation.

2.3.1. YOUR current job may cease to exist.




Figure 12: Funny Because it is True^x

Automation will have a far-reaching impact on the global workforce.

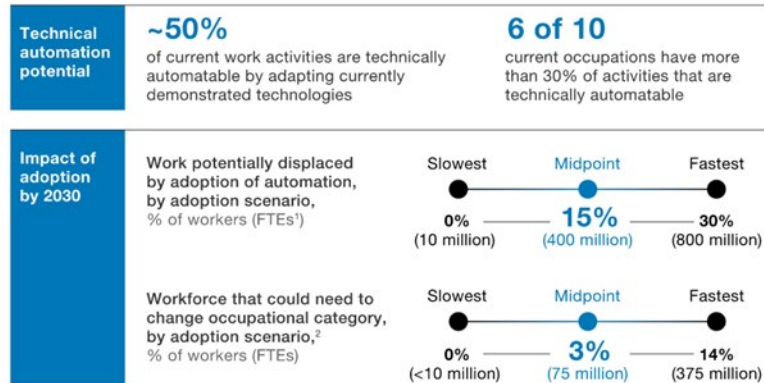


Figure 13: Employment Changes from Automation 2017-2030^{xi}

This same McKinsey study continues “...between 400 million and 800 million individuals could be displaced by automation and need to find new jobs by 2030 around the world...Of the total displaced, 75 million to 375 million may need to ...learn new skills...”. Some respected sources predict over *half of all existing jobs will become obsolete through DAML*—but this can become a source of strength for employees who absorb, employ, and evangelize DAML skills. Every employee must become shrewder in what will disrupt, augment, and improve *every* existing job description, work process, and organizational design.

However, the World Economic Forum predicts that “by 2025, 97 million people will be needed for AI-related jobs, and that doesn't count the numerous other jobs that AI will help produce”. As McKinsey noted, “early, innovation-focused adopters (of AI) are positioning themselves for growth, which tends to stimulate employment.” Moreover, most companies understand that AI's real benefit is not automating people out of jobs but *enhancing their skills, reducing errors, and freeing up employees to do more important tasks*”(italics added)^{xii}.

One reason to not despair: DAML will “*result in the creation of many more (new and unforeseen) jobs.*”^{xiii} Similarly the Institute for the Future predicts that 85% of jobs that will be available in 2030 *have not been invented yet*^{xiv}.

2.4. How can my organization benefit?

The overall benefits of driving an enterprise wide DAML movement are already being embraced by leading corporations across the globe leading to enormous success. The most critical gain from an organizational DAML strategy is to *generate value from information assets* by first improving internal processes, then secondly to increase the value of products and services; both concentrate on efficiency and innovation.

However, the third organizational gain is the wildcard that greatly enhances cultural alignment and cohesion is recruiting and retention. More on this later.

It would be ironic to ignore striking industry statistics that are metricized, collected, and envisioned using DAML processes that themselves indicate the importance of DAML processes:




Figure 14: Impacts of Employee Training Programs

References for Figure 14:

- “Companies that offer”^{xv}
- “A 10% increase”^{xvi}
- “40% of employees”^{xvii}
- “Recent national survey”^{xviii}
- “In a recent”^{xix}
- “According to a study”^{xx}

As stated, the effect of a poor training environment has a profound impact in employee turnover and this impacts recruiting. Employees frequently rank the ability to learn new skills and try different approaches as the key factor in their decision where to work. Intelligent people are also well-connected and talk to others when networking about companies and how employees are treated. The mantra is you tell one person about positive experiences and ten people about the negatives. Therefore, insufficient training impacts recruiting as well as investments in recruiting, hiring, and onboarding a new employee.

Organizational benefits of implementing a comprehensive, DAML-driven, enterprise-wide initiative:

- Connect every aspect of the organization to customers, mission, and vision statements using commonly understood definitions, well-placed metrics, continuous monitoring, and visualization, leading to a reformulation of operating models,
- Use as a valuable tool in cultural change and alignment,
- Create the concept of monetized data; *information* asset values related to actual loaded cost or time value. Also include in values intended/future use, connections to entities/processes/individuals/groups/physical assets, error/believability, security, confidentiality, authenticity, external sales/licensing, update frequency, and expected end of life,

- Improve the prioritization of critical business processes as well as business process efficiencies through elimination of duplicate processes, alignment, optimization, automation, plus continued maintenance, and improvements,
- Combine and accelerate innovation processes to generate new and more meaningful data-driven intellectual property (IP), products, and/or services,
- Focus employees less on tactical execution of repetitive tasks and more on improving the tactical processes. Participating, possibly for the first time, in customer-centric and strategic enhancements,
- Empowers all levels of workers to ask the right questions of data, processes, and machines, build knowledge, make more informed decisions, and better communicate meaning and context to others.

“IDC found that only 36% of their surveyed enterprises succeeded in putting completed (DAML) models into production. Nearly half of the remainder (32%) had not even made it past proof of concept’. By raising the DAML education and experience level of the entire staff, more thoughtful processes will be chosen by the employees that know them best, improved business process flow information will be collected, and more informed visualization and testing will occur leading to higher quantity and especially quality of the resulting information.

2.4.1. How does DAML reduce and eliminate errors?

DAML enables an iterative, experimental approach by personnel closest to customers, opportunities, and obstacles, improving decision speed and *especially quality*, reducing resource waste, and streamlining communications so that all employees can understand and act.

2.4.2. Greater customer insight and retention.

“(DAML) driven companies are twenty-three times more likely to acquire customers than their peers”^{xxi}.

The digital world is built on data and information. Information on user behaviors can be collected and analyzed using DAML methods to understand how people are using and even socially commenting on products and services, enabling customized experiences that better meet their needs, styles, and desires. Using extracted information in context is commonly referred to as hyperpersonalization, and further drives customers to purchase, engage and even evangelize because of experiences tailored only for them.

2.4.3. Retention and recruiting

HR frequently is charged with recruiting and retaining the absolute best, however vague, non-metricized statements are always misleading. It is clear recruiting and retaining competent talent with critical thinking skills—and most especially the most coveted DAML skills—leads to a reduced turnover ratio and therefore a higher profit margin than the competition.

The figure below highlights the key skills required by open roles in several industry segments. Communication, analysis, and collaboration are what every employer needs the most, and these are some of the very skills accentuated by an enterprise wide DAML strategy.




Figure 15: Data Source: Gartner’s TalentNeuron, and represents jobs posted between 4/22/2021 4/22/2022

Most companies pursuing DAML will focus entirely on recruiting a team of highly trained technical experts. However, with the exponential expansion of DAML organizational needs comes a serious shortage of potential employees at this level, forcing many companies to live with extremely high consultant salaries, noticeably shorter tenures, and often very remote or offshore resourcing making direct engagement with the business leaders impossible.

By training *everyone* on DAML, the strain of finding increasingly rare and transient experts to completely lead DAML efforts is reduced—but never eliminated. All employees that become aware of the vocabulary, methods, and potential outcomes of DAML can greatly assist a much smaller and integrated pool of specialists. Employees using their existing business knowledge with their new skills will create a cooperation between them and an atmosphere that exudes excitement both inside and outside the company. When a company understands their value, mission, and customers and becomes emboldened to learn and employ the latest DAML skills to create visible outcomes, recruiting becomes *almost* effortless.

DAML processes are rarely siloed into one small area, but rather benefit from the combination of neighboring departments, subjects, and data sets. By mapping “skill adjacencies,” employees better understand how their skills and job definition connects with their closest working colleges and throughout the company so that their insights can be extended to other areas by both awareness and cross-training.

Use of DAML skills by human resources (HR) *internally* also has great benefits. “(DAML) can influence and improve how human resources departments manage recruitment. Tapping into data can help companies to headhunt the most promising talent by comparing performance metrics. From here, businesses can retain these valuable staff members through the culture of clarity that data-based working environments foster. Essentially, data-driven human resource strategy has significant short- and long-term benefits, creating a productive, motivated staff team who will increase profit margins.”^{xxii}




Figure 16: Every Job Function has Adjacent Skills and Roles. Example Shown for Cloud Engineering Skill Adjacencies

2.4.4. Increased performance through DAML organizational obsession.

DAML drives workforce productivity by increasing efficiencies—whether individually, in work teams, for departments and divisions, or the entire organization. Data selected by the people that best understand the processes and potential outcomes enables more tangible goals by creating a clearer picture of the overall targets, behaviors, and patterns, allowing an agile approach to enhance productivity. It also allows identification of areas for improvement and promotes transparency and accountability.

In a study of over four hundred businesses, it was confirmed that organizations with the most developed analytics capabilities commanded a larger market share, were twice as likely to be in their sector's top 25% for profitability and were five times more likely to make swifter decisions than competitors^{xxiii}. The point being organizations that are not exploiting DAML are certain to lag their competition. Another study confirmed these findings. By analyzing data sets from Fortune 1000 corporations, the impact on profit resulting from DAML was measured. Some notable findings included:

- Companies could increase profit by more than \$2 billion a year by making just 10% of available data usable.
- Return on equity increased by 16% by making data more accessible.
- When advanced reporting was deployed, return on investment increased by 0.7% – which is equal to \$2.87 million in additional revenues.
- Most importantly, a comparably low investment in data analytics was required to produce these significant gains.
- DAML driven companies are 58% more likely to beat revenue goals than those who are not focused on data,

2.4.5. Collaboration and working

in teams.

This is one of those skills that seems obvious for workplace success, so why include it? “Because the nature of collaboration and teamwork will change as teams evolve to include automated processes as well as hybrid workers, fully remote workers, contractors, and other employees from different business units who “float” between projects and teams. With such distributed teams, we need collaborative skills more than ever.”^{xxiv}

Creating a vibrant DAML program will create or enhance the organization’s foundation of cross-disciplinary collaboration and problem solving. More will be mentioned on this later in the HOW section, but employees enjoy brainstorming, goal setting and progressing their own defined initiatives in teams. By composing them of different perspectives and levels instantly adds a diversity element found to be crucial to innovation success, so it is also important for DAML success.

2.5. How can the cable industry benefit?

It is time *the entire cable industry swiftly embraced the implementation of a DAML strategy* to propel it into the future using data-driven initiatives that have been proven by hundreds of leading organizations. DAML is key to not only thrive but exceed the competition and race for the best strategic and forward-looking talent. “Not to decide *is* to decide.”

During a recent virtual conference, SCTE CEO Mark Dzuban made several points in a session labelled “Tackling the 10G Challenge”. 10G is a CableLabs’ series of initiatives to prepare cable operators for symmetric 10Gbit/sec access networks and applications. He described an urgent need to “work with CableLabs to elevate the knowledge base to “bring 10G to life”. He continued to describe “Project Wisdom,” a CableLabs’ training initiative for advanced skills, capabilities, and knowledge through community colleges and universities. “Its goal is to create “unique talents to build 10G” and “AI-driven cable operators” leading to “increased (customer) touchpoints, proactive operations at outages, zero touch for customers, and to connect through all ways (we) communicate”.

Mark closed by saying he would like future job seekers to “think of cable as a career path, especially for automation and AI/ML”^{xxv}. So “(c)able leaders must not only focus on attracting new employees to the industry, but it is imperative to upskill the existing workforce to support the global expansion of connectivity and technological innovation”^{xxvi}.

2.6. Ignoring DAML will doom existing companies.

Failure to embrace this fundamental change will doom an organization’s competitive ability, potentially taking them out of business eventually. “By 2030, AI will lead to an estimated \$15.7 trillion, or 26% increase, in global GDP”^{xxvii}. This increase is greater than both China and India’s current GDP *combined*! Increased productivity accelerated by DAML will contribute to approximately 40% of this increase in GDP growth.

Educating employees to better understand DAML vocabulary, methods, processes, roles, and metrics for both internal and external projects will increase communication, clarity of thought and actions bringing all departments together to speak a universal language that benefits the company.



Figure 17: What You Don't See Can Kill You (and Your Organization)

3. WHAT DAML areas to initially address.

This section gives three initial areas to consider with the goal to model your approach by balancing strategic and tactical priorities of your organization. These are only a few suggestions among thousands, and the tools your organization selects from this DAML toolchest are entirely determined by process selection and prioritization efforts.




Figure 18: Operating Model of a DAML Organization^{xxviii}

3.1. A quick aside on “determinant” vs. “assisted” DAML processes.

DAML is often thought of as determinant, or unsupervised, making complex decisions without any human intervention. Look only to Space Odyssey’s Hal 9000 to see this is naïve; the most critical area to begin with is assisted or *supervised* processes.



Figure 19: Always start with assisted DAML processes

People must still be accountable for the business decisions made, and it is not possible to build everything a human business expert can do well. So, recommendations are output for the human to respond, providing a “feedback mechanism” so that the system can learn from the human response to eventually become *almost* unsupervised; but you can never leave “it” completely alone. The world and organizations are always changing, so system results must always be monitored with key metrics. Documentation must also totally describe the “what and how” of the manual process to enable new employees to take over when manual interaction is required, and automated process for future iterations.

3.2. Focus on data literacy.

“Culture and data literacy are the top two roadblocks for (creating a DAML fueled organization) ... (b)y 2023, data literacy will become an explicit and necessary driver of business value, demonstrated by its formal inclusion in over 80% of data and analytics strategies and change management programs.”^{xxix}

Gartner defines data literacy as the “ability to read, write and communicate data *in context*, including an understanding of data sources and constructs, analytical methods and techniques applied, and the ability to describe the use case, application and resulting value... By 2023, data literacy will become essential in driving business value, demonstrated by its formal inclusion in over 80% of data and analytics strategies and change management programs”.^{xxx}

“Data literacy isn’t a nice-to-have — it must become part of your firm’s DNA.” While most organizations strive to continually enhance decision making for critical tasks, “(t)he problem is not data availability or tools but lack of skills to use the tools and the data to drive business outcomes. Technology executives are creating data literacy programs to bridge this gap. These programs must be comprehensive, companywide efforts with executive support, strategic goals, and established metrics”^{xxxi}.

Tableau rated data skills/data literacy the number one need for Entry Level Skills in three categories: demand, increase in demand, and anticipated importance.




Figure 20: Despite Company Efforts Most Do Not Have Adequate Data Skills^{xxxii}

Goals should include basic data awareness, deeper understanding of insights, how to connect insights to decision-making, continued training of experts, and scaling skills across the organization.

- › Explore the need for data literacy — *Awareness, Comprehension, Expertise, and Scale* — across the entire organization.
- › Understand basic data competencies such as recognizing data and its value to the organization, interpreting data visualizations, data storytelling, data hunting, and insights-driven decision making.
- › Develop a comprehensive data literacy program to increase awareness of data and its value, improve understanding and data decision making, enrich expertise, and scale the program across the entire organization.
- › Identify opportunities for partnerships to expand and enhance data literacy from service providers and data literacy experts.



Figure 21: Forrester Approach to Building a Data Literacy Program^{xxxiii}

Data literacy today suggests the concept of an extended understanding of “Big Data”: literacy that places awareness and critical reflection of big data systems at its center^{xxxiv}. Big data uses extremely large data sets collected from all aspects of a business so that they can be analyzed computationally to reveal patterns, trends, and associations, especially relating to human behavior and interactions.

3.3. “There is No AI without IA (Information Architecture).”^{xxxv}




Figure 22: Most DAML processes fail because of data preparation and access^{xxxvi}

Data collection and preparation is the most time consuming and difficult part of creating meaningful outcomes from DAML. Everyone must begin DAML training on a few key subjects.

3.3.1. *Create data dictionaries that span the organization and align with industry.*

Enterprise culture has been described as “the collective conversations of an enterprise. So, it is impossible to change the behaviors and belief systems of an enterprise without first altering the conversations within it.

“Change your words, change your life.” “Words are containers of power.” Many motivational speakers have used these and similar expressions for good reasons. The way we communicate affects everything. Data literacy harnesses the power of organizational words to provide guidelines for making sure our communications are efficient, constructive, and used to achieve superior results. When companies do not begin with a common and maintained data dictionary (aka glossary), DAML efforts will not achieve optimum outcomes.



Figure 23: “Garbage In, Garbage Out” (GIGO)^{xxxviixxxviii}

3.3.2. *Seeking Common Acronyms*

Just as an organization must have a common vocabulary with identical definitions, they also must agree on acronyms. Common Acronyms used across an entire organization or even industry is not typically explicitly defined. Because they are not common to everyone, they can confuse and alienate unfamiliar audiences. Even well-intentioned writers and speakers may overestimate an audience's familiarity with abbreviations. Abbreviations should not be completely avoided but using them as a default assuming all agree with meaning and context can be problematic.

3.3.3. *Taxono Mies? Taxi please.*

A taxonomy is the discipline of classifying information and are typically created in a hierarchical, parent-child structure. They are extremely well organized by all users of this data and are typically published and maintained by an organization.

The “father of taxonomy” Carl Linnaeus was a Swedish botanist that named, ranked, and classified plants and animals in 1761: Kingdom, Phylum, Class, Order, Family, Genus, Species. Do you remember the helpful mnemonic like “King Philip Came Over For Good Spaghetti”?



Figure 24: Insert Some Humor When Content Appears Boring or Irrelevant

Another quite common taxonomy example that most people are familiar with is the Dewey Decimal System used for information organization, especially in library collections.

001-099	Generalities	PREQUEL TO THE DEWEY DECIMAL SYSTEM WHO AM I?	Encyclopedias, curiosities and wonders, unexplained mysteries
100-199	Philosophy		Books about the self, feelings, dreams, witchcraft
200-299	Religion	WHO MADE ME?	Christianity, Judaism, Buddhism, Hinduism, etc., and Mythology
300-399	Social Science	WHO'S THE GUY IN THE NEXT CAVE?	Customs, cultures, laws, manners, costumes, fairy tales
400-499	Languages	HOW DO I TALK TO THAT GUY?	Dictionaries, parts of speech, sign language, foreign language aids
500-599	Natural Science	LET'S TALK ABOUT THE WORLD WE SEE	Mathematics, earth, astronomy, chemistry, plants, wild animals
600-699	Applied Science	NOW LET'S MAKE STUFF OUT OF WHAT WE SEE	Inventions, robots, transportation, pets, recipe books
700-799	Arts and Recreation	NOW LET'S HAVE SOME FUN	Art, drawing, comics, handicrafts, music, games, sports
800-899	Literature	LET'S TELL OUR CHILDREN HOW WONDERFUL WE ARE	Poetry, plays, classic literature, jokes, riddles
900-999	Geography and History	LET'S TELL OUR FUTURE CHILDREN HOW WONDERFUL WE WERE	Landforms, travel, atlases, exploration, countries
92 and 920	Biography and Collective Biography	FIND OUT ABOUT FAMOUS PEOPLE	Single person: filed by last name of subject Multiple people: by author

Figure 25: Dewey Decimal Library Classification System

Having an organizational taxonomy is critical to begin any DAML project. DAML success only results when the entire organization is speaking the same language—using common definitions and a structured way to hierarchically organize all data and information. Think about it: most organizations continually create their own definitions and information organization, and quite frequently these conflict between groups, departments, divisions, functional areas, and acquired organizations. A company's taxonomy must come together to create a common language in the company and throughout the industry so all can more clearly communicate with each other. This heightened communication is the fruit of a well-versed staff in DAML, increasing productivity and thus contentment among employees.

3.3.4. Aim for INFORMATION CONTEXT.

Organizations must have a thorough and complete understanding and use of common data assets and derived information to maximize their DAML initiative.

When DAML processes connect inventory, data, resources, and people through common data and information communications, not only does the company benefit but the industry benefits as well. If DAML standards are not “set in stone” at the very beginning with pan-organizational definitions of every term, synonym, alternate term, and acronym list, the information *context* will be quickly lost.

It is impossible to create compelling, meaningful DAML outcomes from a misunderstood and fluid data and information foundation.

4. HOW DAML can be understood, embraced, and utilized by all employees.

DAML can be understood through common vocabulary, embraced through clear and simple methods, and utilized by employees who understand the benefits of DAML to themselves, their career, their company, and their industry. The positive DAML outcomes are clear: collaboration, retention, recruiting, mission alignment, external perception, revenues, and the excitement of being part of the future knowing and understand the tools DAML will provide companies in the future.

This section provides an overview of several suggested tactics to get every single person in your organization to understand, adopt, and engage in DAML. By selecting the DAML process priorities based on outcome values, some but not all these methods may make sense.

4.1. Start both “Fast and Furious” and “Slow and Steady”?

Initially the focus is on common DAML-specific vocabulary, basic concepts, methods, and both alignment to the organizational goals and the ability to begin to derive new goals for employees to teach themselves and their teammates. This approach is “Slow and Steady,” involving everyone from the janitor to the board chairperson. This means a tailored agile learning approach for each employee’s personal job responsibility, skill levels, learning style, and pace.

And now for the simultaneous “Fast and Furious” win: creating a one-time, company tiger team. A tiger team is a specialized, cross-functional team brought together to solve or investigate a specific problem or critical issue. People need an example of how to act and a good early success story.

A visible way to display DAML skills is to assemble a “tiger team” to help choose one easy-to-understand, visible-to-all, mundane, repetitive task that if automated would clearly save time and money while reducing errors and risk. This tiger team concept will be explained in more detail later.

4.1.1. Create a grassroot-led effort.



Figure 26: Grass Fed Efforts Produce the Tastiest Results

To succeed in instilling a DAML culture, every senior executive must participate by planning the resource capacity reductions on existing projects so that the entire company can learn DAML at the same time. They must work to identify risks to current projects and determine workarounds, hiring consultants, outsourcing, delaying, combining, and possibly reducing overall number and scope of existing projects., (etc.)

The training itself can take many forms. Bigger, more formal training courses can be used to layout universal guidelines—ensuring everyone is aware of which data is private and which is public, and how to handle it. Small, interactive group training can be an effective way to teach workers how to gain a better understanding of the data they work with every day.

4.1.2. Dub DAML development director.

Adding a new executive function for Chief Data Officer (CDO) really sets an enterprise DAML strategy on fire. The CDO generally oversees a range of data-related functions that may include data management, ensuring data quality, and creating data strategy. They may also be responsible for data analytics and business intelligence — the process of drawing valuable insights from data.

Whatever title is chosen, the CEO/President must personally introduce this person and then work with that leader *directly* on all DAML strategy areas to demonstrate the importance of this initiative.. By beginning from executive request and hands-on leadership, the initiative launches from the crucial *financial value of DAML*, then less so on data quality metrics (accuracy, completeness, scale, and usage), and then even far less measures of the key information and data assets impact of on business processes, such as KPIs.

The CDO should *not* report to the CIO but instead the top business executive (CEO, COO, CFO); it is the only way to demonstrate the value of DAML in the context of the organization's business. CIOs have far too much on their day-to-day workloads to either prognosticate/create the overall organizational future (the CTO's role) or to establish a vibrant DAML center of excellence effort. The marketplace mirrors this: in Gartner's 2018 CDO survey, CDO reports to top business executives increased from 45 to 48% YOY, while those reporting to the CIO declined from 23 to 22% for this same period.

Chief Data Officers (CDOs) or the highest lever executive charged with teaching and utilizing the entire organization should intelligently coordinate the creation of competencies across the enterprise, ensure the efficacy and consistency of distributed practices through metrics and scorecards, and build enterprise capabilities that underlie future success.

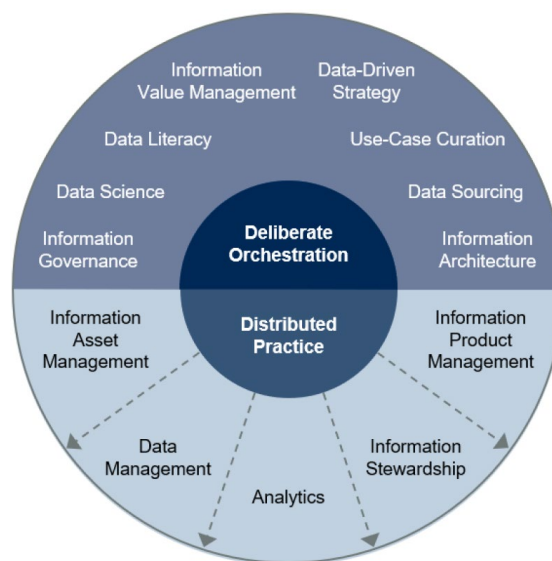


Figure 27: Essential CDO Actions to Create DAML Success Across an Enterprise^{xxxix}

Divide all DAML technical staff into groups that either facilitate enterprise-wide or are assigned to one to three groups of cross-functional staff. These group technical leaders should propose potential collaboration opportunities between members of different teams, and then assist as needed. They also should discuss external organizations that could either lead or supplement the training as well as provide expertise in the DAML developments, but they must be *incredibly careful* about consultant or chosen toolset lock-ins.

4.1. Methods to infuse data driven competencies.

Building a data-driven enterprise is not just about encouraging the use of data in decision making. Data and analytics leaders must lead development of the appropriate competencies and align work to be consistent with their enterprise's ambitions for generating information value.^{xl}

4.1.1. Publicize DAML program and add participation into job evaluation criteria.

Senior management needs to begin continuous messaging to employees about the urgency for ALL employees to know DAML, and why it will benefit them individually in their current work and career plans, as well as their work teams, departments, the entire organization, and even the industry.

Employees should be encouraged to share their ideas. This messaging effort should support employee conversations that are collaborative, assisting both internal and external DAML opportunities with their inside knowledge, especially in their own jobs and on every company level. Management also must connect DAML to what the organization's business is all about, and new direction the organization is heading through their mission and vision statements; these may need to be revisited to accentuate this push. Wherever possible they should use specific examples of planned and currently released product/service features to frame all DAML explanations of new benefits.

To cement this into the company culture, all job descriptions should be updated to reflect expectations for DAML training, competency levels, participation in ideas and methods for products/services, and how these new skills will be used by them personally. This should begin by management first evaluating their own roles to continuously seek the highest ROI applications and then to prominent and visible in designing, testing, and continuously maintaining their applications.

4.1.2. "Tiger Team" taming.

The term "tiger team" originates from the military but was made famous by NASA who deployed a tiger team during the endangered 1970 Apollo 13 mission^{xli}.

Here we begin with top-down approach using a *one-time* tiger team composed of top cross-functional staff to help employees understand by visible example how eliminating a routine task using DAML will allow a focus on higher, more strategic functions of their jobs. To create a tiger team approach, choose an important and well-known business process with established metrics, and then automate this process through a DAML demonstration project that greatly improves the metrics. This shows all employees what they are expected to achieve.

This suggested approach should directly impact one or more people who drudgingly slug through this repeatable task. With this in mind, at least one of these employees needs to help clearly communicate the mundane nature of these job tasks and how they personally will change their approach to their role with the extra emancipated time. They become a DAML "poster child," altering their future by learning new skills that will keep them important in their company role into the future.

Carefully choose staff that take part in this tiger team employees known to work hard, are cross trained to think through various roles/perspectives, can communicate to diverse groups very well, and are not perceived to be favored by executives. This is a one-time shot and must derive successful results and excitement.



Figure 28: Choose Your Tigers Carefully or One Might Dominate

4.1.3. “DAML Academy”?

Creating a DAML Academy (also known as a Center of Excellence) can become the heartbeat of the DMAL strategy. By establishing this as a central organizational theme, all aspects of the company begin to “spoke” around it forcing all processes and procedures to be seen considering DAML opportunities. This DAML Academy should address naming conventions, standardized teaching methods/materials, describe development phases, enable internal DAML strategy refinement, and publicly “show case” the company’s new strategy and positive outcomes.



Figure 29: Make Strange Uniforms and the DAML Salute an Academy Requirement

4.1.4. Form “DAML Translators” to bridge technology and business realms.

“DAML Translators” (or champions) ensure business needs are met and adoption of DAML strategies is smooth. Translators are there to expand(s) the typically siloed IT technical role to become a problem-solving collaborator with an enterprise perspective concerning data and analytics. They enable the formation of communities of practice and DAML competency development.

The best fit will be employees who know every aspect of the department and are known by all key staff members. These employees will excel in their understanding of DAML and will be capable of leading or moderating departmental DAML initiatives. Where qualified employees are lacking, find certifications that can transform them then add that to their job evaluation criteria. Executives may also want to

partner with department leaders to provide clear insight from on the ground as to how DAML can assist and benefit each department and employee directly.

These also could be DAML technical professionals that have excellent empathy, emotional intelligence, and communications skills with a good knowledge of the department and possibly paired with a departmental expert that is lacking DAML skills but will benefit from the paired interaction.

4.1.5. Develop communications and training to inform organization and external stakeholders.

The goal of company communication concerning DAML is to explain the extreme urgency of getting every employee not just *thinking about* but intensely *engaged* in the company DAML strategy.



Figure 30: Best to Outsource for this Communications Approach

Choose one key people management leader in each team so they help communicate training importance, expectations, responsibilities etc. That person will be responsible for creating and participating in a team that will design processes, contacts, roles, responsibilities, project updates, and performance metrics.

Company employees need to hear success stories of top companies and especially their competition to better comprehend the importance for their organization. Employees who envision their current job in an entirely unique way will see that applying data skills in their daily work will eliminate boring and routine actions, giving way to higher levels of engagement and greater company contribution that will insure their own personal growth as well as their future success in the company.

4.1.6. A little DAB'll do ya!



Figure 31: DAML! What the Heck is This Guy Doing?

An external DAML Advisory Board (DAB) can provide profound insight that can save extraordinary savings of time and company expense. Organizations commonly benefit from advisory boards with a combination of business, innovation, and technology capabilities. Here, external professionals are assembled that have DAML experience with the specific business processes of the organization. Meet with DAML technical staff and senior execs to determine initial goals, the best approach to achieve these goals and when, etc. Areas such as reduce operational cost 20% in 3 years can then be exploded into the many sub areas and what each one would take to contribute to this entire org goal.

4.2. Brown bag training.

Begin ongoing sessions for the entire staff led by DAML *teaching* pros with/without professional meeting facilitation or hire an organization to lead this training. Organize series of learning events for each organization and/or team, make it fun and interactive, gauge understanding. Focus on the basics of DAML with definitions and use examples. Describe how staff once free from mundane tasks can contribute at a higher level. Goal is to build in-house capabilities without hiring (as many) business savvy external DAML resources.



Figure 32: Critically Observing People Managers “Change the Face” of Brown Bag Sessions

The author helped plan and facilitate a company-wide brown bag learning and development program at a data center product company. All employees were assigned tables to distribute functional responsibilities, and managers were intentionally kept out. Over the course of several months the knowledge on innovation processes (remarkably like DAML processes) increased allowing the teams to conceive product ideas related to the company's areas of highest competency. These ideas were refined into business proposals, voted on, prioritized, and presented to senior management. Several concepts were then presented to a newly formed corporate advisory board and with their input at least 3 new products emerged that captured customer and industry analyst attention. These combined with the current market leadership led to a 9-figure acquisition.

4.2.1. Find mundane and repetitive tasks and prioritize.

Break into groups of 6-8 on round tables. Larger companies may want to have color-coded badges for each department and/or division. Those that are people managers separate and meet in different rooms from the staff and use data available and stated goals to determine the highest operational cost components. In each case have only one "color" at each table. Goal for each table is to determine 3-5 job activities that are important and could be automated, and then prioritize all or at least the top two.



Figure 33: Frequent, Repetitive, and Error-Prone Processes are Where Initial Automation Excels.^{xlii}

Gather all ideas from each table and discuss as an entire group, then senior management take the top 5-7 ideas and add data on operational costs, direct and indirect benefits, etc. to prioritize by value.

4.2.2. Create a competition.

Can either use hand-selected cross-functional teams to fill in new tables or the same table groupings. Assign each table one process and have them create flow charts to determine normal and variations with all connected processes, systems, and people. Use available information to gauge ROI for choosing and automating this solution and allow them to determine the best business processes where automation can give an immediate benefit.

4.2.3. Present chosen automation processes to DAB.

Opportunity for select team member(s) from each competition table to communicate select information to external stakeholders to increase their understanding of how the organization is evolving. Also, this method forces the speakers to hone their listening and speaking skills to an audience level they may not have addressed before. DAB selects top three processes then senior execs give associated table participants public recognition and some reasonable monetary award (Amazon gift card, extra day of PTO, etc.). Ensure that employees and other stakeholders are working toward common goals, strengthen operations.

4.3. Constantly sharpen the hard and soft skills of technical personnel.

Begin continuous training for your DAML experts, particularly on data science tools and soft skills to enhance collaboration with their fusion teams. Allow unlimited certifications mean something to their peers, but also require proof of competency. Begin weekly seminars led by each DAML technical staff member on his/her focus area, ideas, and extremely focused discussions on challenges they are facing.

4.3.1. Incentivize technical team participation at local DAML groups and events.

Ask each DAML technical staff to begin attending all local relevant technical meetings and consider this part of their job and not “on their time and dime.” Propose and reward speaking at an event on how DAML is changing their organization, and/or taking a visible role in one or more DAML organization.



Figure 34: Try Hard to Actively Network with Meeting Attendees

4.3.2. HOST a DAML event.

Attract like-minded professionals and change the external perception of your organization by hosting a DAML event in the coolest company location. Use this to seed the community for job recruiting, mentors, and other events for the staff to attend. Propose hosting technical group events at the company facility for learning, collaboration and to help get the word out about hiring top talent and changing existing hiring practices.

4.3.3. Crossbow training?



Figure 35: STOP!!! Cross Training NOT Crossbow Training!

Cross training is the practice of training an employee to be able to do the work that another employee does, in addition to their primary job role. Incorporate cross-training initiatives—allow employees to experience the information, processes, methods, and day-in-the-life of 3-5 different departments/divisions, then gather lessons learned feedback.

Cross-Training and significant collaboration lead to increased empathy and awareness that enhance communications leading to meaningfully unique DAML products and services. This improved cooperation between organizational teams also enhances thoughtfully linking their knowledge and datasets to bridge disconnected knowledge and information silos, accelerate cross-training, and improve workforce cohesion.

4.4. Optimizing the learning experience with training and certification programs.



Figure 36: SCTE's Unique Focus on Cable Workforce Development Will Be Key.^{xliii}

DAML requires the addition of new skills, and SCTE will continue to serve the cable industry through existing and new agile, individually tailored training opportunities. SCTE's "Workforce 2027" initiative is an industry outreach to attract the future cable workforce and is aligned with CableLabs' 10G messaging. It includes components of DAML and is focused on enhancing cable employees' knowledge and capabilities in all key areas to learn in the individual's best means paired with question analytics (using DAML methods) to measure competency and help each individual grasp every crucial subject area.



Figure 37: Yes, shameless self-promotion here.

Coordinated with this effort is the "CORTEX[®] VirtuLearn" wholistic learning platform where "learners are getting a clear and concise overview of course content to prepare their minds for training with LightningMods[™] in as little as 10 minutes...eBooks provide in-depth content presented at a pace that

suits learners and provides them with knowledge of how and why...and...correlating games improve knowledge absorption and retention”.^{xliv}

4.4.1. Use DAML to optimize DAML training.

Create an organizational dashboard to track DAML individual progress. Information must include courses that individual employees are enrolled in and answer the following questions:

- What series of courses has been proposed?
- How far within the series of courses at a given point in time?
- How far within a specific course at a given point in time?
- Have they completed the course within the prescribed time or stalled in progress?
- Were there skipped modules along the way
- How long does it take to complete each course’s test?
- How many attempts did it take to pass a test?
- What was performance on the final course test and how it compared to the initial test?
- What was performance on the series of courses final test and how it compared to the initial test?
- Which questions were missed and what was the subject matter?
- Which competencies were mastered?
- Which competencies require additional coursework?

4.4.1.1. Focusing “Fusion Teams” to create breakeven net energy.




Figure 38: Fusion Teams Must Achieve 100 million Kelvin to Become a Self-Sustaining Reaction

Fusion Teams are “a multidisciplinary team that blends technology or analytics and business domain expertise and shares accountability for business and technology outcomes”^{xlv}. This *ongoing* approach, in contrast with the one-time Tiger Team, bridges technology and business areas to reduce failure and improve the value of the DAML project outcomes.




Figure 39: Average value DAML brings to the company based on the organizational model

Fusion Teams overcome the silo effect and finger-pointing between centralized teams, and are typically composed of DAML experts, technical, and business staff working together on a commonly defined, analyzed, and managed problem resolution towards a measurable goal.

One eventual outcome of continuous DAML training and pursuing fusion team projects in every department and division is for every single level of a company to have a ‘data champion,’ a well-known resource to go to for questions related to the area of interest. This more capable expertise inside each work group further catalyzes idea selections for automated processes.

4.5. DAML big picture ideas.

4.5.1. Showcase individual contributors or teams that are exceeding expectations.

Spotlight individuals who are passionate about using data to solve challenges in their role

Simply feature a prominent employee and how they approached a problem or opportunity to create a positive and measurable outcome. Positive peer pressure pushes others to be in that temporary spotlight.

Business leaders need to ensure that data is being used to inform employees’ decisions at every level of the company. An effective way to do this is to routinely spotlight individuals who are passionate about using data to solve challenges in their role. This can be only internal but consider publishing this in your customer newsletter or online current news. This not only celebrates people using data to achieve meaningful outcomes but creates “positive peer pressure” to actively push their peers to do the same.



Figure 40: Gift cards blur tax reporting requirements;)

4.5.2. Publicize DAML effort to all customers, shareholders, on website, and in press releases.

Instrument marketing, sales, and product management for external messaging to improve marketing and brand identity. Alter mission and vision statements on how DAML is now part of the organization mindset. Assist marketing, product management, HR, partnerships, etc. to communicate team's intense focus on continually sharpening their DAML skills. For outreach, do not overestimate the sophistication and literacy of the company's market, may need to provide resources, assistance, and training so they better understand how DAML will automatically optimize resolution of their most common issues.



Figure 41: YES, It's Real! Think We Just Made This Ourselves 🏆

4.5.3. Alter all hiring descriptions to focus on measurable DAML skills and accomplishments.

“According to a recent survey conducted by leading analytics platform, Tableau, half of all knowledge workers in the UK wouldn't want to work for a business that shows no sign of using data to inform decisions”xlvi.

Equip HR for external messaging to improve recruitment of higher-level candidates and ways to measure job evaluations specifically for DAML accomplishments.

4.5.4. DAML Program integrates with organizational change management.

Transforming an enterprise towards broad DAML adoption will require culture change impacting structure, workflows, and ways of working. This is not simply another department. This means to create a buzz by every employee about DAML, how their company leverages data effectively, and how they are

using it in their daily routines allowing them to step up in responsibility. This excitement will translate into a more engaging and questioning culture and necessarily will be felt by customers, analysts, press, and potential employees.

It will also reformulate operating models by enabling an iterative and experimental approach by those closest to the customers, opportunities, and obstacles to improve efficiencies and accelerate revenues.

4.5.5. Tie DAML program into innovation initiatives.

Many organizations have a vibrant innovation initiative in place, and DAML can add to this—or kick start it. Employees that use DAML to experiment with data will uncover new insights and opportunities that can lead to new intellectual property (IP), operational or best practices, products, and/or services.

5. Conclusion.

The crescendo of DAML has been building for decades. It has become a pivotal business strategy giving meaningful outcomes from optimizing, interconnecting, and advancing every occupational process and dataset. The goal of digital transformation is automating every reasonable, critical, and/or error-prone process using technologies like DAML. Ignoring DAML for an industry, organization, or personally is using an umbrella for an avalanche.

There are far too many DAML areas to focus on here, but only your organization can decide the tool based on the challenge to overcome. Three key, underlying areas are described to make sure they are not skipped.

By what means DAML is implemented is nearly exhaustively outlined to give your organization a toolbox of methods to begin, continue, and add to what will become your most important asset.

Abbreviations

DAML	data, analytics, and machine learning
AI	artificial intelligence
ML	machine learning
DAB	DAML advisory board
SaaS	software as a service
IoT	internet of things
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

- ⁱ cf. Marr, Bernard, *The Most In-Demand Metaverse Skills Every Company Will Be Looking For*, July 12, 2022, <https://www.linkedin.com/pulse/most-in-demand-metaverse-skills-every-company-looking-bernard-marr>.
- ⁱⁱ Melfi, Theodore, director. *Hidden Figures*. 20th Century Fox, 2016. 127 minutes. <https://family.20thcenturystudios.com/movies/hidden-figures>.
- ⁱⁱⁱ Marr, Bernard. The 20 Most Important Skills Everyone Needs to Succeed in A Digital World, excerpt from his books at <https://bernardmarr.com/books/>.
- ^{iv} The Economist, Taking Root: The world's leading companies are cultivating a new business culture, <https://dataculture.economist.com/takingroot/>.
- ^v Gartner, AI in Organizations Survey 2021, P-21023, Question: *What is your Organization's Use and Future Plans for AI?*, n=699.
- ^{vi} Gartner, *Fourth Annual Chief Data Officer Study*, ID: 378249, Q09: Which of these areas, if any, are critical to your Data and Analytics team's success?, n=255, Dec. 2018.
- ^{vii} Gartner, *2022 CDO Survey*, Q10: Which of these activities if any are explicitly included in your D&A team's strategy?, n=496 responses.
- ^{viii} Kraemer, Sue. *Bridging the data literacy gap between academia and the workforce*, Forrester commissioned research for Tableau, June 17, 2021, <https://www.tableau.com/about/blog/2021/6/bridging-data-literacy-gap-between-academia-and-workforce>.
- ^{ix} Sallam, Rita, and Ted Friedman, Gartner Webinar, *The Gartner Top Trends in Data & Analytics for 2022*, March 22, 2022.
- ^x Oh, Eunice. Carnegie Mellon University Student Newspaper. *Dystopian future arriving fast as robots steal human jobs*, February 3, 2014. <https://thetartan.org/2014/2/2/forum/robots/classic>.
- ^{xi} Manyika, James et. al. McKinsey Global Institute Report. *Jobs lost, jobs gained: What the future of work will mean for jobs, skills, and wages*. November 28, 2017. <https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages>.
- ^{xii} SOURCE: <https://www.forbes.com/sites/forbestechcouncil/2021/08/13/ai-will-not-replace-you-it-will-make-you-more-valuable/?sh=bb0640c766d5>
- ^{xiii} van der Meulen, Rob, Gartner Conference, *By 2020, Artificial Intelligence Will Create More Jobs Than It Eliminates*, Gartner 2017.
- ^{xiv} Dell Technologies, *Realizing 2030: A Divided Vision of the Future*, <https://www.delltechnologies.com/content/dam/delltechnologies/assets/perspectives/2030/pdf/Realizing-2030-A-Divided-Vision-of-the-Future-Summary.pdf>.
- ^{xv} Robinson, Ryan. Forbes, *3 Ways Leaders In The Workplace Can Create More Time For Deep Learning*, May 3, 2019, <https://www.forbes.com/sites/ryanrobinson/2019/05/03/leaders-workplace-create-time-deep-learning/?sh=618bc0f9b462>.
- ^{xvi} Applebone, Peter, Study Ties Educational Gains To More Productivity Growth, May 14, 1995, <https://www.nytimes.com/1995/05/14/us/study-ties-educational-gains-to-more-productivity-growth.html>.
- ^{xvii} Stevenson, Matson, HR Exchange Network, *7 Stats that Prove Training Value*. August 30, 2019, <https://www.hrexchangenetwork.com/learning/news/7-stats-that-prove-training-value>.
- ^{xviii} Cross, Libby, Learningpool, *Learning & Development Opportunities in the Workplace: A Benefit for a Growing Workforce*, Nov. 1, 2018, <https://learningpool.com/learning-development-opportunities-benefit-growing-workforce/>.
- ^{xix} VerBurg, Steve, Should I Invest in Employee Engagement?, <https://ocdalecarnegie.com/should-i-invest-in-employee-engagement/#>.
- ^{xx} Burvat, Jerome, et. al., CapGemini Digital Transformation Institute, *The Digital Talent Gap: Are Companies Doing Enough?*, https://www.capgemini.com/wp-content/uploads/2017/10/report_the-digital-talent-gap_final.pdf.
- ^{xxi} Performance Improvement Partners, *How to Harness Data Analytics to Drive Business Value and Maximize ROI*, <https://www.pipartners.com/data-analytics-business-value/>.
- ^{xxii} Feliu, Carolota, *HOW ANALYZING BIG DATA CAN INCREASE PROFIT IN YOUR COMPANY*, Datamize, <https://blog.datumize.com/how-analyzing-big-data-can-increase-profit-company>.
- ^{xxiii} Wegener, Rasmus, and Velu Sinha. *The value of Big Data: How analytics differentiates winners*, Bain & Company, September 17, 2013, <https://www.bain.com/insights/the-value-of-big-data>.
- ^{xxiv} Marr, Bernard. The 20 Most Important Skills Everyone Needs to Succeed in A Digital World, excerpt from his books at <https://bernardmarr.com/books/>.
- ^{xxv} Light Reading Webinar, *Tackling the 10G Challenge*, Cable Next-Gen Technologies & Strategies Digital Conference, March 15, 2022.
- ^{xxvi} Bastian, Chris. *Training for the future with an upskilled workforce*, <https://broadbandlibrary.com/workforce-forward/>.
- ^{xxvii} Rao, Anand, PwC, *Sizing the Prize: PwC's Global Artificial Intelligence Study: Exploiting the AI Revolution*, <https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html>.
- ^{xxviii} Gartner, *Fourth Annual Chief Data Officer Study*, ID: 378249, Q08: In 2017, in aggregate, how would you describe the operating model of your D&A organization?, n=255, Dec. 2018.
- ^{xxix} Rollings, Mike, Andrew White, Gartner, *Build a Data-Driven Enterprise*, March 4, 2020.
- ^{xxx} Panetta, Kasey. Gartner Information Technology, *A Data and Analytics Leader's Guide to Data Literacy*, August 26, 2021, <https://www.gartner.com/smarterwithgartner/a-data-and-analytics-leaders-guide-to-data-literacy>
- ^{xxxi} Little, Cindy, et. al., Forrester Trend Report, *Data Literacy Matters to Your Company's Success: Implement a Formal Program to Accelerate Your Journey to An Insights-Driven Business*, August 2, 2021, <https://www.forrester.com/report/data-literacy-matters-to-your-companys-success>
- ^{xxxii} Tableau, *Data Literacy*, <https://www.tableau.com/why-tableau/data-literacy>.
- ^{xxxiii} Belissent, Jennifer. *How Data-Literate Are You? Build A Curriculum Of ACES*, Forrester, February 17, 2020, <https://www.forrester.com/blogs/how-data-literate-are-you-build-a-curriculum-of-aces/>.
- ^{xxxiv} Sander, Ina, Internet Policy Review: Journal on Internet Regulation, *What is critical big data literacy and how can it be implemented?*, May 28, 2020, <https://policyreview.info/articles/analysis/what-critical-big-data-literacy-and-how-can-it-be-implemented>.

-
- xxxv Earley, Seth, *There Is No AI Without IA*, IT Professional (Volume: 18, Issue: 3, May-June 2016), IEEE, p. 58-64, <https://ieeexplore.ieee.org/document/7478581/authors#authors>.
- xxxvi Roder, Klaus, Twitter post from AI Summit Conference, "There is no AI without IA! (Information Architecture)," Sept. 26, 2019, <https://twitter.com/KlausRoder/status/1177285569530888192/photo/1>.
- xxxvii Babbage, Charles, Inventor of the First Programmable Device, FRS, December 26, 1791, subsequently used by George Fuechsel, and IBM technician/instructor.
- xxxviii van Ingen, Kevin, Garbage In. Garbage Out. A Developer Story, December 27, 2016, <http://getgareth.github.io/2016/12/27/Garbage-in,-garbage-out.-A-developer-story.html>.
- xxxix Rollings, Mike, Andrew White, Gartner, *Build a Data-Driven Enterprise*, March 4, 2020.
- xl Rollings, Mike, Andrew White, Gartner, *Build a Data-Driven Enterprise*, March 4, 2020.
- xli Lucid Content Team, Lucidchart, *Understanding the Tiger Team Approach*, <https://www.lucidchart.com/blog/what-is-a-tiger-team#:~:text=A%20tiger%20team%20is%20a,Apollo%2013%20mission%20in%201970>.
- xlii Fry, M. J., *I'm looking for something really dull and repetitive*, Search ID: CS184445, Cartoonstock.com.
- xliii O'Dell, Abigail. *Training Videos Aren't the Only Answer*, <https://broadbandlibrary.com/training-videos-arent-the-only-answer/>.
- xliv Fenton, Robin. *Leading Cable's Continuum for Training*, <https://broadbandlibrary.com/leading-cables-continuum-for-training/>.
- xlv Gartner, *Gartner Glossary*, Information Technology, <https://www.gartner.com/en/information-technology/glossary/fusion-team>.
- xlvi Tableau, *The World's Leading Companies are Cultivating a New Business Culture*, The Economist, July 6, 2020.

Encrypted DNS From Pilot To Production

A Technical Paper prepared for SCTE by

Joe Crowe

Product Development Engineer 5
Comcast
1800 Arch St Philadelphia, PA 19103
215.433.2103
joseph_crowe@comcast.com

Janardhan Bollineni, Comcast

Charlie Helfinstine, Comcast

Thomas Modayil Jacob, Comcast

Table of Contents

Title	Page Number
1. Introduction.....	3
2. DNS at Comcast.....	3
3. First Steps into Encrypted DNS	4
4. Path from the Lab into Production.....	5
5. Encrypted DNS at Comcast	6
6. Comcast's Commitment to Privacy	7
7. Encrypted DNS in the future	7
8. Conclusion.....	9
Abbreviations	9
Bibliography	9

List of Figures

Title	Page Number
Figure 1 – Recent 30-Day Graph of Comcast DNS Queries	3
Figure 2 – DoH Proxy Components	4
Figure 3 – DoH Using a Network Appliance	5
Figure 4 – Comcast DoH Queries Per Day.....	7
Figure 5 – Encrypted DNS Architecture Possibilities.....	8
Figure 6 – Encrypted DNS Using DPUs Block Diagram.....	8

1. Introduction

The domain name service (DNS) is one of the most critical internet services. It is often referred to as “the phonebook of the Internet”, meaning that the DNS facilitates a human-readable fully qualified domain name (FQDN) to be translated to a network IP address, which in turn allows networked devices to communicate to one other and provide content or needed services to allow applications to work as expected. The DNS was first introduced in 1983 by Paul Mockapetris and is one of the original Internet Standards per the IETF since 1986 (https://en.wikipedia.org/wiki/Domain_Name_System).

Since the advent of the DNS, it has been inherently insecure because DNS packets are transmitted in clear text either via the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP). There have been numerous initiatives to secure the DNS, notably DNS security extensions (DNSSEC), which encourage authoritative DNS operators to add extensions and caching DNS operators to perform validations. While this enhances security for the user, it doesn't solve the clear text request and response problems.

More recently, encrypted DNS protocols have been implemented across the Internet, including but not limited to, DNS over HTTPS (DoH), DNS over TLS (DoT), DNSCrypt, and in the near future DNS over QUIC (DoQ). Comcast is one of the first major ISPs to provide DoH and DoT to their customers and has also become a trusted recursive resolver with Mozilla’s browser Firefox.

2. DNS at Comcast

Comcast’s DNS infrastructure currently handles approximately 1.3 trillion queries per day at peak (fig. 1), with a portion of that traffic being encrypted. Decryption is accomplished using a network appliance front end to handle the DoH and DoT translations, which in turn hand off DNS queries to the backend DNS servers. Comcast also implemented DNSSEC validation at the caching layer and DNSSEC signing on most of their zones in 2011, with a commitment to make the DNS infrastructure more secure for their customers. (<https://corporate.comcast.com/comcast-voices/dnssec>)




Figure 1 – Recent 30-Day Graph of Comcast DNS Queries

3. First Steps into Encrypted DNS

In 2017, Mozilla was looking to design a new encrypted DNS protocol to help protect their users' privacy by limiting the exposure to the cleartext DNS packets. A request for comments draft (RFC) was submitted to the Internet Engineering Task Force (IETF), to outline the requirements and goals of DoH. RFC 8484 was published in October of 2018. This RFC helped drive some of the first DoH translators to take a TLS handshake, decrypt the packet, send the DNS query, get a response, re-encrypt, and send back to the client. (Hoffman & McManus)

In 2019, Mozilla partnered with Cloudflare to turn this feature on as a default for all Firefox users. The feature, which was turned on for all US-based customers in 2020, sends encrypted DNS traffic to Cloudflare's implementation of the protocol. Comcast engineers kept a close eye on what was being proposed and started looking into how to implement this protocol without a complete redesign of the whole DNS infrastructure.

Initially, Comcast's goal was to gain a comprehensive understanding of the protocol and how it works before the next steps of designing a solution that could go into production. Comcast hosts "lab weeks" twice a year for their engineers and in 2019 there was a proposal to create a DoH translator and to test it on the DNS infrastructure (fig. 2). This gave the engineers an opportunity to understand how the protocol works, latency measurements, and if it would affect any other services that Comcast offers to their customers.

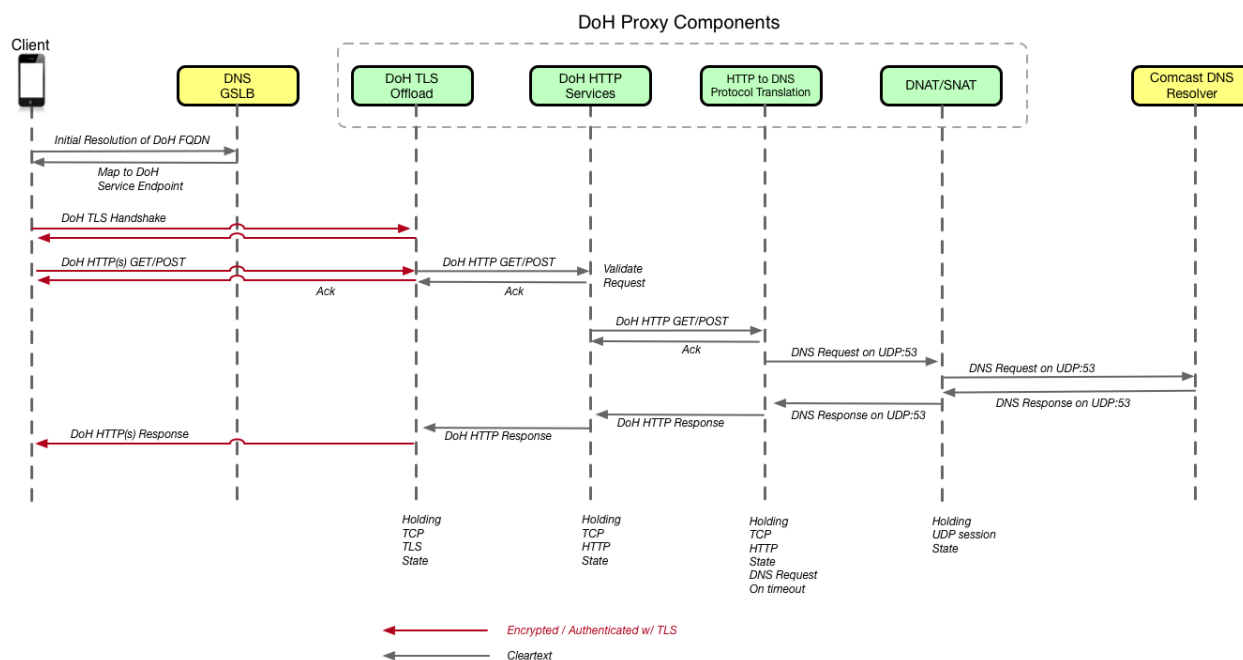


Figure 2 – DoH Proxy Components

DoH traffic utilizes Port 443, the same port that HTTPS uses. DNS traffic is hidden within packets on this port; therefore, a translator is needed to translate from HTTPS to UDP and vice versa. After the lab week in 2019, there was a push to utilize the learnings from the lab week project to build a solution that would work for Comcast. Since the DNS protocol uses mainly UDP and some TCP, the engineers did not want

to install a translator on the current DNS servers, due to unknown and hard-to-test performance capabilities on the current systems. The current server hardware was scoped out for a UDP-centric DNS server application, and the compute required for a TCP application with encryption at scale is drastically different. Therefore, a solution that efficiently reuses current infrastructure was pursued. During some of the initial conversations it was suggested to utilize a network appliance that can handle TLS offloading already, to front-end the DNS servers, and to have a translator on that appliance handle the TLS offloading, encryption/decryption, and forward to the current DNS infrastructure already in place.

4. Path from the Lab into Production

The team determined that using network appliances to perform the encryption and decryption was the most efficient way to introduce a solution to Comcast's infrastructure quickly. The engineering teams engaged with network appliance vendors and internal teams familiar with network appliances operation to develop a workable solution. The opening conversations were around the new DoH protocol and sharing RFC 8484 with the vendors. There was a need to have a network appliance handle the HTTPS connections, store the TLS certificate for the DoH FQDN, decrypt the DoH packet, translate the DoH packet to a DNS packet, forward that packet to backend DNS servers, receive that response, encrypt the packet, and then respond back to the client (Fig. 3).

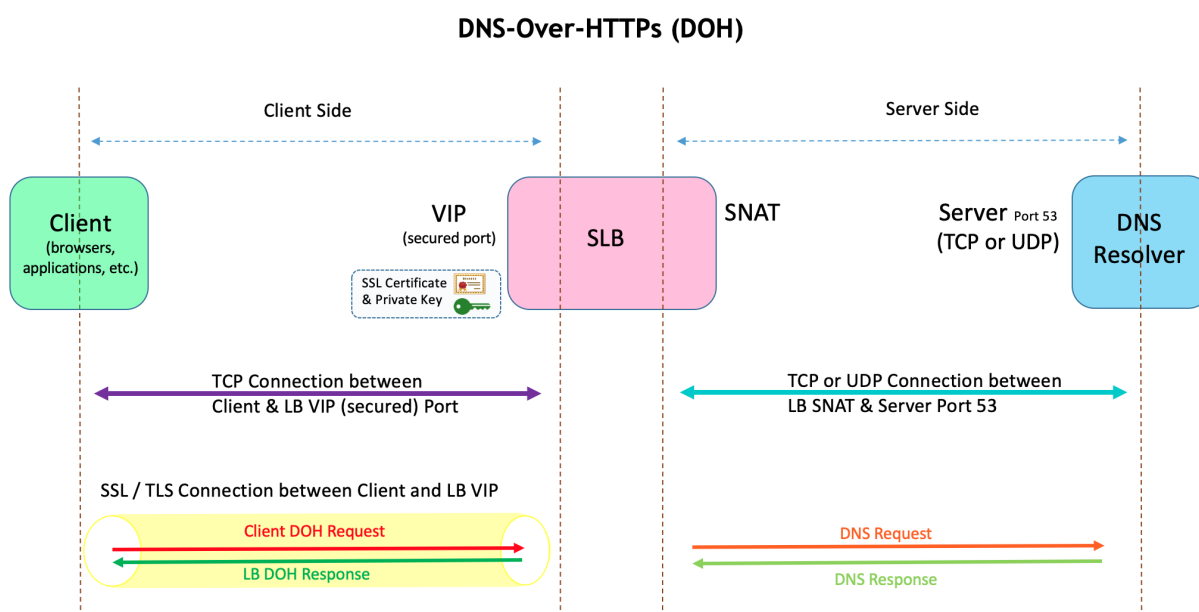


Figure 3 – DoH Using a Network Appliance

Each of the network appliance vendors were able to provide their versions of the specialized software that included the DoH translator and gave Comcast's engineers a chance to test and determine what would work from both a performance and financial perspective. Discussions landed on a specific device that could handle the query load and was relatively cost effective. During this early phase, there was discussion within EDDI and other DNS operator forums on how to stress test this new protocol. The toolsets for DoH operators to test the functionality, latency, capacity, and other unknown scenarios, were very limited in the early stages of testing. One of the toolsets that helped was an open-sourced tool called DoX which can be found at <https://github.com/wttw/dox>. This tool allowed testing against the DoH endpoint set up and comparisons with other known DoH endpoints. A few other tools were used

internally at Comcast to benchmark DoH end-points such as wrk and wrk2. The tools offer parameters that can be tuned to measure important metrics such as requests per second, connections per second, and performance provided by a DoH end-point at different levels of concurrent connection load. After there was comfortability with the network appliances and DoH functionality, there was an opt-in public beta test of the Comcast DoH service in October of 2019. This allowed other DNS operators to test the endpoint <https://doh.xfinity.com/dns-query>, give feedback, and help identify some of the issues that were not found in initial testing. The path to production involved knowing what the architecture would look like and how to provide the best service for customers.

5. Encrypted DNS at Comcast

Launching an encrypted DNS architecture posed a few technical challenges, particularly around ensuring high availability and localizing DNS responses for a given area where DNS is already being served to customers. Currently the way clients are directed to doh.xfinity.com is to either have the FQDN and IP address hard coded into the client or to utilize DNS to get a response. Utilizing geo load-balancing, the DNS lookups to doh.xfinity.com are shaped by where the DNS query originates. The DoH endpoint, doh.xfinity.com, utilizes a canonical name (CNAME) record that is directed to a geo network appliance, to help shape traffic to the correct DoH IP address for a region. Client lookup will go as follows, assuming that the customer is using Comcast's DNS servers provided by the dynamic host configuration protocol (DHCP):

1. Customer's client (browser, application, DNS forwarder..) will do an initial look up for doh.xfinity.com utilizing Comcast's DNS servers.
2. Based on the Comcast DNS resolver that queried for the CNAME endpoint, doh2.gslb2.xfinity.com, the network appliance will give a response for the corresponding caching DNS region's virtual IP (VIP) address.
3. The client will then connect to the encrypted DNS network appliance using the VIP and all subsequent DNS queries are now encrypted, with the VIP acting as the "client" doing queries for that customer.

The last step obfuscates the customers' source IP addresses to Comcast's DNS servers and give an extra layer of privacy. In Q2 2020 Comcast offered encrypted DNS to their customers. The queries per day for DoH currently sits around 90 billion at peak. (fig. 4). There are currently 12 points of presence on Comcast's network that handle the encrypted DNS functions, with 3 more points of presences slated to be deployed in 2022.




Figure 4 – Comcast DoH Queries Per Day

6. Comcast's Commitment to Privacy

During the initial path towards encrypted DNS services, Comcast worked closely with engineers from Mozilla, to find out how Comcast's DoH implementation could be offered to Firefox customers. Mozilla's trusted recursive resolver (TRR) policy had to be followed by any DNS operator to be considered as a DoH endpoint within Mozilla's Firefox browser. (Mozilla) Some of the requirements prompted Comcast to release a privacy statement just for DNS (Xfinity), along with updates to the broader privacy policy for the internet services. (Xfinity, 2021) Comcast has brought support and privacy commitments to Mozilla's TRR program, becoming the first major ISP to be part of that program and advocating for the privacy concerns of their customers (Mozilla, 2020) Along with becoming part of Mozilla's TRR program, work was done in conjunction with Google engineers to help provide Comcast's DoH endpoint to users of the Chrome browser. Furthering the exposure of the commitment to privacy utilizing encrypted DNS. Comcast also contributes to the Encrypted DNS Deployment Initiative (EDDI) (<https://www.encrypted-dns.org/>), a forum for DNS operators to collaborate on their findings of deploying encrypted DNS protocols on their respective DNS infrastructures, while helping shape some of the best practices for encrypted DNS deployments.

7. Encrypted DNS in the future

At Comcast, engineers are constantly looking at ways to improve on solutions in terms of scalability, cost-efficiency, and performance. One such work is harnessing the power of data processing units (DPUs) or SmartNICs to provide DNS encryption. Comcast has built a software solution that utilizes hardware TLS offload components provided by DPUs to efficiently translate between DoH and traditional UDP DNS. The rationale for exploring this method for solving the encrypted DNS problem is illustrated in the following diagram.

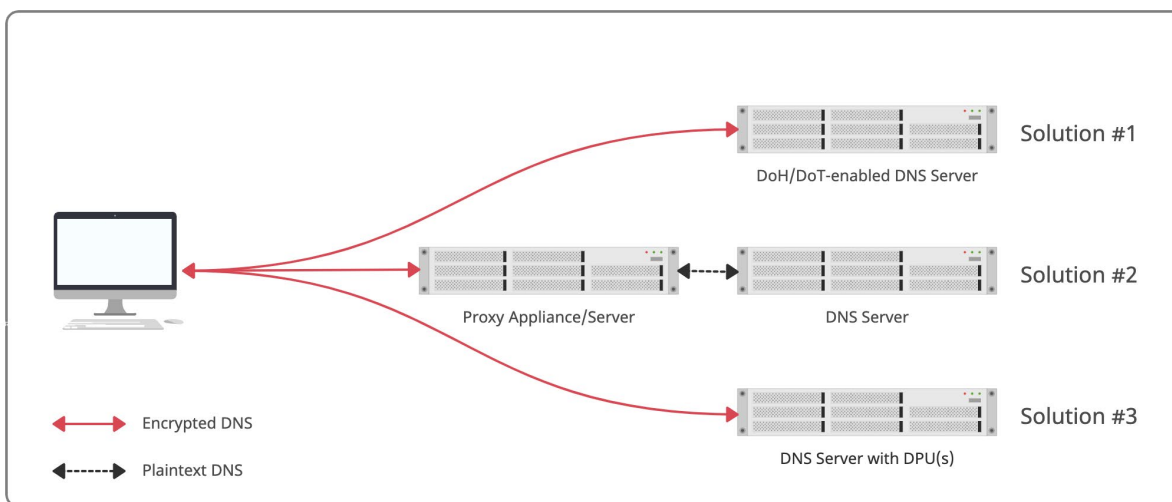


Figure 5 – Encrypted DNS Architecture Possibilities

Solution #1 describes enabling DoH and DoT features on the current DNS server application running on present hardware. As mentioned earlier, current server specifications do not account for DoH/DoT application requirements. In this model, scaling for capacity means adding more servers which results in increased power, space, and licensing costs.

Solution #2 scaling involves increasing the number of expensive hardware proxy appliances in the network.

Solution #3 offers a solution developed at Comcast using open-source components. The components are vendor DPU agnostic and can easily integrate with different DPU vendor offerings. The DPUs are a lot more cost-efficient compared to proxy appliances. Scaling capacity in this model means replacing DPUs with next generation DPUs or adding more DPUs per server. This method provides opportunities to deploy new services at the edge and gives Comcast the ability to cater to the evolving landscape of encrypted DNS standards.

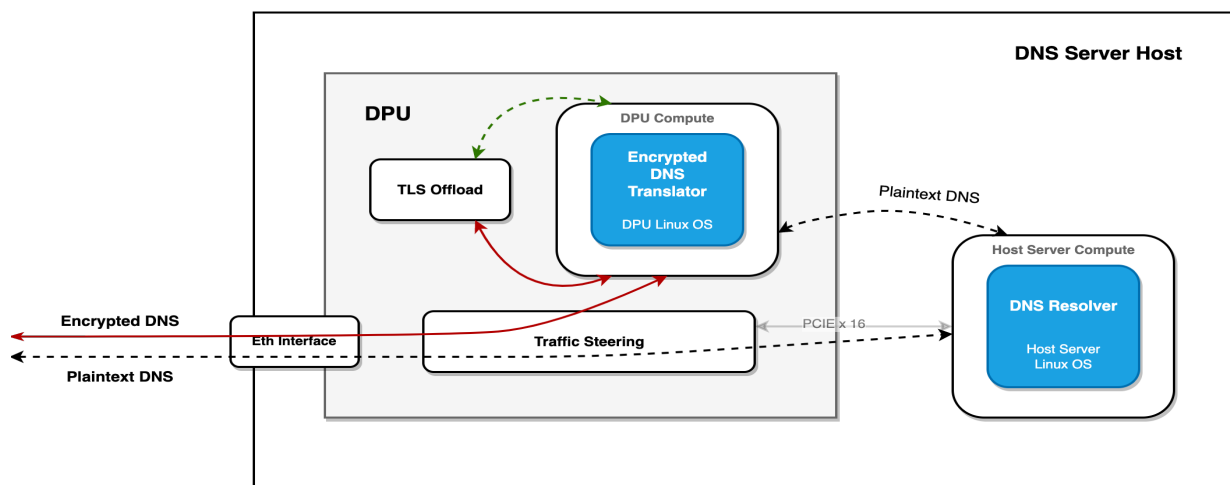


Figure 6 – Encrypted DNS Using DPUs Block Diagram

This solution is being actively benchmarked and tested and is a potential method to add encrypted DNS capacity across the Comcast footprint.

8. Conclusion

As mentioned previously, encrypted DNS continues to evolve, and internet providers have moving targets to hit as new standards are proposed and adopted by the industry. DNS over QUIC (DoQ) is currently a proposed standard for encrypting DNS using QUIC as the underlying protocol. QUIC is designed to reduce protocol induced delays with features such as mitigation of head of line blocking, zero round trip time session resumption (0-RTT) and advanced packet loss and congestion control mechanisms. Additionally, the next version of the HTTP protocol, HTTP/3, is also designed to run on QUIC as the underlying protocol. Comcast is closely monitoring these developments and is working on integrating QUIC into its encrypted DNS capabilities.

Abbreviations

DNS	domain name service
DHCP	Dynamic Host Configuration Protocol
DoH	DNS over https
DoT	DNS over TLS
DoQ	DNS over QUIC
DPU	data processing unit
FQDN	fully qualified domain name
HTTP(S)	Hypertext Transfer Protocol (secure)
QUIC	quick UDP internet connection
UDP	User Datagram Protocol
TCP	Transmission Control Protocol
FQDN	fully qualified domain name
TLS	transport layer security
VIP	virtual IP

Bibliography

- Hoffman, P., & McManus, P. (n.d.). *DNS Queries over HTTPS (DoH)*. Retrieved from RFC 8484, DOI 10.17487/RFC8484: <https://www.rfc-editor.org/info/rfc8484>
- Mozilla. (2020, June). *Comcast's Xfinity Internet Service Joins Firefox's Trusted Recursive Resolver Program*. Retrieved from <https://blog.mozilla.org/en/products/firefox/firefox-news/comcasts-xfinity-internet-service-joins-firefoxs-trusted-recursive-resolver-program/>
- Mozilla. (n.d.). *Security/DOH-resolver-policy*. Retrieved from <https://wiki.mozilla.org/Security/DOH-resolver-policy>
- Xfinity. (2021, October). *Our Privacy Policy explained*. Retrieved from <https://www.xfinity.com/privacy/policy>
- Xfinity. (n.d.). *Xfinity Internet DNS Privacy Statement*. Retrieved from <https://www.xfinity.com/privacy/policy/dns>

Establishing a Strong Security Posture for Open RAN

A Technical Paper prepared for SCTE by

Scott Poretsky
Director of Security, North America
Ericsson
Plano, TX, USA
(508) 261-4429
scott.poretsky@ericsson.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. RAN Architectures.....	3
3. Open RAN Security Posture	5
3.1. Open RAN Security Tradeoffs.....	5
3.2. O-RAN Attack Surface	6
4. Additional Security Considerations	7
4.1. Zero Trust Architecture	7
4.2. Security Controls for 5G Cloud Deployments	8
5. Securing the O-RAN Architecture	9
5.1. Open Fronthaul	9
5.1.1. M-Plane Security.....	9
5.1.2. CUS-Plane Security	11
5.2. RICs and RAN Applications	12
5.3. Service Management and Orchestration	14
5.4. O-Cloud	15
6. Conclusion.....	16
Abbreviations	17
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - RAN Splits	5
Figure 2 - Open RAN Security Advantages and Risks	6
Figure 5 - OFH M-Plane Security.....	10
Figure 6 - Certificates in the Optional M-Plane Models	11
Figure 7 - O-RAN Open Fronthaul Attack Vectors.....	12
Figure 8 - Open RAN RICs and RAN Applications [adapted from 32].....	13
Figure 9 - Mutual authentication on the A1 interface	14
Figure 10 - Mutual authentication on the R1 interface.....	14

List of Tables

Title	Page Number
Table 1 – Alignment of 3GPP 5G Standards to NIST ZTA Tenets.....	7

1. Introduction

5G will deliver significant societal value as it provides critical infrastructure, mission critical applications, smart manufacturing, connected car, and other use cases. As a result, our risk tolerance must decrease due to the increased impact from a cyberattack on a 5G network. Radio Access Networks (RAN) is evolving to Open RAN, including Cloud RAN and O-RAN, characterized by a disaggregated, virtualized, cloud-native, automated, and intelligent network. Open RAN brings many security benefits, including vendor diversity, but it also introduces security risks that must be managed to ensure Open RAN deployments have a strong security posture.

Cloud security risks are not exclusive to Open RAN, but they must be considered during a risk analysis to ensure secure Open RAN deployments. Cloud deployments of Open RAN can offer many security advantages inherent from third-party cloud-based services while also expanding the RAN attack surface due to increased internal threats. Open RAN must be built upon a zero trust architecture (ZTA) to mitigate risks from internal and external threats. This is a new paradigm for securing RAN, where traditional on-premise networks have focused primarily on protection from external threats.

The O-RAN architecture, from the O-RAN Alliance, expands the attack surface by specifying new functions and interfaces built on the 3GPP standardized architecture. The Lower Layer Split (LLS 7-2x) with the Open Fronthaul (OFH) interface, as well as RAN Intelligent Controllers (RICs), with xApps and rApps, and the Open Cloud (O-Cloud) must all be secured to protect O-RAN's network functions, interfaces, and data. The Service Management and Orchestration (SMO) can enhance the Open RAN security posture, but it must also be securely designed and implemented to prevent internal and external threat actors from gaining access and control.

The goal of this paper is to present the security risks and recommend security controls to establish a strong security posture for Open RAN deployments. Section 2 of the paper provides a baseline discussion of evolving Open RAN architectures, including O-RAN from the O-RAN Alliance. Section 3 introduces the Open RAN security posture, presenting the security tradeoffs of Open RAN and expanded attack surface of the O-RAN architecture. Section 4 discusses additional security considerations for a strong Open RAN security posture with focus on ZTA and cloud security to support deployment of 5G critical infrastructure for protection against external and internal threats. Section 5 provides a detailed technical analysis of the security risks across the O-RAN architecture and recommends security controls to mitigate those risks. Recommended security controls are provided with the goal to strive towards a ZTA that protects against internal and external threats, ensuring Open RAN deployments will be secure.

2. RAN Architectures

The 5G network is built from Radio Access Network (RAN) and Core (5GC). RAN uses radio frequencies to provide wireless connectivity to devices for delivery of applications and consists of antennas, radios, baseband (RAN compute), and RAN software enabling high data rates for innovative mobile use cases. Antennas radiate the electrical signals into radio waves and the radio converts digital information into signals that can be transmitted wirelessly while ensuring the transmitted signals are in the assigned frequency bands at the configured power levels. The baseband provides signal processing functions for efficient wireless communication and secure use of spectrum to deliver extremely high data processing speeds.

Traditionally, baseband functionality has been complex software providing intelligence to assign data bits to available frequency and time slots and prioritizing users on a millisecond (or sub-millisecond) basis, running on proprietary hardware deployed at cell sites. With the evolution to virtualization and containerization of network functions, baseband functionality can be implemented in software to operate on Commercial Off the Shelf (COTS) server hardware at edge sites, or central sites, co-located with 5GC components. This means that RAN and 5GC software may be geographically co-located, deployed on the same infrastructure, and managed as a single solution. As a result, the security of 5G RAN has evolved to be as critical and sensitive as security of the 5GC.

Open RAN, including Cloud RAN, such as [1], and O-RAN from the O-RAN Alliance [2], is a general term for open radio access network architectures defined by open and interoperable interfaces, virtualization, cloudification, and intelligence enabled through AI/ML. Open RAN solutions use the 3GPP-specified air interface that provides security features such as signaling confidentiality and integrity protection, user plane confidentiality and integrity protection, and the Subscription Concealed Identifier (SUCI) for subscriber privacy [3].

Cloud RAN is based upon the Third-Generation Partnership Project (3GPP) Release 15 (R15) Higher Layer Split (HLS) having the RAN Compute disaggregated into a Central Unit (CU) and Distributed Unit (DU). The CU and DU use the enhanced Common Public Radio Interface (eCPRI) fronthaul interface between them, as shown in Figure 1. Cloud RAN provides the advantages of open, standards-based cloud-native network functions managed by a SMO, without requiring a LLS.

The O-RAN architecture introduces a LLS disaggregating the RAN's O-DU and O-RU with the Open Fronthaul interface between them, as shown in Figure 1. The O-RAN Alliance's Open Fronthaul interface specifies a Control, User, Synchronization Plane (CUS-Plane) and Management Plane (M-Plane) running over eCPRI. The primary goal of the disaggregation to the O-DU and O-RU is to further increase vendor diversity in the RAN.

The O-RAN architecture also introduces the Near-Real-Time RAN Intelligent Controller (Near-RT RIC) and SMO, with an internal Non-RT RIC, for automation, orchestration, and optimization of RAN functions and performance. The Near-RT RIC and Non-RT RIC are specified to support RAN applications, known as xApps and rApps, respectively, with the goal to enhance RAN innovation and optimization through an ecosystem of purpose-built applications from RIC platform vendors and third-parties. O-RAN specifications also provide an O-Cloud for infrastructure upon which the O-RAN network functions run as applications.

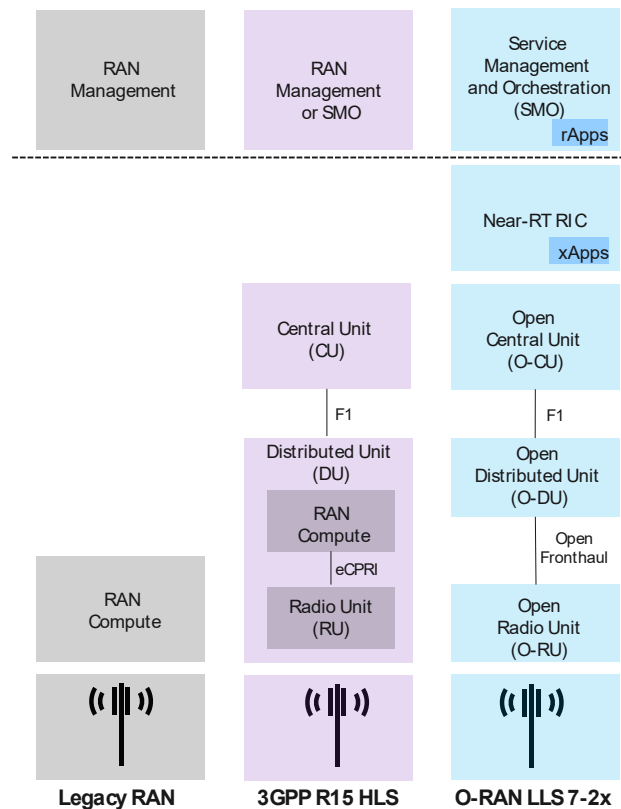


Figure 1 - RAN Splits

3. Open RAN Security Posture

Open RAN solutions and architectures, including Cloud RAN and O-RAN, share common security advantages and disadvantages. This section discusses the security tradeoffs for Open RAN and examines the expanded attack surface specific to the O-RAN architecture.

3.1. Open RAN Security Tradeoffs

A security posture of a telecommunications network is the security status of a network, information, and systems based on security controls *in place* to manage the defense and react to situational changes [4]. Open RAN, including Cloud RAN and O-RAN, provides enhanced RAN security including use of open-source software enabling transparency and common control; open interfaces ensuring transparency and use of standard, interoperable, and secure protocols; disaggregation enabling supply chain security through vendor diversity; and use of AI/ML enabling visibility and intelligence to achieve greater security [5]. However, Open RAN solutions also tradeoff introduction of security risks, such as open-source software vulnerabilities exploited by malicious threat actors, new interfaces with weak security specifications, and architectural modifications that expand the RAN attack surface. These security tradeoffs are summarized in Figure 2. Open RAN security risks were first analyzed in [6] and the O-RAN Alliance Working Group 11 (WG11) (formerly Security Focus Group – SFG) has been evolving O-RAN’s security specifications to support a strong O-RAN security posture.

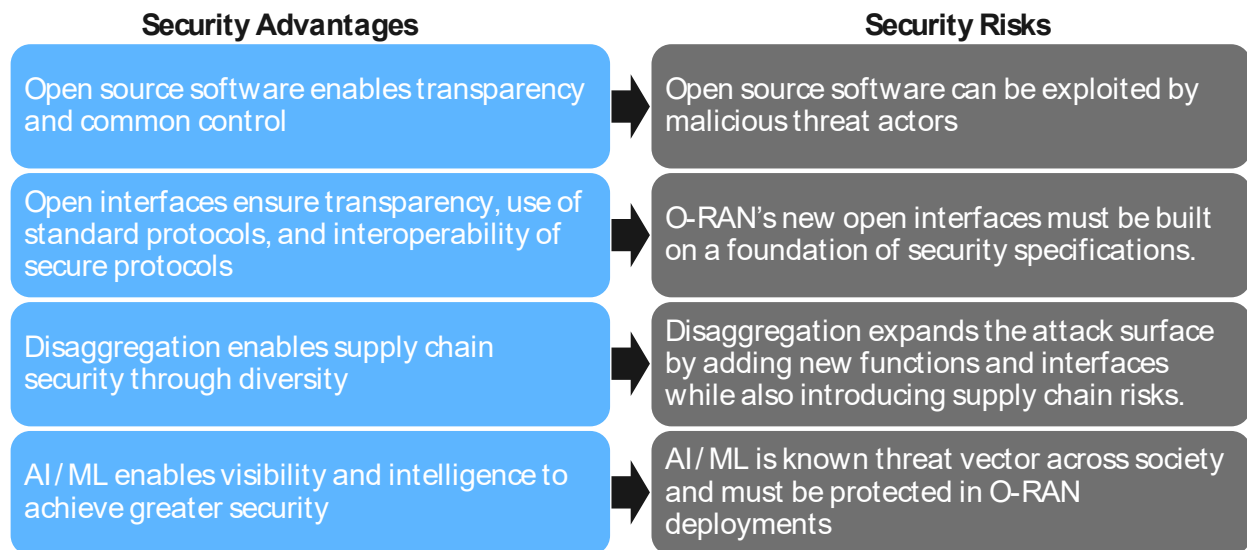


Figure 2 - Open RAN Security Advantages and Risks

3.2. O-RAN Attack Surface

O-RAN introduces architectural changes through disaggregation, opening the ecosystem for increased vendor diversity. The architectural changes that define O-RAN are the LLS 7-2x, OFH interface, RICs, and RAN applications known as rApps and xApps. However, O-RAN's new network functions and interfaces expand the O-RAN attack surface [7], [8], [9]. A strong O-RAN security posture must implement security controls at each layer of the architecture to protect the network functions, interfaces, and data from external and internal threats, as shown in Figure 3. The O-RAN Alliance's WG11 has performed a detailed threat analysis of O-RAN [10] and continues to evolve O-RAN's security specifications to meet the security baseline expected by network operators and their users. The specification effort considers a Zero Trust Architecture (ZTA) in accordance with US NIST SP 800-207 [11] to provide protection from external and internal threats.

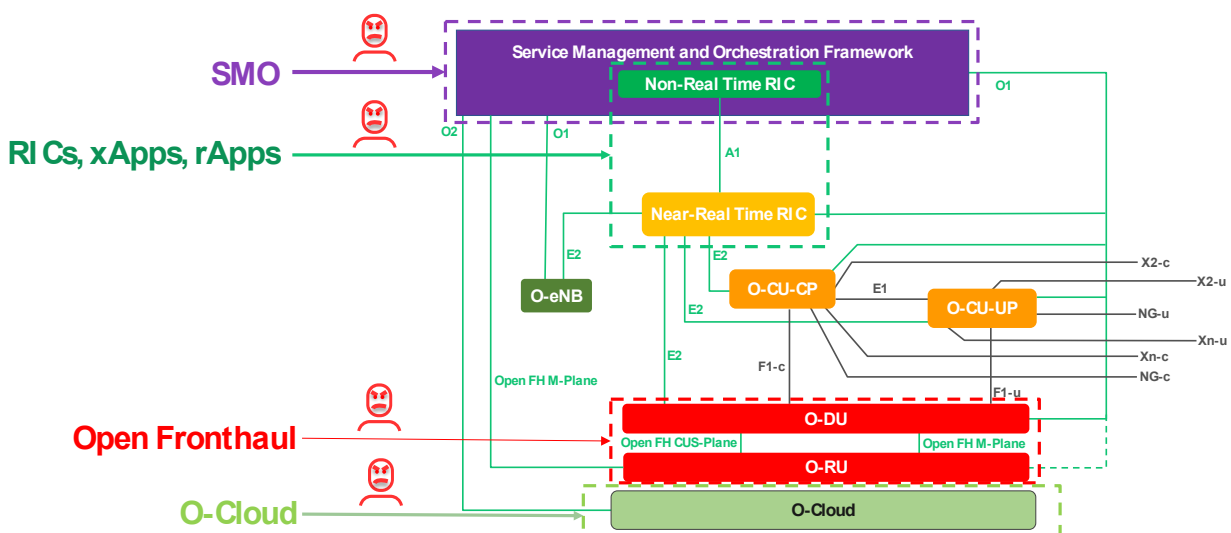


Figure 3 - O-RAN Expanded Attack Surface [adapted from 12]

4. Additional Security Considerations

Additional security considerations should be made to achieve a strong Open RAN security posture built upon a ZTA for cloud deployments. This section discusses ZTA and cloud security to support deployment of Open RAN critical infrastructure that provides defenses against internal and external threats.

4.1. Zero Trust Architecture

5G is the first generation of mobile technology designed for cloud deployments of RAN and Core. Open RAN enables cloud migration of the RAN to leverage its benefits of rapid elasticity, on-demand self-service, broad network access, and multi-tenancy. However, cloud deployments introduce an expanded threat surface due to the increased risk of internal threats, shifting the security paradigm to building a ZTA from the traditional perimeter-based security focused on protecting against external threats. In a ZTA, this is no longer sufficient because we must design for a perimeter-less network [13] that assumes the adversary is already inside the network [14]. This is a new paradigm for RAN security, as RAN has been traditionally secured at the perimeter because it has run on operator hardware in an operator managed network in an operator facility assuming internal trust. 3GPP releases 15 and 16 specify 5G with security features that align well with the NIST seven tenets of a ZTA. Some examples are provided in Table 1 below with additional examples provided in [15].

Table 1 – Alignment of 3GPP 5G Standards to NIST ZTA Tenets (Examples)

#	ZTA Tenet	5G Feature
1	All data sources and computing services are considered resources	The end-to-end 5G network, including UEs, RAN, Transport, Core, Applications, and Services are assets and data sources
2	All communication is secured regardless of network location	Subscriber identity privacy using SUCI . TLS provides confidentiality and integrity protection across the SBI .
3	Access to individual resources is granted on a per-session basis	UE access is granted using 5G-AKA, EAP-AKA', and EAP-TLS. Authentication and authorization between NFs over SBI in the 5GC is provided with certificate-based mutual authentication using TLS
4	Access to resources is determined by dynamic policy	The PCF feeds the AMF with access and mobility policies that affect UE authorization to access 5G network resources
5	The operator monitors and measures the integrity and security posture of all owned and associated assets	NWDAF incorporates standard interfaces from the SBA to collect data and evaluate systems in terms of compliance with security policy rules
6	All resource authentication and authorization are dynamic and strictly enforced before access is allowed	Mutual authentication enables the device to authenticate the network using the AUTH (Authentication Token) returned by the network
7	The operator collects information about the current state of assets, network infrastructure and communications and uses it to improve its security posture	The MNO should have a mature supply chain risk management to ensure NFs are compliant with GSMA NESAS.

NOTE: See acronym list for acronyms used in Table 1.

While Open RAN provides opportunity to deliver an ecosystem of innovative vendors and products, mobile network operators (MNOs) are accountable to evaluate the security posture of its deployment. A threat analysis identifies potential threats, vulnerabilities, and exploits while a risk analysis determines the likelihood and impact of attacks compromising confidentiality, corrupting integrity, or degrading availability. A ZTA approach may increase risk likelihood scores due to consideration of external and

internal threats, while implementation of security controls can decrease risk impact scores. The O-RAN Alliance WG 11 is evolving O-RAN's security specifications to align with industry best practices and meet the security baseline established by 3GPP for 5G, while pursuing a ZTA [16].

4.2. Security Controls for 5G Cloud Deployments

As 5G networks are critical infrastructure, it is important to properly secure cloud deployments to protect against external threat actors at the perimeter and internal threat actors exploiting zero-days, performing lateral movement for reconnaissance, and conducting advanced persistent threats (APTs). With the evolution of 5G to public cloud and hybrid cloud deployments, the 5G attack surface expands due to running on third-party infrastructure in a multi-tenant environment managed by another third-party. The cloud introduces increased internal threats from lateral movement, reconnaissance, and advanced persistent threats (APTs) from nation-state, criminal, and internal threat actors. Open RAN and 5G Core have increased risk of internal threats in the cloud due to increased dependency on cloud service providers, lack of defined security roles across stakeholders, resource sharing with other tenants, greater risk of security misconfiguration, and increased use of open-source software [17].

A risk-based approach must be taken to select the proper security controls for Open RAN deployments to mitigate internal and external threats in the cloud, or O-Cloud. Each layer of the cloud stack must be secured to reduce risk from potential vulnerabilities being exploited by internal or external threat actors. 5G cloud deployments should pursue a ZTA to align with US DHS Cybersecurity and Infrastructure Security Agency (CISA) guidance with the following capabilities [18]:

- prevent and detect lateral movement
- secure isolation of network resources
- data protection
- ensure integrity of cloud infrastructure

Common vulnerabilities, such as misconfigurations, weak authentication and use of open-source software with known vulnerabilities, can be prevented using industry best security practices. Well-known attacks in the cloud, including container escape, host escape, shared resource exhaustion, remote code execution, information disclosure between tenants, distributed denial of service (DDoS), and advanced persistent threats (APT) must be mitigated for Open RAN deployments [19]. Security controls must be provided at each layer of the cloud stack, as shown in Figure 4, to protect data, containers, container runtime engines and orchestration, operating systems, and infrastructure including servers, networks, and storage.

Recommended controls include micro-segmentation, tenant isolation and container isolation, mutual transport layer security (mTLS) 1.2, or 1.3, and X.509 certificates, Identity and Access Management (IAM), Multi-Factor Authentication (MFA), and role-based access controls (RBAC) for user access. To help ensure a secure and trusted runtime environment, Open RAN deployments should operate on a hardware root of trust using hardware security modules (HSM), a purpose-built appliance compliant with 3GPP security standards for hardware-based storage and lifecycle management of cryptographic keys. 5G critical infrastructure deployed in the cloud must have continuous monitoring, logging, and alerting with periodic vulnerability assessments and configuration validation to protect against evolving threats. Security controls for Open RAN deployments in the cloud are discussed further in [20].

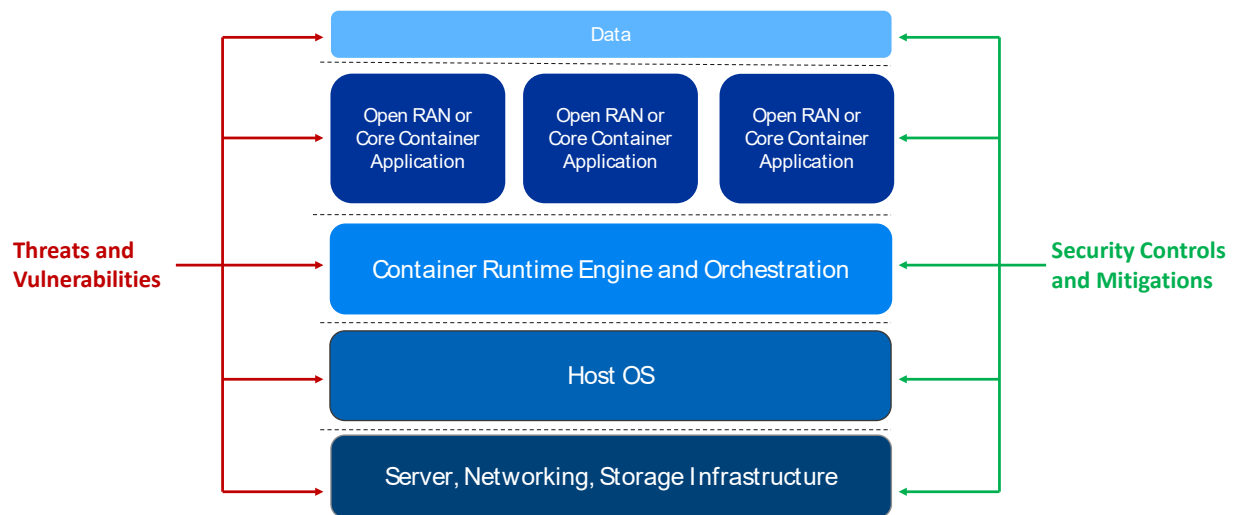


Figure 4 – Multi-Layer Security for 5G Cloud Deployments

5. Securing the O-RAN Architecture

A strong O-RAN security posture provides security controls for protection from external and internal threats introduced by O-RAN's expanded attack surface due to the architectural changes with new network functions, interfaces, and data. This section provides analysis of the threats, risks, and controls for O-RAN's OFH, RICs and applications, SMO, and O-Cloud.

5.1. Open Fronthaul

Fronthaul carries data between the 5G radio and RAN compute nodes using eCPRI, an industry consortium interface specification that utilizes an ethernet for packet forwarding [21]. The design of eCPRI is highly resource efficient and provides the flexibility for different deployment scenarios and functional splits, while enabling use of standard secure IP-based protocols, such as mTLS 1.2, on the fronthaul. The choice of ethernet in eCPRI enables the packet-based fronthaul, COTS, and Open RAN.

O-RAN's OFH interface, using LLS 7-2x, provides O-RAN Alliance specified Control, User, and Synchronization Plane (CUS-Plane) [22] and Management Plane (M-Plane) [23] over eCPRI to provide message exchange between the O-DU and O-RU for coordination. The C-Plane runs over eCPRI to provide message exchange between the O-DU and O-RU for scheduling and beamforming, numerology, and spectrum sharing control. The U-Plane runs over eCPRI to provide uplink and downlink frequency domain IQ data samples. The S-Plane provides timing and synchronization of the O-DUs and O-RUs using Synchronous Ethernet and IEEE 1588 Precision Time Protocol (PTP) [24]. The M-Plane manages and initializes the connection between the O-RU and O-DU. The OFH should have security controls implemented to protect against external and internal threats, consistent with a ZTA. M-Plane and CUS-Plane security are discussed in the sections below.


5.1.1. M-Plane Security

The O-RAN M-Plane specification has mandatory requirement for vendors to support two methods of authentication [25]:

1. O-RU supports Secure Shell version 2 (SSHv2) with password-based authentication of the O-DU and the O-DU supports SSHv2 with simple public key based authentication of the O-RU.
2. The O-RU and O-DU support mutual authentication with TLS 1.2, or higher, and X.509 certificates.


While it is mandatory for the O-RU and O-DU vendors to support both methods of authentication, the operator has the option which to use in production. Password-based authentication is considered weak security for critical infrastructure [26] due to the efficiency in which an attacker can perform a brute force attack upon gaining access to the interface, which can be exploited for lateral movement to northbound functions of the O-RAN architecture to execute broader network attacks. Best security practice, consistent with 3GPP security standards, is for operators to use mTLS with X.509 certificates for 5G deployments [27], including O-RAN deployments for which “it is recommended that operators use NETCONF/TLS and FTPES in production networks” on the M-Plane [28]. This tradeoff for M-Plane security is shown in Figure 5 below.

SSHv2 with Password-based Authentication



- Weak Security
- Does not meet industry best practice
- Violates USG guidance

mTLS with Certificate-based Authentication



- Strong Security
- Meets industry best practice
- Aligns with USG Guidance

**Both are mandatory for vendors to implement
and optional for operators to use**

Figure 3 - OFH M-Plane Security

The M-Plane has two optional deployment models, Hybrid and Hierarchical, as shown in Figure 6, which influences the operator’s approach to certificate management. In the Hybrid model, the O-RU and O-DU have direct IP connectivity to Public Key Infrastructure (PKI) and management systems. The O-RU and O-DU are considered separately managed entities by the SMO platform and enroll their unique operator-signed certificates in a Certificate Authority/Registration Authority (CA/RA) server using Certificate Management Protocol version 2 (CMPv2). PKI provides full lifecycle management of certificates for the O-DU and O-RU and the unique O-RU and O-DU enrolled operator-signed certificates are used to establish the secure mTLS session between O-RU and O-DU.

In the Hierarchical model, the O-DU has direct IP connectivity to PKI and management systems, enrolls its unique operator-signed certificate in the CMPv2 capable CA/RA server, and is considered a managed entity by the SMO platform. The O-RU does not have direct IP connectivity to PKI and management systems and is not considered a managed entity. During start-up installation the O-RU optionally downloads configuration from O-DU, which includes the trust anchor (root certificate) for the O-DU

operator-signed certificate used by O-RU to authenticate the O-DU. The configuration may also include identity of a CMPv2 capable CA/RA server reachable through the O-DU, if supported in the production deployment, for certificate enrollment. The O-RU establishes mTLS sessions with the O-DU by using either an enrolled operator-signed certificate in the CA/RA server reachable through the O-DU or its factory installed vendor-signed certificate when the CA/RA server is not reachable through the O-DU, as shown in Figure 6 below. The O-DU installs the O-RU vendor-signed root certificate used for authenticating the O-RU in the Hierarchical model deployment scenario that the O-DU does not provide the O-RU connectivity to a CMPv2 capable CA/RA.

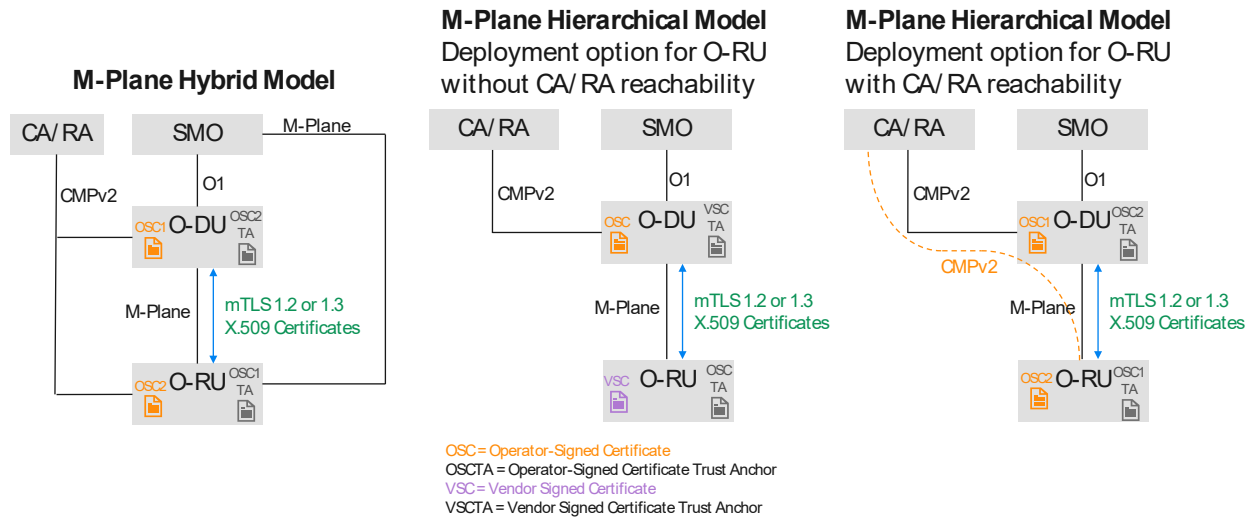


Figure 4 - Certificates in the Optional M-Plane Models

5.1.2. CUS-Plane Security

The OFH interface C/U/S-Plane has control and synchronization messaging that is unauthenticated and in the clear, enabling the man-in-the-middle (MitM) attack vectors shown in Figure 7 and as follows:

- C-Plane - Intercept messages to learn subscriber and network information.
- C-Plane - Message spoofing to inject false information to influence network parameter settings.
- S-Plane - Impersonation of PTP Master Clock or Grand Master, which can be exploited to degrade U-Plane performance and availability.

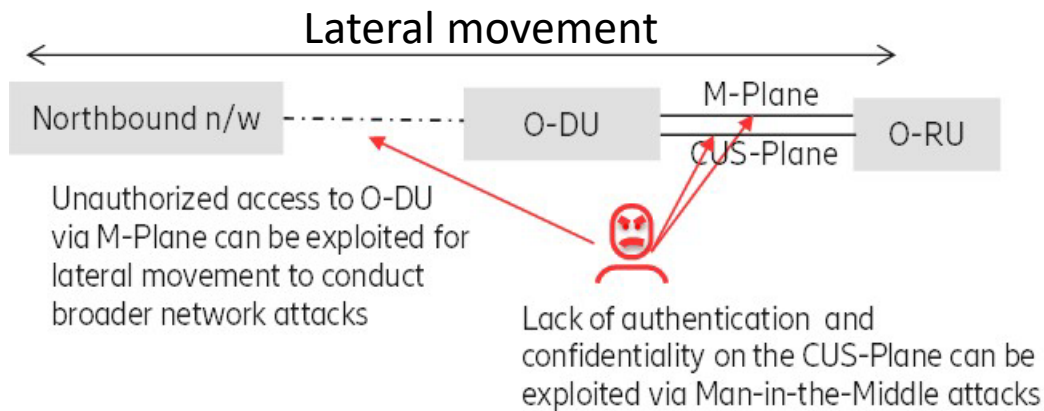


Figure 5 - O-RAN Open Fronthaul Attack Vectors

Security controls for confidentiality and integrity of messages on the CUS-Plane have been limited due to latency requirements on the CUS-Plane [29]. Further study is needed to identify potential security solutions. IEEE 802.1X-2020 Port-based Network Access Control 0 can be configured to provide protection of OFH interfaces at the physical layer for secure network access in point-to-point LAN segments within the Open Fronthaul network [31].

5.2. RICs and RAN Applications

Open RAN architecture describes two new RIC frameworks, Non-RT-RIC and Near-RT-RIC, for hosting automation applications known as xApps and rApps, respectively, as shown in Figure 8. The xApps and rApps enhance RAN innovation and optimization through an ecosystem of purpose-built applications from RIC platform vendors and third-parties. The Near-RT RIC and xApps provide automation and management of use cases with a suggested control loop of 10 msec to 1 second, while the Non-RT RIC and rApps provide automation and management of use cases with a suggested control loop of one second or more. The Non-RT RIC uses its rApps to decide RAN policy that it pushes to the Near-RT RIC, and its xApps, across the O-RAN Alliance specified A1 interface. The R1 interface between the SMO, Non-RT RIC, and rApps enables any rApp, as a Service Producer or Service Consumer, to work with any SMO and other rApps to enable insights from one rApp to serve as input to another, forming more complex decision-making capabilities.

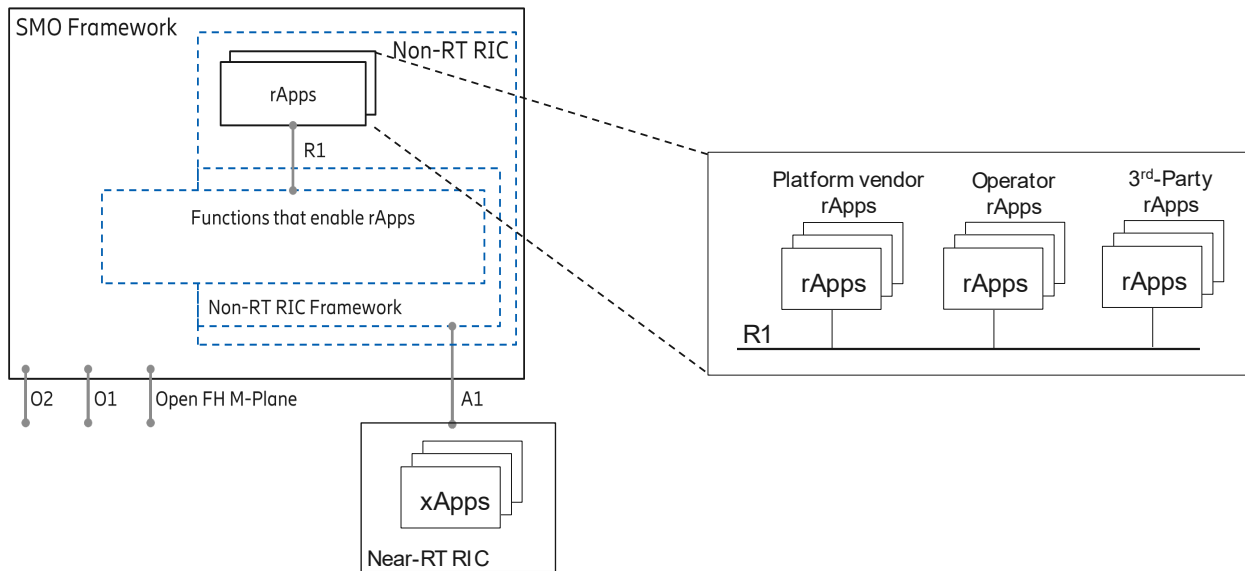


Figure 6 - Open RAN RICs and RAN Applications [adapted from 32]

A primary driver to have xApps and rApps is to broaden O-RAN innovation through smaller best of breed vendors to offer microservices leveraging AI/ML technology for use cases such as spectral efficiency, handover management, network optimization, and network healing. The specification of the A1 and R1 interfaces enables third-party xApps and rApps to efficiently integrate into an Open RAN deployment. However, considerations must be made for architectural and supply chain security risks to the RAN introduced by the RICs and their xApps/rApps. The risks and appropriate security controls are as follow:

- Risk: RICs may decide RAN parameter settings that have direct or indirect conflicts with local gNB decisions, which can degrade performance or availability. The risk of parameter conflicts increases as the number of third-party xApps vendors increases.
 - Solution: The Near-RT RIC supports a conflict mitigation function.
- Risk: Conflicts between rApps from multiple vendors could unintentionally or maliciously push conflicting RAN policies and parameter settings to degrade performance or availability. The risk of parameter conflicts increases as the number of third-party rApps vendors increases.
 - Solution: The Non-RT RIC supports a conflict mitigation function.
- Risk: Use of unsigned, untrusted, or improperly secured third-party xApps or rApps can introduce risks to deployments.
 - Solution: xApps and rApps are digitally signed, penetration tested, and vulnerability scanned prior to delivery. xApps and rApps are securely on-boarded and monitored for anomalous behavior in production. xApps and rApps must support logging and export of logs to the SMO.
- Risk: Malicious Near-RT RIC or Non-RT RIC can attempt to gain access across the A1 interface
 - Solution: The A1 interface supports mutual authentication using mTLS with X.509 certificates.




Figure 7 - Mutual authentication on the A1 interface

- Risk: Malicious rApps can attempt to gain access to other rApps through the Non-RT RIC Framework
 - Solution: The R1 interface supports mutual authentication using mTLS with X.509 certificates.




Figure 8 - Mutual authentication on the R1 interface

- Risk: xApps or rApps could be exploited by internal or external threat actors to gain access to private personal information or sensitive business information.
 - Solution: Near-RT RIC, Non-RT RIC, Non-RT RIC Framework, and rApps support authorization using OAuth 2.0. xApps and rApps provide confidentiality protection for sensitive data at rest. Data in motion across the A1 and R1 interfaces has confidentiality and integrity protection using TLS 1.2, or 1.3.

5.3. Service Management and Orchestration

The SMO is an intelligent automation platform for Open RAN, including Cloud RAN and O-RAN, radio resources that applies automation at scale to simplify the complexity of networks, improve network performance, enhance customer experience, and minimize RAN operational costs. The SMO, as a component of the operational support system (OSS), enables automation and increases the abstraction offered to users by managing Open RAN as a service and intents. The O-RAN Alliance defines technical specifications and interfaces related to the O-RAN's SMO Framework.

The SMO, through its Non-RT RIC, provides policy-based guidance and enrichment information to the Near-RT RIC. The Non-RT RIC is an automation platform that uses rApps to deliver higher layer automation policies, orchestrating the Near-RT RIC and RAN nodes. The rApps provide RAN optimization, with the potential to extend to other RAN functions such as capacity planning or security. The rApps are used in conjunction with AI and ML models, leveraging data sets from other functions in the Open RAN and external sources. A secure, standardized R1 interface enables any rApp to work with other rApps, providing the flexibility to group rApps for more complex use cases and decisions.

The visibility and intelligence of the SMO, such as [33], make it an ideal platform to enhance the security posture of Open RAN cloud deployments, aligning with a ZTA. The SMO's logging capabilities coupled with its AI/ML can provide the awareness, threat intelligence, and automated responses needed for a secure Open RAN. The SMO's intelligence and its support for rApps, as shown in Figure 8, enable an ecosystem of purpose-built security functions providing faster and deeper threat detection, helping to ensure secure Open RAN public and hybrid cloud deployments. As the SMO has network-wide visibility from internal and external data sources, its rApps can be purpose-built to provide RAN-protecting security functions, such as Open RAN anomaly detection, O-Cloud threat detection and response, security configuration validation, and security compliance monitoring [34]. The SMO also provides the flexibility to build-in rApps with Security Information and Event Management (SIEM) and Security Orchestration Automation and Response (SOAR) functionality, plus the ability to integrate with external SOAR or SIEM in the security operations center (SOC).

However, securing the SMO is critical because a vulnerability in the SMO could be exploited to serve as an entry point for attacks against Open RAN and lateral movement across Open RAN interfaces and functions. While the SMO can enhance the O-RAN security posture, the SMO must have built-in security controls implemented with a zero trust mindset, assuming the adversary is already inside the network. It is critical that the SMO implements proper controls to ensure secure access with authentication and authorization of external and internal resources. The SMO must also provide security controls for confidentiality, integrity, and availability protection of SMO functions, interfaces, and data from internal and external threats.

5.4. O-Cloud

While cloud threats and cloud security controls are not exclusive to Open RAN deployments, critical infrastructure deployed in the cloud requires a higher level of due diligence and design with built-in security. O-RAN's O-Cloud inherits the threats and vulnerabilities inherent in the cloud where RAN will run on third-party hardware in a multi-tenant environment managed by a third-party. Recent real-world security events, including Solarwinds, Kaseya, Log4Shell, have demonstrated the potential risks of operating Open RAN as critical infrastructure in the cloud due to external and internal threats, including APTs which could exploit Open RAN vulnerabilities for lateral movement and reconnaissance.

O-Cloud is the cloud computing platform specified by the O-RAN Alliance to host O-RAN network functions, including Near-RT RIC, O-CU, and O-DU. The O-Cloud is a collection of physical infrastructure nodes supporting software components, such as operating system, container runtime, and management and orchestration functions. The O2 interface between the SMO and the O-Cloud provides platform resources and workload management of the cloud infrastructure for support of O-RAN network functions, including discover and administrate, create and delete, dynamic scaling, and Fault, Configuration, Accounting, Performance, and Security (FCAPS).

A risk-based analysis, considering a ZTA to protect against internal and external threats, should be performed to secure the O-Cloud and the O2 interface used to manage it. The O-RAN Alliance WG11, O-RAN Security, has identified O-Cloud threats across the five following threat categories [35]:

- Compromise of virtual network function or cloud-native function images and embedded secrets
- Weak orchestrator configurations, access controls and isolation that can be exploited
- Misuse of a virtual machine or container to attack another virtual machine/container, hypervisor/container engine, or other hosts via shared resources such as memory, network, or storage
- Spoofing and eavesdropping on network traffic to access all O-RAN data processed in the workload

- Compromise to supporting network services

These threat categories are consistent with the Cloud Security Alliance (CSA) identification of the current eleven most important threats to cloud deployments as of 2022 [36], including:

1. Insufficient Identity, Credential, Access and Key Management, Privileged Accounts
2. Insecure Interfaces and APIs
4. Lack of Cloud Security Architecture and Strategy
6. Unsecure Third-Party Resources
9. Misconfiguration and Exploitation of Serverless and Container Workloads
11. Cloud Storage Data Exfiltration

The O-RAN Alliance WG11, O-RAN Security, has a current work item to ensure the O-Cloud and the O2 interface will be securely specified to protect against internal and external threats.

6. Conclusion

5G migration to the cloud provides great opportunity to transition the RAN from proprietary hardware to open software while increasing vendor diversity. While Open RAN, including Cloud RAN and O-RAN, provides security advantages, it also expands the RAN attack surface introducing new security risks requiring a shift in security paradigms from a perimeter-based approach to a ZTA that protects against internal and external threats. The O-RAN Alliance's WG11 continues to evolve the security posture of O-RAN's Open Fronthaul interface, RICs, SMO, and O-Cloud to align with a ZTA. Existing security protocols, including mTLS 1.2 with X.509 certificates, CMPv2, and OAuth 2.0, are valuable tools to protect against external and internal threats. The SMO, along with security rApps, can further enhance the security posture of Open RAN deployments. A secure Open RAN can help fulfill the promise of 5G use cases and ensure secure deployment of 5G critical infrastructure in the cloud.

Abbreviations

3GPP	3 rd -Generation Partnership Project
AI	Artificial Intelligence
APT	Advanced Persistent Threats
CA	Certificate Authority
CMPv2	Certificate Management Protocol version 2
COTS	Commercial Off the Shelf
CU	Central Unit
DDoS	Distributed Denial of Service
DU	Distributed Unit
eCPRI	enhanced Common Public Radio Interface
FCAPS	Fault, Configuration, Accounting, Performance, and Security
FTPES	File Transfer Protocol Explicit Security
gNB	Next Generation Node B
HLS	Higher Layer Split
HSM	Hardware Security Module
IAM	Identity and Access Management
LLS	Lower Layer Split
MFA	Multi-Factor Authentication
MitM	Man-in-the-Middle
ML	Machine Learning
MNO	Mobile Network Operator
mTLS	Mutual Transport Layer Security
Near-RT RIC	Near-Real-Time RAN Intelligent Controller
Non-RT RIC	Non-Real-Time RAN Intelligent Controller
O-Cloud	Open Cloud
O-CU	Open-Central Unit
O-DU	Open-Distributed Unit
O-RU	Open-Radio Unit
OFH	Open Fronthaul
PKI	Public Key Infrastructure
PTP	Precision Time Protocol
RA	Registration Authority
RAN	Radio Access Network
RBAC	Role-Based Access Controls
RIC	RAN Intelligent Controller
SMO	Service Management and Orchestration
SSHv2	Secure Shell version 2
SUCI	Subscription Concealed Identifier
TA	Threat Analysis
TLS	Transport Layer Security
ZTA	Zero Trust Architecture

Bibliography & References

- [1] Ericsson Cloud RAN, <https://www.ericsson.com/en/ran/cloud>
- [2] O-RAN Architecture Description, v6.0, Technical Specification (TS), O-RAN Alliance, WG1, March 2022.
- [3] Security architecture and procedures for 5G System, Technical Specification (TS), 3GPP TS 33.501, Release 16.
- [4] Adapted from NIST Computer Security Resource Center, Glossary, [https://csrc.nist.gov/glossary/term/security_posture#:~:text=Definition\(s\)%3A,react%20as%20the%20situation%20changes](https://csrc.nist.gov/glossary/term/security_posture#:~:text=Definition(s)%3A,react%20as%20the%20situation%20changes).
- [5] ENSIA NIS Cooperative Group, May 11.
- [6] Security Considerations of Open RAN, S. Poretsky and J. Boswell, Ericsson, Aug 2020.
- [7] O-RAN Minimum Viable Plan and Acceleration towards Commercialization, White Paper, O-RAN Alliance, June 2021.
- [8] Open RAN Risk Analysis, Germany BSI, Federal office of Information Security, English Translation, February 2022.
- [9] Report on Open RAN Cybersecurity, ENISA NIS Cooperative Group, May 2022.
- [10] O-RAN Threat Modeling and Remediation Analysis, v3.0, O-RAN Alliance, WG11, March 2022.
- [11] Zero Trust Architecture, NIST SP 800-207, US NIST, Aug 2020.
- [12] Adapted from O-RAN architecture diagram from O-RAN Alliance, O-RAN Architecture Description, v6.0, O-RAN Alliance, WG1, March 2022.
- [13] Zero Trust Architecture, NIST SP 800-207, US NIST, Aug 2020.
- [14] Security Guidance for 5G Cloud Infrastructures, US DHS CISA, Oct 2021.
- [15] Security for 5G, 5G Americas, December 2021.
- [16] O-RAN Minimum Viable Plan and Acceleration towards Commercialization, White Paper, O-RAN Alliance, June 2021.
- [17] Report on the Cybersecurity of Open Radio Access Networks, ENISA NIS Cooperative Group, May 2022.
- [18] Security Guidance for 5G Cloud Infrastructures, US DHS CISA, Oct 2021.
- [19] “OpenRAN – 5G hacking just got a lot more interesting“, Karsten Nohl, MCH (May Contain Hackers), July 2022, <https://www.youtube.com/watch?v=LRQsFTmWa2w&t=13s>.
- [20] Security Considerations for Cloud RAN, S. Poretsky and J. Jardal, Ericsson, Sept 2021.

- [21] eCPRI Specification, v2.0, CPRI, May 2019, <http://www.cpri.info/index.html>.
- [22] CUS-Plane Specification, v9.0, Technical Specification (TS), O-RAN Alliance, WG4, March 2020.
- [23] M-Plane Specification, v9.0, Technical Specification (TS), O-RAN Alliance, WG4, March 2020.
- [24] IEEE Std 1588-2019 "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", November 2019.
- [25] M-Plane Specification, v9.0, Technical Specification (TS), O-RAN Alliance, WG4, March 2020.
- [26] <https://www.cisa.gov/uscert/ncas/current-activity/2021/08/30/cisa-adds-single-factor-authentication-list-bad-practices>, US DHS CISA, August 2021.
- [27] Security Guidance for 5G Cloud Infrastructures, US DHS CISA, Oct 2021.
- [28] M-Plane Specification, v9.0, Technical Specification (TS), O-RAN Alliance, March 2020.
- [29] CUS-Plane Specification, v9.0, Technical Specification (TS), O-RAN Alliance, WG4, March 2020.
- [30] IEEE Std 802.1X-2020 "IEEE Standard for Local and Metropolitan Area Networks--Port-Based Network Access Control", February 2020.
- [31] O-RAN Security Protocols Specifications, v3.0, O-RAN Alliance, WG11, March 2020
- [32] Adapted from O-RAN architecture diagram from O-RAN Alliance, O-RAN Architecture Description, v6.0, O-RAN Alliance, WG1, March 2022.
- [33] Ericsson Intelligent Automation Platform (EIAP), <https://www.ericsson.com/en/ran/intelligent-ran-automation/intelligent-automation-platform>.
- [34] Intelligent Security: Using the SMO to enhance the Open RAN Security Posture, S. Poretsky and J. Jardal, Ericsson, June 2022.
- [35] O-RAN Threat Modeling and Remediation Analysis, v3.0, O-RAN Alliance, WG11, March 2022.
- [36] Top Threats to Cloud Computing: The Pandemic 11, Cloud Security Alliance (CSA), June 2022.

Explainable AI for Data Clean Room Query Validation

A Technical Paper prepared for SCTE by

Srilal Weera PhD
Principal Engineer
Charter Communications
720-699-5079
srilal.weera@charter.com

Table of Contents

Title	Page Number
Table of Contents	2
1. Introduction.....	3
2. Explainable AI – Brief Overview	3
2.1. How it works	3
2.2. XAI algorithms	4
3. Data Clean Room	5
4. Query Validation	6
4.1. Query Relaxation Example	6
5. Machine Learning Module	8
5.1. Machine Learning Model Training Steps	8
5.2. Constrained Optimization	9
5.3. Model specific considerations	10
6. Solution Architecture	11
6.1. Workflow.....	11
6.2. Functional components	12
7. Application to Privacy Mechanisms.....	13
7.1. Query sensitivity score and Privacy	14
8. Conclusion.....	15
9. Abbreviations.....	15
10. Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1 - XAI Example - Titanic survival analysis.....	4
Figure 2 - Data Clean Room concept	5
Figure 3 - Query filter relaxation example	7
Figure 4 - ML classification of graded query sensitivity	8
Figure 5 - Incremental query relaxation	10
Figure 6 - Enhanced query validation workflow	11
Figure 7 - Solution components	12
Figure 8 - Query validator and privacy module	14
Figure 9 - Query sensitivity input to privacy module	14

List of Tables

Title	Page Number
Table 1 - Query relaxation Iteration	7
Table 2 - Sensitivity Analysis - Example	9

1. Introduction

Sophisticated ML models function more or less as black-boxes. A neural network may easily classify a photo of an animal as a cat or dog, but is silent about *why* it made that decision. A recent development is *Explainable AI* (XAI), also called *Interpretable AI*. Dubbed as an enabler for ‘third-wave of AI’, it helps open up the black-box model [1][2]. XAI has found niche applications in many industries. For example, in credit-risk analysis it is common practice to use machine learning models. If a loan application is denied then XAI can further reveal the reasons why it was deemed risky. Another scenario is in product recommendations. XAI could bring to light the contributing factors as to why a certain product was recommended to a specific customer.

In spite of its prowess, XAI applications in cable industry have been lacking thus far. In this paper, we present a timely application that reflects broad global interest in ways to share customer data in a privacy-compliant way.

An emerging solution is the Data Clean Room (DCR) concept [3]. Its goal is to provide a safe place for partnering companies to bring respective data for analysis in a secure manner. Guidelines are established to restrict any sensitive queries to protect the customer identity. However, sensitive querying may occur unintentionally due to micro-targeting. This is ascribed to how the queries are constructed (e.g. too many conditions in the SQL filter). Since the queries have originated from credible sources, blocking them entirely is not desirable. A pragmatic solution would be to assess and relax the query sensitivity which would lead to efficient database querying. This can be achieved with XAI enabled machine learning. Additionally, the query sensitivity scores can be used to fine-tune the privacy mechanisms. This is illustrated with reference to leading privacy technologies.

[Please Note: Charter has a longstanding commitment to protecting the privacy and security of its customers. For example, Charter provides customers with detailed information about its privacy practices, explicitly allows customers to opt out or change sharing preferences at any time, and restricts the collection of information (when enabled by the customer) to what is necessary to provide and optimize service. Learn more at Spectrum.com/Privacy.]

2. Explainable AI – Brief Overview

Explainable artificial intelligence (XAI) attempts to answer the ‘why?’ question about machine learning models. *Explainability* has gained much attention recently as its potential for trustworthy AI is recognized. A common algorithmic technique is to slightly change (perturb) a single feature at a time and measure the impact on the model prediction.

2.1. How it works

To illustrate how XAI works, consider the widely available Titanic dataset listing the survival rate of passengers. XAI analysis results are shown in Figure 1 (see reference [4]). The passenger survival rate is high on the right half of the diagram (positive values) and vice versa.

The XAI (SHAP) analysis reveals the survivors were either mostly female or paid higher fares or belonged to young age group. The ‘embarked destination’ seems to have little impact on survival rate.




Figure 1 - XAI Example - Titanic survival analysis

2.2. XAI algorithms

The well-known XAI algorithms are:

- Local Interpretable Model-agnostic Explanations (LIME)
- Shapley Additive explanations (SHAP)
- Partial Dependence Plot (PDP)

LIME and SHAP are called surrogate models because they attempt to approximate the predictions of the underlying black-box model. The algorithms tweak the input slightly and test the changes in prediction. If there is appreciable change in the predicted value, then that input variable is considered to have a higher impact on the model prediction. Other techniques include gradient-based saliency maps for image analysis. LIME generates a new dataset by randomly turning the data points (pixels/words) on or off. Hence it is a local approximation to the black-box ML model. The surrogate models are model agnostic since they treat the ML model as a black box.

SHAP has origins in game theory. Corresponding to a game, each feature is considered a ‘player’ and the prediction is the prize money. In a game, Shapley values determine how to assign payouts to players in proportion to their contributions to the prize money. In machine learning it is the contribution by each feature to the final model prediction. SHAP algorithm determines the average marginal contribution of a feature to the model prediction.

PDP provides a graphical representation of how each feature affects the prediction in a machine learning model. Each feature value is changed in ascending order and the corresponding change in the prediction is plotted. A partial plot depicts the dependency of the target response over the range of input features. One limitation is that PDP assumes no correlation between input variables, an assumption which is not always realistic.

3. Data Clean Room

Recent privacy regulations and app tracking transparency frameworks reflect a growing trend requiring explicit user consent for tracking. Some major companies are also ending the support of third-party cookies and identifiers, making it harder to run effective campaigns or measure attribution. The challenge for businesses is how to share consumer data for analytics without compromising consumer privacy.

An emerging solution is the Data Clean Room (DCR) model. Its goal is to provide a safe place for partnering companies to bring respective data for confidential analytics. Security and privacy-protection measures are applied, such as data anonymization, obfuscation and differential privacy. All data stays within the data clean room and is not shared with outsiders (Figure 2). Guidelines are established to restrict any sensitive queries to protect the identity of customers in the database [6]. One drawback, however, is that when a query is blocked the analyst (querier) has to reconfigure the filters in the query string and resubmit until the blocking is removed. Since the queries have originated from credible sources, blocking them entirely is not desirable.




Figure 2 - Data Clean Room concept

Another application scenario is programmatic ad-buying, in which the ad-spaces are bid in real-time on ‘ad exchanges’, (*cf.* stock exchanges in finance). The highest bidder wins the auction and places the ad alongside the main content. For more effective targeting, a matching audience need to be identified and is done via querying the publisher database. Multiple advertiser proxies are involved in this process and all such entities may query a publisher database to build profiles. Similar to DCR, a common practice is to block any queries that are deemed sensitive.

Secure multi-party computation (MPC) is a cryptographic technique that secures the privacy of data even in collaborative computing. Some DCRs advocate the use of MPC. Confidential Computing is a related paradigm to protect the data ‘in use’ (compared to data at-rest and data in-transit). Additionally, it may also involve the siloed enclaves with memory partitions and containerized abstractions.

4. Query Validation

Validating a database query prior to execution is done for several reasons:

1. To correct the syntax of a query
2. To prevent hacking attacks targeting security vulnerabilities (e.g. SQL injection)
3. To thwart privacy attacks designed to divulge sensitive data

The third item is the focus of this analysis. While the stored data remain anonymous, sensitive queries could divulge personal data and could pose a privacy risk. As such, a common practice is to block any queries that are perceived as sensitive. This practice however reduces the *utility value* of the database. Also, it disrupts certain business models where a multitude of queries are generated from credible sources.

One such scenario is DCR, where each party query the partner databases to build user profiles. Note that the querier in this scenario is not an adversary but a partnering company. As such, sensitive querying may occur unintentionally due to micro-targeting. This is ascribed to how the queries are constructed, such as having too many conditions in the defined filter. Simply blocking or invalidating such queries is costly and inefficient. It also increases traffic to the database due to repeated attempts of failed queries.

Accordingly, there is a need for an enhanced query validation method to:

1. Notify the querier why a query was deemed sensitive and suggest improvements
2. Determine the conditions to be relaxed for a failed query to be executed successfully
3. Assess the query sensitivity and use it to enhance privacy protection mechanisms

To accomplish this, it is necessary to first identify the factors which contributed to blocking the query from execution. In general, the purpose of a query is to retrieve data records that fulfill a specified criteria. It becomes a *sensitive query* if the query string is constructed in such a way that it could divulge privacy-sensitive data. For example, if a query returns only a handful of matching records, releasing such data may be a privacy risk for the affected individuals or the cohort. Also, in this scenario injecting statistical noise to mask the query response could skew the results. A more appropriate solution would be to devise a way to *relax* the query criteria to enable (unblock) query execution. This applies to other types of querying scenarios such as in graph databases.

In the ensuing sections, a machine learning-based solution is outlined for identifying sensitive queries. Explainable AI techniques are then used to obtain the reasons why the query was deemed sensitive. A distributed workflow is presented for relaxing the query criteria and deriving a quantified query sensitivity, which is then used as input for privacy settings.

4.1. Query Relaxation Example

Relaxing the search criteria constraints will unblock a sensitive query. For example, assume that in the first pass XAI indicated three possible reasons for the blockage.

Blocked by: Ethnic code = Danish, Age group = 50 – 55, Income bracket = 100k – 110k

These conditions are then relaxed iteratively and the query is run recursively. Table 1 shows the contributing factors for blockage at each pass.

Table 1 - Query relaxation Iteration

Query Iteration	ML Outcome	XAI Outcome (Reasons for blocking)
First Pass	query blocked	Ethnicity, Age, Income
Second Pass	query blocked	Ethnicity, Age,
Third Pass	query blocked	Age
Fourth Pass	query validated!	

For example, the age group can be expanded gradually by ± 5 years.




Figure 3 - Query filter relaxation example

Similarly, the ethnic code can be relaxed as in following steps:

- 1st try – Ethnic Code = Danish
- 2nd try – Ethnic Code = Danish + Swedish
- 3rd try – Ethnic Code = all Scandinavian countries
- 4th try – Ethnic Code = all Nordic countries
- 5th try – Ethnic Code = Nordic + Baltic countries

The above steps are continued till the blockage is cleared, or until the max depth is reached (as defined in the rules engine and database schema). By relaxing the conditions, a valid query can be achieved. This is vital to the querier as it avoids the trial and error method of submitting multiple queries and getting a “Your query was blocked” error message. Instead, the ML-XAI based solution provides a more productive response: “Your query was deemed sensitive and blocked. The following conditions may need to be relaxed...” Additionally, (if sanctioned by the contract), the query is run with relaxed conditions on the main database. The results (with noise injected), are supplied back to the querier.

5. Machine Learning Module

ML engines are generally trained on the data gathered externally, such as in transfer learning. In the present case, the training data resides in the main database itself. While database-integrated ML engines already exist (e.g. Amazon Aurora and Google BigQuery), the usage described below is different: In addition to query retrieval, its sensitivity is also analyzed. For this task, a collection of queries over a wide range of conditions is amassed and run on the main customer database. The outputs are then classified and graded according to sensitivity. This action is based on Rules Engine settings, additionally supported by human expert analysis. Database and privacy experts review the queries and classify those based on the perceived sensitivity. The outcome is a gradation based on the level of sensitivity. The neural ML engine is trained (weights adjustment), with the classified queries from the query database. Once trained, the ML engine is able to classify whether a fresh query is sensitive or not.

5.1. Machine Learning Model Training Steps

1. Raw queries are collected from multiple sources.
2. Query database is formed with ‘unclassified’ queries.
3. Queries are run on the main database (which contains user data records).
4. The query outputs are analyzed for sensitivity per rules engine settings.
5. Queries are then classified (graded), according to sensitivity.
6. The sensitive queries are submitted to the XAI module. Reasons for blocking are derived.
7. The ‘unclassified’ queries are updated with sensitivity data.
8. The ‘classified’ queries are used to train the ML engine.

The classified* query strings form the inputs for training the ML engine. Additionally, XAI layer provides the explanation why certain queries were considered sensitive. The latter is used for further analysis and refinement of the model. (*the term *classified* here means, ‘assigned to different classes’.)

Figure 4 illustrates functioning of the ML Engine for the case of a supervised neural network. The categories (fields) in the database form the input layer. The output layer contains the predicted sensitivity of the query. It could be a binary y/n type or a graded value as shown for more granular assessment.




Figure 4 - ML classification of graded query sensitivity

Referring to Table 2, queries form the leftmost column (usually thousands or millions of entries/rows). The next set of columns refer to the types of records in the database, typically ranging from hundreds to several thousand columns. Each row indicates a data record. Column heads are the fields in the database searched by the query. The Sensitivity columns on the right are derived during the training phase as described above. The ‘Y’s in the column indicate blocked queries that were perceived as privacy compromised. The middle column is a more granular representation of sensitivity. Instead of a binary (Y/N) outcome, the sensitivity is represented as a graded score based on the severity of privacy risk. The last column indicates the assessment from XAI for sensitive/blocked queries.

Table 2 - Sensitivity Analysis - Example

Queries	Types of Records in the Database					Sensitivity score (binary)	Sensitivity score (graded) 1- 5	Sensitivity analysis (XAI)
	Income (X1)	Age (X2)	Zip (X3)	Race (X4)	etc. etc.			
Query 1	100-110k	50 - 55	12345	Danish		Y	4	Reason
Query 2								
etc.								
etc.								

5.2. Constrained Optimization

In a typical database, there could be hundreds or even thousands of columns each representing a feature. Without the XAI analysis, it is not easy to pinpoint which feature contributed more for the query sensitivity. Formulating it as a multivariate discrete optimization problem:

If the input variables are denoted by X_i and the target variable (query sensitivity) as ‘y’.

$$X_1 \cap X_2 \cap X_3 \cap \dots X_n < y$$

The constraints for each variable: $a_1 < X_1 < b_1$, $a_2 < X_2 < b_2$, ... $a_n < X_n < b_n$

Then to relax the conditions, each X_i is incremented by $\pm \Delta X_i$.

$$(X_1 + \Delta X_1) \cap (X_2 + \Delta X_2) \cap \dots \cap (X_n + \Delta X_n) \geq y$$

The set of minimum ΔX_i values which meet the constraints would be the optimal result.




Figure 5 - Incremental query relaxation

For example, consider the simplified case of 2 variables.

Age constraint (X_1) $50 < \text{Age-group} < 55$

Income constraint (X_2): $100k < \text{Income-bracket} < 110k$

Assume the sensitivity boundary is ‘ $y=10$ ’. That is, any compound query that returns less than 10 records would be blocked as sensitive.

$$\text{Age-group} \cap \text{Income-bracket} < 10$$

As the field ranges are extended/relaxed, at one point the query will return 10 or more results. i.e. the query is no longer sensitive. It is necessary to find the lowest increments for each variable to meet that criterion. For example, simply expanding the Age-group broadly to 30-70 years might return a large number of records but would be irrelevant. While it may preserve the privacy (query no longer sensitive), the query result will have no utility value to a marketer who is targeting 50-55 age group.

Finding which variable(s) and by how much (ΔX) “minimally” to tweak each is an optimization problem. It gets complicated when dozens or hundreds of variables are involved.

5.3. Model specific considerations

The ML module is based on common classification algorithms such as neural and statistical models. During the training phase, queries with known outcomes (sensitive/not-sensitive) are used to update the weights in the ML engine. (Weights of each link determine the impact of a category). Once trained, the ML engine can classify whether a fresh query is sensitive or not. If the data is volatile, then frequent training would be needed. Statistical ensemble methods such random forests and XGBoost are also options. For more than two outcomes, multinomial logistic regression is another paradigm.

6. Solution Architecture

Process Steps:

1. The Query Controller submits the initial query to the (trained) ML engine.
2. If the query is deemed sensitive, it is run through the XAI module.
3. XAI supplies the reasons for blocking. The querier may be notified.
4. Sensitivity conditions in the string are relaxed per rules engine iteratively.
5. Query is submitted to ML engine recursively until it passes the sensitivity test.
6. The validated query is run on the main database.
7. Query response is assessed for privacy risks and noise injected.

6.1. Workflow

The workflow for the query validation solution are presented below.




Figure 6 - Enhanced query validation workflow

6.2. Functional components

Figure 7 describes the functional components of the distributed solution. The ML engine, XAI module and the rules engine work in conjunction to define, detect and quantify the query sensitivity. While functionally separate, the ML Engine and XAI are integrated components. XAI can be considered a functionality on top of the machine learning layer and is used recursively with the ML Engine. Once XAI is invoked, it parses the input query and supplies the reasoning for the ML classification.




Figure 7 - Solution components

Main Database (Customer Data) – A standard database, such as RDBMS or NoSQL type, either central or distributed. Queries are run on the database and the outputs are classified for sensitivity, then used to train the ML engine.

Query Supply Module – To populate the ‘Queries database’, queries are supplied in several ways.

- Historical data (a collection of previously run queries)
- Queries obtained from an external entity for training purposes. (Transfer learning is a well-known machine learning training paradigm, in which external data is used first to train an ML engine. Then local data is used to fine-tune the algorithm.)
- Boosting and resampling is a standard ML technique to generate data. Also, the use of a sub-module that auto-creates new queries by modifying the existing queries in the database. Suppose that an existing query has the qualifiers: area code, age group, income bracket, vehicle driven and hobby. Then it is possible to auto-change one qualifier at a time and test how the sensitivity changes. The query is run against the actual data in the main database to classify it as sensitive or not.

Queries Database – The queries database is populated by the Query Supply Module, mentioned above. These queries are used for training the machine learning engine.

Rules Engine – Rules are constructed to define instructions, thresholds etc.

Examples:

- A query is considered sensitive if the query response contains less than five records.
- A query is considered sensitive if a combination of certain categories are present in the query string.
- A query is considered sensitive if similar (coordinated) queries originated from the same source.

Query Controller (QC) – QC module interacts with others systems to automate process flows as well as perform decision making based on the Rules Engine settings. It consists of three subsystems described below.

Query Runner subsystem – Queries (retrieved from the ‘Queries database’) need to be run on the Main database to train the ML engine. First, each query result is classified as sensitive or not. This can be done based on the settings of the rules engine or with expert input, as it is done offline. During normal operation (after the ML engine is trained), Query Runner’s function is to run queries submitted by external entities (queriers).

Query Analyzer subsystem – Analyzing the query for sensitivity based on Rules Engine settings.

Query Modifier subsystem – A software construct with automated processing capability. Assume a query was deemed sensitive and the specific reason for blocking was supplied by XAI module. The Query Modifier will relax the constraint(s) (based on Rules Engine settings) and run the query iteratively until it passes the ‘sensitivity test’.

7. Application to Privacy Mechanisms

‘Differential Privacy’ (DP) is a statistical technique for protecting individual privacy during database querying. When a query is run on a database, DP adds a carefully chosen amount of noise/perturbation is added selectively to the result masking the user identity. For example if the query is to find those who subscribed to a certain TV channel, then some of the user responses (yes/no) are flipped randomly. This gives rise to a new concept called ‘*plausible deniability*’. So by looking at the data, it is not possible to establish if that data is truly associated with a person or randomly generated. The flip side of adding noise is the need to strike a balance between *utility* and *privacy*. One cannot be enhanced without compromising the other.

Algorithmic basis: Given two neighboring datasets D and D' differing by one data record, the randomized function K provides ϵ -differential privacy when the following probability condition (denoted by ‘Pr’) is satisfied for all $S \subseteq \text{Range}(K)$.

$$Pr[K(D) \in S] \leq \exp(\epsilon) \times Pr[K(D') \in S]$$

ϵ (epsilon) denotes the privacy loss. For small values, $\exp(\epsilon) \sim 1 + \epsilon$. Note that *differing by one data record* implies that it would not make a difference whether one individual’s data are included in the dataset or not. The query result would still be within the statistical error margin. The definition of sensitivity in DP is tied to this formalism.

Small ϵ (more noise) – better privacy but low utility

Large ϵ (less noise) – low privacy but high utility

Figure 8 shows the query validator functionality (in schematic form) as input to the DP privacy module. Statistical/random noise is added to the outgoing query response. The role of DP controller is to coordinate the epsilon values across a distributed implementation of differential privacy. Such a solution is warranted in the case of targeted advertising, to account for viewership and billing considerations. See reference [5] for implementation details.




Figure 8 - Query validator and privacy module

7.1. Query sensitivity score and Privacy

The leading privacy technologies contain parameter settings to control the level of protection. Examples are the (ϵ) epsilon value in differential privacy, 'k' and 'l' values in k-anonymity and l-diversity, respectively. Calibrating these parameters is a trial and error process. In this regard, the graded *query sensitivity score* described above can be used in tuning the privacy parameters as shown below.

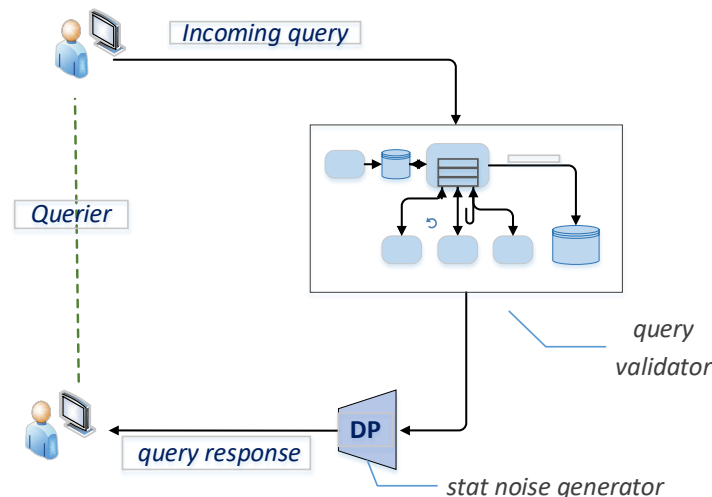


Figure 9 - Query sensitivity input to privacy module

In addition to structured databases, the solution described may also apply to search engines. Instead of a terse message (‘no results to display’), the user would appreciate receiving some form of approximate results. This can be achieved with ML based query relaxation.

8. Conclusion

A timely application of Explainable AI to the cable industry was presented, driven by recent privacy regulations. XAI-enabled machine learning leads to more efficient database querying in specific applications. The quantified query sensitivity scores can be used to enhance privacy mechanisms.

9. Abbreviations

DCR	Data Clean Room
DP	Differential Privacy
LIME	Local Interpretable Model-agnostic Explanations
MPC	Secure Multi-Party Computation
PDP	Partial Dependence Plot
SHAP	Shapley Additive explanations
XAI	Explainable AI

10. Bibliography & References

- [1] *Explainable Artificial Intelligence* (DARPA article) – <https://www.darpa.mil/program/explainable-artificial-intelligence>
- [2] *Explainability won't save AI*; Brookings Institute, (May 19 2021) – <https://www.brookings.edu/techstream/explainability-wont-save-ai/>
- [3] Examples are Google Ads Data Hub (ADH), Facebook Advanced Analytics (FAA), Amazon Marketing Cloud (AMC), Blockgraph DCR, Snowflake DCR, Habu and Decentriq DCR products.
- [4] A sample SHAP analysis for Titanic dataset (reproduced with permission) – <https://meichenlu.com/2018-11-10-SHAP-explainable-machine-learning/>
- [5] *Implementing Differential Privacy for Targeted Advertising*; SCTE Tech. Journal (Mar, 2022) – https://wagtail-prod-storage.s3.amazonaws.com/documents/SCTE_Technical_Journal_V2N1.pdf
- [6] Anonymized data is any information from which the person to whom the data relates cannot be identified, whether by the company processing the data or by any other person.

FMA Cloudification: Methods and Architecture Patterns

A Technical Paper prepared for SCTE by

Jim Huang

Principal Solutions Architect
Telecom Industry Business Unit, Amazon Web Services
2795 Augustine Drive, Santa Clara, CA 95054
jimhuan@amazon.com

Jeff Finkelstein

Chief Access Scientist
Cox Communications
Atlanta, GA
Jeff.Finkelstein@cox.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Flexible MAC architecture overview.....	4
3. Approach to FMA cloudification	6
4. Cloud infrastructure and services for FMA cloudification.....	8
4.1. Cloud edge services and selection	9
4.2. FMA cloud architecture	10
4.3. High Availability	11
4.4. Operation automation.....	13
5. Cloud FMA use case exercise	14
6. Conclusions.....	16
Acknowledgements	16
Abbreviations	16
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1 – FMA Details	5
Figure 2 – FMA Phase 1 Reference Architecture	5
Figure 3 – Cloud Continuum: Region, Edge, and Far Edge	6
Figure 4 – FMA System Architecture	7
Figure 5 – FMA Cloudification Architecture Pattern.....	7
Figure 6 – Types of Cloud Edge Services	9
Figure 7 – FMA on AWS Cloud.....	10
Figure 8 – HA Solution Patterns	11
Figure 9 – FMA Cloud Operation Automation.....	13
Figure 10 – MAC-NE Dial-out Streaming Telemetry Sequence Diagram ^[4]	14
Figure 11 – Streaming Telemetry in Cloud FMA	15

List of Tables

Title	Page Number
Table 1 – Analysis of FMA Function Timing Characteristics and Operation Locality	7
Table 2 – Comparisons of Cloud Edge Services	10

1. Introduction

Over the past decade, Cable operators and communication service providers have been striving towards Distributed Access Architecture (DAA)^[1], which disaggregates Converged Cable Access Platform (CCAP) systems for cable network scaling, operation simplification and cost reduction, as well as resource and space savings at the HFC headend or hub location. DAA deployment also provides a platform for cable operators to continually integrate virtualization elements into their networks for service velocity.

The Flexible MAC Architecture (FMA) specification^[2] released by CableLabs in 2020 is the latest undertaking for DAA evolution. FMA compartmentalizes DAA management plane, control plane, and data plane functions with APIs for vendor product interoperability. It accommodates DAA Remote PHY (R-PHY) and Remote MAC PHY (R-MACPHY) system architectures for flexible DAA solution variants. FMA access technologies include xPON and 4G/5G wireless for evolving last-mile access architectures. The FMA specification further calls out software defined networking (SDN) and network function virtualization (NFV) for future implementation with DAA.

Today, FMA is undergoing vendor product implementation and multi-vendor interoperability demonstrations. But what will FMA NFV look like end to end? How will the FMA NFV be implemented? Will FMA be further optimized for a cable operator's service agility and cost reductions? These questions are all open for exploration and development.

This paper proposes FMA cloudification to advance FMA along its evolutionary path. Many communication service providers (CSP) in the telecommunication industry have recognized Cloud as a mechanism for lowering operator's total cost of ownership (TCO) and gaining operations efficiency. Some CSPs have started migrating their network infrastructures as well as IT workloads to Cloud. Challenges in supporting virtualized network functions (VNFs) in Cloud are often around real-time constraints, high availability, and operation automation. This is particularly true for DOCSIS like products.^[3] Our approach to FMA cloudification is to properly map FMA functions across cloud region, cloud edge, and on-premise environments based on their latency tolerance characteristics and access network technology requirements. Specifically, we allocate FMA control plane functions such as MAC Manager and management plane functions such as Operations Support System (OSS) in the cloud region and control functions of auxiliary cores at the cloud edge, leaving latency-sensitive Remote MAC Devices as well as last-mile access networks AS IS in the field. With this MFA cloud architecture, CSPs can focus their capital and resources on the last-mile technology and deployment into a particular geographic area while simplifying network provisioning and management processes with the cloud.

In the following, we briefly review the FMA architecture in Section 2. Section 3 details our approach to FMA cloudification by introducing a Cloud Continuum model and positioning FMA functions along the cloud continuum based on timing analysis and FMA architecture. In section 4, we describe how the cloud FMA can be realized through a cloud infrastructure and address FMA high availability and deployment automation using cloud-native services. We further demonstrate the FMA cloudification with cloud-native services for an FMA use case – streaming telemetry in Section 5. We conclude the paper in section 6.

2. Flexible MAC architecture overview

The FMA concept began after a number of discussions among technical staff from Cox Communications and the AT&T Foundry as they worked on what became known as “CORD” (Central Office Reimagined as a Data center). FMA was originally called “HERD” (Head-End Reimagined as a Data center), but as it transformed from an architecture for simply managing remote MACPHY devices to an agile service delivery platform that included its original intent of Remote MACPHY Device (RMD) interoperability, the decision to rebrand it was made.

FMA was designed from the beginning with the idea of all components being able to work in a physical and virtual environment. The legacy back-end software, MAC manager, SDN and DOCSIS controllers, and auxiliary cores, would be designed to be placed at the point in the compute network that was consistent with the architectural demands of each individual operator. The interfaces between FMA components were standardized which would allow for the decomposition of functions and their strategic placement in the operator network or public cloud as desired.

The core objectives of the FMA working group are as follows:

1. Define interfaces between a management entity and DOCSIS MAC network element such that the management entity from Vendor A may be interoperable with a DOCSIS MAC network element from Vendor B.
2. Define a DOCSIS MAC network element that contains all necessary DOCSIS MAC layer functions such that a CCAP core network element is not needed for data plane forwarding of customer data traffic.
3. Define a DOCSIS MAC network element for both physical network functions (PNFs) and virtual network functions (VNFs) used for data plane forwarding.
4. Define an architecture such that typical CCAP functions such as management plane, L2/L3 control plane, and DOCSIS control plane may be separate from the DOCSIS MAC network element.
5. Define an open standard interface between operator BSS, OSS, NMS, and orchestration with the management entity or entities which may include management plane, L2/L3 control plane, and DOCSIS control plane used to operate the DOCSIS MAC network element.
6. Leverage previous cable industry specifications as appropriate as well as consider the transition to next generation approaches used in the areas of software defined networking (SDN), network function virtualization (NFV), VNFs, and PNFs, new protocols, new data models, and new telemetry methods as desired by the R-MACPHY MSO steering committee and working group.

Towards the objectives, the FMA working group developed and published the first FMA System Specification in 2020 and since has made two revisions.^[2] A detailed view along with the connections between the FMA components can be seen in Figure 1 below.

The FMA architecture will evolve in three phases.

Phase 1 – Support key components necessary to bring RMDs into existing operator networks without causing disruption to legacy EMS and NMS platforms

Phase 2 – Allow for continued use of Remote PHY devices in the FMA network management paradigm

The diagram illustrates a network architecture centered around a **Converged Interconnect Network**. On the left, a vertical stack of components includes:

- ERM Video Mgmt**
- OOB Video System**
- RF Systems**
- COPS / PCMM**
- Operator (CU)**
- IPDR Collector**
- TOD Server**
- SYSLOG Server**
- SNMP Server**
- DHCP Server**
- File Server**
- NETCONF/Yang Models** (dashed box)
- New Telemetry Data** (dashed box)
- Operator (RESTCONF)** (dashed box)
- Overlay/Underlay Controller** (dashed box)

These components interact with the central network via interfaces labeled **NETCONF,GCP,... New Interface**, **GCP/RCP**, and **R-DEPI/R-UEPI**.

The central **Converged Interconnect Network** connects several key nodes:

- RMD Manager**: Contains a **Management Plane** and an **RPD Management Agent**.
- Auxiliary Video/OOB Core(s)**: Contains an **RPD Management** agent and two planes: **Control Plane Video/OOB** and **Data Plane Video/OOB**.
- Remote MACPHY Device (RMD)**: Contains a **Management Agent**, an **RPD Management Agent**, and three planes: **Control Plane L2/L3**, **Data Plane L2/L3**, **Control Plane (DOCSIS)**, **Data Plane DOCSIS MAC**, and **Data Plane DOCSIS PHY**.
- Remote MAC Core (RMC)**: Contains a **Management Agent**, an **RPD Management** agent, and three planes: **Control Plane L2/L3**, **Data Plane L2/L3**, **Control Plane (DOCSIS)**, **Data Plane DOCSIS MAC**, and **Data Plane DOCSIS PHY**.
- Remote PHY Device (RPD)**: Contains an **RPD Management Agent** and two planes: **Data Plane DOCSIS PHY** and **Data Plane Video/OOB PHY**.

Arrows indicate the flow of management and data traffic between these components through the central network.

Figure 2 shows the FMA Phase 1 reference architecture with components and their functions and interfaces. For details, refer to the FMA System Specification^[2].



This paper introduces how to take Cloud advantages for business agility, operation efficiency, modern application services, and cost savings by running the components of the Flexible MAC Architecture in a cloud environment. The paper focuses on the FMA Phase 1 reference architecture for cloudification.

3. Approach to FMA cloudification

The FMA cloudification builds on a cloud infrastructure model, called Cloud Continuum. Typical cloud workloads such as web applications are centered around Cloud Region which consists of multiple inter-connected data centers. The Cloud Continuum extends the cloud region to cloud edge for time-sensitive network functions and applications. It enables multi-access edge computing (MEC) for network operators' business growth. The Cloud Continuum is further out to the far edge with cloud services like IoT for managing physical devices in the field.

Figure 3 shows the Cloud Continuum model with Cloud Regions, Cloud Edge, and Cloud Far Edge. They are inter-connected and have the same “look and feel” for cloud infrastructure services such as networking, compute, storage, CLI commands, and deployment automation.




Figure 3 – Cloud Continuum: Region, Edge, and Far Edge

With the Cloud Continuum model, FMA is overlayed across the cloud region, cloud edge, and cloud far edge in on-premises environments. This is accomplished through analyzing the real time characteristics and operation locality of FMA systems^[2] and accordingly allocating FMA functions along the Cloud Continuum and remote MAC network elements (MAC-NEs). Table 1 summarizes the analysis with respect to the degree of latency tolerance per major FMA functions, where High means the latency requirement is greater than 50ms roughly, Medium is in-between 20ms - 50ms, Low in-between 5ms - 20ms, and Ultra Low is less than 5ms. It also shows the operation locality of an FMA system or function, i.e., Backoffice, Headend, Hub, or MAC-NE.

Given the timing constraints and operation locality, the FMA functions are now allocated across cloud region, cloud edge, access network, and MAC-NEs. Figure 4 illustrates the FMA System Architecture^[2], where the upper portion shows the management and control plane functions allocated in Data Center, Headend, Hub, and Optical Node and the lower portion depicts the data plane.

The FMA System Architecture can be cloudified in a cloud architecture pattern shown in Figure 5. The FMA functional allocation and architecture cloudification scheme is as follows.

Table 1 – Analysis of FMA Function Timing Characteristics and Operation Locality

FMA System	FMA Functionality	Plane	Latency Tolerance*	Operation Locality
MSO Backoffice	OSS - IPDR Collector, Syslog Server, SNMP Manager, DHCP Server, File Server, AAA, PNM Server, Telemetry	Management	Medium - High	Backoffice
	PacketCable management - PCMM Policy Server, PacketCable CMS, PacketCable PKS	Management	High	Backoffice
	Auxiliary Core Management - OOB Video Mgmt, OOB RF Sys Mgmt, ERM Video Mgmt	Management	High	Backoffice
MAC Manager	Manage MAC-NE through FMA protocols	Management	High	Backoffice or Headend
PacketCable Aggregator	Scale PacketCable Multimedia Backoffice element to RMDs	Management	High	Backoffice or Headend
OOB Core	Manage and control legacy set-top boxes (STBs) as RMDs in substitution of SCTE 55-1 and SCTE 55-2 OOB Cores	Management Control	Medium Low	Headend or Hub
NDx Core	Support Narrowband Digital Forward (NDF) and Narrowband Digital Return (NDR) capabilities	Control Data	Low Low	Headend or Hub
Video Core	Provide video EQAM processing functions, except video PHY and QAM-related processing in RMD	Control Data	Low Low	Headend
PTP System	Clock synchronization across RMDs	Control	Low	Headend
MAC-NE (RMD)	DOCSIS MAC and the upper layer protocols, QAM, digital to/from analog conversion for RF/Ethernet or PON transmission	Management Control Data	High Low Ultra low	Hub or Node

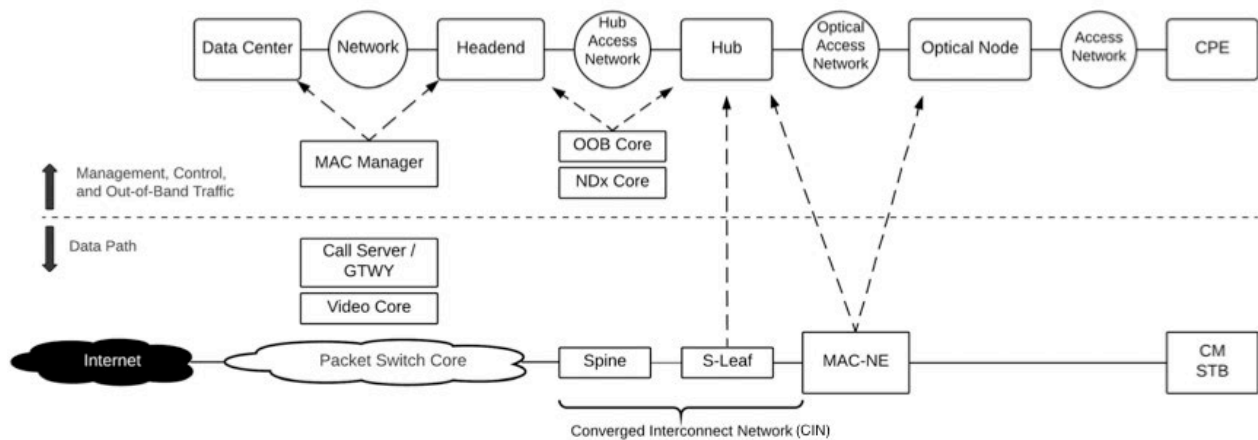


Figure 4 – FMA System Architecture

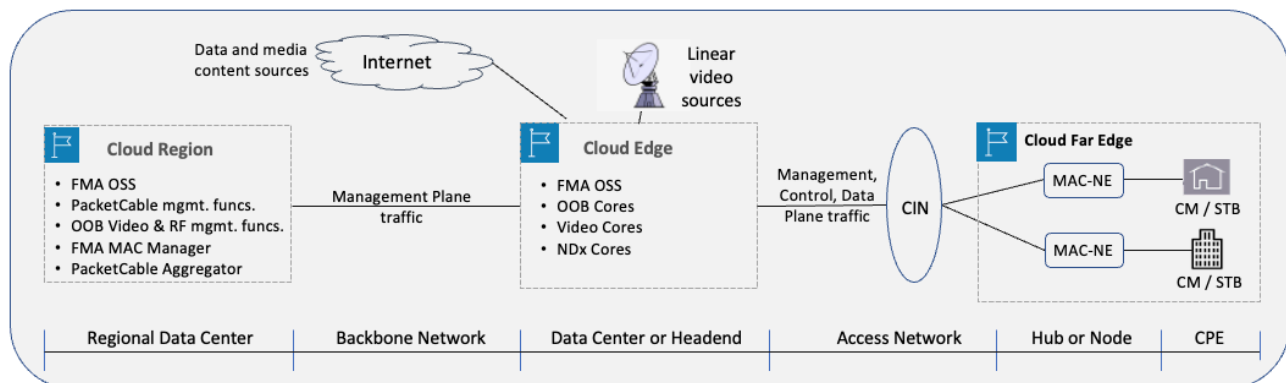


Figure 5 – FMA Cloudification Architecture Pattern

FMA Management Plane consists of FMA OSS functions, MAC Manager, PacketCable Aggregator, PacketCable management functions, and OOB Video & RF management functions. They are High in latency tolerance. Mirroring Multi-System Operator (MSO) Backoffice at regional data centers, MAC Manager, PacketCable Aggregator, most of OSS functions, and PacketCable management functions are placed in the cloud region. The OSS functions such as streaming telemetry that deal with a large amount of data can be allocated at the cloud edge to localize data processing, filtering, and analytics. MAC Manager and PacketCable Aggregator can run at the cloud edge as well because of the operation locality.

FMA Control Plane involves auxiliary Cores, OOB Core, NDX Core, and Video Core, under the FMA Phase 1 Reference Architecture shown in Figure 2. Since they are Low in latency tolerance, they run at the cloud edge. The FMA control plane functional entity resides in MAC-NE for processing L2/L3 control plane traffic. For example, the MAC-NE DOCSIS Control Plane functional entity processes control plane traffic from the PacketCable Aggregator. Hence, the FMA control plane functional entity in MAC-NE is allocated at the cloud far edge.

FMA Data Plane is comprised of data plane functions of MAC-NE (Remote MACPHY Device) as well as NDx Core and Video Core systems. As physical device, MAC-NE performs RF conversion with Ultra Low data process latency outside of cable plant. In addition, MAC-NE may run IoT functions for Proactive Network Maintenance (PNM). Hence, MAC-NE is at the cloud far edge. NDx Core and Video Core data plane functions are Low in latency tolerance and thus placed at the cloud edge.

FMA Networks are comprised of Packet Switched Network Core and Converged Interconnect Network (CIN). The Network Core is for communications between FMA data centers and headend systems as well as Internet traffic. CIN consists of Spine switches and Secure Leaf switches (S-Leaf), connecting MAC-NE (RMD) instances in a hierarchical structure. In the cloud environment, the infrastructure backbone network and SD-WAN in-between the cloud region and cloud edge replace the FMA Network Core, thus relieving FMA operators from the core network management. CIN remains as it is in the field for network technology as well as vendor deployment flexibility.

FMA MAC-NE (RMD) and CPE are vendor and customer choices. Cloud IoT can be added to RMD and CPE for MSO to provide advanced services such as device control, fault diagnostics, or preventive network maintenance which the FMA System Specification^[2] describes.

In summary, the FMA cloudification builds on the cloud continuum infrastructure with FMA functions allocated in the cloud region, cloud edge, and cloud far edge according to their timing characteristics and operation needs.

4. Cloud infrastructure and services for FMA cloudification

This section describes how the FMA cloudification can be realized in cloud. It uses Amazon Web Services (AWS) as an example of a cloud infrastructure. Please note that similar architectures can be built with other cloud providers.

4.1. Cloud edge services and selection

The cloud continuum is comprised of different types of cloud edge technology. AWS provides three types of cloud edge services: Outposts, Wavelength, and Local Zones as shown in Figure 6. Though all provide cloud compute, storage, and networking services for real-time, short-latency or high-throughput virtualized network functions (VNFs) and applications, they are architected for different use cases. MSOs or cable operators need to examine and select a right type of cloud edge technology and services for FMA cloudification.

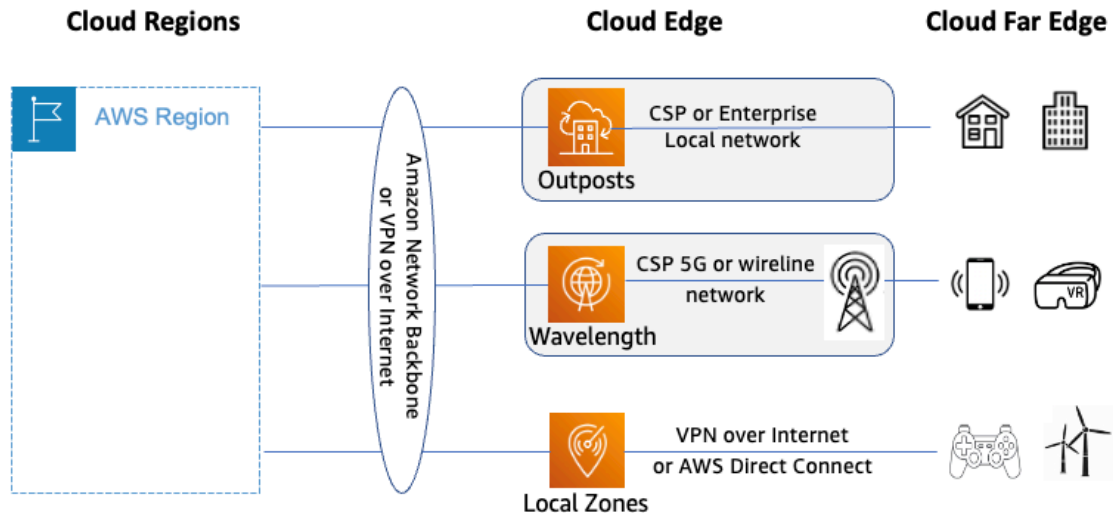


Figure 6 – Types of Cloud Edge Services

Table 2 summarizes the characteristics and use cases of these three types of cloud edge services. Although both AWS Outposts and Wavelength are embedded in CSP networks, the former is dedicated to an operator or enterprise, thus called “private edge cloud,” whereas the latter is shared by different enterprises, organizations, or users, called “public edge cloud.” That is, the Outposts is for private use and the Wavelength is for public. AWS Local Zones is fully managed by AWS with end-user connectivity through CSP wireline or wireless services or enterprise’s direct connection. Hence, the private edge cloud service (Outposts) is suitable for FMA in MSO or cable operator networks. MSO deploys and manages FMA systems and functions in the cloud.

Table 2 – Comparisons of Cloud Edge Services

Type of Cloud Edge Service	Cloud Edge Architecture & Deployment location	Use Case	Type of CSP	Management Responsibility
Outposts	On-premises networks or in CSP networks	<ul style="list-style-type: none"> Private Multi-Access Computing (e.g., smart factory, healthcare operations, real-time ML and analytics) Private Mobile Network MSO network 	<ul style="list-style-type: none"> Telecom operators MSOs 	<ul style="list-style-type: none"> Cloud infra. & cloud services: AWS Local network: CSP or on-premises operator
Wavelength	In CSP networks, especially 5G mobile networks	<ul style="list-style-type: none"> Public Multi-Access Computing (e.g., intelligent vehicles, real-time retails) 	<ul style="list-style-type: none"> Mobile network operators 	<ul style="list-style-type: none"> Cloud infra. & cloud services: AWS Local network: CSP
Local Zones	In AWS infrastructure	Public Multi-Access Computing (e.g., real-time gaming, interactive live video streams)	N/A	<ul style="list-style-type: none"> Cloud infra. & cloud services: AWS Local Zones connectivity: CSP or enterprise

4.2. FMA cloud architecture

Figure 7 depicts an FMA cloud architecture with AWS Region and Outposts-based cloud edge integrated with CIN and MAC-NEs at the far edge.

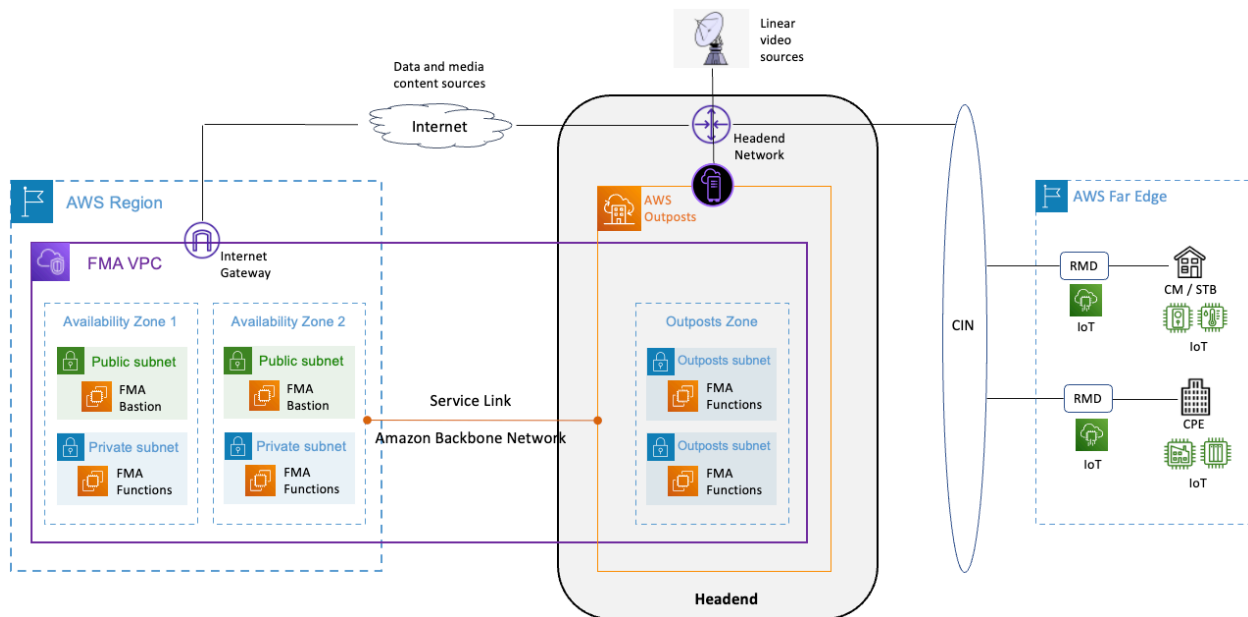


Figure 7 – FMA on AWS Cloud

The FMA cloud portion is comprised of four building blocks:

- *AWS Region* – is a physical location, which is comprised of multiple Availability Zones for high availability and high scalability. Each Availability Zone has multiple datacenters housed in physically separated facilities with redundant power, networking, and connectivity. A region hosts the high latency-tolerance FMA management plane functions such as OSS and PacketCable management.

- *AWS Edge Outposts* – are cloud edge platforms in three form factors, 42-RU, 2-RU, and 1-RU. The platforms provide cloud-native services such as Elastic Cloud Computing Service (EC2) for virtual machines and Elastic Kubernetes Service (EKS) for containers. The Outposts run FMA functions that are low latency tolerance (e.g., auxiliary Cores), require intense local data processing and analytics (e.g., OSS telemetry), or need to be operated at headend (e.g., MAC Manager and PacketCable Aggregator).
- *AWS Virtual Private Cloud (VPC)* – closely resembles a traditional network in a data center, with the benefits of using the scalable cloud infrastructure. For the FMA cloudification, a VPC is extended from a region to one or more edge locations where headend resides. Interconnecting the region and the edge is an AWS Service Link over the Amazon backbone network.
- *AWS Cloud Edge Network* – provides L2 and L3 connectivity with the headend network. It serves FMA traffic between AWS Outposts and three places: CIN for RMD, Internet for data, voice, and media content, and local equipment for linear video.
- *AWS Cloud Far Edge* – provides AWS IoT services on RMD and CPE with device connectivity over CIN to the cloud edge and cloud region. Since the subject of applications such as IoT is out of the FMA scope, the cloud far edge is not discussed below.

This FMA cloud architecture serves the traditional MSO Backoffice, headend, and their Packet Switch Network Core as a whole and leaves to MSO operators the flexibility and choice for CIN connectivity and the last-mile access network technologies (e.g., DOCSIS Remote PHY, DOCSIS Remote-MACPHY, EPON, GPON). The managed AWS infrastructure services (e.g., VPC, Outposts, EC2, and EKS) enable MSO operators to focus on the FMA functionality for technology innovation and business growth.

4.3. High Availability

High Availability (HA) is essential in cable services. Along the cloud continuum, HA can be achieved for FMA systems and functions using cloud infrastructure services. Figure 8 illustrates multiple HA solution patterns in the cloud region and at the edge with AWS as an example.

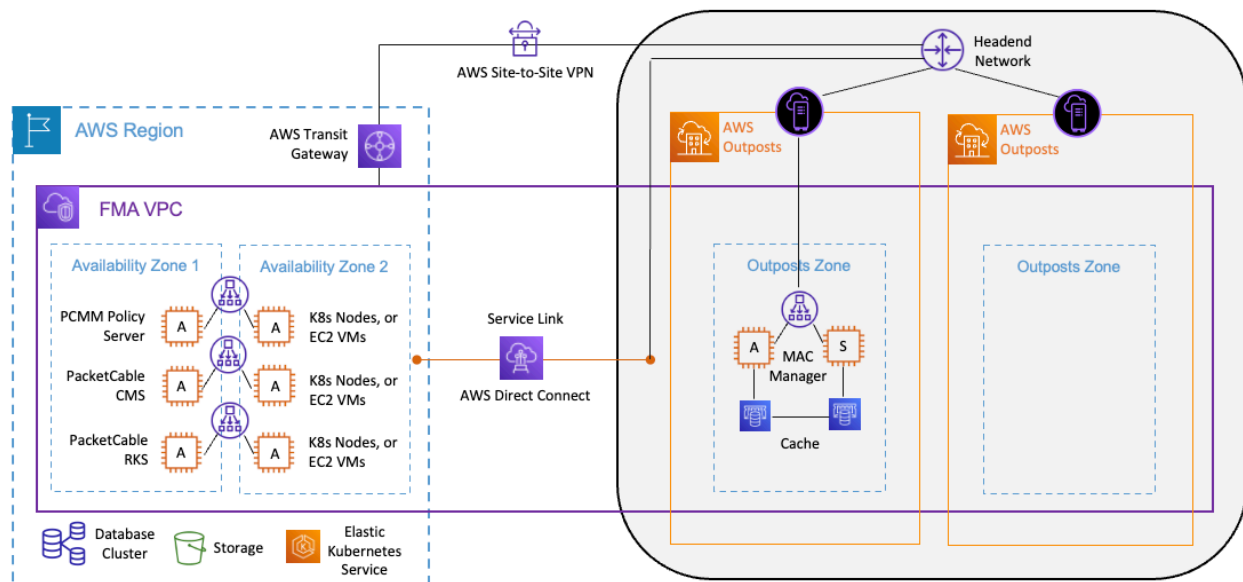


Figure 8 – HA Solution Patterns

Active-Active server HA solution pattern is based on a cloud service mechanism which runs two or more Kubernetes (K8s) nodes with containers in a K8s cluster, or two or more Virtual Machines (VM) instances, across two or more Availability Zones (AZ). The containers or VM instances, are load-balanced through Load Balancers. It is illustrated with FMA PacketCable management functions running in the Active-Active mode in the AWS Region in Figure 5. Failure of an entire Availability Zone, a VM instance, a K8s node, or physical server hosting the VM instance or K8s node will lead to re-routing the traffic to the active resources. Combined with the AWS Auto Scaling service, the Active-Active HA solution delivers the required resource capacity by automatically replenishing lost resources (VM instances or K8s nodes) after the failover.

Active-Standby server HA solution pattern uses two sets of resources (VM instances or K8s nodes) in two AZs in an AWS region or on an Outposts platform, with one set in the Active mode and the other Standby. When the Active fails, the Standby takes over. To simplify the Active-Standby HA solution, the application running on the VM instances or K8s nodes should be stateless with its state data stored in a cache or storage. Figure 5 illustrates the Active-Standby HA model in an AWS Outposts at the edge hosting MAC Manager. For higher availability, two Outposts platforms can be used and the applications are configured for failover over local network paths.

Network HA is supported for the Service Link between the cloud region and the edge with an encrypted set of VPN connections through AWS Direct Connect-based private or/and public connectivity or/and Internet-based public connectivity. In addition, the FMA may add a redundant out-of-band AWS Site-to-Site VPN link over Internet or AWS Direct Connect in-between the cloud region and the edge.

Database, cache, storage, and file system HA are provided as cloud-native services managed by AWS. For example, Amazon Relational Database Service (Amazon RDS) supports High Availability with one standby or two read standbys. AWS ElastiCache provides Redis replication groups within an AZ or across multiple AZs with automatic failover. Amazon Simple Storage Service (S3), an object storage, provides HA and durability through across-AZ replication; S3 can be enabled for cross-region replication as well. Amazon FSx and Amazon EFS file systems can also be deployed across multi-AZs for HA. The FMA cloudification will benefit from these and many other cloud-native database, cache, and storage services.

4.4. Operation automation

The FMA cloudification can utilize cloud services to ease FMA system deployment, provisioning, upgrade, and event monitoring and handling through automation, thus reducing MSO OpEx and in turn TCO. Figure 9 illustrates a procedure for FMA cloud operation automation. It represents an FMA cloud life cycle with AWS services as example.

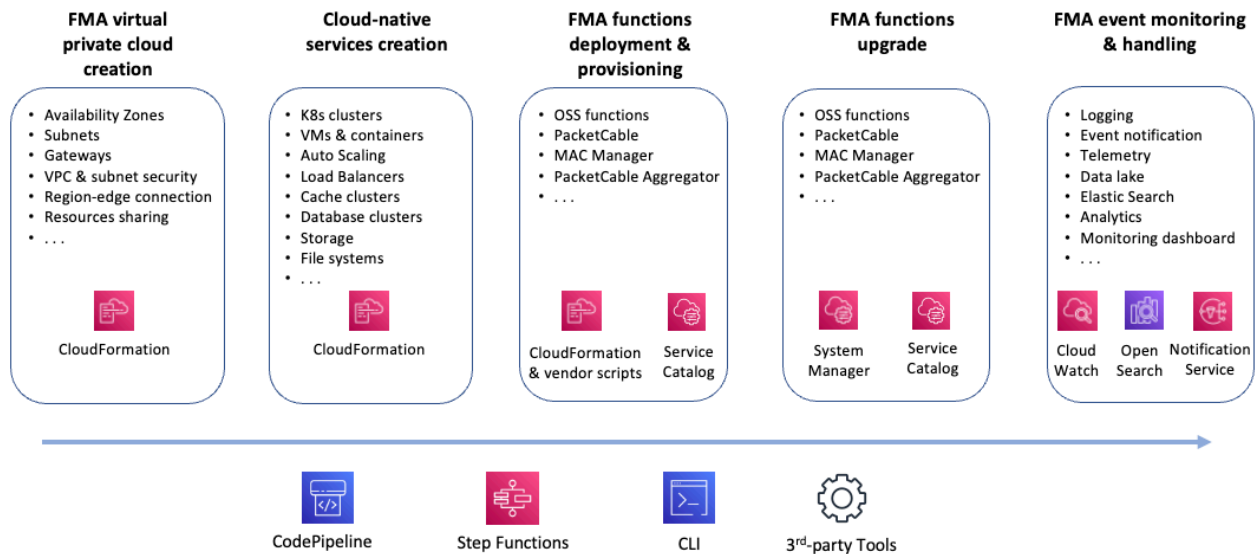


Figure 9 – FMA Cloud Operation Automation

- *FMA AWS VPC creation* – sets the network and security environment for the cloud continuum automatically by AWS CloudFormation templates.
- *AWS cloud-native services creation* – establishes AWS services used by the FMA cloud automatically by AWS CloudFormation templates.
- *FMA functions deployment and provisioning* – are performed in combination of AWS CloudFormation, container services, 3rd-party tools, or scripts with vendor container or VM images. AWS Service Catalog is used to manage products deployed in AWS.
- *FMA functions upgrade* – is part of product life cycle in the FMA cloud. Among different tools, AWS System Manager is a cloud-native service for product update. AWS Service Catalog is used to track the product upgrade.
- *FMA event monitoring & handling* – can be supported by AWS cloud-native services such as AWS CloudWatch, OpenSearch, Analytics, and Simple Notification Service (SNS). These services can implement or integrate some of FMA OSS functions.

The FMA cloud deployment, provision, and upgrade procedures can be driven by AWS CodePipeline and Step Functions with other 3rd-party operation automation tools.

The above content uses AWS cloud infrastructure and services as an FMA cloudification example. Similar solutions can be built with other cloud providers.

5. Cloud FMA use case exercise

The previous sections establish a cloud FMA reference architecture with solution building blocks and provide methods for FMA high availability and deployment automation in the cloud. This section demonstrates how an FMA function – Streaming Telemetry can be implemented across MAC-NE and CIN on premises and MAC Manager and Network Support Services (NSS). It uses AWS cloud infrastructure and services as an implementation example. Similar implementation can be built with other cloud providers.

The streaming telemetry is a PUSH-based mechanism to transport monitored network status data from network elements to external data collectors. Instead of the traditional data PULL model, FMA as well as CCAP utilize the PUSH model to stream data to backoffice applications in near real time^{[4][5]}. Figure 10 shows the MAC-NE dial-out streaming telemetry sequence diagram in the FMA OSS Interface Specification^[4].

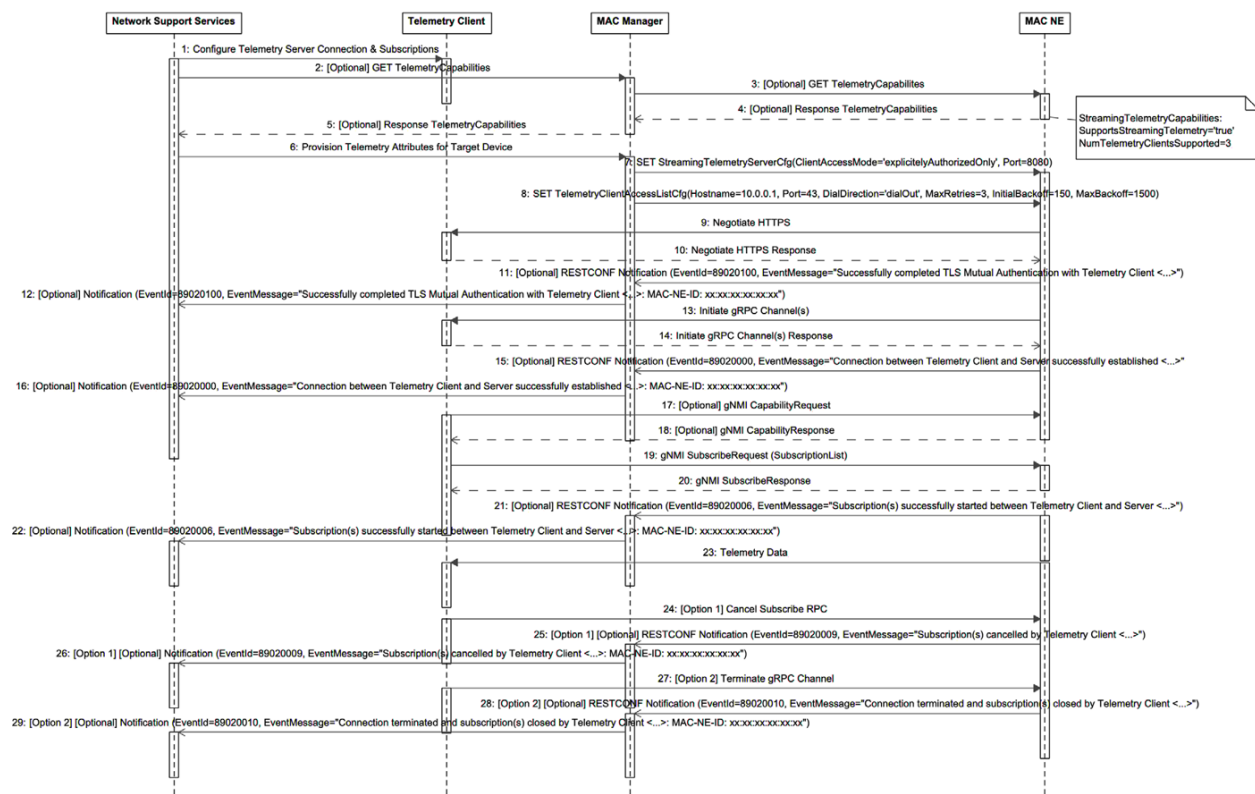


Figure 10 – MAC-NE Dial-out Streaming Telemetry Sequence Diagram^[4]

Figure 11 depicts the streaming telemetry cloudification with cloud-native services on AWS as an example. An AWS Region hosts a Cable operator's backoffice, where the NSS, MAC Manager, and Telemetry Clients are deployed in the FMA VPC subnets. These FMA components can run as VMs on Amazon Elastic Compute Cloud (EC2) or as containers on EC2 orchestrated by either Amazon Elastic Container Service (ECS) or Amazon Elastic Kubernetes Service (EKS). According to the dial-out streaming telemetry protocol, the components communicate with each other through the subnet traffic routing within the VPC. The telemetry

servers in MAC-NEs are connected to the telemetry clients and NSS in the backoffice in the cloud region via Amazon Direct Connect over AWS network backbone. AWS Site-to-Site VPN over Internet is a backup link for connection high availability. Note that the telemetry clients can be implemented at the cloud edge as well if preprocessing, filtering, or analysis of telemetry data should take place locally first.

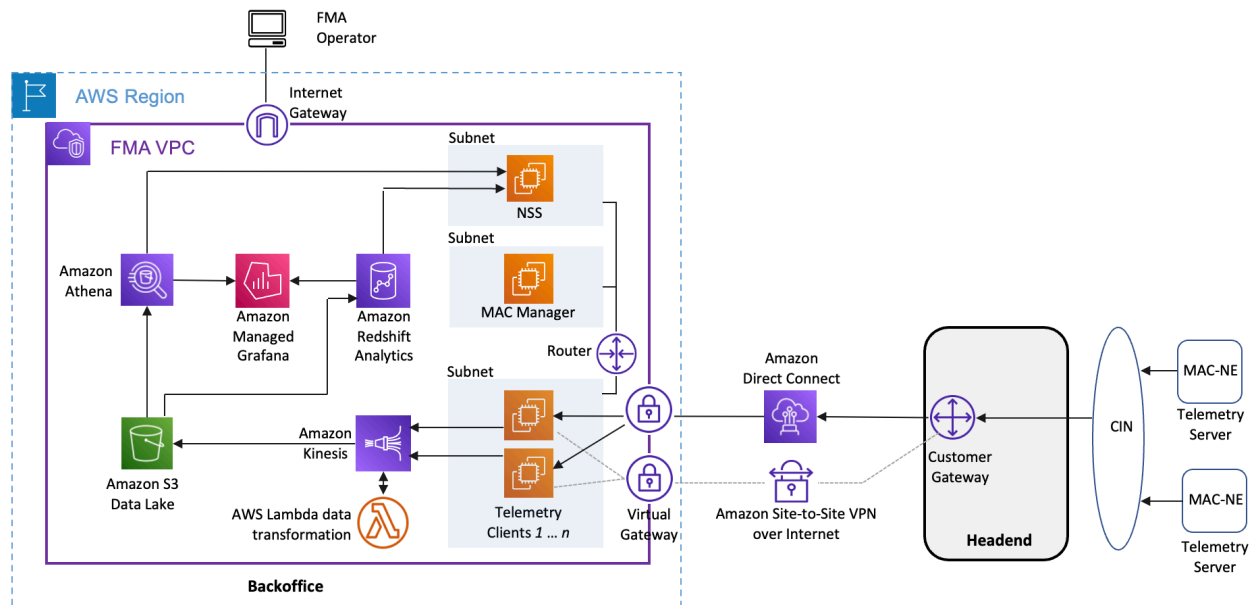


Figure 11 – Streaming Telemetry in Cloud FMA

In addition to the streaming telemetry implementation, the cloud provides several salient services for advanced OSS operations as illustrated in Figure 11.

- *Data lake* – is nowadays a common service provided by cloud service providers to store, process, and secure large amounts of structured and unstructured data. For example, AWS Simple Cloud Storage (S3) is a durable and scalable data storage to host a variety of OSS data for recording, troubleshooting, and predictive equipment maintenance. The telemetry clients further stream the received data to S3 using Amazon Kinesis Data Firehose service.
- *Analytics* – is used by advanced OSS operations for network event correlation, root cause analysis, and outage prevention. Many types of analytics services are provided by cloud service providers. As an example, Amazon Redshift is a fully-managed data warehouse service that can analyze petabytes of telemetry and other OSS data efficiently. Coupled with Amazon Machine Learning service, Amazon Redshift ML makes it easy for data analysts to create, train, and apply machine learning models.
- *Dashboard* – is an essential function of network observability. Grafana is a widely used open-source observability tool. For instance, with Amazon Managed Grafana in the cloud, cable operators can analyze and visualize telemetry and OSS metrics, logs, and traces and configure alerts for OSS event notifications. The Grafana service uses Amazon Athena as data source to access the telemetry and other OSS data stored in the data lake S3.

Note that there are other methods or cloud services for implementing the Streaming Telemetry function for cable operator's data networks, e.g., [6], which are not out of the scope of this paper.

6. Conclusions

The telecom industry has started its journey from network function virtualization to cloudification in order to lower the total cost of ownership, increase operation efficiency, leverage modern application services, and achieve business agility. FMA cloudification is the next undertaking by communication service providers and multi-access operators to further modernize the Distributed Access Architecture.

This paper shows that the FMA can be cloudified along the cloud continuum across the cloud region, edge, and far edge with FMA systems and functions allocated based on their timing characteristics and operation locality. At the same time, the FMA cloudification enables the operators to retain the flexibility of access network technology choices and operations.

It is shown through the cloud infrastructure and services that the private edge cloud is more suitable for cable operator's latency-sensitive FMA functions than other types of cloud edge technology. The virtual private cloud extends from a cloud region ("FMA regional datacenter") to a private cloud edge ("FMA headend"), simplifying FMA OSS operations. Traditional cable HA capabilities can be achieved by the cloud infrastructure and services, including network redundancy, multiple availability zones, load balancing across containers and virtual machines, and database, cache, storage, and file system redundancies. Cloud-native services can not only implement FMA functions but also bring them to modern OSS operations as illustrated through the streaming telemetry use case.

It is time to implement FMA in the cloud by applying the FMA cloudification approach and leveraging the cloud infrastructure and services.

Acknowledgements

The authors would like to thank Dr. Jennifer Andreoli-Fang for her feedback on the writing style for SCTE Cable-Tec Expo and her valuable comments on the paper content.

Abbreviations

CCAP	Converged Cable Access Platform
CIN	Converged Interconnect Network
EMS	Element Management System
FMA	Flexible MAC Architecture
gNMI	gRPC Network Management Interface
MAC-NE	MAC Network Element
MSO	Multi Service Operator
NSS	Network Support Services
RMD	Remote MACPHY Device
VPC	Virtual Private Cloud

Bibliography & References

- [1] Web publication, *Distributed Access Architecture*, CableLabs,
<https://www.cablelabs.com/technologies/distributed-access-architecture>.
- [2] CM-SP-FMA-SYS, *Flexible MAC Architecture System Specification*, CableLabs, January, 2022, <https://www.cablelabs.com/specifications/CM-SP-FMA-SYS>.
- [3] John Chapman and Tong Liu, *Unleash the Power of Cloud Computing for CMTS*, SCTE 2021, Atlanta, GA, October, 2021. <https://www.nctatechnicalpapers.com/Paper/2021/2021-unleash-the-power-of-cloud-computing-for-cmts>.
- [4] CM-SP-FMA-OSSI-I02-220602, *FMA OSS Interface Specification*, CableLabs, June, 2022, <https://www.cablelabs.com/specifications/CM-SP-FMA-OSSI>.
- [5] CM-SP-CCAP-OSSI, DOCSIS 4.0, *CCAP Operations Support System Interface Specification*, CableLabs, June, 2022, <https://www.cablelabs.com/specifications/CM-SP-CCAP-OSSIv4.0>.
- [6] R. Harlin, A. Moustafa, and A. Kalawat, *How Comcast uses AWS to rapidly store and analyze large-scale telemetry data*, August, 2021, <https://aws.amazon.com/blogs/big-data/how-comcast-uses-aws-to-rapidly-store-and-analyze-large-scale-telemetry-data/>.

From Millions to Billions: SCTE Standards Evolve the Smart Grid at Scale

A Technical Paper prepared for SCTE by

Robert Cruickshank
GRIDIoT Power Networks
132 Cruickshank Rd #269, Big Indian, NY 12410
rfciii@cruickshank.org

Derek DiGiacomo
Society of Cable Telecommunications Engineers
140 Philips Rd, Exton, PA 19341
ddigiacomo@scte.org

David Fellows
Fellows Associates
PO Box 76, Bondurant, WY 82922-0076
dmfellows@live.com

Francis Sandoval
Pajarito Technologies
1650 Gaylord St, Denver, CO 80206
francisrsandoval@gmail.com

Tony Werner
Comcast Cable
1701 John F Kennedy Blvd, Philadelphia PA 19103
tony_werner@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Why Evolve The Grid	3
3. Sustainability Financials.....	7
4. Architectural History of Cable Networks and the Grid.....	8
5. The Symbiotic Future	10
6. What Broadband Offers The Grid	11
6.1. ANSI/SCTE 267 Optimum Load Shaping for Electric Vehicle and Battery Charging	11
6.2. ANSI/SCTE 271 Power Sensing in Cable and Utility Networks.....	16
7. Conclusion and Next Steps.....	18
Abbreviations	19
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1 – U.S. Summer 2021-2022 Average Wholesale Electricity Prices.....	4
Figure 2 – MISO summer 2022 Reliability Projections.	5
Figure 3 – 2021 Estimated U.S. Energy Consumption.	6
Figure 4 – U.S. Energy-related CO ₂ Emissions.	7
Figure 5 –U.S. Annual Energy Expenditures.	8
Figure 6 – Similarities in 1-way to 2-way upgrades of Cable and the Grid.....	9
Figure 7 – Optimum load shaping system.	12
Figure 8 – ERCOT Hourly Generation Based on Actual and Daily Optimum Load.....	13
Figure 9 – Electric Vehicle Charging Adapter.....	14
Figure 10 – Hourly Optimum: 1) Load Shape, 2) Charge/Discharge Shape.	15
Figure 11 – A “Local” Load Shape to Mitigate Grid Congestion.	16
Figure 12 – Time Scales for Electric Grid Monitoring and Control.....	17
Figure 13 – Analog Waveform, Utility Capture and Broadband Capture.....	18

1. Introduction

New global initiatives support the electrification of the planet, and the demands for power will continue to rise and stress the already-strained global electric power grid. Electric power directly affects costs for cable broadband, many industries and consumers, and is a root cause of soaring inflation. In response to increasing costs and power outages, federal and state legislatures and regulators, utilities, and energy managers are rapidly reorganizing energy mixes, focusing on the resiliency, sustainability and affordability of electricity generation, transmission, distribution, and storage.

Because utilities are increasingly dependent on the Internet to manage the delivery of electricity, cable broadband providers have a unique opportunity to provide innovative new services to the electric grid and commercial microgrids. These new services can save billions of dollars in fuels and other costs. In Texas, generation fuel savings of \$1 billion dollars a year are possible by time-shifting demand for electricity which would also extend the life and carrying capacity of the end-to-end, generation-to-load power grid.

As societies depend more on non-carbon-based electricity and less on fossil fuels, new broadband standards create valuable opportunities to reduce unnecessary power generation, delivery costs, and outages. Two new standards by American National Standards Institute (ANSI) and the Society of Cable Telecommunications Engineers (SCTE) could ensure seamless integration of broadband-enabled electricity services by all types of users across the globe thereby allowing for rapid, efficient, and effective adoption in transportation, buildings, industry, and agriculture.

The two new standards leverage the decades of investment in broadband networks to manage and monitor the grid and enable its transformation more efficiently. The ANSI/SCTE 267 *grid management* standard improves the efficiency and capacity of generation, transmission, distribution, and storage of electricity by orchestrating electricity demand relative to optimized supply. The ability to continuously manage demand at scale is needed to mitigate the increasing grid operational challenges of distributed energy resources such as fixed and mobile batteries and other flexible electric loads, especially as dispatchable thermal generation is retired and replaced with variable renewable energy. In addition, the ANSI/SCTE 271 *grid monitoring* standard provides a quantum leap in detecting, predicting, and proactively addressing conditions relevant to distribution grid faults, safety, reliability, congestion, and the hosting of renewables and electric vehicles.

Broadband providers may partner with utilities to distribute ANSI/SCTE 267 signals to electricity consumers to optimize electrical load on the grid and may deploy ANSI/SCTE 271 sensors throughout the outside plant to provide utilities with extremely fine-grained telemetry on grid performance. These control and measurement tools can profoundly influence the reliability and affordability of power used by both cable companies themselves and their customers.

2. Why Evolve The Grid

To understand cable broadband's role in services and standards, we first examine the evolving grid. A confluence of factors creates unprecedented challenges in grid operations and business models. One factor is the rising costs of electricity in global markets. For example, at the time of drafting this paper, the Northeast U.S. wholesale prices are forecast to exceed \$100 per megawatt hour (MWh) between June and August 2022, up from an average of about \$50/MWh last summer. Figure 1 shows forecast increases in wholesale electricity prices across the U.S. from the Summer of 2021 to 2022.¹

¹ U.S. Energy Information Agency, Short-Term Energy Outlook (6/16/22), *EIA expects significant increases in wholesale electricity prices this summer.* <https://www.eia.gov/todayinenergy/detail.php?id=52798>

Summer average wholesale electricity prices at selected price hubs (Jun–Aug, 2021–2022)
dollars per megawatt hour

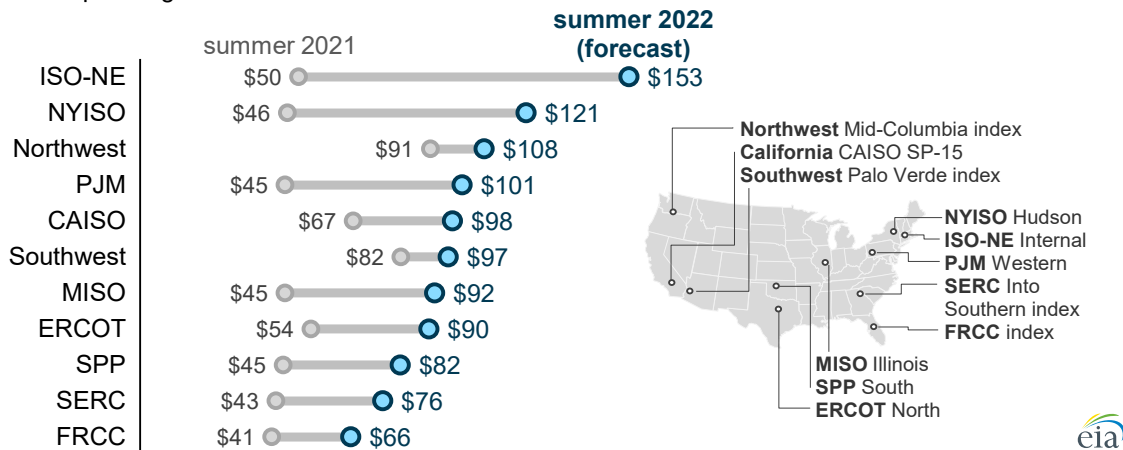


Figure 1 – U.S. Summer 2021-2022 Average Wholesale Electricity Prices.

While there are various reasons for rising wholesale electricity prices, the cost of generator fuel is a primary driver. Across the U.S., the price of natural gas delivered to electric generators is expected to average \$8.81/Million British thermal units (MMBtu) this summer, up 125% from \$3.93/MMBtu last summer.² Price increases in Europe and elsewhere are expected to be much higher.

In the past, when natural gas prices have risen, power providers with natural gas plants have substituted coal-fired generation. More recently however, many coal power plants are less likely to be used because of continued coal capacity retirements³ and lower-than-average stocks at coal plants.⁴ Other industry conditions that can contribute to higher wholesale electricity prices include fuel and water scarcities. For example, restricted contribution of hydropower this summer will likely lead the State of California to generate more electricity from natural gas and to import electricity from neighboring states.

Another factor creating unprecedented challenges in operating the grid is increased outages resulting from declining reliability of near end-of-life grid components, infrastructure frailty in severe storms, and lack of anticipated generation resources. Higher electricity demand coupled with potential supply reductions are raising concerns around the world. For example, the U.S. Midcontinent System Operator (MISO) has predicted outages this summer based on the anticipated shortfall of supply resources to meet normal and extreme demand shown in Figure 2.⁵

² Note the difference in abbreviating a million-watt hours (MWh) which is based on the International System of Units (SI) versus a million British thermal units (MMBtu) which is based on the Imperial System of Units.

³ U.S. Energy Information Agency, Short-Term Energy Outlook (1/1/22), *Coal will account for 85% of U.S. electric generating capacity retirements in 2022*. <https://www.eia.gov/todayinenergy/detail.php?id=50838>

⁴ U.S. Energy Information Agency, Short-Term Energy Outlook (12/7/21), *In September, the U.S. was at its lowest coal stockpiles since 1978*. <https://www.eia.gov/todayinenergy/detail.php?id=50558>

⁵ CleanTechnica (6/5/22), *Potential Electricity Reliability Concern for Central U.S.A.*, <https://cleantechnica.com/2022/06/05/potential-electricity-reliability-concern-for-central-u-s-a/>, U.S. Energy Information Agency, *Today in Energy* (6/3/22), <https://www.eia.gov/todayinenergy/detail.php?id=52618>, NERC (5/22), *2022 Summer Reliability Assessment*, https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_SRA_2022.pdf, MISO (4/28/22), *MISO projects risk of insufficient firm generation resources to cover peak load in summer months*. <https://www.misoenergy.org/about/media-center/miso-projects-risk-of-insufficient-firm-generation-resources-to-cover-peak-load-in-summer-months/>

Midcontinent Independent System Operator (MISO) summer reliability projections (2022)
gigawatts (GW)




Figure 2 – MISO summer 2022 Reliability Projections.

To ensure reliability, MISO and other “balancing authorities” plan to always have more supply available than demand. As shown in Figure 2, in MISO’s summer 2022 reliability projections, of the 143.2 GW of anticipated resources (at left), a reduced capacity of only 114.9 GW of generation may be available to meet between 118.2 GW and 125.2 GW of demand (at right). A shortfall will result in rotating blackout outages, most likely during a heat wave, which will likely cause loss of life and property. To anticipate electricity demand, balancing authorities produce a range of forecasts for average demand and extreme environmental conditions that used to occur only once in 10, 50, or 100 years, but are occurring more frequently. Planned and unplanned maintenance (aka forced outages) of power plants reduces available capacity as does derating generation capacity for factors such as drought, low-wind conditions, or fuel supply limitations .

Yet another factor creating unprecedented challenges in operating the grid are low and declining end-to-end, generation to load efficiencies. Efficiencies are so low that most of the fuel energy used in generating electricity, propelling transportation, and powering buildings and industry is rejected as waste heat as shown in Figure 3.⁶

⁶ Lawrence Livermore National Laboratory (6/16/22), *Energy, Water, and Carbon Informatics*, https://flowcharts.llnl.gov/sites/flowcharts/files/2022-04/Energy_2021_United-States_0.png




Figure 3 – 2021 Estimated U.S. Energy Consumption.

Figure 3 depicts the sources of energy (at left), how they are used, and how much is rejected (at right). Line widths are proportional to quantities of energy flows and all numbers can be roughly interpreted as a percent of total U.S. energy in each flow or sub flow. Complex interrelationships are depicted, such as the fine orange line (at lower right) denoting 0.02% of all energy consumed in the U.S. was used to charge electric vehicles.

In Figure 3, the most important takeaways are the costly inefficiencies (in light gray) that result in unnecessarily rejected energy as waste heat and greenhouse gases from: 1) electricity generation, where almost 2/3 of energy is rejected, 2) transportation, where almost 4/5 of energy is rejected, and 3) industry, where more than 1/2 of energy is rejected. Comparing the totals in light and dark gray at right, the U.S. (like most other developed nations) wastes about 2/3 of all energy used, primarily due to poor inefficiencies in electricity generation and transportation.

An alarming result of not significantly improving energy efficiencies over the last several decades is the emission of carbon continues to be colossal and will not be reduced until we have non-carbon-based electrons propelling transportation and replacing direct fossil fuel use (e.g., propane, heating oil, gasoline) in other applications.

Despite renewables and natural gas being added to the electricity generation mix and the retiring of coal powerplants, U.S. energy-related CO₂ emissions from all energy uses (electricity generation, transportation, and buildings) has not declined below 1975 levels as shown in Figure 4.⁷

⁷ U.S. Energy Information Agency, Today in Energy (5/13/22), *U.S. energy-related CO₂ emissions rose 6% in 2021*. <https://www.eia.gov/todayinenergy/detail.php?id=52380>




Figure 4 – U.S. Energy-related CO₂ Emissions.

In the U.S., the transportation and electric power sectors are the greatest contributors to energy-related CO₂ emissions, each responsible for roughly 1/3 of all greenhouse gas (GHG) emissions. In 2021, transportation emissions rose due to increased petroleum consumption as COVID-19-related restrictions eased. Likewise, electric power sector emissions rose due to increased electricity generation and the use of higher carbon intensity coal-based generation. In the U.S., electric power sector emissions from coal increased for the first time since 2014 (a global trend that is expected to continue for several years given natural gas supply constraints).

Exacerbating all the issues depicted in Figures 1-4 are declining energy efficiencies due to increases in extreme temperatures. On the supply side, as outdoor temperature rises, power plants, transformers, and powerlines become less efficient at moving electrons and rejecting heat. To make matters worse, on the demand side, air conditioners, the largest component of summer demand, also become less efficient as outdoor temperatures rise, causing air conditioners (AC) use more energy to keep buildings cool. This creates a death spiral: The hotter it gets, the more societies suffer by unsustainably declining efficiency, as more and more cooling energy is needed from less and less efficient power plants which are unable to meet demands for electricity. Thankfully, legislatures, regulators and grid operators are making efforts to deploy renewables and storage at utility-scale, community-scale, and premises-scale to help address spikes on the hot days due to AC loads and on cold days due to heating loads.⁸

3. Sustainability Financials

Inefficiencies undermine the sustainability of current grid operations and many other energy uses and business models, especially given the rising costs of energy. In 2019, U.S energy expenditures, the amount of money spent by consumers to purchase energy, was \$1.2 trillion.⁹ Considering more than just electricity, a disaggregation of energy expenses for all types of energy used over the last 50 years is shown in Figure 5.

⁸ Julie McNamara (7/9/2019), *How do power grids beat the summer heat*, The Equation, Union of Concerned Scientists, <https://blog.ucsusa.org/julie-mcnamara/how-do-power-grids-beat-the-summer-heat>

⁹ U.S. Energy Information Agency, Today in Energy (9/9/21), *In 2019, U.S. inflation-adjusted energy expenditures fell 5%*. <https://www.eia.gov/todayinenergy/detail.php?id=49476>

U.S. energy expenditures by source (1970–2019)
trillion real 2019 U.S. dollars




Figure 5 – U.S. Annual Energy Expenditures.

Referencing Figure 5, it is important to consider two significant yet uncharted imminent changes in trajectories: 1) cost increases attributable to the 2022 war in Ukraine have already doubled and could triple the costs of energy, and 2) the process of so-called beneficial electrification, e.g., the introduction of electric vehicles, will double or triple societies' consumption of electricity while decreasing the uses of petroleum, natural gas and coal over coming decades.¹⁰ The product of tripling electricity costs and tripling electricity consumption could yield a near 10x increase in residential, commercial, industrial and agricultural electricity bills, and could be unsustainable in terms of costs and GHG emissions.

4. Architectural History of Cable Networks and the Grid

Today's transformative path of the electric grid is similar to that of the historical evolution of the cable broadband network. Cable networks initially used centralized headends, proprietary systems, and one-way delivery of content. Over time, two-way upgrades and distributed headends and hubs enabled the development and deployment of new services, including telephony and high-speed data. As two-way services became highly penetrated and successful, the threat of network congestion and slowdowns gave rise to traffic engineering and the development of the massively scalable world-standard cable modem.¹¹ Tools to detect and mitigate network congestion were needed, were developed quickly, and continue to evolve.

¹⁰ Beneficial electrification (aka clean electrification, strategic electrification) is a term for replacing direct fossil fuel use (e.g., propane, heating oil, gasoline) with electricity in a way that reduces overall emissions and energy costs. There are many opportunities across the residential and commercial sectors. This can include switching to an electric vehicle or an electric heating system – but only if the end-user and the environment both benefit. Environmental and Energy Studies Institute (Jun 2022), *Beneficial Electrification, An Access Clean Energy Savings Program*. [https://www.eesi.org/electrification/be#:~:text=Beneficial%20electrification%20\(or%20strategic%20electrification,t he%20residential%20and%20commercial%20sectors](https://www.eesi.org/electrification/be#:~:text=Beneficial%20electrification%20(or%20strategic%20electrification,t he%20residential%20and%20commercial%20sectors) .

¹¹ Cable Television Laboratories, Inc. (CableLabs) led the effort to develop the world standard cable modem and certify interoperability. Consumers continue to benefit in the global telecommunications marketplace where vendors compete on price, functionality, and delivery schedules. An example of a powerful standard, interoperable cable modems are a thousand times faster than the proprietary cable modems they replaced and million times faster than dial-up modems.

As depicted in Figure 6, Many striking similarities arise when comparing the challenges of developing and deploying global broadband networks to the existing and anticipated challenges of developing standards-based scalable interoperable smart power grids.

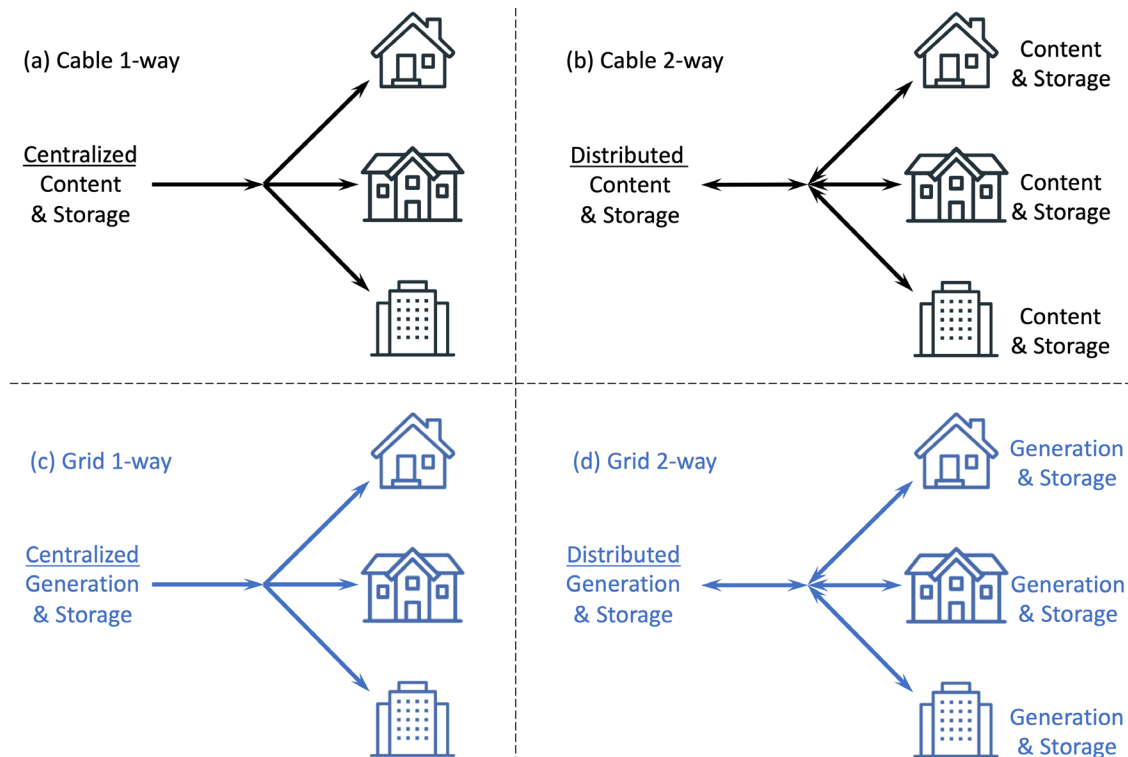


Figure 6 – Similarities in 1-way to 2-way upgrades of Cable and the Grid.

The top row of Figure 6 depicts (in black) the evolution of the cable distribution network from (a) 1-way to (b) 2-way. The bottom row depicts (in blue) a similar transition for the grid from (c) to (d). Cable's experience in transitioning to 2-way, identifying and managing congestion, can be helpful to the grid.

In the grid, the traditional model of central station generation is increasingly augmented by distributed energy resources (DERs) such as wind power, solar power, and other renewable energy technologies. DERs encompass more than electrical generators and include internet-connected batteries and electric loads that are smart in meeting user demands while time-shifting usage to operate now or later to reduce stress on the grid. DERs can choose to operate now, if there is abundant clean energy and no grid congestion issues or can operate later if there will be a more significant benefit to the grid. Some DERs, such as increasingly popular solar plus storage, create two-way flows of electrons that are managed by modulating (aka time-shifting or shaping) demands for electricity.

The process of influencing the timing of when electricity is delivered to the grid by DERs – or consumed from the grid by flexible loads such as batteries – is generally called Demand Side Management (DSM). DSM is facilitated by load shaping information provided by a power supplier in the form of a Demand Response (DR) signal.

Like in the early days of broadband, today's grid is operated with a mix of standards-based and proprietary solutions. Well-established standards include lamp sockets, wall plugs, voltages, and frequencies. Newer standards include managing DER behaviors in specific applications, such as solar

inverters, vehicle chargers and building controls.¹² The most significant challenges with DERs are orchestrating and aggregating their behavior to 1) maximize the resiliency, sustainability, and affordability of the grid, and 2) detect and mitigate congested electron flows that create heat, that shortens the lifespan of grid components and impacts outages and possible loss of property, life, or both.

5. The Symbiotic Future

As consumer-based generation “prosumer” trends increase, community-augmented electricity systems will likely become commonplace for providing resilient, sustainable, and affordable energy mixes. Community-level DER management will benefit from applying standards that rapidly modernize grid technologies to create efficient energy management orchestration synergies.

Modern telecommunications and broadband internet services work smoothly because each industry has adopted open, non-proprietary interoperability standards. Likewise, looking ahead for the next 10-20 years, the grid will perform more smoothly by favoring standards-based DSM solutions over propriety DSM solutions wherever possible. Demand response (DR) programs look to incentivize customers with the ability to inject power, add or shed load at a predictive time to reduce strains on the power grid. Demand response is a grid balancing option with limited deployment and represents untapped potential that has been modeled but not proven at scale. With a total available market of billions of DR devices in the U.S., and only 10 – 15 million deployed over the last 30-40 years, there is tremendous upside for global deployments of standards-based DERs to provide DR. To that end, the broadband industry has work to do in educating utilities, legislators, regulators, and power industry vendors about new scalable standards that accelerate the development of a sustainable and resilient smart grid.

Standards for *managing* demand and *monitoring* the power grid can be used to provide traffic engineering solutions that work remarkably well. Standards-based networked DERs can speed the deployment of renewables, batteries, and smart loads, guiding them to work in unison to shape load to raise the efficiency of the end-to-end generation, transmission, and distribution (G&T&D) infrastructure. By responding to changes in forecast daily grid needs and contingencies such as failures in G&T&D, networked DERs can manage load continuously to avoid costly peak generation and to move as many electrons as possible off-peak through would-be congested areas.

Combining grid monitoring sensor data with easy access to smart meter electricity interval usage data for consumers and authorized third parties will ensure that customer responses to DR signals are recognized and compensated. Compensation will be issued for DERs such as EVs, residential, commercial, industrial, and agricultural uses that follow load shaping signals. Load shaping will be based on forecasts of load and demand as well as observations from broadband-based sensors that monitor and quantify the impact of power anomalies on the grid. Anomalies are measured in terms of power quality issues such as oscillations and spikes in voltages, sustained over/under voltages, overloaded/overheated grid components, and other distribution system behaviors, faults, and their associated risks.

¹² Much work has gone into developing interconnection and interoperability standards that define how to assemble, configure, and connect devices to form a smart modern grid. Nonetheless, a major challenge with standards to date is that they are limited to development within specific technology domains, such as vehicle charging or air conditioning. As such, existing standards lack a complete vision for how distributed energy resources need to interact with G&T&D to deliver resilient and sustainable end-to-end grid services. While some standards are mature, technically robust, and meet the needs within a specific domain, in aggregate, they do not support optimum load shaping. The need for the ANSI/SCTE load shaping standard became apparent when assisting in the review of the U.S. Department of Energy, Grid Modernization Laboratory Consortium (Mar 2021), *Survey of Distributed Energy Resource Interconnection and Interoperability Standards*, <https://www.nrel.gov/docs/fy21osti/77497.pdf>

Combined analyses of customer meter data and grid sensor data will provide complementary insight into DER operations and their impact on the grid. In the near term, sensing and analysis will improve grid situational awareness and help determine and inventory the types, sizes, and locations of DERs operating on the grid today. Furthermore, sensing and analysis will identify normal and anomalous grid and DER operating states and behaviors. In the long term, sensing and analysis will be critical for asset and load capability forecasting. Having this more significant insight into grid operating states will ultimately enhance the availability and reliability of the broadband network.

6. What Broadband Offers The Grid

In 2021 the Society of Cable Telecommunications Engineers (SCTE) completed two synergistic American National Standards designed to reduce grid failures, contain electricity costs, and monetize demand response. The standards specify improved grid sensors and orchestration of electric loads via a simple information model. Separately and together, the standards improve electricity G&T&D, storage, and end-use while accelerating the adoption of electric vehicles (EVs) and resilient local power using batteries and other networked distributed energy resources.

6.1. ANSI/SCTE 267 Optimum Load Shaping for Electric Vehicle and Battery Charging

The SCTE 267 standard specifies end-to-end control of the electric power grid and commercial microgrids from generation to load, using one-way broadcast and two-way interactive signals. The standard defines how to create, transmit, and act upon a forecast optimum load shape (OLS) to manage the charging of EVs and facility batteries, as well as demands for electricity from flexible and discretionary smart electric loads.

An OLS provides grid control with a set of numbers, such as the target load for hours 1-24. The numbers in an OLS can, for example, forecast the cleanest, most efficient, and least cost electrical supply, so that all stakeholders: G&T&D entities, retail electricity providers, and consumers benefit.

Several topics are addressed in SCTE 267: 1) A generation-to-load OLS architecture is specified. 2) Based on inputs of forecast load and forecast generation from renewables, a method for producing a location-specific OLS is specified. 3) A method for managing the charging of electric vehicles is specified as an example of how any smart load can autonomously interpret and take local actions based on an OLS.

The OLS standard was created because existing siloed standards do not provide sufficient control to benefit the different needs of G&T&D. With OLS, stakeholders, including broadband providers, can reduce their electricity costs and carbon emissions by having their smart loads follow the lowest cost forms of supply. If not implemented, broadband providers will have less control over the rising cost of electric power.

Short-term benefits include creating and distributing both near and far-reaching OLS signals quickly and easily that allow intelligent devices that implement the standard to participate by shaping load. Benefits accrue in the short and long-term as more smart devices implement the standard resulting in more significant benefits for most stakeholders in the electricity value chain. The potential impact on the broadband industry is a reducing energy procurement costs and creating new revenue generating units based on managing the charging of cable customers' vehicles and batteries and flexible uses of electricity.

An OLS signal is created, transmitted, and acted upon, as shown in Figure 7. For example, an OLS Producer (typically an entity associated with an electricity supply or control system) ingests forecasts of

load, renewable generation, and costs. The Producer then uses 1-way and 2-way networks to distribute OLS signals to OLS Consumers (devices that manage the consumption of electricity).

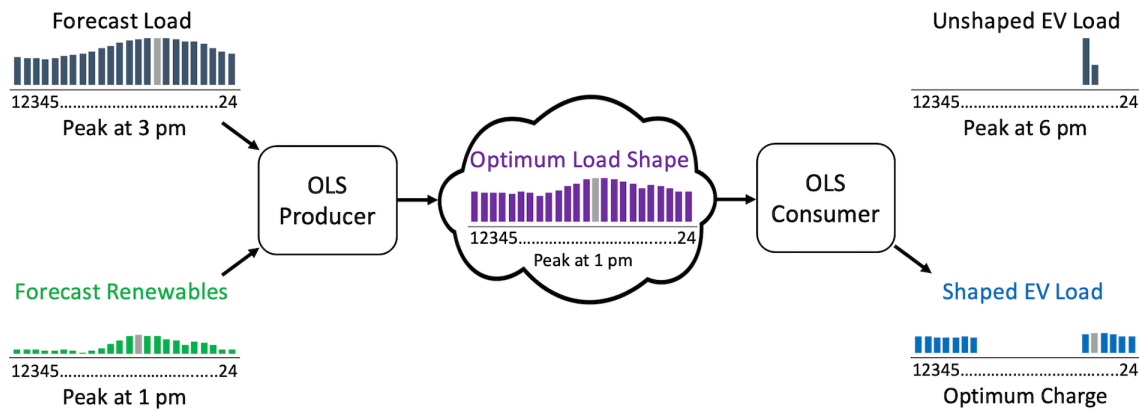


Figure 7 – Optimum load shaping system.¹³

An OLS Producer may employ various techniques to create an OLS signal. An algorithm may minimize costs and optimize the utilization of renewable generation. As depicted in Figure 7, the OLS Producer (at left) is a processor that obtains time-series signals such as load and renewable generation forecasts. The OLS Producer may: 1) subtract the forecast renewable generation from the forecast load to produce a net generation¹⁴ shape, 2) flatten net generation to maximize the efficiency of the mix of thermal generators, and 3) add the forecast renewable generation and the flattened net generation together to create the OLS signal. These steps and the resulting actions of OLS Consumers are described in detail in the SCTE 267 easy-read standard.

A single entity may be a Producer and a Consumer, for example, an entity that consumes an OLS signal from an upstream Producer such as a utility may produce localized OLS signals and share them with several downstream consumers. Continuing the example, a building energy controller may consume an OLS signal from a utility and send a modified OLS signal to water heaters, thermostats, and vehicle chargers to coordinate time-shifting of demand among appliances to reduce energy costs and optimize the performance of a home or business.

The ANSI/SCTE 267 standard makes possible the shaping of load across small and large geographic areas. For example, Figure 8 is based on the simulation of the serving of the Electric Reliability Council of Texas (ERCOT) and shows the hourly generation based on actual load (at top) and daily optimum load (at bottom) on 20–26 Aug 2005.¹⁵

¹³ Source: Society of Cable Telecommunications Energy (2021), *ANSI/SCTE 267: Optimum Load Shaping for Electric Vehicle and Battery Charging*, <https://www.scte.org/standards/library/catalog/>

¹⁴ In this context, net generation refers to the generation required from thermal power plants to meet demand for electricity after accounting for the contribution of renewables.

¹⁵ The primary author's generation-to-load simulation estimated the impact, in terms of production costs and CO2 emissions, attributable to the joint optimization of electric power generation and flexible end uses to support increasing penetrations of renewable energy. Newly conceived, evaluated, and foundational in developing the ANSI/SCTE 267 American National Standard was a transaction-less, yet continuous demand response system based on a day-ahead optimum load shape (OLS) designed to encourage Internet-connected devices to autonomously and voluntarily explore options to favor lowest cost generators – without requiring two-way communications, personally identifiable information, or customer opt-in. Boundary conditions used for model calibration included historical

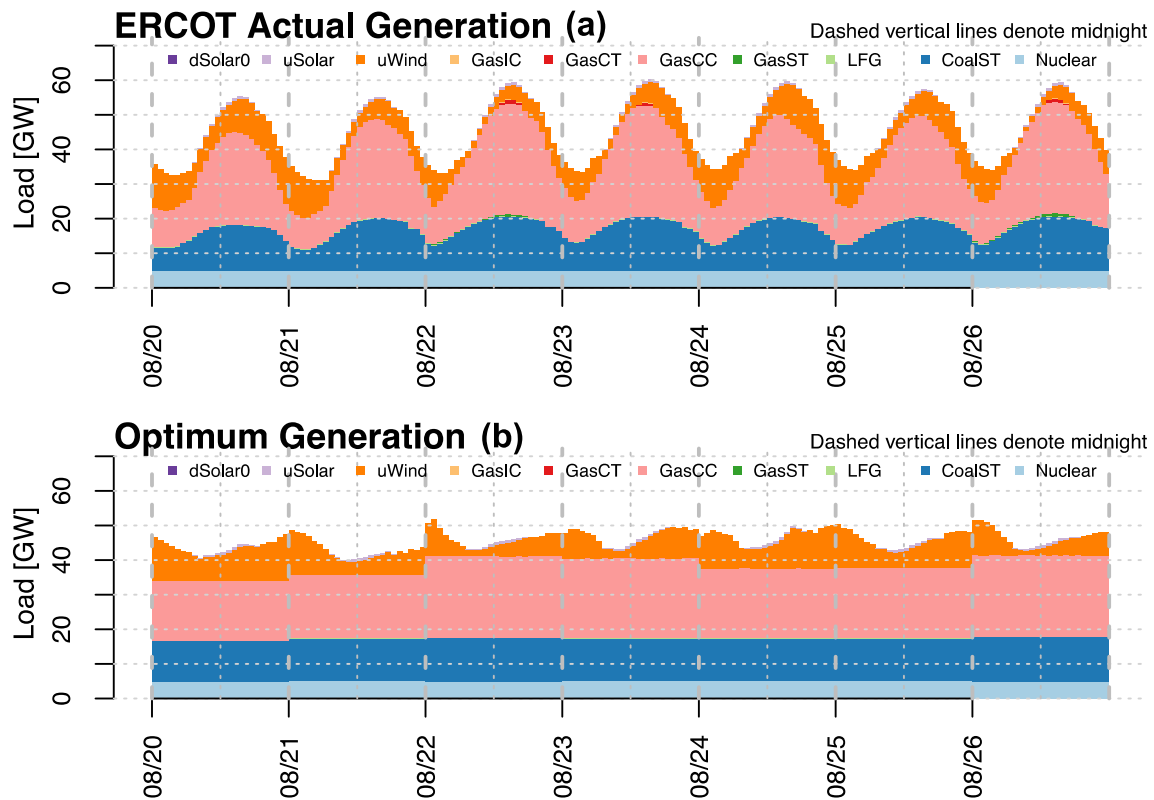


Figure 8 – ERCOT Hourly Generation Based on Actual and Daily Optimum Load.¹⁶

In Figure 8, production cost and CO₂ emissions were calculated for the unshaped actual historical hourly load (at top), and for the same load had it been optimally shaped over seven days. As expected, the minimum production cost was achieved when power plant output was constant, depicted as flat lines for thermal generation, as shown for each hour of 20–26 Aug 2005 as shown at the bottom of Figure 8. The week of 20–26 August is of interest as: (a) it was the hottest week of 2005 with over a TWh of energy delivered each day, and (b) many of the 600+ thermal generation power plants in Texas were active and could benefit from higher efficiencies that would result from higher capacity factors.¹⁷

The primary author's research, funded by the U.S. Department of Energy and conducted in-residence at National Renewable Energy Laboratory (NREL), shows that fixed and mobile batteries are the most

weather, residential building stock construction attributes, home appliance and device empirical operating schedules, prototypical power distribution feeder models, thermal generator heat rates, startup and ramping constraints, and fuel costs. Results of an hourly-based annual case study of Texas indicate a 1/3 reduction in production costs and a 1/5 reduction in CO₂ emissions are possible.

¹⁶ Robert Cruickshank, Ph.D. thesis, University of Colorado Boulder (Sep 2019), *Estimating the value of jointly optimized electric power generation and residential electrical use*.

https://scholar.colorado.edu/concern/graduate_thesis_or_dissertations/x059c851q

¹⁷ In generation, capacity factor is defined as the amount of energy actually produced divided by the maximum amount of energy that could have been produced at full output power. In transmission and distribution, capacity factor is defined as the actual amount of energy moved divided by the maximum amount of energy that could have been moved.

critical DERs to be networked first.¹⁸ The trend of increasing electric vehicle and in-building batteries is an opportunity to be addressed. To that end, e-Radio USA manufactures an SCTE 267-compatible EV charging adapter with various connectors that ensure backward and forward compatibility with all types of EVs as shown in Figure 9. Any kind of electric vehicle can plug into the EV Charging Adapter, which plugs into any EV charger. The adapter's green light indicator illuminates when the EV Adapter is receiving an authenticated FM, satellite, LTE or Wi-Fi signal and is working correctly. The EV charging adapter optimizes charging in response to grid signals, ensuring that wind, solar, and low-cost energy sources are maximized – and can be configure-less when used in broadcast applications. The vehicle owner can use the vehicle's standard control system to override the OLS signal if necessary to assure immediate charging.




Figure 9 – Electric Vehicle Charging Adapter.¹⁹

There are evolving accessibility and information security concerns to be considered and addressed to protect consumer privacy and grid security. To that end, the broadcast one-way delivery method for OLS signals may use multi-network authentication; thus, DERs do not add or shed load until each autonomously confirms the identical OLS signal is being received by multiple sources and/or networks. For example, a DER such as an EV charging adapter would verify that the same guidance signal is being received via FM radio, satellite, and/or the Internet. In all cases, the security of DERs, the data collected, security of software, and algorithm development processes must be addressed.

As part of DER participation, the bulk power system will benefit from a three-level battery hierarchy that manages the flow of electrons within and across transmission, distribution, and consumer networks. Transmission-scale and distribution-scale batteries provide increased system resiliency and pay for themselves with resiliency and with revenues from capacity, energy, and frequency response. While utilities and energy market participants are likely to directly control batteries located in transmission and distribution (T&D) networks, most batteries will be premises-based, consumer-owned, and can be voluntarily controlled by SCTE 267 OLS signals to the benefit of all three G&T&D realms.

¹⁸ Robert Cruickshank, Gregor Henze, Anthony Florita, Charles Corbin & Killian Stone (2021), *Estimating the value of jointly optimized electric power generation and end use: a study of ISO-scale load shaping applied to the residential building stock*, Journal of Building Performance Simulation, DOI: 10.1080/19401493.2021.1998222

¹⁹ Power Networks, LLC, Submission to Public Utility Commission of Texas (Jun 2022). Image courtesy of e-Radio and Xperi <https://drive.google.com/file/d/1CStDFzorzrME8Czf49TqVMsHj7fyu0SZ/view?usp=sharing>

SCTE 267 OLS signals are widely usable due to their ability to guide battery charge and discharge. OLS signals can guide a mobile or fixed battery to modulate charge/discharge over time. Figure 10 depicts battery charging and discharging wherein the charger voluntarily and autonomously uses the OLS signal (in blue) to inform charging and discharging profiles (in orange); a sinusoidal OLS signal illustrates the concept.




Figure 10 – Hourly Optimum: 1) Load Shape, 2) Charge/Discharge Shape.²⁰

In Figure 10, a battery charge controller receives an OLS signal (in blue) and then autonomously draws a horizontal line that divides the signal into charging and discharging intervals (in orange). The timing of charging/discharging is modified daily and hourly based on the forecasted availability of wind and solar power, actual grid contingencies, power system and user needs. When severe weather is expected, batteries can choose to enter “storm mode” and charge up to fully prepare for a potential grid outage.

Congestion detection and mitigation are increasingly crucial as beneficial electrification results in distribution networks connecting to more loads, particularly more numerous and higher capacity EV chargers. Once congestion is detected by location-aware DERs, it can be mitigated using a “local” OLS signal to raise the carrying capacity and extend the lifespan of a distribution network, as shown in Figure 11.

²⁰ Robert Cruickshank Associates (Dec 2021), Presentation to SCTE Microgrid working Group.




Figure 11 – A “Local” Load Shape to Mitigate Grid Congestion.²¹

Local load shaping is something broadband providers can be very good at and accomplish very quickly. Per SCTE 267, some DERs, like EVs, can choose to share their latitude and longitude when checking for load shapes. In Figure 11, mitigation of congestion (segments and locations shown in red at left) is achieved by modulating loads to utilize fully, but not overload, a distribution circuit using a local load shape for loads downstream of T&D congestion points. In the center of Figure 11, the flat purple shape guides DERs to adjust their load to raise the capacity factors of congested T&D network segments.

SCTE 267 and 271 were designed to work separately and together. For example, voltage sags sensed by an SCTE 271 sensor can indicate congested segments – which can deliver more energy when SCTE 267 local load shapes guide DERs to favor off peak electricity usage.

6.2. ANSI/SCTE 271 Power Sensing in Cable and Utility Networks

The SCTE 271 standard specifies provides precision, sampling rate, and configuration requirements if vendors measure and report voltage and current in hardware and software to enable advanced power sensing in cable and utility networks. Included are requirements for sensing and communicating power quality observations from both the 60/75/90 VAC quasi-square wave HFC network and the 120/240 VAC supply from the electric power grid. For systems that support remote communication of measurements, requirements for the control plane and communications security are specified. The specification does not require any particular measurements. Still, if supply voltage and current are measured, it specifies how those measurements are made to benefit the grid and broadband network operations.

The development of the SCTE 271 sensing standard was motivated by successes in the U.S. Department of Energy (DOE) Situational Awareness of Grid Anomalies (SAGA) project underway at the NREL and CableLabs since October 2019. Early in the SAGA project and while developing the Gridmetrics[®] Power Event Notification System (PENS), the NREL and CableLabs team recognized that the existing sensing capabilities specified in the ANSI/SCTE 25-3 standard were 1) out of date, and 2) could be updated to enrich grid and broadband Proactive Network Maintenance by finding congestion and loose seizure screws, high-impedance faults, and other voltage and current glitches and anomalies that could lead to grid and HFC service outages.

²¹ Source: R Cruickshank, A Silverstein, A von Meier (Jun 2022), The Society for Standards Professionals, Standards Engineering Journal, *Broadband standards to manage and monitor the grid*. <https://www.ses-standards.org>

For grid providers, a timeless and ever-important question is, “How many EVs and batteries can a distribution segment or transformer support before there is an increased risk of congestion and premature equipment failures and outages due to a thermal overload? For example, what would happen if many consumers in a community traded in their internal combustion engine vehicle for an EV? Would the community be unable to charge the new load profile with today’s “plug it in and take it” model? Furthermore, how could signaling and sensing comfortably help answer that question with a resounding YES? Both of the SCTE standards address these needs.

Today, congestion visibility in sensor-starved distribution networks can be achieved by deploying power quality sensors, such as SCTE 271-based sensors. Once detected, congestion can be mitigated using local OLS signals to avoid slowdowns/failures on the grid.

Many forms of grid measurement work together to ensure a resilient and sustainable power grid. As shown in Figure 12, Advanced Metering Infrastructure (AMI) provides kW demand and kWh consumption at customer meters, typically reported at 15-min resolution. Supervisory Control and Data Acquisition (SCADA) provides voltage or current magnitudes, reported at a resolution on the order of several seconds. Phasor Measurement Units (PMUs) provide voltage or current magnitudes and phase angles, frequency, and derivative quantities reported roughly every cycle (25-120 Hz). Point-on-Wave (POW) sensors, provide voltage or current magnitudes and phase angles, frequency and derivative quantities, and 256 to 1 million samples/sec of voltage or current waveform, reported for a short duration or on a continuous monitoring basis. SCTE 271-compliant sensors implement POW functionality to provide unprecedented visibility to grid issues.




Figure 12 – Time Scales for Electric Grid Monitoring and Control.²²

Referencing Figure 12, POW sensors provide the finest temporal resolution of power quality information to communications providers, utilities, and others. POW sensors stream data in real-time, enabling the rate of change of frequency and sine-wave goodness-of-fit to be calculated remotely from the measuring network element. This approach is favored by utility engineers and operations staff as it is robust in situations with high electrical noise and distortion, where traditional assumptions about sinusoidal waveforms are not met.

A key feature of SCTE271-compliant sensors is their ability to stream uncensored waveforms for cloud-based analysis aided by operations experts and machine learning. Utility engineers are reacting positively

²² International Council on Large Electric Systems (CIGRE) Webinar Academy, Alexandra von Meier (Oct 2020), *AI on the grid: Understanding PMU data*. <https://www.youtube.com/watch?v=qRAPYVtC2zM> and <https://iclesunc.memberclicks.net/assets/AI%20on%20the%20Grid%20%28Day%201%29.pdf>

to the availability of POW streaming waveforms that aid in detecting vegetation strikes and other grid failure signatures. By using human-in-the-loop machine learning, failure signatures can be rewound in time to improve early detection of grid and broadband network physical layer issues and provide a quantum leap beyond current best practices for identifying and predicting network issues.

SCTE271-compliant sensors do not over compress or distort data before backhauling to the cloud. With voltage and current precision of 0.02% per unit and sampling rates of 10,000 samples per second, SCTE-217 sensors preserve anomalies such as the voltage spike depicted in top panel of Figure 13.




Figure 13 – Analog Waveform, Utility Capture and Broadband Capture.²³

Figure 13 shows an actual analog wave form (at top), a phasor approximation commonly used in utility best practices (in middle), and a SCTE 271-based POW capture (at bottom). Note the glitch in the actual waveform is not observed by utility best practices that assume the underlying alternating current signal is a perfect sine wave –but is accurately observed by POW measurement.

7. Conclusion and Next Steps

Around the world, the aging and frail bulk power grid is weakened by a confluence of factors that include inefficiencies, fuel scarcities, unbridled demand, and increasingly severe and unpredictable weather.

²³ North American Synchrophasor Initiative (3/20/2020), Technical Report: High-Resolution, Time- Synchronized Grid Monitoring Devices, Alison Silverstein, Alison Silverstein Consulting, Dr. Jim Follum, PNNL

Temperature extremes add enormous demand while lowering the efficiency of generation, transmission, distribution, and end use. Costs are skyrocketing, and the impacts of outages are more cataclysmic due to our growing dependency on old and new electrically powered technologies. Moreover, there is a lack of consumer adoption of energy efficiency and automation.

The cable broadband industry has enormous opportunities to monetize critically essential capabilities for the evolving grid. New SCTE standards for managing and monitoring the grid enable relatively rapid transformational opportunities that improve resiliency and sustainability on a global scale. The power and broadband industries should build out SCTE 267 and 271 grid infrastructure to maintain business continuity, contain energy costs, and create new business opportunities.

Rising financial and environmental costs can be thoughtfully contained. New operations and business models can provide additional revenues from delivering load shaping signals and grid sensing as a service. Cable's established and evolving playbook in traffic engineering informs the roadmap for modernizing the grid. A symbiotic future of the grid and broadband networks is possible and beneficial.

Likely outcomes point the way for collaboration, such as: 1) renewables will continue to be added to the generation mix, 2) the retirement of dispatchable thermal generation and the rise of renewables will create a growing void in controls that balance electricity supply and demand., and 3) balancing supply and demand on the grid will be managed and monitored across multiple timescales. As opportunities continue to mature, the SCTE 267 standard should be pro-actively updated to support load shaping in seconds and minutes to complement the hourly and daily timescales already implemented.

Abbreviations

ANSI	American National Standards Institute
DERs	distributed energy resources
DR	demand response
EIA	U.S. Energy Information Agency
EV	electric vehicle
G&T&D	generation, transmission and distribution
GHG	greenhouse gas
MISO	U.S. Midcontinent System Operator
MMBtu	Million British thermal units
MWh	Megawatt hour, a million-watt hours
SCTE	Society of Cable Telecommunications Engineers
T&D	transmission and distribution

Bibliography & References

1. U.S. Energy Information Agency, Short-Term Energy Outlook (6/16/22), *EIA expects significant increases in wholesale electricity prices this summer.*
<https://www.eia.gov/todayinenergy/detail.php?id=52798>
2. U.S. Energy Information Agency, Short-Term Energy Outlook (1/1/22), *Coal will account for 85% of U.S. electric generating capacity retirements in 2022.*
<https://www.eia.gov/todayinenergy/detail.php?id=50838>

3. U.S. Energy Information Agency, Short-Term Energy Outlook (12/7/21), *In September, the U.S. was at its lowest coal stockpiles since 1978.*
<https://www.eia.gov/todayinenergy/detail.php?id=50558>
4. CleanTechnica (6/5/22), *Potential Electricity Reliability Concern for Central U.S.A.*,
<https://cleantechnica.com/2022/06/05/potential-electricity-reliability-concern-for-central-u-s-a/>
5. U.S. Energy Information Agency, *Today in Energy* (6/3/22),
<https://www.eia.gov/todayinenergy/detail.php?id=52618>, NERC (5/22), *2022 Summer Reliability Assessment*,
https://www.nerc.com/pa/RAPA/ra/Reliability%20Assessments%20DL/NERC_SRA_2022.pdf
6. MISO (4/28/22), *MISO projects risk of insufficient firm generation resources to cover peak load in summer months.* <https://www.misoenergy.org/about/media-center/miso-projects-risk-of-insufficient-firm-generation-resources-to-cover-peak-load-in-summer-months/>
7. Lawrence Livermore National Laboratory (6/16/22), *Energy, Water, and Carbon Informatics*,
https://flowcharts.llnl.gov/sites/flowcharts/files/2022-04/Energy_2021_United-States_0.png
8. U.S. Energy Information Agency, *Today in Energy* (5/13/22), *U.S. energy-related CO₂ emissions rose 6% in 2021.* <https://www.eia.gov/todayinenergy/detail.php?id=52380>
9. Julie McNamara (7/9/2019), *How do power grids beat the summer heat*, The Equation, Union of Concerned Scientists, <https://blog.ucsusa.org/julie-mcnamara/how-do-power-grids-beat-the-summer-heat>
10. U.S. Energy Information Agency, *Today in Energy* (9/9/21), *In 2019, U.S. inflation-adjusted energy expenditures fell 5%.* <https://www.eia.gov/todayinenergy/detail.php?id=49476>
11. Environmental and Energy Studies Institute (Jun 2022), *Beneficial Electrification, An Access Clean Energy Savings Program.*
[https://www.eesi.org/electrification/be#:~:text=Beneficial%20electrification%20\(or%20strategic%20electrification,the%20residential%20and%20commercial%20sectors](https://www.eesi.org/electrification/be#:~:text=Beneficial%20electrification%20(or%20strategic%20electrification,the%20residential%20and%20commercial%20sectors)
12. U.S. Department of Energy, Grid Modernization Laboratory Consortium (Mar 2021), *Survey of Distributed Energy Resource Interconnection and Interoperability Standards*,
<https://www.nrel.gov/docs/fy21osti/77497.pdf>
13. Source: Society of Cable Telecommunications Energy (2021), *ANSI/SCTE 267: Optimum Load Shaping for Electric Vehicle and Battery Charging*,
<https://www.scte.org/standards/library/catalog/>
14. Robert Cruickshank, Ph.D. thesis, University of Colorado Boulder (Sep 2019), *Estimating the value of jointly optimized electric power generation and residential electrical use.*
https://scholar.colorado.edu/concern/graduate_thesis_or_dissertations/x059c851q
15. Robert Cruickshank, Gregor Henze, Anthony Florita, Charles Corbin & Killian Stone (2021), *Estimating the value of jointly optimized electric power generation and end use: a study of ISO-scale load shaping applied to the residential building stock*, Journal of Building Performance Simulation, DOI: 10.1080/19401493.2021.1998222
16. Power Networks, LLC, Submission to Public Utility Commission of Texas (Jun 2022). Image courtesy of e-Radio and Xperi
<https://drive.google.com/file/d/1CStDFzorjrME8Czf49TqVMsHj7fyu0SZ/view?usp=sharing>
17. R Cruickshank, A Silverstein, A von Meier (Jun 2022), The Society for Standards Professionals, Standards Engineering Journal, *Broadband standards to manage and monitor the grid.*
<https://www.ses-standards.org>
18. International Council on Large Electric Systems (CIGRE) Webinar Academy, Alexandra von Meier (Oct 2020), *AI on the grid: Understanding PMU data.*
<https://www.youtube.com/watch?v=qRAPYVtC2zM> and
<https://iclesunc.memberclicks.net/assets/AI%20on%20the%20Grid%20%28Day%201%29.pdf>

From Support Calls to Insights

Using Automatic Speech Recognition and Natural Language Processing to Drive Product Roadmaps

A Technical Paper prepared for SCTE by

Jing Qing

Principal Data Scientist I
Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
Jing.Qing@charter.com

Veronica Bloom

Director of Data Science
Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
+1 720-699-3798
Veronica.Bloom@charter.com

Michael Addonisio

Lead Data Scientist
Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
Michael.Addonisio@charter.com

Christy Gearheart

Data Scientist IV
IntePros Consulting, Charter Communications
6380 S Fiddlers Green Circle, Greenwood Village, CO 80111
C-Christy.Gearheart@charter.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Technical Approach	3
2.1. Call disposition Flow Overview	3
2.2. Machine Learning Pipeline Overall Architecture Review	3
2.3. Data Ingestion	4
2.4. Audio Preprocessing	4
2.5. Convert audio into text transcripts.....	4
2.5.1. Custom vocabulary	5
2.5.2. Channel identification.....	5
2.5.3. Redacting transcripts	5
2.5.4. Word Error Rate (WER)	5
2.6. Classification	5
2.6.1. Identify call disposition categories	5
2.6.2. Multi-label classification approach	6
2.6.3. Create labels through LabelStudio.....	6
2.6.4. Train Comprehend classification model and model evaluation	7
2.6.5. Inference	7
2.7. Topic Modeling	8
2.7.1. Text normalization.....	8
2.7.2. Training the Topic Model	8
2.7.3. Topic interpretation	9
2.8. Productionalization and Monitoring	9
2.8.1. Data Ingestion	9
2.8.2. Audio Preprocessing.....	9
2.8.3. Convert audio into text transcripts	9
2.8.4. Classification	10
2.8.5. Configuration Details.....	10
2.9. How AI/ML Insights Drive Product Roadmaps.....	10
3. Future work	10
4. Conclusion.....	11
5. Acknowledgements	11
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 – ML Pipeline Architecture Overview.....	4
Figure 2 – Label Studio call labeling interface	6
Figure 3 – Agent vs. ML classification F1 score performance	7

1. Introduction

The goal of this project is to utilize artificial intelligence (AI) and machine learning (ML) techniques to gain insights into the drivers of support calls. These insights are used to determine areas of improvements to both processes and products to enhance our customer experience.

In this paper, we discuss the end-to-end automated pipeline to go from call center data to actionable insights utilizing machine learning techniques such as Automatic Speech Recognition (ASR) and Natural Language Processing (NLP). Discussion of the pipeline architecture, use cases of automated call disposition, and call topic trend analysis are also included. Broadly, the pipeline converts recorded call audio conversations to text transcripts using automatic speech recognition, then uses those transcripts to train classification models to predict call disposition. Unsupervised topic modeling extracts new trends and topics from the call-volume data over time to identify new or emerging call drivers. These call drivers subsequently drive new feature development and product roadmaps.

Please note: Charter has a longstanding commitment to protecting the privacy and security of its customers. To provide our customers with technical support and high quality customer service as well as to determine areas of improvement, customers are informed that communications between customer service agents and customers may be recorded when authorized by the customer for quality and training purposes.

2. Technical Approach

2.1. Call disposition Flow Overview

Current call dispositions are collected through a manual process in a call tracking tool, where call center agents select a call disposition and enter notes into the system. If a call is transferred to another agent, it is split into call segments where each agent selects a disposition for their corresponding segment. Dispositions are predefined, and selected via drop-down menu.

While this approach enables quick selection and consistent responses, there are some limitations. Selection of disposition is subjective and can vary from agent to agent. Only one disposition can be selected per call, even when multiple issues are discussed. Changes to the disposition response list requires both a software update and agent training, and cannot be applied retroactively. Furthermore, due to the nature of a manual dispositioning process, some call segments are missing a recorded disposition in the tracking tool.

These limitations can be addressed through a machine learning approach. Using an automated machine learning approach enables scalable, consistent, and reliable call disposition labeling. Nuances in how an individual agent labels a given call segment are removed, as all call segments are systematically evaluated. Furthermore, assignment of multiple labels to a single call can bring to light underlying factors affecting multiple call reasons.

2.2. Machine Learning Pipeline Overall Architecture Review

We have developed a pipeline for data processing and modeling, in combination with cloud AI services such as cloud transcribe and cloud NLP services for an end-to-end call disposition architecture. Additional call insights are gained through the implementation of our topic modeling framework. The pipeline is implemented in Python 3 through boto3 (Amazon Web Services 2015). See Figure 1.

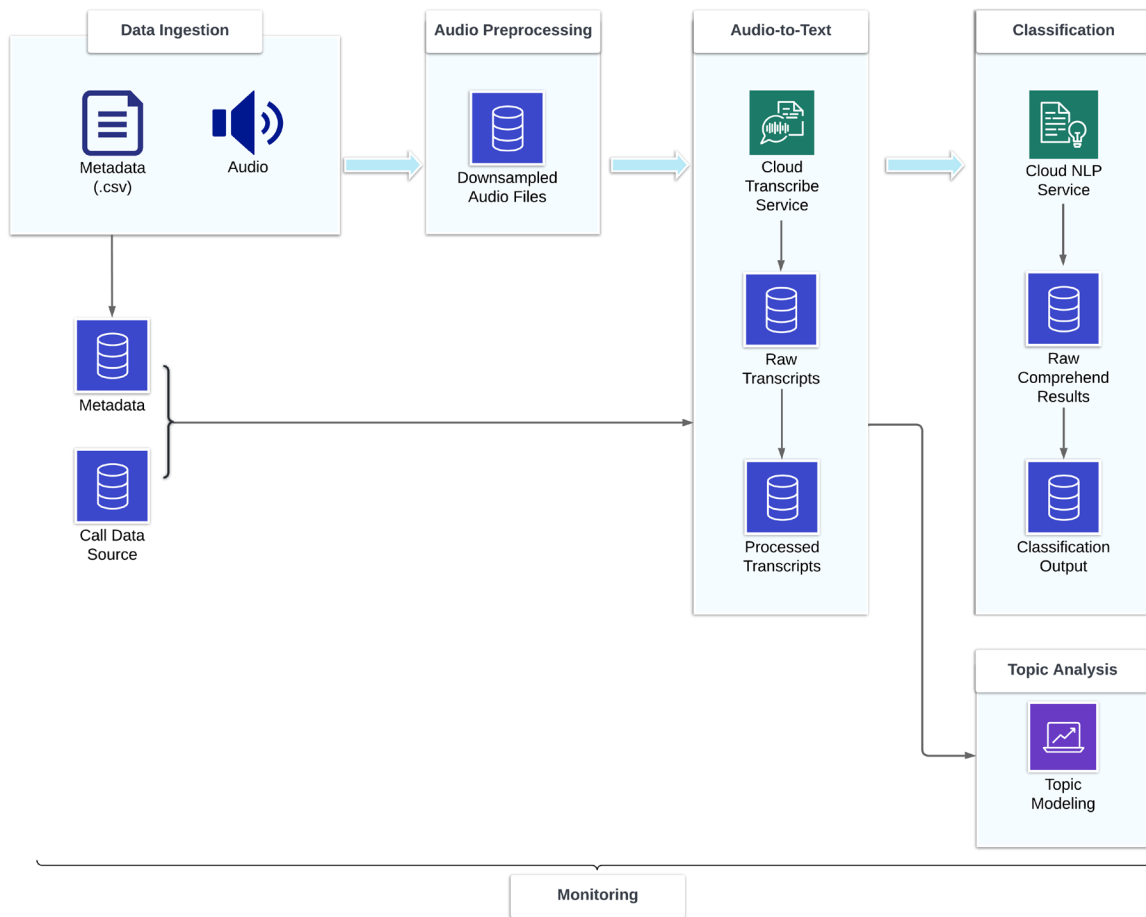


Figure 1 – ML Pipeline Architecture Overview

2.3. Data Ingestion

Every morning, call audio recordings and associated metadata from the previous day are ingested into the data science cloud storage.

2.4. Audio Preprocessing

To prepare the audio data for the cloud transcribe service for the best transcript output, we preprocess audio recordings after the audio files are ingested. In this step, we downsample the audio to a sampling rate of 8,000 Hz, and convert the audio files into .wav format through a parallel batch process. The downsampling and conversion are implemented through the open-source library librosa (McFee, et al. 2022). This pipeline is also batch parallelized and implemented through joblib (Joblib developers 2020).

2.5. Convert audio into text transcripts

There are multiple automatic speech recognition services available on the market to convert audio to text transcripts. In our pipeline, we utilize a cloud transcribe service to convert downsampled audio files into text transcripts. We harness built-in features such as custom vocabulary to improve the transcript quality, channel identification to add data dimensionality for analysis, and transcript redaction for data privacy and compliance. We experimented with different settings and decided on the configuration for production

based on the best word error rate (WER) evaluation metric. We describe these features and the WER evaluation in the next sections.

2.5.1. Custom vocabulary

Using custom vocabulary can help to improve transcribe accuracy for domain specific words and phrases. We worked with subject matter experts (SME) to identify technical vocabulary words and phrases commonly used in call center conversations for our use case. With sound-alike and phonetic annotations, we created a custom vocabulary for the cloud transcribe process to ensure Spectrum specific terms are represented correctly in the final transcripts.

2.5.2. Channel identification

Since the audio recordings contain two audio channels, we are able to generate transcripts for the full conversation, as well as separated transcripts by each audio channel. The transcripts by channel could be used in further analysis, such as associating topics in different channels to call issues and resolutions (Pratik K. Biswas 2021).

2.5.3. Redacting transcripts

We enabled redaction in the cloud transcribe service, where Personally Identifiable Information (PII) such as name and address are removed from the transcript output.

2.5.4. Word Error Rate (WER)

Speech-to-text is quite challenging for call center audio recordings. The audio quality is usually low, and there could be background noise such as a TV playing, dogs barking, or the customer asking questions. There are at least two people on the call, sometimes talking at the same time, and the speaking style is conversational and usually loosely structured. Thus, it is crucial that we measure the quality of the speech-to-text translations.

The most common metric for speech recognition accuracy is called word error rate (WER) (Seyfarth and Zhao 2020). WER counts the number of incorrect words identified during recognition, and divides the sum by the total number of words provided in the human-labeled transcript (N). $WER = (I + D + S)/N$, where I stands for Insertion Error: words incorrectly added in the transcript; D stands for Deletion Error: words undetected (deleted) in the transcript; S stands for Substitution Error: words substituted between reference and hypothesis. The lower the WER, the more accurate the system.

We randomly selected 15 transcripts and labeled them to evaluate WER. The WER ranged between 5% and 43%, with an average WER of 19%. By comparison, a benchmark analysis of nearly 3,000 call center conversations across five domains transcribed with automatic speech recognition (ASR) tools versus transcribed by two professional annotators found WER for call center conversations within the telecommunications domain to be between 17.62% and 23.31% (Szymański, et al. 2020).

2.6. Classification

2.6.1. Identify call disposition categories

The data science team worked with SMEs to identify eight high-level categories to train a classification model to predict call dispositions.

2.6.2. Multi-label classification approach

In the multi-label classification, individual classes represent different categories, but these categories are somewhat related and not mutually exclusive. As a result, each document has at least one class assigned to it, but can have more. In call center calls, a conversation can discuss multiple topics, justifying the use of multi-label classification.

We implemented the multi-label classification model through a cloud NLP service. For each input transcript, the model returns three labels with the highest predicted scores. Note, in some cases, even the highest predicted score is low if the model is not confident in the classification of any of the eight categories. Therefore, we applied category-specific score thresholds to the predicted labels, and only keep labels where the score is higher than the threshold as the final predicted label(s).

2.6.3. Create labels through LabelStudio

To train a classification model, it requires us to provide the cloud NLP service with examples of call transcripts with their appropriate classification labels. In order to do this, the data science team worked with SMEs to understand what to look for in calls in order to provide example labels. This team listened to approximately 1,000 calls in order to provide examples of about 100 calls for each category.

To support labeling the call data, we set up a cloud instance to run Label Studio (Tkachenko, et al. 2020), an open source annotation tool. Label Studio allowed us to create a custom interface that includes both audio and text media for annotation. The annotator was required to label each document with at least one label and optionally up to three labels. Each document was labeled at least twice, and, if there was no consensus, a third time. An example of the annotation layout can be seen in Figure 2. We consider the call categories labeled in Label Studio through this process as the “ground truth” data.

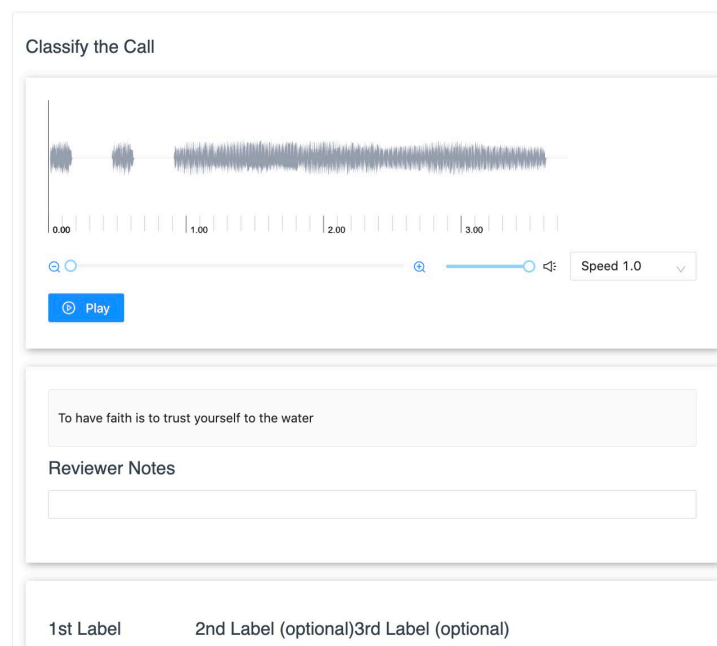


Figure 2 – Label Studio call labeling interface

2.6.4. Train Comprehend classification model and model evaluation

Once we obtained approximately 100 example calls for each category, we split our transcripts into 675 files for training and 191 files for validation. We experimented with various transcribe and comprehend configurations in the pipeline to determine precision and recall scores, and combined them into a single representative micro F1 score. A micro F1 score balances precision and recall across individual categories and provides a single quality metric for multi-label binary problems. The best performing comprehend model has a micro F1 score of 60% on both the training and validation set. Finally, we evaluated the fit of our model based on a random selection of 90 transcripts of previously unseen data, and examined the model's eight categorical predictions compared to the manually-curated call dispositions supplied by our call center agents.

We found the classification model trained with the cloud NLP service obtained a micro F1 score of about 59% on the evaluation set across the categories, with some categories performing better or worse than others. For example, one category obtained an F1 score of about 84%, while another category obtained an F1 score of about 25%. We also found the classification model performed very similarly to agent-based dispositions (58%) against labeled ground truth data, with similar performance across categories. See Figure 3. Future work will focus on improving the comprehend model with additional labeled data and refining the classification categories.




Figure 3 – Agent vs. ML classification F1 score performance

2.6.5. Inference

After training the classification model, we ran the model on a month of call transcripts from one month to understand which of these eight categories are driving the majority of calls. Recall that calls can be classified in multiple categories. We found the majority of calls for our use case, 40%, came from requiring support in one specific area. We also see a large portion of these calls, 27%, fall into the “Other” category. In the next section, we show how utilizing topic modeling can give us more granular

insights into the call drivers in the Other category, and will break down the top category into actionable subcategories.

2.7. Topic Modeling

When we shared the ML disposition results with stakeholders, they immediately wanted to know if we could find out more granular call reasons and if we could identify improvements to features and workflows to provide improved product experience. To answer these questions, we applied topic modeling to further analyze the call transcripts.

Topic modeling is an unsupervised method for discovering topics in a large collection of documents. Topics can then be used to find high-level summaries of these documents, search for documents of interest, and group similar documents together. In practice, we implemented the open-source library Top2Vec (Angelov 2020) for our analysis. Some benefits of the Top2Vec algorithm include automatic identification of the number of topic clusters and the ability to identify topic hierarchies.

The topic modeling process includes the following steps: text preprocessing such as normalization, training the topic model, interpreting topics, and visualizing and analyzing trends.

2.7.1. Text normalization

Text normalization is the process of transforming text into a single canonical form, so it is “cleaner” or less random. We found this step necessary to obtain more interpretable topics. As mentioned earlier, the audio-to-text task is challenging for call center audio files. Even after significant efforts to optimize AWS Transcribe, there are still consistent words or phrases incorrectly represented in the transcripts. Additionally, there are text formats in the transcripts that interfere with the topic model’s tokenization process. For example, “DVR” is often displayed as “D. V. R.” in the transcript. Since we also use “.” as a text delimiter, it would have been split into three separate letters “D”, “V”, “R”, and the individual letters lose the meaning compared to “DVR” as a whole. Through normalization, we are able to fix consistent transcribe errors, and normalize words and phrases so they will be represented correctly for topic modeling.

In order to handle the text normalization task on a large volume of call transcripts, we needed a process that is optimized and parallelizable. We chose HuggingFace’s tokenization pipeline (Wolf, et al. 2019) to normalize the data, and trained a Byte-Pair Encoding (BPE) Tokenizer on call transcripts to generate cleaned input documents for the Top2Vec model. We chose this tokenization pipeline because it is an efficient and optimized implementation, allowing us to use all cores on the data science cluster and process the text in parallel, which significantly reduced the processing time.

2.7.2. Training the Topic Model

The Top2Vec algorithm includes the following steps: the first step is to create a joint embedding of document and word vectors. Once documents and words are embedded in a vector space, the goal of the algorithm is to find dense clusters of documents, then identify which words attracted those documents together. Each dense area is a topic and the words that attracted the documents to the dense area are the topic words (Angelov 2020).

For experiments, we tested multiple embedding models, including Doc2Vec (Rehurek 2022), and pretrained embedding models such as universal-sentence-encoder (Cer, et al. 2018) and all-MiniLM-L6-v2 (Reimers and Gurevych 2019). We have found using Doc2Vec to generate the document and word embedding vectors produced the most interpretable topics. This is because call transcripts contain many

Spectrum-specific terms, and the Doc2Vec embedding is trained from scratch using all the transcripts as part of the topic modeling process. Therefore, the Doc2Vec embedding is able to represent Spectrum-specific terms well. While pretrained embedding models, even though they are trained on massive language datasets using complex neural networks, since they have never “seen” Spectrum call data before, they require additional fine tuning to perform well with our call data.

2.7.3. Topic interpretation

We trained the Top2Vec model on a month of call data. Because topic modeling produces unlabeled clusters of call transcripts, we supplemented the model by labeling each topic cluster with the most frequent keywords and bigrams. This provided some contextual meaning to each of the clusters.

2.8. Productionalization and Monitoring

Given the number of moving parts in our pipeline, we have established a monitoring framework to automatically measure data availability, reliability, and operational performance against our baseline. This enables us to identify and correct any arising issues early. Our monitoring framework has been integrated into the pipeline so every execution produces an associated monitoring report.

2.8.1. Data Ingestion

Each day, we receive call audio recordings in a cloud storage repository. Additionally, we receive a daily manifest file listing each audio file expected and its associated metadata, such as unique call identifier, agent information, and source of call. As an initial step to our pipeline, we validate the files received match the files listed in the manifest. Independently, we validate the manifest file contains the calls we expect to receive based on a separate internal call data source. By utilizing automated data validation, we are able to quickly identify and address discrepancies when these numbers deviate from one another. As a result of this monitoring, we are able to quickly identify and resolve this issues that come up during normal course of business.

2.8.2. Audio Preprocessing

We downsample and convert audio files to WAV format for ingestion into a cloud transcribe service. We monitor the number of audio files that failed to convert, the total count of files successfully converted, and the minimum, maximum, average, standard deviation, and percent quartiles for both call lengths and file sizes following downsampling. This enables us to capture any issues in file transfer or conversion in our hybrid environment, and to provide insights into the size and length of the audio files received.

Monitoring became crucial in identifying an underlying issue in our downsampling process. We were finding roughly 100 files were being produced with a zero-length audio file. Interestingly, the files producing the zero length audios were different from run to run. Through troubleshooting, we identified that the pydub python library (Robert, Webbie and others 2018) would intermittently fail to downsample a given file when run as a parallel process. In the end, this was overcome by our transition to the librosa python library.

2.8.3. Convert audio into text transcripts

As files are transcribed, we capture the total number of audio files that fail as well as the individual errors and corresponding number of files per failure reason. Capturing the reasons associated with the failures enables us to see if there are underlying errors requiring remediation. The most common error observed

stems from the file being less than .5 milliseconds, the minimum audio duration required for transcribe. Because we cannot process these files, they are logged and removed, and the pipeline continues execution.

After transcription, we parse the individual files based on their audio channels, join to their associated metadata, and combine into a single dictionary in preparation for classification. In addition to monitoring the total number of processed transcripts, we further monitor the individual word count statistics produced. For example, we monitor the mean word count across the transcribed files; when the word count drops substantially, we are able to identify if the current run has failed to appropriately transcribe the calls.

2.8.4. Classification

At the completion of classification, we validate the output file has been created and its location documented. Summary results are logged indicating the distribution of audio files across our defined classification labels. Insights are delivered to our stakeholders for actionable business intelligence.

2.8.5. Configuration Details

Finally, to ensure we have repeatable results, we log the configuration metrics, including run time, file locations, custom vocabulary file, and any process-specific parameters utilized in the run.

2.9. How AI/ML Insights Drive Product Roadmaps

With quantitative metrics in hand, our stakeholders are empowered with empirical evidence to better guide and prioritize their roadmaps. Broadly, we were able to utilize topic modeling to identify call drivers and use those call drivers to identify areas where improvements would increase efficiency. As such, these topic modeling results are driving current and future product roadmaps and supporting continual product improvement.

3. Future work

Future work is focused on improving the accuracy and efficiency of our pipeline. Using audio transcripts for call disposition is just the beginning. We want to enrich our call disposition data with additional data sources which can open up our analysis to understanding patterns, and will enable us to better understand the broader picture. Improvements to our visualization dashboard will better highlight topic insights derived from the model.

We also plan to experiment using the same tokenization pipeline on call transcripts to train a transformer-based model, such as BERT (Devlin, et al. 2018). As a result of the training process, BERT learns contextual embeddings for words. We can then use the custom trained model to produce embeddings for topic modeling. These BERT embeddings should contain more contextual information and might be able to improve the topic model.

As an unsupervised learning method, topic modeling is great in clustering semantically-similar documents together to show trends; however, it still requires human input for topic interpretation and requires effort to measure output accuracy. Conversely, supervised classification modeling, is easy to set up, but requires high-quality labels for meaningful training and evaluation. As such, we plan to augment our topic modeling process with a human-in-the-loop to create new training labels for the classification model.

4. Conclusion

We have discussed an automated end-to-end pipeline from call center data to actionable insights. In our pipeline, we developed a flow to ingest audio data, performed audio downsampling, converted audio to text, ran a classification model on high-level dispositions, and performed topic modeling for more granular insights. This pipeline is monitored at each step with measurable metrics for quality control.

While the manual call tracking tool collects high-level call dispositions, the machine learning and natural language processing approach is fast, reliable, and scalable. Machine learning elevates us from statistical metrics to deriving deep analytical insights and identifying new trends. These insights help shape product roadmaps and drive new feature development.

5. Acknowledgements

We would like to extend our special thanks to Jonathan Stribley, Sean O'Donnell, and Keara O'Brien for their partnership in these endeavors; the Call Center Technology team for their partnership in data ingestion, Sandeep Kumar Sarikonda, Pablo Lopez, and the Data Science Machine Learning Operations team for their assistance in setting up and maintaining the pipeline environment; the Data Platform team for establishing the development and production environment; Steven Naumann and the Data Engineering team for their assistance and insights with the internal call data; Laurie Porter for listening to hundreds of calls and helping us to understand the business process so we can establish a labeling guideline; Rohan Khataavkar and Aaron Osher for their help in labeling call dispositions; Cody Zeigler and Pierre Dumas from the cloud service team for their expertise and partnership in the proof of concept (POC) of the pipeline and productionalization; Mike Baldino for his support and leadership; and last, but not least, Marjorie Sedillo and Brock Bose for their support, expertise, and guidance in the development of the pipeline and model, as well as their editorial feedback on our manuscript.

Abbreviations

AI	artificial intelligence
ASR	automatic speech recognition
BERT	bidirectional encoder representations from transformers
BPE	byte-pair encoding
ML	machine learning
PII	personally identifiable information
POC	proof of concept
SME	subject matter expert
WER	word error rate

Bibliography & References

Angelov, Dimi. 2020. *Top2Vec: Distributed Representations of Topics*. arXiv. {<https://arxiv.org/abs/2008.09470>}.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin. 2017. *Attention Is All You Need*. <https://arxiv.org/abs/1706.03762>.

Cer, Daniel, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, et al. 2018. *Universal Sentence Encoder*. {<https://arxiv.org/abs/1803.11175>}.

- Devlin, Jacob, Ming-Wei Wang, Kenton Lee, and Kristina Toutanova. 2018. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." Oct. <https://arxiv.org/abs/1810.04805v2>.
- Joblib developers. 2020. *Joblib: Running Python Functions as Pipeline Jobs*. {<https://joblib.readthedocs.io/>}.
- McFee, Brian, Alexandros Metsai, Matt McVicar, Stefan Balke, Carl Thomé, Colin Raffel, Frank Zalkow, et al. 2022. *librosa. librosa*. {<https://github.com/librosa/librosa>}.
- Pratik K. Biswas, Aleksandr Iakubovich. 2021. *Extractive Summarization of Call Transcripts*. {<https://arxiv.org/abs/2103.10599>}.
- Rehurek, Radim. 2022. *Gensim: Topic Modeling for Humans - Doc2Vec Paragraph Embeddings*. <https://radimrehurek.com/gensim/models/doc2vec.html>.
- Reimers, Nils, and Iryna Gurevych. 2019. "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks." *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. {<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>}.
- Robert, James, Marc Webbie, and others. 2018. "Pydub." GitHub. {<http://pydub.com/>}.
- Seyfarth, Scott, and Paul Zhao. 2020. *Evaluating an automatic speech recognition service*. October 5. {<https://aws.amazon.com/blogs/machine-learning/evaluating-an-automatic-speech-recognition-service/>}.
- Szymański, Piotr, Piotr Żelasko, Mikolaj Morzy, Adrian Szymczak, Marzena Żyła-Hoppe, Joanna Banaszczyk, Lukasz Augustyniak, Jan Mizgajski, and Yishay Carmiel. 2020. "WER we are and WER we think we are." *Findings of the Association for Computational Linguistics: Empirical Methods in Natural Language Processing (EMNLP) 2020*. Punta Cana: Association for Computational Linguistics. 3290–3295.
- Tkachenko, Maxim, Mikhail Malyuk, Nikita Shevchenko, Andrey Holmanyuk, and Nikolai Liubimov. 2020. *Label Studio: Data labeling software*. {<https://github.com/heartexlabs/label-studio>}.
- Wolf, Thomas, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, et al. 2019. *HuggingFace's Transformers: State-of-the-art Natural Language Processing*. arXiv. {<https://arxiv.org/abs/1910.03771>}.

Gaming Isn't Homework: Predicting Demand at the Neighborhood Level

A Technical Paper prepared for SCTE by

Christopher Stump

Senior Engineer, Predictive Analysis
Comcast Cable
1701 JFK Blvd Philadelphia, PA
christopher_stump@cable.comcast.com

May Tan

Senior Manager, Traffic Forecasting
Comcast Cable
1701 JFK Blvd Philadelphia, PA
may_tan@cable.comcast.com

Adam Gruchevsky

Senior Manager, Network Planning
Comcast Cable
1701 JFK Blvd Philadelphia, PA
Adam_gruchevsky@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Current State.....	4
3. Model Evolution.....	7
4. The Model	10
5. Future State.....	11
6. Conclusion.....	12
Abbreviations	13

List of Figures

Title	Page Number
Figure 1 - Gaming is changing when SGs peak	3
Figure 2 - Forecasting layers	4
Figure 3 - SG traffic allocation.....	5
Figure 4 - Site to SG example.....	6
Figure 5 - Time Series model results	7
Figure 6 - Time Series model issues with irregular patterns.....	8
Figure 7 - sMAPE of NN on bps.....	8
Figure 8 - NN on bps example output.....	9
Figure 9 - Where NN on bps does not forecast well.	9
Figure 10 - NN on growth rate results.....	11

1. Introduction

Efficient network planning requires accurate forecasting of customer behaviors in the near and long term. The goal is to anticipate the location and magnitude of capacity augments months, and even years, into the future using forecasted monthly bits per second demand (Mbps). This prevents outages, increases reliability and prepares the network to handle faster speeds. However, no longer does a blanket growth rate apply to every neighborhood. Network architecture is transitioning to run fiber deeper into the network and make Service Groups (SGs) smaller, decreasing households passed while increasing speeds. This adds further challenges as smaller SGs are more sensitive to individual household patterns and faster speeds create more usage spikes. As an example, single releases of popular games downloading to a console are enough to create network congestion as are increasingly frequent live sports streaming events.

In the existing analog node world, an office park could hang off the same SG as a housing complex. The difference between day versus night peak, and weekday versus weekend, are washed away due to the size of the SG and augment decisions could be made by the macro trend to consume more traffic at faster speeds. Now, in the digital node world, SGs are smaller. An office park may have a dedicated SG and a residential community may have its own dedicated SG. If we sum their traffic at the headend level, we will certainly see the macro trend, but now the individual SGs have their own driving forces. The office park is driven by weekday, daytime peaks with seasonal dips for vacations and holidays, not to mention pandemics. The residential component is driven by nighttime and weekend peaks that spike heavily in work-from-home situations or when popular games have new releases (and homework lay forgotten). In fact, popular gaming releases can more than double the number of SGs that peak on a single day. With the popularity of gaming and the increase of sporting events dedicated to streaming, these spikes are anticipated to become more frequent. As such, the trends of these SGs cannot be accurately predicted using a headend or national level CAGR (Compound Annual Growth Rate)

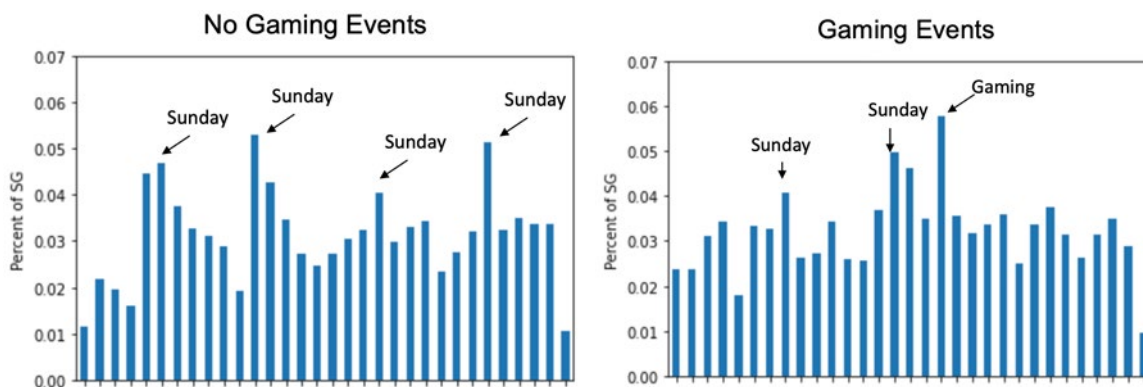


Figure 1 - Gaming is changing when SGs peak

The share of SGs experiencing their monthly peak on a given day is predictable when there are not large gaming download events.

At Comcast we are enhancing our forecast efforts to predict traffic deeper into the network. A small team within the CONNECT organization has worked for over one and a half years to develop a machine learning (ML) solution that produces individual patterns for SGs and Bonding Groups (BGs). The resulting Neural Network uses the growth patterns of SGs and BGs to accurately predict monthly 98th

percentile bps regardless of age or technology across the entire network representing hundreds of thousands of individual time series allowing for targeted network planning months and years in advance .

2. Current State

For more than 10 years, Comcast has been developing a comprehensive process to provide network traffic forecasting. Models are developed at the national, regional, and site level. Our access network component model is built to take the site-level aggregate and apply it to each active SG on the network. The model forecasts when SG augments will be required to maintain healthy SG utilization.

At the national level, there are macro level growth analysis models which look at different contributors, drivers, and industry trends. We also evaluate various assumptions and insights across Comcast product categories.

At the regional and site levels, time series focused Machine Learning (ML) has been developed to predict DOCSIS 3.0 (D3.0) and DOCSIS 3.1 (D3.1) traffic growth. We also apply reconciliation and a quality control process to reconcile the bottom-up (site level) model outputs with top-down (national level, regional level) models to ensure consistent forecasting output throughout. A Unified Demand Model (UDM) is the last step in the process to integrate bottom-up and top-down, reconciliation, quality control model outputs, products, device assumptions, and facility information. The output of the UDM model for access network is site level D3.0 and D3.1 traffic.

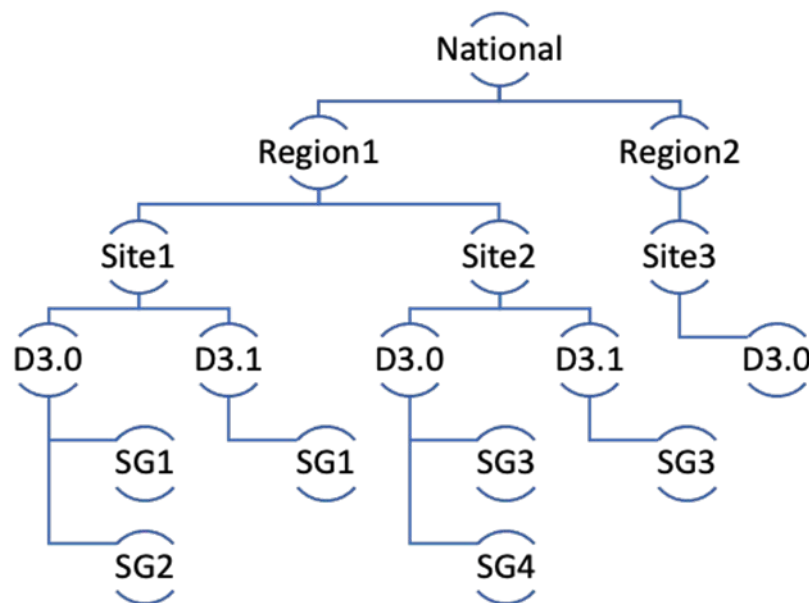


Figure 2 - Forecasting layers

National, Regional and Site level models are reconciled, then product, device and facility assumptions are added through the UDM to produce D3.1 and D3.0 traffic forecasts .

The next layer is SG. However, due to the number of the SGs, the scale and volatility of the data, we were not able to generate SG level forecast from the ML models which were used in site-level forecasting. At the SG level, our traffic growth process applies ratio-based site-level utilization to the individual SG (e.g., SG-01's last known actual bps utilization was 30% of the site-level aggregate, so in the next forecast period we will add 30% of the incremental bps site-level traffic to SG-01). We then look at the number of DOCSIS QAMs currently configured on each SG. If the number is less than the current year's target, we will add D3.0 or D3.1 channels to ensure the SGs will have the targeted capacity for that time period. QAM configurations will also be adjusted based on the amount of D3.0 versus D3.1 traffic. If an augment needs to occur in the model, a new SG is created; traffic is split among old and new SGs, and the traffic growth process continues. The ratio-based traffic growth process applies to upstream SGs as well.

n98th Mbps Aggregate		
	Site Actual	Site Forecast Month 1
Site A	5000	6000

	Service Group Actual	Percent of Actual	Service Group Forecast Month 1
SG-01	1500	30%	1800
SG-02	2000	40%	2400
SG-03	1500	30%	1800
Total	5000		6000

Figure 3 - SG traffic allocation

n98th Mbps Aggregate		
	Site Actual	Site Forecast Month 1
Site A	5,000	6,000

	Service Group Actual	%'age of Actual Mbps	Service Group Forecast Month 1
SG-01	1,500	30%	1,800

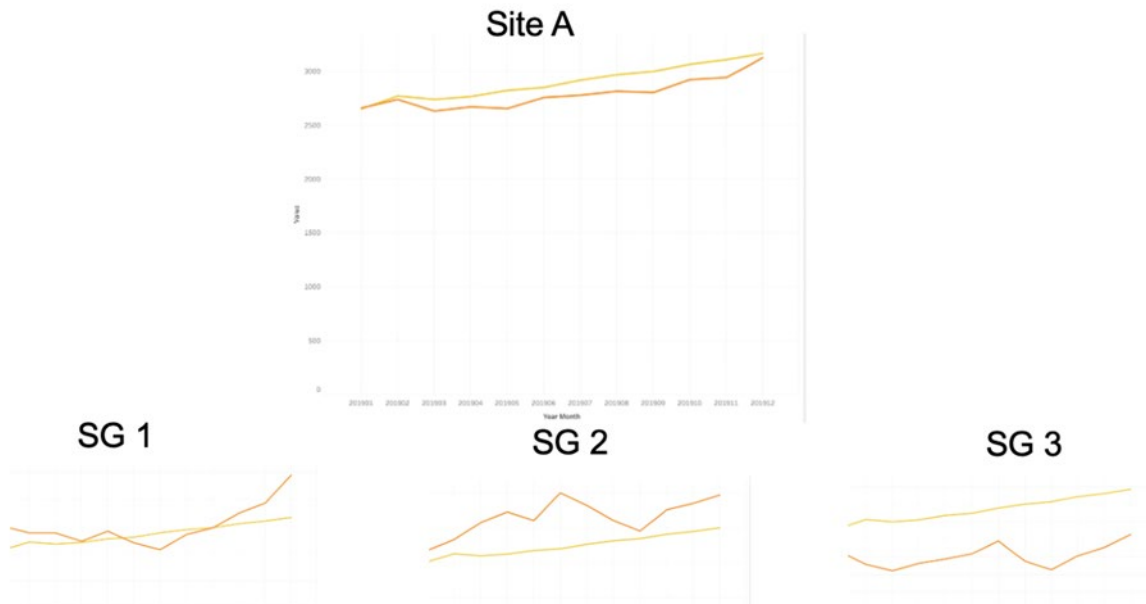


Figure 4 - Site to SG example

The site level forecast is very accurate, but when applied to the SGs underneath it the pattern no longer holds. This is because each individual SG has its own growth drivers and the smaller sample size leads to higher variance in the pattern.

This network simulation method has effectively helped with our long-range planning and budget needs for the next calendar year. The challenge comes as we get closer to

and into those budget years – how do we help our partners in the field find the right SGs to augment six to twelve months out? Our site-level-aggregate forecasts are only run at certain times of the year which

leads to stale data. Customers churn, and release events for gaming and streaming video occur, leading to site-level forecasts that do not translate to in-year planning for the individual SGs. As we focus on proactive network augments, how can we direct construction teams to the right areas before utilization exceeds capacity and impacts customers?

3. Model Evolution

There is currently a time-series solution at both the regional and site (headend) levels. These solutions use ARIMA (Autoregressive integrated moving average) and Exponential Smoothing models to produce results for each unique region and site. So, the first step in the SG modeling effort is to see if there is an application of what we are already familiar with at the SG level. Using R on Spark and later Databricks we built a solution that produces unique ARIMA and Exponential Smoothing Models for each SG. This model produces results for all SG regardless of age and the sMAPE (Symmetric mean absolute percentage error) for the DS traffic was improved.




Figure 5 - Time Series model results

This demonstrated improvement in the forecasting process, but the model still did not produce a scalable solution because: (1) the sMAPE nationally was not improved, (2) the model behaved very poorly for SGs with inconsistent trends, and (3) it took 28 hours to run on Spark. So, our first pass was not our final answer, but we did learn how to handle the massive data sets and gained an appreciation for variance in the observed data.



Figure 6 - Time Series model issues with irregular patterns

The next step was to explore more advanced machine learning models. Based on the deficiencies mentioned above, we needed a solution that could scale, handle high variance, and run in less than a day. We experimented with multiple options of ensemble methods and Neural Networks and landed on a Sequential Neural Network. We trained the neural network on 2 years of raw bps data per SG. The results included improvement in the sMAPE for downstream (DS), D3.1 and upstream (US) demand.

Average sMAPE by Month by Model




Figure 7 - sMAPE of NN on bps



Figure 8 - NN on bps example output

However, once again certain inconsistent trends, like SGs being collapsed and seasonal patterns did not perform well. Furthermore, the model could only be run on SGs with perfect data. This meant that we could not produce forecasts for SGs with data issues or seasonality, nor provide guidance on newer SGs including traffic over RPDs (Remote PHY Devices) on the new vCMTS locations that only have a few months of observed data.




Figure 9 - Where NN on bps does not forecast well.

Collapsing sites and seasonal locations did not produce valid models.

To produce a model that meets our needs we need improved accuracy, fast runtime, and solutions for all SGs regardless of age. The discussion then veered to isolating problematic SGs since they were the gap in the Neural Network solution. The largest class of problematic SGs is seasonal locations. That is, SGs serving a geography that is primarily a vacation area such as shore and mountain locations. Identifying an appropriate model for these SGs first requires identifying them at scale. We currently identify site level seasonality through a mix of time series analysis and exception lists developed over years of site level modeling. However, seasonal sites can have non-seasonal SGs, and we need to be able to capture that in an automated way.

We know a seasonal site has a ‘look’ when you are scanning visually - you can see the peaks and valleys. This pattern can be described as a monthly growth rate with high positive growth in the season ramp-up followed by large negative growth as the season ends. Using this approach, we can put all the SGs on the

same scale - size will not be visible. We can identify similar growth patterns and find seasonal locations. We can also more universally describe all SG growth regardless of size and pattern. This highlights customer behavior as opposed to size driven by engineering decisions. We can increase forecast accuracy by feeding all SG data into the model as a monthly growth rate. Since the scale of the input is the same, the neural network can focus on learning a pattern, not both size and pattern.

4. The Model

Zooming back out to the national view we can create monthly growth rates from the monthly 98th traffic for each SG/BG. To accurately predict seasonal trends, it is ideal to have 24 monthly growth rates. Additionally, we will need a target growth rate for our first prediction month. This means that we need 26 months of data to produce 25 growth rates. Many of our SGs have missing data so we can fill in those gaps with 1s. We also have the issue of hyper-growth or data driving high monthly increases. Although it is impractical to perfectly clean a huge data set, we can cap monthly growth and decline to anywhere between 15%-50%. These options seek to balance variance vs natural seasonality. The most recent month is used as the target and the previous 24 months are the features. We then train that on a sequential neural network with 4 layers. Twenty-four months of data predicts 1 month. To reach a full year, we must modify the model input to include recent forecasts as features in the trained model.

P_month1: X includes Growth Rate 1 - 24

P_month2: X includes Growth Rate 2-24, P_month1

P_month3: X includes Growth Rate 3-24, P_month1, P_month2

These growth rates are then restated as bps by referencing the most recent month's bps and multiplying by the following month's predicted growth rate.

Once the result is stated as bps, sMAPE is used to guide model tuning and showed that breaking out the model by geography (North, South, East, West) and direction (downstream, upstream, OFDM) produced the most accurate result. Hyperparameter tuning used sMAPE and Loss to select layer count, optimization functions and epoch count to optimize the model.

The resulting model predicts 12 months of traffic for each SG and BG regardless of age.



Figure 10 - NN on growth rate results

To determine if this is a superior model to our previous SG model iterations or the legacy solution, we once again look at the sMAPE which shows that the NN produces the lowest error of any model for both the US and DS. So, this NN solution produces a more accurate model for all SG/BG regardless of age and can be run in a few hours. We can now move to the step of implementing these results in network planning models to inform future budgeting.




Figure 11. sMAPE graph comparing all models.

Lower is better for sMAPE. Across the time forecasted the NN using growth rate (green) outperforms all other models.

5. Future State

We now have a working time series model. Our immediate next step is to insert this model into the network planning process and facilitate decision-making from its output. We will have to slightly re-factor the network planning models as well as set up a quality control process. The quality control process will have multiple visualizations in 2 groups. The first group will be based on model accuracy to

show how close we are in aggregate and compare model options to target the lowest error rate. The second group will highlight SGs with high variance and questionable forecasts.

Once we have a working model, we will continue to improve the neural network. We are striving to get more accurate with seasonal sites by incorporating Holt Winters seasonality in those specific locations. We are also looking at options to add other data into the model including geography, node type, customer mix, device type, and SG age. We will measure these additions by the level that sMAPE and MAPE improve when compared with the historical time series.

Additional improvements to the algorithm have already been identified. Currently, we are producing a point estimate for each SG for each month. In the future we can produce a range of values based on the probability of scenarios. Additionally, we can use this information to inform probabilistic risk of needing a split for each SG within a period. Being able to define algorithms and probabilities at this level also allows us to build scenarios using optimization modeling techniques. Improved forecast accuracy, probability scenarios and split risk can be extended to build cost functions for needing a split. We can then enter the next frontier of establishing the cost of NOT splitting by adding customer experience metrics like speed tests, latency, and packet loss. At which point we will have a set of ML driven functions defining risk of split/non-split which can be solved using a Linear Program.

6. Conclusion

Growth in internet demand is currently driven by residential gaming, streaming video, and business applications. To ensure the highest customer experience we need to forecast network traffic to ensure we maintain sufficient capacity to meet user demand. In the past national, regional and site level forecasts were sufficient and development cycles were dedicated to them. Now, in the digital node world we need more accurate and granular forecasting to the neighborhood level, using ML to predict trends at the SG level.

Our team has developed an ML solution that has demonstrated improvement in time series accuracy. Additionally, it can handle new SGs regardless of size because it relies on monthly growth rates instead of bps to predict demand. Currently we are working closely with our internal partners to add this level of detail to the network planning process. This includes visualizing results to both aid model selection and highlight SGs that need further analysis.

A working model that aids financial budgeting and long-range planning decisions will mark the completion of our first phase. Next, we need to enhance the accuracy of our forecasting models using seasonal models and include additional data about the SG's behavior. This will allow us to explore other modeling techniques to cluster based on device, geography, or customer behavior. A further goal is to incorporate optimization modeling techniques that build on our ML models to maximize both the customer experience and financial performance.

We are in the beginning phases of an evolution in forecasting and network planning. Our combination of ML, time series analysis and collaboration will ensure our success at accurately predicting network demand so that all games can continue to be streamed and homework can (eventually) be done.

Abbreviations

SG	Service Group
BG	Bonding Group
ML	Machine Learning
DS	Downstream
US	Upstream
bps	bits per second
DOCSIS	Data over cable service interface specification
UDM	Unified demand model
ARIMA	Autoregressive integrated moving average
sMAPE	Symmetric mean absolute percentage error
vCMTS	Virtual cable modem termination system
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1

Gaming Latency vs. Engagement and Potential Impacts of Lower-Latency DOCSIS & Wi-Fi

A Technical Paper prepared for SCTE by CableLabs

Stephen Froehlich

Sr. Strategist

CableLabs

858 Coal Creek Cr. / Louisville, CO 80027

+1 303 661 3708

S.Froehlich@cablelabs.com

Jacob Malone

Dir & Principal Strategist

CableLabs

858 Coal Creek Cr. / Louisville, CO 80027

+1 303 661 3775

J.Malone@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Data Source & Methodology	4
2.1. Data Source	4
2.2. Data Limitations	4
2.3. Analysis Methodology	5
3. Interaction of Latency and Game Engagement	7
3.1. Latency Overview.....	7
3.2. Relationship Between Active Unique Days & Latency.....	7
4. Wired Latency	9
4.1. Overall Wired Latency.....	10
4.2. Routing Delay.....	11
4.3. Queueing Delay	12
5. Wi-Fi Queueing Delays	16
5.1. Wi-Fi Share	16
5.2. Wi-Fi Delay.....	17
5.3. Performance by Operator & Gaming Platform + WMM Packet Marking	18
6. Conclusion.....	22
Acknowledgements	22
Abbreviations	22
Bibliography & References.....	23
Appendices	23
Appendix 1 – Regression of Unique Active Days vs. Latency & Operator.....	23

List of Figures

Title	Page Number
Figure 1 – Per-User Latency by Measurement Quantile	7
Figure 2 – Latency vs. Unique Active Days by Measurement Quantile.....	8
Figure 3 – Correlation of Latency by Quantile to Unique Active Days.....	9
Figure 4 – 99%ile Wired Latency by Operator.....	10
Figure 5 – 99%ile Wired Latency U.S. Map.....	11
Figure 6 – Routing Delay by Operator	12
Figure 7 - % of Latency due to Queueing	13
Figure 8 – Queueing Delay Dist. by Operator.....	14
Figure 9 – Queueing Delay Dist – LLD Counterfactual.....	15
Figure 10 – Est. 99%ile Latency w/ Full LLD Deployment.....	16
Figure 11 – Wi-Fi Share	17
Figure 12 – Wi-Fi Delay by Operator	18
Figure 13 – Wi-Fi Performance by Operator and Gaming Platform.....	19

Title	Page Number
Table 1 - Distributional Comparison of latency _{NN} + jitter _{NN} to Raw Latency, Duration: 700 Seconds.....	6
Table 2 - Regression of Wi-Fi log(latency) vs. ISP & Gaming Platform	21
Table 3 - Regression Output of Unique Active Days per user vs. Z-score of Latency & Operator by Measurement Quantile	24

1. Introduction

As the industry gives monitoring and optimizing latency more attention, there is an increasing interest in understanding the value created by that network investment. Gamers are often mentioned as a key user segment that would value such investment because of the improvements to their user experience when playing games online.

Using a dataset from Network Next that includes 10-second latency data for every session of a multiplayer, competitive game during the first 2 months after its launch, we estimate that a one standard deviation improvement in 99th percentile (abbreviated as “99%ile” for the remainder of the paper) latency results in an 8% increase in the number of unique days of game play on average (approximately 0.4 extra unique days).

In addition, when a 7ms ceiling is applied to users’ queuing delay, to estimate the impact of Low Latency DOCSIS[®] technology, we find that some cable operators could see up to a 50ms improvement in 99%ile latency, while the improvement for other operators is smaller due to variation in queuing delay across cable operators. However, the user experience of gamers on all cable broadband networks are likely to benefit from the greater consistency in observed latency.

Lastly, the Network Next data allows analysis of Wi-Fi performance by gaming platform (PS4, Xbox One, Switch, Windows) on a per operator basis. Wi-Fi connections are used by 68% of the players of this game. Given that Wi-Fi is one of the largest sources of latency on the network, any complete solution to improving latency must address home Wi-Fi networks.

One potential technology to improve Wi-Fi latency is Wi-Fi Multimedia (WMM) tagging. For this game, the Nintendo Switch uses WMM. We find that for the Nintendo Switch, Wi-Fi latency is 46% lower than that of a PlayStation 4 (PS4), a console that does not use WMM. We also find there is substantial variation in Wi-Fi performance across operators.

2. Data Source & Methodology

2.1. Data Source

For the first two months following the launch of a game in the summer of 2021, Network Next provided CableLabs latency observations for each 10-seconds of activity. Each observation included the minimum latency observed during that 10-second window, jitter (defined as the standard deviation from the minimum latency observed during that 10-second window), and packet loss. Each 10-second observation is a summary of 100 pings.

The metadata include the user ID, a session ID, a postal code for the user and for the server, a timestamp, the gaming platform, whether the session was over the wired or wireless interface on the device, and the ISP name. The global data set includes 6.4 million players of the game, playing 260 million sessions, and 9.8 billion valid 10-second observations. When filtered down to the top 14 operators in the lower 48 states of the U.S., the data set includes 1.6 million players of this game who played for at least an hour.

2.2. Data Limitations

The Network Next data set comes with two notable limitations. First, the relationship between engagement and latency that we observe are not externally valid for other games or applications. This is because the sample only contains data on a single game, which limits how generalizable our results are,

and it has a unique relationship between latency and engagement, which is discussed in Section 3.2. In short, maximum engagement for this game peaks where there are more players (i.e., the latency experienced by *most* players), not at the *minimum* latency observed; this relationship creates an upside down “U” shape that is shown below. Network Next says they have verified this result but say it is unique to this game based on their experience studying other games. Other games tend to have flat engagement out to approximately 100ms of latency, and then it decreases. However, more generally, we hypothesize that the relationship between lower latency and engagement are related to a game’s characteristics. For example, players of real-time puzzle games may value latency improvements differently than those who play fighting games.

Second, some believe (Broadband Internet Technical Advisory Group, 2022) that online gaming is sensitive to the 99th percentile packet latency. This sample only provides minimum latency and standard deviation every 10 seconds. By comparison, we try to estimate the relationship between the measured minimum latency in the Network Next data to packet latency percentile, using empirical data collected at CableLabs, but this relationship is imprecise. Also, as shown in Table 1, the relationship between Network Next and actual latency measurement quantiles includes a couple of step functions that could easily move around based on network conditions, adding uncertainty to the meaning of the lower quantiles.

2.3. Analysis Methodology

We considered four different variables to proxy for game engagement:

1. Total game time played
2. Unique active days
3. Time between first and last session
4. Average time per Unique Active Day

Each of these four variables correlate well with latency, but unique active days was chosen because it had the strongest correlations, consistent region-to-region, and closely relates to how Network Next tracks retention. Using unique active days also mutes the effect of differences in free time between gamers. Finally, we propose, without empirical evidence or market research to confirm, that choosing to play a game again on a different day likely speaks better to a gamer’s enjoyment from playing more so than a longer playing session on a single day.

Based on modelling against lab data and input from Network Next, we find *latency + jitter* for each 10-second observation (summary of 100 pings) is the best stand-in for actual latency. For each user, the primary input variable used is their quantiles of 10-second (*latency + jitter*) observations.

As discussed above, the Network Next measurement system is based on 10-second windows, where 100 “ping” packets are sent and summarized in terms of minimum measured round-trip time and the root mean square deviation from that minimum. For the round-trip time (*rtt*) of the 100 pings in each 10-second observation:

$$latency_{NN} = \min(rtt)$$

$$jitter_{NN} = \sqrt{\sum \frac{(rtt - \min(rtt))^2}{100}}$$

To estimate the relationship between a Network Next's $latency_{NN} + jitter_{NN}$ measurements (i.e., individual summaries of a 10-second window ping) and the raw latency distribution over a period of observation, we measured raw per-packet latency for a cable modem connection carrying a mix of traffic for 700 seconds and calculated the $latency_{NN} + jitter_{NN}$ on 10-second windows. We then compare percentiles of the $latency_{NN} + jitter_{NN}$ distribution to percentiles of the raw latency distribution. The results are summarized in Table 1 below. For example, we find that the 99th percentile $latency_{NN} + jitter_{NN}$ value is roughly equivalent to the 96th percentile raw latency value.

Table 1 - Distributional Comparison of $latency_{NN} + jitter_{NN}$ to Raw Latency, Duration: 700 Seconds

Network Next Quantile	4.0%	14.0%	30.0%	45.0%	56.0%	75.0%	90.0%	95.0%	99.0%
Est. CL Latency Measurement WG Quantile	25.0%	52.3%	57.1%	78.1%	90.1%	94.0%	95.4%	95.7%	96.1%

Regression analysis is used in Section 3 to estimate correlations between latency and game engagement. As described earlier in the first part of Section 2.2 and more completely below in Section 3.2, this game has a unique upside down “U” shape in the relationship between latency and engagement. The upward slope approaching the peak (i.e., the portion of the curve that describes how the lowest latencies in the sample relate to engagement) is filtered out, so only data from the peak and after is used in the regressions. More exactly, mean engagement is calculated for gamers (minimum 10 gamers) by operator and state and the maximum value is found (L_{pe}). User data between L_{pe} and $3*L_{pe}$ are used in the regressions. All regressions include operator fixed effects to account for any time invariant operator-specific variance in the data.

To determine wired queuing delay, the mean routing delay (r) was determined for each user connection to a given server from a given location over a given ISP. R is estimated by the mean of the 10-second minimum latency for each combination of user, game server, and ISP. Then queuing delay for each 10-second observation was calculated as $latency + jitter - r$. Reported quantiles in Section 4 are of this distribution.

Cleanly isolating Wi-Fi delay was difficult. For a single user, differentiating jitter due to wired queuing from jitter due to Wi-Fi queuing is impossible. Therefore, we matched pairs of wired and wireless users from the same postal code connecting to the same server from the same ISP at the same time using the same gaming platform. As there were only two wired Switch players in the entire data set, Switch wireless gamers were matched with Windows wired (Dowle & Srinivasan, 2021) gamers. This resulted in a set of 42,252 user pairs with at least an hour of overlapping activity. The latency ($latency + jitter$) of the wired user was subtracted from the latency of the wireless user to get an estimate of latency from Wi-Fi for the wireless user. After that, quantiles of Wi-Fi latency ($latency + jitter$) were calculated on a per-wireless-user basis.

An exponential regression of ISP and gaming platform dummy variables was performed to see the effects of different ISPs, gaming platforms, and the interactions between the two on Wi-Fi latency.

3. Interaction of Latency and Game Engagement

In this section, the focus is on latency and game engagement. In particular, how different latency percentiles compare to each other and whether the relationships with engagement are significant. Overviews of latency and its relationship with game engagement are first, followed by the regression analysis.

3.1. Latency Overview

Figure 1 shows the cumulative distribution of latency per user measured at different quantiles (note the x-axis is log scale). While quantiles 75% or below are generally under 70ms, latency grows quickly at higher quantiles, especially the 99%ile.

Per-User Latency by Measurement Quantile, Cumulative Distributions




Figure 1 – Per-User Latency by Measurement Quantile

3.2. Relationship Between Active Unique Days & Latency

As shown in Figure 2, the game used in this study has a unique response curve between latency and unique active days playing the game. The peak unique number of active days is wherever most players are and is near the average latency for each operator, causing occasional double peaks or at least visible “shoulders”. The shoulders in the 56%ile, 75%ile, 95%ile, and 99%ile curves are caused by the presence of an operator whose mass of latency is at that point. The plot uses a logarithmic x-axis to emphasize the slope of each curve.

According to Network Next, other games have flat engagement to about 100ms and then it falls off. However, in this game, those with better-than-average 95%ile latency play significantly fewer days than

those with average latency. While the cause remains unknown, Network Next is working with the game developer to test several hypotheses.

As you can see from the plot, the steepest falloff is in the 56%ile and the 75%ile, and this will be reflected in the regressions. However, given the step-function nature of the correlations found in Table 1 this could be an artifact of the measurement system instead of actual engagement. Nevertheless, all measurement quantiles seen here have a significant correlation to gamer engagement both above and below the peak engagement.

Latency vs. Unique Active Days

Gamers with less than 200 ms median latency + jitter




Figure 2 – Latency vs. Unique Active Days by Measurement Quantile

Following the regression analysis methodology outlined above, we estimated three different regressions to better understand different parts of the latency-engagement relationship. In all of the regressions, unique number of active days is the y-variable. Our x-variable includes three different latency calculations. The first is latency in 10ms increments, the second is median latency, and the third is user z-scores. The results of these three regressions are summarized below in Figure 3, where each panel is a different regression. Each dot represents a estimate coefficient in the regression on the latency variable. Note that the negative correlations in Figure 3 are inversely related to engagement: the more negative a coefficient is, the more engagement there is. The primary findings of these regressions is the following.

1. The positive slope in the top panel is misleading. While we do find that improvements in lower latency quantiles is larger, the next two regressions add perspective. That is, 99%tile latency does matter. However, the key result is intuitive. Median latency, for example, is experienced much more frequently by a gamer than 99%tile latency.
2. When engagement is measured relatively for each user, where their median value for each percentile is used, we find the reduction in engagement from worse latency measurements across

percentiles flat until the 99%tile. This suggests the importance of higher percentiles is more important than the first panel suggests.

3. Finally, a z-score is a useful way to summarize how extreme a value is relative to a mean using the distributions standard deviation. In the third panel, the results suggest that extremely large latencies at all percentiles similarly impact engagement, even out to the 99%tile.

Based on these results, we will focus on 99%ile measurements for the remainder of the paper.

Change in Unique Active Days ...

Correcting for Operator-level Effects




Figure 3 – Correlation of Latency by Quantile to Unique Active Days

4. Wired Latency

This section begins by quantifying the latency distribution of wired gamers and separating that latency into routing latency and 99%ile queuing latency. Next, improvements to 99%ile queuing latency are

examined and with estimates of what improvements to wired latency might be expected from a full-scale deployment of Low Latency DOCSIS technologies.

4.1. Overall Wired Latency

This section uses PS4, Xbox, and Windows players with a wired connection to their home gateway.

Total 99%ile latency L is separated into a routing delay (R) and a queuing delay (Q).

$$R_{99\%ile} + Q_{99\%ile} = L_{99\%ile}$$

The 99%ile queuing delay (Q) is the 99%ile latency (L) minus an estimate of routing delay (R).

$$Q_{99\%ile} = Quantile_{99\%}(L - R)$$

As depicted in Figure 4, with one exception, cable and telco latencies were generally comparable during May & June of 2021. 99%ile latency generally ranges from 30 to 200ms for cable operators with a somewhat wider range for most telcos.

99%ile Wired Latency by Operator, Cumulative Distributions




Figure 4 – 99%ile Wired Latency by Operator

Figure 5 is a map of latency within the lower 48 states of the U.S. It shows 99%ile wired latency by location. The sample is biased away from dense urban centers to avoid overplotting, so it is not representative. The line from each user to a game server represents the game server most often used by that gamer while each gamer's icon is colored according to his or her latency.

Two things stand out on this map:

1. Some gamers primarily use a server other than their nearest one. In fact, outside of a few states such as Illinois, Michigan, Georgia, and Florida this seems to be the norm rather than the exception. This result suggests that game matching algorithms are using more criteria than just the closest server location (at least, sometimes as shown by this game). *(Note: Lines that appear to go from one server location to another such as from Chicago to Omaha are an illusion and are virtually passing through “Chicagoland” from users in Northern New England.)*
2. The worst latency performers are generally telco customers in rural locations (likely DSL) such as the Mississippi Delta, east Texas, eastern Arizona, and central Nevada.

99%ile Latency per Wired User

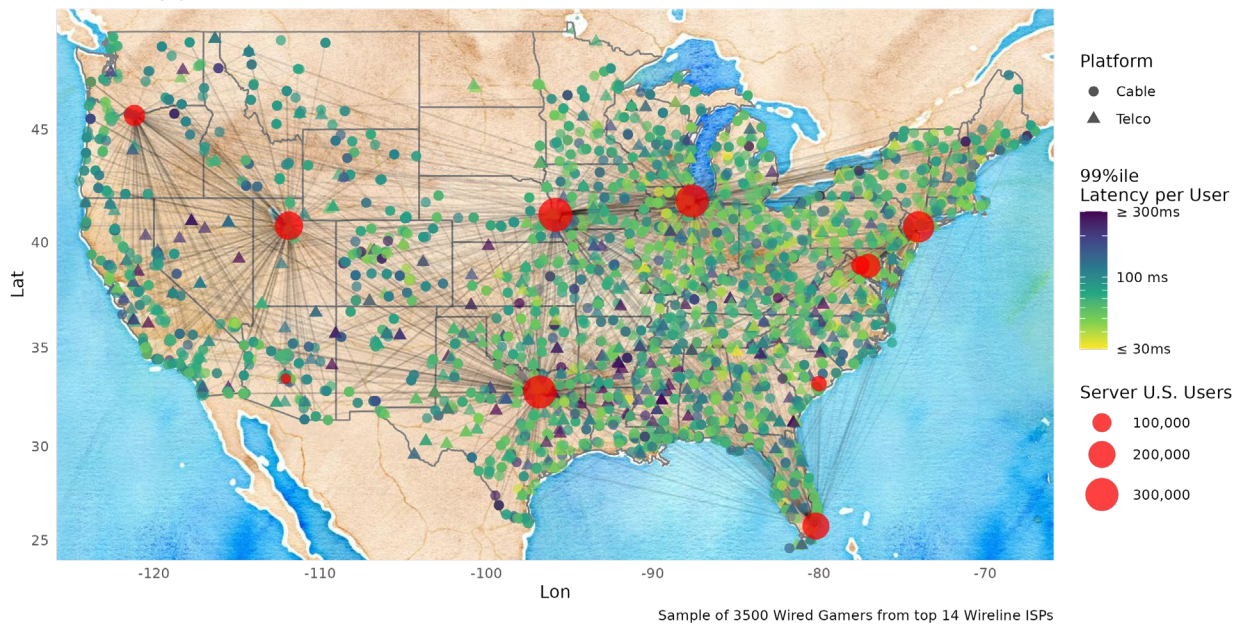


Figure 5 – 99%ile Wired Latency U.S. Map

4.2. Routing Delay

This section is presented to help readers compare the size and shape of routing delay distributions to the wired queuing delay distribution in Figure 8, the Low Latency DOCSIS counterfactual in Figure 10, and the Wi-Fi queuing delay in Figure 12 & Figure 13.

In general, routing delay ranges from 10 to 100ms in the lower 48 states of the U.S., and as the map in Figure 5 implies, is primarily determined by geographic proximity to the relevant game server for that operator’s user population.

Routing Delay by Operator, Cumulative Distributions




Figure 6 – Routing Delay by Operator

4.3. Queueing Delay

While routing structures are unique to each cable operator, the access network and the queueing delays that happen there are more a function of DOCSIS technology. Therefore, queueing delay and the potential impact of Low Latency DOCSIS technology as a way to improve queueing delays is of high interest. This section will examine how much of an impact wired network queueing delays (assumed to be primarily access network queueing delays) typically have and then the impact Low Latency DOCSIS could have on them.

To give an idea of the relative scale of routing and queueing delay for different operators: For cable operators, 25% of gamers in the study had at least 45% of their total 99%ile latency come from transient, assumably queueing-based delays. For telcos, the top 25% starts at 34% of 99%ile latency coming from queueing. That is, for these gamers, queueing delay is a larger share of overall latency on cable networks compared to telco. The results are summarized on a per operator basis in Figure 7.

% of 99%ile Latency due to Queuing Top 10 Operators by Wired Gamer Count




Figure 7 - % of Latency due to Queueing

Figure 8 shows the cumulative distributions of queuing delay for each of the top 10 U.S. wireline Internet service providers and provides some guidance on what Low Latency DOCSIS[®] (LLD) should be able to provide. The queuing delay experienced by users typically ranges from 2ms to more than 200ms. In general, the best telco customers experience less than 3ms queuing delay while the best cable customers experience less than 11ms. However, with one exception, all the telcos fall behind the cable curves in the top quartile. The 6 largest cable broadband operators fall into three distinct groups with a singular queuing delay leader, a close pack of three in the middle, and then two others.

99%ile Queuing Latency by Operator, Cumulative Distributions




Figure 8 – Queuing Delay Dist. by Operator

The expectation of Low Latency DOCSIS is that it should make cable operators queuing delay much more competitive to telcos. Without the Proactive Grant Service (PGS), LLD should be able to deliver queuing delays reliably 7ms or below. With PGS, that falls to 2ms but at the cost of allocating some upstream bandwidth to the proactive grants. This is reflected in the annotations in Figure 8.

Taking that a step further, Figure 9 is a *simplistic* imagining of what might happen if the top 6 cable broadband providers in the U.S. enabled LLD across their entire installed base. Suddenly the queuing delay cumulative distributions would appear as they do in Figure 9, resulting in an average 38% improvement in 99%ile latency for cable broadband customers.

Low Latency DOCSIS Counterfactual

99%ile Queuing Latency by Operator, Cumulative Distributions




Figure 9 – Queuing Delay Dist – LLD Counterfactual

Finally, Figure 10 estimates what the overall 99%ile latency curves would look like for the top 6 cable operators if LLD were fully enabled on each of their networks. As you can see, the results vary wildly with the current queuing delay leader (top left) making relatively modest gains but the 99%ile queuing delay of other operators improving by 50ms or more for a significant fraction of their customers. However, the improved consistency in latency from lower jitter is understated by this figure. On today's network, latency + jitter can exceed 100ms with too much regularity to support the requirements of future real-time applications and user experiences.

Cable Low Latency DOCSIS Counterfactual

Est. 99%ile Cumulative Latency Distributions w/ LLD by Operator




Figure 10 – Est. 99%ile Latency w/ Full LLD Deployment

5. Wi-Fi Queueing Delays

5.1. Wi-Fi Share

For this game, 68% of players connect via Wi-Fi and 32% on wired. A full breakout of gaming platforms can be found below in Figure 11. Xbox One is the most popular Wi-Fi platform, followed by the PS4, then the Switch, and last is Windows. Wi-Fi delays can be larger than all wired delays combined. Therefore, improving Wi-Fi must be considered in any end-to-end network latency plans.

Share of Game Players by Platform & Connection




Figure 11 – Wi-Fi Share

5.2. Wi-Fi Delay

For the 42,252 Wi-Fi gamers where we could find a wired match with at least one hour of overlap, their 99%ile Wi-Fi delay ranges all the way from 4ms to near 300ms. Most operators are grouped together, but two telco operators stand out for the worst 45% of their matched gamers. Each operator in this plot has at least 400 matched gamers.

Wi-Fi Delay by Operator, Cumulative Distributions




Figure 12 – Wi-Fi Delay by Operator

5.3. Performance by Operator & Gaming Platform + WMM Packet Marking

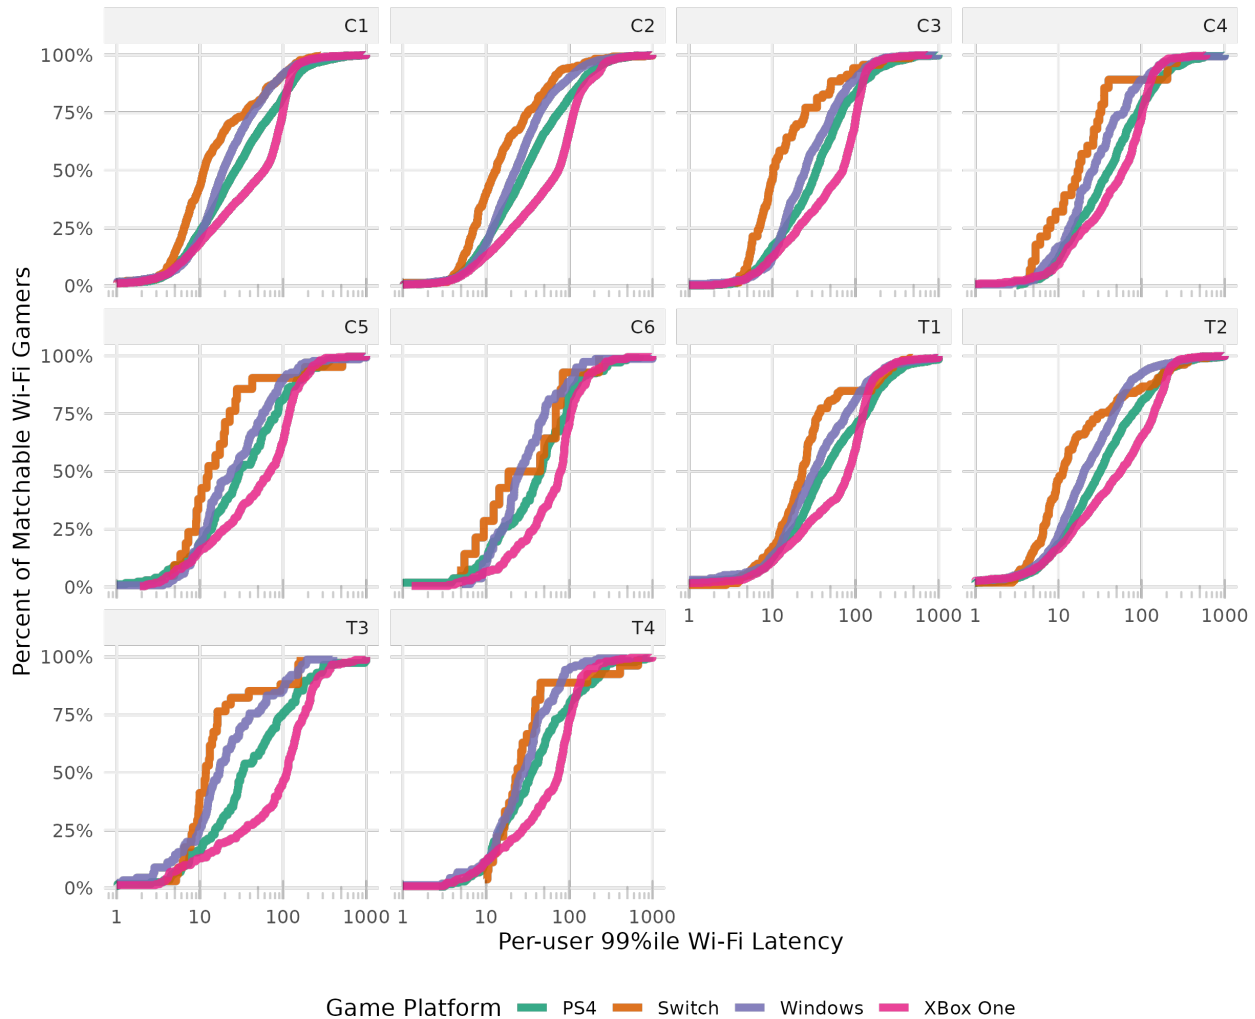
One unique feature of this data set is that Network Next's Wi-Fi stack included WMM packet marking on the Nintendo Switch and Windows Wi-Fi clients while the network stacks on the Xbox One and the PS4 did not. Figure 13 provides insights into the impact of this and other differences in the Wi-Fi implementation on different devices on 99thile Wi-Fi latency performance. This gives us a chance to compare the performance of each of these network stacks as they interact with the home gateways used by customers of different Internet service providers.

In general, the Nintendo Switch has the best Wi-Fi performance, followed by MS Windows clients, followed by PS4 clients, with Xbox having the worst performance.

Other things to note:

- For T1, T2, and C4, it appears that there are a few gateways that have noticeably worse performance with the Switch than other gateways used by customers of the same operator. (Step function in the middle of the figure.)
- In general cable operators are outperforming the telcos in Wi-Fi delay, but not by much.

Wi-Fi Delay by Operator & Game Platform, Cumulative Distributions Switch & Windows use WMM Packet Marking; Consoles Do Not



Top 10 U.S. Wireline Broadband Operators by Matchable Gamer Count

Figure 13 – Wi-Fi Performance by Operator and Gaming Platform

Finally, referring to Table 2, we run a dummy variable regression of the top 10 ISPs interacting with the four gaming platforms against log of latency. The base case in this dummy variable regression is ISP C1 + PS4.

While the other measurement quantiles are presented in Table 2, the 99thile alone will be discussed:

- The geometric mean intercept is 28ms, and each ISP except C5 results in a statistically significant increase of latency over C1 (who is the Wi-Fi latency leader) ranging from a small 11% increase for C2 to a 61% increase for T1.
- The Nintendo Switch has 46% lower latency than the PS4 and Windows Wi-Fi has a 23% improvement, meaning that WMM packet marking has a significant and large effect on Wi-Fi latency compared to the unmarked PS4 packets (and even more so compared to the Xbox One). We note, however, that WMM works well on packets like those studied in this sample (i.e.,

gaming), but may not be sufficient for other types of traffic (e.g., real-time video). Therefore, other technologies like L4S AQM implemented over Wi-Fi are things to be considered.

- The Xbox One Wi-Fi implementation results in a 34% worse latency over the PS4
- ISP C2 has a unique issue with the Xbox One (it does quite well otherwise), seeing a 21% increase in latency compared to other operators serving the Xbox One. ISP T3 seems to have a similar problem.

Table 2 - Regression of Wi-Fi log(latency) vs. ISP & Gaming Platform

	4%ile	14%ile	30%ile	45%ile	56%ile	75%ile	90%ile	95%ile	99%ile	Max
Geometric Mean (Intercept - ISP C1 + PS4)	1.43 ms ***	1.83 ms ***	2.41 ms ***	3.03 ms ***	3.6 ms ***	5.16 ms ***	8.5 ms ***	12.4 ms ***	27.71 ms ***	106.71 ms ***
ISP C2	-17% ***	-21% ***	-20% ***	-17% ***	-13% ***	-6% **	3%	5%	11% ***	6% *
ISP C3	1%	7% *	11% **	14% ***	16% ***	19% ***	22% ***	25% ***	20% ***	0%
ISP C4	16% ***	24% ***	26% ***	30% ***	33% ***	40% ***	47% ***	47% ***	43% ***	21% **
ISP C5	-13% **	-14% **	-11%	-8%	-3%	2%	10%	10%	14%	1%
ISP C6	-1%	7%	9%	7%	8%	14%	25% *	31% *	40% **	32% *
ISP T1	26% ***	33% ***	35% ***	36% ***	37% ***	38% ***	45% ***	49% ***	61% ***	51% ***
ISP T2	13% ***	9% ***	6% *	5%	5%	5%	5%	5%	13% **	16% **
ISP T3	1%	10%	13%	9%	8%	5%	10%	14%	37% **	73% ***
ISP T4	1%	6%	18% **	20% **	20% **	23% **	38% ***	36% ***	36% **	31% **
Switch	-3%	-3%	-6%	-8%	-10% *	-15% **	-25% ***	-35% ***	-46% ***	-48% ***
Windows	-11% ***	-16% ***	-19% ***	-21% ***	-23% ***	-24% ***	-27% ***	-29% ***	-23% ***	-31% ***
XBox One	-9% ***	-13% ***	-16% ***	-18% ***	-19% ***	-21% ***	-22% ***	-20% ***	34% ***	0%
ISP C2 + Switch	-1%	-3%	5%	4%	2%	-4%	-7%	-3%	-4%	-6%
ISP C3 + Switch	-16% *	-21% *	-22% *	-25% *	-25% *	-29% *	-27% *	-26%	-17%	-16%
ISP C4 + Switch	-16%	-27% *	-27%	-30% *	-32% *	-34% *	-32%	-21%	-13%	-13%
ISP C5 + Switch	31% *	29%	43%	40%	30%	23%	9%	4%	-1%	-17%
ISP C6 + Switch	16%	18%	32%	35%	32%	28%	18%	23%	28%	-26%
ISP T1 + Switch	-12% *	-10%	-8%	-11%	-12%	-11%	-1%	9%	9%	-9%
ISP T2 + Switch	-22% ***	-27% ***	-24% **	-21% *	-19% *	-16%	-11%	-3%	-4%	11%
ISP T3 + Switch	3%	-4%	-3%	1%	1%	-3%	-17%	-18%	-23%	-51% **
ISP T4 + Switch	22%	18%	4%	10%	14%	12%	12%	19%	46%	48%
ISP C2 + Windows	7% **	10% ***	9% **	8% *	7% *	5%	4%	5%	1%	8%
ISP C3 + Windows	-5%	-8%	-8%	-6%	-4%	0%	4%	6%	5%	15%
ISP C4 + Windows	-13% *	-19% **	-24% ***	-29% ***	-30% ***	-30% ***	-27% **	-24% *	-10%	28%
ISP C5 + Windows	6%	7%	5%	2%	0%	1%	0%	8%	9%	8%
ISP C6 + Windows	0%	-6%	-5%	1%	0%	2%	4%	-3%	-1%	-13%
ISP T1 + Windows	-3%	-1%	4%	7%	8%	8%	11%	13%	-4%	-4%
ISP T2 + Windows	-4%	-4%	-8%	-10% *	-11% *	-10%	3%	3%	-11%	-11%
ISP T3 + Windows	1%	-11%	-18%	-20%	-23%	-27% *	-25%	-16%	-34% *	-32% *
ISP T4 + Windows	7%	11%	-1%	-3%	-1%	-2%	-10%	-8%	-9%	-13%
ISP C2 + XBox One	8% ***	10% ***	12% ***	13% ***	13% ***	15% ***	20% ***	25% ***	21% ***	32% ***
ISP C3 + XBox One	-3%	-5%	-5%	-6%	-5%	-3%	-1%	0%	2%	14%
ISP C4 + XBox One	-11% **	-14% **	-15% **	-15% **	-17% **	-20% **	-18% **	-11%	-16% *	-11%
ISP C5 + XBox One	9%	11%	11%	12%	10%	18%	19%	20%	12%	35% *
ISP C6 + XBox One	-1%	0%	2%	9%	11%	16%	27%	32% *	16%	3%
ISP T1 + XBox One	-1%	1%	1%	-1%	-2%	-2%	1%	2%	-8%	-2%
ISP T2 + XBox One	1%	3%	3%	4%	4%	8%	13% *	12% *	0%	12%
ISP T3 + XBox One	0%	-10%	-12%	-5%	1%	18%	39% **	69% ***	40% *	9%
ISP T4 + XBox One	21% ***	24% **	17% *	18%	19%	21%	16%	18%	0%	10%
Adj. r ²	4%	4.10%	3.50%	2.90%	2.60%	2.20%	2.30%	2.60%	6%	3.40%

As is typical * = p < 5%, ** = p < 1%, *** = p < 0.1%.

6. Conclusion

There are several opportunities for cable operators to make more enjoyable gaming experiences for their customers by reducing latency.

The first priority is to be able to understand how to measure latency in a way that correlates well to a gamer's enjoyment of this game, and the takeaway is that all of the measurement quantiles up to and including the 99%ile are well correlated to users' engagement in this game with an across-the-board one standard deviation improvement in latency likely resulting in between 0.4 and 1 (7% - 18%) more unique active days played of this game.

The second priority and biggest opportunity for differentiation compared to telcos is Wi-Fi. While there is hope of substantial improvements with Wi-Fi 6E, which will feature new (greenfield) 6 GHz channels where all the clients and base stations will include major latency-improving features such as OFDMA. However, the state of the game console cycle combined with the fact that the honeymoon period where the 6 GHz channels are unpopulated will come to an end leads us to emphasize that improvements here will require full ecosystem collaboration. For most cable operators, broader WMM adoption could result in halving the Wi-Fi lag for the Xbox One and the PS4.

Low Latency DOCSIS will have a significant effect on wired queuing delays measured at the 99%ile with some operators able to save many of their customers up to 50ms 99%ile round trip time.

Finally, routing latency should not be neglected. While it is beyond the scope of this research, there are indications in the data that one telco has been able to achieve results in routing latency vs. distance-to-server over and above the other ISPs, and there is therefore a real opportunity to bring this number down.

Acknowledgements

We would like to thank Greg White of CableLabs for his immense assistance modelling the 10-second summarization technique used by Network Next on lab latency data he has and using that model to help us adjust the quantile range we were looking at. This paper would be far harder to understand without his help.

We would also like to thank Glenn Fielder, Tapan Desai, and Andrew Baumbach of Network Next for sharing several of their current hypotheses and their excellent support of this data set.

Abbreviations

DOCSIS [®]	Data Over Cable Service Interface
DSL	digital subscriber line
LLD	low latency DOCSIS
ms	millisecond
NN	Network Next
%ile	percentile
PGS	proactive grant service
PS4	PlayStation 4
rms	root mean square

StDev	standard deviation
WMM	Wi-Fi multimedia

Bibliography & References

- Broadband Internet Technical Advisory Group. (2022, 01 10). *Latency Explained*. Retrieved from BITAG: <https://bitag.org/latency-explained.php>
- Dowle, M., & Srinivasan, A. (2021). data.table: Extension of `data.frame`. Retrieved from <https://CRAN.R-project.org/package=data.table>
- Network Next. (2020, 10 22). *Introducing the Network Next Analytics Portal*. Retrieved from Network Next: <https://www.networknext.com/post/introducing-the-network-next-analytics-portal>
- R Core Team. (2022). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- Subspace. (2021, 05 03). *How to Ensure Your Matchmaking Experience for Multiplayer Games is Fast and Full*. Retrieved from Subspace: <https://subspace.com/resources/multiplayer-matchmaking>

Appendices

Appendix 1 – Regression of Unique Active Days vs. Latency & Operator

Table 3 includes a summary of the full regression coefficients and significance levels for each measurement quantile. (The regression is run separately at each measurement quantile.) The highest r^2 occurs at the 75%ile, but the strongest effect for improvement (such as by deployment of Low Latency DOCSIS[®]) occurs in the 99%ile measurements, so that is what is discussed here.

The inclusion of operator effects (that some operators have more engaged gamers than others) was necessary to get a clean correlation for the main variable.

Notably:

- The average player on operator C1 plays 5.1 unique active days of this game.
- For every standard deviation of increase in 99%ile latency, a gamer plays 0.4 unique days less.
- ISPs C7 and T6 both have significantly higher engagement than the base case operator (C1), while T3, T5, and possibly C2 and T4 have lower base case engagement.

Table 3 - Regression Output of Unique Active Days per user vs. Z-score of Latency & Operator by Measurement Quantile

	4%ile	14%ile	30%ile	45%ile	56%ile	75%ile	90%ile	95%ile	99%ile	Max
Mean Unique Active Days (Intercept)	5.6 ***	5.6 ***	5.7 ***	5.1 ***	5.1 ***	4.6 ***	4.4 ***	4.6 ***	5.1 ***	8.5 ***
Per StDev Increase in Latency	-0.5 ***	-0.4 ***	-0.5 ***	-0.7 ***	-0.7 ***	-1 ***	-0.8 ***	-0.7 ***	-0.4 ***	0.4 ***
ISP C2	0.3 ***	0.1 ***	0.1 ***	0.4 ***	0.3 ***	-0.4 ***	-0.1	-0.2 *	-0.2 *	-0.3 ***
ISP C3	0	0.1 ***	-0.1 *	0.1	0.1	-0.1	-0.2	-0.3 **	-0.1	-0.4 **
ISP C4	-0.5 ***	0.3 ***	-0.1 *	0.9 **	0.8 *	0.4 *	0.5 *	0.3	-0.1	-0.4
ISP C5	-0.3 ***	-0.2 **	-0.6 ***	-0.2	-0.3 **	-0.4 **	-0.1	-0.1	-0.1	-1.2 ***
ISP C6	-0.4 ***	-0.3 ***	-0.1	0.2	-0.6 ***	-0.6 ***	0.1	-0.1	-0.1	-1.6 ***
ISP C7	-0.2 ***	-0.3 ***	-0.1	0.1	0	-0.1	0.1	-0.1	0.7 ***	1.1 *
ISP C8	0	-0.1	-0.3 **	0	0.3 **	-0.2	0.1	0.1	0.1	-1 ***
ISP T1	-0.1 *	-0.1 ***	-0.2 ***	0	0.1	-0.2 *	0.1	0.1	0.1	-1.4 ***
ISP T2	-0.4 ***	-0.4 ***	-0.9 ***	-0.5 *	-0.4	-0.7 **	-0.3	-0.4 *	-0.1	0.1
ISP T3	-0.3 ***	-0.2 ***	-0.4 ***	-0.3 ***	-0.3 **	-0.5 ***	-0.1	-0.1	-0.3 ***	-1.7 ***
ISP T4	-0.6 ***	-0.2 *	-0.5 ***	0.3 ***	-0.3 *	-0.5 ***	0	-0.3 **	-0.3 *	-1.5 ***
ISP T5	-0.5 ***	-0.7 ***	-0.7 ***	-0.3 *	-0.7 ***	-0.7 ***	-0.2	-0.3 *	-0.5 ***	-2.5 ***
ISP T6	-0.4 ***	-0.2	-0.3 *	-0.1	0.2	-0.1	0.3	0.3	0.5 **	0
Adj. r ²	0.8%	0.7%	1.2%	1.9%	2.5%	4.9%	3.4%	2.6%	1.1%	1.4%

As is typical * = p < 5%, ** = p < 1%, *** = p < 0.1%.

Holographics Over 10G

Paving the Way for the Immersive Future

A Technical Paper prepared for SCTE by

Austin Pahl¹

Software Engineer, Immersive Media Experiences
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
303-661-3867
a.pahl@cablelabs.com

Dr. Abhinav Kshitij¹

Senior Engineer, Emerging Technologies
Charter Communications
6360 S Fiddlers Green Cir, Greenwood Village, CO 80111
480-376-4115
Abhinav.Kshitij@charter.com

Dell Wolfensparger

XR Architect, Emerging Technologies
Charter Communications

Logan Cho

Software Engineering Intern, Immersive Media Experiences
CableLabs

Thomas Alder

Software Engineer
Formerly OTOY

¹ These authors contributed equally.

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Challenges to Wider Adoption of Immersive Media	3
2.1. Immersive Media offers Depth-Based Perception	3
2.2. Growing industry demand	4
2.3. Rasterized transmission will require prohibitive bandwidths.....	5
3. Vectorized Content Delivery over Scalable Networks.....	5
3.1. 3D Streaming of Interchangeable Media Format	5
3.2. System Architecture	6
3.3. Intelligent Buffering over the 10G Network	7
3.3.1. Variables and Metrics	8
3.3.2. Asset Quality Selection	9
3.3.3. Basic Heuristic Example	9
4. Analysis	9
4.1. Performance Metrics	10
4.2. Compared to Conventional 2D Streaming	11
5. Discussion	12
5.1. Related Work.....	12
5.2. Future Work.....	12
6. Conclusion.....	13
Abbreviations	14
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 - 3D Streaming Architecture built with an Asset Delivery Pipeline	6
Figure 2 - Intelligent Buffering over the 10G Network.....	7
Figure 3 - Timing of Asset Transmissions in Example 3D Stream	10
Figure 4 - Measured Latency and Bandwidth of Assets in Example 3D Stream	10
Figure 5 - Measured Latency and Bandwidth of Assets in 3D Stream of a Sample Scene	11

1. Introduction

Depth-based media is an emerging market for users and enterprise that has recently witnessed a sharp uptick in growth and investment. With the rising demand for remote communication in virtual spaces, automation in transportation, maintenance, supply chain, and visualization techniques in healthcare, defense and simulation industry, startups and large companies are competing in this nascent market with the launch of fixed and wearable display units. A host of ecosystems are making efforts to integrate and co-ordinate accelerating development efforts among software professionals and industry experts. Despite growing support, challenges to content generation and transmission include developing an interchange format to support compatibility and an evolved network infrastructure to satisfy bandwidth and latency requirements to reliably and securely deliver immersive content.

The Immersive Digital Experiences Alliance (IDEA) was formed in 2019 to solve the twin problem of *media compatibility* and *media-aware transmission* over a 10G network. IDEA² developed 3D Streaming and Intelligent Buffering with the overall objective of enabling optimal immersive content delivery at minimal bandwidth consumption, while preserving viewing experience on multiple classes of immersive display units. The main benefit in bandwidth savings comes through offloading rendering from the core network and moving to the client-side.

The network architecture is robust enough to deliver assets of varying quality to the wide range of available compute resources on fixed and wearable units. The key assumption behind adaptive streaming being that display on a small screen requires fewer details, and therefore lower quality assets would reasonably allow for a satisfactory user experience on mobile and AR glasses. Larger displays however require greater amount of detail for objects closer to the viewer, which is captured and represented in higher-quality assets.

The present work briefly introduces media format interchange and proceed to explain the media-aware network enabled by 3D Streaming and Intelligent Buffering. Section 2 presents current bandwidth challenges to immersive media streaming to a host of different immersive platforms of varying screen sizes. Section 3 explains the 3D Streaming network architecture and Intelligent Buffering over a 10G network, along with a heuristic implementation of asset scheduling contained within the client-side logic. Section 4 evaluates bandwidth usage savings of queue-forming traffic flows of vectorized asset streaming over non-queue forming streaming of rendered frames. Latency measurements on the client-side demonstrate asset scheduling effectiveness in gaining maximum concurrency while fetching assets from remote asset servers. The conclusive sections describe algorithmic improvements made to asset scheduling and integration of current 10G capabilities into the existing network architecture.

2. Challenges to Wider Adoption of Immersive Media

2.1. Immersive Media offers Depth-Based Perception

Immersive media refers broadly to a variety of media that involves depth-based perception and accurately accounting for parallax differences at multiple depth levels. Immersive video allows the viewer to perceive the distance to depicted objects, with the viewer's own eyes as if the objects are physically present in the real world. *Parallax* is one of the main drivers of real-world perception: With viewer perspective shifting, objects nearer to a viewer appear to move relatively faster than objects in the background. VR and AR might be the most well-known examples of immersive media today. In contrast, an image on a computer monitor does not qualify as immersive media because the viewer can only

² IDEA Webinar on 3D Streaming, May 4th, 2022 [1]

perceive a flat image on the screen from any perspective and location. The present form of 3D movies also does not qualify as immersive media as parallax is not correctly accounted for by a pair of 3D glasses that stereoscopically superimpose a pair of rendered frames to create an illusion of depth perception.

While immersive media could reasonably be considered an emerging technology, there is a growing support on a host of platforms available in the market today. VR headsets are perhaps the longest running example of an immersive display available today, beginning with the high-profile founding of OculusVR in 2012 [2]. Augmented reality (AR) and mixed reality (MR) are also supported by several products on consumer and enterprise markets today [3], [4]. Together, VR, AR, and MR are collectively known as extended reality (XR).

Volumetric displays, supporting depth-based video, have recently emerged as a new class of display with form factors of a television or a desktop monitor. Volumetric displays are divided into two categories: *eye-tracking displays*, which track a single viewer's eye movements to create the illusion of depth on a 2D screen, and *light field displays*, which send different images out at different angles to produce perceivable depth for multiple viewers at once. Although few volumetric displays have reached the public to date, there are a number of companies already working in this space [5]–[10].

Immersive content generation occurs via digital content creation (DCC), through live captures from the real-world events using advanced camera technologies or using a combination of the two approaches. Digital content creation is the most common approach today, because for the most part it is straightforward to extend existing digital workflows to support immersive displays. Immersive display manufacturers publish free-to-use plugins that integrate with popular graphics toolsets and game engines. In contrast, live capture methods have lagged behind in technological development and adoption for immersive content generation and streaming, although this has been an area of significant innovation in both industry and academia over recent years. Depending on the use case, application requirements, and access to compute and network resources, live capture approaches extend from 2D photo conversion pipelines to depth cameras built into the latest smartphones to high-precision specialized camera systems: Neural Radiance Fields (NeRF) uses a neural representation to achieve exceptionally high fidelity from a sparse set of 2D input images [11]. For a deeper dive, refer to the IDEA white paper [12] on live capture methods and representations.

2.2. Growing industry demand

Over recent years, interest in immersive media has risen significantly. Some have argued that prior to the pandemic, AR and VR reached the “trough of disillusionment” along the Gartner Hype Cycle, which occurs after a product reaches peak inflated expectations and fails to deliver on the hype [13]–[15]. Then during the pandemic, people became acutely aware of the limitations of video conferencing as opposed to face-to-face exchanges. The reasons are numerous – lack of copresence, removal of spontaneous, random encounters, and perhaps worst of all, the newly dubbed “Zoom fatigue” that people experience after extended periods of time spent on video calls [16]–[18]. AR and VR witnessed a significant growth throughout the pandemic, at least in part because of the increased time spent working from home and the prospect of overcoming the limitations of video conferencing [19].

Then the concept of the metaverse came into mainstream attention with Facebook's rebranding to Meta at Facebook Connect 2021³, while showcasing their progress on building a new platform for rich, immersive social experiences online. Google's Project Starline also investigated immersive media technologies for enhanced telepresence, bringing live 3D video to the video call format [16]. NVIDIA announced its

³ CEO Mark Zuckerberg's letter to Meta employees [20]

launch of Omniverse platform that allows real-time collaboration on digital twins, architecture, education, and facilities maintenance [21]. Looking Glass Factory recently secured the CIA's venture capital funding to provide immersive displays for intelligence and defense applications [22]. Hollywood movies are increasingly being produced using game engines [23].

Industry giants and startups continue to expand upon XR development – creating applications, utility tools and supporting hardware to allow immersive content generation, streaming and consumption. Display manufacturers, game engine developers, network operators, chip manufacturers and application developers continue to invest capital and participate in evolving ecosystems surrounding immersive technologies. We believe this trend will continue to dominate as demand for immersive content finds an increasing use in improving learning and productivity while enhancing entertainment experiences.

2.3. Rasterized transmission will require prohibitive bandwidths

Newer displays and media pose several challenges that need to be considered to enable the ideal vision of an immersive future. First, the massive variety of methods for capture, encoding, and display of immersive media leads to a “many-to-many” problem when developing workflows and processes related to immersive media. Each method, format, and display has its own advantages and drawbacks, and conversion between representations runs the risk of significant information loss. Even today, conversion among existing 3D scene description formats can lead to problems like missing materials, untranslatable logic, and subtly altered rendering behaviors. This gets worse as the number of features and formats continues to grow.

Bandwidth requirements for immersive displays are expected to grow at an unprecedented rate. State-of-the-art VR headsets today reach resolutions above 4K [24], but light field displays are anticipated to be orders of magnitude higher resolutions than anything available today. To provide a 3D effect without eye tracking, light field displays attempt to mimic the behavior of rays of light bouncing off a real, physical subject. While a pixel on a 2D display unit encodes a single color at 24 bits in total, a *holographic pixel* would need color encoding for each ray emanating from angles discretized in the azimuthal and altitudinal directions. A holographic pixel supporting 90 different angles horizontally (azimuthal) and 90 different angles vertically (altitudinal) could enable viewers to experience a few inches of depth [25], but would require a total of 8100 color encodings. A holographic still image on a UHD-4K display (3840×2160) using these 90×90 holographic pixels would require 67 gigabytes of uncompressed data. While compression reduces data requirements, their application on compressing immersive media for light field displays is in early prototyping phases, and public data and literature remains scarce.

3. Vectorized Content Delivery over Scalable Networks

3.1. 3D Streaming of Interchangeable Media Format

IDEA has published a suite of royalty-free specifications establishing a baseline for interchange of immersive media, known as the Immersive Technologies Media Format (ITMF). The format was initially intended to be used for interchange amongst industry-standard digital content creation (DCC) tools, i.e. for the packaging and creation of 3D synthetic, computer generated, and natural media, including audio and visual media. As a baseline format primarily for use with DCC tools, assets described by ITMF are agnostic to the specific type of device on which they may be presented. For example, visual media will be display-agnostic, so that a subsequent rendering step in a media- and application-aware distribution system can reformat the visual media to match the capabilities of the client display.

While streaming rendered frames require massive bandwidths, real-time streaming of immersive content followed by rendering 3D assets on the client within display units could mitigate the challenges associated with the delivery of immersive media. Local rendering of 3D assets on a game engine runtime has the advantage of asset reuse over multiple scenes, thereby eliminating bandwidth redundancy that comes with streaming rasterized frames. We present **3D Streaming** — a system architecture supporting real-time streaming of immersive content to clients by transmitting ITMF scene graphs and associated assets to clients. In addition to reducing light field display bandwidth requirements, 3D Streaming simplifies content distribution to heterogeneous immersive display units. This system also generalizes to other scene graph formats – Universal Scene Description (USD), Graphics Language Transmission Format (glTF) – accommodating diverse application-level requirements and use cases.

The general framework of 3D Streaming finds its form in a network architecture that could be scaled on-prem, in cloud or the edge, dictated by network flows in streaming 3D assets depending on the use case. The key components described here would be considered essential to the overall implementation.

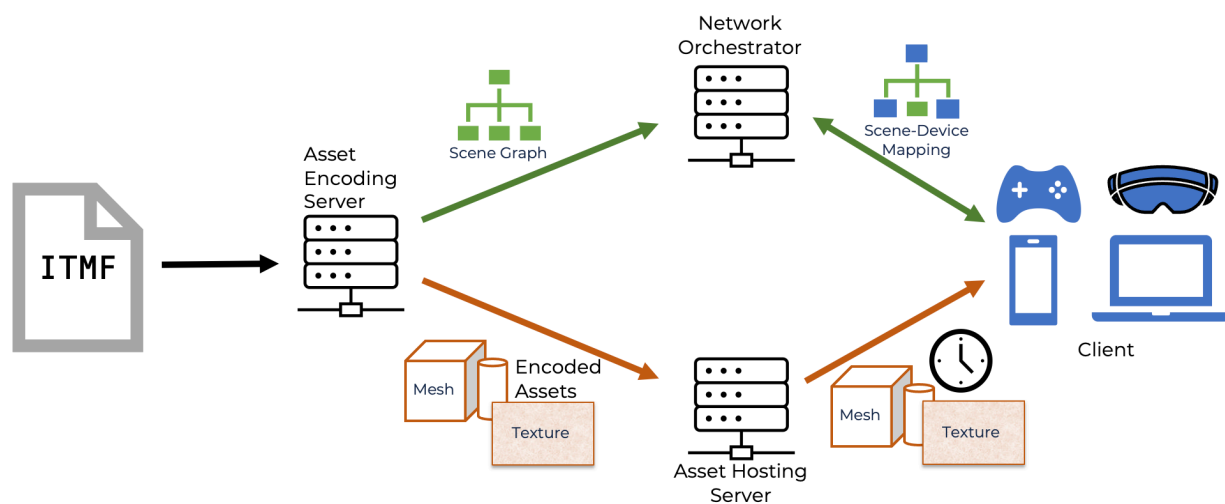


Figure 1 - 3D Streaming Architecture built with an Asset Delivery Pipeline

3.2. System Architecture

Figure 1 describes the architecture of the 3D Streaming demonstrated in IDEA's May 2022 webinar [1]. The end-to-end pipeline of conveying an ITMF scene to an immersive display unit consists of the Asset Encoding Server (AES), the Network Orchestrator (NO), and the Asset Hosting Server (AHS). Initially an ITMF container is ingested by the Asset Encoding Server (AES), which extracts various files from the ITMF container, namely (1) the *assets* (meshes, textures) that spawn in scenes, and (2) the *scene graphs* (XML, JSON) describing the layout and properties of the assets and the scenes (lighting, animation).

The AES pushes the scene graphs to the Network Orchestrator (NO), which is mainly responsible for mapping client requirements to the *asset quality*, and control plane operations related to individual streams. Asset quality (AQ) is defined as an abstraction of mesh compression, texture compression and levels-of-detail (LODs). The AQ-client mapping allows for a real-time adaptive asset streaming to heterogeneous platforms (clients) of widely ranging compute requirements and operating under dynamic network conditions (latency, jitter).

The assets are pushed to the Asset Hosting Server (AHS)⁴, defined as an abstraction layer of content delivery network (CDN) hosted on-prem, in public/private cloud and on edge resources. The AHS distributes assets of varying asset quality (AQ) to meet the compute and network requirements of the application (Section 3.3.2). Alternate encodings and/or levels of detail (LODs), i.e., AQ, may be generated by the AES and included in the distribution to the AHS to support performance and hardware constraints on the client.

When a client initiates a 3D Stream, it establishes a connection with the NO, which ensures that any constraints known up front are applied to the scene, such as support for specific asset encodings or display-specific content. The modified scene information is then sent to the client. The client reads the scene information, sends data plane asset requests to the AHS, and finally the client renders the scene in real time on a game engine (Unreal, Unity) runtime. The client makes several asset fetch calls to the AHS at different times during an interactive session (multiplayer gaming, XR application), and a non-interactive session (live streaming, playback) to buffer immersive content on an immersive display unit. Therefore, this *scene awareness* allows for buffering content time-to-time, distributing large downloads over multiple smaller downloads over time, as and when the required assets are relevant to the scene.

3.3. Intelligent Buffering over the 10G Network

Intelligent Buffering refers to the use of network and scene awareness to fetch 3D assets from the AHS such that fetch times and the impact of adverse network traffic conditions are minimized while ensuring best possible QoE for the client. *Network awareness* refers to the consideration of latency and bandwidth and *scene awareness* refers to the consideration of 3D scene properties like asset placement and timing. Together, network and scene awareness enable cloud orchestration for scalable, adaptive streaming of 3D assets of different levels of detail and compression. We consider two different facets of intelligent buffering that can be controlled at runtime to support QoE: (1) Queueing asset fetches according to the time they appear in the scene; (2) selecting asset LODs when prioritization is not sufficient.

Intelligent Buffering

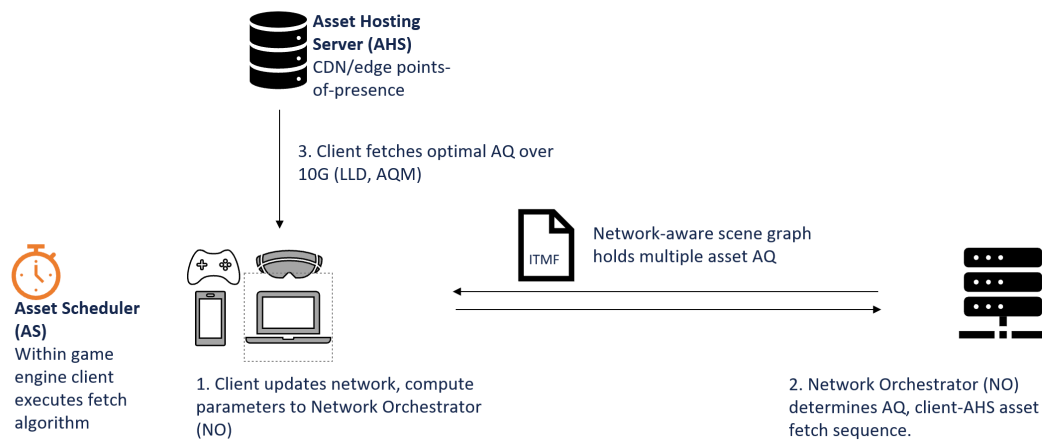


Figure 2 - Intelligent Buffering over the 10G Network

⁴ In prior webinars, this was called the “Asset Server”.

As noted in Section 3.1, the main advantage of 3D Streaming comes in the form of asset reuse over multiple scenes, therefore mitigating redundant content streaming requirements over several rasterized frames. For example, assets representing natural background objects (forest, rocks, foliage) remain largely static in a scene, and could be fetched once from the AHS and reused on multiple scenes. The uncompressed rendered assets (highest AQ) preserve the content quality under favorable network and compute conditions, much like adaptive bitrate streaming can deliver the highest quality of video content available, and adaptively adjusts frame bitrate as network conditions degrade during video streaming.

The client-AHS connection forms the main network bottleneck of 3D Streaming due to larger bandwidths (compared with client-NO, AES-NO connection) and rendering latency requirements on the client. In most cases, the scene graph is significantly smaller than the assets that fill the scene. During a 3D Stream, it is essential that the client fetches and renders all of the necessary assets on its memory to allow for compute and rendering latencies to display when needed. Failure to display assets during a scene playback would result in loss of necessary details or may temporarily pause the playback while the assets are transferred and loaded (akin to “buffering” on video streams). Such issues arise due to adverse impacts of network congestion, packet loss, and poor memory management.

This critical queue-forming traffic can be optimized by leveraging maturing 10G capabilities of Low-Latency DOCSIS[®] (LLD) technology with Active Queue Management (AQM) being incorporated into the working design of Intelligent Buffering. While traditional video streaming is download-heavy, immersive traffic would require the increased upstream capability of DOCSIS 4.0 networks to enable live capture and streaming. Mobility considerations include Low-Latency Wi-Fi and Low Latency Mobile Xhaul over converged networks.

The following sections describe intelligent buffering using its runtime variables and metrics, along with an overview of asset selection processes and algorithms as key enabling tools to optimize asset delivery to multiple clients on heterogenous platforms.

3.3.1. Variables and Metrics

In our initial intelligent buffering system, we incorporated the following variables and metrics available at runtime to support effective scheduling decisions. A list of assets from the 3D scene carries the following information:

- Asset type (mesh, texture, animation),
- URLs for multiple AQs of each asset,
- Asset file size associated at each AQ,
- Start and end times of asset visibility in a scene.

While the first three metrics constitute asset metadata available on the AHS and the scene graph, asset visibility time range is not explicitly described in typical scene graph representations. Rather it may be computed in a preprocessing step on the client. Content that is also generated in real-time, such as that from a source game engine, could be instrumented to log and transmit asset visibility, but this is not yet implemented.

From the network connection, the round-trip time (RTT) and download throughput is measured for every asset fetch call from the client to the AHS to determine current network conditions that dictate AQ selection for next fetch call or subsequent fetch calls over a time window.

3.3.2. Asset Quality Selection

The asset quality selection algorithms determine fetch sequencing and AQ. In traditional 2D video streaming, a high bitrate video is typically encoded on the server with multiple bit rates, enabling clients to perform adaptive bitrate: switching between encodings to maximize the content visual quality without experiencing interruptions like buffering. The analog to this in 3D Streaming is asset selection: the AHS provides multiple AQs for data intensive assets like textures and meshes. This way, the intelligent buffering algorithm may react to network conditions to provide similar assurances to those applied by adaptive bitrate.

Most common asset types, especially those that tend to have large file sizes, have existing facilities for lossless and lossy compression. For example, textures are often ingested into game engines using widely used formats like PNG or JPEG. However, the engine may convert image files to specialized, lossy texture compression formats like DXT1 which reduce game package size with minimal or zero impact on decode latency [26], [27]. With respect to meshes, games often include multiple AQs of the same mesh because different amounts of detail are needed when an object is close or far from the camera. In practice, mesh AQs may be generated either automatically or manually based on how much control is needed [28].

3.3.3. Basic Heuristic Example

Here, we demonstrate a simple heuristic approach for network-aware intelligent buffering to illustrate its fundamental application. If a designated latency threshold is met, say, over 45 milliseconds, a “high latency” mode is triggered which means subsequent asset fetches are made using smaller asset variants. In practice, there are many directions that could be taken to improve performance (see Section 5.2 for further discussion).

Consider a client that initiates a 3D Stream of an ITMF file containing a list of N assets. For each asset, a high detail and low detail variant are hosted on the AHS. Once the client receives the scene information, the client sorts the asset list according to time of first appearance. The first asset’s larger variant is fetched, and round-trip time (RTT) is measured. If this time exceeds 45 milliseconds, the high latency mode is toggled so that the next asset fetched uses its smaller variant. Fetch the next asset, measure RTT, and update high latency mode if needed. Repeat this process until all assets have been fetched. Rendering can begin as soon as all assets that appear at the very start of the scene have arrived, at which point rendering and streaming may continue in parallel.

While simple, this approach is primarily intended to illuminate the problem space that is occupied by intelligent buffering. The overwhelming majority of the logic occurs on the client, which facilitates low-cost, high throughput server deployments while achieving the objective of maximal QoE for clients. As previously mentioned, the AHS requires no more than static file hosting, so that component may be deployed on a traditional CDN.

4. Analysis

This section presents early analysis of intelligent buffering, including a description of key network performance metrics and a qualitative comparison to traditional video streaming. Measurements were taken on 3D Streams of small sample scenes.

4.1. Performance Metrics




Figure 3 - Timing of Asset Transmissions in a 3D Stream




Figure 4 - Measured Latency and Bandwidth of Assets in a 3D Stream

Figure 3 shows a waterfall chart of the asset transmissions that occurred over time during a stream. Assets 0 and 1 were transferred first for the scene to begin rendering, then Assets 2-4 were transferred afterwards. This demonstrates that asset transmissions in a 3D Stream need not occur synchronously: much like a web browser, the client can utilize multiple connections at once to provide a smoother experience where possible. Our goal with intelligent buffering is to develop an approach that minimizes the horizontal length of this plot (total time elapsed transmitting assets).

Figure 4 shows, for each asset transmitted in a stream, the latency and average bandwidth measured from the HTTP response. The vertical dashed line represents the latency threshold from our heuristic example in Section 3.3.3: assets whose latency surpasses 45 milliseconds, Asset 4 in this case, trigger “high latency mode”. Moving beyond the basic heuristic example, bandwidth measurements are a key metric for understanding network conditions to enhance QoE.

4.2. Compared to Conventional 2D Streaming




Figure 5 - Measured Latency and Bandwidth of 3D and 2D Streaming

3D Streaming can be considered a complementary solution to traditional video streaming for media delivery: While 3D Streaming does not completely replace the use cases enabled by video streaming, it performs better in various contexts. At a high level, the biggest differences between 3D and video streaming are the use of 3D rendering on the client and the patterns of network transfer. 3D Streaming requires sufficient client resources to perform the real-time render, but this is increasingly common today due to widespread consumer use of graphics processing units (GPUs).

The difference in network behavior is illustrated by Figure 5. This plot shows bytes transferred each second on a 3D Stream and a video stream. In this case, both streams coincidentally transfer about 40 megabytes of data in total. However, the 3D Stream transmits most of that data in one short burst up front, while the video stream transmits small amounts of data steadily over time. This is because the 3D Stream transmits all assets as soon as possible, while the video stream must consistently transmit video frames over the full duration of the content playback.

With the rise of heterogeneous immersive displays, we anticipate that content creators will increasingly need to develop media that supports many different types of displays at once, such as VR headsets, volumetric displays, and 2D displays. Ideally, that media should be tailored as best as possible for each display type. This would be challenging with traditional video: a different version of the content would need to be produced for every single display type, and it would be hard to accommodate the specific advantages of different displays, such as VR's higher degrees of freedom, without a focused, manual effort. Real-time rendering simplifies this process because the display-specific experience can be generated at runtime, often via a specialized plugin or API provided by the display manufacturer. The streaming application may also support custom enhancements like interactivity, which are much more complex to execute in a video streaming context.

As mentioned earlier, immersive displays are leading to increases in effective video resolutions at unprecedented rates, particularly light field displays which increase by an order of magnitude for each degree of freedom introduced. While today's display resolutions are generally well supported by modern network infrastructure and video codecs, it is unclear how long today's systems will remain effective. 3D Streaming is unaffected by this problem because the size of the scene is not correlated with the resolution of the display like a video is. Furthermore, rising adoption of 10G facilitates the delivery of larger scenes faster and more reliably.

As long as the client has sufficient storage, assets only need to be transmitted to the client once each. This means that reuse of assets over time is rewarded with less network chatter. In Figure 5, this is apparent because all the assets were transmitted at the beginning, then nothing else needed to be transmitted for the remaining playback time, keeping the network silent for the rest of the trace. Asset reuse is also useful for scenes that have repetitive content, such as trees, bushes, or buildings in a cityscape. This is already a commonly practiced technique in game development, as it allows the use of *GPU instancing*, which is a performance optimization that renders multiple copies of an object in a scene at one time [29].

5. Discussion

5.1. Related Work

Several online video games and geographic applications have developed systems for the real-time delivery of 3D content over the network [30]–[33]. Similarly, efforts in cloud and distributed rendering have implemented related functionality such as real-time collaboration on shared 3D scenes or scene delivery over the network [34], [35]. All these systems excel at their respective use cases, but do not generalize beyond that. In contrast, 3D Streaming is a general-purpose system for 3D content delivery over the network. In the future, it should be possible to build new networked 3D applications, whether games or content creation tools or otherwise, by leveraging the architecture presented here.

Petrangeli et al. [36] presented a system for streaming AR objects in real time with a mechanism for heuristically adapting LODs according to network condition and scene placement, significantly reducing startup latency and data requirements compared to predownloaded AR scenes. In the context of our work, their approach would be an effective drop-in solution for network and scene awareness in intelligent buffering, likely to be included in future analyses of intelligent buffering methodology. Our architecture also generalizes to broader use cases, including heterogeneous immersive display units with support for tailored experiences.

5.2. Future Work

Enhanced network awareness. We intend to develop a robust analysis of intelligent buffering algorithms. Petrangeli et al.’s work on AR streaming [36] provides one option to analyze, but there may be ways that we can incorporate enhanced network and scene awareness for further improvements.

Device awareness. Another factor to consider in optimizing the QoE of a 3D Stream is the client hardware. We refer to the usage of client hardware conditions and specifications for ensuring the best possible QoE as *device awareness*.

One example of device awareness is to consider the client as a cache comprised of three layers: its GPU memory, CPU memory and storage. In situations where the scene is particularly large or the client is resource-limited, such as embedded or mobile devices, it is possible that the entire scene would not fit on the client at one time. When an asset is delivered to the client over the network, we would store it on an available layer, beginning with GPU memory, falling back to the next when out of space. Coupled with an eviction algorithm that takes into consideration the available space in each layer, along with network and scene conditions, to intelligently free up space, this approach would allow for lower latency, local retrievals of previously seen assets into the scene. Another opportunity for device awareness is to factor in the client’s screen resolution into our asset quality selection algorithm. Lower asset quality is less likely to harm QoE on lower resolution displays.

Real-world capture support. We are interested in exploring the deployment of methods for viewing real-world 3D data on the client. Many real-world capture encodings can be embedded into a 3D scene graph [12], [37], and some of them, including NeRF, can render in real time [38]. As real-world 3D capture becomes more accessible, this will likely become a core use case for immersive displays. Embedding these captures into scene graphs will also enable new forms of mixed content: for example, one could imagine a virtual gallery filled with 3D scans of real art.

6. Conclusion

The current paper explores recent progress made on the development of a network architecture for scalable, vectorized content distribution to multiple platforms, rendered on client-side game engines. An end-to-end testing of a content delivery pipeline demonstrates significant bandwidth savings, while preserving content quality during transmission and allowing for asset reuse over multiple scenes. The architecture leverages modularity and scalability to allow for high availability of content over core network and cloud deployment. By streaming content over a 10G network, queue-forming traffic of immersive content can be delivered over reasonable times.

A key challenge for the present architecture lies in compute requirement on the client-side, especially with the growing demand for lighter wearable XR platforms to improve user experience. Future testing of the asset scheduler will determine algorithmic effectiveness in improving AQ adaptability, while *tail-end latency* ranges are expected to be curtailed primarily by deploying 10G capabilities on the existing platform without the need for making significant hardware changes.

As newer versions of game engines feature photorealism with even greater detail, vectorized content streaming could be the preferred choice of streaming immersive content. Although it may satisfy the demands of latency-sensitive applications like gaming, live event streaming (sports, concerts) will deliver live-captures (large bandwidths) but driven by latency requirements. The current architecture presents a general framework that can be adapted for live-capture and streaming by using low-latency techniques, delivering traditional video frames and 3D assets on separate queues.

Abbreviations

2D	two-dimensional
3D	three-dimensional
AES	Asset Encoding Server
AHS	Asset Hosting Server
AQM	Active Queue Management
AR	augmented reality
CDN	content delivery network
DCC	digital content creation
glTF	Graphics Language Transmission Format
GPU	graphics processing unit
IDEA	Immersive Digital Experiences Alliance
ITMF	Immersive Technologies Media Format
LLD	Low Latency DOCSIS
LOD	level of detail
NeRF	neural radiance field
NO	network orchestrator
QoE	quality of experience
RTT	round-trip time
UHD-4K	ultra-high-definition 4K
USD	Universal Scene Description
VR	virtual reality
XR	extended reality

Bibliography & References

- [1] *ITMF 3D Streaming Demo Webinar*, (May 04, 2022). Accessed: Jul. 14, 2022. [Online Video]. Available: <https://www.immersivealliance.org/videos/>
- [2] G. Kumarak, “A Brief History Of Oculus,” *TechCrunch*. <https://social.techcrunch.com/2014/03/26/a-brief-history-of-oculus/> (accessed Jul. 13, 2022).
- [3] “Microsoft HoloLens | Mixed Reality Technology for Business.” <https://www.microsoft.com/en-us/hololens> (accessed Aug. 03, 2022).
- [4] “Enterprise augmented reality (AR) platform designed for business | Magic Leap.” <https://www.magicleap.com/en-us/> (accessed Aug. 03, 2022).
- [5] “Simulated Reality | 3D Display technology,” *Dimenco*. <https://www.dimenco.eu> (accessed Jul. 13, 2022).
- [6] “Sony Spatial Reality Display | Sony US,” *Sony Electronics*. <https://electronics.sony.com/more/spatial-reality-display/p/elfsr1> (accessed Jul. 13, 2022).
- [7] “Leia Inc. – 3D Lightfield Experience Platform.” <https://www.leiainc.com/> (accessed Jul. 13, 2022).

- [8] “Looking Glass Factory: The Hologram Company,” *Looking Glass Factory*.
<https://lookingglassfactory.com> (accessed Jul. 13, 2022).
- [9] “Light Field Lab.” <https://www.lightfieldlab.com/> (accessed Jul. 13, 2022).
- [10] “Avalon Holographics Inc.,” *Avalon Holographics Inc.* <https://www.avalonholographics.com>
(accessed Jul. 13, 2022).
- [11] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis,” in *Computer Vision – ECCV 2020*, Cham, 2020, pp. 405–421. doi: 10.1007/978-3-030-58452-8_24.
- [12] “Photographic Live Action Capture for Immersive Media,” Immersive Digital Experiences Alliance. Accessed: Jul. 14, 2022. [Online]. Available:
<https://www.immersivealliance.org/download/download-photographic-live-action-capture-for-immersive-media/>
- [13] M. Brenner, “The Resurgence of AR and VR Content in a Post-Pandemic World,” *Marketing Insider Group*, Jul. 27, 2021. <https://marketinginsidergroup.com/content-marketing/the-resurgence-of-ar-and-vr-content-in-a-post-pandemic-world/> (accessed Jul. 13, 2022).
- [14] J. Pace, “XR and the Self-Inflicted Trough of Disillusionment,” *Medium*, Nov. 07, 2019.
<https://arvrjourney.com/xr-and-the-self-inflicted-trough-of-disillusionment-e2177c6b33fe> (accessed Jul. 14, 2022).
- [15] “Gartner Hype Cycle Research Methodology,” *Gartner*.
<https://www.gartner.com/en/research/methodologies/gartner-hype-cycle> (accessed Jul. 14, 2022).
- [16] J. Lawrence *et al.*, “Project starline: a high-fidelity telepresence system,” *ACM Trans. Graph.*, vol. 40, no. 6, pp. 1–16, Dec. 2021, doi: 10.1145/3478513.3480490.
- [17] C. Morris, “This startup wants to replace your Zoom meetings with holograms,” *Fast Company*, Mar. 11, 2022. <https://www.fastcompany.com/90730176/startup-matsuko-zoom-meetings-holograms>
(accessed Jul. 14, 2022).
- [18] V. Ramachandran, “Four causes for ‘Zoom fatigue’ and their solutions,” *Stanford News*, Feb. 23, 2021. <https://news.stanford.edu/2021/02/23/four-causes-zoom-fatigue-solutions/> (accessed Jul. 14, 2022).
- [19] S. Vardomatski, “Council Post: Augmented And Virtual Reality After Covid-19,” *Forbes*.
<https://www.forbes.com/sites/forbestechcouncil/2021/09/14/augmented-and-virtual-reality-after-covid-19/> (accessed Jul. 21, 2022).
- [20] M. Zuckerberg, “Founder’s Letter, 2021,” *Meta*, Oct. 28, 2021.
<https://about.fb.com/news/2021/10/founders-letter/> (accessed Jul. 14, 2022).
- [21] “NVIDIA Announces Omniverse Open Beta, Letting Designers Collaborate in Real Time — from Home or Around the World,” *NVIDIA Newsroom*. <http://nvidianews.nvidia.com/news/nvidia-announces-omniverse-open-beta-letting-designers-collaborate-in-real-time-from-home-or-around-the-world> (accessed Aug. 01, 2022).

- [22] L. Fang and J. Poulson, “The Brooklyn Hologram Studio Receiving Millions From the CIA,” *The Intercept*, May 27, 2022. <https://theintercept.com/2022/05/27/metaverse-cia-military-hologram-looking-glass-factory/> (accessed Aug. 01, 2022).
- [23] M. Seymour, “Art of LED Wall Virtual Production, Part One: ‘Lessons from the Mandalorian,’” *fxguide*, Mar. 04, 2020. <https://www.fxguide.com/featured/art-of-led-wall-virtual-production-part-one-lessons-from-the-mandalorian/> (accessed Aug. 01, 2022).
- [24] “Varjo XR-3,” *Varjo.com*. <https://varjo.com/products/xr-3/> (accessed Jul. 21, 2022).
- [25] “SO, WHAT IS A HOLOGRAPHIC DISPLAY?,” *Avalon Holographics Inc.* <https://www.avalonholographics.com/resources/what-is-a-holographic-display> (accessed Jul. 21, 2022).
- [26] “Textures in Unreal Engine.” <https://docs.unrealengine.com/5.0/en-US/textures-in-unreal-engine/> (accessed Jul. 17, 2022).
- [27] “Texture Format Support and Settings.” <https://docs.unrealengine.com/5.0/en-US/texture-format-support-and-settings-in-unreal-engine/> (accessed Jul. 17, 2022).
- [28] “Geometry Best Practices for Artists.” <https://developer.arm.com/documentation/102496/0100/Level-of-Detail---LOD> (accessed Jul. 17, 2022).
- [29] “Unity - Manual: GPU instancing.” <https://docs.unity3d.com/Manual/GPUInstancing.html> (accessed Jul. 22, 2022).
- [30] “3D Tiles,” *Cesium*. <https://cesium.com/why-cesium/3d-tiles/> (accessed Jul. 18, 2022).
- [31] “VRChat,” *VRChat*. <https://hello.vrchat.com> (accessed Jul. 18, 2022).
- [32] F. Brown, “Microsoft wants to bring back Flight Simulator to show it supports PC,” *PC Gamer*, Jun. 11, 2019. Accessed: Jul. 18, 2022. [Online]. Available: <https://www.pcgamer.com/microsoft-wants-to-bring-back-flight-simulator-to-show-it-supports-pc/>
- [33] “Second Life - Virtual Reality, VR, Avatars, and Free 3D Chat.” <https://secondlife.com/> (accessed Jul. 18, 2022).
- [34] “Omniverse Platform for Virtual Collaboration,” *NVIDIA*. <https://www.nvidia.com/en-us/omniverse/> (accessed Jul. 18, 2022).
- [35] “The Render Network.” <https://render.x.io> (accessed Jul. 18, 2022).
- [36] S. Petrangeli, G. Simon, H. Wang, and V. Swaminathan, “Dynamic Adaptive Streaming for Augmented Reality Applications,” in *2019 IEEE International Symposium on Multimedia (ISM)*, Dec. 2019, pp. 56–567. doi: 10.1109/ISM46123.2019.00017.
- [37] Q.-A. Chen, “nerf_Unity.” Jul. 21, 2022. Accessed: Jul. 21, 2022. [Online]. Available: https://github.com/kweal23/nerf_Unity
- [38] A. Yu, R. Li, M. Tancik, H. Li, R. Ng, and A. Kanazawa, “PlenOctrees for Real-time Rendering of Neural Radiance Fields,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, Montreal, QC, Canada, Oct. 2021, pp. 5732–5741. doi: 10.1109/ICCV48922.2021.00570.

How Broadband Customers Can Benefit from Newfangled Wi-Fi Multiple User Features

A Technical Paper prepared for SCTE by

David John Urban
Distinguished Engineer
Comcast
One Comcast Center Philadelphia PA
610-476-2596
david_urban@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Begin with the fundamentals	3
3. Types of multiuser techniques in Wi-Fi 6e 802.11ax	8
4. DL MU-MIMO	9
5. UL MU-MIMO	12
6. DL OFDMA.....	13
7. UL OFDMA.....	13
8. Spatial Reuse	19
9. Putting It All Together With Real Devices and Applications	20
10. Conclusion.....	23
Abbreviations	23
Bibliography & References.....	24

List of Figures

Title	Page Number
Figure 1 - Unit circle illustrating quadrature signals	4
Figure 2 Impulse response for sending high speed information in amplitude	4
Figure 3 - Impulse response in the frequency domain illustrates orthogonality.....	6
Figure 4 - DL MU-MIMO Example Block Diagram	11
Figure 5 - Dynamic Range in the uplink	14
Figure 6 - Traffic mix	20
Figure 7 - DL MU-MIMO and UL OFDMA traffic mix	21
Figure 8 - DL MU-MIMO and UL OFDMA Measured traffic mix	22

List of Tables

Title	Page Number
Table 1 kHz tones and ns	7
Table 2 - DL MU-MIMO Measurement 1200 Mbps PHY Stations	10
Table 3 - DL MU-MIMO Measurement Two 2400 Mbps PHY Stations	11
Table 4 – DL MU-MIMO Measurement Three 2400 Mbps PHY Stations.....	12
Table 5 – UL OFDMA Measurement Two 2400 Mbps PHY Stations 6 GHz band.....	14
Table 6 – UL OFDMA Measurement Two 1200 Mbps PHY Stations 5 GHz band.....	15
Table 7 – Receiver noise floor and channel width for 3 dB noise figure at room temperature.....	16
Table 8 – EIRP regulation in the US for 6 GHz low power indoor access points	16
Table 9 – EIRP regulation in the US for 6 GHz low power indoor stations.....	17
Table 10 – Relationship between frequency and wavelength for 2.4, 5, 6 GHz bands.....	17
Table 11 – Distance, receive level for 160 MHz channel width uplink low power indoor station EIRP - 1 dBm per MHz	19
Table 12 - Statistics of Throughput with Four MU devices	22

1. Introduction

Historically, Wi-Fi is a polite protocol. Users will take turns in the time dimension to access the radio interface, sending and receiving messages. As the number of devices in the home has grown, more efficient multiple-user technologies have been introduced. The objective is to utilize these advanced multiuser features to benefit broadband customers in a tangible way. Techniques to improve the speed of a single user can be readily observed by customers with a simple speed test. The benefits of multiple user technologies require several devices and applications working in tandem. Laser focus on how technology benefits customer experience rather than technology for technology's sake is key. A traffic model that accurately reflects the situation in customers' homes is critical in applying multiuser technology that improves customer experience.

For multiuser techniques such as Multiple User Multiple Input Multiple Output (MU-MIMO) and Orthogonal Frequency Division Multiple Access (OFDMA) to work, a confluence of devices and applications must meet. The customer's devices must have multiple-user capability. Not all devices can take advantage of multiuser technology. Sometimes, the probability of several devices using applications that need access to the Wi-Fi radio signal at the same time is low; when low probabilities multiply the overall probability gets exponentially lower. The trick is to weed out the low probability use cases and focus on the commonly occurring use cases. This paper identifies the commonly occurring use cases for multiuser Wi-Fi technology and describes the theory behind multiple user techniques of multiple antenna spatial streams, spectrum resource unit allocation, and spatial frequency reuse. The paper sorts out aspects that are mostly for show from the technologies that lead to better user experience.

2. Begin with the fundamentals

The cosine of 0 degrees is 1 while the sine of 0 degrees is 0. The cosine of 90 degrees is 0 while the sine of 90 degrees is 1. The sine and cosine wave functions are in quadrature, simply offset in phase by ninety degrees, $\frac{\pi}{2}$ radians. This property allows information to be sent on the amplitude of the cosine wave and sine wave at the same time. The trick is to read the amplitude of the cosine wave at the phase of 0 degrees when the sine wave is nulled out, and the amplitude of the sine wave at phase 90 degrees when the cosine wave is nulled out. Figure 1 shows a unit circle as an aid to visualize the quadrature property of sine and cosine waves. This is the familiar quadrature amplitude modulation (QAM). In the cable world we are so familiar with these signals that we give them a nickname, "QAMs". QAMs carry a multiplex of video streams in a 6 MHz channel width. Many QAMs are combined, each with different center frequency. This is frequency division multiplexing, FDM. Wi-Fi uses QAM and FDM as well, yet in a slightly different way.

OFDM, orthogonal frequency division multiple access, is a form of FDM with the restriction that the signals are orthogonal. The FDM used with DOCSIS and video QAMs contain each multiplexed signal within an assigned 6 MHz block of spectrum. The signals do not exceed their 6 MHz block. By contrast, the frequency multiplexed signals in OFDM are not contained in frequency at all, even though they do tend to die off further away from the center frequency. In fact, each signal goes on in frequency forever.




Figure 1 - Unit circle illustrating quadrature signals

A periodic signal can be represented by a Fourier series expansion, an infinite sum of sine and cosine waves with discrete frequencies that are multiples of the inverse of the period. A periodic signal in the frequency domain consists of uniformly repeating spectral lines. Periodic functions can be useful, such as clocks, but since they last forever in time, they do not make for high-speed information transfer.




Figure 2 Impulse response for sending high speed information in amplitude

A good signal for fast information transfer is a quick impulse function as shown in Figure 2. The impulse function, unlike the periodic function, is contained in a finite chunk of time. Being finite in time makes the impulse response go on forever in frequency. The frequency response of the time impulse function is determined by the Fourier transform shown in Equation (1).

$$F(\omega) = \int_{-\infty}^{+\infty} f(t)e^{-j\omega t} dt \quad (1)$$

$$e^{j\theta} = \cos(\theta) + j \sin(\theta) \quad (2)$$

The time domain impulse function as shown in Figure 2 is even and symmetric. Euler's formula shown in Equation (2) allows the exponential function of an imaginary argument to be broken up into real cosine and imaginary sine parts. The sine function is odd. When integrated with an even function the imaginary sine part of $e^{-j\omega t}$ will vanish leaving only the real cosine part of $e^{-j\omega t}$ to integrate. The cosine function is even and symmetric as is the impulse function. The integration can be performed only from 0 to $\tau/2$ and the result multiplied by 2.

$$F(\omega) = 2 \int_0^{\tau/2} \cos(\omega t) dt \quad (3)$$

The sine function has a slope of 1 at 0 degrees phase just as the cosine of 0 phase is 1. The sine function has a slope of 0, flat top, at phase 90 degrees just as the cosine function is 0 at 90 degrees. The derivative of the sine function is the cosine function. The anti-derivative or integral of the cosine function is the sine function. Knowing the integral of the cosine is the sine function allows the Fourier transform of an impulse response to be evaluated shown as a sinc() function in Equation (4).

$$F(\omega) = 2 \int_0^{\tau/2} \cos(\omega t) dt = 2 \frac{\sin(\omega t)}{\omega} \Big|_0^{\tau/2} = \frac{\sin\left(\frac{\omega\tau}{2}\right)}{\frac{\omega}{2}}$$

$$\omega = 2\pi f \quad \tau = \frac{1}{f_0}$$

$$F(\omega) = \frac{\sin\left(\frac{\pi f}{f_0}\right)}{\pi f} \quad (4)$$

At $f=0$ the Fourier transform of the impulse function goes to zero in both the numerator and denominator. L'Hôpital's rule is needed to determine that the Fourier transform of the impulse response is $1/f_0$ when f is zero. This is the peak amplitude of the frequency response. At any frequency that is a non-zero positive or negative integer of f_0 the frequency response of the impulse is zero. We can take advantage of these periodic nulls by frequency shifting other impulses by an integer multiple of f_0 .

Figure 3 plots the frequency domain response of the impulse determined in Equation (4). The frequency response of the impulse peaks at a frequency of zero and nulls whenever the frequency is a positive or negative integer multiple of f_0 .




Figure 3 - Impulse response in the frequency domain illustrates orthogonality

By multiplying the impulse by $e^{j\omega_0 t}$ the Fourier transform is shifted in frequency. The frequency shift property of the Fourier transform is derived in Equation (5). Multiplying the impulse by $e^{j\omega_0 t}$ is a form of QAM as revealed in Equation (2). The impulse multiplied by $e^{j\omega_0 t}$ can be thought of as a tone with a frequency of f_0 . In OFDM and OFDMA, tones are QAM modulated.

The Fourier transform of the impulse response has a peak at zero and periodically nulls every multiple of f_0 both positive and negative as shown in Figure 3. Another impulse can be shifted in frequency by f_0 and added. The two impulses will overlap in both the frequency and time domain. However, in the frequency domain the first impulse peaks at a null in the second impulse while the second impulse peaks at a null in the first impulse. The two impulses are said to be orthogonal. This is the “O” in OFDM. The information contained in the amplitude of the two impulses can be demodulated by taking the Fourier transform of the combined time domain waveform and sampling the result at 0 and f_0 . We can add as many orthogonal impulses as we wish and still demodulate with a Fourier transform and sampling every f_0 .

$$\int_{-\infty}^{\infty} (f(t)e^{j\omega_0 t})e^{j\omega t} dt = \int_{-\infty}^{\infty} f(t)e^{j(\omega-\omega_0)t} dt = F(\omega - \omega_0)$$

(5)

OFDM does this slightly differently, but the principle is the same. First, the frequency domain subcarriers are QAM modulated. An inverse fast Fourier transform (FFT) converts the signal to the time domain. The receiver demodulates the multiplexed subcarriers with an FFT and sampling. The impulse function and Fourier transform used to understand the principle of orthogonal signals in an OFDM symbol are continuous in both the time and frequency domain. They represent the waveforms after the transmitter digital to analog converter and before the receiver analog to digital converter. In the digital domain both time and frequency representations are finite and discrete. The discrete Fourier transform (DFT), is a version of the Fourier transform. The fast Fourier transform, is an efficient method of calculating the DFT. The DFT is shown in Equation (6).

$$X_k = \sum_{n=0}^{N-1} x_n \cdot e^{-j\frac{2\pi}{N}kn} \quad k = 0, 1, \dots, N-1$$

(6)

Table 1 kHz tones and ns

kHz	tones	ns
20,000		50
10,000	2	100
5,000	4	200
2,500	8	400
1,250	16	800
625	32	1600
312.5	64	3200

A 64-point FFT with a time sampling of 50 ns has a subcarrier spacing of 312.5 kHz covering 20 MHz in the frequency domain and 3.2 μ s in the time domain. Table 1 illustrates the relationship of channel width, tones, and time duration. A 6 Mbps signal used for beacons, probes, and other control and management frames uses a 64-point FFT with a 3.2 microsecond time duration in a 20 MHz channel width. Forty-eight of the subcarriers carry one half of a data bit. 800 ns guard interval is inserted to prevent inter symbol interference. The bit rate of a symbol, often referred to as PHY, can be calculated $(48*0.5)/(3.2+0.8) = 6$ Mbps.

The orthogonal subcarriers can be created all at once by an access point (AP) or station client wi-fi network adapter (STA), this is OFDM. Many stations can create orthogonal subcarriers that combine at the AP to form a complete OFDM symbol. This is called uplink (UL) OFDMA. The “A” stands for access since the resource that is being shared for multiple access is tones in the frequency domain of an OFDM symbol. OFDMA was introduced in 802.11ax along with a change in the subcarrier spacing. The subcarrier spacing for an IEEE 802.11ax high efficiency, he mode, signal is one fourth that of earlier modes. With a channel width of 160 MHz the FFT size of an he mode symbol is 2048. The subcarrier spacing is 78.125 kHz and the FFT duration is 12.8 μ s. The normal guard interval is 800 ns so that the symbol time is 13.6 μ s. The details of 802.11ax subcarriers, guard time, modulation, and coding are found in reference [1] and not repeated in this paper.

Multiple access can use the time and frequency domain. The space domain can also be used for multiple user access to the radio channel in several ways. Simply separating networks in space is one way. As the coverage falls off in one wireless network another wireless network can reuse the frequency and time resources. Spatial reuse employs transmit level control and coordinated transmissions that help multiple wireless networks to use shared time and frequency resources.

When multiple transmitters and receivers are used in a channel with strong reflections from scattering objects, multiple spatial streams can send information using the same time, spectrum, and space. With two transmitters and two receivers, two spatial streams can be sent. There are four paths from the two transmit antennas to the two receive antennas. This forms a 2x2 channel matrix. If the channel matrix can be inverted the receiver can calculate the two input signals by multiplying the output signal with the inverse of the estimated channel matrix.

Likewise, with four transmitters and four receivers four spatial streams can be sent. This poses a bit of a dilemma since most phone, tablet, and notebook wireless adapters are 2x2. A 4x4 AP can send four spatial streams and yet only two spatial streams to a 2x2 STA. This is a bit of a waste of the AP capability.

In general, this is not so bad. Four spatial streams at the highest PHY rates with 4x4 AP and 4x4 STA turn out to be quite rare. A 4x4 AP and a 2x2 STA turns out to be a performance sweet spot. The combination of two spatial streams and beamforming with a 4x4 AP to a 2x2 STA works at high PHY rate at good distance in residential multipath channels. Even so, MU-MIMO comes to the rescue by allowing a 4x4 AP to send four spatial streams to 2x2 STAs. We just need to have two or more 2x2 STAs with MU-MIMO capability.

3. Types of multiuser techniques in Wi-Fi 6e 802.11ax

The wireless access point for a residential broadband service connects many customer devices. Homes with broadband service may have multiple users with each user having several devices. A family of four active users, each with both a notebook computer and smart phone, along with two television sets and many home security and automation and health devices, provides a useful canonical example. There are three dimensions in which users can share access to the wireless network: space, time, and frequency.

In the time dimension each user gets a slice of time to use the wireless network. In the downlink from AP to STA all users receive the same signal and just need to know which time slice is meant for them. This is referred to as time division multiplexing (TDM). In the uplink from STA to AP the user must contend for a time slice or be assigned a time slice so that only one STA transmits to the AP at a time. This is referred to as time division multiple access (TDMA).

In the frequency domain each user is assigned a different portion of the channel width. The channel is broken up into tones and each tone is individually modulated with data. A discrete Fourier transform implemented with a fast Fourier transform that is a mathematical technique which makes modulating many tones over a channel-width very efficient. The 160 MHz channel width of a Wi-Fi 6 signal is created with a 2048-point FFT. The tones are spaced by 78.128 kHz with an FFT duration of 12.8 μ s. In the uplink from STA to AP, OFDMA allows multiple devices to transmit at the same time with each STA assigned different tones of the FFT symbol. In the downlink from AP to STA, OFDMA transmits a full FFT symbol by the AP that is received by all STAs with each STA assigned tones. Equation (7) shows the calculation of the subcarrier spacing and FFT duration for a 160 MHz channel width and 2048-point FFT forming a Wi-Fi 6/6e he mode 802.11ax OFDMA symbol.

$$B = 160 \times 10^6 \text{ Hz} \quad N_{FFT} = 2048 \quad \Delta f = \frac{B}{N_{FFT}} = 78.128 \times 10^3 \text{ Hz} \quad T_u = \frac{1}{\Delta f} = 12.8 \times 10^{-6} \text{ s} \quad (7)$$

It is important to realize that there is no theoretical difference between time division and frequency division multiplexing regarding channel capacity or throughput. Allowing users to take turns in time or allowing users to use the channel at the same time over portions of the channel width results in the same

efficiency. The details of the use case and the channel characteristics will determine if one technique has advantages over the other for a given situation. Accessing the channel in the time domain requires quite a bit of set up. To address the hidden node problem, a request to send signal is sent, followed by a clear to send reply, followed by the data transmission, and wrapped up with an acknowledgement. An interframe spacing is required between each of these bursts. A channel busy assessment, and potentially a backoff, is required before each transmission. The transmissions themselves have a preamble so that the payload data symbols can be demodulated. Even with a PHY rate of 2400 Mbps, there is a lot of set up that makes the channel usage inefficient for small amounts of data throughput. Here is where OFDMA may have an advantage over TDMA. In the case where many users have a small amount of data to transmit, rather than go through all the set up for each one at a time, efficiency is achieved by setting up all the devices to work at the same time on small chunks of the channel width, OFDMA.

The space dimension is a bit different as it depends heavily on the reflections due to scattering objects in the home. Multiple spatial streams of information can be sent at the same time and same frequency. This is possible when there is difference in the space dimension that allows the information streams to be sorted out. In the downlink direction when the STAs are in different physical locations, beamforming may be possible to create constructive interference to the desired station and destructive interference to the undesired station. With beamforming an information stream can be sent to one station with a null towards another station. Then a second information stream can be sent to the other station with a null to the first station.

The uplink space dimension multiple access is more of a true MU-MIMO. Here each station can send an information stream to the access point. The access point receives the signal from all the stations. Multiple inputs from the station's transmitters are received by the multiple antennas of the access point. The access point then makes a channel estimate and determines the many input signals by multiplying the channel inverse matrix by the receive vector. The key difference between uplink and downlink MU-MIMO is that the downlink requires a priori beamforming while uplink all processing can be performed after signal reception.

Five multiuser techniques in Wi-Fi 6e 802.11ax work together when many users are trying to access the wireless network. OFDMA and MU-MIMO work both in the downlink from AP to STA and the uplink from STA to AP. OFDMA and MU-MIMO work together. OFDMA and MU-MIMO are key building blocks of spatial reuse. Orthogonal frequency division multiple access, OFDMA, divides the OFDM symbol with many subcarriers over the channel width into resource units assigned to individual stations.

4. DL MU-MIMO

The maximum PHY rate with four spatial streams and 160 MHz channel width is 4800 Mbps. Each spatial stream with a 160 MHz channel width at MCS11 is 1200 Mbps. Four spatial streams require a minimum 4x4 AP and 4x4 STA. Four transmit antennas and four receive antennas are needed for four spatial streams. Most wireless adapters in phones, tablets, and notebook computers are 2x2, some with 80 MHz channel width and others with 160 MHz channel width. These 2x2 devices are limited to two spatial streams. Downlink (DL) MU-MIMO enables the 4x4 AP to fully utilize four spatial streams by sending information streams to many 2x2 wireless adapters at the same time. A 4x4 AP and two 2x2 STAs has the same multiple input multiple output structure as a 4x4 transmitter sending spatial streams to a 4x4 receiver. In both cases there are four transmit antennas and four receive antennas and 16 paths from antenna to antenna.

A critical difference between a 4x4 transmitter sending MIMO spatial streams to 4x4 receiver and DL MU-MIMO is that the stations do not receive all four signals. The receivers cannot make a channel

estimate and invert the four-by-four channel matrix and then multiply by the four-vector output to determine the four input signals. Instead, the stations provide beamforming feedback based upon channel sounding that the AP can use to send multiple spatial streams that add at one station and subtract at another.

In practice we find that with a 4x4 AP four spatial streams do not often happen while three spatial streams at the highest MCS rate are common at reasonable distances. With four vertically polarized half wave dipoles spaced 5/4 wavelength apart, the optimum DL MU-MIMO is in the azimuthal direction. With stations spaced in various adjacent rooms from the AP, three spatial streams at the highest MCS rate can be measured. Table 2 shows such an example. This is a DL MU-MIMO measurement with a 4x4 AP and two wireless adapters in notebook computers with 80 MHz channel width, 2x2 MIMO configuration, and 1200 Mbps maximum PHY rate. With a 1200 Mbps maximum PHY rate, a single device can download at most between 900-1000 Mbps. The maximum PHY rate of a single he spatial stream with 80 MHz channel width is 600 Mbps. Three spatial streams have a maximum PHY rate of 1800 Mbps. With 1800 Mbps PHY rate a TCP throughput of 1200 Mbps is possible. The measurement shown in Table 2 finds a total TCP throughput of 1285 Mbps. This is an important use case in delivering Gbps broadband service to customers. For example, if the broadband service has a peak speed of 1.2 Gbps then a single 2x2 80 MHz 802.11ax wireless adapter cannot by itself deliver the full broadband speed. However, several devices together can combine to exceed 1.2 Gbps Wi-Fi speed. This works well with two, three, and four DL MU-MIMO stations. We find that three DL MU-MIMO stations are the sweet spot, having slightly higher aggregate download speed than two or four stations.

The access point is 4x4 160 MHz 4800 Mbps PHY. The stations are 1200 Mbps PHY, 2x2, 80 MHz. The 80 MHz channel width 2x2 STAs used in the measurement of DL MU-MIMO in Table 2 provide beamforming feedback consisting of 10 angles every fourth subcarrier over the full 80 MHz. This powerful beamforming results in high PHY rate and throughput in both SU and MU operation. DL MU-MIMO shown in Table 2 is 5 GHz band with 80 MHz channel width devices.

Table 2 - DL MU-MIMO Measurement 1200 Mbps PHY Stations

STA	PHY Mbps	Data Mbps	bw MHz	mcs	Nss	mu-mimo
0	617.6	450.9	80	11	1	96.9%
1	1136.7	834.1	80	10.5	2	100.0%
sum		1285.0				

In Table 3 the stations are 2400 Mbps PHY, 2x2, 160 MHz, operating DL MU-MIMO in the 6 GHz band with 160 MHz channel width. Table 3 shows a measurement sample with two 160 MHz wireless adapters using DL MU-MIMO to download TCP at over 2 Gbps. Each device alone can download at around 1.6 to 1.8 Gbps TCP with a maximum of 1.92 Gbps. DL MU-MIMO increases the number of spatial streams to three, resulting in a capacity increase and an aggregate download for two devices over 2 Gbps, something not possible in single user, SU mode.

Table 3 - DL MU-MIMO Measurement Two 2400 Mbps PHY Stations

STA	PHY	Data	Channel width	MCS	Spatial streams	MU-MIMO
	Mbps	Mbps	MHz			
0	1310.0	755.8	160	11	1	90.9%
1	2401.0	1400.8	160	11	2	99.0%
sum		2156.6				

From Wireshark captures of the beamforming feedback from the 2x2 160 MHz 802.11ax stations used for the DL MU-MIMO measurement, only a single set of ten angles is reported. The beamforming for both SU and MU do not seem to work as well for these 160 MHz devices compared to 80 MHz devices with richer beamforming feedback over the full channel width. The PHY rates and throughput of the 160 MHz channel devices are still greater than the 80 MHz devices given the doubling in the number of data subcarriers. Yet, the 80 MHz channel width devices work better than the 160 MHz devices when normalized for spectrum use.

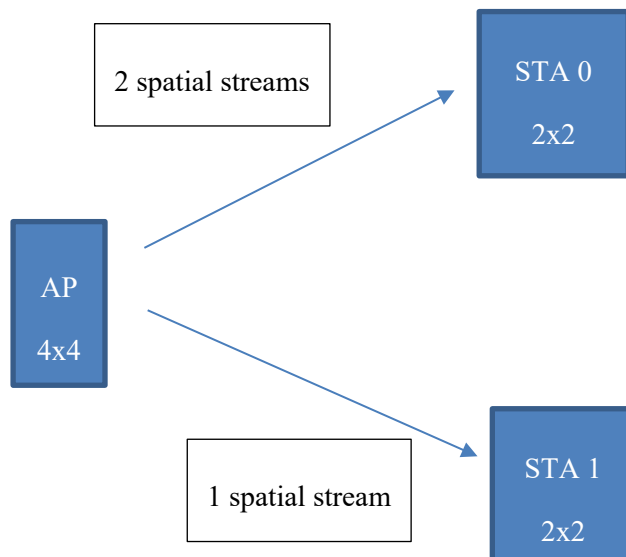


Figure 4 - DL MU-MIMO Example Block Diagram

Table 4 shows a sample measurement of DL MU-MIMO with three 160 MHz wireless adapters. In this sample, the devices were further away from the AP and the MCS rate dropped to between 8 and 9. As a result the sum of the data rates downloaded by the three devices was 1.5 Gbps, less than the maximum at the highest PHY rate, yet still an improvement over SU operation at these distances. DL MU-MIMO increases capacity of the wireless networks only when several devices have large files to download at the same time from Internet servers capable of Gbps downloads with all devices in high SNR conditions with good multipath reflections.

Table 4 – DL MU-MIMO Measurement Three 2400 Mbps PHY Stations

STA	PHY	Data	Channel width	MCS	Spatial streams	MU-MIMO
	Mbps	Mbps	MHz			
0	1696.8	782.4	157.5	8	2.0	97.7%
1	1246.3	595.3	157.9	8.7	1.3	69.2%
2	959.2	180.5	159.2	9.0	1.0	99.35
sum		1558.2				

DL MU-MIMO requires quite a bit of traffic demand from several capable devices at close range. Applications do not often generate enough traffic for DL MU-MIMO packets to be sent. An application downloading a large movie on a smart phone would seem to be ideal for MU-MIMO. However, when tried, the application downloaded the movie to the phone at 10 Mbps, not enough for DL-MU-MIMO.

5. UL MU-MIMO

Uplink multiuser multiple-input multiple-output allows several stations to transmit spatial streams at the same time. The access point collects the signals from each station and sorts out the spatial streams. UL MU-MIMO is pure MIMO whereby the receiver can process all the streams after reception rather than relying on a priori adaptive beamforming as in the case for DL MU-MIMO.

UL MU-MIMO has the advantage of significant spatial diversity and angle of arrival of spatial streams as the stations transmitting the streams can be in different locations relative to the access point.

Multiple stations can each transmit a spatial stream to the AP at the same time. The AP has multiple antennas and receivers. Each of the AP antennas receive a combination of all the signals from the stations. The spatial streams of independent information are all mixed together at each AP receiver. Without strong multipath in the channel and good channel estimation at the AP, the information streams from each station would be hopelessly jumbled. Each antenna of the stations to each antenna of the AP has an impulse response defined by the attenuation and delay of the direct path and various reflections from scattering objects.

The Fourier transform of the impulse response from antenna to antenna in the presence of multipath reflections reveals a frequency response that would need to be characterized in a single carrier system. OFDM comes to the rescue for UL MU-MIMO since each tone is narrow and can be considered as a flat frequency response. Thus, each antenna-to-antenna path is treated on a tone-by-tone basis and characterized by a single complex number per tone representing the Fourier transform of the impulse response. In practice the channel matrix is adequately characterized by the phase differences between the various paths.

UL MU-MIMO the AP estimates the channel matrix, inverts the channel matrix, and multiplies the inverted channel matrix by the received vector to determine the input signals from each of the stations. Thus, all the information streams from the stations can be decoded.

6. DL OFDMA

Downlink orthogonal frequency division multiple access divides the spectrum into tones and assigns resource units consisting of continuous tones to multiple stations. In general, with heavy traffic, DL OFDMA does not increase capacity. DL MU-MIMO increases channel capacity when stations are in a high signal to noise ratio (SNR), zone and enough traffic demand is there to fill the pipe. DL OFDMA is complementary in the sense that it helps when stations are in a low SNR zone with small amounts of data to transmit. Here SNR is defined as the average power of the transmitted signal at the receiver demodulator divided by the average noise and interference level from all sources expressed in dB. DL OFDMA does not improve the link budget in the same way as UL OFDMA. While stations in theory could take advantage of the lower noise floor of a smaller resource unit allocation, devices tend to demodulate the full downlink channel width and do not measure improved sensitivity due to DL OFDMA.

DL OFDMA is observed when devices are far enough away to forfeit DL MU-MIMO capacity increase and when traffic demand is low. One of the big advantages observed with DL OFDMA relates to the scheduler. With SU and many devices competing for wireless resources there is a tendency for one device to dominate the throughput. With several devices downloading at the same time one device may have very high throughput while the other devices are stuck at low throughput. DL OFDMA is much better at sharing the channel with each device downloading at close to the same throughput.

DL MU-MIMO takes a heavy traffic demand from multiple devices able to utilize DL MU-MIMO and a considerable amount of time. Typically, about 200 Mbps or more from several devices for over one minute is required to see DL MU-MIMO packets. By contrast, DL OFDMA packets make up a large proportion of traffic even with small amounts of traffic and the setup is very quick.

7. UL OFDMA

Uplink orthogonal division multiple access allows multiple stations to transmit at the same time with each using a different part of the channel width. UL OFDMA improves the link budget in the uplink. Each station can only transmit so much power. The maximum transmit level in the 2.4 GHz band for a typical wireless adapter in a notebook computer ranges from +17 to +21 dBm per chain with two chains in a 20 MHz channel width. Restrictions on out of band radiated emissions often results in different transmit power on some channel settings. 5 GHz band station transmit level typical ranges from 10 to 18 dBm per chain with two chains. The wireless adapter in the notebook computer used for 160 MHz channel width measurements in this paper is a type accepted with 5 dBi antenna gain in the 5 GHz band and transmit level of about +15 dBm per chain with two chains for a 160 MHz HE0 signal at 5570 MHz, +17 dBm per chain for an 80 MHz HE0 signal at 5530 MHz, +17 dBm per chain for a 40 MHz HE0 signal at 5510 MHz, and +18 dBm per chain for a 20 MHz HE0 signal at 5500 MHz. Note that the power spectral density and thus the received SNR, increases when the channel width decreases. The 160 MHz channel width case has 12 dB lower signal to noise ratio at the AP receiver than the 20 MHz channel width, 9 dB due to higher receiver noise level at 160 MHz compared to 20 MHz, and 3 dB due to the increase in transmit level at 20 MHz compared to 160 MHz. Therefore, the uplink channel width will drop to 20 MHz at the cell edge in the 5 GHz band. In the 6 GHz band for low power indoor operation the effective isotropic radiated power (EIRP) of the power spectral density is regulated and the behavior at cell edge is different. Figure 5 shows a measurement of the uplink dynamic range with PHY rate improving with the received signal strength indicator (RSSI).




Figure 5 - Dynamic Range in the uplink

Table 5 shows a sample of measurement results with two computers uploading at the same time in the 6 GHz band. Both the AP and the STA fall below the restrictions for low power indoor operation. The coverage is good throughout the home used for the measurement. Shown in Table 5 are cell edge conditions with stable connection and throughput, but connectivity will be lost at further distance between the AP and STA or more obstructions. Both stations are 2x2 wireless adapters in notebook computers with 160 MHz channel width capability and 2400 Mbps maximum PHY rate. The channel width remains at 160 MHz since under the regulatory limit of -1 dBm EIRP per MHz the uplink does not benefit from reduced channel width. UL OFDMA accounts for all the uplink packets with about 1000 tones allocated to each station. The PHY rate of each station is about 70 Mbps and the total upload rate was measured at 49 Mbps. Both stations sent two spatial streams at the lowest MCS rate of 0.

Table 5 – UL OFDMA Measurement Two 2400 Mbps PHY Stations 6 GHz band

STA	rssi	PHY	Data	Channel width	MCS	Spatial streams	ofdma	tones
	dBm	Mbps	Mbps	MHz				
0	-85	72.5	32.9	160	0	2	100%	1002.7
1	-85	73.5	16.0	159	0	2	100%	1014.3
sum			48.9					

Table 6 shows a sample measurement in the 5 GHz band with two 2x2 80 MHz stations having a maximum PHY rate of 1200 Mbps. Both stations are notebook computers using a common Unix based operating system. The sample measurement shown in Table 6 is the lowest observed UL OFDMA RSSI levels at the AP while making measurements of the two computers uploading data throughout a two-story residential home as well as outside the home with coverage up to 100 meters. Both stations have an uplink RSSI at the AP of -86 dBm. The channel width was 80 MHz with each station assigned about 484 tones with 100% UL OFDMA packets. The total throughput was 26.4 Mbps with PHY rates of 35 and 47 Mbps. Both stations transmitted two spatial streams at the lowest MCS level of 0.

Table 6 – UL OFDMA Measurement Two 1200 Mbps PHY Stations 5 GHz band

STA	rssi	PHY	Data	Channel width	MCS	Spatial streams	ofdma	tones
	dBm	Mbps	Mbps	MHz				
0	-86	35.2	12.7	80	0	2	100%	491
1	-86	47.4	13.7	79	0.3	2	100%	497
sum			26.4					

The 160 MHz stations in the 6 GHz band utilized UL OFDMA down to a receive level at the AP of -86 dBm. The 80 MHz stations in the 5 GHz band utilized UL OFDMA down to a receive level at the AP of -86 dBm. Is this good? Why were these the lowest observed levels? To answer these questions and better understand what is going on, let us review noise floor and Boltzmann constant. The thermal noise floor of the receiver is determined by Boltzmann constant and the noise figure of the low noise amplifier.

Ludwig Boltzmann developed statistical mechanics to understand the physics of steam engines. Boltzmann constant relates temperature to energy. Boltzmann constant applies equally well for determining the thermal noise in a receiver due to the vibrations of electrons. The thermal noise density is -174 dBm per Hz at room temperature of 300 Kelvin, equal to 25 degrees Celsius and 77 degrees Fahrenheit.

$$E = k_b T \quad (8)$$

$$k_b = 1.380649 \times 10^{-23} \quad (9)$$

$$N_0 = 10 \log_{10}(1000 \cdot 1.380649E - 23 \cdot 300) \quad (10)$$

$$N_0 = -174 \frac{\text{dBm}}{\text{Hz}} \quad (11)$$

The receiver noise floor referenced to the antenna input terminals depends upon the thermal noise density, the noise bandwidth, and the noise figure of the low noise amplifier. A 3 dB noise figure is a reasonable expectation for an AP allowing for a 1 dB noise figure low noise amplifier and 2 dB of losses through the switch, filters, and traces to the antenna.

$$N = N_0 + F + 10 \cdot \log_{10}(B) \quad (12)$$

$$B = 20 \text{ MHz}$$

$$N = -174 + 3 + 10 \cdot \log_{10}(20E6)$$

$$N = -98 \text{ dBm} \quad (13)$$

The receiver noise floor at room temperature with a 3 dB noise figure receiver in a 20 MHz channel width is -98 dBm.

Table 7 – Receiver noise floor and channel width for 3 dB noise figure at room temperature

B	N
MHz	dBm
20	-98
40	-95
80	-92
160	-89

In the 6 GHz band for low power indoor operation the AP is restricted to +5 dBm EIRP per MHz and the STA is restricted to -1 dBm EIRP per MHz. The AP can transmit +16 dBm per chain with four chains, 2.15 dBi half wave dipole antenna elements, 3 dB beamforming gain with two spatial streams.

Table 8 – EIRP regulation in the US for 6 GHz low power indoor access points

EIRP per MHz	5.00	dBm/MHz
channel width	160.00	MHz
channel width	22.04	dBMHz
EIRP	27.04	dBm
Antenna Gain	2.15	dBi
Beamforming Gain	3.00	dB
total transmit level	21.89	dBm
chains	4.00	
transmit level per chain	15.87	dBm per chain

The station EIRP per MHz in the 6 GHz band for low power indoor operation is -1 dBm per MHz. A 2x2 station with half wave dipole antenna element gain is limited to +16 dBm per chain transmit level for a 160 MHz channel width signal.

Table 9 – EIRP regulation in the US for 6 GHz low power indoor stations

EIRP per MHz	-1.00	dBm/MHz
channel width	160.00	MHz
channel width	22.04	dBMHz
EIRP	21.04	dBm
Antenna Gain	2.15	dBi
Beamforming Gain	0.00	dB
total transmit level	18.89	dBm
chains	2.00	
transmit level per chain	15.88	dBm per chain

Most stations transmit 6 to 9 dB lower than the levels shown in Table 9. The phones used to make the measurements of MU-MIMO and OFDMA in the 6 GHz band operate at +7.5 dBm per chain with two chains and antenna element gain of -6 dBi. The notebook computers used to measure MU-MIMO and OFDMA operate at +10 dBm per chain with two chains in the 6 GHz band with 160 MHz channel width.

The wavelength is calculated by dividing the speed of light by the frequency. In Equation 14, c is the speed of light, f is the frequency in Hz, and λ is the wavelength in meters. Table 10 shows the wavelength in mm for 2.4 GHz band channel 1, 5 GHz band channel 100, and 6 GHz band center frequency of a 160 MHz channel width signal having primary steering channel 69.

$$c = 299,792,458 \text{ m/s speed of light} \quad (14)$$

$$\lambda = \frac{c}{f} \quad (15)$$

Table 10 – Relationship between frequency and wavelength for 2.4, 5, 6 GHz bands

frequency	wavelength
MHz	mm
2412	124.3
5500	54.5
6345	47.2

The free space path loss at a reference distance is denoted A_0 measured in dB. The reference distance is denoted by d_0 measured in meters. A_0 is calculated as a function of d_0 and λ as shown in Equation (16).

$$A_0 = 20 \cdot \log_{10} \left(\frac{4\pi d_0}{\lambda} \right) \quad (16)$$

A log normal path loss model can be applied for distances beyond the reference distance. L is the log normal path loss in dB. A_0 is the free space path loss at the reference distance d_0 in meters. The distance between the AP and STA in meters is d . L_{floors} is the loss through floors and L_{walls} is the loss through walls. N_σ is a log normal probability distribution with standard deviation σ , typically around 3 dB. Equation 17 shows the IEEE 802.11 log normal path loss model for indoor residential channel.

$$L = A_0 + 35 \cdot \log_{10} \left(\frac{d}{d_0} \right) + L_{floors} + L_{walls} + N_\sigma \quad (17)$$

The Friis transmission equation was derived in 1945 at Bell Labs by Harald T. Friis. R is the received level in dBm, $EIRP$ is the effective isotropic radiated power equal to the transmitter power level plus the transmitter antenna gain, L is the log normal path loss in dB, G is the antenna gain of the receiver. A version of the Friis transmission equation is shown in Equation (18).

$$R = EIRP - L + G \quad (18)$$

Table 11 shows the distance d in meters for a 6 GHz band signal with channel width 160 MHz for each MCS level along with the PHY rate in Mbps and the maximum TCP throughput in the uplink for a low power indoor station at -1 dBm EIRP per MHz. This table does not include wall and floor loss or the log normal probability distribution factor. As can be seen the distance is quite far at 132 meters.

The wall and floor loss is a wild card. Often the wall and floor loss is not that much. Drywall attenuation may only be around 1 dB. Plaster with wire mesh can be 20 dB attenuation or more. Brick and concrete can have high attenuation. Even some glass sliding doors can have high attenuation. The indoor residential IEEE 802.11 path loss model uses a floor attenuation that is quite high at 20 dB for a single floor. This is much higher than often observed in practice. The wall attenuation used in the calculation references the IEEE 802.11 model of 5 dB per wall. Using these factors for the case of one floor and two walls the additional 30 dB attenuation reduces the coverage distance to 18 meters for the lowest MCS rate. Eighteen meter coverage is still not bad considering the high wall and floor attenuation.

An important point to notice is MCS 6 at a distance of 43 meters. The receive level in the uplink is -69 dBm. The PHY rate is 1297 Mbps and the TCP throughput is 1037 Mbps. This is the lowest level that will still deliver over 1 Gbps speed to the customer.

Table 11 – Distance, receive level for 160 MHz channel width uplink low power indoor station EIRP -1 dBm per MHz

MCS	Rx	PHY	TCP	d 6 GHz
	dBm	Mbps	Mbps	m
0	-86	144	115	132
1	-83	288	230	108
2	-81	432	345	95
3	-78	576	461	78
4	-74	864	691	60
5	-70	1152	922	46
6	-69	1297	1037	43
7	-68	1441	1152	40
8	-63	1729	1383	29
9	-61	1922	1538	25
10	-58	2161	1729	21
11	-56	2401	1920	18

8. Spatial Reuse

There are two different types of spatial reuse. PD uses power detection. SRP uses spatial reuse parameters contained within an OFDMA trigger.

Power detection spatial reuse lowers transmit level based upon neighbor's RSSI. The channel busy threshold for an 802.11 signal is -82 dBm. The energy detection threshold for a non-802.11 signal is -62 dBm. Oftentimes real devices deviate from the threshold levels called out in the IEEE 802.11 standards.

If the RSSI of a neighbor device, be it AP or STA, exceeds -82 dBm then the channel is determined to be busy. After an exponential time back-off the transmitter will be checking for channel busy once more, hoping for an opportunity to transmit when no non-802.11 signal exceeds -62 dBm and no 802.11 signal exceeds -82 dBm. Then the AP or STA will transmit.

Power detection spatial reuse allows the AP or STA to transmit even when an 802.11 neighbor signal exceeds -82 dBm. To reuse the time and frequency resources efficiently in two different locations the transmit power is reduced in proportion to the neighbor RSSI. If the neighbor RSSI is -72 dBm then the

transmit level is reduced by 10 dB. If the neighbor RSSI is -62 dBm then the transmit level is reduced by 20 dB.

In effect, power detection spatial reuse keeps the interference toward a neighbor wireless network no worse than full power with -82 dBm threshold while allowing for more transmit opportunities. The power detection method will work with any device, even devices without spatial reuse capability.

The SRP method is much more interesting, and powerful. SRP requires all participants to have spatial reuse capability. The UL OFDMA trigger of a neighbor contains the information needed to reuse time and frequency resources in different locations without interference. The path loss to the neighbor can be calculated by subtracting the measured RSSI from the neighbor transmit level read in the trigger. The trigger will advertise an acceptable level of interference to the UL OFDMA transmission. A transmit level can thereby be calculated that can work over top the neighbor UL OFDMA signal without interference.

9. Putting It All Together With Real Devices and Applications

Consider an illustrative use case with a phone, a notebook computer, and two television sets. All four of these devices have the same Wi-Fi capabilities, including multiple user. The four stations have a maximum PHY rate of 1200 Mbps. There are many different types of Wi-Fi clients. Some have 160 MHz channel width while others only operate at 80 MHz channel width. The devices tested for MU operation with a mix of traffic were 80 MHz channel width wireless adapters, which is very common. Each device has different application needs and different throughput capabilities. A mix of devices is shown in Figure 6.

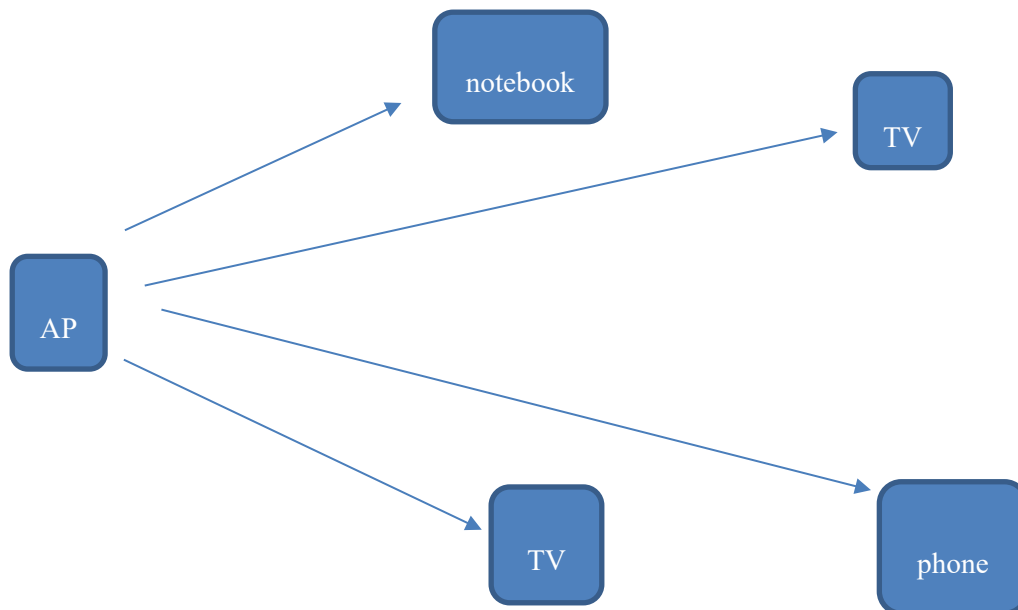


Figure 6 - Traffic mix

The measurement procedure begins with verifying broadband service from each of the devices used for the test. The cable modem wireless router used has four Ethernet LAN ports, three 1 Gbps speed and one 2.5 Gbps speed. A computer with a 2.5 Gbps USB-C network adapter is connected to the 2.5 Gbps Ethernet LAN port of the router. A speed test taken on the computer connected to the 2.5 Gbps Ethernet LAN port of the router measured 1.4 Gbps download. This verifies that this computer has a broadband connection at the full speed of the service.

Next, a computer with a 6 GHz band, 2x2 MIMO, 160 MHz channel width, 2400 Mbps PHY rate is connected to the 6 GHz band radio of the cable modem wireless gateway and a speed test is run at close range. The locations necessary for the highest PHY rate of 2400 Mbps are in the same room, and adjacent rooms on the same floor, or directly upstairs or downstairs of the AP. The speed test, although not as consistent as the DOCSIS to Ethernet speed, measured 1.3 Gbps download speed. This verifies the full broadband speed from DOCSIS WAN to 6 GHz radio.

While it is convenient and sometimes useful for testing purposes to measure from Ethernet LAN to Wi-Fi, it is imperative to include DOCSIS to Wi-Fi testing and verification since this accounts for almost all customer traffic. The phone and computer used in the mixed traffic MU testing measured download speed tests between 800 and 900 Mbps. The phone and computer wireless adapter were 2x2 MIMO, 80 MHz channel width, 1200 Mbps PHY. The set up shown in Figure 7.

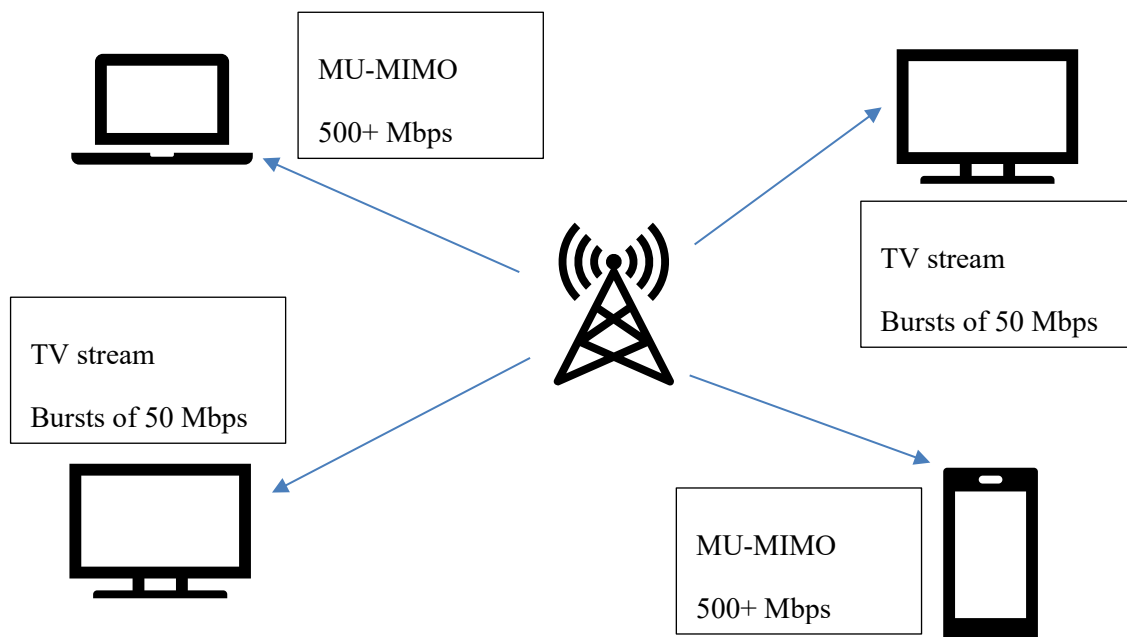


Figure 7 - DL MU-MIMO and UL OFDMA traffic mix

Generating traffic demand to form DL MU-MIMO required using iperf3 on the client phone and notebook computer with a server connected to the 2.5 Gbps Ethernet port of the cable modem wireless router. Traffic exceeding 1 Gbps was measured even though no individual station was capable on its own to download at 1 Gbps. Thus, DL MU-MIMO allowed the total of multiple devices working together to exceed that of any one device.

In addition to the download of large files to the computer and phone, two television sets with the same wireless adapters were used for streaming video. The statistics of the measured throughput listed in Table 12 and the graph of throughput for each device over time is shown in Figure 8. The television sets had bursty traffic with peaks of about 50 Mbps and average of 10 Mbps. The computer and phone had aggregate download speed that exceeded 1 Gbps employing DL MU-MIMO and UL OFDMA at the same time as the two video streams operating without impairment.

Table 12 - Statistics of Throughput with Four MU devices

Stat	phone	computer	TV1	TV2	Sum
average	455.9132	407.8673981	15.32821	8.530408	887.6392
std	126.2928	208.0782331	9.354155	14.13454	152.5777
max	917.4	785.2	49.7	64.5	1212

While these measurements focused on DL MU-MIMO, UL OFDMA was used for the lower bandwidth upload response. This is a good example of the various multiple user techniques working together for a better overall user experience.




Figure 8 - DL MU-MIMO and UL OFDMA Measured traffic mix

As seen in Figure 8 the sum of traffic often exceeds 1 Gbps and even peaks at 1.2 Gbps. This is a good example of MU techniques helping to deliver Gbps broadband service when devices work together. Alone, none of these devices can use the full broadband service. Yet working together, the customer receives the benefit of the full speed of the broadband service.

10. Conclusion

UL MU-MIMO and spatial reuse are promising technologies customers can look forward to with newer devices and upgraded drivers with Wi-Fi 6e/7 wireless adapters. Common Wi-Fi 6 wireless adapters in notebooks, phones, and tablets have 2x2 MIMO with 80 MHz channel width and a maximum PHY rate of 1200 Mbps with drivers that have enabled DL MU-MIMO and both UL and DL OFDMA. When more than one of these devices uses the wireless network at the same time these multiple user technologies kick in. At close range with high traffic demand from more than one such device downlink frames are MU-MIMO while uplink frames are OFDMA. Further away, under low signal to noise ratio conditions, the downlink and uplink frames are OFDMA. The beamforming feedback in these devices provides the ten angles every fourth subcarrier over the full 80 MHz channel width. DL MU-MIMO provides aggregate throughput that is not possible with SU operation. OFDMA improves the uplink link budget, not at the very channel edge where connectivity is soon lost, but increased PHY rates and aggregate throughput in the upload for signals as low as -86 dBm in an 80 MHz OFDMA channel.

Wireless adapters are readily available for notebook computers and some phones and tablets with a maximum PHY rate of 2400 Mbps. These devices have 160 MHz channel width in the 5 and 6 GHz band with 2x2 MIMO. Getting the full broadband service of 1.2 Gbps or higher is readily delivered in SU mode. With a maximum PHY rate of 2400 Mbps, there is a healthy margin of error to deliver 1.2 Gbps service, the MCS rate can drop to 6 when the channel is unused by neighbors or the PHY rate can be close to the maximum and still deliver 1.2 Gbps while sharing the spectrum with neighbors. These devices take full advantage of DL and UL OFDMA when several devices are active at the same time. These devices also can improve download aggregate throughput with DL MU-MIMO. Download speed from two 2400 Mbps PHY devices with DL MU-MIMO exceeded a throughput of 2.1 Gbps. This is not possible with SU operation. Still, this only worked with both computers in the same room and even then, only occasionally. The beamforming feedback from the 160 MHz channel width devices was not sufficient for robust DL MU-MIMO. Still, with the SU mode able to deliver 1.2 Gbps speeds with plenty of margin and both DL and UL OFDMA working well, this does not seem to be much of a concern for overall user experience. Certainly, customers will always benefit from 160 MHz channel width devices.

Finally, the benefits of DL MU-MIMO and UL OFDMA were illustrated with two television sets streaming video while a phone and a tablet downloaded TCP at maximum throughput. The video streaming works flawlessly with periodic bursts of throughput of about 50 Mbps averaging out at about 10 Mbps. Neither the phone nor the computer is capable of downloading over 1 Gbps since their maximum PHY rate is 1200 Mbps. Yet, the sum of the throughput of all four devices quite often exceeded 1 Gbps thanks to DL MU-MIMO working in tandem with UL OFDMA.

Abbreviations

AP	access point
bps	bits per second
DFT	discrete Fourier transform
DL	downlink
DOCSIS	Data over cable service interface specification
EIRP	Effective isotropic radiated power
FDM	Frequency division multiplex

FFT	Fast Fourier transform
HD	high definition
he	High efficiency
Hz	hertz
K	kelvin
Mbps	Mega (million) bits per second
MIMO	multiple input multiple output antennas
MU	multiple user
OFDM	Orthogonal frequency division multiplex
OFDMA	orthogonal frequency division multiple access
PD	Power detect
PHY	Physical the bit rate of a single OFDM symbol
RSSI	Received signal strength indicator
SNR	Signal to noise ratio
STA	Station client Wi-Fi network adapter
SCTE	Society of Cable Telecommunications Engineers
SRP	Spatial reuse parameter
SU	Single user
TDM	Time division multiplex
TDMA	Time division multiple access
UL	uplink

Bibliography & References

[1] The Importance of Wi-Fi 6 Technology for Delivery of Gbps Internet Service, David John Urban, Comcast, Cable-Tec Expo 2019 SCTE ISBE 2019 Fall Technical Forum

How the New ANSI/SCTE 275 Grounding and Bonding Standard Can Improve Your Network Resiliency and Continuity

A Technical Paper prepared for SCTE by

Mike Glaser

Principal Engineer – Critical Facilities

Cox Communications

6305 Peachtree Dunwoody Road, Atlanta GA 30328

404-269-0143

mike.glaser@cox.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Development And Benefits Of The Standard	3
3. Main Purposes Of Grounding And Bonding.....	4
4. Exterior Grounding	5
4.1. Soil Resistivity	6
4.2. The Ground Ring.....	6
5. Interior Grounding	7
5.1. Various Grounding System Methods/Theories/Approaches.....	7
5.2. Common Grounding And Bonding System Requirements	8
5.3. The Master Ground Bar	9
6. Surge And Lightning Protection	10
7. Humidity And Electrostatic Discharge	11
8. Commissioning And Maintenance.....	11
9. Conclusion.....	12
Abbreviations	12
Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 – Exterior Ground Ring Conductor Exothermically Welded To A Ground Rod	7
Figure 2 – Acceptable Interior Grounding Methods - Network Equipment Racks	8
Figure 3 – Master Ground Bar With PANI Configuration	10

1. Introduction

Network reliability and resilience are usually only as effective as the weakest component in the entire network system. The words ‘reliability’ and ‘resilience’ are used in various industries, and have been used more frequently in recent years. They may have slightly different meanings depending on the context, but with respect to network systems, reliability means “it’s there when we expect it to be”. Resilience is a bit more multi-faceted; it means “how well are we prepared for, and how easily and quickly can we bounce back from extreme or prolonged events”. Both require significant investments.

Another way to think of the relationship is this: reliability is the outcome and resilience is how you achieve that outcome. Reliability can be expressed as an availability key performance indicator (KPI), say 99.999% uptime, while resilience is all the detailed technical planning and design needed to achieve that, and for a critical facility, one of the main components is a good grounding and bonding system.

Knowing that critical facilities are foundational to the network and the delivery of lifeline telecommunications services, it is essential that all of their infrastructure systems are fully functional at all times. Cable operators always focus on maintaining robust and redundant systems, but how often do they overlook the proper grounding and bonding of these infrastructure systems?

A single electrical storm or utility surge can wreak havoc on network equipment in a critical facility, often creating significant performance issues that cost valuable time and money to repair, and can potentially result in a loss of customers. Human error also contributes significantly to these unfortunate, but largely avoidable consequences.

The proper grounding and bonding of electrical systems and network equipment in a critical facility is an indispensable element in the assurance of business continuity in the broadband and telecommunications industry, therefore high priority must be given to the grounding and bonding system from design and installation to commissioning and ongoing maintenance.

ANSI/SCTE 275 2021 brings together the best practices and cable industry expertise and experience, as it endeavors to guide cable operators in the deployment of grounding and bonding systems that will enable them to meet their reliability targets and ensure resilience of the network.

This paper will give an overview of the development and contents of the ANSI/SCTE 275 2021 standard, explain how it can improve network resilience, and provide some guidance on how best to deploy a grounding and bonding system that will ensure just that.

2. Development And Benefits Of The Standard

In early 2018, it was determined that cable operators could greatly benefit from an aggregation of grounding and bonding best practices at the very least, and perhaps even an industry standard in this area. A grounding and bonding task force was formed to explore the opportunities around various aspects of telecommunications grounding and bonding practices.

The task force met over a period of three years, and reviewed nearly 100 existing grounding and bonding standards currently being used by cable operators, as well as those from electrical, fire safety, and telecommunications carrier industries. There were contributions to the standard that eventually became ANSI/SCTE 275, “Electrical Grounding and Bonding for Cable Broadband Network Critical Facilities” from Cox, Comcast, Rogers, Shaw, and Charter.

There were a number of challenges, benefits and opportunities that came to light while this project was ongoing:

- Due to the varying histories, organizational compositions and customer services both common and unique to these cable operators, there were consequent variations in intent and approach, as well as the necessary level of detail, within the existing cable operator grounding and bonding standards. For example, there are inherent differences in the AC power and DC power grounding and bonding systems design, data center and headend design, vendor-driven requirements, and the hodgepodge of practices found in legacy acquisition sites. Additionally, the presence of a telephony switch would necessitate particular practices around isolated grounding specifications that are not present (or needed) in other critical facilities.
- To assist with collecting and aggregating the similarities and differences between cable operator specifications, a ‘matrix’ document comparing the specific practices and requirements (or lack thereof) in each grounding and bonding category (or sub-category) for each cable operator was developed and maintained.
- ANSI/SCTE 275 is the first grounding and bonding standard specifically developed for cable operators in the telecommunications industry. It brings together the best practices from various industry sectors, and outlines minimum standards for ensuring network reliability and continuity of cable operators service delivery to customers.
- Additional benefits of ANSI/SCTE 275 are as follows:
 - It lays the groundwork for safety and reliability, and is a key component to the resilience of critical facilities in all grounding and bonding system lifecycle phases - from the engineering design phase, through installation, then on to commissioning and periodic maintenance. By following this standard, cable operators can help assure business continuity.
 - The standard can be used to ensure vendors are using a standardized approach to grounding and bonding systems through each of the lifecycle phases.
 - A commissioning checklist is included in the appendix, which can assist engineering and operations staff with ensuring the grounding and bonding system is properly installed.
 - Proper grounding and bonding are foundational to personnel safety and network reliability, yet it is often the least understood system in a critical facility. ANSI/SCTE 275 will be able to be used as an educational tool for engineers and technicians, and better enable them to ensure continued reliable operation of all infrastructure systems within the critical facility.

3. Main Purposes Of Grounding And Bonding

The main purposes of grounding and bonding in general, and at a high level, are all about safety of personnel and reliability of network and infrastructure equipment systems. The essential elements or qualities for a good grounding and bonding system are primarily that it has a low impedance, and secondarily that it be dependable throughout the lifetime of the critical facility. This is typically ensured through properly sized grounding and bonding cables, short cable runs without sharp turns, and connections that are resistant to corrosion or loosening over time.

There are five principal objectives for providing a dependable low impedance grounding and bonding system. These include:

- 1) Personal Safety: From a “safety first” perspective, we certainly want to minimize the development of any electrical potential that could create a shock hazard to personnel, within and between metallic frames and structures.
- 2) Equipment Protection: It is critical, and therefore required, that adequate fault current paths be provided so any installed overcurrent devices can disconnect faulted circuits to reduce the possibility of fire and limit damage to the critical equipment. System design and operational procedures for electrostatic protection of equipment is also required.
- 3) Equipment Operation: Proper grounding provides an equalized ground reference to electronic communication circuits connected to the ground plane.
- 4) Electrical Noise Reduction: Proper bonding assists in the reduction of electrical interference by maintaining low impedance paths between ground points throughout the communication system.
- 5) Reliability: Naturally, to protect our investment, we should install a grounding system that resists deterioration, requires minimal maintenance, and ensures network service continuity.

Cable operators must of course meet the basic requirements of applicable national and local electrical and safety codes, and comply with any additional requirements set forth by the local authority having jurisdiction (AHJ).

Telecommunications services providers - including cable operators - recognize that these minimum requirements are often not stringent enough to fulfill the industry-specific protection needs of their modern critical facilities, and so go far beyond the minimum requirements as needed to meet the principal objectives and assure business continuity.

4. Exterior Grounding

Although unwanted electrical disturbances may come from both internal and external sources, the most egregious events come from outside the critical facility. Primarily for that reason, and to provide a commonsense flow to the outline and description of the various elements of a grounding and bonding system, the ANSI/SCTE 275 standard is organized from an “outside-to-inside” perspective. This paper will address the important aspects of grounding and bonding in the ANSI/SCTE 275 in the same fashion.

The exterior grounding system establishes a direct path of known low impedance for all power, communication, and signal systems. The exterior grounding system also provides a path to earth for the discharge of lightning strikes, restricts step and touch potential gradient in the area accessible to persons, and assists in the control of electrical noise in signal and control circuits by minimizing voltage differentials, ensuring network continuity through any of these events. Currents associated with power faults and noise typically use this path to return to their source, whereas currents associated with lightning use this path to equalize the charge established on the earth’s surface during storms.

Any ground system installation, whether commercial, industrial, or residential, requires a certain target resistance-to-earth value be met. The industry standards vary on acceptable and maximum values, and

environmental or geographic conditions will affect the application of some of these grounding systems. In general, the lower resistance the better, and the target for cable operators is typically 5 Ohms or less, and 25 Ohms would be considered a maximum level.

4.1. Soil Resistivity

The first task in the design of a good ground system is to test the resistivity of the soil surrounding the critical facility. The ground system design and associated components can only be determined after we know what we have to work with. For the soil resistivity, as is the case for the ground system resistance, a lower value is better.

The resistivity (typically measured in Ohm-m) of a soil depends largely on the types of materials that are contained in the soil, as would be the case for metals and any other materials commonly considered conductive. Soils can contain various materials including clays, sands, rocks, shale, etc. It is this combination of materials that largely determines how conductive the soil will be to electric current flow. In addition, the resistivity of a specific type of soil will vary seasonally because of its moisture, temperature and chemical content:

- The greater the moisture content, the lower the resistivity.
- The higher the temperature, the lower the resistivity.
- The greater the salt content, the lower the resistivity.

Though there are several acceptable methods for measuring soil resistivity, the one most commonly used is the 4-point method. In this method, four small-sized (typically 18-inch) electrodes are driven into the earth at the same depth, in a straight line and equidistant from each other. A known current from a constant current generator is passed between the outer electrodes. The voltage drop, which is a function of the resistance, is then measured across the two inner electrodes. Greater detail on the application of best practices for measuring soil resistivity are described in the latest IEEE Standard 81, “Guide for Measuring Earth Resistivity, Ground Impedance, and Earth Surface Potentials of a Ground System”.

4.2. The Ground Ring

Once the soil resistivity has been measured, the external ground system can be designed to the specific resistance-to-earth target value. This is accomplished mainly by the ground ring and driven rod electrodes, and depending on the characteristics and limitations of the surrounding earth, may include flat plates, rebar, Ground Enhancement Material (GEM) or chemical ground rods, which are typically hollow copper rods filled with an electrolytic salt mixture.

ANSI/SCTE 275 spells out the best practices and minimum standards for cable operators to go by in specifying an external ground system. The ground ring is typically a 2/0 AWG (minimum #2 AWG) copper conductor buried 30 inches below grade or below the frost line, whichever is greater. The ground rods should be ¾ inch x 10 ft long, buried to the same depth as the ring, spaced 20 feet apart, and exothermically welded to the ground ring. Additionally, there should be several inspection wells or hand holes at key locations around the ring.

The external ground system is typically best designed by engineering companies that specialize in this, and there are quite a few good software programs that are utilized for this purpose; however most electrical engineering firms should be capable of designing an adequate ground system.

Finally, ANSI/SCTE 275 gives direction for the proper grounding of exterior objects like towers, guy wires, satellite dishes, masts, fencing and other metallic objects, all of which should be tied to the external ground ring.



Figure 1 – Exterior Ground Ring Conductor Exothermically Welded To A Ground Rod

5. Interior Grounding

The interior grounding system is the most complex portion of a critical facility grounding and bonding system. It also has the most variations in application between the cable operators, and between the site types (i.e. headend, data center, edge hub site), not to mention the variances introduced by legacy and acquisition sites. While there are good established references there was no detail specific to this industry that could be seen as a standard prior to ANSI/SCTE 275 being released.

Still it was possible to detail minimum standards as well as recommended targets for best practices, knowing that cable operators must weigh costs to upgrade against benefits received. For example, if a legacy site has a #2 AWG ground conductor running along the rows of racks (known as an aisle feeder), and the company specification calls for a 2/0 AWG aisle feeder (the size difference between #2 and 2/0 is not insignificant, but neither is it a ‘show-stopper’; it is a commonly used size – especially in legacy sites, and is acceptable as a minimum size in ANSI/SCTE 275), it may be better to wait until an equipment refresh or new racks are installed to do the upgrade, whereas if there is no aisle feeder at all, it should be rectified more immediately.

5.1. Various Grounding System Methods/Theories/Approaches

Most of us have heard the anecdote involving “ask three experts, and you’ll get four opinions”. In the world of grounding, it is no different. Grounding and bonding - is it art, is it science, is it both? The answer is “Yes”.

When you walk into any cable operator headend, telecommunications carrier central office, commercial or military grade data center or critical facility, you will find a variety of grounding methods, and a mixture of some in the same facility, on purpose or by accident, and there are benefits and drawbacks to each.

The main types are Star and Mesh bonding networks, which may or may not be Isolated. What we are mainly talking about is the bonding of equipment racks to the interior grounding system, which ultimately ties to the Master Ground Bar (MGB). Here is an easy way (I hope) to understand it:

- Star means each item (network equipment rack or cabinet) is independently grounded to a dedicated Frame Ground Bar (FGB) which taps into a conductor that goes to a single “point”, typically an area ground bar. Most cable and telecom critical facilities use this method, primarily, but not exclusively.
- Mesh means multiple items are bonded to each other in a sort of “grid” fashion. This is typically done in a raised floor data center.
- By Isolated we really mean a “single-point ground” connection to a ground bar, or perhaps to a specific section of a ground bar is involved. A Mesh or Star sub-“area” may have a connection of this type.

These types may be mixed together in a critical facility, the proper use of which may be dictated by the specific grounding needs of the types of equipment, whether they are AC or DC powered, whether they have the chassis bonded or isolated from the neutral or grounded conductor, etc.




Figure 2 – Acceptable Interior Grounding Methods - Network Equipment Racks

5.2. Common Grounding And Bonding System Requirements

The ANSI/SCTE 275 goes into great detail about best practices of interior grounding and bonding to racks and cabinets, exterior cable entry points, the utility service entrance connection, DC plants and batteries, and multi-floor building applications.

Regardless of the application within the building, there are a few practices that are required for nearly all of them wherever possible:

- Each grounding or bonding connection should be dedicated if possible; no “daisy-chained” or back-to-back connections are permitted in a critical facility.
- Mechanical connections are not permitted to be used for grounding or bonding in a critical facility. All connections should be made using crimp lug or tap connectors.
- Connection of dissimilar metals is prohibited in critical facilities.
- Two-hole lugs should be used wherever possible to prevent rotation and loosening of the lug during or after installation. For some small ground bars within racks or for equipment chassis that provide only a single-hole lug connection, serrated washers should be used.
- Any paint should be scraped off surfaces and anti-oxidation grease should be used for all grounding and bonding connections.
- Ground cables shall “gracefully” flow in the direction of the MGB, as though downstream, with no sharp bends.
- Ground cables of all sizes shall have a minimum bending radius of 8 inches, and the angle of any bend shall not be less than 90 degrees.
- All ground cables that terminate to a ground bar shall be appropriately labeled at both ends.
- Grounding conductors shall be sized per National Electrical Code (NEC) table 250.122

5.3. The Master Ground Bar

The MGB is the primary element of the critical facility grounding and bonding system - it is the single point connection, or main common interface between all interior and most exterior grounding systems of the critical facility. It will typically have connections to the utility ground, the ground ring (two connections, if possible), building steel, cold water pipe, all building interior network equipment area auxiliary ground bars, the DC plant positive return bus, DC equipment frame ground bar, and finally, any cable entry ground bars.

It is very important for all of the conductors that are terminated to the MGB to be labeled. This is helpful for both administrative documentation, maintenance, and troubleshooting purposes when any ground anomalies are discovered.

Additionally, some cable operators employ what is called a PANI configuration of the conductors terminated to the MGB, in order to arrange them according to surge potentials by specific classifications of the grounding system elements. Each section of the MGB is labeled accordingly, and the PANI classifications are as follows:

P - Surge Producers: Typically conductors coming from Cable Entrance Ground Bars (CEGB), as well as generator, UPS, and transformer enclosures

A - Surge Absorbers: Typically contains connections to building steel, the exterior ground ring, the cold water pipe, and the incoming main AC service ground bar – also called the multi-grounded neutral bond

N - Non-Isolated ground zone equipment grounds: Typically includes Secondary Ground Bars (SGB) which are for equipment room ground conductors that include the cabinet/rack grounds, ladder rack, etc.

I - Isolated ground zone equipment grounds: Typically equipment associated with a digital telephony switch, logic grounds and the DC Plant reference ground

Finally, proper connections and cable sizing of the conductors terminated to the MGB help to achieve network reliability by increasing the resilience of the grounding system to ensure it will protect the operational integrity of the equipment and infrastructure within the critical facility.




Figure 3 – Master Ground Bar With PANI Configuration

6. Surge And Lightning Protection

Surge and lightning protection systems complement each other and represent the most efficient and cost effective way to ensure continuity of the network by preventing or mitigating any deleterious effects from voltage instability and high frequency electromagnetic disturbances which are known to cause the most damage to network equipment and infrastructure.

Lightning strikes and electric utility company switching can cause voltage surges on the incoming AC power feeders. To protect the critical facility from these voltage surges, a surge-protective device (SPD) meeting current UL standards is required to be installed at the main service switchboard (MSS) for each facility. The connecting leads should be kept as short as possible and there should be no severe bends in any of the leads. Additional SPDs may need to be installed on downstream AC panels as required by special circumstances for a specific facility, such as those located in areas with high lightning activity or unstable utility power. Care should be taken when specifying these additional devices to ensure they are properly coordinated with the upstream SPDs according to their permissible energy.

A roof mounted lightning protection system is designed to protect a facility by intercepting lightning flashes and conducting the lightning currents into the earth in a safe and efficient manner. A properly designed lightning protection system should protect the facility from structural damage that can be caused when lightning currents flow through a building's concrete and metallic superstructure. The roof mounted lightning protection system should connect to the facility buried ground ring through down lead conductors spaced in accordance with NFPA 780, "Standard for the Installation of Lightning Protection Systems".

A lightning protection system consists of air terminals, conductors, interconnecting conductors, down leads and all the fittings, brackets and support devices required to complete the system. The final installation should be inspected, and UL certified to indicate that it has been properly installed and is compliant to all national and local standards. All down leads should be tested to verify that they are making a low impedance connection with the buried ground ring.

7. Humidity And Electrostatic Discharge

Static electricity is created by the accumulation of stationary electrical charge on a body or conducting medium. As such it is a common phenomenon created by physical motion. Even circulating air currents can cause a charge build-up, especially during low humidity conditions. This electrostatic charge build-up discharges whenever the charge storing medium meets ground. These static charges can store up to as much as 40,000 Volts during the normal course of a human body walking across a nonconductive floor in a low humidity environment. If exposed to electrical static charges of these magnitudes, the static sensitive electronic components contained in a critical facility can become permanently damaged.

Some of the best practices detailed in ANSI/SCTE 275 are as follows:

- To avoid damage of electrostatic dissipative (ESD) sensitive devices the relative humidity (RH) should be ideally kept within the limits specified by ASHRAE TC9.9 and all manufacturer's requirements.
- Wrist straps and anti-static bags should be used by all personnel handling printed circuit boards and blade server type cards.
- There should be wrist strap test stations at all critical facility sites, and wherever sensitive electronic components are handled.
- ESD-rated grounded conductive floor tiles or mats should be installed near the base of each equipment rack.

By employing these best practices, cable operators will be better able to achieve reliability targets and ensure network continuity.

8. Commissioning And Maintenance

The commissioning of a critical facility's grounding and protection system installation shall be performed to ensure that the final installation meets or exceeds the original design intent so the facility will provide a safe and reliable environment for personnel and sensitive electronic equipment to operate, and help to ensure network resilience and thus meet reliability KPIs.

Commissioning entails the visual inspection of the interior and exterior grounding systems, lightning protection system, and SPD installations to verify that the electrical contractor has built these systems in accordance with the design engineer's drawings and specifications, as well as compliance to all applicable standards.

In addition, current measurements are made in all conductor connections to the MGB to identify any unwanted ground current loops in the system, and impedance measurements are made in the lightning protection down leads to verify they are connected to the buried ground ring.

Grounding system preventive maintenance is required periodically, to protect equipment and facility investments, as well as to ensure continuity of network resilience. Improper grounding and bonding, combined with the inevitability of internal and external power surges are the most likely reasons for network equipment failures from "unknown causes" and are often incorrectly attributed to equipment defects.

Replaced equipment will likely fail repeatedly unless the root cause is determined. Ongoing maintenance of the grounding and bonding systems helps to identify the potential causes for ground and surge protection issues before they happen.

9. Conclusion

As we are all aware, these are unusual times, and regardless of the political, environmental or biological climate that cable operators find themselves having to navigate and make adjustments for, there are certain non-negotiable elements of doing business in an uncertain world. The business bottom line is primary of course, and second to that is the customer, without which there is no bottom line.

The cable operator customer of the third decade of this millennium expects a reliable network, and continuous, seamless delivery of services, whether it be broadband cable, internet, telephony, cellular, wireless, home security, telemedicine, etc.

One primary way that network resilience can be attained, risk to service delivery continuity mitigated, and reliability targets realized is to invest the time and energy and dollars to deploy and maintain a solid grounding and bonding network, and the ANSI/SCTE 275 standard is the best way to achieve that objective.

Abbreviations

AC	alternating current
AHJ	authority having jurisdiction
ANSI	American National Standards Institute
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
AWG	American wire gauge
CEGB	Cable Entry Ground Bar
DC	direct current
ESD	electrostatic dissipative
FGB	frame ground bar
GEM	ground enhancement material
IEEE	Institute of Electrical and Electronics Engineers
KPI	key performance indicator

MGB	master ground bar
MSS	main service switchboard
NEC	National Electrical Code
NFPA	National Fire Protection Association
Ohm-m	Ohm-meters
RH	relative humidity
SCTE	Society of Cable Telecommunications Engineers
SGB	secondary ground bar
SPD	surge protective device
UL	Underwriters Laboratories

Bibliography & References

ANSI/SCTE 275 2021: *Electrical Grounding and Bonding for Cable Broadband Network Critical Facilities*

Available: https://wagtail-prod-storage.s3.amazonaws.com/documents/ANSI_SCTE_275_2021.pdf

IEEE 81-2012: *IEEE Guide for Measuring Earth Resistivity, Ground Impedance, and Earth Surface Potentials of a Grounding System*

Available: <https://standards.ieee.org/standard/81-2012.html>

NFPA 780 2020: *Standard for the Installation of Lightning Protection Systems*

Available: <https://www.nfpa.org/codes-and-standards/all-codes-and-standards/list-of-codes-and-standards/detail?code=70>

How Will Proactive Network Maintenance Change Under DOCSIS[®] 4.0?

A Technical Paper prepared for SCTE by

Ron Hranac
rhranac@aol.com

Jason Rupe, Ph.D. CableLabs[®]

Dan Torbet, CommScope

Brady Volpe, The Volpe Firm

Table of Contents

Title	Page Number
1. Introduction.....	4
2. PNM Overview	4
3. What is DOCSIS 4.0?	5
3.1. Frequency division duplexing.....	6
3.2. What is full duplex?	7
3.3. PNM in DOCSIS 4.0 networks	7
4. Plant Preparation and Transition path	8
5. Impairment Management and Other Challenges	8
5.1. Challenges in legacy plants	8
5.2. Potential sources of ingress at higher frequencies	9
5.3. Distribution and drop impairment impacts on FDD and FDX.....	10
5.4. Managing total composite power	10
5.5. OUDP leakage detection.....	11
5.6. Outside plant amplifiers and impacts on PNM	11
5.6.1. FDX amplification and PNM.....	11
5.6.2. Smart amplifiers and FDD.....	12
6. DOCSIS 4.0 Technology	13
6.1. Triggered RxMER	13
6.2. Interference group information	13
6.3. Echo cancellation in the node	14
6.4. Echo cancellation in the CM.....	15
7. Telemetry	15
8. Test and Query.....	16
8.1. Challenges and Opportunities with FDD.....	17
8.2. Challenges and Opportunities with FDX.....	17
8.3. DOCSIS 4.0 Impact on Standard PNM Tests	18
8.3.1. Summarizing the gaps	21
8.3.2. What is required?	23
9. Conclusion and Future Outlook.....	24
Abbreviations	25
Bibliography & References.....	26

List of Figures

Title	Page Number
Figure 1 - DOCSIS 3.1 "test points" for an HFC network.	5
Figure 2. FDD frequency maps.	6
Figure 3. FDX frequency maps.	7
Figure 4: End-to-end test results for a legacy tapped feeder. Graphic courtesy of CableLabs.....	9
Figure 5. Over-the-air frequency allocations in the U.S. from about 900 MHz to 1.85 GHz.....	10
Figure 6. Examples of active device output signal level versus frequency in an FDD network.....	11
Figure 7: Example diplex filter (a), diplex filter transition band (b).....	11

List of Tables

Title	Page Number
Table 1. Impact of FDD and FDX on DOCSIS 4.0 PNM tests	22

1. Introduction

DOCSIS[®] 4.0 specifications introduced two important changes beyond DOCSIS 3.1 specifications: extended spectrum, and full duplex transmission. As a result, the assumptions change around the interpretation of proactive network maintenance (PNM) data, how the PNM tests and queries may work, and the sensitivity of some frequencies to certain impairments. As such, PNM tools will likely evolve, and operator use of PNM will need to increase to assure service reliability is met.

This paper outlines expectations around how the network will change, and as a result how network operations may change. It can serve as the foundation for an industry project plan to develop network and service operations solutions to keep pace with new DOCSIS 4.0 technology.

2. PNM Overview

During SCTE's 2008 Cable-Tec Expo in Philadelphia, CableLabs' Alberto Campos, Eduardo Cardona, and Lakshmi Raman presented a paper titled "Pre-equalization Based Pro-Active Network Maintenance Methodology" [1]. The authors proposed using cable modem (CM) upstream transmitter adaptive pre-equalization coefficients to detect and localize plant impairments.

The basic idea involved (1) deriving complex frequency response signatures from pre-equalization coefficients, (2) looking for responses indicative of the presence of linear distortions,¹ and (3) overlaying CM location information from the cable company's customer database on a system topology display of some sort – for instance, digitized outside plant maps.

In 2009 CableLabs formed a PNM working group to implement the ideas presented in the Expo '08 paper. The output of the working group's efforts was a PNM best practices document published by CableLabs in 2010 (updated versions of the best practices document have since been published [2]), followed by a reference implementation.²

Using the CableLabs PNM best practices recommendations and sometimes also the PNM reference implementation, several cable operators and third parties were able to create software-based PNM applications. The PNM applications allowed operators to remotely identify and locate plant and drop impairments using data from CM upstream pre-equalization coefficients.

When the DOCSIS 3.1 specifications were created, a decision was made to incorporate provisions and "hooks" for PNM in those specs. PNM was revamped for DOCSIS 3.1 specifications from the ground up to provide downstream and upstream "test points" in the cable modem termination system (CMTS) and cable modem, allowing operators to characterize and troubleshoot hybrid fiber/coax (HFC) plant and subscriber drops; support remote proactive troubleshooting of plant faults; and improve reliability and maximize throughput in well-maintained plants. As shown in Figure 1, from [3], the cable network can be thought of as a device under test (DUT), and PNM measurements are virtual test equipment. For more information, see Section 9 of the DOCSIS 3.1 PHY Specification [3], which details PNM support and requirements.

¹ Linear distortions in cable networks include micro-reflections, amplitude ripple, and group delay distortion.

² SCTE's Network Operations Subcommittee Working Group 7 (NOS WG7), created in 2017, also handles PNM. The CableLabs and SCTE PNM working groups collaborate on the subject, and each group's efforts complement the other's.

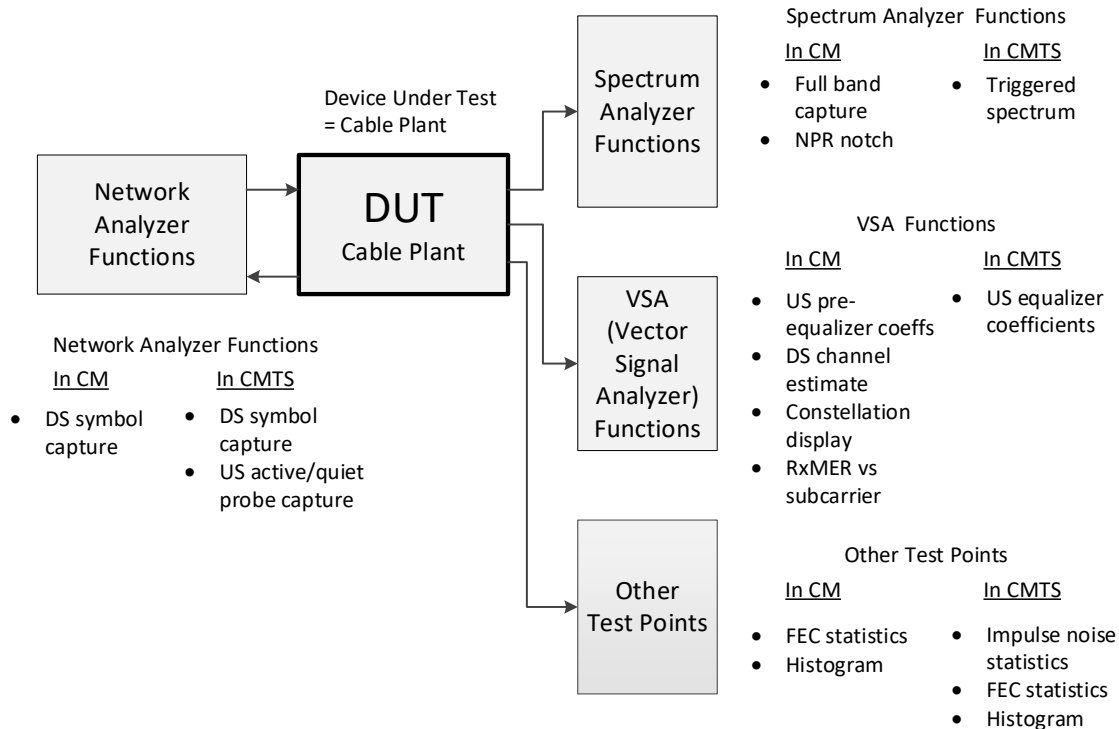


Figure 1 - DOCSIS 3.1 "test points" for an HFC network.

3. What is DOCSIS 4.0?

DOCSIS 4.0 specifications, released in 2019, are the latest in the DOCSIS family. The following description from the introduction in the DOCSIS 4.0 Physical Layer Specification [4] provides an overview:

This generation of the DOCSIS specifications builds upon the previous generations of DOCSIS specifications (commonly referred to as the DOCSIS 3.1 and earlier specifications), leveraging the existing Media Access Control (MAC) and Physical (PHY) layers. It includes backward compatibility for the existing PHY layers in order to enable a seamless migration to the new technology. Further, the DOCSIS 4.0 specifications introduces Full Duplex (FDX) DOCSIS PHY layer technology as an expansion of the OFDM PHY layer introduced in the DOCSIS 3.1 PHY specification to increase upstream capacity without significant loss of downstream capacity versus DOCSIS 3.1. The DOCSIS 4.0 specification also builds upon DOCSIS 3.1 OFDM and OFDMA technology with an extended Frequency Division Duplex (FDD) DOCSIS alternative. DOCSIS 4.0 FDD supports legacy high split and also provides extended splits up to 684 MHz in an operational band plan which is referred to as Ultra-high Split (UHS). DOCSIS 4.0 FDD also introduces expansion of usable downstream spectrum up to 1794 MHz. Both the FDX and FDD DOCSIS 4.0 alternatives are based on OFDM PHY. Many sections refer to basic OFDM sublayer definitions described in [DOCSIS PHYv3.1].

Cable operators have for decades designed their networks to use sub-split band plans. A sub-split band plan is one that has most of the usable radio frequency (RF) bandwidth allocated to downstream signal transmission. A small portion near the lower end of the usable spectrum is allocated to upstream transmission. For example, a common sub-split band plan used in North America and elsewhere has the upstream operating from 5 megahertz (MHz) to 42 MHz, and the downstream operating from about 54 MHz to the highest downstream frequency limit (e.g., 750 MHz). In an effort to increase upstream capacity and data throughput, the industry has been migrating to mid-split and high-split band plans, with the former using 5 MHz to 85 MHz for upstream transmission, and the latter using 5 MHz to 204 MHz for upstream transmission. For more information on band splits and their history, see [5].

Introduced at the 2019 CES, the cable industry's 10G Platform [6], [7] will deliver speeds of 10 gigabits per second (Gbps) with improved reliability, security, and lower latency, using DOCSIS 3.1 and DOCSIS 4.0 technologies, passive optical networks (PON), coherent optics, dual channel Wi-Fi³, and more.

In particular, the 10G Platform will take advantage of DOCSIS 4.0 technology's expanded spectrum usage – to 1794 MHz (aka 1.8 gigahertz, or GHz) or higher – and more efficient use of parts of the RF spectrum with FDX operation.

3.1. Frequency division duplexing

Originally called “extended spectrum DOCSIS” (ESD), the term frequency division duplexing (FDD) is used in the DOCSIS 4.0 specifications. The reason it's called FDD is because, just like DOCSIS 3.1 and earlier technology, downstream signals operate in one frequency range and upstream signals operate in a different frequency range. The DOCSIS 4.0 upstream RF spectrum can operate to as high as 684 MHz, and the downstream to as high as 1.8 GHz or more. Figure 2, from [4], shows the configurable FDD upstream allocated spectrum bandwidths.

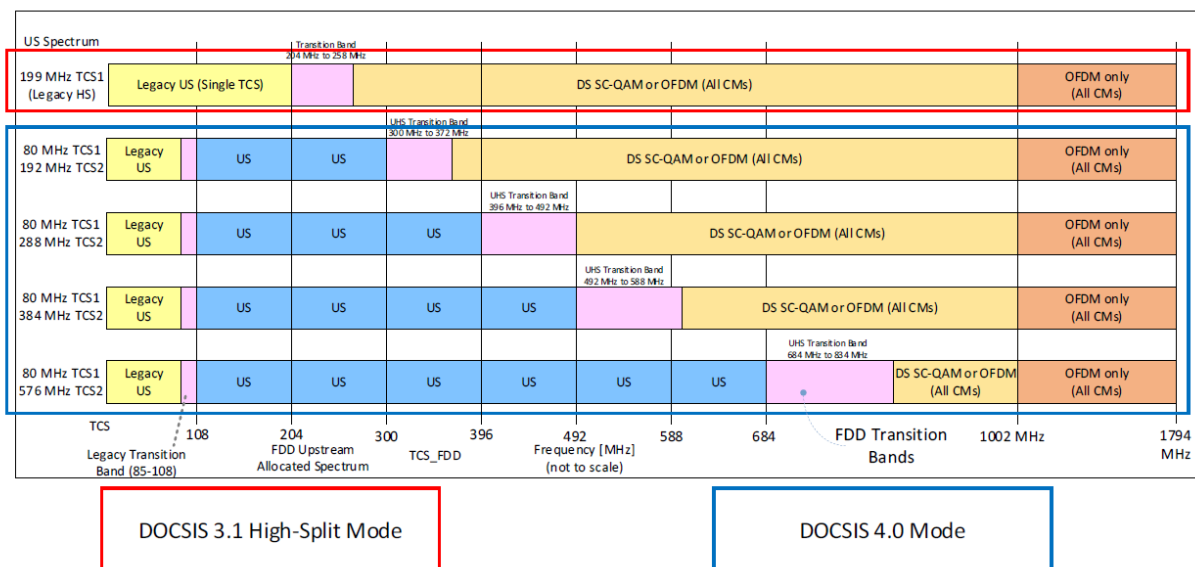


Figure 2. FDD frequency maps.

³ Wi-Fi[®] is a registered trademark of the Wi-Fi Alliance[®]. Wireless local area networks (WLANs) are commonly called Wi-Fi.

3.2. What is full duplex?

FDX – commonly known as “FDX DOCSIS” – was originally introduced as an annex in DOCSIS 3.1 specifications, and is now part of the DOCSIS 4.0 specifications. Through the magic of echo cancellation (EC) and other technologies, FDX allows the carriage of downstream and upstream signals on the same frequencies at the same time. The graphic in Figure 3, from [4], shows configurable FDX allocated spectrum bandwidths, including what is called FDX allocated spectrum. The latter comprises the frequency ranges where downstream and upstream signals can simultaneously occupy the same frequencies, allowing increased data speeds in both directions.




Figure 3. FDX frequency maps.

3.3. PNM in DOCSIS 4.0 networks

As mentioned previously, [3] includes a full section (section 9) covering PNM. “Section 9 PROACTIVE NETWORK MAINTENANCE” in [4] simply says “See [DOCSIS PHYv3.1] section 9.”

While the PNM parameters in [3] are for the most part directly applicable to DOCSIS 4.0 technology deployments, there are some important differences and some new challenges. For example, full band capture (FBC) in cable modems will have to support a higher upper frequency limit in the downstream, to 1.8 GHz in FDD applications. Indeed, all of the downstream PNM parameters described in [3] and referenced in [4] will need to support operation to 1.8 GHz, and the upstream PNM parameters will need to accommodate operation in all of the supported frequency ranges to as high as 684 MHz.

Cable network operation on higher frequencies in both the upstream and downstream will be susceptible to new sources of ingress, as well as services with which signal leakage can interfere. Other challenges include such things as management of total power at active device outputs; isolation requirements for FDX; additional attenuation at higher frequencies; PNM test and query; and more. Developers and users of PNM tools and applications will need to understand these challenges, many of which are discussed in subsequent sections of this paper.

4. Plant Preparation and Transition path

The best proactive network maintenance is that which happens when the network is being prepared for DOCSIS 4.0 technology, well before an impairment occurs.

As networks are being upgraded, consider drop tests as well, to find potential leakage, poor drop performance, and potential sources for non-linear impairments. Passive intermodulation (PIM) distortion is anticipated to be worse in DOCSIS 4.0 networks with higher operating levels. Plant preparation is a convenient opportunity to find and remove any older distribution and drop passives that will not perform well at higher frequencies, find and remove bad or detrimental filters (including in-line equalizers), and find and remove any house amplifiers that will impact service. Degraded and poor performing drops may also have difficulty carrying signals at higher frequencies.

In-depth guidance on plant preparation is beyond the scope of this paper, but the authors acknowledge its importance for PNM. In particular, operators must address plant quality before upgrading to and deploying DOCSIS 4.0 technology.

The next section discusses some of the challenges related to managing impairments, and the potential impacts of those impairments on DOCSIS 4.0 technology deployments. Other potentially impacting topics, such as FDX-capable amplifiers and smart amplifiers, are also discussed.

5. Impairment Management and Other Challenges

5.1. Challenges in legacy plants

Figure 4 shows the results of end-to-end testing of a tapped feeder leg (active device output to last tap) using legacy components designed for a maximum downstream frequency of 1 GHz to perhaps 1.2 GHz. The coaxial cable has a standalone attenuation at 1 GHz of about 24 decibels (dB), typical of just under 1000 feet of 0.500 hardline coax. Looking at the S_{12} and S_{21} traces,⁴ the combined insertion loss (cable plus passives) is about 45 dB at 1 GHz, typical of a span of feeder with taps and other passives. Of particular concern is the sharp frequency response rolloff starting at about 1.3 GHz in the S_{12} and S_{21} traces (circled in red in the figure), indicating that attenuation at higher frequencies is substantial. That rolloff is caused by the passives. The S_{11} trace, from which return loss can be derived, also indicates poor performance above about 1.3 GHz. From this example, operation above 1.3 GHz would be impossible using the legacy passives.

Cable operators contemplating operation to 1.8 GHz will need to evaluate their networks to determine to what extent upgrades or changes will be necessary to support higher downstream frequencies. PNM tools will need to support operation at the higher frequencies, too.

⁴ The S_{11} , S_{12} , S_{21} , and S_{22} parameters in the figure are scattering parameters, or S-parameters. For more on S-parameters, see [13]




Figure 4: End-to-end test results for a legacy tapped feeder. Graphic courtesy of CableLabs.

5.2. Potential sources of ingress at higher frequencies

Cable operators are already familiar with sources of ingress and over-the-air signals affected by signal leakage in the 5 MHz to 1 GHz frequency range. Most operators have little or no experience with ingress and leakage at frequencies above 1 GHz, though. Figure 5, from [10], shows over-the-air frequency allocations in the United States from about 900 MHz to 1850 MHz (frequency allocations in other countries may be different). This frequency range includes the 902 MHz to 928 MHz industrial, scientific, and medical (ISM) band (shared with amateur radio); the 23 centimeters amateur radio band (1240 MHz to 1300 MHz); six aeronautical radio navigation bands (960 MHz to 1215 MHz, 1300 MHz to 1350 MHz, and four smaller bands from 1559 MHz to 1626.5 MHz); some long term evolution (LTE) bands; among others. GPS frequencies⁵ are in the 1100 MHz to 1600 MHz frequency range, too. Signals on some of the aforementioned frequencies are potential sources of ingress interference to the cable network, and can be interfered with by signal leakage.

PNM's full band capture and receive modulation error ratio (RxMER) will continue to be valuable for identifying and helping to locate potential ingress, especially at the higher frequencies discussed here.

⁵ Global Positioning System (GPS) frequency L1 is 1575.42 MHz (15.345 MHz bandwidth); L2 is 1227.6 MHz (11 MHz bandwidth); and L5 is 1176.45 MHz (12.5 MHz bandwidth). See <https://www.nist.gov/pml/time-and-frequency-division/popular-links/time-frequency-z/time-and-frequency-z-g>



A comprehensive list of plant and drop impairments and their impacts on FDD and FDX operation could easily be the basis for a standalone paper. The following are some of the more important considerations.

- Proper attenuation/insertion loss, frequency response, return loss, port-to-port isolation (where applicable), and other characteristics of network and drop actives, coaxial cable, passives, connectors, etc., across the full operating bandwidth are critical. Out of spec performance for any of the aforementioned could negatively affect FDD and FDX operation.
- Ingress in FDX bands causes errors in channel characterization, and the RF bandwidth of the spectrum affected by noise funneling can be larger in an FDX architecture than in others because of its wider upstream bandwidth. Noise funneling remains a problem because one source impacts all. That is, severe ingress from just one drop can significantly impair a node's upstream performance, regardless of the size of the node's service area or the number of homes passed – decreasing the size of the serving area does not necessarily decrease the noise problem. As well, drop ingress in an FDX band could impair downstream (in the drop) and upstream performance.
- Common path distortion (CPD), originating from inside customers' homes, might be increased by high transmit levels of FDX and FDD CMs.
- An FDX or FDD CM located on unconditioned house wiring, rather than the point of entry, could experience more problems on average. This outcome is due to additional complexity of the inside wiring, presence of drop passives and actives, the number of connectors, etc.

To overcome increased attenuation at higher frequencies in FDD networks, active device output power – including output total composite power (TCP) – will be higher. Figure 6, from [11], illustrates three examples of signal level-versus-frequency in an FDD network. PNM tools can be an important part of the management of active device signal levels and TCP. Section 30.14 and Appendix J of [11] include additional discussion about TCP.




Figure 6. Examples of active device output signal level versus frequency in an FDD network.


5.5. OUDP leakage detection

While signal leakage detection is generally not part of a PNM toolset, ingress detection is, and the two are often related. That is, if leakage exists ingress usually does, too. In high-split and ultra-high-split band plans, the 108 MHz to 137 MHz aeronautical band overlaps part of the cable network's upstream spectrum. Leakage detection and measurement are more challenging, since a continuous downstream leakage test signal cannot easily be transmitted in or near aeronautical band frequencies. One promising method is to use OFDMA upstream data profile (OUDP) for leakage detection and monitoring. This approach is discussed in [12], and recent lab and field test results are encouraging.


5.6. Outside plant amplifiers and impacts on PNM

5.6.1. FDX amplification and PNM

RF amplifiers are an integral part of any HFC network. Their primary purpose is to amplify and condition RF signals so that they may propagate through subsequent spans of coaxial cable. Coaxial cable attenuates RF signals in a non-uniform manner (that is, cable attenuation is greater at higher frequencies than it is at lower frequencies), thus requiring the next RF amplifier to again amplify and condition the RF signals. Rinse and repeat. RF amplifiers and fiber nodes have one or more devices in them called a diplex filter. Two diplex filters are shown in Figure 7(a) and a simulated transition band of the frequency response of a sub-split diplex filter is shown in Figure 7(b). The purpose of the diplex filter is isolate the downstream from the upstream path in the active device, helping to prevent problems such as oscillation.



(a)



(b)

Figure 7: Example diplex filter (a), diplex filter transition band (b).

Diplex filters are not compatible with FDX operation. This is because upstream and downstream signals are present on some frequencies at the same time. The diplex filter prevents this from occurring, so FDX networks require one of two considerations:

- A network design with a fiber node and no active devices after the fiber node (aka Node + 0).
- A design with a fiber node and amplifiers capable of amplifying in both directions without diplex filters, or dedicated FDX-capable amplifiers with echo cancellation.

In the case of FDX amplification and PNM, a Node + 0 network in theory makes for a much simpler network to maintain and troubleshoot. Since it lacks any actives after the fiber node, one only need be concerned with the node (and perhaps power supply), the hardline coax, passive devices, and subscriber drops and their components. When considering troubleshooting using such things as correlation groups and echo cavities, the following are important:

- The number of subscribers available in a correlation group will typically be smaller due to smaller service groups.
- Echo cavities will always be formed between the fiber node and passive devices and/or damaged coax/connectors. Actives after the node will no longer play a role in echo cavities.
- The number of impairments in a given plant segment will decrease because the number of connectors, coax, etc. will be fewer.
- Node + 0 will not fix pre-existing distribution coax and drop issues. RF failure group sizes may be reduced, and some causes of failure eliminated, but the number of overall service issues may or may not be greatly reduced.

While the general maintenance requirements for a Node + 0 network are not substantially expected to go down, the overall performance is expected to improve. This is because RF downstream signals from the fiber node will be nearly equivalent to those in the headend from a quality perspective. Further, upstream signals from the CM will be received and demodulated at the fiber node, assuming a DAA deployment with digital fiber links.

PNM will be a key troubleshooting tool in an FDX deployment to monitor performance, impairments and EC performance. It is assumed that for FDX to be successful some amplification may be required in HFC networks. It will be essential for PNM to have visibility into FDX amplified networks, especially at the amplifier level. See the next section on smart amplifiers.

5.6.2. Smart amplifiers and FDD

FDD is very different than FDX in its requirements for diplex filters and amplification, but there is one commonality discussed this section: smart amplifiers.

FDD will still require diplex filters to separate the downstream from the upstream in actives and certain passives. However, diplex filters in an FDD environment will need to at a minimum be upgradable and ideally programmable (or remotely switchable). Similarly, the RF conditioning circuitry in RF amplifiers should also be programmable. Vendors are producing new amplifiers with diplex filters and conditioning circuitry which can be remotely configured or locally configured with a mobile app.

This leads into the concept of smart amplifiers. Smart amplifiers not only include the ability to eliminate legacy plug-in pads and equalizers, but they are also adding PNM functionally such as full band capture so that one can remotely see the output of the amplifier to configure its padding and conditioning. PNM can take advantage of FBC in the fiber nodes and amplifiers as yet another monitoring point in the field to identify and localize RF impairments. This is useful for both FDD and FDX deployments.

Because the FBC capability in smart amplifiers is built on CM technology, other PNM tests can use the CM in the amplifier. Now the amplifier can be used for all PNM tests supported by the CM chipset. Of

particular importance is that this functionality operates over the extended upstream and downstream RF bandwidths.

6. DOCSIS 4.0 Technology

6.1. Triggered RxMER

As stated in [4], FDX modems can measure RxMER over all subcarriers at the same time. This measurement is triggered by events including time triggers, echo canceller training (ECT) probe triggers, and OUDP sounding triggers.

Time triggers could be the most useful for PNM because they are triggered at a specific time, and that time could be correlated to certain events by the CMTS or converged cable access platform (CCAP). For example, a CM can measure RxMER per subcarrier while a specific test downstream signal is sent, and while another upstream signal is sent by another CM. This could be used to find nonlinear impairments. Also, it could be used to synchronize data collection on multiple CMs at once, or to coordinate with a test device signal or measurement, or to time an external radio signal to potentially look for ingress sources, for example.

ECT RxMER probe triggers measure a modem's receive capabilities during worst case conditions, used for setting bit-loading after echo cancellation training. As such, it could be useful for PNM as an indication of the environment for transmission, and will provide a peek at the bit loading that the CM is capable of achieving.

OUDP sounding triggers allow measurement of the interference between CMs. With a measure and a sounding CM pair, the information could be very useful for PNM, particularly for fault localization. The information may even be useful to verify relative location information.

While these alternate forms of RxMER are described for the DOCSIS 4.0 protocol in both the PHY [4] and MULPI [8] specifications, the OSSI [9] specifications have not yet outlined how the data would be reported for fault management purposes.

6.2. Interference group information

The CMTS or CCAP determines a given FDX modem's participation in an interference group (IG) and transmission group (TG) after modems have performed sounding. This process helps the CMTS to identify which CMs interfere with other CMs on a given HFC plant segment. During sounding one CM is granted time to transmit and surrounding CMs are told to measure the per-subcarrier RxMER for the FDX sub-band being used for the sounding. Future enhancements in the OSSI specifications are planned to address the reporting of which CMs have been placed into specific IGs and TGs. Because CM membership in a given IG or TG is not operator-configured, the CMTS/CCAP is the source of truth for these associations.

Additionally, FDX manages the usage of a given FDX sub-band through the use of what is called a resource block (RB). It is possible to assign a given RB in a given TG as either static (always one direction, usually upstream) or dynamic (scheduled by the CMTS or CCAP, the block operates in both directions in a given sub-band for a specific TG). In the dynamic resource block assignment (RBA) scheduling, RBA switching can happen at an extremely fast rate which makes tracking and reporting of the RBA metrics challenging.

6.3. Echo cancellation in the node

To maintain increased capacity, and to meet 10G goals, greater reliance on echo cancellation is to be expected. When echo cancellers do not perform as needed, RF impairments may impact service.

FDX-capable nodes support simultaneous upstream and downstream communications over each FDX channel. FDX-compliant CMs will operate in FDD mode, where on any FDX channel or sub-band, the CM is either transmitting in the upstream or receiving in the downstream direction. To avoid the risk of co-channel interference (CCI) and adjacent channel interference (ACI) between CMs, the CMTS schedules transmissions and grants such that a CM does not transmit at the same time as other CMs that are susceptible to interference from the transmitting CM. CM to CM interference susceptibility is measured through a sounding process that is defined in the specification. Even with the CMTS determining interference groups, there is still a need to manage the impacts of upstream and downstream signals at the node. The FDX node has to employ echo cancellation methods to help remove the reflection of the downstream transmitted signals that are reflected back from components in the node and in the plant which impact the reception of upstream signaling. This echo cancellation can be done in both the analog and digital domains. In the analog domain, traditional techniques can be utilized that copy the DS carriers and then manipulate phase and magnitude and then apply that as a filter on the receiver path. EC conducted in the digital domain allows for near signal regeneration, depending on the node design. At the node, the DS signals are at their strongest so this method can be effective at cancelling out echoes in the FDX node itself.

From a PNM point of view, more work and definition need to be performed to determine how to measure and report echoes, and the signals after echo cancellation has been performed. As more FDX plant and nodes are deployed, these challenges will be met, and new management objects will be created which will aid in better performance of these plant segments.

A pre-EC RxMER measurement, possible at the node, could provide some information about the signal before EC, which may be partially informative toward gauging the EC effectiveness.

Generally, there are four approaches for providing this information, which could be useful for PNM and operations tools and applications:

- a measure in the specification that is based on what is described here;
- a measure in the specification that is based on some combination of US and DS RxMER and maybe more to describe some equivalent of the effort spent on echo cancellation and how much more could be corrected, without getting into manufacturer-specific intellectual property;
- a best practice based on one of these approaches, doing what is best to use external testing and available measures to provide equivalent information; and
- new requirements to collect and hold the information for future use.

A fiber node typically feeds up to four legs of coax. Management of the FDX spectrum and EC process in nodes with multiple legs can vary among vendors, with potentially different implementations. Actual capabilities and details are beyond the scope of this paper. However, from a PNM perspective there would be value if per-leg EC data were available.

The authors suggest that the DOCSIS specifications teams develop an engineering change to provide management objects that describe the limits and performance of ECs.

6.4. Echo cancellation in the CM

As with the node EC, a couple of parameters should be reported for the FDX CM EC. However, the FDX CM is a lower cost device and may not have the horsepower to provide all the measurements that a node could.

- EC is trained or not trained—In the OSSI and MULPI specifications, a trained EC is sufficiently converged. Reporting on the state of training is important for PNM. For instance, if the EC is not converged, reception of the channel(s) in the sub-band might not be usable, and knowing that the EC is not converged provides valuable information for troubleshooting why the channel/sub-band is not usable.
- Echo before and after cancellation—Operators need a way to express this so they can characterize the plant and find when and where changes happen in the network.
- Margin remaining before EC begins to have problems—Operators need a way to express this parameter. The specifications development organizations could suggest measurement and reporting methods.
- Any indication of why an EC cannot train or is on the margin—For example, if there is a particular echo that is too big, then data on that echo could be used to provide the echo's distance, which would allow an operator to troubleshoot. The industry will need to determine how that information would be communicated and what the data would look like.

7. Telemetry

With increased complexity and the expectation of new service-impacting failure modes being revealed in the network, new telemetry, and more frequent telemetry in some cases, will be needed. PNM fault management requires identification of faults from the telemetry, and then use cases for localizing the fault must follow. Fault identification requires a broad scope covering the entire network, with an initial granularity for the next step of fault management; fault localization requires several different groupings and finer resolution of telemetry. PNM fault management will require ways for operators to manage their operations and maintenance costs, so improvements in tools will follow the improvements in telemetry delivery.

PNM telemetry today consists of queries (polling data that are intermittently collected) and tests (requiring configuration to enable the data collection). In DOCSIS 3.0 networks, Simple Network Management Protocol (SNMP) was the primary method in use. The need for larger data sets occurred with DOCSIS 3.1 technology. Trivial File Transfer Protocol (TFTP) was introduced to enable more optimal data transfer than what is possible with SNMP. TFTP was adopted by CM vendors, but had limited adoption by CMTS vendors.

DOCSIS 4.0 technology brings new possibilities with telemetry from CMTS equipment. With R-PHY architectures, there are new tunneling protocols such as Layer 2 Tunneling Protocol (L2TP) pseudowires presenting a packet streaming protocol (PSP). DOCSIS 4.0 technology brings requirements for more data, more measurements, more often. SNMP will continue to increase our “technical debt” that limits the full capabilities of monitoring platforms. L2TP is one such protocol that is in use by CMTS vendors for transporting large amounts of data fast.

With R-MACPHY, YANG data models describe the methods for acquiring telemetry, delivered through TFTP, HTTP, and other means. Like L2TP, YANG models eliminate limitations with SNMP and bring data aggregation from DOCSIS devices into the 21st century. YANG models are an interesting discussion

topic but are not a requirement for CMTS vendors to support, therefore, similar to TFTP it could be unlikely that CMTS vendors will adopt YANG models.

Vendors have developed proprietary solutions, based on protocols and techniques such as L2TP, CMTS SSH direct READ access, Kafka bus access, and more. Proprietary solutions have proven to be substantially more effective than SNMP and deliver critical data in near-real time. Examples of this are upstream spectrum analysis and OFDMA RxMER per subcarrier data. Vendor implementations of streaming telemetry using proprietary solutions enable upstream spectrum analysis which rival hardware-based spectrum analyzers in terms of trace update response time. Further, OFDMA RxMER data can be obtained on a fully loaded CMTS chassis every 15 minutes for every active CM connected to that CMTS. This can be done simultaneously while running other PNM tests. For cable operators considering upstream profile management application (PMA), this is a game changer. Those familiar with standard PNM tests know that running OFDMA RxMER typically prevents one from running any other PNM test, such as upstream triggered spectrum capture (UTSC).⁶

The advent of smart amplifiers will add a new telemetry opportunity. This paper previously discussed the FBC capability in smart amplifiers, but there is far more to it. DOCSIS 4.0 amplifiers are expected to run hotter, provide more gain, operate at higher frequencies, and contain sophisticated internal electronics. Having an on-board CM enables the ability to monitor the modem's internal temperature sensor(s), voltages, and any other on-board sensor the vendor may choose to include. All this data gets communicated directly back to the CMTS and the monitoring system. Vendors could, for example, give access to mainline power supply monitoring, RF probing, and more. Each amplifier becomes another telemetry point in the network. How the telemetry is retrieved is again up to DOCSIS 4.0 specifications and vendor implementation. Ideally it will not use SNMP.

What does all this mean? There are many solutions being tested on the open market for DOCSIS 4.0 data collection. Further, new solutions are being proposed by CableLabs. It will ultimately be up to the cable operator community to drive which technologies are adopted and implemented. This paper has identified the need that DOCSIS 4.0 technology has for PNM. However, it should be apparent to the reader that a gap exists. That gap is the lack of a clear path between how vendors and operators will align to a consistent PNM implementation and how PNM will function across multiple vendor platforms to meet the expectations of cable operators. It is recommended that operators understand this gap and align to address it.

8. Test and Query

As discussed previously in this paper, DOCSIS 4.0 technology brings new opportunities and challenges from a PNM perspective. New test methods and data analytic queries must be created and optimized for new spectrum changes and new technologies, such as echo cancellation. As data speeds increase, bandwidths expand and complex technologies are introduced, PNM will be increasingly more important to ensure continued quality of experience (QoE) to subscribers. There is no doubt that subscribers will continue to be more dependent on high-speed data, and competition will continue to apply pressure to improve network quality.

⁶ Rather than track per-modem OFDMA RxMER per subcarrier (and preclude the use of UTSC by field personnel), some operators monitor port average RxMER and other data at the upstream input to the CMTS/CCAP. If a problem is detected, then data from individual modems can be looked at more closely.

8.1. Challenges and Opportunities with FDD

In FDD, the upstream band can extend to as high as 684 MHz. In DOCSIS 3.1 technology, the highest upstream frequency is 204 MHz. Higher frequencies mean much more data to aggregate and store in databases from cable modems and CMTSs. Further, large data sets require significantly more CPU (or GPU) processing power from an analytics standpoint when identifying impairments. Today, most modems and CMTSs still rely on SNMP to retrieve data from them. While SNMP has been a great protocol for the cable industry for more than a decade, it is a very slow and outdated protocol. Fortunately, CableLabs specifications are moving towards other methods of obtaining large data sets, such as TFTP and streaming telemetry. It is critical that adoption of these methods by vendors occurs quickly.

The downstream spectrum in FDD will also be increasing from 1.2 GHz to 1.8 GHz with visions of one day supporting up to 3 GHz (or higher!). As in the upstream, this will require cable modems to support FBC up to the highest frequency supported in the network. Currently, DOCSIS 3.1 modems support FBC up to 1.2 GHz while DOCSIS 3.0 modems only support FBC up to 1 GHz. One can see the disparity of a 1.8 GHz network supporting a mixture of DOCSIS 3.0, 3.1 and 4.0 cable modems. As DOCSIS 4.0 modems are initially deployed, one will have limited visibility to impairments in the RF spectrum above 1.2 GHz, assuming significant deployment of DOCSIS 3.1 modems. It is expected that many passive devices, and subscriber drop cables and components, will have various impairments above 1.2 GHz because the 1.2 GHz to 1.8 GHz spectrum has never been widely tested. It may sound trivial when speaking in terms of “GHz,” but this is 600 MHz of largely untested spectrum that PNM will be essential in analyzing and testing. DOCSIS 4.0 modems with FBC capabilities up to 1.8 GHz are essential. Further, a method of quickly obtaining the FBC spectrum from 5 MHz (or lower) to 1.8 GHz will be critical.

Upstream spectrum analysis is a “meat and potatoes” feature of any PNM application. Technicians rely on it every day to identify and resolve return path ingress and other impairments. The state-of-the-art return path upstream spectrum analysis relies on a CableLabs-based specified measurement called UTSC. UTSC enables compatibility across vendors and platforms whether it is integrated CCAP (iCCAP) or distributed access architecture CCAP (dCCAP). CCAP vendors and PNM vendors will need to ensure their platforms support upstream spectrum analysis up to the highest frequencies supported in FDD. Further, vendors must be able to support fast refresh speeds on upstream spectrum analysis over a much wider bandwidth in order to capture transient noise events, many of which may occur at higher frequencies not previously seen. The current state of DOCSIS 3.1 UTSC across the vendor space is non-optimal in that each CCAP vendor has partial adoption of the CableLabs UTSC specification. This state creates challenges for adoption by cable operators, and a lack of feature sets with some vendors means that not all tests are supported. It will be important that vendors fully adopt UTSC in DOCSIS 4.0 networks so that operators are able to troubleshoot more complex problems as frequency expansion will certainly bring unanticipated complexities.

8.2. Challenges and Opportunities with FDX

Like its counterpart FDD, FDX has similar upstream and downstream frequency expansion challenges for PNM. However, FDX adds more technical hurdles which PNM will be critical to help solve. For instance, downstream FBC at the CM may contain upstream and downstream transmissions within the same interference group in the FDX band (108 MHz to 684 MHz). Visualization and troubleshooting of simultaneous upstream and downstream will lead to new challenges for both vendors and technicians.

Having visibility into the level of EC and reserve EC capacity will be essential. As previously discussed, the CMTS and CM have EC functionality. The authors believe that it will be possible to determine some amount of impairment between the CMTS and the CM by utilizing the information provided by the EC operation in the CM and FDX node.

In general, PNM tests for FDX will be more challenging overall than FDD. A general summary of this can be seen later in Table 1.

8.3. DOCSIS 4.0 Impact on Standard PNM Tests

As defined in the DOCSIS 3.1 and 4.0 specifications, there exists a standard set of PNM test and query features designed to enable cable operators and vendors to obtain optimal troubleshooting data from the CMTS and CMs. Those specifications are designed in order to establish consistent interoperability among vendors of CMTS, CM and PNM software. This section provides a brief description of each PNM test followed by table that summarizes the gaps for full support of DOCSIS 4.0 FDD and FDX.

DS Symbol Capture (CM and CCAP)

Description:

- The `DsOfdmSymbolCapture` object provides partial functionality of a network analyzer to analyze the response of the cable plant. A symbol is generated at the CCAP and also captured at the CM, and then the results compared.

DsOfdmNoisePowerRatio (CCAP/Spectrum)

Description:

- The purpose of downstream NPR measurement is to view the noise, interference and intermodulation products underlying a portion of the OFDM signal. As an out-of-service test, the CCAP can define an exclusion band of zero-valued subcarriers which forms a spectral notch in the downstream OFDM signal for all profiles of a given downstream channel. The CM provides its normal spectral capture measurements per [PHYv3.1], or symbol capture per [PHYv3.1], which permit analysis of the notch depth. A possible use case is to observe LTE interference occurring within an OFDM band; another is to observe intermodulation products resulting from signal-level alignment issues. Since the introduction and removal of a notch affects all profiles, causing possible link downtime, this feature is intended for infrequent maintenance.

DS CM Spectrum Analysis Full Band Capture

Description:

- This test allows for the full band capture of the DS RF spectrum that the modem is configured to use.

CmDsOfdmChanEstimateCoef

Description:

- The purpose of this table is for the CM to report its estimate of the downstream channel response. The reciprocals of the channel response coefficients are typically used by the CM as its frequency-domain downstream equalizer coefficients. The channel estimate consists of a single complex value per subcarrier. The channel response coefficients are expressed as 16-

bit two's complement numbers using 2.13 nibble format. The CM samples are scaled such that the average power of the samples is approximately 1, in order to avoid excessive clipping and quantization noise.

- Summary metrics (slope, ripple, and mean) are defined in order to avoid having to send all coefficients on every query. The summary metrics are calculated when the corresponding MIB is queried. A Coefficient filename and trigger are provided to obtain the channel coefficients.
- The CM will report these metrics for each OFDM channel it has been assigned.

CmDsConstDispMeas

Description:

- The downstream constellation display provides received QAM constellation points for display. Equalized soft decisions (I and Q) at the slicer input are collected over time, possibly with subsampling to reduce complexity, and made available for analysis. This measurement is intended for data subcarriers only. Up to 8192 per OFDM channel samples are provided for each query; additional queries can be made to further fill in the plot.

ModulationOrderOffset

Description:

- This attribute specifies an offset from the lowest order modulation for the data subcarriers in any of the profiles in the downstream channel. If the lowest order modulation order that the CM was receiving was 1024-QAM and the ModulationOrderOffset was zero, then the CM would capture the soft decision samples for all of the subcarriers which were using 1024-QAM. If the ModulationOrderOffset was 1, then the CM would capture the soft decision samples for all of the subcarriers using the next highest modulation order in use for the profiles in the downstream channel.

CmDsOfdmRxMer

Description:

- Provides measurements of the RxMER for each subcarrier.

CmDsOfdmMerMargin

Description:

- Provide an estimate of the MER margin available on the downstream data channel with respect to a modulation profile. The profile may be a profile that the modem has already been assigned or a candidate profile. This measurement is similar to the MER Margin reported in the OPT-RSP Message [MULPIv4.0].

CmDsOfdmFecSummary

Description:

- The purpose of this item is to provide a series of codeword error rate measurements on a per profile basis over a set period of time.

CmDsHist

Description:

- The purpose of the downstream histogram is to provide a measurement of nonlinear effects in the channel such as amplifier compression and laser clipping. For example, laser clipping causes one tail of the histogram to be truncated and replaced with a spike. The CM captures the histogram of time domain samples at the wideband front end of the receiver (full downstream).

Upstream Histogram

Description:

- The upstream histogram provides a measurement of nonlinear effects in the channel such as amplifier compression and laser clipping. For example, laser clipping causes one tail of the histogram to be truncated and replaced with a spike. When the upstream histogram enable attribute is set to 'true', the CCAP will begin capturing the histogram of time domain samples at the wideband front end of the receiver (full upstream band). The histogram is two-sided; that is, it encompasses values from far-negative to far-positive values of the samples. The histogram will have a minimum of 255 or 256 equally spaced bins. These bins typically correspond to the 8 MSBs of the wideband analog-to-digital converter (ADC) for the case of 255 or 256 bins. The histogram dwell count, a 32-bit unsigned integer, is the number of samples observed while counting hits for a given bin and may have the same value for all bins. The histogram hit count, a 32-bit unsigned integer, is the number of samples falling in a given bin. The CCAP will report the dwell count per bin and the hit count per bin. When enabled, the CCAP will compute a histogram with a dwell of at least 10 million samples at each bin in 30 seconds or less. The CCAP will continue accumulating histogram samples until it is restarted, disabled or times out. If the highest dwell count approaches its 32-bit overflow value, the CCAP will save the current set of histogram values and reset the histogram, so that in a steady-state condition a complete measurement is always available.

US Impulse Noise

Description:

- The UsImpulseNoise object provides statistics of burst/impulse noise occurring in a selected narrow band. A bandpass filter is positioned in an unoccupied upstream band. A threshold is set, energy exceeding the threshold triggers the measurement of an event, and energy falling below the threshold ends the event. An optional feature allows the threshold to be set to zero, in which case the average power in the band will be measured. The measurement is time-stamped using the DOCSIS 3.0 field of the 64-bit extended timestamp (bits 9-40, where bit 0 is the LSB), which provides a resolution of 98 ns and a range of 7 minutes.
- The CCAP provides the capability to capture the following statistics in a selected band up to 5.12 MHz wide:
 - Timestamp of event
 - Duration of event
 - Average power of event
- The CCAP provides a time history buffer of up to 1024 events. In steady state operation, a ring buffer provides the measurements of the last 1024 events that occurred while the measurement was enabled.

Us OFDMA Active and Quiet Probe

Description:

- The purpose of upstream capture is to measure plant response and view the underlying noise floor, by capturing at least one OFDMA symbol during a scheduled active or quiet probe. An active probe provides the partial functionality of a network analyzer, because the input is known, and the output is captured. This permits full characterization of the linear and nonlinear response of the upstream cable plant. A quiet probe provides an opportunity to view the underlying noise and ingress while no traffic is being transmitted in the OFDMA band being measured.
- When enabled to perform the capture, the CCAP selects a specified transmitting CM, or quiet period when no CMs are transmitting, for the capture. The CCAP sets up the capture as described in [MULPIv3.1], selecting either an active SID corresponding to the specified MAC address or the idle SID, and defining an active or quiet probe. The active probe symbol for this capture normally includes all non-excluded subcarriers across the upstream OFDMA channel, with pre-equalization on or off as specified in the MIB. The quiet probe symbol normally includes all subcarriers, that is, during the quiet probe time there are no transmissions in the given upstream OFDMA channel. For the quiet probe, the CCAP captures samples of at least one full OFDMA symbol including the guard interval. The CCAP begins the capture with the first symbol of the specified probe. The sample rate is the FFT sample rate (102.4 megasamples per second).

Us OFDMA MER per Subcarrier

Description:

- This item provides measurements of the upstream RxMER for each subcarrier. The CCAP measures the RxMER using an upstream probe, which is not subject to symbol errors as data subcarriers would be. The probes used for RxMER measurement are typically distinct from the probes used for pre-equalization adjustment. For the purposes of this measurement, RxMER is defined as the ratio of the average power of the ideal QAM constellation to the average error-vector power. The error vector is the difference between the equalized received probe value and the known correct probe value. If some subcarriers (such as exclusion bands) cannot be measured by the CCAP, the CCAP indicates that condition in the measurement data for those subcarriers.

Us Triggered Spectrum Capture

Description:

- Capture of upstream spectrum through a number of triggering means including free run, time stamp value, mini-slot number, MAC-SID, idle SID, symbol, event trigger, and IUC.
- Note that reliable US triggered spectrum capture is a top priority for PNM in general, as this has not yet been implemented following the specifications.

8.3.1. Summarizing the gaps

With a high-level understanding of each of the PNM test queries, Table 1 provides an overview of the needed support in DOCSIS 4.0 tools for FDD and FDX as of the writing of this paper. As can be seen in Table 1, it is expected that PNM testing on FDD channels will be less impacted than PNM testing on FDX channels due to the intrinsic complexities of FDX. In general, when PNM tests are run on a channel configured in the FDX band, to perform downstream PMN tests like DS Symbol Capture and

NoisePowerRatio, the RBA for the sub-band must be set in the downstream direction, and upstream PNM tests will need the sub-band to be configured in the upstream direction while the test is performed. Cable operators deploying FDD or FDX will experience new challenges. Having proper tools, especially proper PNM tools, will enable cable operators to be better positioned to effectively and quickly troubleshoot complex problems in their HFC networks.

Table 1. Impact of FDD and FDX on DOCSIS 4.0 PNM tests

DOCSIS PNM Test	FDD Impact	FDX Impact
DS Symbol Capture (CM and CCAP)	None	RBA configured for DS, Testing required – investigation required
DsOfdmNoisePowerRatio (CCAP/Spectrum)	None	RBA for sub-band used on target DS – investigation required
Spectrum Analysis Full Band Capture	More bins, more data	Dual direction, more bins, more data, more complexity, filters in modems may differ by vendor – investigation required
CmDsOfdmChanEstimateCoef	None	Only possible when RBA for TG is set in DS direction, other dependencies involved
CmDsConstDispMeas	None	Uncertain. There may be an ability to capture I and Q values in two directions – investigation required
ModulationOrderOffset	None	None expected
CmDsOfdmRxMer	None	None Expected
CmDsOfdmMerMargin	None	None Expected
CmDsOfdmFecSummary	None	Test runs for several minutes which may be impacted based on RBA scheduling – investigation required
CmDsHist	None	Undetermined what happens with this test when in FDX operation – investigation required

Upstream Histogram	None	Uncertain how to measure the FDX band from 108 MHz to 684 MHz and how to account for any co-channel interference and echo cancellation
Us Impulse Noise	None	Recommend that this test does not apply to the FDX band
Us OFDMA Active and Quiet Probe	None	Multiple issues to address such as configuring all RBAs for the TG and configuring active probes – investigation required
Us OFDMA Rx Power	None	None expected
Us OFDMA RxMER per Subcarrier	None	If other transmission groups are operating in a DS direction, the RxMER values for the tested OFDMA channel could be lower – investigation required
Us Triggered Spectrum Capture	Wider spectrum, more bins, more data	Wider spectrum, more bins, more data, in addition, for SID filtering all TGs and channels must be sync'd to same TG to get a valid measurement

8.3.2. What is required?

As shown in Table 1, several PNM tests are impacted when FDX has been configured for operation. The most obvious impact is seen in the FDX allocated spectrum and the need for the test to be performed when the specific sub band that the channel is configured to use is set in the correct direction. As more FDX segments are brought into service over the next 12 to 24 months, these issues will be overcome and the specific impacts, and their workarounds, will be better understood through additional testing and use. FDX also has the concept of sounding, and the usage of that data for PNM related activities has yet to be explored fully and will be studied once enough DOCSIS 4.0 FDX modems and nodes are deployed.

FDD, in contrast, does not have as significant of an impact to the PNM testing because these channels are all in place today; the issue here is that there are more of them and a greater frequency span to cover for tests including FBC in the modem.

Another class of products will be employed for DOCSIS 4.0 deployments: smart amplifiers. These updated components in the plant will be more bi-directional and have capabilities for sampling and some PNM testing as well, which will allow the operator to have another valuable testing point in the network for measurement and troubleshooting analysis.

There will be challenges that will be met during this period. Some of the challenges are outlined as follows, and are expected to serve as the foundation of an industry project plan to resolve these challenges:

- Vendors of PNM tools must adapt some of the testing to accommodate the FDD and FDX impacts and the implementation of smart amplifiers that are added to the plant.
- Cable company operations and back-office teams dealing with the increased amount of data coming from devices in the field.
- Scheduling of testing for FDX channels.
- Compliance of FDD- and FDX-capable CMTSs and CMs with the DOCSIS 4.0 specifications
- Interoperability among vendors' products (CMTSs, nodes, CMs, etc.) with both FDD and FDX
- Clearly defined FDD and FDX PNM test and query specifications from the standards and specifications development organizations
- Adoption of FDD and FDX PNM test and query DOCSIS specifications by vendors
- Standards and specifications development groups exploring further the usage of FDX sounding data for PNM testing; the addition of test capabilities in smart amplifiers and other plant equipment is an area ripe for study and requirements creation that will likely see more activity as more FDX plant and modems become available.

In order to address the challenges identified above and develop the proposed industry project plan, the following groups will need to collaborate as has historically been done in specifications development:

- Chipset vendors
- CMTS vendors
- CM vendors
- PNM tool vendors
- Standards and specifications development organizations
- Cable operators

Input and collaboration from all parties are essential for bridging the gaps identified in this document.

9. Conclusion and Future Outlook

The past several years have seen significant progress in the use of equipment and PNM tests. This has provided a new level of capability for cable operators to offer more services at higher data rates to more subscribers. As the industry looks to the near future with DOCSIS 4.0 technology and FDX and FDD updates to the plant, these well-known tests will continue to provide significant insight into the health and operation of our cable networks. Operators can rest a little easier knowing that the same applications and methods being used today can be extended, with modifications, into the DOCSIS 4.0 networks that will soon be deployed. While there are still areas for continued innovation and standards work, the authors feel optimistic that the groundwork that has already been implemented will continue to provide operators with actionable data that can help to keep our networks healthy.

However, there is an opportunity for improvement. Many PNM tests were defined in the DOCSIS 3.1 specifications and implemented by equipment vendors. Some of those tests have been invaluable, such as downstream RxMER per subcarrier. However, other tests, such as UTSC giving insight into return path noise and upstream RxMER per subcarrier have been inadequately adopted by vendors. In preparation for

deployment of DOCSIS 4.0 technology, this paper has identified gaps in PNM which are needed to support DOCSIS 4.0 deployments. While specifications can include recommendations, it is up to the cable operator community to decide if PNM test functionality, such as return path monitoring (e.g., UTSC) and extended frequency FBC are valuable tools or not. It is incumbent on cable operators to hold discussions with vendors and determine the priorities. Should enhancing PNM be a priority or not? This is a question cable operators must determine and communicate to their vendor partners.

Abbreviations

ADC	analog-to-digital converter
CCAP	converged cable access platform
CCI	co-channel interference
CM	cable modem
CMTS	cable modem termination system
CPD	common path distortion
CPU	central processing unit
DAA	distributed access architecture
dCCAP	distributed CCAP
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	downstream
DUT	device under test
EC	echo canceller
ECT	echo canceller training
ESD	extended spectrum DOCSIS
FBC	full band capture
FDD	frequency division duplexing
FDX	full duplex [DOCSIS]
FEC	forward error correction
FFT	fast Fourier transform
Gbps	gigabits per second
GHz	gigahertz
GPS	Global Positioning System
GPU	graphics processing unit
HFC	hybrid fiber/coax
HTTP	Hypertext Transfer Protocol
I	in-phase
iCCAP	integrated CCAP
IG	interference group
ISM	industrial, scientific, and medical
IUC	interval usage code
L2TP	Layer 2 Tunneling Protocol
LSB	least significant bit
LTE	long term evolution
MAC	media access control
MER	modulation error ratio
MHz	megahertz
MIB	management information base

MULPI	MAC and upper layer protocols interface
NPR	noise power ratio
ns	nanosecond
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OSSI	operation(s) support system interface
OUUDP	OFDMA upstream data profile
PHY	physical layer
PIM	passive intermodulation
PNM	proactive network maintenance
PMA	profile management application
PON	passive optical network
Q	quadrature
QAM	quadrature amplitude modulation
QoE	quality of experience
RB	resource block
RBA	resource block assignment
RF	radio frequency
R-PHY	remote PHY
R-MACPHY	remote MAC PHY
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers
SID	service identifier
SNMP	Simple Network Management Protocol
SSH	secure shell
TCP	total composite power
TFTP	Trivial File Transfer Protocol
TG	transmission group
US	upstream
UTSC	upstream triggered spectrum capture
WLAN	wireless local area network
YANG	yet another next generation

Bibliography & References

[1] Campos, A., E. Cardona, L. Raman, “Pre-equalization Based Pro-Active Network Maintenance Methodology.” Proceedings of SCTE Cable-Tec Expo, Philadelphia, PA, 2008.

[2] *PNM Best Practices: HFC Networks (DOCSIS 3.0)*, Cable Television Laboratories.
<https://www.cablelabs.com/specifications/proactive-network-maintenance-using-pre-equalization>

[3] Data-Over-Cable Service Interface Specifications DOCSIS[®] 3.1 Physical Layer Specification CM-SP-PHYv3.1-I19-211110, Cable Television Laboratories. <https://www.cablelabs.com/>

[4] Data-Over-Cable Service Interface Specifications DOCSIS[®] 4.0 Physical Layer Specification, CM-SP-PHYv4.0-I02-200429, Cable Television Laboratories. <https://www.cablelabs.com/>

[5] Hranac, R., “Understanding Band Splits in Two-Way Networks,” *Broadband Library*, Spring 2020, <https://broadbandlibrary.com/understanding-band-splits-in-two-way-networks/>

[6] <https://www.10gplatform.com/>

[7] <https://www.cablelabs.com/10g>

[8] Data-Over-Cable Service Interface Specifications DOCSIS[®] 4.0 MAC and Upper Layer Protocols Interface Specification CM-SP-MULPIv4.0-I05-220328, Cable Television Laboratories. <https://www.cablelabs.com/>

[9] Data-Over-Cable Service Interface Specifications DOCSIS[®] 4.0 CCAP[™] Operations Support System Interface Specification CM-SP-CCAP-OSSIV4.0-I07-220629, Cable Television Laboratories. <https://www.cablelabs.com/>

[10] U.S. DEPARTMENT OF COMMERCE, National Telecommunications and Information Administration, Office of Spectrum Management “United States Frequency Allocations: The Radio Spectrum” frequency allocation chart. <https://www.ntia.doc.gov/page/2011/united-states-frequency-allocation-chart>

[11] SCTE 270 2021 Mathematics of Cable

[12] Chrostowski, J. et al, “Leakage in a High Split World: Detecting and Measuring Upstream Leakage Levels in a One Gpbs Symmetrical High Split Hybrid Fiber Coax Network.” Proceedings of SCTE Cable-Tec Expo 2020. <https://www.nctatechnicalpapers.com/Paper/2020/2020-leakage-in-a-high-split-world>

[13] Hranac, R., “A Quick Look at S-Parameters,” *Broadband Library*, Winter 2019, <https://broadbandlibrary.com/a-quick-look-at-s-parameters/>

Hybrid Fiber Coaxial (HFC) Spectrum Efficiency and Quality

Systematically Evolving Networks using Profile Management Platform (PMA)

A Technical Paper Prepared for SCTE by

Jay Liew, Advanced Analytics Architect/Charter Communications

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Analyzing Plant Conditions	4
2.1. Using data to determine network performance	4
3. Profile Management Application	6
3.1. Determining gains and losses.....	7
3.2. Potential gain from a live port.....	8
4. High-split Roll-off	9
4.1. Initial feedback	9
4.2. Utilizing OFDM channel on a high-split deployment	11
5. Autonomia Platform.....	11
5.1. Data Ops	12
5.2. Model Ops	13
6. Conclusion.....	15
Abbreviations	17
Bibliography & References.....	17
Acknowledgements	18

List of Figures

Title	Page Number
Figure 1 – Med Var plot of Mac Domain A.....	4
Figure 2 – Med Var plot of Mac Domain B.....	4
Figure 3 – LTE impaired cable modem.....	5
Figure 4 – Cable Modem Exhibit 1.....	7
Figure 5 – SNMP Polling Architecture	12
Figure 6 – Modern Streaming Telemetry Data Arthitecture.....	13

List of Tables

Title	Page Number
Table 1 – DS RxMER to QAM Level Mapping.....	6
Table 2 – Lab Performance of Cable Modems using Flat-Configured Profiles versus Customized Profiles.....	8
Table 3 – Potential Gain/Loss comparison using profiles with non-FEC and FEC adjustment.....	9

1. Introduction

During the [COVID-19 Pandemic](#), the surge of broadband bandwidth utilization grew 16% for downstream (DS) consumption while upstream (US) increased 33% from March to May 2020. The resiliency of the hybrid fiber-coaxial (HFC) network delivered nationwide service as subscriber demand continued to increase. As demand for broadband connectivity services continues to grow, Charter Communications and other multiple service operators (MSOs) must utilize intelligent automation (IA) driven by artificial intelligence (AI) and machine learning (ML) technology to address HFC network capacity growth and issues.

Plant upgrades such as high-split are underway and, eventually, DOCSIS 4.0. To defer the high costs of future enhancements and improve capacity, MSOs can extract maximum capacity from existing plants. They can also use this same solution and approach to augment bandwidth capacity for new deployments and network upgrades.

The HFC network is complex. Every plant is unique and requires customization to address individual plant conditions. Using Charter's network data-driven platforms, we can determine how to design, build, or evolve our network. For example, in-home modem telemetry is necessary to fix impairments in the plant. Telemetry collected from devices in the plant transmits the device's health and reports the condition of the plant.

In this paper, we'll discuss the techniques used to evaluate plant conditions and how Charter's Autonomia is used to repair them. Autonomia is one of Charter's data platforms used to standardize data and software engineering by delivering Charter's network data needs under one program and making all network data and models available companywide. Profile Management Application (PMA) is a use case running on Autonomia that aims to focus on the entire data life cycle, including:

- Streaming and data acquisition capability;
- The ability to analyze data using data engineering for governed analytics data sets with ML capability;
- An automation platform that focuses on closed-loop implementations for various use cases.

We'll also discuss optimizing spectrum usage on plants with ingress and roll-off impairments on high-split deployments.

¹ Analytics from Charter Data Technologies Group.

2. Analyzing Plant Conditions

Telemetry data determines the state of the network and provides information that's not easily detectable. Often, subscribers may not be aware of network impairments. However, the network operator must be aware of impairments in order to fix issues, ensure subscriber quality of experience, and meet subscriber expectations.

Proactive Network Management (PNM) data contains critical telemetry information for cable operators to proactively manage and mitigate plant issues. Receive modulation error ratio (RxMER) data provides the status of a subscriber's home network and describes the condition of a service group.

2.1. Using data to determine network performance

Figure 1 and **Figure 2** are sample snapshots of media access control (MAC) domains from different cable modem termination systems (CMTSes). Each dot represents a single cable modem that represents the associated median RxMER value and the variance for all subcarriers on the orthogonal frequency division multiplexing (OFDM) channel. This methodology allows operators to see the current condition of the network, which they can trend over time.




Figure 1 – Med Var plot of Mac Domain A

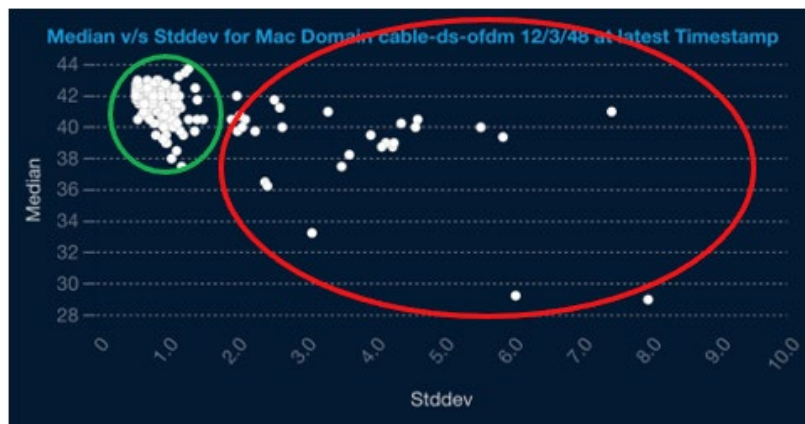


Figure 2 – Med Var plot of Mac Domain B

Using the Med Var plot observations, operators can anticipate a scatter plot pattern to be created from OFDM channels operating in the Long Term Evolution (LTE) frequencies. The OFDM channels shown in the above samples range from approximately 680 MHz to 772 MHz.

Note: This pattern is only identified on plants running on LTE interfered OFDM spectrum and may not apply to other kinds of impairment such as high-split roll-off.

The clusters in the top-left green circles in **Figures 1 and 2** denote healthy cable modems with high median RxMER and low variance. The cable modem on the bottom right denotes unhealthy cable modems that are most likely in partial service.

As the level of variance increases, the severity of impairment increases, as represented by the clusters in the red circles. Our analysis of the RxMER variance for CMs allowed us to classify the modems to a degree of impairment.

Below are our field data observations:

- Standard deviation of <1.0 dB represents cable modems with very consistent RxMER values across the channel. Cable modems lock on the highest-profile signal quality supports.
- Standard deviation of >1.0 dB and <3.0 dB represents cable modems with moderately inconsistent RxMER values across the channel. Cable modems may not be able to lock on the highest profiles.
- Standard deviation of >3.0 dB and <5.0 dB represent cable modems with highly inconsistent RxMER values across the channel. Cable modems may not be able to lock onto more than one profile.
- Standard deviation of >5.0 dB represents cable modems with extremely high inconsistent RxMER values across the channel. Cable modems will not be able to lock onto more than one profile. In many instances, cable modems will exhibit partial-service characteristics.

Most impairments from the data (680 MHz-772 MHz) are LTE-related, as portrayed in **Figure 3**. Using the RxMER data, we correlated the profile posture of cable modems from the Mac Domains. We also collected profile postures for each cable modem CMTSes in **Figures 2 and 3**.




Figure 3 – LTE impaired cable modem

We configured default fixed profiles at 4096 QAM, 1024 QAM, and 256 QAM. Data from live plants indicated the dynamic nature of these profile changes was visible from the hottest to the coolest part of the day, from day to day, season to season.

3. Profile Management Application

Intelligent automation can change the configuration of the CMTS many times a day and should be evaluated to ensure minimal impact on the network. For example, a radio frequency (RF) plant can lose capacity instead of gaining capacity when using PMA profiles without considering Forward Error Correction (FEC).

The PMA software solution optimizes the use of the spectrum by automating the generation of DOCSIS 3.1 profiles and customizing profiles to the condition of the plant. Ultimately, using PMA optimized with FEC data, will enable gains in capacity.

The RxMER and bit-loading thresholds determine the profile a CM runs. The **DS RxMER to QAM Level Mapping** in **Table 1** is specified in the DOCSIS PMA Technical Report, and guides RxMER and bit-loading thresholds.

Table 1 – DS RxMER to QAM Level Mapping

Constellation / Bit loading	CNR / MER (dB)
16 QAM	15.0
64 QAM	21.0
128 QAM	24.0
256 QAM	27.0
512 QAM	30.5
1024 QAM	34.0
2048 QAM	37.0
4096 QAM	41.0
8192 QAM	46.0
1634 QAM	52.0

Using lab tests and data from Charter's network, we concluded that low-density parity check (LDPC) is highly efficient for correcting errors. Moreover, the RxMER thresholds can several dBs below the specification, as shown in Table 1, and continue to work at the highest-available profile (4096 QAM). We performed comprehensive testing in Charter's lab and verified it in production using RxMER readings and the profile posture of those devices in the network.

Figure 6 shows a cable modem with a median of 35.75 dB and a standard deviation of 1.43 dB from the latest RxMER readings. The modem exhibited profile changes from 12 to 10-bit flat profiles during the day, and the other way around in the early morning hours. The graphic also shows that LDPC can correct error packets even if the variance of the signal quality is moderate, with parts of the channel dipping below the 35.0 dB threshold.




Figure 4 – Cable Modem Exhibit 1

The use of FEC to baseline RxMER relative to QAM Level Mapping is critical to maximize the gain, by calculating profiles, as aggressive as FEC can handle. Conservative RxMER to QAM Level Mapping can lead to loss of capacity in plants with relatively clean environments.

The data from CMTSes determine that 70-95% of cable modems operate on the highest bit rate of 4096 QAM at any given time. Varying conditions in the plant from day to day and hour by hour cause a disparity in performance. Unless cable modems can operate at higher bit rates, there's no capacity gain for CMs running on the highest bit-modulation available. However, capacity increases on cable modems not running at the highest bit rate.

This evidence leads us to ask the following questions:

- How many modems can be removed from partial service using customized profiles for each service group?
- How much extra capacity can be gained from cable modems not operating on the highest profile?

3.1. Determining gains and losses

We can determine cable modem speed gains and losses using PMA versus flat-configured profiles. Our configuration includes five 6-MHz, 8 bits-per-hertz, Single Carrier QAM (SCQAM) channels, and one 96 MHz-wide OFDM channel.

We injected noise into the channel spectrum on 12 cable modems, running on 12-bit profiles for our experiment. By injecting noise into the channel, we changed the 12 modems operating 12-bit profiles to three partial-service modems and changed nine modems to run on 8-bit profiles. We hypothesized the best speed gains would impact modems with partial service.

As portrayed in **Table 2**, traffic passed through the modems that used the SCQAM and OFDM channels. The impaired modems passed traffic entirely on the SCQAM channel. The PMA-generated profiles increased speed and capacity for impaired modems and modems running on lower fixed-modulation profiles.

Table 2 – Lab Performance of Cable Modems using Flat-Configured Profiles versus Customized Profiles

Cable Modems	Non-impaired		Impaired		Impaired PMA	
	Profiles **	Speed (Mbps)	Profiles **	Speed (Mbps)	Profiles ***	Speed (Mbps)
CM1	0,1,2,3	907	0,1,(2),(3)	665	0,1,2,3	836
CM2	0,1,2,3	907	(CH)0,1,2,3	175	0,(1),(2),3	655
CM3	0,1,2,3	907	0,1,(2),(3)	665	0,1,2,3	836
CM4	0,1,2,3	907	(CH)0,1,2,3	175	0,(1),(2),3	655
CM5	0,1,2,3	908	0,1,(2),(3)	665	0,1,2,3	836
CM6	0,1,2,3	907	0,1,(2),(3)	665	0,(1),2,3	811
CM7	0,1,2,3	907	0,1,(2),(3)	665	0,(1),2,3	811
CM8	0,1,2,3	903	0,1,(2),(3)	665	0,(1),2,3	811
CM9	0,1,2,3	907	(CH)0,1,2,3	175	0,(1),(2),3	655
CM10	0,1,2,3	907	0,1,(2),(3)	665	0,(1),2,3	811
CM11	0,1,2,3	907	0,1,(2),(3)	665	0,1,2,3	836
CM12	0,1,2,3	906	0,1,(2),(3)	665	0,1,2,3	836

** Flat profiles where 0 = 6 bits, 1 = 8 bits, 2 = 10 bits, 3 = 12 bits, CH = impaired

*** Customized profile where 0 = 6.26 bits, 1 = 10.76 bits, 2 = 10.38 bits, 3 = 7.63 bits

NOTE: PMA profiles generated using FEC-adjusted RxMER and bit-loading thresholds

The data also indicated cable modems removed from partial service gained an average of 7.36 bits per hertz across the OFDM channel, a capacity increase of 480 Mbps for partial-service modems. For modems running on 256 QAM, the capacity increased from 8 bits per hertz to profiles utilizing an average of 10.76 bits per hertz and 10.38 bits per hertz, gaining more than 2 bits per hertz with speeds between 146 Mbps and 171 Mbps. This change represents a significant capacity increase for cable modems operating on low-modulation profiles and partial-service modems, thereby increasing modem speed for those cable modems.

3.2. Potential gain from a live port

Actual field experience works differently from what we simulate in the lab. In the field, a more diverse set of impairments of varying degrees require classification and clustering. Using the Mac Domains in **Figure 2** (above), we calculated the gain loss potential, identified cable modems that are impaired.

The CMTS can run on four profiles with a maximum of four segments. We tested two scenarios, one without FEC and one using FEC, optimizing RxMER bit-loading thresholds.

The results are summarized in **Table 3** as follows:

Table 3 – Potential Gain/Loss comparison using profiles with non-FEC and FEC adjustment

Cable Modems (206 total)					
Flat Profile Configuration		Non FEC PMA Profile Configuration		FEC PMA Profile Configuration	
Bit Load (bits per hertz)	# CM	Bit Load (bits per hertz)	# CM	Bit Load (bits per hertz)	# CM
12	182 (88.3 %)	10.4	182 (88.3 %)	12	182 (88.3 %)
10	14 (6.76%)	10.0	2(0.9%)	11.6	8 (3.8%)
8	5 (2.4%)	7.8	18 (8.7%)	10.4	10 (4.8%)
		5.5	4 (1.9%)	7.1	6 (2.9%)
0	5 (2.4%)	0	0	0	0

When calculating PMA without FEC to adjust RxMER bit load mapping, we discovered:

- Capacity loss by 1.6 bits per hertz for CMs operating on the highest profile;
- General loss of total capacity by 14%;
- Partial-service cable modems are taken out of partial service.

When calculating PMA using FEC to adjust RxMER bit load mapping, we discovered:

- Cable modems running on the highest profiles don't lose capacity;
- Partial-service cable modems are taken out of partial service;
- Increased capacity for cable modems that are not on the highest QAM level.

4. High-split Roll-off

The path to 10G is underway. Using trials, we proved the viability of frequency-division duplexing (FDD) DOCSIS 4.0 by demonstrating the ability to deliver multi-gigabit symmetrical service. Network evolution is complex requiring expensive plant hardware configurations and upgrades, necessitating network operators to take the necessary steps to evolve the network.

4.1. Initial feedback

The **Figure 8 – (Lab) Sample Cable Modem from Cluster A** mimics a live field environment consisting of a node with a cascade of five amplifiers (1.2 GHz), with 32 passive taps (1.0 GHz), with 20 devices on five taps. The most salient pattern is the three distinct clusters (A, B, and C). The cable modem data from the RxMER data indicates the roll-off behavior is different between cable modems and can be completely different within each cluster.

A is the best performing cluster, demonstrating moderate impairment with cable modems located on the second and third amplifiers, with four-seven passive devices.




Figure 7 – (Lab) 1.0-1.2 GHz Spectrum Med Var Plot

Figure 8 is a sample cable modem from *Cluster A* in **Figure 7** and demonstrates a usable spectrum through the entire 192 MHz OFDM channel.




Figure 8 – (Lab) Sample Cable Modem from Cluster A

Cluster B (**Figure 7**) demonstrates extreme impairment with high median RxMER, with cable modems on the fourth amplifier with 11-13 passive devices (shown in **Figure 9**), with usable OFDM spectrum through 1.18 GHz.




Figure 9 – (Lab) Sample Cable Modem from Cluster B

Cluster C (Figure 7) is a single cable modem with extremely low signal quality with extreme impairment as shown in **Figure 10**. The cable modem is on the fifth amplifier with 20 passive devices. Some cable modems deep in the network cannot report OFDM telemetry due to channels in partial service.




Figure 10 – (Lab) Sample Cable Modem from Cluster C

4.2. Utilizing OFDM channel on a high-split deployment

The signal quality in the Med Var charts and RxMER data from cable modems indicates variability for every modem in the service group. The deployment variance is caused by the roll-off, which is caused by the 1.0 GHz taps. This is a unique situation for how an operator can create a profile that works on the channel without sacrificing spectrum and capacity. Using this analysis, we can systematically develop a roadmap for high-split deployments as follows:

- Flat profiles don't work where the roll-off region renders higher parts of the spectrum unusable for most cable modems.
- PMA profiles must operate at 1.0-1.2 GHz to work with the roll-off region.
- Operators can incrementally upgrade plant components such as taps using PMA profiles.

To mitigate the roll-off region entirely for an OFDM channel running on 1.0-1.2 GHz, replacement of taps in the plant will be required. By using PMA we can defer the initial costs of the deployment. Using the data within our lab, we can learn and derive a systematic plan to upgrade a high-split plant. This will give us the ability to postpone upgrades and reduce the initial costs of deployments, simultaneously boosting capacity until upgrades are necessary.

5. Autonoma Platform

As mentioned earlier, PMA is running on Autonoma, that focuses on Data Ops, by building fast and reliable data pipelines, feeding Model Ops, by turning data into action. As the network continues to evolve, data has become prominent in architecting the modern network. The evolution applies to the physical access network architecture and how we design data and software. It isn't practical for large MSOs like Charter to operate approximately 5,000 CMTSes in a 54 million homes-passed footprint using traditional polling such as Simple Network Management Protocol (SNMP) or command-line interface (CLI).

There is cost associated with supplementing and accommodating polling capability and increasing the number of devices. Demand is extensive for data, since various consumers have differing

requirements. It is impossible to control multiple data-polling agents, polling disparate management information bases (MIBs) at disparate frequencies during various times of the day.

5.1. Data Ops

“Growing scale and the increasing use of automation in next-generation enterprise networks require a modern, more efficient approach to network data capture and analytics” – Bo Lane, VP Global Engineering, Kudelsky Security

One of the challenges for operators is how to promptly acquire network data from devices. Two classes of data are required to operate the network:

- **Bulk data:** Telemetry data that is collected at predefined frequencies from all devices on the network;
- **Real-time polling:** Data collected from devices, on an ad-hoc basis, that identifies the devices' conditions at a particular moment.

As the network modernizes, it becomes more complicated to operate, and requires telemetry to operate the network. **Figure 11** is an example of SNMP polling architecture. As the number of devices increase, so does the complexity, and reliability decreases for bulk-data ingestion.

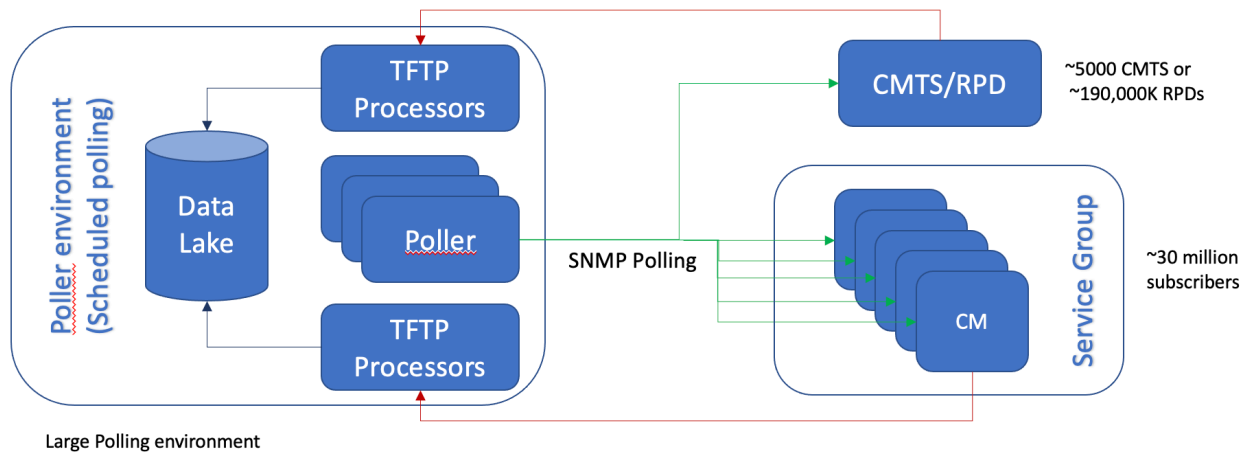


Figure 5 – SNMP Polling Architecture

Alternatives to polling methods like push-based data acquisition are becoming more relevant as real-time data becomes more critical. For example, if the signal quality for a particular channel degrades or experiences latency issues on a specific device, details of the event close to real-time are required. The responsibility to acquire data from remote polling mechanisms shift to the device to provide the telemetry at a granular level.

Figure 12 (below) illustrates a modern data architecture.

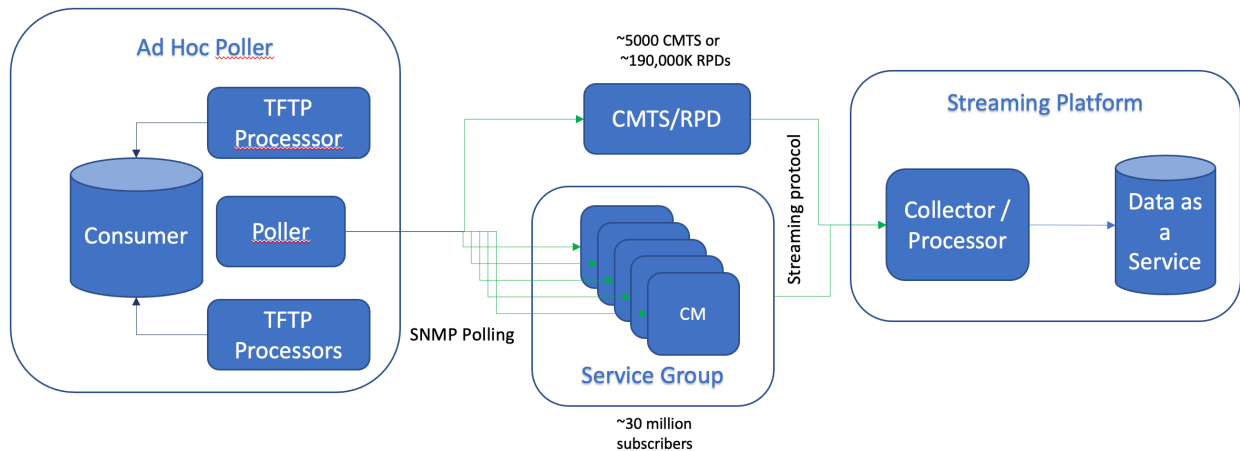


Figure 6 – Modern Streaming Telemetry Data Arthitecture

5.2. Model Ops

The DOCSIS 3.1 PHY specification provides many capabilities to gather telemetry data to conduct PNM. RxMER, a signal-to-noise ratio, is one of the most helpful signals for diagnosing the condition of individual modems on the Charter network. RxMER provides measurements for each subcarrier in the OFDM (downstream) or OFDMA (upstream) channel. We can determine the quality of the signal subscribers receive, including a range of potential impairments that can impact service experience, by analyzing the data's level, shape, and spread.

Some of these impairments, shown in **Figure 13**, include conditions such as:

- **Amplitude Ripple:** Caused by improper network alignment, micro-reflections, or missing end-of-line terminators;
- **Suckout:** Caused by unterminated cable, smashed cable, or repeating divets;
- **LTE Ingress:** Caused by signal leakage (the passage of an outside signal into a coax cable), which can be remediated using Exclusion Bands;
- **Roll-off:** A result of improper balancing, bad amplifiers, or exceeding amplifier specification;
- **Standing Wave:** Can be used to determine the distance to a fault and is the initial premise behind PNM.

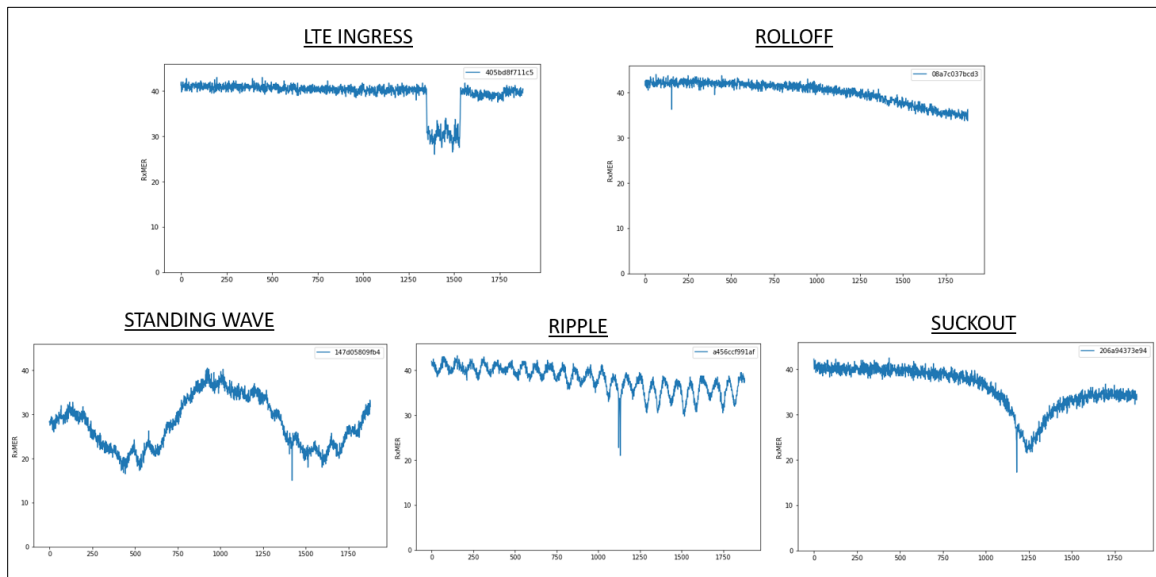


Figure 13 – Impairment Classifications

The classification impairments can identify potential network faults with a known solution. Using the data patterns, we can score and rank-order potential candidates to resolve issues in the field. The challenge is these patterns aren't self-evident, and manually reviewing the available data can be difficult and time-consuming.

We leveraged heuristics and unsupervised learning using AI and ML methods to acquire sufficient sample impairments. Paired with subject-matter-expert consultation, we provided a suitable training set for supervised machine-learning model development.

Using ML, we can monitor data signals 24x7x365, automatically identifying and classifying various impairment types. We can also detect if several modems are displaying the same impairment—a clear indication of an outside plant (OSP) issue—to isolate the impact in homes versus OSPs using clustering and nearest-neighbor modeling, illustrated **Figure 14** (below).




Figure 14 – ML Nearest-neighbor Algorithm Identifying a Cluster of Impairments

6. Conclusion

As the HFC network changes, data is necessary for the evolution of the network. Furthermore, we can use data to determine how and when to make changes to the network.

Software solutions are cost-effective ways to troubleshoot and address impairment issues in the plant. While software solutions may not solve plant issues, they provide time for operators to address issues without immediately sending a technician to a site.

PNM data allows the operator to assess plant conditions in real time by the subscriber, service group, CMTS, and region. We can use various methods to understand a service group. For example, we can learn traits from a collection of devices using clustering. PMA also allows operators to fix impaired cable plants.

We tested plants with LTE Ingress impairments and determined adjustments to the RxMER and QAM-level mappings were required. Our findings are as follows:

- RxMER and QAM-level mappings that are too conservative will result in a capacity loss on the OFDM channel, and FEC is needed to optimize PMA Engine;
- We optimized RxMER and QAM-level mappings for maximum capacity gain;
- Optimizing PMA can gain 30-40% (aggregate capacity) for cable modems not operating on the highest QAM level;
- 3-5% capacity gain for the overall channel;
- We removed cable modems from partial service using PMA;
- Cable modems operated at higher QAM levels when using PMA.

PMA can assist with transitions such as high-split deployments. Our lab-simulated plant indicated a relationship between signal quality, the amplifier, and tap cascades. We based the following assumptions on plants with 1.2 GHz amplifiers working with 1.0 GHz taps, with the cable modem spectrum exhibiting roll-off:

- Fixed modulation profiles don't work on high-split deployments because of spectrum roll-off;
- Dynamic profiles are required to direct spectrum roll-off to utilize a 1.0-1.2 GHz spectrum;
- PMA provides operators the ability to use a 1.0-1.2 GHz spectrum based on the subcarrier RxMER;
- PMA allows operators to strategically execute plant upgrades.

The evolution of modern networks requires ML to process the increasing volume of data and the number of data sets from devices which include:

- Transitioning from pull-based polling methods (e.g., SNMP) for bulk data collection;
- Modernizing bulk data collection using push-based methods, such as gRPC network management interface (gNMI), message queueing telemetry transport (MQTT), Kafka, and other stream-based protocols;
- Only using pull-based polling methods, such as SNMP, for limited, ad-hoc situations.

Machine learning is vital for how networks operate in the future, including:

- Leveraging heuristics and ML to get an adequate number of sample impairments to develop and train models;
- Data-signal monitoring on a 24x7x365 basis to automatically identify and classify various impairment types;
- Clustering classified impairments and using nearest-neighbor modeling to localize impairments.

Abbreviations

AP	Access Point
CLI	Command Line Interface
CM	Cable Modem
CMTS	Cable Modem Termination Service
DS	Downstream
FEC	Forward Error Correction
GHz	Gigahertz
gNMI	gRPC Network Management Interface
gRPC	Google Remote Procedure Call
HFC	Hybrid Fiber Coax
bps	Megabits per second
MHz	Megahertz
MSO	Multiple System Operator
LDPC	Low Density Parity Check
LTE	Long Term Evolution
QAM	Quadrature Amplitude Modulation
OFDM(A)	Orthogonal Frequency Division Multiplexing (Access)
OSP	Outside Plant
PMA	Profile Management Application
PNM	Proactive Network Management
RxMER	Receive Modulation Error Ratio
SCQAM	Single Carrier QAM
SCTE	Society of Cable Telecommunications Engineers
SNMP	Simple Network Management Protocol
SNR	Signal Noise Ratio
US	Upstream

Bibliography & References

Data-Over-Cable Service Interface Specifications Technical Reports. DOCSIS 3.1 Profile Management Application Technical Report. CM-TR-PMA-V01-180530 (May 30, 2018)

Diana Goovaerts, [*Charter hits 8.5 Gbps in DOCSIS 4.0 test with Vecima*](#): Fierce Telecom, January 6, 2022.

John Edwards, [*Streaming telemetry challenges SNMP in large, complex networks*](#): Network World, September 29, 2020.

Acknowledgements

The research and analysis as discussed in this paper would not have been possible without the work of Charter's Network Technology Group and Data Technologies group, including:

Roger Stafford, Principal Engineer, Access Network Architect, Network Technology Group
Vinod Dani, Principal Engineer, Access Network Architect, Network Technology Group
Derik Johnson, Principal Data Scientist, Data Technologies Group
Karl Harrison, Principal Engineer, Access Network Architect, Network Technology Group
Justin Stiles, Principal Engineer, Access Network Architect, Network Technology Group
Diana Linton, Network Engineer, Access Network Architect, Network Technology Group

I Didn't See It Coming: The Rise of The Bot

A Technical Paper prepared for SCTE by

Don Jones

Director, Strategic Fraud Intelligence
Comcast Cable
Denver, Colorado
303-712-3588
Don_Jones@Comcast.com

Claire Nobles

Project Manager 5, Technical Fraud Management
Comcast Cable
Cheyenne, WY
303-246-1188
Claire_Nobles@Comcast.com

Andrew Frederick

Principal Engineer, Comcast Technology Solutions
Comcast Cable
Andrew_Frederick@Comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. What is a Botnet?.....	3
3. History of Botnet Technology.....	3
4. Infection Methods.....	4
4.1 Phishing.....	4
4.2 SMSHING.....	5
4.3 Malvertising.....	5
4.4 Game and Program Cracks and Cheats.....	6
5. Payload Types.....	6
5.1 Info Stealers.....	6
5.2 RATs & RESIPs.....	7
5.3 Crypto Jacking.....	7
5.4 Ransomware.....	8
6. Indicators of Risk.....	8
7. Conclusions.....	9
Abbreviations.....	10
Bibliography & References.....	10

1. Introduction

A single click of the mouse can cost you your house. A 2021 news article by network solution provider Barracuda measures “bad bots mak[e]ing up nearly 40% of all traffic” (N/A, 2021). It all starts with a single click. Every day, millions of people are asked to click on a link that could cost them everything they own. Malicious links are presented by email, by phishing websites, by dangerous ads on less than reputable websites, by gaming cheats and cracks, and even SMS (Short Message Service) text messaging.

With the exponential growth of the Internet, the threat of these kinds of attacks is supported by a burgeoning underground economy that has only increased the complexity and frequency of their attacks. Tactics are not limited to offering something for nothing. Recent phishing frauds focus on telling the victim something will happen unless they opt out and will often attempt to appeal to their victims at an emotional level. Malicious botnets have different purposes ranging from identity theft to distributed denial-of-service attacks against critical infrastructure. Victims of identity theft based on botnet infections rarely know how they were compromised, leaving the door open for victims to continuously be re-compromised. The end goal of most botnets is monetary gain through identity theft, but the proliferation of botnets also lends itself well to their use as a cyber weapon. The threat vectors vary for users, Internet Service Providers (ISP), retail companies or governments. Though there are legitimate uses of botnets, malicious payloads span a range from questionably legal tactics to blatantly malicious activity. This paper shares ways to identify the initial signs of danger, minimize exposure to these threats and to help bring focus to the recognizable indicators of malicious links.

2. What is a Botnet?

A bot, short for robot, is a software executable application that performs tasks that are automated and repetitive, typically via scripts on the internet. A botnet is a network of bots working together to complete a set of automated tasks that are managed and monitored by a command-and-control system (C2). This C2 is also referred to as a botmaster or botherder. Botnets act like many soldiers being issued orders and reporting back to their general. Botnets are typically used because they can complete tasks quicker and more efficiently than a human.

Botnets have both legitimate and illegitimate purposes. A few examples of both types of botnets are:

- Legitimate botnets
 - o Search engines
 - o Chat bot communications
- Illegitimate botnets
 - o Info-stealers
 - o RAT's (Remote Access Trojan/Tool)
 - o Ransomware
 - o Crypto Miners

Legitimate botnets are essential in enabling the discovery of added content like new websites. This paper focuses on infection methods and payload types of illegitimate botnets. Malicious botnets seek to subvert security controls and to generate revenue at the expense of their victims.

3. History of Botnet Technology

One of the first bots was ELIZA, a chatter bot, that was created at Massachusetts Institute of Technology (MIT) by Joseph Weizenbaum in 1966. According to an eBook Machines Who Think: A Personal Inquiry

into the History and Prospects of Artificial Intelligence, “ELIZA was intended to simulate – or caricature, as Weizenbaum himself suggests – the conversation between Rogerian psychoanalyst and patient, with the machine in the role of the analyst” (McCorduck, 1979). Since then, botnets have evolved drastically in sophistication and purpose.

Botnets use a variety of communication protocols. One of the original protocols is Hypertext Transfer Protocol (HTTP), which was until recently, the most used protocol by many websites. Many of these websites have now moved to Hypertext Transfer Protocol Secure (HTTPS), adding in a security encryption layer. One of the more popular protocols created and still being used is Internet Relay Chat (IRC), which was introduced in 1988 by Jarkko Oikarinen of the University of Oahu. Creators of more advanced botnets added the use of P2P (Peer 2 Peer) in addition to encrypted protocols like The Onion Router (TOR) SOCKS5 proxies and Simple Mail Transfer Protocol (SMTP). Most modern botnets rely on blockchain for a more distributed command and control communication network. Botnet communication keeps pace with all technological evolution.

As the technology of botnets has evolved so has the intent of the botnet. Since 1998, there have been many other malicious botnets that have been unleashed on the internet with names like Melissa, Code Red, Storm Worm, Mirai, and Smominru. A white paper titled The Historical Perspective of Botnet Tools best expresses the growth of botnets from a historical perspective. “As researchers continue to unveil the botnet trend and its mode of propagation within networking platforms the botmaster would continue to create different techniques meant to surpass the earlier botnet tool version used” (Osagie, Enagbonma and Inyang, 2019). These botnets have expanded their targets from attacking computers to cell phones, smart cars, and all types of Internet of Things (IoT) devices.

Overall, malicious botnets continue to evolve in their communication methods, target platforms and intentions over the decades. While botnet technologies have evolved, so have the infections methods.

4. Infection Methods

Today, most malicious bots are installed through several primary channels, email, text messages, malicious advertising and program/game cheats or cracks.

4.1 Phishing

Phishing refers to email sent with malicious links or attachments or fake websites that attempt to appear legitimate. Using social engineering tactics to gain the readers confidence or trust, carefully worded email with malicious attachments or links to malicious or fake websites designed to defraud the victim is one of the most effective cyber-attacks that exist. Attachments such as Adobe PDF (Portable Document Format), Microsoft Excel spreadsheets and even pictures can contain distinct types of malicious payloads as categorized below. Opening any attachment effectively “executes” the code inside that attachment with the associated program. Although anti-virus programs catch some of this malicious code, often the malicious code escapes detection through code obfuscation or encryption. More often, a phishing email contains a link to a malicious or cloned website.

Since most email uses Hypertext Markup Language (HTML), hyperlinks, or links to websites, which allows the text to be different than the underlying link to a website. For example, the text that reads “<http://GoodGuyWebsite.xyz/>” appears to point to GoodGuyWebsite.xyz. However, hovering the mouse reveals the link to the target Uniform Resource Locator (URL) is actually “www.BadGuyWebsite.xyz”. Another way of using formatting to forge URLs of familiar websites is the use of character manipulation. An example might appear to be a link to “www.maybeLegitimate.xyz/” but could be easily confused with

a link to “rnybeLegitimate.xyz” depending on the font. Careful examination reveals that the “M” has been replaced with an “R” and an “N”. Malicious actors often use the original company’s artwork and logos to increase the appearance of the email’s validity and the links to malicious sites. Phishing web sites, using genuine appearing artwork, will ask the victim to login and provide their credentials, which is the goal of this type of attack. A shift in tactics was observed in May 2021 when security firm Proofpoint discovered that a malicious actor set up a fake site that appeared to be a video piracy site called BravoMovies.

Instead of encouraging the victim through reward or fear of missing out, the criminals claimed the victim’s free trial they had signed up for had expired and their credit card was going to be charged, unless they cancelled. According to the Proofpoint article, BazaFlix: BazaLoader Fakes Movie Streaming Service, the email contained only a phone number, where a human operator directed the victim to the malicious website. The cancellation instructions were only available through a downloaded Microsoft Excel Spreadsheet which contained the malicious code. Phishing schemes tend to lose efficacy with each additional victim action, but as Proofpoint pointed out, “...despite being counterintuitive, the techniques used by the threat actors in this, and similar, campaigns help bypass fully automated threat detection systems” (Larson and Mesa, 2021). SMShing, is like phishing in that the actors attempt to persuade the victim to perform an action, however, the risk surface of mobile devices is far greater than a personal computer.

4.2 SMShing

SMShing is the act of enticing victims to click on hyperlinks delivered via SMS messaging on their mobile device or by opening a malicious attachment such as a PDF, Microsoft Excel Spreadsheets, or pictures. Often, malicious actors impersonate large businesses that most people are familiar with. The efficacy of SMShing attacks is often remarkably high, given several varied factors. Gartner Research published the paper called Tap Into the Marketing Power of SMS, which quantified the efficacy of SMS messaging as a sales medium. “Various sources report SMS open and response rates as high as 98% and 45%, respectively — in contrast to corresponding figures of 20% and 6% for email” (Pemberton, 2016). URL shortening services help to create the illusion of legitimacy to links (i.e., the link tinyurl.com/amazon could be redirected to any unknown website), further complicating the ability to identify a SMShing attempt. Compounding the risk level, most mobile devices lack antivirus protection. SMShing is a form of social engineering, which often uses a time-based pretext.

A common technique is to employ a call to action that expires quickly, such as a limited time exclusive offer that expires in one hour. This prompts the targeted user to suspend judgment about risk/reward since there is often a time component that can cloud judgment. The objective is to convince the victim to open an attachment or visit a website that appears legitimate, that either steals the information the victim enters or installs malicious code. Another social engineering technique used is to convince the victim that not responding to the message will have consequences. For instance, a message saying a credit card will be charged if the victim does not act. While SMShing attacks grow in popularity, malvertising has matured significantly.

4.3 Malvertising

Malvertising (Malicious Advertising) is the act of spreading malware through web or application-based advertisements. Advertisements are one of the most profitable areas of online commerce. Criminals monetize on the advertising platforms by distributing malicious payloads through legitimate ad delivery platforms. The threat to online and application users increases in proportion to the reputation quality of the website or application.

Brand protection companies like Trustworthy Accountability Group (TAG) and White Bullet help to combat malicious ads from being placed on reputable websites. This concerted effort to protect specific brands shifts the delivery of malicious ads to websites or in applications that are less reputable. For instance, White Bullet and Digital Citizens Alliance (DCA) published a study in 2021 called Breaking B(ads) measuring the volume of malicious ads in video piracy websites and applications. “White Bullet reviewed 664 billion ad impressions and found that roughly one in three piracy websites and apps have risky advertising that exposes consumers to fraud and malware” (N/A, 2021). The more reputable a company is, the less of a likelihood that malicious ads will be placed on their website or application due to the corporate diligence in brand protection. The advantage of employing malvertising to deliver malicious payloads is that it does not require compromised hosts to spread a malicious payload and the audience reach is limited only by budget.

4.4 Game and Program Cracks and Cheats

Game and application activation cracks and cheats are created to bypass free trial periods on many different applications and games and to add hidden features like in game invincibility and hidden levels. Often packaged together with the trial application files available on torrent sites, the directions for applying the patch or crack explains to disable any antivirus protection to avoid “false” positives. In doing so, the victim purposely shut down their protection long enough to be infected. Game cheats can differ in that the malicious payload is not immediately installed but often lies in wait for the victim to invoke. Game cheats often bypass the antivirus identification by creating in game links to malicious payloads rather than delivering that payload itself.

5. Payload Types

5.1. Info Stealers

The Information Stealer (info stealer) is the most prolific malicious botnet payload type being distributed, often through malvertising, game/application cracks and cheats. Info-stealers focus on stealing the victim's information and exfiltrating it to the C2 servers. The information it extracts can be defined by the person running the campaign but usually includes system and browser information by default. Computer information, which includes operating system and system level information, antivirus's running, IP (Internet Protocol) address, machine name and all installed browsers stored information are usually default targets. The browser information consists of all cached URLs, all stored (or remembered) usernames and passwords, all cached form information and stored (or remembered) credit card information in clear text from all installed browsers. Most stealers (Redline, LokiBot, Dark Crystal RAT etc.) can also be configured to exfiltrate all PDF's, Microsoft Word documents, pictures, or text files in specific folders, like “My Documents,” “Desktop” etc. These payloads can be configured to search for cryptocurrency wallets, social media applications and Multifactor Authentication (MFA) keys as well. Exfiltration of the information takes mere seconds. The explosion in volume of these infections goes completely unnoticed by the victims but can be measured independently.

A 2021 white paper called An Analysis and Investigation of InfoStealers Attacks during COVID'19: A Case Study, attempted to measure the growth of info stealers in the window of the COVID-19 pandemic. “Resultantly, a significant increase in the quantity and variety of cyber-attacks is observed since <the> emergence of COVID-19. Cybercriminals promptly leveraged this pandemic premise to rebrand general attack vectors. These attacks are typically info stealers and vary from attacks of minimal intricacy such as 404 Keylogger to the latest and more frequent attacks such as Lokibot” (Sharma, et al, 2021). The paper explains that the second largest country attacked, “...USA (40%) ...” (Sharma, et al, 2021), is the focus of this case study. The gravity and scope of the impact is multilayered and compounded.

The scope of the credentials is most often extremely broad. Because botnets can extract data from all the victims' browsers, all accounts are compromised, not just one. When a victim realizes their email has been hacked, they may change their password, but if their social media password is different, they may not change all account passwords. To make matters worse, the victim rarely knows their machine is still currently infected. Each time they change the password on the account they know has been compromised and then allows the browser to store that information, the new passwords are captured in an endless loop of re-compromise. Though they are the most predominant payloads deployed, Remote Access Trojans (RAT) serve a diverse set of criminal intentions with equally devastating effects, but the degree impact varies depending on the victim.

5.2. RATs & RESIPs

RATs allow criminal actors to remotely control a victim's computer or mobile device. Unlike info-stealers, RATs allow criminal actors to use the victim's device in any way they would like, exactly as if they were physically using the victim's device. Modern uses of RATs include, but are not limited to, advertising click fraud, credential stuffing, Business Email Compromise (BEC) and Distributed Denial of Service (DDoS) attacks. A paper published in IEEE's (Institute of Electrical and Electronics Engineers) 2020 European Symposium on Security and Privacy Workshops titled Growth and Commoditization of Remote Access Trojans attempts to measure the growth of this type of bot, states "Remote Access Trojans (RAT) are a special type of remote access software commonly used for malicious purposes, where (i) the installation is done without user consent, (ii) the remote control is done secretly, and (iii) the program hides itself in the system to avoid detection" (Valeros and Garcia, 2020). Accurate in the strictest sense, but, due to the profitability of this type of botnet, a new type of RAT has emerged that is offered as a legitimate commercial service. Residential IP Proxy as a Service (RESIP) companies offer the ability to use a consumer's device as a proxy, allowing them to tunnel through the device, assuming its digital identity.

Hola VPN (Virtual Private Network) is a free VPN software/browser that provides users the ability to spoof their location and device type to avoid content access restrictions. The free version comes with a separate software package from Bright Data (formerly Luminati) that hosts the RESIP proxy service on the person's device in lieu of paying for the premium service. Bright Data is one of over a dozen different companies in the RESIP space. The difference between a Remote Access Tool (like Remote Desktop) and a Remote Access Trojan is consent. A white paper called RESIP Host Detection: Identification of Malicious Residential IP Proxy Flows, at the time of its writing in 2021, measured the number of hosts offered by these companies at over 300 million, collectively. "Once a user installs the free HolaVPN, she is recruited as one of Luminati's (Bright Data's) exit nodes" (Tosun et al, 2021). When criminals use these exit nodes for fraud or cybercrimes, the Internet Service Providers (ISP) logs only show that the exit node (the consumer's device) generated the actions of the criminal actor, making it appear the victim is the criminal actor. RATs allow criminals to use other people's digital identities to defraud others. RESIP companies charge by the gigabyte for their victims mobile and home internet usage. Other payloads, such as cryptocurrency miners, have a more direct relationship to the infected victim's device.

5.3. Crypto Jacking

A crypto jacking (cryptocurrency hijacker & miner) solves computational math problems in exchange for cryptocurrency. The explosive increase in value of crypto currency drove the development of malicious deployments of crypto hijacking miners that target high performance computer environments like virtual machines (VM) and cloud computers. A 2017 paper called Mining on Someone Else's Dime: Mitigating Covert Mining Operations in Clouds and Enterprises describes the suitability of crypto miners in

commercial computational environments. “The sheer amount of resources needed for a covert cryptomining operation are readily available in a cloud setting” (Tahir et al, 2017). The costs to the victims are manifested as higher power consumptions, more heat generation as well as the loss of availability of Central Processor Unit (CPU) & Graphics Processor Unit (GPU) resources for legitimate use. Though most profitable in high performance computing environments, crypto jackers are content with exploiting the average consumers devices as well. Often times, crypto jackers run leverage javascript in browsers to launch the covert mining. Ransomware is another botnet payload type that directly impacts the victim's device.

5.4. Ransomware

Ransomware is malicious software which limits or fully prevents a person's access to their infected device (computer, tablet, phone, or other devices) by encrypting the contents of a hard drive. Ransomware also exfiltrates the data it has encrypted to the C2 server. To regain access to the infected device and the information stored on it, the victim is required to pay a ransom. Ransomware is the technologically advanced way to perform traditional crimes like theft or coercion. “In 2013, scareware arrived in the form of Reveton, which locked and prevented access to the infected devices (known as a Locker). After locking the computer (,) the ransomware falsely alleged that the computer (and user) had engaged in unlawful activities and needed to pay a fine to unlock the computer.” (O’Kane, Sezer and Carlins, 2018). The scope and impact of ransomware attacks depends on the victim.

When ransomware infects computer systems in companies, the ransom tends to be much larger and usually based on the company's value. When a botnet detects it is in an enterprise environment, the botnet may try to infect other computers (lateral movement) on the network before activating. Infecting additional computers on the company network allows the ransomware to encrypt and exfiltrate more data. Defending the risks of ransomware in the business environment includes backing up data, limiting access to data, isolating/segmenting systems, anti-malware software, education/training of staff related to malware, phishing tactics, and a dedicated cyber security team monitoring for threats on the network. Ransomware not only encrypts the contents but sends a copy of the files to the ransomware group or actor. Since most businesses use data backup systems, the loss of access to their data is usually limited to the last good backup. This includes information like contracts, emails from high-level executives, market research, intellectual property, and other valuable information. Since most companies have access to their back up data, bad actors shift their tactics to extortion, threatening to publish the information, not just deny access to it. Ransomware has transformed from a platform that denies access to information to an extortion as a service business model. Ransomware is indiscriminate, it impacts individual consumers just as often as large companies, though the impact is quite different.

On personal devices, this equates to the loss of all data like personal pictures and stored documents. The use of reputable anti-virus software is one level of protection against ransomware, another protective measure is to backup data to other locations. This could be one or two hard drives that are used solely for the purpose of data backups that are not connected to the internet. While this will help in recovering some of the information, all usernames, passwords, all credit cards, and all form information (like ship to/bill to addresses) that are retained by browsers are exposed to the criminal actor. All these various payloads are significant threats to all connected devices. Protection should always be applied in layers. In addition to having industry recognized anti-virus protection, there are some obvious signs of the dangers that lie ahead.

6. Indicators of Risk

Whether using a mobile device, personal computer, or web connected television, downloading, and opening any attachments should only be done when the sender is known and trusted. Malicious actors

depend on the victim's curiosity being piqued or reaction to being threatened in some way. Though filters in email and antivirus programs identify some malicious attachments, often many variants make it through with the victim's assistance. Fake URLs and links that make it through automated protection can often be identified with a closer examination.

A URL starts with a protocol, like HTTP or HTTPS. As the name indicates, entering any information into a website without the “S” is unsecured and readable to an attacker. The “S” should also be accompanied by a locked padlock symbol in the address bar. Following the protocol and the forward slashes, the domains are defined. Consider the example “https://subdomain.domain.top level domain (TLD)”. Most legitimate websites are in the .com, .org and .gov TLD’s. Unrecognized TLD’s should always bear scrutiny. The domain, along with the TLD are the most important part of the URL. Any typos, missing letters or characters, or unfamiliar variations of domain names should also inspire closer scrutiny. The subdomain is malleable and should be considered informative but not authoritative. For instance, Google.com is vastly different than Google.hackexample.com. Protecting devices and systems requires both behavioral effort as well as reputable automated solutions.

Antivirus programs are not foolproof but, as the adage goes, you get what you pay for. Free antivirus programs are often viruses in disguise. All companies monetize their products, the RESIP example demonstrates the unseen price a victim pays for a free product. Applying the same level of protection across all devices, like mobile and tablets that connect to the internet, especially those that are used to access secure websites ensures having a baseline level of automated protection of those devices. Anti-virus, anti-malware and automated solutions can provide some protection, but exercising prudence minimizes the reliance on automated solutions. Before downloading an attachment or clicking a link, be certain the source is someone or some company that is known and trusted. Double check the spelling of the sending email address and the contents of the email. Often, foreign malicious actors use broken English. Hover over links in email to examine underlying hyperlinks. Free products are often paid for in ways that are unknown. Game and application cracks and cheats often contain info stealer payloads. Protect mobile devices with antivirus programs as diligently as a personal computer. Use of MFA, passwordless authentication and password managers minimize account password exposures. Never allow web browsers to store passwords or credit card information. Info stealers can extract them all at once. Never reply to an unknown sender of unsolicited SMS messages. Phishing attacks often use the pretext of an offer that is limited time only, time is critical. They also use previous data breaches and open-source intelligence to gather information on the victim. Never share SMS text codes with anyone else.

7. Conclusions

The rise of the volume of botnets is only dwarfed when compared to the risks that the botnet payloads create. The persistent exposure of all account and credit card data stored or saved in browsers that info stealers access ensures continual re-compromise until the infection is eradicated. RATs can impersonate a victims’ network identity and be wielded as a massive network attack weapon. Ransomware has evolved to include extortion, not just denying access to the victims’ data but publishing it as well. Crypto jackers exploit computer resources rather than information so the cost to the victim often goes unnoticed. Free products come with costs that are not as direct as paying with money. In a digital world, every character is important. Reputable automated protection solutions are only one layer, applied prudence and critical assessment of all digital communications and transactions is still the primary line of defense.

Abbreviations

C2	Command and Control
CPU	Central Processor Unit
DCA	Digital Citizens Alliance
GPU	Graphics Processor Unit
HTML	Hypertext Markup Language
HTTP/HTTPS	Hypertext Transfer Protocol/Hypertext Transfer Protocol Secure
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
IP	Internet Protocol
IRC	Internet Relay Chat
ISP	Internet Service Provider
MIT	Massachusetts Institute of Technology
MFA	Multifactor Authentication
P2P	Peer to Peer
PDF	Portable Document Format
RAT	Remote Access Tool or Trojan
RESIP	Residential Internet Protocol Proxy as a Service
SMS	Short Message Service
SMTP	Simple Mail Transfer Protocol
TAG	Trustworthy Accountability Group
TLD	Top Level Domain
TOR	The Onion Router
URL	Uniform Resource Locator
VPN	Virtual Private Network

Bibliography & References

McCorduck, P. (1979). *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. San Francisco, CA: W.H. Freeman, pp. 293.

N/A, (2021). Barracuda research reveals skyrocketing levels of bot traffic. Retrieved from:
<https://www.barracuda.com/news/article/833#.YjHmrnrMliuU>

Sharma, R., Sharma, N., and Mangla, M. (2021). An Analysis and Investigation of InfoStealers Attacks during COVID'19: A Case Study, 2021 2nd International Conference on Secure Cyber Computing and Communications (ICSCCC), pp. 1.

Valeros, V. and Garcia, S. (2020). Growth and Commoditization of Remote Access Trojans, 2020 IEEE (Institute of Electrical and Electronics Engineers) European Symposium on Security and Privacy Workshops, pp. 1.

Tosun, A., De Donno, M., Dragoni, N., & Fafoutis, X. (2021). RESIP Host Detection: Identification of Malicious Residential IP Proxy Flows. IEEE (Institute of Electrical and Electronics Engineers) International Conference on Consumer Electronics, p. 4.

Tahir, R. *et al.* (2017). Mining on Someone Else's Dime: Mitigating Covert Mining Operations in Clouds and Enterprises. In: Dacier, M., Bailey, M., Polychronakis, M., Antonakakis, M. (eds) Research in Attacks, Intrusions, and Defenses. RAID 2017. Lecture Notes in Computer Science, vol 10453. Springer, Cham.

Larson, S and Mesa, M (May 26, 2021). BazaFlix: BazaLoader Fakes Movie Streaming Service. Proofpoint Threat Insight Blog

Osagie, M., Enagbonma, O. and Inyang, A. (2019). The Historical Perspective of Botnet Tools. Current Journal of Applied Science and Technology, Benson University, Nigeria. p. 7.

O'Kane, P., Sezer, S., and Carlin D. (2018). Evolution of Ransomware. The Institution of Engineering and Technology, UK. p.3.

Pemberton, C. (2016). Tap Into the Marketing Power of SMS, Gartner Research. Retrieved from <https://www.gartner.com/en/marketing/insights/articles/tap-into-the-marketing-power-of-sms>

N/A, 2021. Breaking B(ads): How Advertiser-Supported Piracy Helps Fuel A Booming Multi-Billion Dollar Illegal Market, Digital Citizens Alliance. Retrieved from: <https://www.digitalcitizensalliance.org/clientuploads/directory/Reports/Reports/Breaking-Bads-Report.pdf>

Impacts of Legacy and Next Generation Cooling Technologies on Power Demands and Environmental, Social and Governance Strategies

A Technical Paper prepared for SCTE by

John Dolan

Senior Guideline Specialist
Rogers Communications Canada
8200 Dixie Road
Brampton, ON L6T 0C1
519-852-5666
john.dolan@rci.rogers.com

Arnold Murphy, President, SCTi

Michael Glaser, Principal Engineer, Cox Communications

Allison Richards, Sr. Energy Analyst, Charter Communications

Ken Nickel, Quest Controls, Inc.

John Teague, Mechanical Engineer, Worldwide Environmental Services

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Cooling Technology and GHG	5
2.1. Refrigerants.....	5
2.2. Cable Critical Facilities Electricity – Cooling Systems	5
2.3. Legacy Cooling Systems.....	5
2.4. Cooling Systems Today	6
3. Improvements to Reduce Energy Consumption in Legacy/Existing Facilities and Cooling Systems.....	6
3.1. Floating Head CRAC Retrofit.....	7
3.2. Room or Perimeter Cooling.....	9
3.3. In-Row Cooling.....	9
3.4. In-Rack or Close Coupled Cooling.....	10
3.5. Investigating bulk rack exhaust fans vs server muffin fans for rack containment.....	12
3.6. Why not Switch To CO ₂ (R744) Refrigerant Based Cooling Units?.....	12
3.7. Monitoring and Control Systems.....	13
3.7.1. Background	13
3.7.2. Current Monitoring and Control Technologies	13
3.7.3. Data Aggregation	14
3.7.4. Data Analytics	14
3.7.5. Next Gen Controller Technologies.....	14
3.7.6. Tie to Carbon	16
4. Next Generation	16
4.1. Next Generation Refrigerants	16
4.2. Next Generation Cooling Technologies	17
4.2.1. Phase Change Material (PCM).....	17
4.2.2. Liquid Cooling	17
4.2.3. Internal Server Cooling	18
5. Conclusion.....	18
Abbreviations	18
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1- Footprint of a Typical Cable Operator (from SCTE 208 2021).....	5
Figure 2 - ECM Estimated Savings.....	7
Figure 3 – Map of Temperatures for selecting Floating Head as an ECM	8
Figure 4 - Example of in-row cooling where air conditioning units fit inside rack rows. Courtesy: PRWeb	10
Figure 5 - Rear Door Cooling (Diagram credit: OptiCool).....	11
Figure 6 - Air Diverter System (Courtesy of Chillirack).....	12
Figure 7 - Economizer Setpoint Change over Time.....	15
Figure 8 - Control System Communicates to Smart Equipment and Smart T-Stats Connected to Legacy Equipment	16
Figure 9 - Phase Change Material and Free Air Cooling (Diagram Courtesy of Energy Cool [®])	17

Figure 10 - GWP Life Cycle Climate Performance 18

List of Tables

Title	Page Number
Table 1 - Floating Head CO ₂ Reduction	9
Table 2 - R744 Advantages and Disadvantages	13

1. Introduction

Environmental, Social and Governance (ESG) is a framework within the umbrella of sustainability. It is used to describe and measure an entity's behavior with respect to environmental issues, which includes greenhouse gas emissions and impact on natural resources, its engagement with and impact on society – both local and global, and the strength of its governance and the ethics behind its decision and policy making.

Companies around the world are being pressed to provide more transparency around their ESG risks and meaningful progress toward the mitigation of those risks.

ESG matters for a number of reasons, not least because the criteria are a set of standards that potential investors use to screen and evaluate companies. When a company's ESG "score" (a measure of a company's exposure to long-term environmental, social, and governance risks) goes up, its capital costs are reduced, and the company valuation improves. It is like a sustainability credit rating.

A good sustainability program drives business growth and enhances the brand, while cutting costs and reducing risk.

The most important reasons for focusing on ESG, however, are the long-term impacts on society and the planet.

Energy use is a key component of the Environmental element of ESG. For Cable Operators energy use for critical facilities operation represents one of, if not the largest operating expense. Heating, Ventilation and Air Conditioning (HVAC) systems typically account for 35% to 45% of the cable facility energy use. The Information Technology Equipment (ITE) heat loads represent approximately 45% to 50% and miscellaneous lighting is +/- 5%. By achieving reductions in energy use, related Greenhouse Gas (GHG) emissions can be reduced resulting in an improved ESG score.

The Environmental part of ESG is categorized in three scopes.

- Scope 1 - direct GHG emissions occurring from sources that are controlled by an organization - eg emissions associated with furnaces or vehicles
- Scope 2 – indirect GHG emissions are due to the organizations use of electricity
- Scope 3 – are indirect emissions occurring in the supply chain.

This paper will explore how legacy sites with older cooling systems can achieve significant energy and carbon footprint reduction and will also step into new technologies that will lead to further energy efficiency improvements.

Scope 2 emissions will be the primary focus of this paper and more specifically emissions related to cooling in legacy facilities. Reducing GHG is best done within a management lifecycle process and must include a focus on all three emissions scopes. Meaningful measurement needs to be performed, where baselines and targets are set, and progress is tracked.

Due to variations between operators' approach to preventive maintenance (ie. internal or external business partners) and operators not having readily available emissions data for purchased goods, Scopes 1 & 3 impacts are not considered in this paper. See also Note 1.¹

¹ Note 1: This paper represents the opinions of the authors and is the product of professional research. It is not meant to represent the position or opinions of Rogers Communications, and Rogers Communications does not accept any responsibility or liability for the accuracy, content, completeness, legality, or reliability of the information contained in this paper. The information contained in this paper is not intended to be relied upon for any specific application without independent verification and assessment of suitability.

2. Cooling Technology and GHG

Guidance in collecting and managing data to determine GHG and carbon footprint is found in SCTE 208 2021 – Cable Operator Greenhouse Gas Emissions Data Collection Recommended Practices.

As can be seen from Figure 1- Footprint of a Typical Cable Operator (from SCTE 208 2021) Cooling Technology consists of *Scope 1 (Direct Emissions)*, Refrigerants, 4.2%, and *Scope 2 (Indirect Electricity Emissions)*, Cable Critical Facilities Electricity of 22.2%.




Figure 1- Footprint of a Typical Cable Operator (from SCTE 208 2021)

2.1. Refrigerants

For Refrigerants, *Scope 1 (Direct Emissions)*, the Environmental Protection Agency (EPA) provides a method to estimate emissions from: Installation, Operation and Disposal to determine total emissions in this category. This is covered in detail by the EPA in: *Greenhouse Gas Inventory Guidance: Direct Fugitive Emissions from Refrigeration, Air Conditioning, Fire Suppression, and Industrial Gases*.

Refrigerants used in HVAC systems affects their efficiency and therefore can have an effect on *Scope 2 (Indirect Electricity Emissions)*.

2.2. Cable Critical Facilities Electricity – Cooling Systems

Cable Operators Critical sites commonly consist of a combination of older and newer cooling systems.

Cooling systems that are over 10 years old are generally considered “Legacy” Cooling Systems with higher electrical consumption hence higher emissions.

Legacy Cooling Systems have a number of limitations which result in contributing significantly to the Scope 2 reported emissions of 22.2%.

2.3. Legacy Cooling Systems

Legacy HVAC systems operate at a fixed capacity, or at best in a rudimentary step function to provide some variable capacity, and therefore have a high mechanical Power Usage Effectiveness (PUE) for a

given heat load. Energy saving features such as: Free Air Cooling; Pumped Refrigerant, were not available.

If more than one unit is used to cool a given heat load, quite often they operate independently which means they may not share the heat load causing one unit to operate at a much higher cooling level than another. In some cases, networking, sequencing and ‘teamworking’ the units to operate together, if applied properly, can improve their combined efficiency.

The fans in the evaporator and condenser are fixed speed and so have a fixed energy consumption, regardless of heat load.

Monitoring capability is rudimentary and does not provide any significant data on “real time” energy consumption to enable effective energy management.

Air cooled HVAC installed in Critical Facilities prior to 2010 used the refrigerant R-22. Although the GHG (Measured by Global Warming Potential (GWP)) of R-22 was similar to refrigerants used today, for example R-407C or R410A, they are not chlorinated substances which are classified as an Ozone Depleting Substance (ODP). That makes R-22, a chlorinated refrigerant, a much less desirable refrigerant. New R-22 refrigerant is no longer available which means cooling systems using R-22 are reliant on recycled refrigerant if top ups are required. The supply of R-22 will diminish over the next few years resulting in higher costs and will eventually not be available.

2.4. Cooling Systems Today

HVAC systems in use today have variable cooling capacity, using scroll type compressors or the equivalent to gain over 2000 points to adjust the evaporator. Evaporators have a higher capacity and Electronically Commutated (EC) capability enable variable speed fans to be used on evaporators and condenser. Pumped Refrigerant, or direct free air cooling are also available to reduce energy consumption. This has allowed a significant energy reduction for a given heat load.

Controls and monitoring are much more sophisticated but represent a training challenge for operations teams. There is room for significant improvement both in the operation of units and in right sizing for a given heat load.

Air Cooled units now use more environmentally friendly refrigerants such as R407C or R410A which have zero ODP and a reasonable GHG/GWP.

3. Improvements to Reduce Energy Consumption in Legacy/Existing Facilities and Cooling Systems

Older critical facilities suffer from poor cooling efficiency due to older cooling systems being used and because the sites were not designed for the heat load levels being experienced today. As a result, the racks and cooling designs were not well laid out for effective heat removal. This resulted in poor air flow management requiring excess cooling and air flow to meet the growing heat load demands.

The SCTE Facilities Cooling Technology Optimization Working Group has addressed methods to improve and optimize cooling in legacy facilities in the papers highlighted below.

- Improving air flow management – resulting in less air flow required and higher return air temps – SCTE Journal SCTE-EM-V5N1 – Rightsizing Network Cooling – Getting Ready for 10G
- Increase in set points – SCTE 253 2019 Cable Technical Facility Climate Optimization Operational Practice: Understanding Set Point Values, Part 1
- Air containment – SCTE 274 2021 Cable Operator Critical Facility Air Containment Operational Practice
- Control Systems and Networking cooling units Reference SCTE 184 SCTE Energy Management Design, Construction and Operational Practices for Cable Facilities
- Operational Practice to improve air flow and climate conditions in Critical Facilities – SCTE 219 2021 – Technical Facility Climate Optimization Methodology

Pilots and research trials by multiple Cable Operators have measured and verified the possible energy savings potential of the listed Energy Conservation Measures (ECMs) as shown in Figure 2 - ECM Estimated Savings. The scalability and magnitude of these measures will be unique to the site and providers as they identify, through assessments, where their baselines are, and what would have the most impact to their sites.

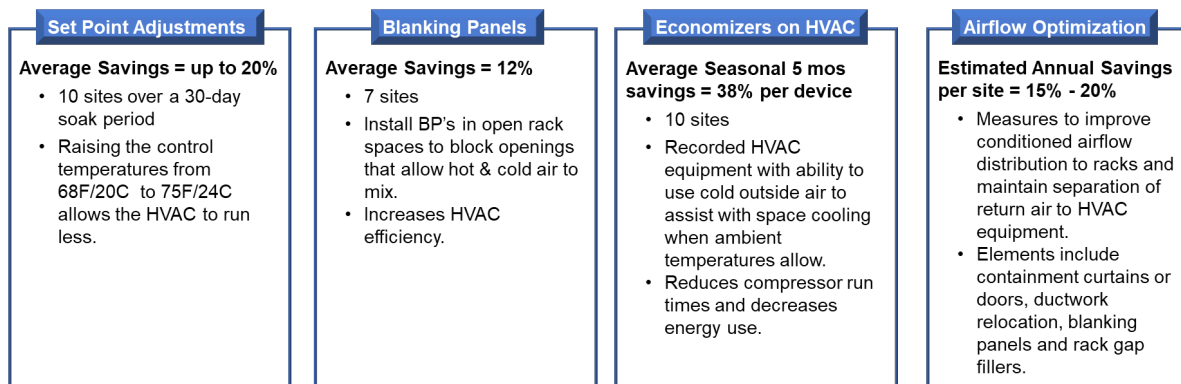


Figure 2 - ECM Estimated Savings

As the savings above indicate, several low-cost implementations can result in measurable energy reductions. Additionally, the replacement of existing aged HVAC units can result in immediate efficiency gains resulting in improvement of Scope 2 by 20% to 30%.

3.1. Floating Head CRAC Retrofit

Floating head retrofit is applicable to Computer Room Air Conditioning (CRAC) units that are Direct Expansion (DX) based and that are 5 years or older. This retrofit will improve energy efficiency and reduce carbon footprint. The mechanical expansion valves in these older units are replaced with electronic expansion valves, enabling the unit to adapt to ambient temperatures and regulate condensing temperatures accordingly.

Traditional fixed head pressure systems run at a condensing temperature of 40°C (105°F). The energy savings for retrofitted low condensing systems begins to accrue whenever the ambient temperature is

below 24°C (75°F). Figure 3 below shows the percentage of time the respective area is below 24°C (75°F) geographically.

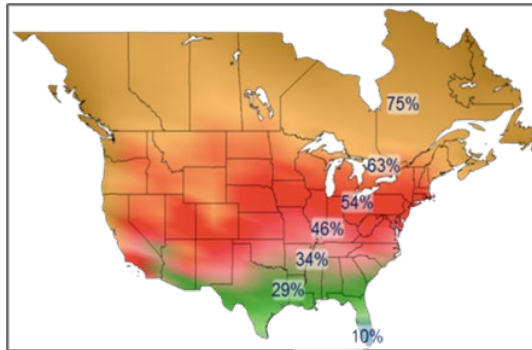


Figure 3 – Map of Temperatures for selecting Floating Head as an ECM

(SCTE Journal of Energy Management 2016 VINI, New A/C System Architecture Promises Significant ROI in Data Centers, Floating Head Pressure Technology Reduces Energy Costs and Consumption)

Temperatures below 10°C (50°F) ambient represents the level at which the maximum energy savings from low condensing operation can be achieved as the compressor wattage decreases. A small increase in cooling capacity also occurs.

Energy consumption can be reduced by up to 45%, and the life expectancy of the unit will be extended as there is less wear on the main components.

A floating head retrofit won't make the legacy unit as energy efficient as the newest economizer models. However, comparing the cost and short payback, typically 2 years or less, of this retrofit, with the capital cost and operational disruption caused by replacing a cooling unit with an economizer system this option is very attractive. The retrofits do qualify for energy incentives if available.

In addition to energy savings two other benefits are captured with the retrofit. If the unit is using R22 refrigerant, that is units produced before 2010, this can be replaced and upgraded to R407C. The second significant benefit is less refrigerant is required for the retrofitted unit. Approximately 35% of the refrigerant is removed. The amount of refrigerant removed is dependent on the cooling capacity of the unit and the physical layout, distance from condenser.

The table below shows a case of retrofitting two Liebert DH-290's and one Liebert DS105. In this example the R22 refrigerant in the DH-290's was removed and appropriately disposed and replaced with R-407C. CO₂ equivalence was calculated using the EPA Greenhouse Gas Equivalencies Calculator.

Table 1 - Floating Head CO₂ Reduction

Floating Head Retrofit - example			
Cooling Units	2 - Liebert DH290; 1 - Liebert DS105		
	Pre-Retrofit	Post-Retrofit	Reduction
Cooling Power kW	67	34	33
Cooling Energy (kWh)	585,416	299,201	286,215
Energy Cost (\$0.12/kWh)			\$34,346
CO ₂ Equivalence			
Refrigerant Removal (kg)			211.47
Refrigerant Removal (lbs)			466.2
CO ₂ Equivalence Reduction			Metric Tons
	Due to kWh avoided		203
	Due to refrigerant removal		383
	Total		586

3.2. Room or Perimeter Cooling

Room scale cooling is where interior perimeter CRAC or Computer Room Air Handlers (CRAH) are used to remove heat from the space and provide conditioned air into the room. These systems are spaced around the exterior and in larger rooms near the center. This system typically provides bulk conditioned air into the room via raised floors, direct air discharge from the unit or through overhead duct work on a slab floor. Because of this bulk delivery approach several inefficiencies are evident due to hot and cold air mixing, by-pass airflow, high volume over provisioned air, and older fixed fan speed and staged compressor speeds. ECMs can be applied to improve efficiency.

3.3. In-Row Cooling

As rack heat loads increase a more efficient cooling method must be implemented to manage high density heat zones. In-row cooling is the deployment of multiple smaller units selectively placed between the cabinets within the rows of server racks. With the assistance of containment doors, curtains and sometimes row caps, a more controlled environment can be created. Having available space for the added in-row devices can be an obstacle for a retrofit application in legacy sites. As the density within the server racks increase so too will the added requirement for redundancy.



Figure 4 - Example of in-row cooling where air conditioning units fit inside rack rows.
Courtesy: PRWeb

3.4. In-Rack or Close Coupled Cooling

In-rack or close coupled cooling places the cooling much closer to the source of heat . The extremely short distance between the server exhaust and rear door coils is measured in inches vs feet for the room and row applications thus eliminating the opportunity for mixing and other efficiency losses. .

The rejection of heat is accomplished with pumped refrigerant or chilled water running through coils at the rear doors, to an in-room heat exchanger using outside condensers. The exterior components have either economizer or mechanical cooling to complete the cooling cycle. The refrigerant pump within the main heat exchange chassis can also be DC powered. The rear door heat exchangers, as shown below, can handle 10 kilowatt (kW) per fan pair and have redundancy when all 3 pairs are installed.

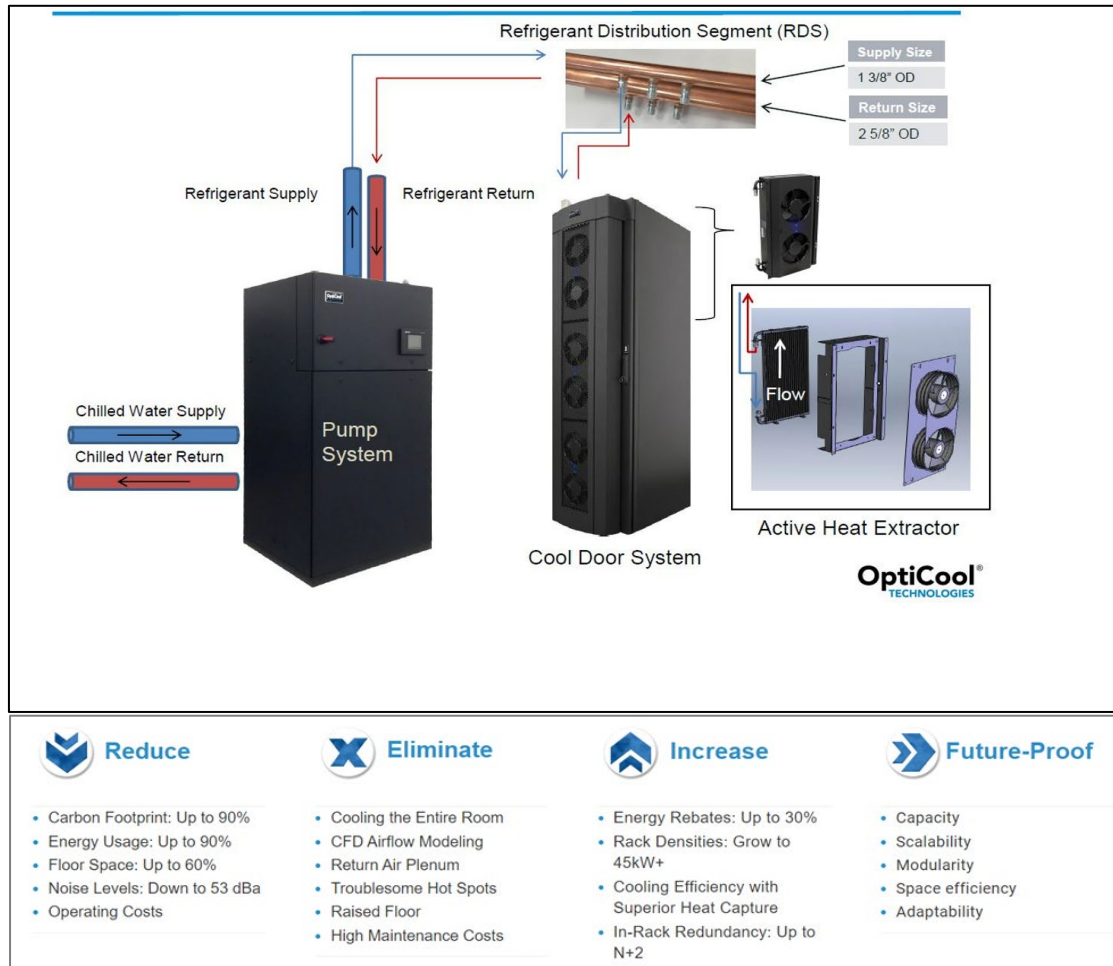


Figure 5 - Rear Door Cooling (Diagram credit: OptiCool)

Additionally, these systems require a minimal load working best on a high kW density (typically 10 kW to 30 kW per rack) applications. Overhead space is required for manifold piping and each rack depth is extended approximately 12-14 inches for the rear doors.

3.5. Investigating bulk rack exhaust fans vs server muffin fans for rack containment

Another innovative In-Rack cooling approach being developed is to utilize an existing raised floor environment or overhead ducted supply system to bring the cooling more efficiently where it is needed. This approach uses a clear solid door on the racks with a separate supply and return ducting with dual in-line fans reducing the need and cost of the server fans. Additional in-rack airflow manifolds have been developed to direct airflow for equipment intake variations. As shown, the doors would be solid and individual servers would be fed and exhausted via side discharge.




Figure 6 - Air Diverter System (Courtesy of Chillirack)

3.6. Why not Switch To CO₂ (R744) Refrigerant Based Cooling Units?

High pressure Carbon Dioxide, CO₂ or R744 systems were developed at the end of the 19th century. It has an GWP of 1 so it would seem to be the ideal refrigerant. There are real advantages to using R744 but also some drawbacks holding it back from general adoption in the industry.

The following table summarizes some of the advantages and disadvantages of R744.

Table 2 - R744 Advantages and Disadvantages

Advantages	Disadvantages
Refrigeration capacity is high, approximately 5 times that of R404A (smaller compressor displacement but the same motor size)	Operating and standing pressures are high with resulting higher leak potential. Components specific to R744 required
Pressure drops in pipes allowing longer lines	Compressors are R744 specific, steel, or stainless-steel welded fittings required
Evaporators and condensers have high heat transfer	System complexity resulting in higher costs
System pressure drops, compression ratios all result in higher system efficiency	Complexity can result in poor performance and reliability because commissioning and operation need to be done correctly
Thermosyphoning of the refrigerant is possible reducing compressor run time.	Not suitable for high ambient temperature areas
R744 is inexpensive to manufacture, has low toxicity so release, often just by venting, or leakage is not an issue as there are no disposal restrictions	Listed as an asphyxiant so leak detection is required in small, enclosed spaces and release needs to be in well ventilated areas.

As can be seen from Table 2 - R744 Advantages and Disadvantages although there are significant advantages to R744, the disadvantages, particularly the complexity that operational staff will need to deal with, and the specialized materials, have resulted in a low adoption rate of these systems.

3.7. Monitoring and Control Systems

3.7.1. Background

Mechanical thermostat controls are the most traditional form of control for HVAC systems that can still be found today in headend and hub locations performing basic on/off control of the HVAC systems. Unfortunately, these types of controls do not lend themselves to modern facilities and certainly do not provide the necessary data needed to monitor and control HVAC systems to ensure optimum performance. Fortunately, there are options available today to quickly and reliably upgrade existing HVAC to smart thermostats that network into a cohesive system.

The Building Management System (BMS) have been a powerful tool to help critical facility managers monitor and control key facility infrastructure support systems. These typically include main facility electrical and mechanical systems such as: electrical power distribution systems, central air conditioning and ventilation, fire protection and lighting.

BMS systems will tend to be a localized system consisting of a software and hardware. Older BMS systems however typically do not offer a more granular control for specific critical support equipment necessary for today's demands from critical processing servers.

3.7.2. Current Monitoring and Control Technologies

Critical facility management relating to energy consumption, costs, downtime as well as its carbon footprint contribution can be challenging to most organizations without the proper tools. For this reason, monitoring and control systems play a crucial role in assisting facility management. Although current systems such as BMS, have been a valuable tool for facility management, these systems are becoming dated and lack sophistication for optimizing energy consumption management. These existing systems, albeit provide critical information to the operator, and primarily serve as a historical data archive and

alarming system. These systems serve a limited capability with respect to data collection and usage as well as critical metrics and features for control optimization. As the industry looks towards updated systems offering enhanced tools and features, current systems employed are pushed to their limits of operation and obsolescence.

Systems are becoming more intelligent, providing access to significant amounts of data from their sensors regarding the performance of the equipment.

3.7.3. Data Aggregation

Most control and monitoring systems today have some form of a local user interface to view trend data such as room temperature and the ability to make adjustments. While this is necessary at a site level, it is important when managing multiple sites to have a central server-based monitoring software that can gather trending and alarm data from all of your facilities. Modern monitoring and control systems offer comprehensive monitoring and data acquisition and data storage capabilities. Many of these systems are now capable of integrating services and solutions across multi-vendor platforms for data centers and other complex mission critical facilities.

These systems have the capability to aggregate data from separate and disparate systems to provide a useful database for system comparisons, interactive intelligence and enhanced control decisions for energy savings, optimization, and efficiency. Incorporating data from ITE Assets and existing BMS systems into a single cohesive platform will provide the most comprehensive intelligence for monitoring critical facilities.

This provides the ability to identify outlier/trouble sites and prioritize maintenance resources.

3.7.4. Data Analytics

Central monitoring software is being used to provide operators with information that can predict operation of equipment and serve early warnings to impending failures which improves network reliability. These tools provide true real time situational awareness and insight to present operators with critical information to address system availability across a single facility or facilities throughout the enterprise. This approach integrates advanced energy analytics, asset management and facility monitoring tools with advanced analytics to increase visibility, improve overall efficiency, reduce costs, and mitigate risk.

This ability to create actionable data can be used to further optimize operation to reduce energy consumption of equipment and provide verification of carbon reduction. For example, the software can analyze excessive runtime with HVAC equipment at all sites. It can include frequency of ON / OFF occurrence and shortest duration between ON / OFF (To indicate unhealthy short cycling). The report can show historical performance, live data overlays, suggested targets, and remediation. Another example, the central monitoring software can poll the local control systems to gather the current cooling operation and then generate a report indicating zones where temperature setpoint is not maintained when cooling is on. Time between triggered cooling event and time to satisfy zone temperature are tracked to indicate degradation in cooling effectiveness.

3.7.5. Next Gen Controller Technologies

Control systems can use machine learning for adaptive control strategies. This is leading to control systems that can automatically adjust their settings to changes in the facility and better matching of the heat load. An example of this is the ability to adjust the economizer control setpoints based on successful cooling using outside air. Traditionally the economizer settings are fixed thresholds for when outside air can be used. Intelligent control systems can learn what outside air temperatures will provide necessary

Economizer Setpoint Change over Time
March 9th to April 9th

The graph displays the Economizer Setpoint (L1) over time. The y-axis represents the setpoint value, ranging from 55 to 73 in increments of 2. The x-axis represents time, from March 9th to April 9th. The setpoint starts at 65, drops to 61, then fluctuates between 61 and 67 before rising to 71, peaking at 72, and ending at 70.

Date	Economizer Setpoint (L1)
March 9th	65
March 10th	65
March 11th	61
March 12th	61
March 13th	61
March 14th	62
March 15th	63
March 16th	64
March 17th	65
March 18th	66
March 19th	67
March 20th	65
March 21st	65
March 22nd	66
March 23rd	67
March 24th	67
March 25th	67
March 26th	69
March 27th	69
March 28th	70
March 29th	70
March 30th	70
March 31st	70
April 1st	72
April 2nd	71
April 3rd	71
April 4th	71
April 5th	71
April 6th	71
April 7th	72
April 8th	72
April 9th	70

For example: A site in Denver, CO with the ability to float up the economizer setting to 55 °F (13°C) compared to a fixed economizer setting of 50 °F (10°C) and would see a 35% increase in the available opportunity to use economization.²

² Based on degree day calculations for a year using <https://www.degreeddays.net/>. Calculations used weather stations ID KDEN




Figure 8 - Control System Communicates to Smart Equipment and Smart T-Stats Connected to Legacy Equipment

3.7.6. Tie to Carbon

With the advancements of control and monitoring systems, organizations are able to realize significant reductions in direct energy usage and carbon footprint. Monitoring and control systems for today's facilities are at their best when they can take dissimilar or standalone intelligent systems and coordinate their control to ensure the facility is sufficiently cooled for their intended use. Intelligent controls are able to take advantage of alternate cooling methods such as outdoor economization to further reduces compressor run time. Charter Communications, Inc. reported in 2021 that "evaporative free-cooling avoided over 5,700 megawatt-hours of electricity usage"³. This translates to a CO₂ reduction of approximately 4,039 Metric Tons⁴. Another MSO has recently completed a free air economization proof of viability at a Northeast facility where a reduction in compressor run time provided a savings of approximately 19,162 kilowatt-hours (kWh) during a five-month period available for economization use. This will yield a CO₂ reduction of 13.6 metric tons.

The ability to remotely monitor the performance of a portfolio of facilities and make necessary changes prevents unnecessary travel to sites which reduces carbon emissions and saves on fuel expense. Through the advancements of the Internet of Things (IoT), additional sensing is economically being deployed to remote facilities. The ability to gather information and turn it into actionable data will further increase an operator's ability to document, categorize and manage their carbon footprint.

4. Next Generation

4.1. Next Generation Refrigerants

Refrigerants available are being updated to reduce their GWP. One of the drivers for this is the Kigali Amendment to the Montreal Protocol. This primarily affects R22 for Cable Operators at this time. Refrigerant properties are available in ASHRAE Standard 34 Designation and Safety Classification of

³ Charter-2021-ESG-Report.pdf, <https://corporate.charter.com/esg-report>

⁴ Carbon reduction calculated using the EPA calculator. <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator>

Refrigerants which provides refrigerant numbering, toxicity, and flammability ratings. The pursuit of other refrigerants may take place over time and depends on low pressure, medium pressure, or high-pressure systems. The best choice today may change tomorrow. Although complex, the key decision for selecting a new refrigerant needs to be made on efficiency.

Nonetheless for GWP the HVAC system has a direct effect from the chemical refrigerant, and indirect effect from the energy use adding up to the total climate performance.

The biggest factor that can be controlled by the Cable Operator is then the indirect effect from energy use and the associated GWP of that energy source.

4.2. Next Generation Cooling Technologies

4.2.1. Phase Change Material (PCM)

The effectiveness of cooling systems that use outside air (free air-cooling systems), can be improved considerably with the use of (PCM). This can cool the incoming air reducing compressor run time, extending the free air-cooling time, and be recharged at night when the temperature is lower.




Figure 9 - Phase Change Material and Free Air Cooling (Diagram Courtesy of Energy Cool[®])

4.2.2. Liquid Cooling

As ITE heat loads increase, the effectiveness of air-cooling systems will reach a limit with reduced efficiency. In order to cool the higher heat loads and drive PUE down below 1.1, closer coupling of heat removal is required using some form of liquid cooling. This paper will not go into detail of the technologies that are available, but they are:

- Liquid cooling - A cabinet or plate to remove heat directly from the surface of the chips used by the ITE heat load.
- Immersion cooling - Removes all fans and submerges the ITE heat load directly in a dielectric cooling liquid.

Currently both are used for higher density computing above 50 kW per cabinet.

4.2.3. Internal Server Cooling

Internal server cooling may improve using hybrid two phase cooling using passive low height thermosyphons to dissipate the heat generated in high power components of servers. This augments traditional cooling of other server components.

5. Conclusion

Through the careful use of ECMs, and selection of cooling technologies that are appropriate, Critical Facility energy use for cooling can be significantly reduced as the ITE heat load increases. This in turn will help to reduce Scope 2 emissions and overall GHG impact.

Additionally the use of more environmentally friendly refrigerants can help to reduce GHG's. Refrigerants are being updated to reduce their GWP but must be selected carefully as this can in turn affect the HVAC performance.

It can be seen in Figure 10 - GWP Life Cycle Climate Performance, that the major climate impact of HVAC is the electrical power used, and how it is generated, over the equipment lifetime. The key factor in refrigerant selection then is its effect on the efficiency in the HVAC system. A refrigerant that results in overall lower HVAC cooling effectiveness per kW of heat load will result in a higher Scope 2 emission for the system as it will consume more electrical power.

With the source of the electrical power being a key factor in the critical facility life cycle, climate performance consideration must be given to reduce that impact through use of microgrids, Demand Energy Response (DER) and other solutions for greener energy sources.

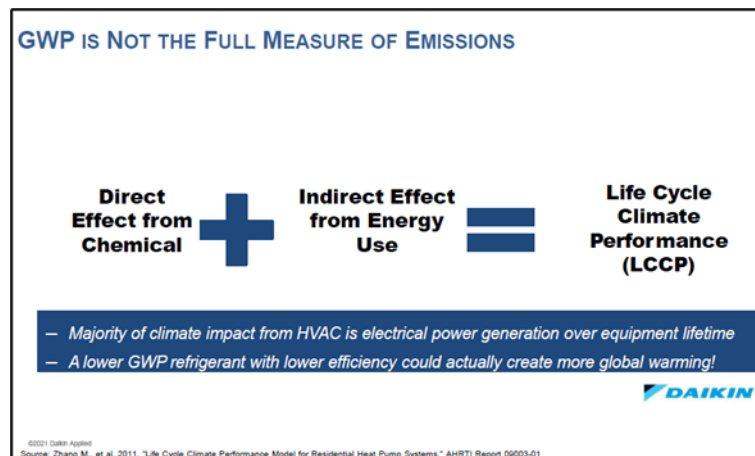


Figure 10 - GWP Life Cycle Climate Performance

Abbreviations

BMS	Building Management System
CRAC	Computer Room Air Conditioner
CRAH	Computer Room Air Handler
DER	Demand Energy Response
DX	Direct Expansion
EC	Electronically Commutated

ECM	Energy Conservation Method
EPA	Environmental Protection Agency
ESG	Environmental, Social and Governance
GHG	Green House Gas
GWP	Global Warming Potential
HVAC	Heating, Ventilation and Air Conditioning
ITE	Information Technology Equipment
ODP	Ozone Depleting Potential
kW	kilowatt
kWh	kilowatt-hours
PCM	Phase Change Material
PLC	Programmable Logic Controller
PUE	Power Usage Effectiveness

Bibliography & References

ASHRAE Standard 34 Designation and Safety Classification of Refrigerants, available at:
<https://www.ashrae.org/technical-resources/bookstore/ashrae-refrigeration-resources>

EPA Greenhouse Gas Equivalencies Calculator - <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator>

Greenhouse Gas Inventory Guidance: Direct Fugitive Emissions from Refrigeration, Air Conditioning, Fire Suppression, and Industrial Gases - <https://www.epa.gov/sites/default/files/2015-07/documents/fugitiveemissions.pdf>

SCTE Journal of Energy Management 2016 V1N1 - New A/C System Architecture Promises Significant ROI in Data Centers, Floating Head Pressure Technology Reduces Energy Costs and Consumption - https://wagtail-prod-storage.s3.amazonaws.com/documents/SCTE-ISBE-EM_Journal_V1N1.pdf

SCTE Journal SCTE-EM-V5N1 – Rightsizing Network Cooling – Getting Ready for 10G - https://wagtail-prod-storage.s3.amazonaws.com/documents/SCTE-ISBE-EM_Journal_V5N1.pdf

SCTE 184 SCTE Energy Management Design, Construction and Operational Practices for Cable Facilities - <https://www.scte.org/standards/library/catalog/scte-184-cable-facilities/>

SCTE 208 2021 – Cable Operators Greenhouse Gas Emissions Data Collection Recommended Practices - <https://www.scte.org/standards/library/catalog/scte-208-cable-operator-greenhouse-gas-emissions-data-collection/>

SCTE 219 2021 – Technical Facility Climate Optimization Methodology - <https://www.scte.org/standards/library/catalog/scte-219-technical-facility-climate-optimization-methodology/>

SCTE 253 2019 Cable Technical Facility Climate Optimization Operational Practice: Understanding Set Point Values, Part 1 - <https://www.scte.org/standards/library/catalog/scte-253-cable-technical-facility-climate-optimization-understanding-set-point-values-part1/>

SCTE 274 2021 Cable Operator Critical Facility Air Containment Operational Practice -
<https://www.scte.org/standards/library/catalog/scte-274-cable-operator-critical-facility-air-containment/>

Improve Routing Security by validating BGP (Border Gateway Protocol) with RPKI (Resource Public Key Infrastructure)

A Technical Paper prepared for SCTE by

Tony Tauber

Distinguished Engineer
Comcast
Cambridge, MA
tony_tauber@comcast.com

Courtney Smith

Principal Engineer
Comcast
Mt. Laurel, NJ
courtney_smith@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. BGP Background.....	3
3. BGP “Hijack” Risks.....	4
4. RPKI and ROV Overview	5
5. Deployment - Reading and Writing (Vaildating and Publishing).....	7
5.1. Validating.....	7
5.1.1. Relying Party software	7
5.1.2. Router configuration.....	8
5.1.3. Deployment considerations.....	11
5.2. Publishing.....	11
5.2.1. Hosted vs. Delegated (vs. Hybrid) model	12
5.2.2. Risks	12
5.2.3. Bottom-up ROA creation	13
5.3. Learning	13
5.3.1. Bugs.....	13
5.3.2. Monitoring and instrumentation.....	14
5.4. Future work	14
6. Conclusion.....	14
Abbreviations	16
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1: BGP Hijacks steal and divert Internet traffic to attackers.....	4
Figure 2: 2008 Pakistan Telecom BGP incident	5
Figure 3: RPKI (Reseource Public Key Infrastucture) Delegation Structure	6
Figure 4: ROA (Route Origin Authorization) Creation.....	7
Figure 5: RPKI ROV (Route Origin Validation) Data Flows	9
Figure 6: RPKI ROV Analysis of IPv4 prefix-origin pairs	10

1. Introduction

In this paper, we will discuss the basic operation of BGP and inter-provider Internet routing including some vulnerabilities of the system. We will then describe RPKI, a set of technologies developed by the IETF (Internet Engineering Task Force) to help address a sub-set of these vulnerabilities. Deployment of these tools is not without risk and complication, and we will describe how we went about assessing and enabling RPKI in a large MSO network including design tradeoffs and lessons learned.

This is a complex set of technologies which have matured over the course of a decade or more of implementation and deployment experience. It operates across many operational and software domains with many actors and interacting systems. Hence decisions whether to deploy it and exactly how to do so while minimizing risk and disruption must be balanced carefully. These risks and complexities are magnified as function of the number of resources, interactions, and systems involved.

Making use of RPKI to improve security for an MSO (Multi-System Operator) service provider network is just such a complicated scenario and special care must be taken. In this paper, we describe the various considerations we assessed and decisions we made in the course of this enablement.

2. BGP Background

In order to communicate with given resource on the Internet, it's necessary to send packets to the destination IP (Internet Protocol) address which can service that request. BGP (Border Gateway Protocol) is the method different networks that make up the Internet use to communicate what endpoints (IP addresses) can be reached on their infrastructure. As put in [RFC 4271](#), “[t]he classic definition of an Autonomous System is a set of routers under a single technical administration”. BGP speakers (routers configured with BGP) within an AS can communicate reachability information for IP prefixes which are reachable within the infrastructure of that AS to neighboring ASes.

As a BGP message (sometimes also called an “advertisement” or “announcement”), leaves an AS border router to an neighboring AS, if the first AS originated the announcement for a given IP address block, that AS puts its ASN (Autonomous System Number) in the BGP message as the “origin AS”. Each AS that conveys that reachability information onward via BGP appends its own ASN in a chain known as the “AS-Path”. The AS-Path has a few functions including preventing loops in the topology and also operationally tracing responsibility for handing routing and traffic for a given destination.

To take an example, if AS150 originated a BGP route for 2001:db8::/32 which was then received by AS280 and passed along from there to AS320, we would expect an AS-Path of “320 280 150” read from left to right as “nearest” to “furthest”, ultimately to “origin” AS.

[Include references and pictures.]

3. BGP “Hijack” Risks

The basic BGP protocol was specified almost three decades ago ([RFC 1654](#) superseded by later RFCs) and for more than two decades the operational community has recognized some security risks with the protocol.

One important aspect of the Internet routing system is that routing information flows not only between directly adjacent ASes but also, transitively, to distant ASes. The complexity and dynamic nature of the global Internet routing system historically has not had any “ground truth” declaration of the proper topology and relationships among the various ASes and IP address blocks connected to them.

BGP mis-originations, sometimes called “hijacks”, refer to cases where a network other than the legitimate address-holder incorrectly advertises reachability for a given IP address block. The term “hijack” suggests malicious intent however in most cases the intent of the actors involved is not disclosed or definitively known. Some cases appear to be the result of a configuration error where some characters in the address block may have been transposed or another typographical error has occurred. In some other cases, a bug in a router or other piece of network equipment seems to have caused the problem. The risk of such events can be visualized by this diagram:




Figure 1: BGP Hijacks steal and divert Internet traffic to attackers.

(Courtesy USA National Institute of Standards and Technology)

Perhaps the most well-known case is the 2008 Pakistan Telecom incident which appeared to involve a BGP redirection within the Pakistan Telecom network to possibly enable censorship of a popular Internet video provider in country. The bogus BGP information leaked out to the broader Internet and ended up disrupting legitimate traffic intended for said provider as illustrated in the diagram below

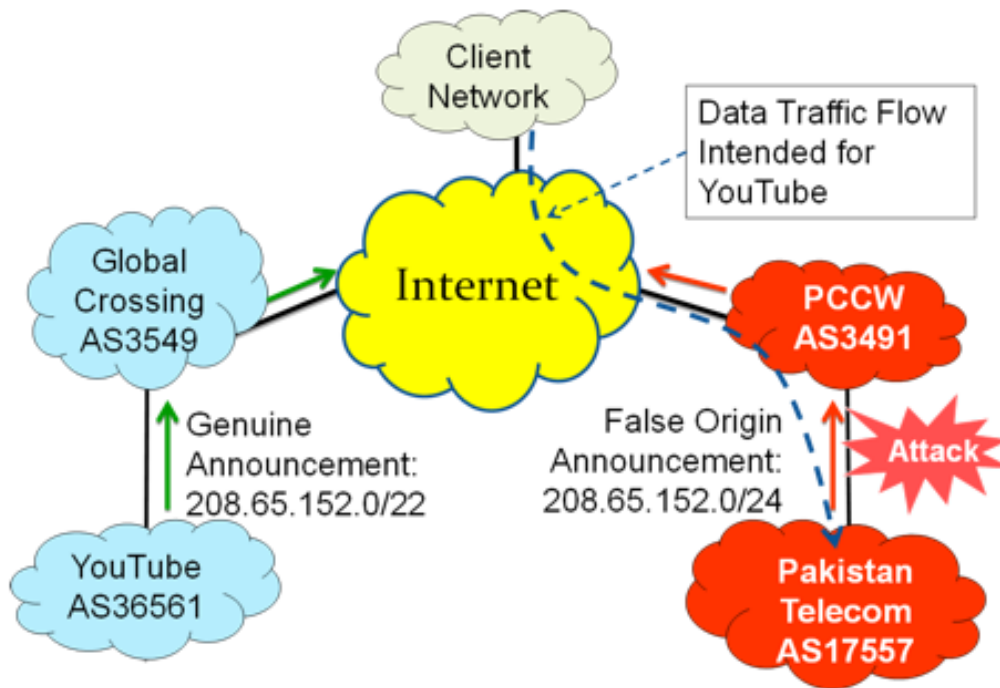


Figure 2: 2008 Pakistan Telecom BGP incident

(Courtesy USA National Institute of Standards and Technology)

A [catalogue of public BGP hijack incidents](#) on Wikipedia includes approximately twenty episodes spanning many years and that list is not exhaustive.

To address this particular weakness in the routing system, the [IETF SIDR](#) (Secure Inter-Domain Routing) Working Group worked to develop new enabling technologies. We will discuss the first one of those below and what Comcast did to utilize them.

4. RPKI and ROV Overview

RPKI (Resource Public Key Infrastructure) defined in RFCs 6480-6493, is a method which follows the IP address assignment hierarchy. The root of Internet addressing rests with IANA (the Internet Assigned Number Authority) which delegates addresses to the 5 Regional Internet Registries (RIRs) assign IP address resources to different network operators or delegate further to LIRs (Local Internet Registries which can be on a country (e.g., China, Brazil) or Service Provider level.

The structure of the RPKI uses X.509 digital certificates to enable cryptographic verification of the chain of authority and particularly to issue ROA (Route Origin Authorization) objects which describe a mapping of an IP address range and the origin AS (Autonomous System) which has authority to announce reachability via BGP.

The figures below illustrate the delegation hierarchy manifested in the RPKI as well as the components needed to create a ROA (Route Origin Authorization) object.

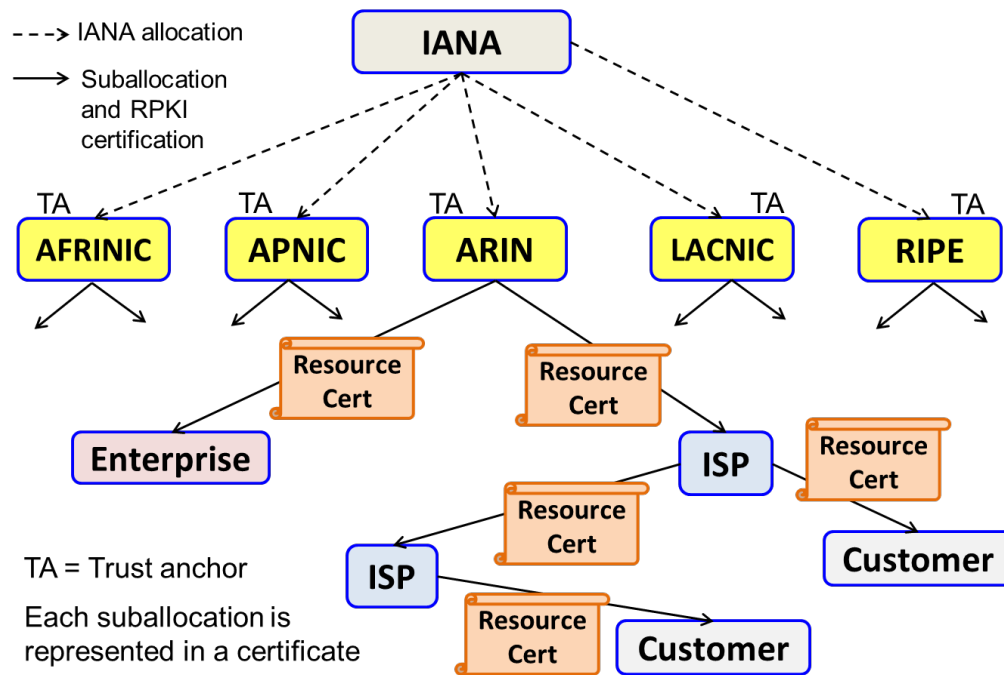


Figure 3: RPKI (Resource Public Key Infrastructure) Delegation Structure

(Courtesy USA National Institute of Standards and Technology)

A ROA consists primarily of three elements:

- IP prefix (IPv4 or IPv6) or range
- Maximum valid length for a prefix within that range
- The AS authorized to announce (“originate”) that prefix in BGP

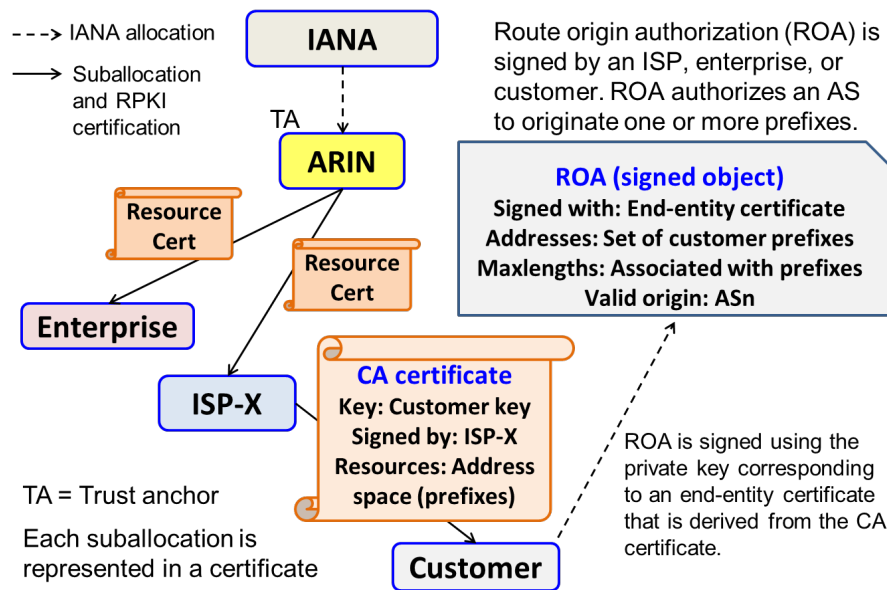


Figure 4: ROA (Route Origin Authorization) Creation

(Courtesy USA National Institute of Standards and Technology)

These objects and associated artifacts, e.g., manifests and CRLs (Certificate Revocation Lists) are published so that they may be verified using any of several compliant RP (Relying Party) software packages. Once the verification is complete, the VC (Validating Cache) system renders the output mappings in a form that can be consumed by routers using the RTR (RPKI To Router) protocol.

A compliant router can then be configured to use the information to validate incoming BGP announcements and take action based on comparing them to this set of authorizations. This process is known as ROV (Route Origin Validation).

5. Deployment - Reading and Writing (Validating and Publishing)

There are two aspects to RPKI ROV; validating and publishing RPKI data which can be thought of as “reading” and “writing”. It is not required to deploy these at the same time, in a given order, or for a given operator to do both, necessarily. We will discuss each below.

5.1. Validating

In order to perform validation of BGP information, the source data must be collected and loaded into the routers. In this section, we will outline the components and process.

5.1.1. Relying Party software

A number of freely available [open source software implementations](#) of the RP (Relying Party) function have been developed over the spread of several years. Some of them have been maintained more vigorously than others.

RP software is run on VC (Validating Cache) servers and must have access to reach any Internet destination via at least the *rsync* protocol (TCP port 873) and RRDP (RPKI Repository Delta Protocol) in order to collect the RPKI data from the various publication points which cannot be known a priori and are communicated via URIs, including DNS (Domain Name System) names, embedded within RPKI data; hence, DNS-based and not IP-address based and not amenable to IP-based firewall rules or access controls.

After fetching all the RPKI data and performing cryptographic and other consistency checks against it, the VC produces a list of VRPs (Validated ROA Payloads) consisting of an IP prefix (IPv4 or IPv6), a maximum length (often referred as “maxlen”), and an origin AS authorized to originate such a prefix up to the “maxlen” value.

Additionally, [RFC 8416](#) defines SLURM (Simplified Local Internet Number Resource Management) by which the operator of a VC can inject “local overrides” which will augment the public RPKI data as may suit the network operator’s needs.

5.1.2. Router configuration

The routers consume VRP data from the VC servers via the RTR (RPKI To Router) protocol, defined in [RFC 8210](#), which runs via a TCP connection by default port 323 but can be some other user-defined port. Multiple RTR servers can be defined, and, for most router vendors, the router considers the union of all VRPs received via RTR. Refer to the diagram below for illustration of the data exchanges involved.

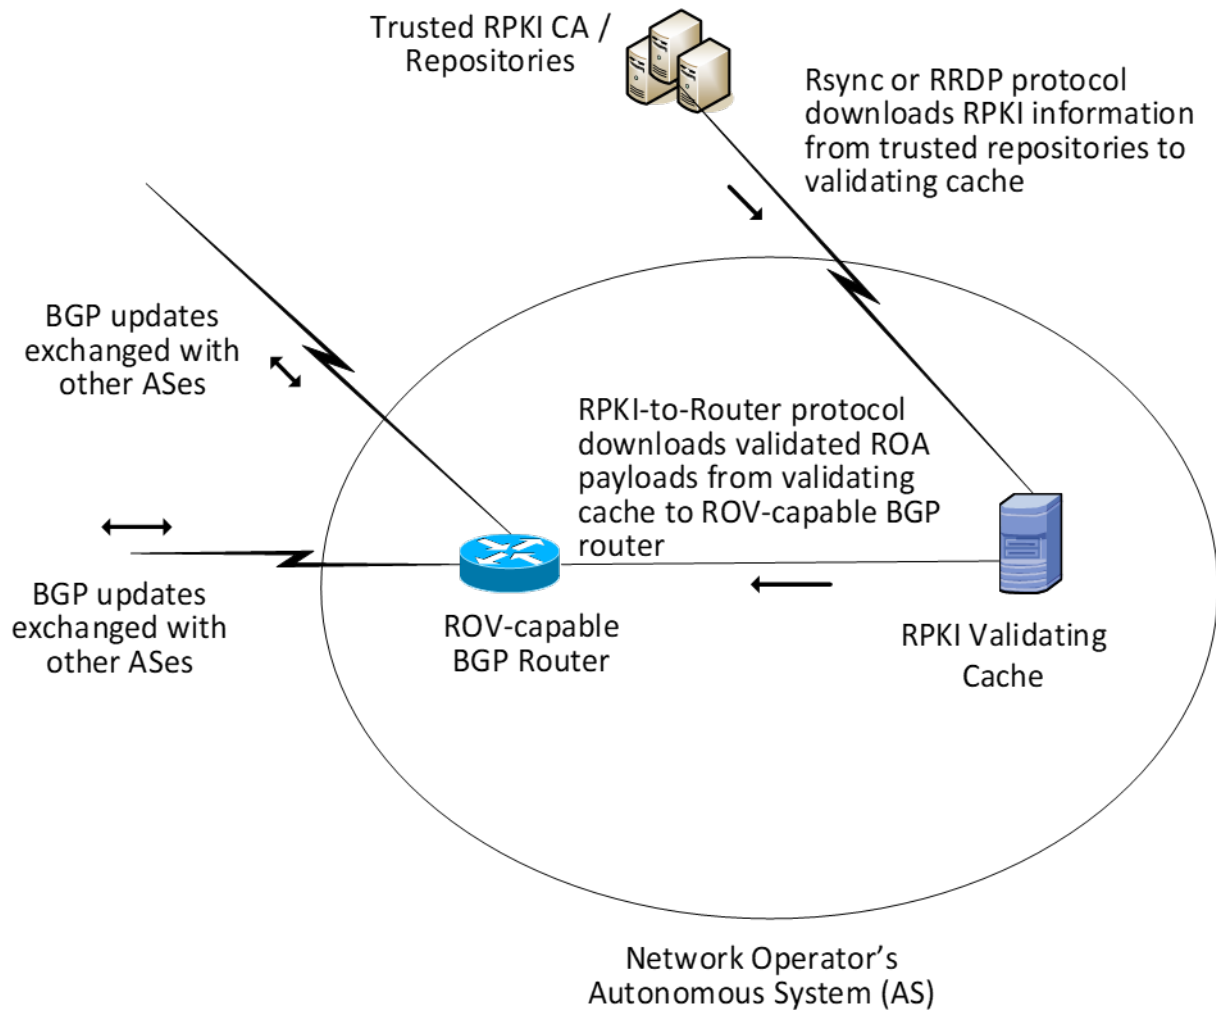


Figure 5: RPKI ROV (Route Origin Validation) Data Flows

(Courtesy USA National Institute of Standards and Technology)

Then the router is configured to perform ROV on the BGP table. The configuration options and implementation details among the vendors may vary somewhat but for our purposes, we decided that the most important risk of erroneous BGP announcements was from other networks; hence we would focus on eBGP sessions to other networks (in our case we have eBGP sessions between different ASes that are under our administration so not truly “external”). With the configuration we selected, each BGP announcement received over these eBGP sessions would be evaluated against a local table of VRPs (Validated ROA Payloads) received from the VCs. The validation of a BGP route against a ROA has one of three possible outcomes:

- **NotFound (a.k.a. Unknown)**
 - BGP route does not match any ROA
- **Valid**
 - BGP route matches a ROA – same Origin AS and same length or w/in “maxlen”
- **Invalid**

- The ROA and route announcement differ either of these ways:
 - Originating ASN
 - Maximum length ("maxlen")

It is very important to note that the operational model for RPKI ROV is that if a route is in the "not-found" category (no matching ROA in the routers local cache), the route should be accepted in the routing table and traffic forwarded to that destination. Routes which are "valid" should, of course, be accepted and "invalid" routes should be dropped. In this sense ROV is said to follow a "fail open" model which means that if the connection to the VC server is lost or the VC server loses its feed of RPKI information, traffic will continue to flow and not cause an outage.

At the time of this writing, (June 2022), below is a view of the global IPv4 BGP table and what routes are covered by ROAs. As this graph shows most of the Internet IPv4 routing table is not covered by any ROA but an ever-growing slice is covered, and a very small amount seems covered but not valid per the logic above. At the time of our deployment the portion covered by a valid ROA was closer to 30%. The net effect of these factors meant that, if there were to be some failure in the RPKI components as we were deploying them, we would end up in no worse a situation than we had been prior to the rollout.




Figure 6: RPKI ROV Analysis of IPv4 prefix-origin pairs

(Courtesy USA National Institute of Standards and Technology)

5.1.3. Deployment considerations

Given the fairly large number of variables that we had to confront, some basic guiding principles made it possible to narrow the world of possibility to a smaller set of options. We sought to reduce the potential risks involved with a new technology and these concepts helped:

- Diversity
- Redundancy
- Incremental rollout (and rollback)

5.1.3.1. Diverse RP software packages

One key decision was to utilize two different [RP software packages](#). RPKI is still a maturing technology with many different deployment insights and considerations coming to light over time. By using two different implementations, if one had a problem either transient or perhaps somewhat longer lasting, that package could be taken out of production temporarily.

Initially, we chose:

- [Routinator](#) (written in the Rust programming language by NLnet Labs)
- [RIPE Validator](#) (written in Java by RIPE NCC)

When RIPE Validator was [declared end-of-life in 2021](#), we replaced that with:

- [rpki-client](#) (written in C as part of the OpenBSD project)

5.1.3.2. Diverse data centers

Since a given data center could conceivably to offline or become isolated from some part of the network, we further refined our redundancy plan utilize one two VC servers in each of two diverse data centers (e.g., “East” and “West”). Each data center then would feature one of each of the above RP packages running.

Each edge router then would be configured with RTR sessions to each of the four VC servers. The final piece is to enable validation within BGP on a per-neighbor basis.

5.1.3.3. Incremental rollout which can be incrementally rolled back

Once a router is configured with RTR sessions with VCs, it is then ready to have validation enabled on eBGP sessions with other networks. We chose to enable ROV first on BGP sessions with a given partner network and went through a field trial working first with one, then another, and then began to expand the list. This deployment approach allowed for incremental activation and, if needed, incremental roll-back of the changes should any problems be encountered.

We started with a few partner networks and one router platform to be able to monitor for any possible ill effects and gradually expanded this circle to include others.

5.2. Publishing

In order to get protection from the RPKI-ROV system, an address-holder must publish ROAs to cover their IP address space indicating which AS(es) are authorized to originate BGP announcements for said address space. Comcast’s address space has been issued almost entirely

within the ARIN (American Registry for Internet Numbers) region so that was the focus of our considerations.

5.2.1. Hosted vs. Delegated (vs. Hybrid) model

At the time we were planning our deployment, there were two models for publishing RPKI data offered by ARIN and we will discuss each of those options here.

In the “[hosted](#)” model an organization generates a public-private key pair and sends the public key to ARIN and receives a “Resource Certificate” which covers the resources that ARIN has issued the organization. The organization is then able to sign ROAs with its corresponding private key and load them each into the ARIN system.

In the “[delegated](#)” model, the “parent” registry (ARIN in this case) issues a resource certificate to the “child” which hosts its own CA (Certificate Authority) and can issue ROAs for said resources or even further sub-delegate if desired.

There were (and are at the time of this writing) [two open-source software packages](#) which can be used as a CA and to issue ROAs:

- [Krill](#) (written in Rust by NLnet Labs)
- [rpkid](#) (written in Python2 and C by Dragon Labs)

In considering which model would make the best choice for us, the relative immaturity of these software packages and lack of broad operational experience suggested the best approach would be to use the “hosted” model for the initial rollout at least. Similarly, there had been a few instances operational problems of RIR repositories or CAs and should those happen, it would be preferable to share fate with a large number of RPKI users such that the most operational attention and expertise could be brought to bear on a solution.

One complication to the delegated model is that, as initially conceived, the publication of the ROAs and other RPKI material (manifests, CRLs – Certificate Revocation Lists) is done by the operator of the delegated CA. However, this represents an operational risk and has different availability requirements than the CA itself. Hence the “hybrid” or “publish in parent” approach has been developed and is offered by a number of RIR or LIR issuers. This option was not available from ARIN at the time of our deployment so could not be considered.

In either model, software mechanisms can be brought to bear to automate the issuing and analysis of ROAs against other information such as IP address management systems and routing tables. In both scenarios well-featured APIs (Application Programming Interfaces) are available for these purposes.

5.2.2. Risks

Before we started issuing ROAs for our address space, all our routes would be in the “not-found” status. As we start issuing ROAs, they must ensure they properly match the BGP routes that we advertise to other networks.

For many network operators, particularly many or most Enterprises and perhaps CDN (Content

Distribution Network) operators keeping these elements (BGP announcements and ROAs) aligned might be quite straightforward. However, for Comcast's network, the situation is rather complex. We have something over one hundred IP address blocks issued from ARIN however these are split up across more than twenty ASes, networks, and services. There is also a certain amount of dynamism in the addressing.

Our addressing design, generally, features the large IP address blocks being advertised from our Backbone network (AS7922) and different parts issued among our regional, data center, enterprise and other networks. Hence many different more-specific routes within our blocks would show up to other networks with different origin AS numbers.

5.2.3. Bottom-up ROA creation

Given this routing and addressing design, if we were to issue ROAs for the large blocks with an origin of AS7922, immediately all other routes which were properly visible for reasons of particular traffic flow would become "invalid" to other networks that were doing ROV which was a significant list by the time of our rollout.

Instead, the proper course was to issue ROAs for the various smaller-sized BGP announcements that were expected to be visible externally. (It is normal for a network to carry more-specific announcements internally and Comcast has thousands of these.) We gradually filled in ROAs for these pieces over the course of several weeks, starting very slowly until we had a good level of confidence we could properly detect if there was any problem caused. That detection included monitoring the level of traffic destined for a prefix using NetFlow data collected at our network edges.

Once all the more-specific routes in a given block were complete, we could "top off" by issuing a ROA for that large block and the over-arching advertisement from our Backbone AS. Issuing ROAs to deal with incremental changes after this first large deployment would be simpler and bear little risk.

All this work was considerably aided by software automation for the analysis and ROA issuance. Some was internally developed as well the [ARIN ROA request](#) script developed by Rich Compton of Charter Communications which, as the name suggests, uses ARIN APIs to effect bulk issuing (or deletion) of ROAs in ARIN's hosted system.

5.3. Learning

In the course of our deployment, a number of realizations came to light that we share here in hopes that it might be useful in the considerations of others who contemplate a similar course.

5.3.1. Bugs

Almost inevitably, bugs and unforeseen error conditions are discovered in the course of implementing new technology, particularly one as complex as RPKI. We had the good fortune to not be an early adopter and that strategy proved beneficial as many early growing pains had already been found and fixed by others and we appreciate their role in the development.

5.3.1.1. RP software

In some cases, one or another of the RP software packages would yield significantly different results from the other. Generally, these cases resulted from some differences in processing or implementation details sometimes including how they interacted with other infrastructure (e.g., publication points).

There have also been a few vulnerability reports (though none were catastrophic) and patches issued at different rates.

5.3.1.2. Router software

We worked with our router vendors for advice on best current code in the software trains that we use prior to enabling ROV. If there were outstanding bug reports, we examined the circumstances and impacts to help assess the relative risks.

5.3.2. Monitoring and instrumentation

From time to time, on one vendor's router platform we would see occasional situations where the number of prefixes received from the VC servers via RTR sessions would drop to zero and only slowly get repopulated as new VRPs came in. Monitoring the RTR sessions on each router which is configured with them is crucial for detecting and remedying the situation.

Each of our RP software packages comes with the ability to publish metrics for consumption within Telegraph, Prometheus or some similar monitoring solution and then to visualize the data over time using Grafana. This can allow for tracking any anomalous behavior to a given start time.

5.4. Future work

As of this writing, ARIN has begun supporting a [“hybrid” or “publish in parent”](#) ROA deployment model where the CA server is housed and managed by the network operator but the publication can be done on ARIN's high-availability Publication Point servers. We will likely consider making use of this which could also simplify administration of ROAs for address space issued by other RIRs.

The IETF [SIDROPS](#) (Secure Inter-Domain Routing OperationS) Working Group is where engineers from industry and other interested parties (e.g., academia, government, etc.) continue to convene to monitor the expanding deployment of RPKI and related technologies and discuss incremental changes or advice that may be appropriate. Also, under development are expanded uses of the RPKI architecture to further improve the security and resilience of the global Internet routing system.

6. Conclusion

The global Internet routing system is critical to the services, commerce, education, and entertainment of the world's population. Over the course of decades it has, through careful engineering and collaboration, managed to grow and scale to meet the changing objectives of its users. The fundamental security model is one piece that needed to evolve but also required

backward compatibility for incremental rollout while still realizing benefit with partial implementation. RPKI is the enabling foundation and ROV is the first realization of such improvements.

The operational model of RPKI is reasonably flexible to accommodate many different operational and deployment scenarios, however this malleability also means that deciding exactly how to deploy and in what order can seem daunting. By gradually narrowing options by adoption of guiding principles, it becomes easier to plot the way forward.

Abbreviations

AfriNIC	African Network Information Center
API	Application Programming Interface
APNIC	Asia-Pacific Network Information Center
ARIN	American Registry for Internet Numbers
AS	Autonomous System
ASN	Autonomous System Number
BGP	Border Gateway Protocol
CA	Certificate Authority
CRL	Certificate Revocation List
eBGP	External Border Gateway Protocol
IANA	Internet Assigned Numbers Authority
iBGP	Internal Border Gateway Protocol
IETF	Internet Engineering Task Force
IP	Internet Protocol
LACNIC	Latin America and Caribbean Network Information Centre
LIR	Local Internet Registry
PP	Publication Point
RIPE (NCC)	Réseaux IP Européens (Network Coordination Centre)
RIR	Regional Internet Registry
ROA	Route Origin Authorization
RP	Relying Party
RPKI	Resource Public Key Infrastructure
RRDP	RPKI Repository Delta Protocol
RTR	RPKI To Router protocol
SIDR	Secure Inter-Domain Routing
SIDROPS	Secure Inter-Domain Routing OPerationS
SLURM	Simplified Local Internet Number Resource Management
VC	Validating Cache

Bibliography & References

Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006

J. Hawkinson and T. Bates. "Guidelines for creation, selection, and registration of an Autonomous System (AS)", [RFC 1930](#), [doi:10.17487/RFC1930](#), March 1996

B. R. Smith and J. J. Garcia-Luna-Aceves, "Securing the border gateway routing protocol," *Proceedings of GLOBECOM'96. 1996 IEEE Global Telecommunications Conference*, 1996, pp. 81-85, doi:10.1109/GLOCOM.1996.586129.

S. Kent, C. Lynn and K. Seo, "Secure Border Gateway Protocol (S-BGP)," in *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582-592, April 2000, DOI: 10.1109/49.839934

Murphy, S., "BGP Security Vulnerabilities Analysis", [RFC 4272](#), DOI 10.17487/RFC4272, January 2006

Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", [RFC 6811](#), January 2013

K. Zetter, "[Revealed: The Internet's Biggest Security Hole.](#)," Wired, August 2008.

N. Anderson, "[How China swallowed 15% of 'Net traffic for 18 minutes](#)," ArsTechnica, November 2010.

D. Godin, *Traffic Misroutes through China Telecom* <https://arstechnica.com/information-technology/2018/11/strange-snafu-misroutes-domestic-us-internet-traffic-through-china-telecom/>, ArsTechnica, November 2018

Wikipedia, *BGP Hijacking – Public Incidents*
https://en.wikipedia.org/wiki/BGP_hijacking#Public_incidents

NIST, *Robust Interdomain Routing*, <https://www.nist.gov/programs-projects/robust-inter-domain-routing>, April 2022

Haag W., Montgomery D., Barker William C., Tan A., [Protecting the Integrity of Internet Routing: Border Gateway Protocol \(BGP\) Route Origin Validation](#), NIST Special Publication, NIST-SP-1800-14, July 2018

Sriram K., Montgomery D., [Resilient Interdomain Traffic Exchange: BGP Security and DDoS Mitigation](#), NIST SP 800-189, December 2019

Hannachi L., Sriram K., Borchert O., Montgomery D., [NIST RPKI Monitor](#), NIST Software Release, April 2021

Improving Operational Intelligence for Maintaining Cable Networks

A Technical Paper prepared for SCTE by

Mike Spaulding

Vice President, Plant Maintenance
Comcast Corporation
1701 John F Kennedy Blvd Philadelphia, PA 19103
303-917-4003
Michael_Spaulding@cable.comcast.com

Larry Wolcott

Engineering Fellow
Comcast Corporation
183 Inverness Drive West, Englewood, CO 80122
303-726-1596
Larry_Wolcott@cable.comcast.com

Jason Rupe

Principal Architect
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
303-661-3332
j.rupe@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Failure Mode Effects and Criticality Assesment (FMECA).....	3
2.1. Background	3
2.2. Definition	3
2.2.1. Failure Mode	3
2.2.2. Effects	4
2.2.3. Criticality.....	4
2.2.4. Probability	4
2.2.5. Causality and Decay	4
2.3. Application.....	4
2.3.1. (FM - Failure Mode) Event Management.....	5
2.3.2. (E - Effect) Localization and Service Impact.....	6
2.3.3. (C - Criticality) Severity and Decay.....	7
2.3.4. (A - Analysis) Benefit Analysis	8
2.3.5. Feedback Loop	9
3. Use Cases	10
3.1. Water Damaged Cables.....	11
3.1.1. Specification.....	11
3.1.2. Detection	11
3.1.3. Localization	12
3.1.4. Criticality.....	13
3.1.5. Effects	14
3.1.6. Benefits	15
4. Conclusion.....	15
Abbreviations	16
Bibliography & References.....	16

List of Figures

Title	Page Number
Figure 1 - Failure modes, effects and criticality analysis - United States Army.....	5
Figure 2 – Example of symptomology-based event management.....	6
Figure 3 – Example of FMECA-based event management	6
Figure 4 – HFC Network topology represented as a graph model	7
Figure 5 – Example of benefit FMECA input, including failures, incedents and criticality	8
Figure 6 – Example of benefit FMECA output, icluding network reliability	9
Figure 7 – Closed loop detection, reccomendation and repair validation.....	10
Figure 8 – Downstream RF failure mode examples from SCTE Industry Reference 208.....	10
Figure 9 – Upstream RF failure mode examples from upcoming SCTE Industry Reference.....	11
Figure 10 – Water impaired frequency response.....	12
Figure 11 – Water signature localized to a single location indicating damaged drop cable	13
Figure 12 - Example degradation of cable with exposure and water penetration.....	14
Figure 13 – Field trial, water wave before and after repair	15

1. Introduction

Cable operators have vast amounts of network performance data available to detect and measure defects within the cable plant. Transforming this data into actionable intelligence can be a daunting task, especially having many existing systems and process already in place.

Failure Mode, Effects, and Criticality Analysis (FMECA) was first developed in the 1940s (Summy) and is now widely used in industries including aerospace, automotive, and electronics. The role of FMECA is to identify potential problems that may occur in a system or component, define how to detect the failures, and measure the effect. This analysis provides a consistent way to measure the criticality of common network failures and help prioritize repair efforts.

Authors Spaulding, Rupe, and Wolcott will present the jointly developed progress of the CableLabs and SCTE working groups for Proactive Network Maintenance (PNM), with an operator's perspective from the field. This paper will demonstrate how FMECA can be used to improve customer experience, reduce trouble calls, and increase operational efficiency.

2. Failure Mode Effects and Criticality Assessment (FMECA)

2.1. Background

FMECA is a highly detailed, disciplined process for understanding the impact and mitigation value of critical failure modes. Related to FMEA which is a similar analysis without the criticality component present in the analysis. This process has been around for several decades most notably used in the aerospace industry. The process was used in the aerospace industry to understand the needs for redundancy based on identified critical failures. In other industries such as healthcare, this same FMECA model has been used to identify high risk processes. In an American Society for health care Risk Management (ASHRM) White Paper, the FMECA process is described as “proactive examination of what could go wrong and the opportunity to fix it before it fails.”(Summy)

FMECA has been the underpinning of the pattern detection work in the PNM working group. Identifying failure modes based on distinguishable patterns has been one of the most widely used applications of FMECA within our industry. According to co-author Jason Rupe FMECA is defined as “identifying a system and determining if it performs as designed.” (Spaulding and Rupe) He further explained that if you are looking at a power network and a fuse blew, did the fuse do what it was designed to do? Yes. The question then becomes what failed leading to the fuse blowing and how do we prevent it. FMECA takes us through the process of understand the *effects, criticality, probability, and decay*. These pieces of information allow us to understand the value and timing of mitigation.

2.2. Definition

2.2.1. Failure Mode

Identifying a condition that signifies an abnormal operation of a system or component in a network, system, or process. These conditions can take the form of a process within a system that no longer meets the requirements of the desired outcome. In a hardware application, failure modes would indicate when components are no longer working as designed. The failure mode itself does not represent severity just the abnormal operation or outcome.

2.2.2. *Effects*

The meaning of effect in this document will be any empirical measurement of system performance. In Data-Over-Cable Service Interface Specifications (DOCSIS) systems, there are a number of measurements available including registration state, packet loss, latency, forward error correction (FEC), modulation error ratio (MER) and others. Each of the failure modes typically have gradient errors that can be detected, depending on the severity and capacity available for mitigation. For this reason, each failure mode can have discrete effects measurements.

2.2.3. *Criticality*

In FMECA, criticality assessment may be qualitative or quantitative. For qualitative assessment, a mishap probability code or number is assigned and entered on the matrix. For quantitative measurements, ratios may be applied. This can be useful in network analysis to scale problem severity with the number of nodes effected.

2.2.4. *Probability*

While not called out in FMECA, the probability of a failure mode is an important consideration when deciding what to do about it. Highly likely failure modes with high criticality should be aggressively mitigated through operations. These are the drivers of cost, friction, and many undesirable outcomes. But even high probability and lower criticality failure modes deserve attention, so they can be properly addressed through fault management, mitigation, and repair measures. On the other hand, highly critical failure modes must be addressed when they happen, no matter how unlikely.

2.2.5. *Causality and Decay*

Causality, describing degradation or decay, is key to how we can use this in our PNM efforts. In today's environment, we have become experts in identifying faults once an impairment starts to impact DOCSIS performance. We detect the impact of a failure mode through corrected and uncorrected packet errors, unbonded channels, etc. When these effects meet certain thresholds, we dispatch technicians to fix the network failure (DM, or demand maintenance). As we move the thresholds on these metrics further from severity, we are getting ahead of the most severe customer impacts (PM, or preventative and proactive maintenance).

2.3. Application

FMECA is a methodology and process, rather than an application. However, the elements of FMECA map nicely to our cable domain which can easily be implemented as practical applications. In this section, we will decompose the elements of FMECA including cable-specific adaptations and demonstrate how it can apply to real cable systems. Figure 1 shows an example of the FMECA process.

After defining the components of FMECA we will further consider how to operationalize it with data which is readily available to cable operators. Using the model below we combine the components together to determine a priority of events based on decay and criticality. Assigning severity rankings based on criticality and decay rate drives prioritization of how we should address the problems.

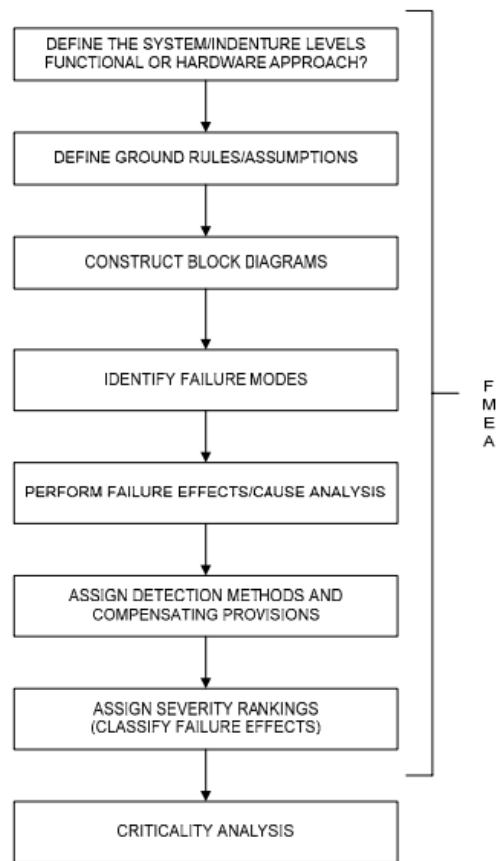


Figure 1 - Failure modes, effects and criticality analysis - United States Army

2.3.1. (FM - Failure Mode) Event Management

The most fundamental aspect of FMECA is identification of failure modes, or the places and reasons that components and/or systems can break. This might seem intuitive but many of our fault management systems have evolved based on symptomology, not necessarily specific cause-and-effect. For example, cable operators might have some form of network monitoring based on packet loss or FEC performance. However, these systems often lack in effective root cause analysis, delegating the troubleshooting process to technicians in the field or operations centers. Figure 2 shows an example of this symptomology-based event management, where potentially hundreds of individual errors are presented to technicians. Those technicians are then subject to analyze and repair based on their individual experience and training. Figure 3 shows this same system failure, modeled in FMECA which provides a root cause analysis, pointing the technician to the exact system component failure, including common repair recommendations. In this example, an amplifier module was experiencing a common form of failure associated to ground plane corrosion. The result was hundreds of symptomatic events which were difficult to interpret and repair efficiently.

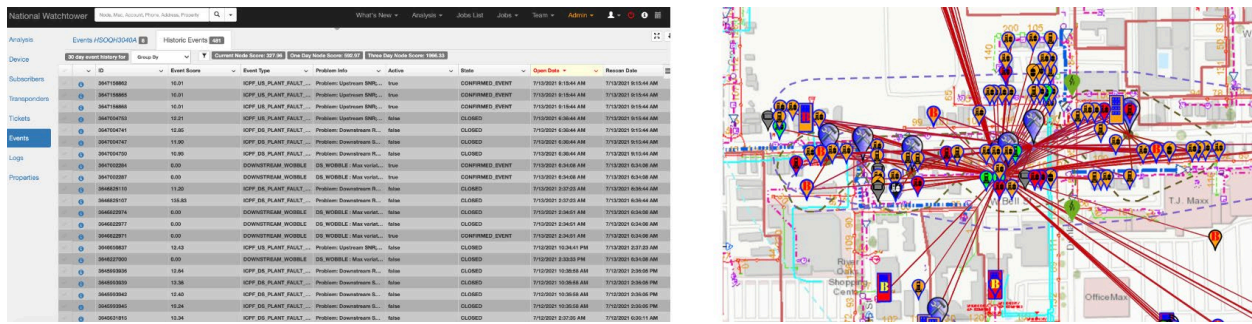


Figure 2 – Example of symptomology-based event management

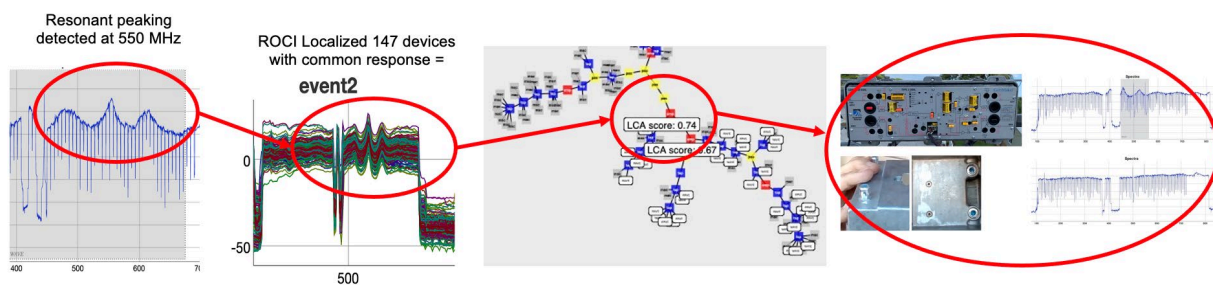


Figure 3 – Example of FMECA-based event management

2.3.2. (E - Effect) Localization and Service Impact

The first cable-specific adaptation is to determine the effect, which has additional complexity over some other component-based systems. In cable networks, HFC plant can be thought of (although not entirely correct) as a large, shielded antenna. The cable segments are largely passive with bidirectionality achieved with diplex filters. The performance sensors of our cable networks are typically at the transmitters and receivers (cable modems and CMTS) and often lack knowledge of the component chain between them, such as drops, taps, feeders, splitters, couplers and amplifiers. Because any one of these could be the point of failure, additional localization is usually required. In our example, a graph topology database and radio frequency (RF) signature analysis will be used to help provide fault location. Once the fault is localized, the tree-and-branch nature of hybrid fiber-coaxial (HFC) networks facilitates the cascading effect of service impact analysis (Figure 4).




Figure 4 – HFC Network topology represented as a graph model

Going beyond empirical measurements, we can also try to understand the subjective aspects of impaired service. For cable access network services, any friction incumbering the use of the service is an undesirable effect. For a user, friction builds up over time, but memory is not forever. If over a short period of time a customer experiences a lot of friction, their impression of the reliability of a service is damaged. If the friction to switch providers becomes lower, they may become a lost customer.

For network capacity, any signal impairment is an undesirable effect. But clearly, a complete loss of all services is a more significant effect than say noise or signal attenuation in a few frequencies which can be addressed through DOCSIS resiliency mechanisms, or an application that won't authenticate properly to work as intended. But some effects are early indicators of more significant issues to come. If any undesirable effect becomes significant, then it may rise in criticality.

2.3.3. (C - Criticality) Severity and Decay

The challenge in today's operating environment is showing the business benefit as we push to become more proactive. We look for optimization in initiatives on the basis of trouble call reduction and call-in rates; but what happens when you fix the network before the truck roll and before the calls? That is where modeling the causality as degradation or decay and adding that information to the methodology helps us to create the business and customer value. Through observation, we can determine the rate at which a particular failure mode will decay before it fails and creates the need for demand maintenance. By measuring the severity of degradation or decay, we can predict the point of failure and repair the failed component before the first call. This modeling doesn't work without a component of timing which is given through decay. With the decay timing we can then assign the value of fixing certain failure modes before they impact customers, which eventually generate the truck rolls or calls we've traditionally used to determine customer impact.

In DOCSIS access networks, there are a number of resiliency capabilities that help protect our services from network failures. A failure mode and its effects become critical if the effect becomes significant. That significance can be determined in a number of ways. But in systems and networks, significance is often estimated through a combination of the number of users impacted, the duration of the impact, and the severity of the impact from the effect. Complete loss of service is the most severe, whereas an

impairment that does not impact service is likely very low in severity from an end user perspective but may rise in importance if it signals an opportunity as it serves as an indicator of other issues. For example, an outage for an entire node is critical because it impacts all services ($S=\max$) for a large number of customers ($N \gg 1$); thus, there is strong incentive to return customers to service immediately, to reduce the outage duration (T). Criticality in this case could be estimated as $S*N*T$. On the other hand, a damaged drop is much less critical because it impacts a single customer ($N=1$) and may partially impact their services at worst case ($S \ll \max$); so, there is much less urgency with repair in these cases, thanks in part to the resiliency of DOCSIS.

Note that we don't need precise models to predict what will happen and when. It is often good enough just to recognize we found an opportunity to improve operational efficiency and delight customers.

2.3.4. (A - Analysis) Benefit Analysis

One of the outcomes of the FMECA analysis is to provide improved operational intelligence that helps transition from reactive to proactive network repairs. By modeling the number of faults with their severity and decay (Figure 5), the outcomes can be projected in new ways, such as network reliability (Figure 6).

Failure Subsystem	Failure Mode	Occurrence	Severity	Duration Days	Decay Rate	Decay Days	Criticality Model	Number of Subs
Home	Wiring	1	0.1	1	0.3	30	-20.00	1
Drop	Cut	1	1	1		60	40.00	1
Drop	Other Damage	1	0.3	0.5	0.025	180	-165.00	1
Drop	Water Damage	1	0.25	0.25	0.083	60	-53.75	1
Drop	Ingress - Customer	1	0.9	0.25		0	22.50	1
Drop	Ingress - Hot Drop	1	0.9	0.25		0	22.50	1
Tap	Damaged / Other	1	0.05	1		0	20.00	4
Tap	Damaged / Water	1	0.2	0.5		0	40.00	4
Feeder	Cut	1	1	1		0	1200.00	12
Feeder	Cracked	1	0.25	0.05		0	15.00	12
Feeder	Water Damage	1	0.6	0.3		0	216.00	12
Amplifier	Power Failure	1	1	1		0	2500.00	25
Amplifier	Grounding Fault	0.05	0.25	0.25		0	7.81	25
Amplifier	Failing Module	0.1	1	0.1		0	25.00	25
Hardline	Cut	1	1	1		0	5000.00	50
Hardline	Damaged	0.2	0.25	0.2		0	50.00	50
Hardline	Shielding Separation	1	0.25	1		0	1250.00	50
Sm Node	Power Failure	1	1	1		0	15000.00	150
Med Node	Power Failure	1	1	1		0	30000.00	300
Large Node	Power Failure	1	1	1		0	60000.00	600
Headend	ACP (Channel Alignment)	1	0.025	1		0	12500.00	5000

Figure 5 – Example of benefit FMECA input, including failures, incidents and criticality

Probability of Repair	Repair Disruption Mins	Benefit of Cost (BCR)	Benefits of TCs	Benefit of Churn	Benefit of Reliability
0.30	60.00	-0.09	-0.20	-0.01	-20.00
1.00	0.01	0.19	0.40	0.01	40.00
0.90	0.05	-0.74	-1.65	-0.04	-165.00
0.70	0.05	-0.22	-0.54	0.00	-53.75
0.70	0.05	0.09	0.23	0.00	22.50
0.70	0.05	0.09	0.23	0.01	22.50
0.95	0.05	0.29	0.80	0.02	80.00
0.75	0.25	0.27	1.60	0.04	160.00
1.00	0.01	22.15	144.00	3.60	14400.00
1.00	0.01	0.24	1.80	0.05	180.00
1.00	0.01	3.46	25.92	0.65	2592.00
0.50	0.01	83.33	625.00	15.63	62500.00
0.50	0.01	0.26	1.95	0.05	195.31
0.50	0.01	0.83	6.25	0.16	625.00
1.00	0.01	277.78	2500.00	62.50	250000.00
0.70	0.01	2.31	25.00	0.63	2500.00
0.70	0.01	57.87	625.00	15.63	62500.00
0.70	0.01	2083.33	22500.00	562.50	2250000.00
0.70	0.01	8333.33	90000.00	2250.00	9000000.00
0.70	0.01	33333.33	360000.00	9000.00	36000000.00
0.70	0.01	181159.42	625000.00	15625.00	62500000.00

Figure 6 – Example of benefit FMECA output, including network reliability

2.3.5. Feedback Loop

This is not described in FMECA systems and methodologies but is important to maintain a modern machine learning (ML) and artificial intelligence (AI) based system. As our cable access networks continue to evolve, the management systems need to become more adaptive. Especially as these systems become more reliant on ML and AI systems, the models need feedback to improve. By providing feedback to the system, they may continue to improve and adapt (Figure 7).

1. Failure mode detection using ML and AI pattern matching
2. FMECA knowledgebase provides repair recommendations to technician
3. Technician makes informed repairs, reducing the analysis time
4. System automatically validates repair and potentially solicits feedback
5. Pattern matching models and recommendations adapt

Ultimately, as the systems continue to adapt and improve, gains in repair times and network reliability become realized.




Figure 7 – Closed loop detection, recommendation and repair validation

3. Use Cases

In both the CableLabs and SCTE Network Operations Subcommittee working groups for PNM, a number of failure modes have been identified and documented. The following Figure 8 shows examples of these documented downstream RF failure modes; these failure modes to RF transmission are often referred to as impairments or faults, to differentiate from the network failure mode which causes the RF failure modes. The working groups are continuing to make progress in this area including understanding and troubleshooting upstream RF impairments and developing a repair matrix. Additional examples of this upcoming work can be seen in Figure 9.

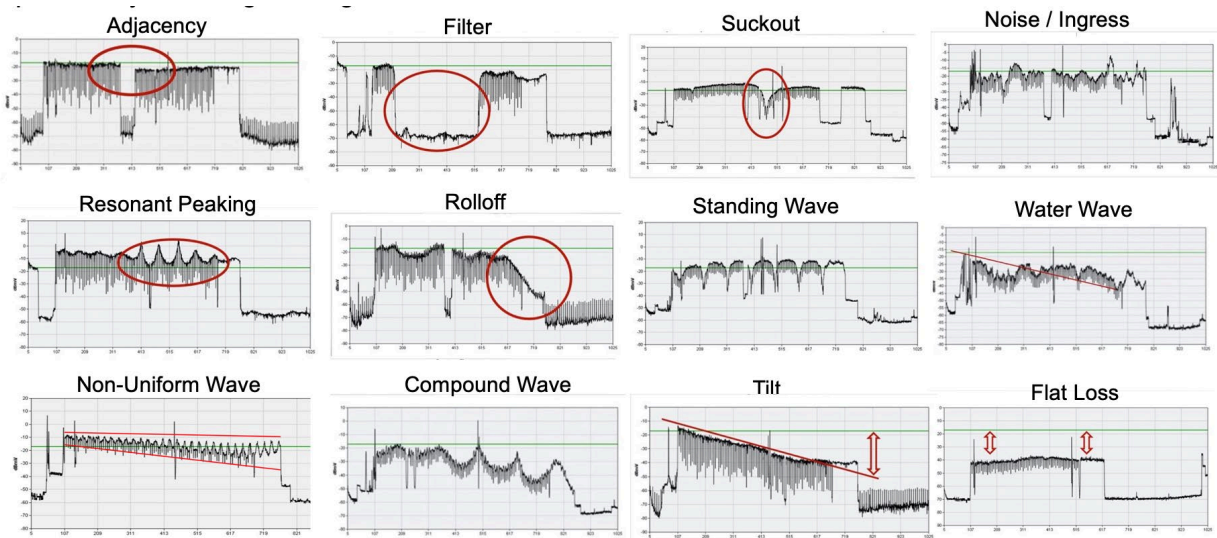


Figure 8 – Downstream RF failure mode examples from SCTE Industry Reference 208

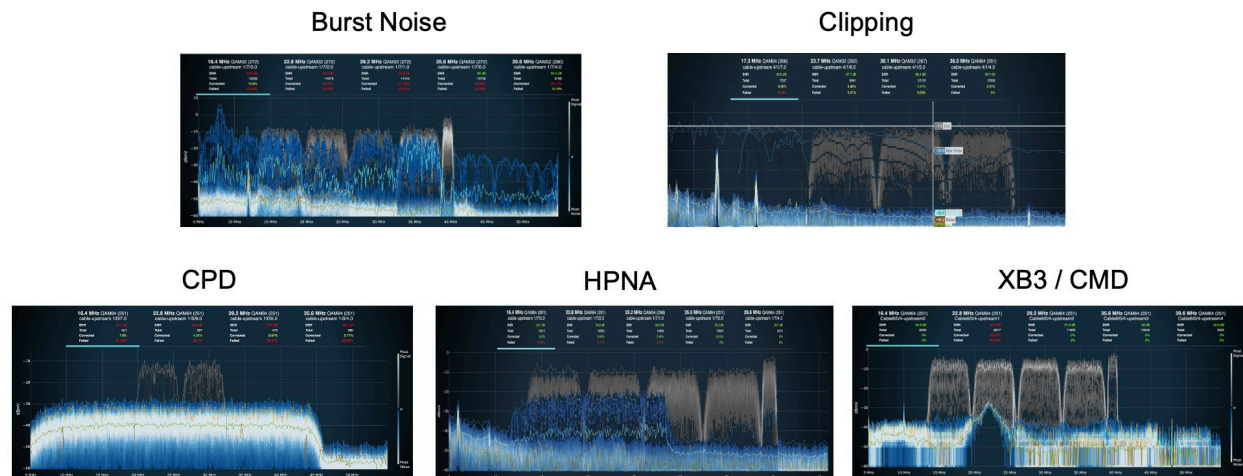


Figure 9 – Upstream RF failure mode examples from upcoming SCTE Industry Reference

3.1. Water Damaged Cables

In the previous material, we reviewed the fundamentals of FMECA and how the process could be adapted to work in a cable access network. In this section, we will apply what we've learned to a specific use case which has been a recurring theme in recent SCTE PNM working groups: water damaged cables.

3.1.1. Specification

The first step in our water damaged cable use case is to define the system and failure mode. Our implementation will rely on recent PNM advancements which allow cable operators to accurately detect, and measure cable failure associated to water ingress. This failure mode has been well documented and specified in previously published SCTE material, found in the bibliography.

3.1.2. Detection

Impairments in DOCSIS RF are easy to detect, but that is only the first step in efficient maintenance. A simple spectrum capture or RxMER per subcarrier plot will reveal an impairment in the signal, and the signature of the impairment indicates the type of fault. These and additional PNM tests and queries help us measure the impact on the RF signal, which we can translate to impact on service through a model.

Detection of the impairment and identifying its type allows us to do something about it and know how important it is to address the fault. Each fault type can behave differently over time, as it is exposed to elements such as heat, dryness, water, ice, and cold.

There have been many different methods devised for identifying one type of impairment from another, such as the CableLabs Spectrum Impairment Detection, or the CableLabs Anomaly Detector. (See [Zhu, Sundaresan, Rupe] for details on the machine learning based Anomaly Detector available for use by CableLabs members.) These methods all essentially match patterns in bin data, which are indicators of the fault type.

As explained in an Expo paper from 2021 [Fox, et. al.], water damaged cables are indicated from spectrum data with a signature of a few factors: a random (aperiodic) pattern of fluctuating attenuation

over frequency bins, and often a general trend of more attenuation toward higher frequencies. As a result, a pattern matching solution can be used to identify water in cables. As described in the paper mentioned here, a two phased approach works well: identify a standing wave, which is indicated by highly variable attenuation over frequency bins, and then use a secondary test that separates standing waves (which are periodic) from water waves (which are aperiodic).

Possible approaches, some described in that paper, include transforming the data further to find how periodic the variability may be, or calculating and removing any slope in the bin data. For example, taking the bin data, calculating and removing the slope trend from the data, and then using an FFT to transform the bin data into a new frequency domain to find the frequencies that might describe the variability in bin signal levels; a single strong frequency would indicate a standing wave, whereas several nearly equally strong frequencies would indicate water is the culprit.

We have also learned that water in a coax cable causes variability in the phase of the signal as well. With the ability to capture complex (I,Q) bin data, we can further confirm the presence of water in the cable.

All this information helps us localize too, as we'll explain next.




Figure 10 – Water impaired frequency response

3.1.3. Localization

As previously discussed in Section 2.3.2, localization is conducted through a cable-specific adaption to the FMECA model. Because of the tree-and-branch structure of our HFC networks, additional topology or location information is required to determine the effects. Fortunately, water-soaked cable localization is straightforward. Cables with water ingress are almost always outside, so that limits the localization to the drop network (Figure 11) or distribution cable such as feeders and trunk. The simplest, most effect method is to determine if the water signature is isolated to a single location, or multiple. There are a number of basic pattern-matching routines that can accommodate this type of localization. There are some examples where incorrect localization can occur with this method, but these should be considered edge-cases.

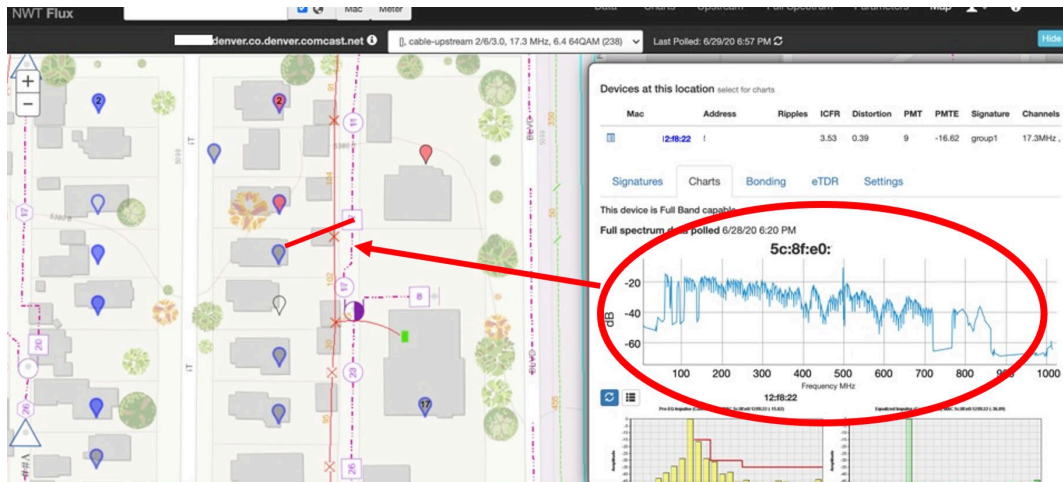


Figure 11 – Water signature localized to a single location indicating damaged drop cable

3.1.4. Criticality

One of the most intriguing aspects of our water damaged cable use case is that it has a predictable decay model. In the case of RG drop cables, the water damage always gets worse with little-to-no improvement over time, other than the transient effects of temperature. In the case of water-soaked RG6 drop cable, water can migrate or drain, but mostly it pools within the dielectric. Once the dielectric has been compromised, a chain of decay becomes inevitable. Figure 12 illustrates this degradation process.

As we explained earlier, the criticality of an impairment is estimated by the number of customers impacted, the severity of the impact (in terms of impacted services, potential service impacts, etc.), and the duration of the impact. While DOCSIS provides strong resiliency in RF signals, severity is estimated proactively in terms of potential impact to service. The potential impact is estimated through a model describing the degradation or decay path that the wet cable will take if left to degrade.

Again, see Figure 12, which describes the degradation path of a coax cable, such as an exposed drop. At first, the cable is protected from the elements. After time, that protection weakens, and the elements can begin to enter the shield. After more time, water gains access to the braiding of the shield, and fills the gaps between the wires. After some amount of freezing and thawing, the insulation begins to fail, and the distance between the shield and center conductor can change, leading to changes in the dielectric constant and multiple impedance changes. After more time, the water can get to the center conductor, and corrosion can occur throughout this degradation process. As this degradation progresses, PNM telemetry can indicate worsening degradation by showing a stronger water signature in the bin data.

Field observation and testing of recovered cables provides various snapshots of this degradation path, so we know it happens, and we can compare the impairment patterns to the failure modes found in the field to align our telemetry to the likely network failure modes. All this allows us to know that PNM is meeting its intent.

With additional work, we could develop prediction models to estimate the time to the next level of failure. But we already know enough. Once we find a failure mode, and we can estimate its severity in the future, we can determine its criticality, and therefore know the problem is worth addressing well before the cost of failure is inflicted.

Through PNM, we delight customers and save on operations costs. With water-soaked cables, the motivation is clear, as is the link from impairment to fault to failure.




Figure 12 - Example degradation of cable with exposure and water penetration

3.1.5. Effects

Empirical performance measurements can often be a good lead indicator of experience, which is subjective. In early field trials of our water detection, the effects were easily validated with before-and-after signal analysis (Figure 13). While the criticality model discussed in section 3.1.4 is a work-in-progress, and we intend to identify other degradation patterns to extend our knowledge, the merits of fixing these types of problems are agreeable. Starting with the worst problems first is obvious, and eventually we'll get ahead of the critical failures.

"We ran on this high-variance water in cable address over the weekend. The customer is internet only for the past 4 years and no trouble calls in history. This customer's service has definitely been suffering."

Tech Notes

- Yes, the recommended fix was correct
- Drop had had water damage for a very long time. Corrosion on center conductor and powder when fitting cut off.
- Drop had quite a bit of age to it, and I can tell it had come down and been driven over multiple times in the past. A lot of flattened sections along the line.
- Also, there was a return noise filter at tap for noise.
- Customer had, for quite some time, many intermittent issues with service dropping out for long periods of time.
- Home performance test was failing and is now passing.

- We contacted customer to make arrangements for access.




Figure 13 – Field trial, water wave before and after repair

3.1.6. *Benefits*

In the water damaged cable field trials, the benefits have been modeled to help provide a way to quantify and rationalize the value of proactive vs. reactive repairs. To support the field trial, a year's worth of drop replacements were analyzed. The analysis yielded interesting results in the repair time of typical drop replacement activity.

On average, by removing the need to troubleshoot and diagnose, repairing these water damaged drop cables resulted in 1 hour of improved operational efficiency. When multiplied by the number of instances, the operational benefits are compelling. However, there are additional benefits that are realized:

- Our proactivity delights our customers rather than forcing them to call us
- Improves network reliability in a measurable way
- Removes friction from customers that may be silent, but unhappy with their service

4. Conclusion

Most cable operators have network monitoring and management systems including PNM. In many cases, we've become excellent stewards of our network performance and aim to provide the best possible customer experience. FMECA provides an easy to understand and intuitive methodology for maintaining a consistent knowledgebase of failure modes, severity assessment, and repair prioritization.

By extending FMECA with DOCSIS PNM, we have access to network telemetry which can indicate an impairment. The impairment signature indicates the type of fault in the RF signal, which we can link back to a network failure of a particular type. We trace impairment to fault to failure mode in the network, and that helps us know what to look for and where, and what to do to fix it. When coupled with cable's PNM, FMECA helps us transition from "find failure and fix it" to "anticipate failure and prevent it".

Abbreviations

AI	artificial intelligence
ASHRM	American Society for health care Risk Management
CM	cable modem
CMTS	cable modem termination system
dB	decibel
DM	demand maintenance
DOCSIS	Data-Over-Cable Service Interface Specifications
FEC	forward error correction
FFT	fast Fourier transformation
FMECA	failure mode, effects and criticality analysis
HFC	hybrid fiber-coaxial
IQ	in-phase and quadrature
MER	modulation error ratio
ML	machine learning
N	number of occurrences
PNM	preventative and proactive maintenance
PM	proactive network maintenance
RF	radio frequency
RG	residential grade
S	services
SCTE	Society of Cable Telecommunications Engineers
T	time

Bibliography & References

"Proactive Network Maintenance Using Fast, Accurate Anomaly Localization and Classification on 1-D Data Series." *IEEE Xplore*. Web. 12 Aug. 2022.

SCHOOMAKER, PETER J. "Failure Modes, Effects and Criticality Analysis ... - United States Army." United States Secretary of the Army. Web. 12 Aug. 2022.

Summy, Elizabeth. *WHITE PAPER White Paper - ASHRM*. American Hospital Association, <https://www.ashrm.org/sites/default/files/ashrm/FMEAwhitepaper.pdf>.

Rupe, Jason. "Interview on Failure Mode Analysis." Online interview. 1 June 2022.

Keeping The Lights On

Protecting Your Remote Active Fiber Elements Against Unplanned Outages

A Technical Paper prepared for SCTE by

Rob Anderson

Director of Product Management and Strategic Technologies

EnerSys

3767 Alpha Way, Bellingham, WA

360-392-2293

rob.anderson@enersys.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. PON is Growing Brighter.....	4
2.1. Strand-Mount R-OLT.....	6
2.2. About RDOF.....	6
3. Anatomy of a PON.....	7
3.1. PON and 10G.....	8
4. Powering PON.....	9
4.1. Powering HFC vs. PON.....	10
4.1.1. PON Example 1.....	11
4.1.2. PON Example 2.....	12
4.1.3. PON Example 3.....	12
4.2. RDOF PON: Accommodating Rural Fiber Deployments.....	13
4.3. Sizing the UPS for Small Loads.....	14
4.4. UPS Status Monitoring for R-OLT Powering Applications.....	16
4.5. UPS Status Monitoring for EDFA Powering Applications.....	18
5. Outdoor Powering Hazards.....	21
5.1. Ferroresonant Overview.....	22
6. Conclusion.....	25
Abbreviations.....	26
Bibliography & References.....	28

List of Figures

Title	Page Number
Figure 1 – Broadband Operators' Poll (Omedia, 2021).....	5
Figure 2 – Broadband Equipment Providers Poll (Omedia, 2021).....	5
Figure 3 – Example Strand-Mount R-OLT.....	6
Figure 4 – Sample Broadband Network Diagram.....	7
Figure 5 – Cox Communications FTTX Architecture.....	8
Figure 6 – PON Beyond 10G, CableLabs.....	9
Figure 7 – PON Powering Examples.....	11
Figure 8 – Typical Broadband UPS.....	11
Figure 9 – High Capacity PON Enclosure.....	12
Figure 10 – Small Broadband UPS for Dedicated R-OLT Powering.....	13
Figure 11 – Sample EDFA (Optical Amplifier).....	14
Figure 12 – Ferroresonant 15 Ampere UPS Efficiency Graph.....	15
Figure 13 – Example Pluggable Optical SFP Compatible Module.....	16
Figure 14 – Example of Pluggable ONU for UPS Monitoring.....	17
Figure 15 – EDFA UPS Monitoring.....	18
Figure 16 – UPS Event Log.....	20
Figure 17 – Ferroresonant UPS Transformer.....	22
Figure 18 – Sine Wave vs. Quasi-Square Wave Energy.....	23
Figure 19 – Ferro Output Fold Back Overcurrent Protection.....	24

List of Tables

Title	Page Number
Table 1 – HFC & PON Powering Differences	10
Table 2 – UPS Event Log Review.....	20

1. Introduction

Cable operators worldwide are deploying passive optical networks (PON), realizing high-capacity, low latency performance with low maintenance costs. Although PONs are essentially passive, current architectures make use of outside plant (OSP) active optical components. Strand- or pedestal-mounted remote optical line terminals (R-OLT) can replace rack mounted OLTs of a few years ago. When fiber runs extend beyond optical transceivers' effective range, midspan erbium doped fiber amplifiers (EDFA) boost optical levels and extend fiber reach. Evolving PON architectures are designed to support higher speeds and extended service areas. Next-generation PON implementations will likely include high speed coherent optic links from headends to new OSP aggregation nodes, enabling even faster services to extend deeper into the network. Each active optical component in our PON architecture requires reliable, uninterrupted power to keep the network running.

This paper reviews aspects of powering the OSP active components of modern PONs including power levels, utility grid isolation, utility backup duration and power system status monitoring. Unique requirements for powering expanded rural fiber networks are discussed as well as considerations for addressing an aging utility grid and the impact of extended outages. Roles and methods for power monitoring and remote management are discussed as well as the need for predictive and restorative maintenance.

2. PON is Growing Brighter

Today's consumer broadband networks need high-speed symmetrical bandwidth and low latency to deliver superior customer experiences for applications like augmented reality/virtual reality, over-the-top streaming video, cloud gaming, videoconferencing, 5G backhaul and other peer-to-peer networking services.

All roads lead to fiber. Fiber optic cable as a data transport medium has no equal with respect to performance and total cost of ownership (TCO). However, DOCSIS over coaxial cable has outpaced consumer broadband demands exceptionally well. Most experts believe DOCSIS still has many years of life before running short on performance. So, there's no immediate or urgent need to replace the majority of installed coax with fiber. The low TCO and almost unlimited bandwidth of fiber are well understood by operators. As a result, operators have started migrating greenfield and network extensions to all fiber based. In 2021, global research firm Omdia polled broadband operators, asking which type of network upgrades do their companies plan to carry out by fall 2026. The poll results in Figure 1 clearly indicate FTTH deployments as the highest priority.




Figure 1 – Broadband Operators’ Poll (Omedia, 2021)

In a separate poll, broadband equipment providers were asked about access equipment revenue forecasts for North America. The poll results shown in Figure 2 clearly identify PON equipment sales leading the ranks and increasing dramatically through 2026.




Figure 2 – Broadband Equipment Providers Poll (Omedia, 2021)

There’s no doubt that investment in cable broadband access solutions will continue for many years. The strategy for most cable operators is not replacing coaxial cable with FTTH. The strategy for most

operators involves when and where to add FTTH. Another indicator of the trend towards FTTH from the Fiber Broadband Association reports that at the end of 2021, 43% of U.S. households and 60% of Canadian households had access to fiber services. [1]

2.1. Strand-Mount R-OLT

Silicon advancements in recent years have enabled R-OLT modules to fit within strand-mount optical node enclosures or housings. Distributed access architecture (DAA) node housings can be used for this application. The modules and node housing are designed such that a R-OLT or remote physical layer (R-PHY) module could be installed. Some equipment providers advertise the ability for their housing to support both R-OLT and R-PHY modules simultaneously, servicing PON and quadrature amplitude modulation (QAM) customers from the same node. This application has been nicknamed DAA-OLT. An example of a strand mount R-OLT module and housing is shown here.




Figure 3 – Example Strand-Mount R-OLT

Two large North American (NA) operators reported in June 2022 that strand-mount R-OLT equipment will be a primary fiber-to-the-X (FTTX) upgrade tool for the next few years. One R-OLT equipment vendor has reported that revenues for DAA-PON modules for NA increased from about \$300,000 in 2020 to over \$10 million in 2021. It looks like PON is here to stay. [2]

2.2. About RDOF

A major contributing factor to accelerated broadband fiber builds is the Rural Digital Opportunity Fund (RDOF). RDOF will provide \$20.4 billion in funding over a 10-year period to support broadband networks in rural communities across the US. Among the top RDOF awards, the U.S. based operator Charter qualified to receive \$1.2 billion in federal government support for deployment plans spanning some 1.1 million locations. Tom Rutledge, Charter's chairman and CEO, said that the company has RDOF work underway in all 24 states where it won bids in phase I of the RDOF auction. All told, Charter plans to build fiber and gigabit services to more than 1 million rural, unserved locations. Through RDOF, the company will add more than 100,000 miles of new network infrastructure to its existing 800,000 miles of infrastructure in the coming years. [3]

The RDOF program does not dictate specific technologies for new broadband services. However, the FCC has proposed creating performance tiers for RDOF funding. High bandwidth, low latency fiber solutions will be most favored.

3. Anatomy of a PON

During a recent CableLabs conference, Stephanie Mitchko-Beale, Charter CTO, stated “the speed of change keeps me up at night” [4]. This sentiment is likely shared by anyone responsible for the accelerated FTTX deployments in recent years. FTTX and PON are simple in principle but the popular cliché “the devil is in the details” was never truer than with PON deployments. For more details let’s review the sample broadband network block diagram in Figure 4.




Figure 4 – Sample Broadband Network Diagram

Network architectures vary with operators, geography and many other factors. This block diagram does not represent a specific installation but only a few of the more common network building blocks. Moving from the headend on the left through the access network on the right, the network elements shown in this diagram include:

- A headend based converged interconnect network (CIN) delivers transport and aggregation functions for diverse networking elements such as DOCSIS, PON OLT, DAA, 4G/5G, and so on. The CIN is the glue enabling various disparate elements to interact and interoperate.
- Fiber distribution is physically routed from the headend to hubsites, nodes and OLTs. Some operators refer to this fiber segment as optical trunk fiber. Today, this fiber would likely carry 10 Gbps Ethernet which is fast replacing the CWDM and DWDM analog fiber transmissions of the past.
- DAA node service groups can support QAM video and data using DOCSIS. R-OLT modules can be co-installed in node housings with DAA modules or installed nearby to support PON.

- In densely populated areas, multiple shelf mounted R-OLTs may be housed in curbside equipment enclosures along with fiber distribution, powering equipment and batteries.
- Remote OLT modules installed in strand or pedestal mounted node enclosures are an alternative to cabinet (rack or shelf) mounted OLTs.

Let's review a specific FTTX example presented at the 2021 Cable-Tec Expo by Cox Communications. [5].

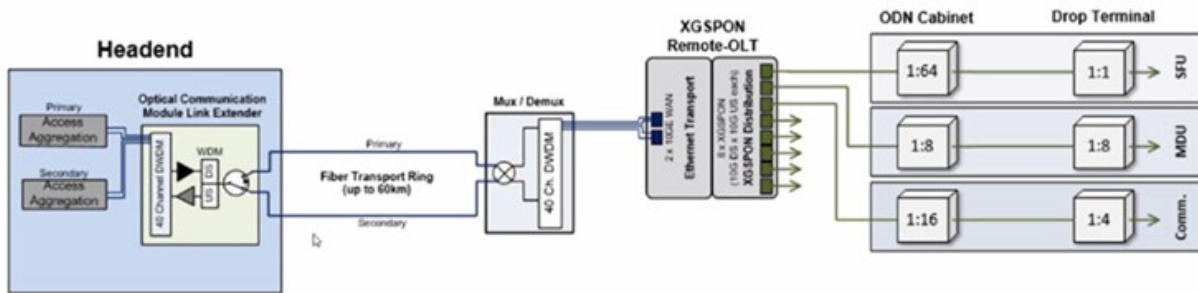


Figure 5 – Cox Communications FTTX Architecture

A few noteworthy details from this FTTX architecture include:

- The R-OLT is the only active element in the access network. When practical, the OLT is located near the HFC plant to utilize coax for both power and power supply telemetry.
- In greenfield installations where there is no nearby coax to power the OLT, a curbside enclosure may be used to house multiple OLTs, fiber distribution equipment and a backup power system for the OLTs.
- Headend-based EDFAs facilitated a fiber transport ring of up to 60 km to the R-OLT (either strand-mount or cabinet based). With a 1:64 split, the reach from the R-OLT to the premise is up to 20 km. From ODN cabinet to premise is 8 km maximum reach.
- The Cox network utilized XGS-PON optical signaling.

3.1. PON and 10G

The CableLabs 10G Platform promotes 10 Gbps symmetrical services using any number of technologies. Two predominant PON protocols are in use today: XGS-PON and 10G-EPON. Both can directly support the full intent of the 10G Platform. XGS-PON is an ITU standard with roots in the telecommunications industry while 10G-EPON is an IEEE standard more common in data centers. Both protocols offer approximately 10 Gbps symmetrical service. When it comes to powering PON elements, it doesn't matter which optical communications standard is used.

With change being inevitable, PON architectures beyond 10 Gbps are anticipated, and in some cases already here. Powering new and upcoming PON devices should be considered to assure that today's powering approach will support future upgrades. CableLabs has presented a next generation PON

architecture utilizing a new OSP active component: an optical aggregation node. This next generation PON is shown in Figure 6.

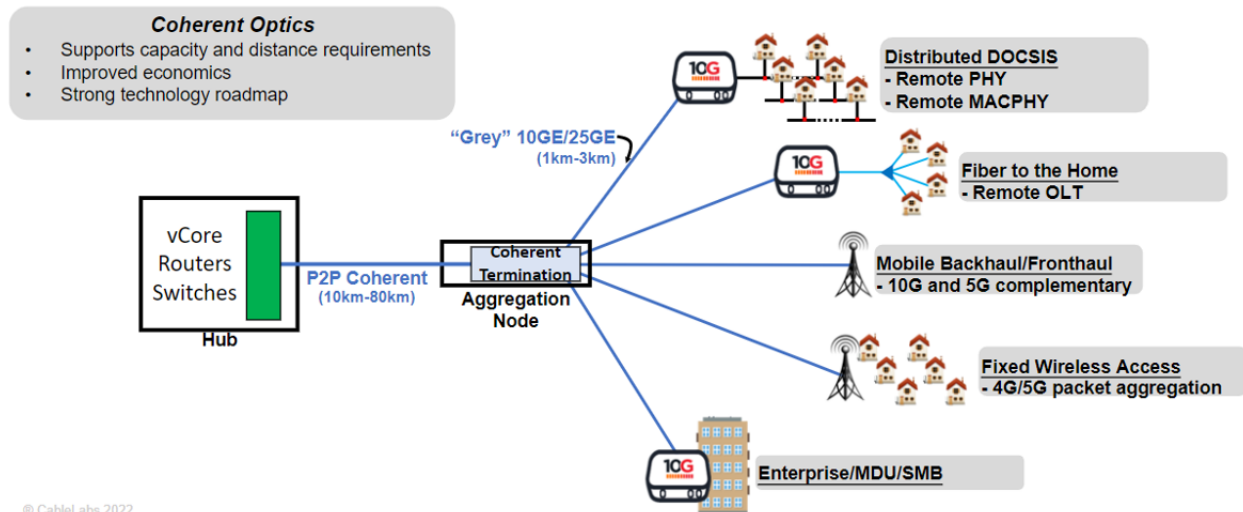


Figure 6 – PON Beyond 10G, CableLabs

This near-future PON enhancement contains a new active component, an aggregation node (AN). In this architecture, new high speed fiber communications are pushed deeper into the network by use of high-speed coherent optics between the hub and the new AN. This link could be 50 Gbps, 100 Gbps or faster. [6] The AN will act as an optical Ethernet switch providing lower speed (10 Gbps for example) optical links to downline elements such as DAA nodes and OLTs. The result is higher capacity services deeper into the network and closer to subscribers. From our PON powering vantage point, we've simply added another active element to the mix. Once AN elements begin to be deployed, power requirements will need to be analyzed. Given the current state of coherent pluggable modules combined with known high speed switching elements, we anticipate that an AN may require approximately 200 W in a fully configured state. This estimate remains highly speculative until actual networks using AN devices are deployed.

In general, backup power is fiber protocol agnostic. Optical wavelength, communication protocols (10G-EPON vs. XGS-PON), distribution methods and so on, don't affect power. We need to understand the location and power needs of the few but critical active components. Powering these active elements are discussed in the next section.

4. Powering PON

Powering broadband PON is similar, and often simpler, than powering traditional HFC networks. Optical actives can be powered from the coax network when an HFC plant segment is nearby. More often, PON equipment needs a dedicated uninterruptable power supply (UPS) and batteries for uninterrupted power. A side-by-side review of HFC and PON powering should illuminate noteworthy differences.

Table 1 – HFC & PON Powering Differences

Powering Consideration	HFC	PON
Powered equipment	Optical nodes, amplifiers, line extenders	OLTs, EDFAs and ANs (future)
UPS to equipment ratio	One-to-many	One-to-one (typically)
Equipment voltage	90 VAC	90 VAC (strand-mount R-OLTs); 48 VDC (shelf mount R-OLTs)
Distance from UPS to equipment	From co-located to 1,000s of feet via powered coax.	Co-located (1)
Total power required	HFC network designs often utilize full UPS capacity. Typical broadband UPS are rated from 1350 W to 1620 W	Powering a single OLT may require ~140 W.
Utility grid protection	Required	Required
Backup power duration	4 to 72 hours (2)	6 to 72 hours
UPS status monitoring	DOCSIS	Fiber

(1) OLTs may be coax powered in situations where HFC coax is nearby.

(2) The CPUC 72-hour telecommunication backup requirement is an extreme case. Operator backup time policies typically vary from two to eight hours.

Let's review some of the important differences between powering HFC and PON.

4.1. Powering HFC vs. PON

Traditional HFC networks include optical nodes, amplifiers, and line extenders. These active components are powered from a UPS through the coax. The number of actives that can be powered from a single UPS depends on a combination of the equipment's total static load and the combined coax span lengths between powered devices. Longer coax spans have higher loop resistance, creating greater joule losses (i.e., power loss due to resistance in coaxial cable, electrical energy is converted to heat). Cable resistance and current flow through the coax results in a voltage drop across each coax span and a lower end-of-line voltage. In most cases, cable plants will run short of power or end-of-line voltage before they run out of devices to be powered. Also, operators' internal design policies guide network designs to limit UPS loads to a percentage of the UPS' rated capacity, often in the 80% range. [7]

Powering PON equipment takes a slightly different approach than powering HFC. The broadband network in Figure 7 shows three different PON segments, each with unique powering schemes. A review of these PON segments and powering options will prove illuminating.

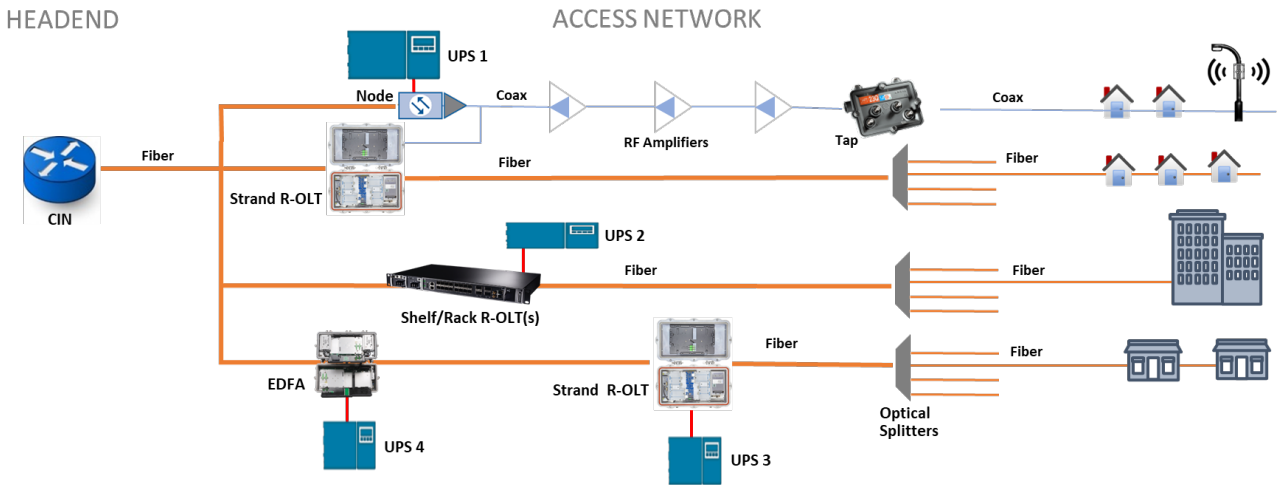


Figure 7 – PON Powering Examples

4.1.1. PON Example 1

The network diagram shows UPS 1 powering an optical node, downline amplifiers and a R-OLT. In this scenario the PON R-OLT is installed in the vicinity of a powered coax segment and can be powered from that coax. One operator's rule-of-thumb for coax powered R-OLTs is that if powered coax is available within one-half mile of the R-OLT location, they will extend coax to power the R-OLT. Beyond that distance, it's more cost effective to add a dedicated UPS to power the PON equipment. A typical R-OLT of this style (see Figure 3) would require up to 140 W to power and could be configured to support up to 512 FTTX customers through downline optical splitters. A typical broadband UPS used to power nodes, amplifiers and our R-OLT is shown in Figure 8. These UPS systems may be pole or ground mounted and typically consist of an enclosure, the UPS and one or more strings of three 12-volt batteries. If UPS 1 has sufficient power and battery capacity to accommodate the added R-OLT load, no additional powering considerations are needed.



Figure 8 – Typical Broadband UPS

4.1.2. PON Example 2

In service areas with a higher population density, PON architectures may dictate supporting many customers from a single fiber distribution point. One option is to utilize outdoor, curbside enclosures housing multiple R-OLTs. UPS 2 in our network diagram represents this scenario. An example of this type of enclosure is shown in Figure 9.




Figure 9 – High Capacity PON Enclosure

The multiple shelf or rack mount R-OLTs in this example will service far more subscribers than our earlier strand mount example. These cabinets would include fiber distribution as well as backup power for the R-OLTs. Shelf mounted OLTs are typically powered from 48 VDC. The cabinet would contain a 48 VDC rectifier system, likely in an N+1 redundant configuration, and batteries. The cabinet is connected directly to utility power.

Cabinet-based OLT systems can be scaled to specific applications. One North American operator has standardized on a curbside cabinet capable of servicing up to 8000 homes. This cabinet contains multiple OLTs, fiber distribution panels and a 48 VDC rectifier system with batteries to provide six hours of backup power during utility outages.

4.1.3. PON Example 3

The third PON example from Figure 7 represents a strand-mount R-OLT powered from UPS 3. This differs from the UPS 1 example in that there is no nearby coax for R-OLT power. For this application a UPS is dedicated to a single R-OLT or potentially multiple co-located R-OLTs. Figure 10 shows a typical broadband UPS sized appropriately for this application.



Figure 10 – Small Broadband UPS for Dedicated R-OLT Powering

This broadband UPS is similar to the UPS from example 1 with a few exceptions that are reviewed after a note about RDOF PON.

4.2. RDOF PON: Accommodating Rural Fiber Deployments

As previously discussed, RDOF is bringing broadband, specifically fiber broadband, to many un-served and under-served rural locations in the United States. Does rural fiber deployment differ from its urban or suburban counterparts? The answer is yes. There are two differences to consider. First, rural areas have lower population density than urban or suburban, which goes without saying. Rural areas will likely never need the service group density available from dedicated enclosures housing multiple high-capacity R-OLTs. Also, rural areas will most often be considered greenfield with no incumbent fiber or coax to leverage. This means that PON example 3 from Figure 7 will likely account for most of our rural fiber installations.

The second consideration with rural fiber deployments deals with fiber span distances. One operator mentioned that their rural deployments are seeing fiber lengths up to 160 km from headend to end users. To facilitate these distances the optical signal must be amplified. For this operator, the solution is to install optical amplifiers or EDFAs mid-span in the fiber run. An EDFA is shown in Figure 7 under PON example 3. Figure 11 is a typical EDFA used for this type of application.



Figure 11 – Sample EDFA (Optical Amplifier)

Like the R-OLT example, this EDFA is housed in an outdoor hardened aluminum housing and can be strand, pedestal or underground vault mounted. This EDFA requires about the same amount of power as our R-OLT and would be powered by a similarly sized UPS system. The following sections discuss UPS details that specifically apply to R-OLT and EDFA powering.

4.3. Sizing the UPS for Small Loads

Using coax to transport both power and RF has been a CATV practice for decades. This fundamental approach has not changed much over the years. A single stage UPS using a ferroresonant transformer is extremely robust and provides exceptional protection for active network elements from outdoor utility grid hazards. A later section discusses more about ferroresonant technology.

One important characteristic of the ferroresonant UPS is that it operates most efficiently when loaded to its maximum rating. Most traditional broadband UPS systems are rated for use between 1350 W to 1620 W. That ends up being 15 A to 18 A at 90 V. Using a simple Ohm's Law calculation: $\text{current} = \text{power} / \text{voltage}$ (we'll ignore power factor for now to keep things simple). In our PON example 3 from Figure 7, the R-OLT required about 140 W or about 1.5 A at 90 V. The efficiency for several similar 15 A ferroresonant UPS models is shown in Figure 12.




Figure 12 – Ferroresonant 15 Ampere UPS Efficiency Graph

If this 15 A UPS were dedicated to powering a single R-OLT drawing 1.5 A, the UPS efficiency would be under 65% as shown. The lost efficiency would be energy converted to heat in the ferroresonant transformer. This energy loss would occur both under utility power (lost energy means higher utility cost from consumed kW/h) and during battery backup (resulting in wasted battery energy and less available backup time).

The obvious solution is to use a smaller UPS system for smaller loads, like our R-OLT. Ferroresonant UPS systems come in many ratings down to 5 A or less. Additional benefits of using a smaller, dedicated UPS include:

- Smaller broadband UPS systems use a single 12 V battery instead of a 36 V string of three 12 V batteries.
- The UPS is smaller, lighter and can be mounted in a smaller enclosure designed for the UPS and a single battery.
- The system cost will be lower due to the size and battery configuration.

The available backup runtime during utility outages varies greatly with the type and condition of the battery used in the UPS system. For example, if our R-OLT UPS system uses a single case size 31, high quality TPPL (thin plate pure lead) battery designed for outdoor broadband applications, the UPS system would provide over six hours of backup runtime for our R-OLT during a utility power outage.

4.4. UPS Status Monitoring for R-OLT Powering Applications

Monitoring broadband UPS systems provides several benefits. The most common benefit is that operators receive real-time status. During extended utility outages it's critical that operators have advance warning with sufficient time to take action prior to battery depletion, resulting in dropped loads and customers losing service. UPS monitoring also notifies operators of equipment concerns. Notifications may indicate that immediate attention is required or issues that should be addressed during an upcoming scheduled maintenance visit.

Traditional broadband UPS systems incorporate a purpose-built DOCSIS modem, called a transponder, for status communications between the UPS and the monitoring software. UPS monitoring over DOCSIS has been used for over 20 years. Operators have well defined processes in place for installing, provisioning and operating DOCSIS modems for OSP UPS monitoring.

For PONs where DOCSIS is not available, an alternative UPS monitoring approach is needed. To address this need, modern broadband UPS systems include a SFP interface (small form-factor pluggable). SFP is a standard interface compatible with pluggable, optical transceiver modules from multiple vendors. A typical pluggable SFP module is shown in Figure 13.



Figure 13 – Example Pluggable Optical SFP Compatible Module

For status monitoring of a UPS system that is powering a R-OLT, a specialized SFP module would be installed in the UPS's SFP socket. This SFP would function as an optical network unit (ONU) or an optical network terminal (ONT), supporting the specific optical protocol used by the R-OLT. An ONT and an ONU are the same things. ONT and ONU both refer to the consumer end equipment in an optical fiber (FTTX) communication link. ONT is an ITU-T term, whereas ONU is an IEEE term. The term ONT would be used if the R-OLT is configured for ITU-T protocols including GPON and XGS-PON. The term ONU would be used when the R-OLT is configured to IEEE protocols including EPON and 10G-EPON. For simplicity, we will use the term ONU going forward to mean either ONU or ONT.

The pluggable SFP ONU would be provisioned and managed by the R-OLT in the same way a home-based ONU operates. This configuration is represented in Figure 14.




Figure 14 – Example of Pluggable ONU for UPS Monitoring

In this example the R-OLT provisions and manages all ONUs in this PON segment, including the ONU installed in the UPS.

This UPS monitoring approach seems straightforward. However, there is more to this story. UPS status monitoring must be routed to the operator's OSP power supply monitoring software. The operator would likely associate the power supply MAC and assigned IP address to the monitoring software in a router or Layer 3 switch in their back office. Also, the UPS needs to be configured with the IP address of the monitoring software to enable the UPS to send SNMP traps or alerts. The UPS would often be configured with a list of IP addresses of destinations where SNMP notifications are to be sent. Other UPS parameters may need to be configured as well. In a traditional HFC network where UPS systems use DOCSIS modems for status communications, the DOCSIS configuration files are customized to configure UPS systems, including trap destination addresses. Unfortunately, PONs have no equivalent to a DOCSIS configuration file, at least not one that can be used for UPS configuration.

At first glance, a CableLabs specification called DOCSIS Provisioning of EPON (DPoE) sounds like a promising option for UPS configuration. However, any hope in DPoE provisioning the UPS is short lived. DPoE enables DOCSIS configurations and service tiers to be passed on to the PON. For example, a PON customer could subscribe to a 100 Mbps downstream and 10 Mbps upstream service tier. DPoE can use existing DOCSIS services to configure the PON to support this service tier for this subscriber. The R-OLT and ONU supporting DPoE work together to implement this type of DOCSIS service. DPoE has no mechanism for configuring equipment behind the ONU on the subscriber side of the network. Within our UPS ONU configuration, the power supply resides behind the ONU in the network, much like a subscriber's computer resides behind a home-based ONU.

To date, two methods have been used to configure the UPS for monitoring and alarm reporting in a PON. First, some UPS systems support the use of a configuration file containing TLV (type-length-value) style information, similar to DOCSIS configuration file TLVs. When the UPS is provisioned and receives an IP address from the network DHCP server, the provisioning process typically provides the IP addresses for a TOD (time-of-day) server and a TFTP (Trivial File Transfer Protocol) server. The UPS will search the root directory of the TFTP server for a specific configuration file to be used for UPS configuration. This approach is similar to how a DOCSIS configuration file is used to configure cable modems.

The second method available for configuring UPS parameters is by using vendor specific DHCP options. A DHCP offer contains a variety of information fields beyond just an IP address. The standards for DHCP define various options that can be part of the DHCP offer, which the DHCP server can provide for all or selected requesting devices. For DHCP on IPv4, this information is provided using DHCP option 125. For DHCP on IPv6, this information is provided using DHCPv6 option 17. Options can specifically be used to configure SNMP general settings, access settings, and SNMP notification settings (issuing traps to specific destinations).

There are currently no industry standards defining either of these provisioning schemes. They are proprietary per individual vendors. Operators should proceed accordingly to ensure the selected configuration approach is well defined and meets their specific requirements.

4.5. UPS Status Monitoring for EDFA Powering Applications

In a prior section we discussed that powering EDFAs and powering R-OLTs are very similar. Status monitoring the UPS that is powering the EDFA is less straightforward than monitoring the R-OLT UPS system. Our monitoring solution using a pluggable ONU doesn't work without the corresponding R-OLT. If we're powering an EDFA somewhere midspan in a fiber-run, there's likely no nearby R-OLT or coax to for implementing familiar status monitoring methods.




Figure 15 – EDFA UPS Monitoring

Operators have tried multiple approaches to solve this quandary. Unfortunately, there is no one size fits all solution. A few options that have been explored so far are described here.

Dedicate DWDM or CWDM wavelengths. DWDM or CWDM (depending on the fiber traffic) transmit and receive wavelengths could be set aside for UPS traffic. A DWDM/CWDM SFP would be installed in

the UPS and the appropriate demux, filter and splicing mechanism would be installed near the EDFA and UPS to facilitate the fiber jumper physical connection.

Dedicated dark fiber. A dedicated fiber or fiber pair could be assigned for UPS traffic. A splice tray, patch panel or other appropriate mechanism to patch into the fiber would be installed near the EDFA and UPS. A SFP would be installed in the UPS using the desired optical format.

Cellular data modem. A cellular modem could be installed in the UPS system to implement an out-of-band (OOB) communication interface. A copper Ethernet SFP installed in the UPS would connect to the cellular modem via an RJ-45 (copper) Ethernet cable. A cellular antenna would be mounted on the outside of the UPS enclosure using weather tight techniques. The modem would be powered from a power converter connected to the UPS output, providing the modem's required voltage.

EDFA proprietary management port. Some EDFA manufacturers support a proprietary interface for configuring and managing each EDFA in the operator's network. Data from each EDFA is routed through the fiber and back to the headend where it is available at a central management console. If the EDFA supports a physical maintenance port to access this management interface then it's feasible for a fiber jumper from the UPS to connect to the port, providing a communications pathway from the UPS to the headend. This would enable UPS traffic to be routed to the UPS management system. A proprietary UPS monitoring system of this type was developed between a remote hub vendor and a UPS vendor and initially deployed in 2015.

Don't monitor the UPS. If none of the current options for EDFA UPS monitoring are desirable, operators may elect to leave UPS systems unmonitored. This could be a conscious choice or may be the result of a deferred decision due to lack of an ideal monitoring option. Either way, the results will be the same. Operators will have no advance warning of utility outages that are at risk of exceeding the backup runtime of the UPS system. This risk can be mitigated by installing additional batteries to extend runtime, combined with regular maintenance visits to maximize the probability that the UPS system will perform as anticipated during utility outages. Another tool operators should make use of, especially for unmonitored UPS systems, is the utility event history typically stored in nonvolatile memory in the UPS. Figure 16 shows the history log from a typical broadband UPS.

Standby Events - XM3.1-918-HP

XM3.1-918-HP
Standby Events

Alarms Battery English Log In...

Overview

- Hardware
- Network
- Management
- History
 - Power Supply Events
 - Power Supply Configuration
 - Battery Events
 - Standby Events**
 - Modem Log
 - Network System Log
 - Network Event Log
 - Firewall Log
 - Alarm Log
- Tools

Displayed time zone: UTC-7:00

<<Prev 1 2

ID	Time	Event	Outage Time	Down Time
448	2022-02-30 07:01:14	Utility outage	0h:01m:03s	0h:00m:00s
447	2022-02-29 07:00:31	Utility outage	0h:00m:52s	0h:00m:00s
446	2022-02-28 06:59:42	Utility outage	0h:00m:41s	0h:00m:00s
445	2022-02-27 07:02:19	Utility outage	0h:01m:07s	0h:00m:00s
444	2022-02-26 07:01:51	Utility outage	0h:00m:56s	0h:00m:00s
443	2022-02-18 03:32:51	Remote self-test	0h:12m:02s	0h:00m:00s
442	2022-01-18 03:21:06	Remote self-test	0h:09m:12s	0h:00m:00s
441	2021-12-28 18:26:42	Utility outage	3h:42m:09s	0h:00m:00s
440	2021-12-18 03:52:22	Remote self-test	0h:11m:31s	0h:00m:00s
439	2021-11-21 04:46:11	Utility outage	9h:03m:56s	3h:12m:10s
438	2021-11-04 02:31:20	Utility outage	2h:08m:15s	0h:00m:00s

<<Prev 1 2

Figure 16 – UPS Event Log

This type of utility power history is typically available through the UPS monitoring system. However, for unmonitored UPS systems, a log of this type is only available on-site and may be the only indicator of the utility power activity at the UPS location. Operators should schedule regular site visits to retrieve and analyze these utility events. As an example of the value of this type of data, the entries from the log in Figure 16 are reviewed in Table 2.

Table 2 – UPS Event Log Review

ID	Event	Description & Implication
444-448	Utility outage	Short duration utility outages occurring at approximately the same time each day. This may indicate an external event affecting the utility grid.
443, 442, 440	Remote self-test	UPS test cycles appear to be initiated by the UPS monitoring system each month at approximately the same time of day.
441	Utility outage	A utility outage lasting 3h:42m. The UPS has sufficient battery capacity to support the network during this event.
439	Utility outage	A utility outage lasted 9h:03m. The UPS powered the network for approximately six hours before depleted batteries forced dropping the load.
438	Utility outage	A utility outage lasting 2h:08m. The UPS has sufficient battery capacity to support the network during this event.

If this UPS were remotely monitored, a network outage and service interruption could have been avoided (event no. 439) through advance notification of the extended outage. Also, repeated, short direction

outages (event no. 444-448) would be identified for investigation. During one example of these types of short duration regular outages, an investigation identified a nearby factory's daily start-up processes were the cause of utility transients that triggered the UPS to switch into standby operation.

Our list of EDFA UPS monitoring options did not mention that UPS data must be received in the headend, converted to IP traffic over Ethernet, and routed to the operator's UPS monitoring software. This step must occur regardless of the data transport method.

The viability of implementing any of these options must be evaluated by operators. In one recent example, an operator determined that the only viable options for their network was using dedicated DWDM wavelengths. They further determined that the negative impact from dropped customers outweighed the cost of implementing dedicated wavelength style monitoring. In another example the operator chose to leave the UPS unmonitored, mitigating risk by stepping up scheduled maintenance to unmonitored locations.

5. Outdoor Powering Hazards

Network operators understand the harsh environmental conditions that outdoor equipment must endure. Temperature extremes, wind driven rain, snow and ice can wreak havoc on equipment not designed to endure these conditions. The electrical utility grid presents a laundry list of hazards poised to dismember the best engineered gear. Broadband UPS systems are powered directly from the outdoor utility grid. These UPS systems must be engineered to endure a litany of power grid anomalies including:

Interruption, the complete loss of voltage for thirty (30) cycles or longer.


Transients, a temporary, rapid fluctuation in the measured quantity of total power.

Surges, voltage greater than 110% above normal

Spikes, sudden voltage peaks that can reach thousands of volts

Sags, under-voltage conditions where fluctuation exceeds allowable thresholds for at least a cycle.

Brownouts, intentional or unintentional voltage drop for an extended period.



For over four decades, broadband UPS systems have provided reliable network power through every condition imaginable. Like any technology, these UPS systems have evolved over the years. However, there is one component in the broadband UPS that has remained mostly unchanged throughout the years, the ferroresonant transformer.

5.1. Ferroresonant Overview

Broadband networks have used ferroresonant (ferro) based UPS systems since the early days of CATV. Figure 17 shows a diagram of a ferro transformer.




Figure 17 – Ferroresonant UPS Transformer

What makes the ferro transformer ideal for outdoor powering and how has it remained mostly unchanged throughout years of technology advancements? Let's review some transformer basics to answer these questions.

Transformers, both ferro and the more common linear variety, consist of two primary components: windings and cores. Windings are (usually) copper wire, wound onto a core material such as steel. Energy is transferred by magnetic induction from one set of windings to another by means of varying magnetic flux. In the more common linear transformer, the output voltage is determined by a combination of the input voltage and the ratio of primary windings to secondary windings. In contrast, the ferro transformer uses nonlinear magnetic properties and a resonant circuit to provide a stable output voltage over a wide range of input voltage. During each AC cycle, energy is stored in the resonant circuit (an inductive-capacitive “tank circuit”) and then consumed by loads connected to other windings. The tank circuit operates in magnetic saturation, resulting in the output waveform being square or trapezoidal, often referred to as a quasi-square wave (QSW), instead of the more familiar sinusoidal waveform produced at the output of a linear transformer. The ferro's QSW results in more available energy than the sinusoidal counterpart. The QSW vs. sinusoidal energy relationship is illustrated in Figure 18.




Figure 18 – Sine Wave vs. Quasi-Square Wave Energy

The QSW per cycle voltage rise time is greater than the equivalent linear sinusoidal risetime (dV/dt has a higher slope value). Practically, this means that within every 60 Hz AC cycle, the QSW output has more energy available for a higher percentage of each 60 Hz cycle than the sine wave, assuming equivalent power factors. Also, the QSW rise time results in more time per cycle above the usable threshold voltage. The threshold voltage is the minimum voltage required for powered equipment to operate. As an example, consider an RF amplifier powered from the ferro UPS. During each AC cycle there is a period around the voltage zero-cross threshold where the UPS output voltage is too low for the amplifier's power supply to operate. Since the ferro's QSW output voltage rises faster than a linear UPS sine wave output, the amplifier can draw energy from the ferro UPS for a larger percentage of time within each AC cycle.

Greater output energy is one benefit of a ferro UPS. Other important benefit is output short circuit protection. The ferro can operate for extended periods of time at 150% of its rated capacity. When this level is exceeded, the tank circuit can no longer keep up with the amount of energy being pulled from the output. The ferro will then "fold back" or drop its output voltage to zero. A ferro fold back is shown here.




Figure 19 – Ferro Output Fold Back Overcurrent Protection

HFC networks consist of active equipment separated by spans of fiber optic and coaxial cable. Damage and wear to equipment, power-carrying coax, termination components and technician mishaps can result in electrical short circuits or “faults” in network power. A fault causing the UPS output load to exceed 150% of its rated capacity will result in a UPS output fold back condition, dropping the output voltage. When the fault is remedied, the ferro output is restored and normal network powering resumes. Neither the UPS nor the powered equipment are damaged because of the fold back condition.

Another characteristic of ferro based UPS systems is extremely high transient and noise filtering. The ferro is one of the best-known mechanisms for filtering utility line noise and transients from reaching the output and affecting or damaging network equipment. The ferro boasts an impressive 1000:1 isolation ratio. The nature of the ferro design makes it extremely immune to electrical surges. This means that a 1000-volt surge on the input winding would produce only a 1-volt difference on the output winding. This isolation offers very robust surge protection for powered network equipment.

Ferro based UPSs also implement a single stage conversion design. Voltage from the input winding is transferred to the output winding and directly to the network without any additional power conversion steps. When operating on battery backup, the input utility winding is disconnected using a relay and a separate inverter winding is energized by an inverter circuit utilizing battery current to produce energy. The UPS backup function is only required to operate when utility power is lost. In most systems the inverter operates less than 1% of the life of the UPS. This means that for over 99% of the time, the UPS is providing output power by using only the ferro transformer. Since no electronic components are in the critical powering path, the reliability of the ferro based UPS is exceptional.

Since the ferro is constructed of copper and steel, it will typically only fail from physical damage or from catastrophic electrical events such as a direct lightning strike.

6. Conclusion

OSP active PON elements are subject to extreme outdoor conditions, including utility grid anomalies. The ferroresonant based UPS has been used for many years, effectively isolating network equipment from electrical utility grid hazards. An HFC network segment, including nodes and amplifiers, is typically powered from a single UPS with the coax transporting both power and payload. OSP PON equipment, specifically R-OLTs and EDFAs, are most often powered with dedicated UPS systems. PON UPS systems must be sized for the smaller load of a single active component for the UPS to operate efficiently. PON UPS systems can be housed in physically smaller enclosures with fewer batteries and smaller components than their HFC counterparts. Operators may also decide to house R-OLTs and EDFAs with the UPS system in a single enclosure.

Status monitoring UPS systems powering R-OLTs have recently become straightforward with the use of a pluggable ONU installed in the UPS. Monitoring UPS systems powering EDFAs is not yet as cut and dry. Several monitoring options are possible. Operators need to weigh the pros and cons of each option and determine the method that works for them.

UPS status monitoring provides multiple benefits. First, operators receive advance warning of utility outages that may extend beyond usable battery runtimes. This advance warning enables timely action to keep extended utility outages from becoming service affecting events. Unmonitored UPS systems provide no advance warning, leaving operators to glean outage data from other sources or to just hope for the best. Another benefit from status monitoring includes notifications of imminent and potential service affecting conditions. This could include advance indicators that batteries have reduced capacity or that UPS service is needed. Preventive maintenance should be scheduled to correct these items. Unmonitored UPS systems cannot report these conditions, leaving operators to identify them during scheduled maintenance visits, hopefully before a utility outage becomes service affecting.

Abbreviations

10G-EPON	10 Gbps Ethernet passive optical network
A	ampere
AN	aggregation node
CATV	Originally community antenna television, now cable TV.
CEO	chief executive officer
CIN	converged interconnect network
CTO	chief technology officer
CWDM	coarse wavelength division multiplexing
DAA	distributed access architecture
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data-Over-Cable Service Interface Specifications
DPoE	DOCSIS provisioning of EPON
DWDM	dense wavelength division multiplexing
EDFA	erbium-doped fiber amplifier
EPON	Ethernet passive optical network
FCC	Federal Communications Commission
FTTH	fiber-to-the-home
FTTX	fiber-to-the-(home, curb, premise, etc.)
Gbps	gigabits per second
GPON	gigabit passive optical network
HFC	hybrid fiber/coax
Hz	hertz
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ITU	International Telecommunication Union
km	kilometer
MAC	media access control
NA	North America
Mbps	megabits per second
OLT	optical line terminal
ONT	optical network terminal
ONU	optical network unit
OOB	out-of-band
OSP	outside plant
PON	passive optical network
QAM	quadrature amplitude modulation
QSW	quasi-square wave
RDOF	Rural Digital Opportunity Fund
RF	radio frequency
R-PHY	remote physical layer
R-OLT	remote optical line terminal
SCTE	Society of Cable Telecommunications Engineers

SFP	small form-factor pluggable
SNMP	Simple Network Management Protocol
TCO	total cost of ownership
TFTP	Trivial File Transfer Protocol
TLV	type-length-value
TOD	time of day
TPPL	thin plate pure lead
UPS	uninterruptable power supply
V	volt
VAC	volts alternating current
VDC	volts direct current
W	watt
XGS-PON	10 Gbps passive optical network

Bibliography & References

- [1] "Fiber Broadband Association," [Online]. Available: <https://www.fiberbroadband.org/>.
- [2] J. Baumgartner, "Light Reading," 24 March 2022. [Online]. Available: <https://www.lightreading.com/opticalip-networks/remote-olt-sales-activity-spotlights-cables-growing-focus-on-fiber-/d/d-id/776270>.
- [3] J. Baumgartner, "Light Reading," 1 February 2021. [Online]. Available: [https://www.lightreading.com/opticalip/fttx/charter-plans-\\$5b-broadband-expansion-to-1m-locations-/d/d-id/767052](https://www.lightreading.com/opticalip/fttx/charter-plans-$5b-broadband-expansion-to-1m-locations-/d/d-id/767052).
- [4] Mitchko-Beale, Stephanie, Charter Communications, "Cable Next-Gen Technologies & Strategies, CableLabs," 2022.
- [5] Brian Yarbough, Cox Communications, "FTTX PON Architecture Considerations," in *SCTE Cable-Tec Expo*, 2021.
- [6] Curtis Knittle, CableLabs, "Cable Next-Gen Technologies & Strategies, CableLabs," 2022.
- [7] R. Anderson, "Access Network Powering, Network Power Considerations for 10G Enhancements," *SCTE Standards Journal*, pp. 23-37, 2021.

©2022 EnerSys, all rights reserved. EnerSys and the EnerSys logo are trademarks and/or registered trademarks of EnerSys and/or its affiliates in the United States and/or other countries. Other brand and product names may be trademarks or registered trademarks of their respective holder(s). The information contained herein is subject to change without notice. EnerSys shall not be liable for technical or editorial errors or omissions contained herein.

Leakage Detection In Full Duplex DOCSIS

Identification and Measurement of Leakage Levels in a Multigigabit Symmetrical Full Duplex DOCSIS Network

A Technical Paper prepared for SCTE/ISBE by

Benny Lewandowski

Engineering Architect IV

Comcast

Benny_Lewandowski@Comcast.com

Greg Tresness

President

Arcom Digital, LLC.

tresness.greg@arcomlabs.com

Jon-En Wang

Executive Director Quality Assurance

Comcast

jon-en_wang@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
Capabilities of Full Duplex DOCSIS.....	3
Differences in leakage detection for Full Duplex DOCSIS (FDX) versus tradition standard/mid/high split systems	5
2. Leakage detection techniques appropriate for FDX.....	6
3. Monitoring of potential FDX LTE egress	8
Simulated Full Duplex DOCSIS Leakage Detection Test Results	9
Radiated Three Meter Open Air Sensitivity Testing in an Anechoic Chamber	10
1 Source Signal Details.....	10
2 Baseline Test Setup	11
3 Baseline Test Procedure	11
4 FDX Leakage Detector Test Setup	12
5 FDX Leakage Detector Meter Test Procedure.....	12
Conclusions.....	13
Abbreviations	13

List of Figures

Title	Page Number
Figure 1 - FDX Resource Block Assignments ("RBAs")	4
Figure 2 - Configurable FDX Allocated Spectrum Bandwidths with Aeronautical and LTE Bands Shown	4
Figure 3 - FDX Cable Modem Upstream Output Power Spectral Tilt (64.5 dBmV TCP)	5
Figure 4 - FDX Network Downstream to Upstream Signal Level Deltas	6
Figure 5 - OFDM harmonics used for DS leakage detection-.....	7
Figure 6 - Spectrum of OUDP burst generated by a CM	8
Figure 7 - FDX test setup block diagram	9
Figure 8 - FDX meter configuration.....	9
Figure 9 - FDX DS and US detection results	10
Figure 10 - Baseline Leakage Detection Setup	11
Figure 11 - Leakage Detection Meter Test Setup	12

1. Introduction

Operators of cable networks are very familiar with the mandated requirements for monitoring signal leakage. In standard and mid-split deployments, leaks are discovered using a variety of techniques such as direct Quadrature Amplitude Modulation (QAM) measurement, direct detection of Orthogonal Frequency-Division Multiplexing (OFDM) signal components, and detection of Continuous Wave (CW) signals generated at the headend or Node Remote PHY Device (RPD) in the downstream direction. In high split deployments, leaks in the aeronautical band are discovered by monitoring the upstream for a specific test burst generated by the cable modem under control of the Cable Modem Termination System (CMTS.)

With the introduction of Full Duplex DOCSIS (FDX), some of the status quo practices for leakage detection will need to be changed due to the inherent differences between the traditional network and FDX which pushes the upstream (US) frequency up to 684 MHz. This paper explores several subjects related to detection and measurement of leakage in the FDX network. We will discuss how this will now necessitate leakage detection methods in both the downstream (DS) and US direction simultaneously in the aeronautical band, and we will discuss the potential benefit to perform additional Long-Term Evolution (LTE) band frequency measurements for upstream leakage when using full bandwidth FDX implementations.

We will additionally share lab test results demonstrating concurrent upstream and downstream leakage detection and we will propose a new test methodology for use in a controlled environment which will allow meter sensitivity testing for the new simultaneous upstream and downstream leakage detection.

Capabilities of Full Duplex DOCSIS

Full Duplex DOCSIS is targeted at significantly increasing upstream capacity by using coaxial spectrum for synchronized upstream and downstream communications by means of full duplex techniques. As part of the industry “10G” initiative, the primary objective of DOCSIS 4.0 FDX is to enable multi-Gigabit symmetric data services over the Hybrid Fiber Coax (HFC) network, while also adding significant new upstream bandwidth for continued capacity growth. The DOCSIS 4.0 FDX specification in conjunction with FDX Node and Amplifier specifications are intended to provide the requirements that enable DOCSIS 4.0 Full Duplex performance over “Node+x” HFC architectures. The overlapping DS and US spectrum, which is the key new physical (PHY) component of FDX, necessitates the need for Echo Cancellation capability in the node and amplifier, as well as other digital processing blocks to manage two-way signal passage, while maintaining platform stability and signal fidelity.

Building on a long history of delivering significant capacity improvements – most recently powered by advances in the implementation of DOCSIS 3.1 -- operators continue to develop and build on customers’ future 10G needs leveraging DOCSIS 4.0.

FDX allows blocks of channels, via “Resource Block Assignments,” as shown in Figure 1 to be used in both the upstream and downstream direction dynamically, significantly increasing the upstream throughput capability of the HFC network. Figure 2 displays the spectrum grid definition called out in the DOCSIS 4.0 FDX specification along with the overlapping Aeronautical and LTE frequency bands. In addition to managing signal overlap at the physical layer, FDX operation also requires new scheduler techniques that groups users based on their relative isolation characteristics to prevent co-channel interference from one FDX US Transmit (Tx) onto another’s FDX DS Receive (Rx). The vCMTS (virtual Cable Modem Termination System) creates “Interference Groups” and “Transmission Groups” to ensure compatible devices are given access to the spectrum when it is shared. Finally, Figure 3 shows that due to

cable loss at higher frequencies, the cable modem will now be responsible for providing up tilt in the upstream direction out to 684 MHz to account for cable loss with a nominal tilt of 10 dB that can be adjusted to 8-12 dB (1 dB relaxation for 12 dB tilt).

Sub-Band 1	Sub-Band 2	Sub-Band 3
DS	DS	DS
DS	DS	US
DS	US	DS
US	DS	DS
DS	US	US
US	DS	US
US	US	DS
US	US	US

108 MHz
 300 MHz
 492 MHz
 684 MHz

Figure 1 - FDX Resource Block Assignments (“RBAs”)




Figure 2 - Configurable FDX Allocated Spectrum Bandwidths with Aeronautical and LTE Bands Shown




Figure 3 - FDX Cable Modem Upstream Output Power Spectral Tilt (64.5 dBmV TCP)

Differences in leakage detection for Full Duplex DOCSIS (FDX) versus tradition standard/mid/high split systems

In standard and mid split networks, aeronautical band signals are only transmitted in the DS. As such, DS leakage detection is required to comply with mandated leakage regulations. In high split networks, signals transmitted in the aeronautical band are exclusively in the US. As such, for high split US leakage detection where leakage test signals are generated from the cable modem (CM) becomes a necessity.

Figure 4 provides an illustration of the US and DS signal levels in an FDX network at 108 MHz and 684MHz. When looking at the levels the inverted nature of signal levels in the network becomes obvious, where US signal levels are at a maximum in the 'Drop', and DS signal levels are at a maximum in the 'Feed' section of the plant. This makes it such that legacy aeronautical DS leakage detection techniques are insufficient to maintain visibility to US leakage.

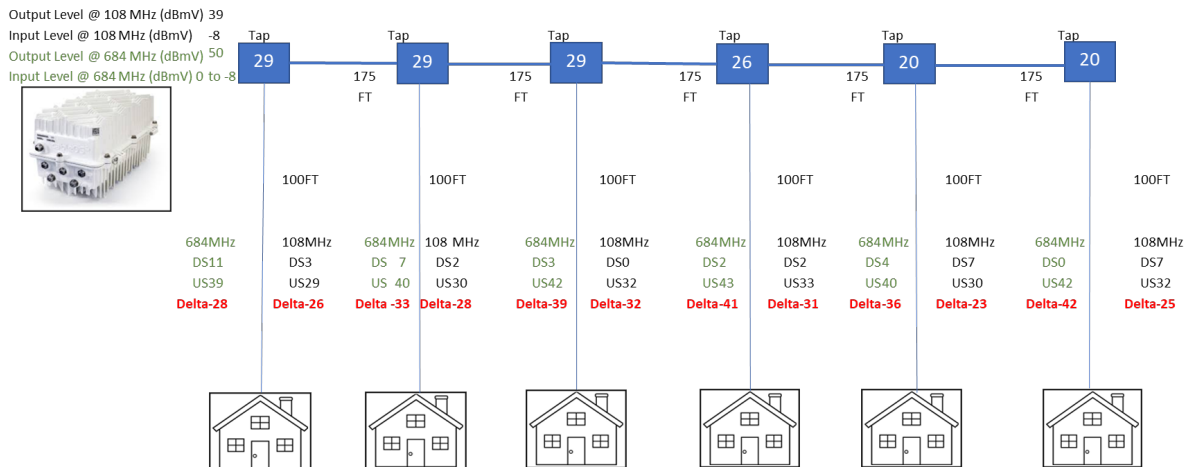


Figure 4 - FDX Network Downstream to Upstream Signal Level Deltas

With FDX, this drives a need for US leakage monitoring in the aeronautical band; in addition to the traditional DS aeronautical band leakage monitoring. This simultaneous monitoring provides full visibility to all aeronautical band network egress. If only DS detection was enabled, there would not be visibility to US leakage. And if only US detection was enabled, there would not be visibility to DS leakage.

2. Leakage detection techniques appropriate for FDX

There are many different techniques currently employed for leakage detection. Many of the techniques are optimized and directly tied to the signal format transmitted at the desired detection frequency. For example, DS detection at frequencies where OFDM signals are present is most easily performed by direct detection of existing signal components of the OFDM channel. This technique has the benefit of no extra signals being inserted into the network and no corresponding additional headend equipment required. It has an additional benefit that the amplitude of detected signal is significantly higher than alternative inserted low level carrier approaches, so depending upon the vendor implementation, the realized meter sensitivity can be superior to alternative approaches. Figure 5 illustrates the OFDM signal harmonics utilized for DS signal detection in one implementation.




Figure 5 - OFDM harmonics used for DS leakage detection-

The US signal format for FDX is Orthogonal Frequency-Division Multiple Access (OFDMA). The 2020 SCTE paper, Leakage in a high split world describes a methodology where an OFDMA upstream data profile (OUDP) test burst is used for aeronautical band leakage detection. The paper details how this approach provides ample sensitivity for Federal Communications Commission (FCC) compliance. This OUDP approach has now been accepted as a standard as the methodology used for US signal detection. With the OUDP approach, each CM on a node, under control of the CMTS – will sequentially generate an OUDP test burst with known signal characteristics. The leakage detector is then able to lock onto this OUDP test burst and detect signals egressing the coaxial network.

The spectrum of an OUDP burst generated by a CM is shown in Figure 6.




Figure 6 - Spectrum of OUDP burst generated by a CM

With FDX, regardless of the FDX bandwidth, the DS signal format will be exclusively OFDM. Similarly, the US signal format is exclusively OFDMA. As such, for DS leakage detection the direct detection of OFDM signal component techniques should be employed; and for US leakage detection the OUDP technique should be employed.

3. Monitoring of potential FDX LTE egress

As a best practice, in addition to FCC mandated aeronautical band leakage detection, cable operators today regularly monitor downstream LTE frequencies for leakage ingress. Benefits resultant from this effort include hardening of the plant which improves network quality, and it additionally helps to ensure that signals egressing from the plant do not adversely affect the licensed spectrum utilized by mobile operators. Figure 2 shows the spectrum bandwidths configurable for FDX and highlights the overlap with the aeronautical band and a lower LTE band. In the full 576 MHz FDX implementation, there is an FDX bandwidth (BW) overlap with the 600 MHz LTE band which is widely utilized by a mobile operator in the United States. As such, if a network is utilizing the full 576 MHz FDX bandwidth, attention should be paid as to not cause harmful interference to this licensed spectrum.

The 2012 Society of Telecommunications Engineers (SCTE) paper, Another Look at Signal Leakage, explored the lack of correlation between signal leakage in the aeronautical band and signal leakage at 700MHz LTE frequencies. A conclusion of the paper was that monitoring solely at the aeronautical band provides inadequate visibility to signal leakage at or near the LTE band. This was an impetus for the high frequency DS leakage detection performed today. This conclusion holds true with FDX – if utilizing full 576 MHz FDX, the OUDP US detection performed at the aeronautical band will not provide visibility to any US egress in the 600MHz band. Depending upon the detection frequency and upon antenna BW, the DS detection at or near the LTE band should provide adequate coverage of DS egress at the 600 MHz LTE band, but it will not provide visibility to US egress at this band.

What will be required in such a situation in order to effectively monitor US egress is a 2nd OUDP US detection session, performed at some frequency below 684 MHz.

An additional factor which supports the conclusion that aeronautical band US detection does not correlate with egress in the 600MHz band is related to the significant tilt between 108 MHz and 684 MHz. This tilt which is illustrated in Figure 3, will result in any signal leakage at the higher FDX frequencies having correspondingly higher levels as compared to leaks at lower frequencies such as the aeronautical band.

Simulated Full Duplex DOCSIS Leakage Detection Test Results

A signal leakage meter was designed and built for FDX operation, capable of simultaneous US and DS detection. Testing was performed within the CTA channel 16 bandwidth. To verify the concurrent US and DS aeronautical band detection, a test setup was built as shown in Figure 7.

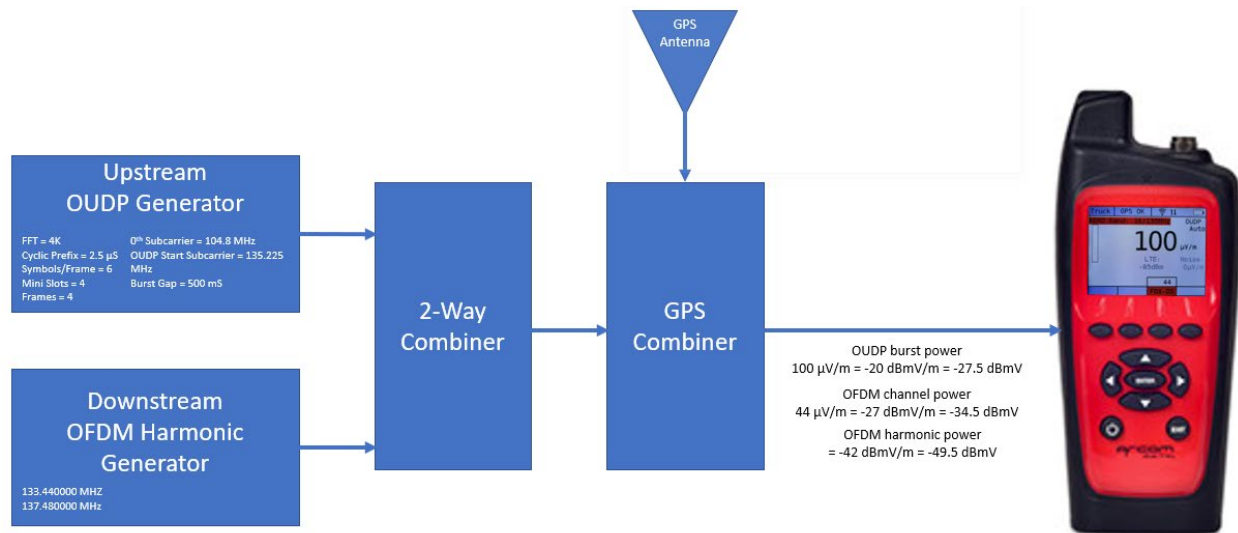


Figure 7 - FDX test setup block diagram

Channel type: FDX					
AERO band channel:	16/135.0MHz	Cyclic prefix, µs/Ts:	2.5/256	Number of minislots:	4
OFDMA:	4K, pattern 11	Symbols per frame:	6	Number of frames:	4
OFDMA zero subcarrier freq., MHz:	104.800	OUDP start subcarrier freq., MHz:	135.225	Tuner freq., MHz:	136.0125
OFDM F1, MHz:	133.440	F2, MHz:	137.480		

Figure 8 - FDX meter configuration

The detection results are shown on Figure 9, where the meter displays the measured signal level for both US and DS. Here, the US OUDP showed the expected 100 $\mu\text{V/m}$ and the DS OFDM detection showed the expected 44 $\mu\text{V/m}$, which were the measured signal level input into the meter as illustrated in Figure 7. Two views of the display are provided, each with the prominent US/DS detection location alternated.

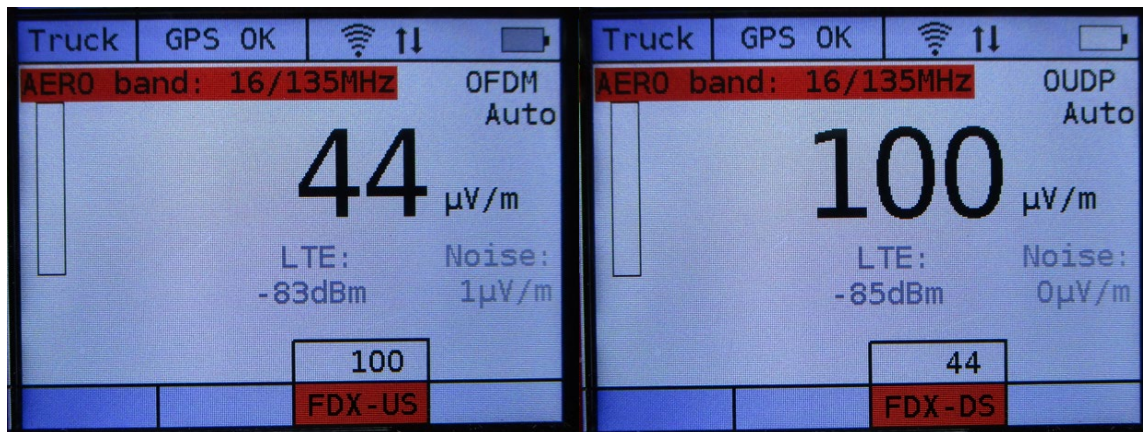


Figure 9 - FDX DS and US detection results

Radiated Three Meter Open Air Sensitivity Testing in an Anechoic Chamber

The test procedures described below can be used for all HFC Standard, Mid and High split leakage meter testing as well as Full Duplex DOCSIS testing. Legacy downstream OFDM and upstream OUDP bursts will be retained as the source signals as defined in previous technical papers and specifications. By simultaneously transmitting an upstream OUDP and downstream OFDM signal, FDX functionality of the leakage detection meter can be evaluated.

1 Source Signal Details

US OUDP Signal Details:

1. Symbols Per Frame (K) = 9
2. Modulation Order = 256 QAM
3. Pilot Pattern = 11
4. Center Frequency of OUDP Signal = 136.0125 MHz, 603 MHz
5. 4 Mini-slots (1.6 MHz Upstream Bandwidth with the 4 adjacent mini-slots to the center frequency above)
6. Number of Frames = 8 and 2.16 mS in transmit time duration. (For 256 House Holds Passed roundtrip time = 563.2 mS)
7. 4K Fast Fourier Transform (FFT) = 40 μS per symbol + Cyclic Prefix
8. Cyclic Prefix = 1.5625 μS
9. Window Roll off Period = 0.9375 μS

DS OFDM Signal Details:

1. Channel 16, 88
2. 135 MHz, 609 MHz
3. Cyclic Prefix = 2.5 μ S
4. 4K Fast Fourier Transform (FFT)

2 Baseline Test Setup




Figure 10 - Baseline Leakage Detection Setup

3 Baseline Test Procedure

1. Set the Radio Frequency (RF) Signal Generator to output a CW carrier at each of the center frequencies given in the “Source Signal Details” section.
2. Connect the RF signal generator, all cables used for the test, and the RF variable attenuator to the power meter. Confirm the RF power level with the RF power meter and account for any cable loss and 0 dB attenuator setting using the Amplitude offset on the power meter.
3. Verify that all of the switches in the RF variable attenuator are within 5% tolerance of their expected values.
4. Connect the RF Signal Generator to the Very High/Ultrahigh Frequency (VHF/UHF) transmit antenna.
5. Measure and verify the distance between the VHF/UHF transmit antenna and Electric Field Strength Probe is 3 meters.

6. Connect the RF Power Meter to the Electric Field Strength Probe and measure the RF level of the CW carrier.
7. Confirm the measured field strength is correct taking into account the gain/antenna factors of the antennas and Free Space Path Loss (FSPL) during calibration.
 - a. 3m Free Space Path Loss is ~24.65 dB at 136 MHz
 - b. 3m Free Space Path Loss is ~36.68 dB at 609 MHz

4 FDX Leakage Detector Test Setup




Figure 11 - Leakage Detection Meter Test Setup

5 FDX Leakage Detector Meter Test Procedure

1. Configure the FDX cable modem/gateway to output continuous OUDP signals defined in “US OUDP Signal Details” and the FDX node to output downstream signals defined in sections “DS OFDM Signal Details.”
2. Confirm each RF power level with the RF power meter and account for any cable loss and RF variable attenuator insertion loss.
3. Verify that all of the switches in the RF variable attenuator are within 5% tolerance of their expected values.
4. Connect the FDX cable modem and Node to the VHF/UHF transmit antenna through the RF variable attenuator.
5. Measure and verify the distance between the VHF/UHF transmit antenna and leakage detection meter is 3 meters.

6. Make sure the meter is configured properly to the settings given in “Source Signal Details” and measure the RF level of the Aeronautical Band Signals in both the upstream and downstream direction.
7. Make sure the meter is configured properly to the settings given in “Source Signal Details” and measure the RF level of the LTE Band Signals in both the upstream and downstream direction.
8. Add 6 dB of attenuation to change the distance to ~6m. Repeat Steps 6-7.
9. Continue to add attenuation in 6 dB steps doubling the distance each time until you get to ~192m (36 dB total attenuation). Repeat Steps 6-7 each time. Record all Measurements.
10. Configure the FDX cable modem to output an OUDP burst 2 times per second with the parameters given in “US OUDP Signal Details” and repeat steps 6-9.

Note 1: The Device under Test (DUT) should rotate 360 degrees while the antenna raises and lowers in 1 or ½ meter steps up to 3 meters to account for antenna patterns.

Note 2: A more advanced test setup can be created to simulate mobile detection in the anechoic chamber by connecting a function generator to the RF variable attenuator and simulate signal strength increasing to a peak and then decreasing in a real-world scenario where the leakage detector meter’s position is changed relative to the source of leakage.

Conclusions

With cable networks advancing, the demand for higher download and upload speeds continues to be met. FDX is a great way of utilizing existing bandwidth in the network transmitting data in both the downstream and upstream directions at the same time. These state-of-the-art FDX networks require leakage detection meters that can detect and record leakage levels simultaneously in the DS and US direction to stay in compliance with FCC regulations, as well as keeping the entire network integrity in good working order. In FDX implementations utilizing the full 576 MHz bandwidth, there is the potential for egress affecting licensed 600MHz LTE – so additional US leakage monitoring using a second high frequency OUDP test session should be considered. Test equipment capable of the concurrent US and DS monitoring was demonstrated. Lastly, with the test procedure given in this paper, operators can reliably assess the sensitivity of FDX and all HFC leakage detection meters in a precise setting allowing for continued leakage monitoring that not only allows for sustained regulatory compliance, but also reliable plant monitoring and maintenance for optimal system performance.

Abbreviations

BW	bandwidth
CATV	community antenna television
CM	cable modem
CMTS	cable modem termination system
CPE	customer Premises equipment
CW	continuous wave
DAA	distributed access architecture
dBc	decibel from Carrier
dBmV	decibel Millivolt
DOCSIS	data over cable service interface specification
DS	downstream
DUT	Device under test

FCC	Federal Communications Commission
FEC	forward error correction
GPS	global positioning system
HFC	hybrid fiber-coax
HSD	High Speed Data
f	frequency
FDX	full duplex DOCSIS
FSPL	Free space path loss
Gbps	gigabit per second
Hz	hertz
ISBE	International Society of Broadband Experts
K	maximum number of cable modems on a node
K	symbols per frame
K (dBc)	boosting gain of CW test signal
kHz	kilohertz
LTE	Long Term Evolution
M	number of subcarriers in 6 MHz
MAC-PHY	media access control channel physical layer
Mbps	Megabit per second
MHz	megahertz
ms	millisecond
MSO	multiple system operator
N (dBc)	coefficient of recalculation level of OFDMA signal in BW= 6 MHz to the measured level of CW test signal.
NCTA	National Cable Television Association
OOB	out of Band
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
ODUP	OFDMA upstream data profile
QAM	quadrature amplitude modulation
RF	radio frequency
RPD	remote physical device
R-PHY	remote physical layer
RX	receive
s	Second/s
SC-QAM	single carrier quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers
STB	set-top box
STD	standard
TCP	total composite power
TDMA	time division multiple access
TX	transmit
UHF	Ultrahigh Frequency
μV/m	microvolt per meter
μs	microsecond
US	upstream
vCMTS	Virtual Cable Modem Termination System
VHF	Very High Frequency

Wi-Fi	wireless fidelity
-------	-------------------

Bibliography & References

- 1) [PHYv4.0] Physical Layer Specification, CM-SP-PHYv4.0-I04-210826, August 26, 2021, Cable Television Laboratories, Inc
- 2) *Another Look at Signal Leakage, The Need to Monitor at Low and High Frequencies*; Ron Hranac, Greg Tresness, SCTE EXPO '12
- 3) Code of Federal Regulations, Title 47, Part 76 MULTICHANNEL VIDEO AND CABLE TELEVISION SERVICE
- 4) *Leakage in a High Split World*; John Chrostowski, Dan Rice, Greg Tresness, Benny Lewandowski, SCTE EXPO '20

Machine Learning and Telemetry Improves Outside Plant Power Resiliency for More Reliable Networks

A Technical Paper prepared for SCTE by

Stephanie Ohnmacht

Sr. Director, Product Management
Comcast
1800 Arch St, Philadelphia, PA 19403
Stephanie_Ohnmacht@comcast.com

Matthew Stehman

Senior Data Scientist
Comcast
1800 Arch St, Philadelphia, PA 19403
Matthew_Stehman@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Identifying Power Supplies.....	4
2.1. Geolocation Match via Excel Algorithm	5
2.2. Algorithm 2.0	5
2.3. Machine Learning to Identify Power Supply through Open-Source Mapping Tools (Find My Power)	5
2.4. Manual Intervention.....	5
3. Asset Inventory Interface	6
3.1. Asset Inventory User Interface.....	6
3.2. Alarming Interface	7
3.3. Impacted Actives (nodes, cable modems) on Plant.....	7
4. Power Supply Telemetry	8
4.1. Telemetry Overview	9
4.2. Streaming Infrastructure and Real Time Eventing	10
4.3. Big Data Pipelines for Efficient Storage and Analysis	11
5. Modeling Power Supply Health.....	12
5.1. Modeling Objectives.....	13
5.2. Data Sets Used for Battery Runtime Prediction.....	14
5.3. Machine Learning for Battery Runtime Prediction	14
6. Conclusion.....	17
Acknowledgements	17
Abbreviations	18
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - Architecture.....	4
Figure 2 - Geolocation Matching Algorithm.....	6
Figure 3 - Example Output from Find My Power	6
Figure 4 - Inventory/Alarming Tool Architecture	7
Figure 5 - PS Outage Impact Analysis Example.....	8
Figure 6 - Power Supply Overview	8
Figure 7 - Big Data Processing Pipeline Architecture.....	12
Figure 8 - Proactive vs Demand Maintenance.....	13
Figure 9 - Machine Learning Run Time Prediction Architecture.....	15
Figure 10 - Random Forest Performance. a.) Feature Importances, b.) Prediction Errors Compared to Baseline.....	16
Figure 11 - Example Random Forest Prediction of Battery Run Time	17

List of Tables

Title	Page Number
Table 1 - Relevant Power Supply Telemetry Metrics.....	9
Table 2 - List of Real Time Alerts.....	11

1. Introduction

As the network becomes more powerful through 10G technology, resiliency in our extensive power supply network is essential to ensure our customers are always connected. The stability of outside plant (OSP) power supplies (PS) is key to keeping the network online and serving customers. Maintaining accurate location and active telemetry information is vital to keeping the power supplies in optimal health. Machine learning is employed to analyze this massive amount of telemetry data and provide actionable insights related to operating conditions and the overall health of the power supplies and batteries.

The authors, Stephanie Ohnmacht & Matt Stehman, will present a multi-tier solution that was developed to address this at scale. The solution includes the integration of mainstream mapping technology, with machine learning (ML) routines to optimize location probabilities and use of big data pipelines and advanced data science techniques to inform proactive and demand maintenance activities. Predictive models are built on top of this large dataset to help detect long term variations in power supply performance as well as real-time performance evaluations during active outages. This paper will illustrate a proven approach to improving PS resiliency, leading to increased network reliability, and satisfying the requirements of a 10G network.

2. Identifying Power Supplies

Historically Comcast has had to maintain multiple information sources and databases to manage the network of power supplies. These separate databases included those used for network designers, technical operations, and asset inventory management. This disparate data provided an opportunity to find a solution to enable all systems to talk to each other with one unique key. Creating an innovative system to maintain PS information in the future. To start, a team was formed to begin a 4-step process to validate and correct location information for Comcast's PS across databases.

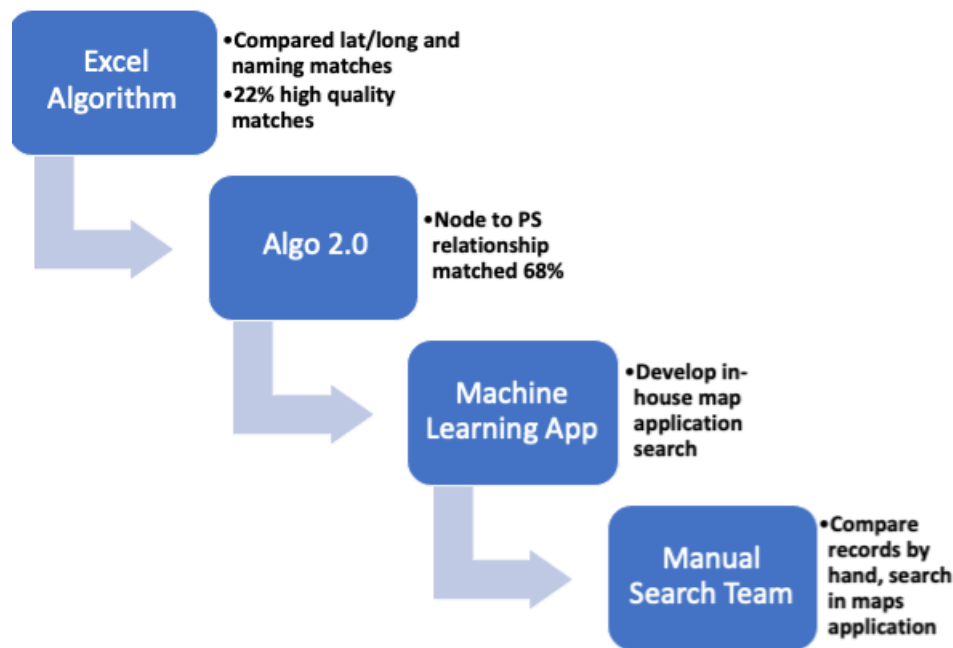


Figure 1 - Architecture

2.1. Geolocation Match via Excel Algorithm

The goal of the geolocation match was to begin to narrow down location discrepancies across databases. To determine a single PS referenced in multiple databases, the team developed an algorithm to compare the delta of linear feet between the two locations in the databases. If the PS locations were within 300 ft, then it was marked as the same location. In situations where the distance was greater than 300 ft, an index was applied to show a level of confidence of a match based on the linear feet differences. This level of geolocation matching resulted in 22% of the total population identified as having a high-quality match within the total population.

2.2. Algorithm 2.0

Once completing the first algorithm, the next phase of the analysis was to narrow down outstanding matches based on other known links within the databases. This approach focused on using node-to-power supply relationships to solve for additional location matches and increase level of confidence. This analysis increased matched records to 68% of the PS population.

2.3. Machine Learning to Identify Power Supply through Open-Source Mapping Tools (Find My Power)

A tool was developed that would identify PS in its natural environment (street or pole) using existing mapping applications to confirm the actual location. To initiate the project, power supply photos were taken in the field, then loaded into the application to begin training the system to identify different power supply types and installation environments. These examples were integral to the ability of the tool to accurately detect a PS in all situations. Once the right training model was developed, an open-source visual mapping tool was integrated into the tool that captured real street view imagery and provided associated latitude and longitude coordinates for the exact location of the identified power supply. After an attempted PS match, a level of confidence was determined by the tool and applied to the record which was then confirmed manually by a team.

This new system was able to find power supplies for 65% - 70% of the locations. The mapping tool did not have images for about 10% of the locations (given easements, landscape, and gated communities). About 45-50% of the power supplies predicted were correctly identified by the artificial intelligence and significantly boosted the team's productivity by 2X when verifying manually.

2.4. Manual Intervention

Once the team confirmed proposed matches from the machine learning tool, outstanding PS matches were set aside to be reviewed by field teams and updated.

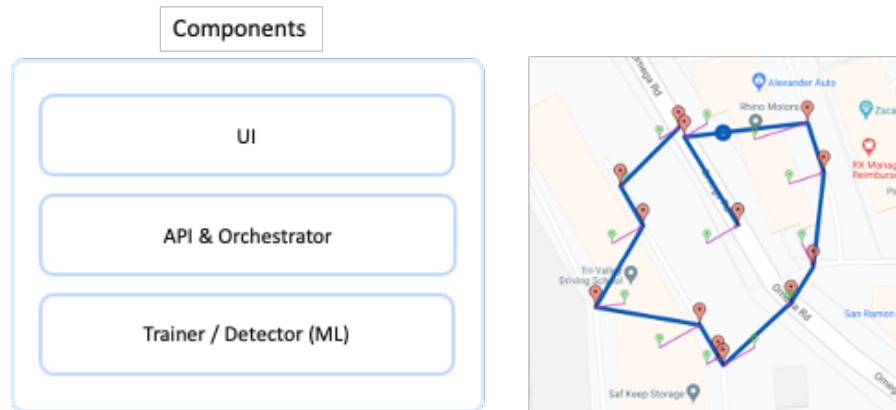


Figure 2 - Geolocation Matching Algorithm

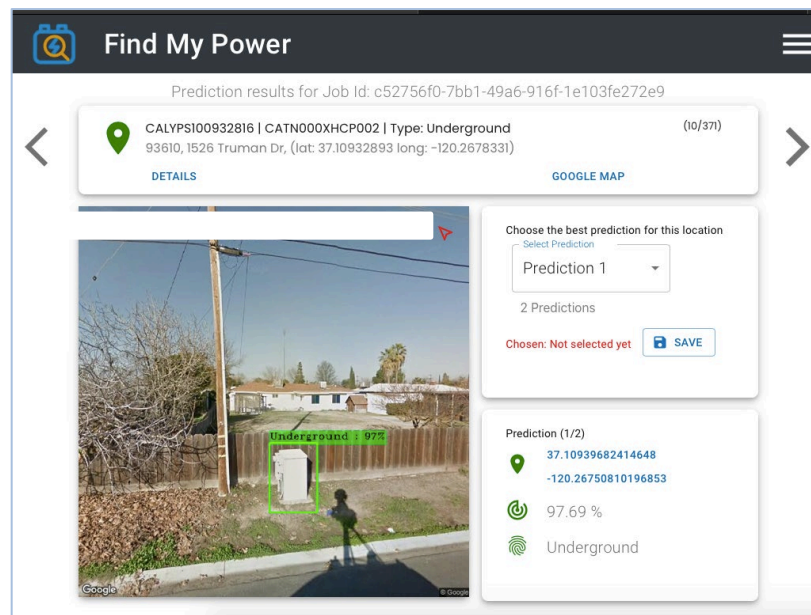


Figure 3 - Example Output from Find My Power

3. Asset Inventory Interface

Updated PS location information and node associations provided the opportunity to create a tool for management, field techs and operations called Power Supply Notebook (PSNB).

3.1. Asset Inventory User Interface

The team created a source of truth User Interface (UI) for all PS asset details. Including PS make/model/serial number, location, batteries, transponder, installation date, and maintenance records. For new PS the asset information is loaded once the PS is built and online. Ongoing updates and changes can be completed by administrators, technicians, business partners and the network help desk. Depending on the user's permission level, changes are added to a queue and approved by appropriate staff. All changes are tracked and recorded for historical purposes and investigation needs.

3.2. Alarming Interface

PS are polled every minute by in-house software. This data is monitored by the Correlation Engine (CE). As changes in telemetry flow through CE, anomalies are monitored to determine if an event is in progress. CE references all active devices on plant, downstream from the power supply and can correlate the impact of an event to number of devices impacted. Based on pre-determined threshold levels, by event type, an alert is sent to PSNB. PSNB is the presentation layer that displays current and historical alarms and events. Telemetry can be viewed by battery string, as well as the historical trends in battery voltage, input voltage, output current, temperature, output voltage, and output power. Events that need dispatching are sent to an automated dispatching system to be processed according to business rules. Alerts include PS on stand-by (on battery), PS loss of communication, and low battery.

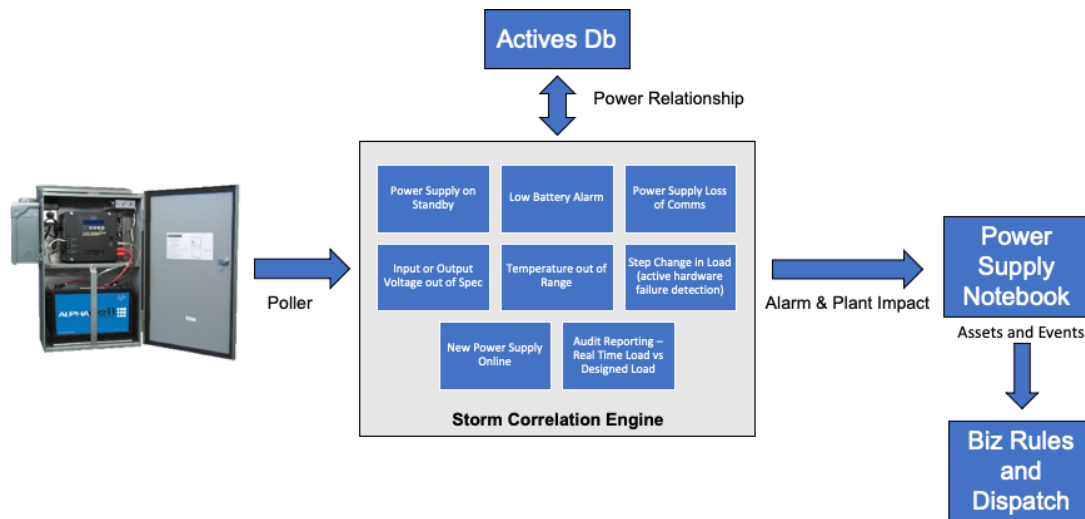


Figure 4 - Inventory/Alarming Tool Architecture

3.3. Impacted Actives (nodes, cable modems) on Plant

The data output from CE and passed to the PSNB UI, provides insightful data about outage impacts. In the past, it was assumed that any PS power outage affected all actives. Based on preliminary data, that is not the case. Data suggests that on average, 30-40% of customers stay online (have power) during a PS on discharge scenario. The geography of the power grids from the utility company and broadband provider do not overlap. As you can see in the example below, an outage from a PS only impacted a portion of the cable modems supported by that power supply. Empowered by this additional layer of information, teams can make more informed choices about how, when, and where to send a technician.

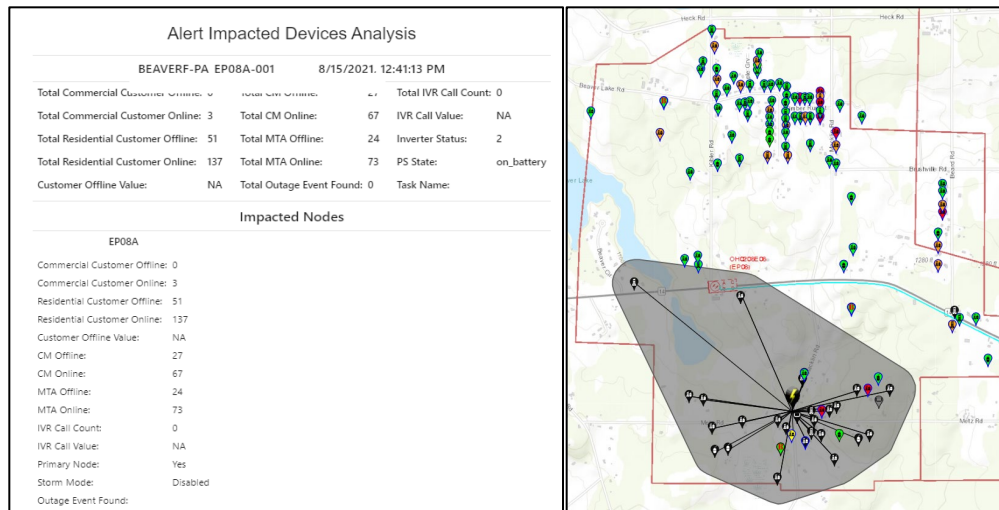


Figure 5 - PS Outage Impact Analysis Example

4. Power Supply Telemetry

The previous sections discussed the effort to build a source of truth that is foundational to the overall tracking and management of a power supply network. Once the global tracking and visibility for each power supply is established, live telemetry from individual power supplies can be incorporated to provide the status of each power supply as well as the entire network of power supplies. This section will discuss the telemetry gathered from each power supply in Comcast's plant and how the raw data can be used to make informed decisions.

Each power supply is made up of several sub-components where each is responsible for various aspects of overall operation. Understanding the function of the different components as well as the various telemetry metrics captured is key to modelling the health of the power supplies. See Figure 6 for a simplified summary of the power supply with components relevant to this paper.

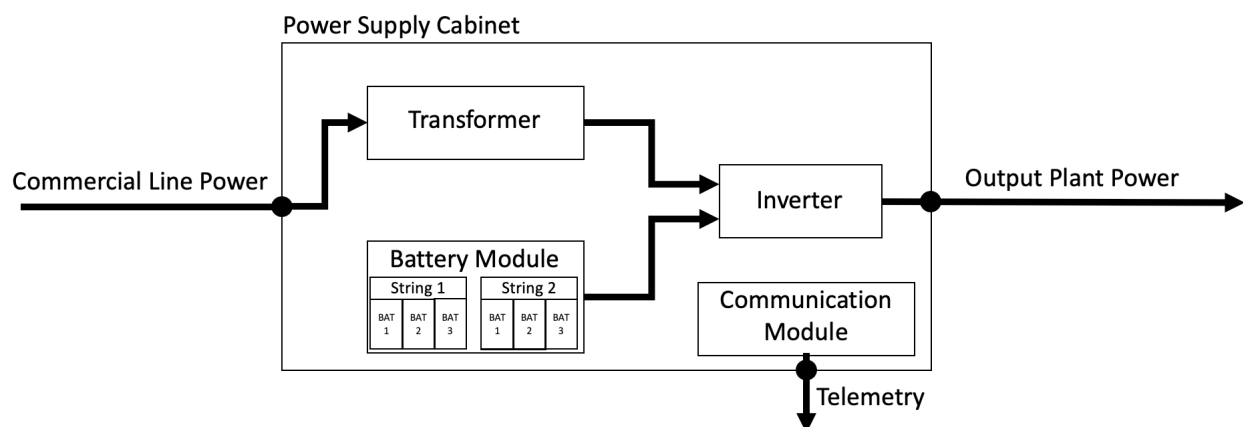


Figure 6 - Power Supply Overview

The transformer component is responsible for converting the commercial high voltage AC power down to operational range and format for the power supply. The battery module houses the strings of backup

batteries that are used to power the plant in case of a commercial power outage. The Inverter is used to convert the DC voltage from the backup batteries into AC as required by the elements of the RF plant (RF nodes, RPDs, amplifiers, etc..). The inverter also controls which power source to use in the event of a change in input voltages. The communications module is responsible for relaying all relevant information about the power supply through DOCSIS[®] protocols back to Comcast HQ for continuous monitoring.

The following sections will discuss the different telemetry metrics and data pipelines to consume the streaming data which allow for analysis and insights.

4.1. Telemetry Overview

The communications module is responsible for combining all measurements of the power supply state from the individual sub-components and packaging them into a telemetry stream that can be read through DOCSIS protocols. The PS telemetry is polled via simple network management protocol (SNMP) at two sampling intervals: every 60 seconds for high frequency engineering telemetry and hourly for configuration and system level information. The PS can also push alarms in between SNMP polls in the case of specific conditions being outside of acceptable values. Table 1 contains a simplified list of the telemetry available from the communications module.

Table 1 - Relevant Power Supply Telemetry Metrics

Frequency	Context	Telemetry Metrics	Description
Minute	Input to Power Supply	Voltage, Current, Power	Measurements of electrical input to the power supply transformer from commercial line power
	Power Supply Operation	Inverter Status	State of power supply: commercial input, battery backup/standby (outage), battery backup (self-test)
		Tamper Status	Indication of whether PS has been tampered with
		Internal Temperature	Environmental temperature of the power supply cabinet
		Alarm	Internal PS Alarms. Major: service impacting, Minor non-service impacting
	Battery Operation	Individual Voltages	Measured voltage at each battery terminal
		Total Voltage	Total voltage from all battery strings at the input to the inverter
		String Charge Current	Measured charging current at the input to each battery string
		String Discharge Current	Measured discharge current from each string of batteries
		String Float Charge Current	Measured floating charge current for each string of batteries

	Power Supply Output	Voltage, Current, Power	Measurements of electrical output from the power supply to the outside plant
Hourly	Power Supply Info	Vendor, Model Number, Serial Number	General information about the individual power supplies
		Configurations and supported functions	Overview of various configurations for each PS and which functionalities are supported
		System	Software versions, device ids, Ip/Mac addresses, etc....

The table indicates the general families of live telemetry captured from the communication modules of each power supply in the field. The minute polls tend to include more engineering specific information about voltages, currents, and statuses relevant to the functional operation of the power supply. The hourly polls contain more stable information useful for understanding the configurations and settings of the power supply.

From an operations standpoint, the power supply input telemetry gives an indication of the health of the underlying commercial power network. The inverter status is used to determine whether the PS is being powered from commercial power or back-up batteries. The output telemetry is useful for quantifying the size of the plant the power supply is powering and is a key feature in validating the power supply against design specifications. The output telemetry is also a key feature in outage runtime prediction models (will be discussed in Section 5), since how long a battery will last is largely dependent on how much current/power draw the plant requires. The PS temperature field is also useful in this regard because it is well known that batteries tend to perform worse in colder conditions due a resulting reduction in the chemical energy in each cell.

The individual battery metrics are very useful for monitoring batter health. For instance, inspection of the voltage discharge and charge curves can contain a lot of information about the health of the batteries. The durations and boundaries of these individual events contain a mass of information related to battery performance and health. Since voltage data is stored for each individual battery, a determination can be made about whether a specific battery may be performing sub-optimally as opposed to assuming the entire string of batteries needs to be addressed.

This section introduced the individual power supply telemetry metrics and gave some indication of how they can be used to monitor the health of the power supplies. The next sections will discuss the data pipelines built to consume and analyze the real-time telemetry data to provide more insights and help guide the field with resource allocation and maintenance activities.

4.2. Streaming Infrastructure and Real Time Eventing

The telemetry from each individual power supply is streamed through a real time alerting architecture where alarms and warnings can be generated based on a variety of conditions that may require action to be taken. The data is then landed in a short-term storage solution to support further deep dive analyses on recent events if needed. Due to the shear amount of telemetry being generated by the network of power supplies the raw data is only stored for a short time to minimize overhead. More information on how this challenge is handled will be provided in Section 4.3.

The streaming architecture allows for real time alerting and eventing to be performed as the raw telemetry is collected. The events are then landed in a front-end user interface for prioritization and determining the appropriate action. A general overview of the type of alarms currently implemented is given in Table 2.

Table 2 - List of Real Time Alerts

Context	Name	Description
Inverter Status	No Commercial Power	Power supply has no commercial power and is operating on battery
	Power Supply Running Test	Power supply is running a scheduled self-test to check battery performance
	Power Supply Test Failed	Power supply self-test did not last the intended duration (requires investigation)
Telemetry Variations	Output Current, Voltage, Power	Alert when measurement varies significantly from its nominal value
	High Temperature	Power supply temperature is reaching unsafe levels
	Low Battery Voltage	Power supply is operating on battery and voltages are reaching depletion levels
Communications	Loss of Communications	SNMP can no longer communicate with communications module
	PS Internal Alarm	Forwarding of internal power supply alarms

These alerts are meant to be an indication that a given power supply is currently experiencing an event of interest and requires an assessment. These alerts are currently logic based and developed from an in-depth knowledge of the operating conditions of the power supplies. A history of these alerts (including the duration the alert was active) are kept for analysis of long-term trends so analysts can evaluate trends and provide appropriate courses of action when needed.

These types of alerts are industry standard and allow for a live view of activity on the network. These alerts are only aware of a local time window around the present time stamp and thus are limited in their abilities to find complex patterns in the raw data. Thus, there is a large opportunity in this space to adopt a machine learning approach that can gain more complex insights from historical data in combination with real time data.

4.3. Big Data Pipelines for Efficient Storage and Analysis

In order to build predictive models that can learn from historical events, a historical database of training data is required. However, due to the size of Comcast's power supply network which typically generates hundreds of billions of data points a day, it is cost prohibitive to store all the raw telemetry metrics for a long duration. To overcome this obstacle, Comcast has developed a staged strategy for identifying and storing events relevant to modeling the health of the power supplies. Figure 7 outlines the approach to only storing data that is useful for modeling.

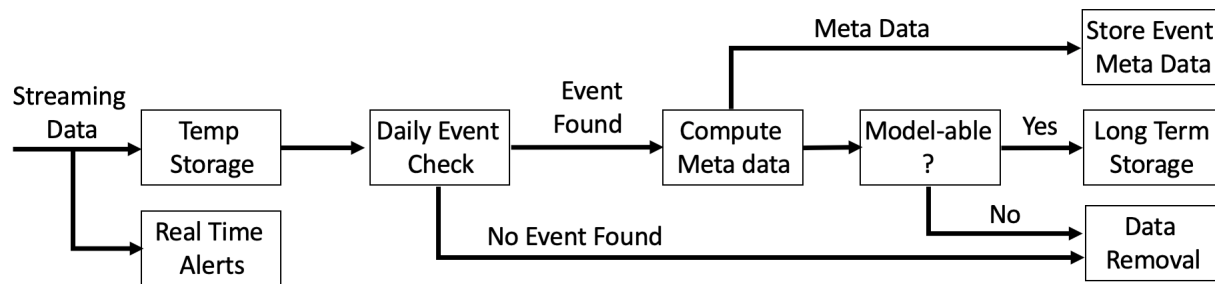


Figure 7 - Big Data Processing Pipeline Architecture

The big data pipelines consume the raw telemetry from the temporary storage solution discussed in Section 4.2. A daily job is then run that processes a single day's worth of data for the entire footprint in one go. This was implemented because some meta data that is calculated for the individual events requires more information than a local time window would afford and thus event meta data cannot be calculated through the streaming event/alerting process discussed previously. The daily job processes the raw data and determines if the criteria for any model-able events were found. The set of events that are checked against are events that contain information related to the health of the power supply, namely: battery discharge and charge events. By looking at how batteries charge and discharge, one can infer the health of the batteries and the different approaches for this health prediction will be discussed in Section 5.1.

If an event check does determine that a model-able event has occurred for a given power supply, the entire set of raw data for the duration of that event is stored indefinitely in a long-term storage solution. All raw data outside of these model-able events are then removed as part of the standard temporary data retention policies. When the event is processed, there is a litany of meta data calculated that is stored to help describe the characteristics of the event which is placed in an event database that has no retention policy. This table is light weight since every event is only a single row in a table and the meta data allows for insights into the physics of the events computed from raw data. This approach results in storing only 0.25% of the original raw data for an equivalent period of time. Obviously, this number can change based on the number of events but by implementing this staged storage approach Comcast has significantly cut down on the cost of storing data while not missing anything relevant to long term modeling efforts.

The model-ability of an event is determined by another set of criteria on the event meta data, a few examples are event duration, data quality, etc. For example, while a power supply may need to go into battery mode for a split-second commercial power outage, that event might not be useful for modeling the health of power supplies due to the very short duration. The next section will discuss more about the modeling objectives and how the raw telemetry values can be used in building models that can deliver on those objectives.

5. Modeling Power Supply Health

The previous section introduced the power supplies and telemetry metrics as well as the data processing pipelines that consume the raw data to provide real time alerts and pre-process for future modeling efforts. This section will discuss the different objectives of predictive models in the power spaces as well as the different modeling approaches and some initial results from this relatively new initiative.

5.1. Modeling Objectives

In the cable domain, a positive customer experience is the main objective of all efforts. For outside plant power that means keeping the plant online and customers connected at all costs. Thus, a predictive model should use all available data to determine if a customer impacting event is imminent or the impact assessment of any live customer impacting events. These types of scenarios in the cable world are typically referred to as proactive maintenance and demand maintenance. Many initiatives exist in the RF network maintenance realm such as Wolcott, et al. (2016) and Wolcott, et al. (2018) and many practices have even become industry standard as documented in CableLabs PNM Best Practices Primer (2020). For the power supply domain most efforts have been focused on improving the engineering models related to the power supplies such as Anderson and Burgett (2014). The authors believe that there is a large opportunity to enhance the overall field of power supply management and maintenance with machine learning which can consider many more inputs outside of raw engineering data. Figure 8 showcases the distinction between proactive and demand maintenance.

		Customer Impact	Resolution Window
Customer Impact	Resolution Window	<p>Now: Experiencing Impact Future: Impact Imminent</p>	<p>(Now, Now) Demand Maintenance</p>
	Resolution Window	<p>(Future, Future) Proactive Maintenance</p>	<p>Now: ASAP Future: Months, Weeks</p>

Figure 8 - Proactive vs Demand Maintenance

Proactive maintenance is the ability to proactively detect and fix any potential issues on the network before a customer is impacted. For power supplies a simple but effective passive technique would be to do a standard workup on every power supply on a regular schedule, say annually. While this approach is generally good at keeping power supplies operational in the long term a more active approach can help address any potential issues in a more targeted fashion. This is where the telemetry-based models will be able to thrive. For example, by monitoring the charge and discharge curves during all events for each power supply, models can be built that attempt to catch any long-term degradations in power supply performance. This type of model can start to predict slow variations in the performance over many months of run time and help alert the field that certain power supplies may need to be visited sooner than others.

In a demand maintenance scenario, there is an active situation, and a model would need to predict the outcome of a given event to help the field teams allocate resources in real time. In this instance the data related to a specific event is heavily relied on in addition to any relevant historical data to predict and score potential outcomes of the current event. For example, say a severe weather event results in a loss of commercial power and the power supply needs to operate on battery backup to keep the plant and customers online. A predictive model in this case would need to estimate how long the batteries will run

until they will reach end of discharge. Based on the model outputs, the field teams can prioritize which power supplies may need a backup generator if the run time is expected to be less than the commercial power outage duration.

In either case, the output of the respective model will be consumed by field teams and prioritized and worked accordingly. Therefore, any additional information that can be supplied to the teams performing the work will help increase efficiency and add confidence to how the work will be prioritized. For this reason, the models that will be explored in this effort need to have some avenue for explainability of the given prediction. Examples of explainability not only include a level of confidence in each prediction but also elements like global and local feature importance to indicate what features the model believes to be most important in making a prediction.

Comcast is currently focusing on a demand maintenance use case using a machine learning based run time prediction model for power supplies operating on battery backup during commercial power outages. Comcast sees many potential avenues for using the aforementioned data sets and methodologies to enhance the outside plant power domain on many fronts, from long term budget planning and global resource allocations to using the power supply telemetry in other models related to access network root cause analysis.

5.2. Data Sets Used for Battery Runtime Prediction

For the goal of predicting how long the power supply can power the plant in the absence of commercial input power a machine learning model will need access to data that contains predictive power in this instance. It is apparent the telemetry related to the electrical circuitry discussion in Section 4.1 would be required to describe the physical principals occurring. In fact, the state-of-the-art modeling in the battery domain pretty much solely rely on the electrical data to make run time prediction, i.e., physics based and even more modern time series-based approaches.

However, as discussed in previous sections, Comcast has put in a significant amount of effort to create an accurate inventory of all the power supply assets which can provide extra enhancements to the modeling efforts. Items of particular interest are the battery manufacturer, age, type, install date and number of batteries in each power supply. This battery specific information will help enhance a predictive model using global patterns from various meta data fields related to the identity of the batteries and power supplies.

The final class of data under consideration as input to a machine learning model are features that can be engineered from the aforementioned data sources. The process of feature engineering has proven to be one of the most important aspects in developing high performing machine learning models. Features that are derived from the raw input data can produce significantly more accurate results than the raw features alone. While Comcast is still exploring the different approaches to feature engineering in this capacity, some immediately obvious features of interest for a run time prediction could be: information about the last discharge (test or outage) of a given power supply like time since, duration of, ending voltage, variance of individual battery voltages; other relevant statistics as well as historical information including number of discharge cycles in last n months, number of cold temperature cycles, time since last maintenance visit.

5.3. Machine Learning for Battery Runtime Prediction

The requirements for a run time prediction model are straightforward: when commercial input power is lost the model should predict how long the power supply can run on battery power before it can no longer power the plant. Figure 9 outlines the overall architecture of the desired model where the model takes

inputs from the aforementioned data sources and makes a run time prediction at the moment that the commercial power outage begins, as well as a confidence in the prediction and some form of explainability in what led the model to making that prediction. The 3 outputs combined will be the key aspect of how the field will prioritize and ultimately address predictions that indicate intervention will be needed prior to commercial power coming back online.

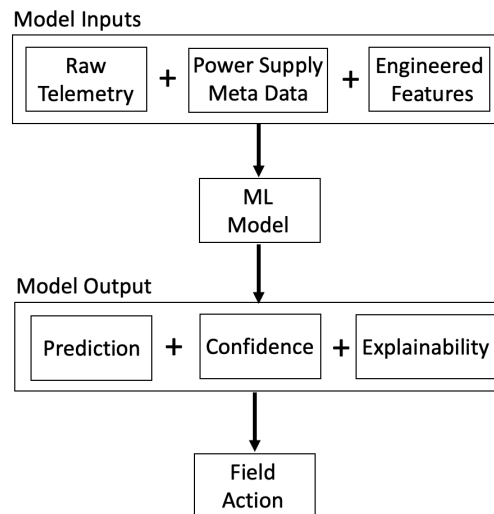


Figure 9 - Machine Learning Run Time Prediction Architecture

Given the variety of input features and the desire to have a form of model explainability, the initial machine learning architecture is a random forest (RF). Random forests have proven to excel in finding nonlinear patterns in data with limited initial assumptions and they also have the added benefit of the ability to add some sense of explainability. The results of the random forest model will be compared against a baseline physics-based model to determine the efficacy and potential enhancements over the baseline model that is implemented in production. More information on the baseline model can be found in Lin and Nispel (2022).

The random forest was trained on approximately 2500 examples where the power supplies operated on battery backup until it reached end of discharge and safely shuts downs. The features used in this initial model include the initial conditions of the physical properties at the beginning of the event and some meta data on battery information (type, age, count). The model then makes run time predictions on a test data set of 500 examples compared against the predictions from the baseline model. The results of the random forest are shown in Figure 10 where Figure 10a shows what features the model found most important during model training and Figure 10b shows the distribution of individual prediction errors in minutes compared to the baseline model.

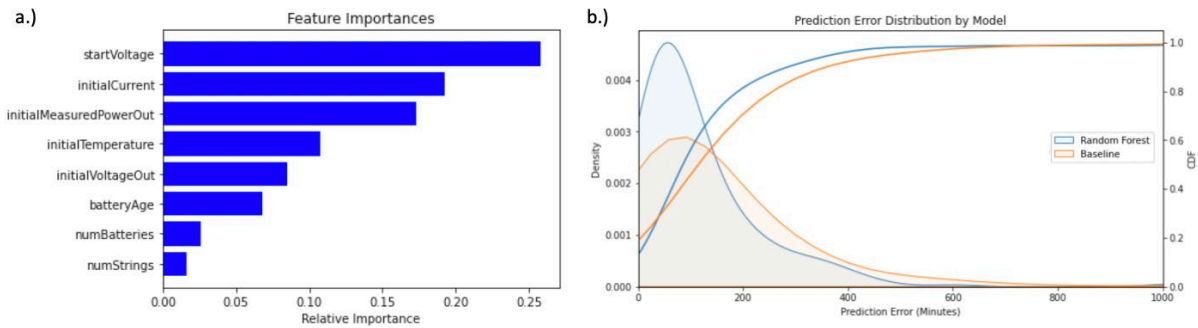


Figure 10 - Random Forest Performance. a.) Feature Importances, b.) Prediction Errors Compared to Baseline

As expected, the random forest model finds the most important features to be the average battery voltage in the first minutes of the event as well as the initial current draw and power consumption of the plant. These features combined essentially indicate the current capacity of batteries and how demanding it is to power the plant. The next group of features which are also found to be important are the ambient temperature at the beginning of the event, the output voltage of the power supply and the battery age in years. The model indicates that the number of batteries and the number of battery strings are marginally important for predicting the run time, likely due the majority of samples having a similar configuration. The random forest model tends to produce smaller errors compared to the physics-based baseline model as shown by the smaller tail in the distribution. As the model is enhanced with more features, the distribution of errors is expected to continue to shift lower indicating improved performance.

As discussed previously a random forest can also provide a measure of explainability for individual predictions. Figure 11 displays an example prediction from a particular sample. One can see that the output includes the prediction, confidence interval and a form of insight into the prediction. In this case the predicted run time was 270 minutes while the actual run time was 222 minutes. Along with the raw run time prediction, the 50% confidence interval is derived from the distribution of all the individual decision tree predictions. The width of this confidence interval gives an indication of how confident the model is in its raw prediction. Finally for a random forest, the raw prediction can be thought of as mean prediction plus a contribution from all the feature values for a particular sample. In the case below the model's mean prediction from training data is 239 minutes and features with blue bars increase the predicted duration over the mean and red decreases the predicted duration relative to the mean. In this case the most significant contribution is due to the battery age field, in this case the batteries are only 1 year old, so relatively new, which is why it is increasing the predicted duration. These types of supplementary information in addition to a run time prediction can help the field action and prioritize more efficiently.




Figure 11 - Example Random Forest Prediction of Battery Run Time

This section presented an initial proof of concept model using random forests for predicting the battery run time for power supplies which performs on par if not slightly better than the current physics-based prediction model. More advanced features are being developed which are expected to significantly improve the predictive accuracy over the current models. Comcast hopes to use the results of this model (once fully validated) over the long term as an input to a proactive maintenance model. For example, this model could be run periodically during nominal conditions (in the absence of a commercial outage) to determine if a power supply would perform to spec in the event of an outage and therefore open a window to perform proactive maintenance on suspect power supplies.

6. Conclusion

Spending time and resources to go deep into the power supply environment and build a product to support and improve network powering reliability has the potential for huge resource savings. We now have visibility into customer impact during power outages by correlating backend systems and building powerful a front-end interface. The new systems and interfaces combined with real time power supply telemetry open the doors for machine learning models to significantly enhance the proactive and demand maintenance capabilities of field teams. Machine learning models can not only provide more accurate predictions but also allow for explainability of given predictions to further support planning efforts. These initiatives are rooted in providing an amazing customer experience and having a more resilient power supply network is helping Comcast continue towards that goal.

Acknowledgements

The authors would like to acknowledge our colleagues that contributed significantly to the success of this effort. Chris D'Andrea was instrumental to setting up the big data pipelines and exploratory data analysis of the events in preparation for machine learning models. Ilana Weinstein developed the first iteration of the random forest model and is continuing to make enhancements.

Cory Thompson was a key resource during the development of PSNB, CE and Find My Power to ensure the functionality, benefits, and accuracy were developed and working in production as designed. Alex Falcon led the strategy and prioritization of the tool user story, minimum viable product outcome and change management in the field.

Abbreviations

DOCSIS	Data Over Cable Interface Specifications
ML	machine learning
OSP	outside plant power
PS	power supply
PSNB	Power Supply Notebook
RF	random forest
SCTE	Society of Cable Telecommunications Engineers
SNMP	simple network management protocol
UI	User Interface

Bibliography & References

Anderson, R. & Burgett, J. (2014). *Improving Network Reliability Through Effective Power Management*. SCTE Cable-Tec Expo 2014, Denver, Co.

PNM Best Practices Primer: HFC Networks. CM-GL-PNM-3.1-V01-200506. Cable Labs, 2020.

Lin, K., & Nispel, M. (2022). *A New Model for Plant Power and Health Estimation*. SCTE Cable-Tec Expo 2022, Philadelphia, Pa.

Wolcott, L., Heslip, J., Thomas, B. & Gonsalves, R. (2016). *A Comprehensive Case Study of Proactive Network Maintenance*. SCTE Cable-Tec Expo 2016, Philadelphia, Pa.

Wolcott, L., O'Dell, M., Kuykendall, P., Gopal, V., Woodrich, J. & Pinckernell, N. (2018). *A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact*. SCTE Cable-Tec Expo 2018, Atlanta, Ga.

Managing the Data Firehose to Put Out Network Fires

A Technical Paper prepared for SCTE by

Jingjie Zhu

Senior Engineer
CableLabs

j.zhu@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Background	5
2.1. The Complexity of DOCSIS 3.1 Data Collection.....	5
2.2. CCF Architecture.....	5
2.3. CCF Technology Stack	7
2.4. Parallelization	8
3. Analysis of Pain Points and Challenges	8
3.1. Estimated Data Collection Performance	9
3.2. API Performance	9
3.3. Data Store Performance	10
3.4. Configuration.....	11
3.5. Deployment Challenges	11
4. Resource Usage and Efficiency of Data Collection.....	11
4.1. Computation	12
4.2. Networking	13
4.3. Data Collector's Storage	14
4.4. Data Collector's APIs	15
4.5. Parallelization	15
5. The Next-Generation CCF (ng-CCF).....	16
5.1. Technology Stack.....	16
5.2. Architecture	18
5.3. Data Collection Functions	20
5.4. Packaging.....	21
5.5. Configuration.....	22
5.6. Scaling.....	22
5.7. Performance Testing.....	24
6. Conclusion.....	28
Abbreviations	28
Bibliography & References.....	29

List of Figures

Title	Page Number
Figure 1 – CCF Architecture	6
Figure 2 – CCF's role in today's data-driven system	7
Figure 3 – PNM procedure runtime measurement	12
Figure 4 – gofiber's API benchmark (requests per second) https://docs.gofiber.io/extra/benchmarks	17
Figure 5 – ng-CCF's architecture	19
Figure 6 – ng-CCF deployment architecture (horizontal scaling)	23
Figure 7 – ng-CCF and CCF's API performance (requests per second).....	24
Figure 8 – ng-CCF and CCF's API performance (request latency)	25
Figure 9 – ng-CCF and CCF's data collection performance.....	26
Figure 10 – ng-CCF's CPU usage (16-core) with different numbers of tasks	27
Figure 11 – ng-CCF's memory consumption with different numbers of tasks	27

List of Tables

Title	Page Number
Table 1 – Downstream OFDM RxMER SNMP network load.....	13
Table 2 – Network load estimations for concurrent data collection of OFDM RxMER from 1 million CMs (duration: 5 seconds).....	13
Table 3 – PPS estimations for concurrent data collection of OFDM RxMER from 1 million CMs (duration: 5 seconds).....	14
Table 4 – ng-CCF data collection functions.....	20

1. Introduction

In today's network, data plays a very important role in helping the operators gain more visibility into their networks to make their networks more reliable and more performant. For example, for Cable Modems (CMs), by calculating robust DOCSIS 3.1[®] profiles using Profile Management Application (PMA), impairments are mitigated, and the OFDM/OFDMA channel's performance is maximized. The profiles calculated by PMA can also be used to target low performing Cable Modems.

For other remote devices such as the Remote PHY Device (RPD) and Remote MAC Device (RMD), YANG modeling language is used in developing their north-bound data models and interfaces, which allows Google Remote Procedure Call (gRPC) Network Management Interface (gNMI) Streaming Telemetry to be easily implemented for advanced real-time device monitoring, largely reducing the probability and duration of service disruption events.

The complexity of data collection itself is also increasing as new devices support more sophisticated measurement functions. For instance, comparing to DOCSIS 3.0 data collection, collecting data from DOCSIS 3.1 CMs requires following much more complex procedures, increasing the amount of resources that the data collector uses and the number of states the data collector needs to track during the data collection process. Therefore, five years ago, CableLabs developed the first Common Collection Framework (CCF) [1],[4] as the initial work of diving into DOCSIS 3.1 data and shared its source code with the industry. The goal was to provide a reference implementation to converge the south-bound data collection interfaces, which are the interfaces for working with devices such as CMs and CMTSs, automate the data collection procedures and the handling of states, and provide standard interfaces to applications on the north-bound of the data collector. Its original goals have been achieved as it has provided operators and vendors a well-documented reference implementation and has served as the go-to data collector in many small-scale trials such as lab trials and limited field trials [2].

However, in the past field trials, CCF's performance wasn't impressive. When collecting data from around 8,000 DOCSIS 3.1 CMs, CCF spent more than 1 hour to complete the tasks even when it's multi-processed and was occupying 100 percent of the CPU resources on the data collection server. While acceptable for trials and development sandboxes, and could be scaled, it was not as scalable as we wanted.

Considering the scale of the whole network where some operators may have many millions of CMs deployed, and the increase in data collection frequency and the number of measurements, a better reference design of CCF is much desired.

In this paper, we use CCF as an example to identify and analyze potential performance bottlenecks and other pain points that could exist in today's network data collectors. And we share the experience of building the Next-Generation CCF (ng-CCF) to tackle each pain point and overcome scalability challenges. We hope the experience we share could provide references to others who are looking for such a data collection tool, on their way of building their own, or seeking for ways to improve the performance of their existing data collection tools. And we would like to share the new software, ng-CCF, with the cable industry as a reference design.

2. Background

2.1. The Complexity of DOCSIS 3.1 Data Collection

Since DOCSIS 3.1 technology was developed, multiple advanced measurements have been added to the CMs to collect and upload comprehensive physical layer metrics. Such data include:

1. CM downstream orthogonal frequency-division multiplexing (OFDM) symbol capture
2. CM downstream OFDM channel estimation coefficients
3. CM downstream OFDM constellation display
4. CM downstream OFDM receive modulation error ratio (RxMER) per subcarrier
5. CM downstream histogram
6. CM upstream orthogonal frequency-division multiple access (OFDMA) pre-equalization
7. CM upstream OFDMA pre-equalization last update
8. CM downstream OFDM forward error correction (FEC) stats
9. CM downstream spectrum analysis (full-band capture)

These test and query results provide rich information of the physical layer of the access network and are the fundamental requirements of advanced applications such as PMA. They require the data collector to perform multiple sequential simple network management protocol (SNMP) set steps on CMs for each test, and integration with trivial file transfer protocol (TFTP) servers for reading data uploaded by the devices.

On the CMTS side, DOCSIS 3.1 PNM results often require information that are challenging to gather or configure manually, such as the interface index numbers of OFDMA channels. These interface index numbers are often used as unique channel identifiers but are usually a reference number pre-determined by the CMTS and are offered through SNMP only, which suggests that an ideal data collector should automatically collect and prepare such intermediate information prior to data collection tasks that have dependencies on it.

The above aspects make data collection of these advanced PNM measurements significantly more sophisticated than data collection of traditional, common metrics where the collector usually queries the devices with one SNMP get or SNMP walk step for each data type and does not need to manage states.

As of today, this problem is solved as there are existing data collectors that are capable of handling the complexities, such as the first generation of CCF. But it has an impact that it encourages the data collectors to be highly concurrent for simplicity and scalability and to be microservice-like applications for the ease of scaling horizontally.

2.2. CCF Architecture

To allow flexible deployment, CCF version 2 was designed to consist of two microservices at a high-level: the Rest Agent (RA) and the Workflow Controller (WC). For the actual data collection tasks, modular “drivers” are introduced as plugin-like programs in CCF to provide extensibility and support rapid development. The RA provide hypertext transfer protocol (HTTP) APIs for the applications to use, and the WC handles the lifecycles and states of each individual data collection “driver” in parallel; the communication between RA and WC are done through

RESTful API calls. With this architecture, CCF can be scaled horizontally by applying multiplications of RA, WC, and coupled RA and WC instances.

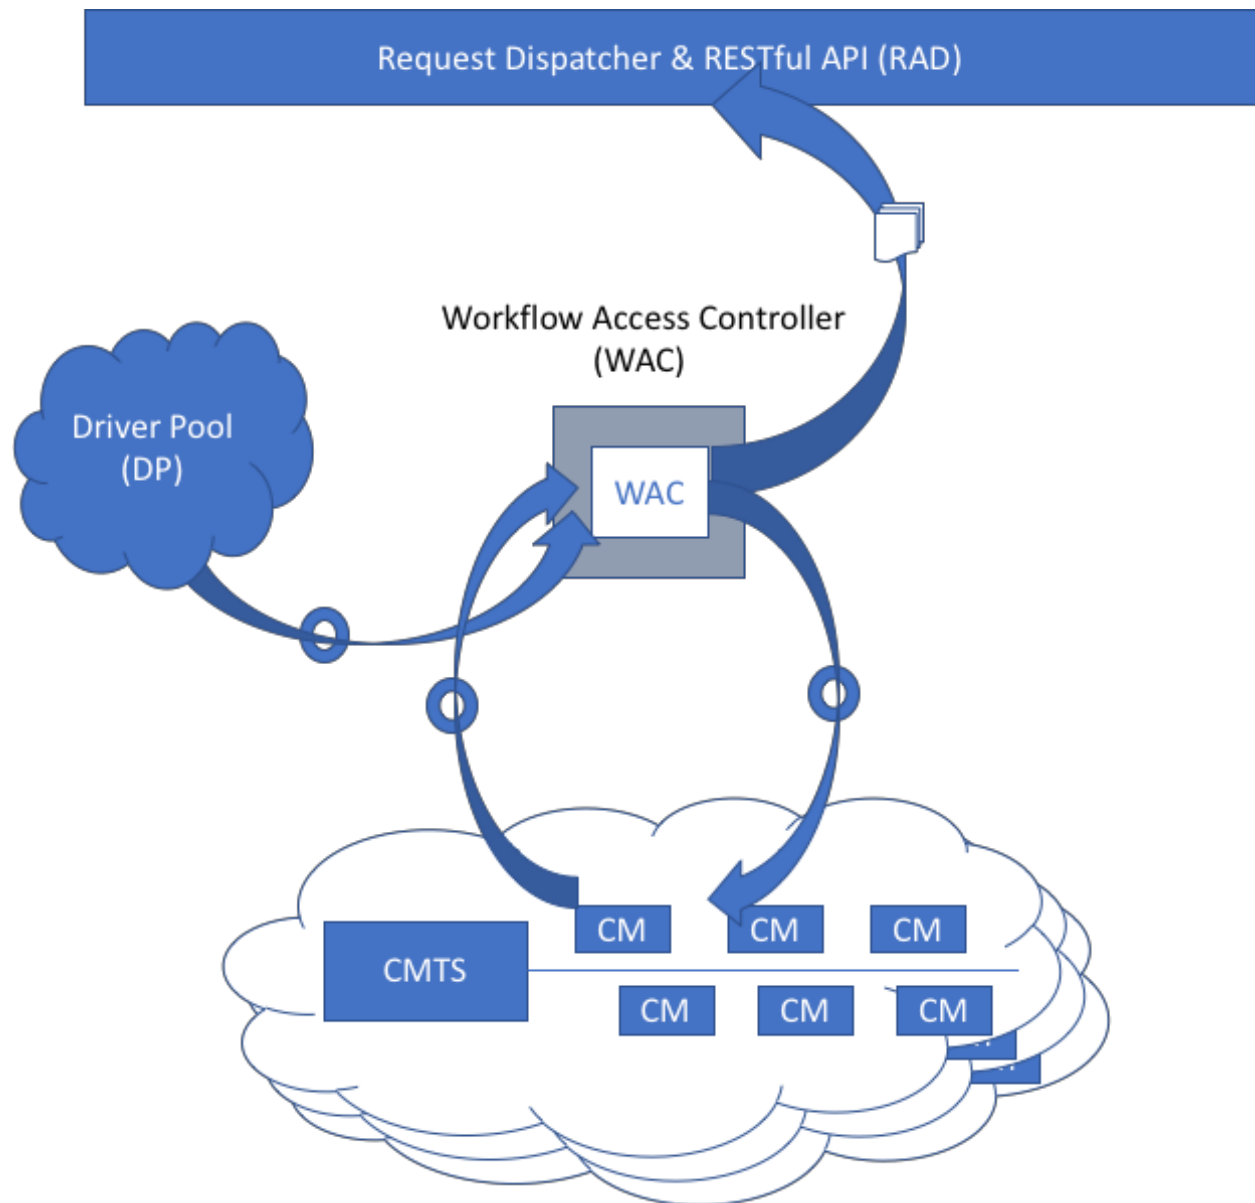


Figure 1 – CCF Architecture

CCF also integrates with external resources such as a local/remote TFTP server for gathering CM and CMTS uploaded PNM files, and an in-memory or filesystem datastore for storing data collection states and results. This architecture has been proven to provide benefits as applications can be easily built upon CCF's abstraction layer and its common APIs, while CCF scales the underlying interfacing activities with the network elements and provides protection to the devices' resources as the duplicated data collection is avoided and the same data is reused by different applications.

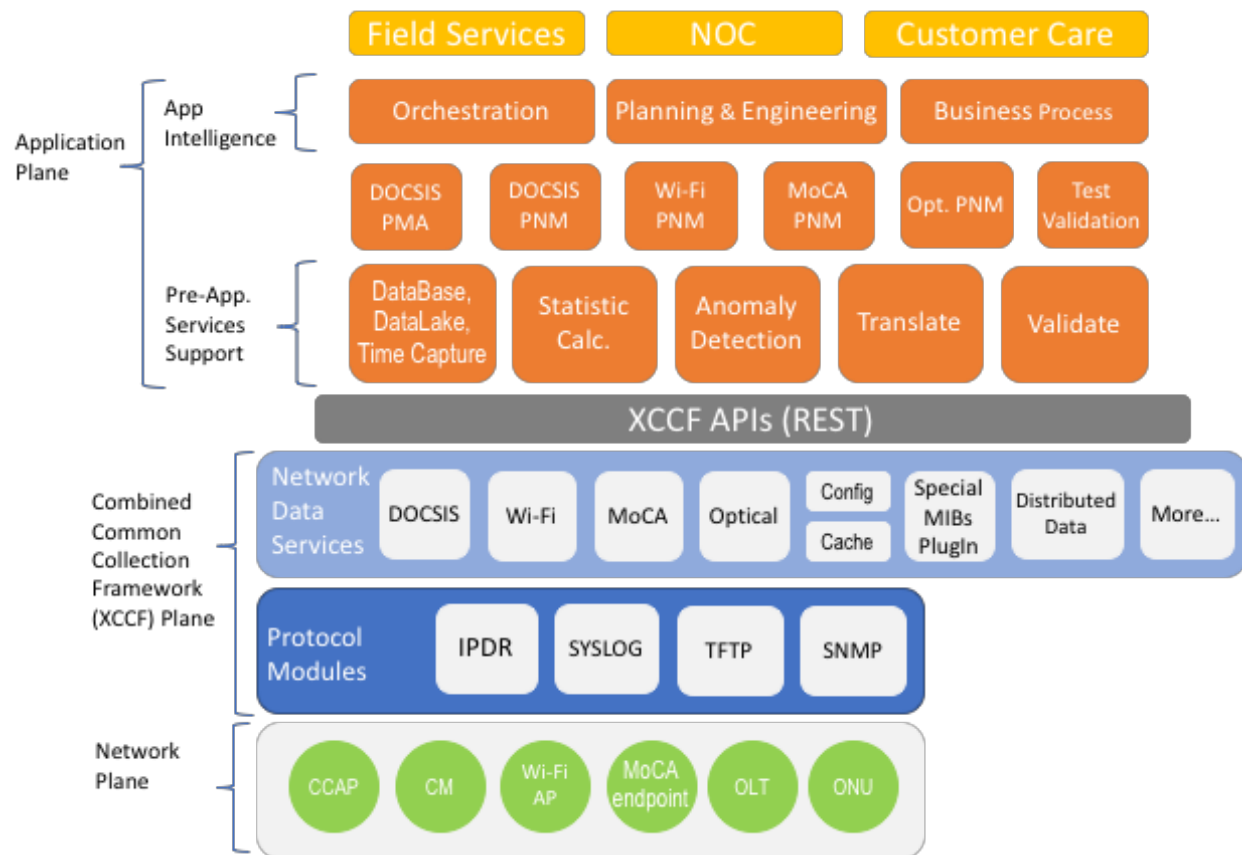


Figure 2 – CCF's role in today's data-driven system

Since the CCF is not intended to be coupled directly with a database, it's often deployed with a data collection scheduler to manage data collection cycles asynchronously, retrieve and decode the data collected by the CCF, and store the decoded data to a data service for applications to use. By doing this, the CCF only stores temporary data from the network and becomes straightforward to manage and maintain.

The drivers of the CCF perform individual data collection tasks that are designated for different devices and different measurements. The driver layer provides a simplified framework for users to easily add or modify data collection processes because the complexity of handling worker states, multiprocessing, and storing the data etc. are handled by the upper layers within the CCF. The drivers are also easy to develop and test as they can be tested individually from the command line.

2.3. CCF Technology Stack

The first two major versions of the CCF were developed in Python3. Python allowed us to quickly develop a working prototype and demonstrate the architecture and functionalities of the CCF. The HTTP APIs of the CCF were implemented with the Flask[9] Python library which provides a simple way to define and develop RESTful APIs. For communications between the

CCF's microservices, namely the RA and WC, the CCF uses Python3's requests [10] library to perform HTTP API calls.

For data storage, the CCF uses an in-memory cache to store temporary data as the default option. As an alternative, the CCF can use the Linux filesystem as a persistent data storage. This approach makes it easy to configure for a quick setup, but also provides persistent data storage options to the users.

For SNMP functionalities, the CCF uses a library python3-netsnmp which provides Python3 bindings for the NET-SNMP C library. Because of this, the SNMP dependencies are not portable and have to be pre-compiled for the system or compiled on the running system.

Last but not least, because DOCSIS 3.1 PNM measurements require the devices to upload the encoded measurement results to a TFTP server, the CCF integrates with an external TFTP server through the filesystem. This requirement allows the CCF to flexibly integrate with any TFTP server by pointing to their upload file directories. However, this method may have potential performance costs as the uploaded files are searched and identified using their filenames on the filesystem.

2.4. Parallelization

Because the CCF was designed to handle many data collection tasks in parallel for efficient data collection, multithreading/multiprocessing/concurrency is required in CCF's implementation. To simplify the driver layer, each driver runs single-threaded tasks, and multithreading/concurrency is handled by the WC.

Because Python has a Global Interpreter Lock (GIL) and it prevents the threads in Python from being "real" threads that use computing resources from multiple CPU cores, multiprocessing is needed for true parallelization that can use multiple cores from modern CPUs. However, in Python, this comes with significant CPU and memory overhead which prevents the CCF from efficiently handling a large number of concurrent data collection tasks using limited machine resources. The Python CCF implementation worked around this by starting a limited number of subprocesses and concurrently handling tasks assigned to each individual subprocess in event loops. This approach significantly helped improve the CCF's resource efficiency when it comes to parallelization. However, based on the trial results, this approach couldn't scale as optimally as desired, so higher efficiency is needed for handling a very large amount of concurrent data collection tasks.

3. Analysis of Pain Points and Challenges

During the past 5 years of using the CCF, helping operators and vendors configure and use the CCF, and building applications around the CCF at CableLabs, we've identified several pain points and challenges that are worth sharing and may be helpful for others to identify similar issues in their data collectors. These pain points and challenges are discussed in the following sections.

3.1. Estimated Data Collection Performance

The first challenge we identified is the CCF's ability to handle data collection for a large number of network devices. To understand the performance challenges at the high level, we can start with estimating the resource requirements of the data collector with the data collection requirements. An example of today's data collection requirements is:

1. the number of devices is around 10 million (DOCSIS 3.1 CMs)
2. a 6 hour data collection interval is required
3. 1 or 2 PNM data metrics are collected

However, in the foreseeable future, the following requirements may be asked given the increasing demand for field data, and the increasing number of deployed CMs, internet of things (IoT) devices, and more.

1. the number of devices is around 50 million
2. the desired data collection interval is 1 hour
3. multiple PNM data metrics are collected

With these future requirements, one can estimate that the data collection system will be required to collect data from around 56,000 devices every second on average. Based on the CCF's performance we observed in the previous field trials where 1 CCF instance spent more than 1 hour to collect Receive Modulation Error Ratio (RxMER) per subcarrier data from around 8,000 CMs, it infers that around 25,000 CCF instances will be required to handle this target workload to complete the data collection tasks within the required interval. This estimated number is overwhelmingly large, and it's asking for the computing power of a data center for a straightforward task of collecting data from network devices in parallel.

Although the CCF is considered a reference design, its performance is not ideal. Even if we consider a potential 10 to 100 times performance improvement for CCF, it will still ask for hundreds if not thousands of servers or virtual machines (VMs) on the cloud to be dedicated to data collection tasks. As the data collection demands continue to ramp up, it could only become more challenging for data collectors such as the CCF to catch up. Not only the resource consumption and cost of such data collection applications is significant, but it also increases the workload and complexity of managing such a large number of servers or VMs, not to mention databases.

3.2. API Performance

Another important aspect for microservice performance analysis is the application programming interface (API) performance. The CCF APIs are the interface for north-bound applications, these RESTful APIs are often called frequently during the data collection sessions as the north-bound applications may continuously check the data collection status of each individual measurement or each batch of measurements. This often results in the APIs being called thousands of times if not more, during the data collection sessions.

The baseline of how many API calls a performant microservice should handle varies by the context and requirements of the application. However, from the API performance testing results, the CCF could only handle up to 300 requests per second, which is significantly less than an ideal number.

The limited performance of the CCF's representational state transfer (REST) APIs can cause the north-bound applications to hang on measurement status checking requests. Low performance APIs could also cause high CPU or input/output (I/O) usage by the application. In CCF, the CPU usage spiked to 100% during the API performance test, which indicates that the CCF's HTTP APIs have a low CPU resource efficiency. When a large number of data collection tasks is being run, and when the north-bound application checks the measurement status frequently, the inefficient APIs can introduce issues to the entire data collection system by competing with other processes in the data collection application on CPU resources.

3.3. Data Store Performance

The CCF comes with an internal data store for saving and managing states, configurations, restricted amounts of collected data, and any other intermediate information. Usually, the CCF is configured to work with the following 2 types of data stores:

- filesystem
- in-memory

Because most web services are I/O bound, it's important to understand the CCF's data stores' performance and identify potential disadvantages of them.

Both data stores keep track of historical data and are implemented with in-memory file indexes. The use of in-memory file indexes makes data operations, such as insertion, deletion, and searching to be efficient. However, accumulating historical data increases the data operation costs over time. When working with hundreds of thousands of data entries, the performance impact is significant. This could further affect the CCF API's performance because slow data operations could shift the CCF APIs from CPU bound to I/O bound, further reducing the number of APIs the CCF can support during data collection sessions. We've observed this causing performance issues on long-running CCF instances in the lab setups and field trial setups.

Because the CCF's historical data storage has not showed value in multiple practical lab and field trials as the CCF has never been used as a primary data storage service itself, it is suggested that the fundamental design of the CCF's data store should be changed to offer improved data store performance. Removing historical data entries could be one of the improvements. In addition, compared to modern databases and caches, such as postgresql and redis, it doesn't provide benefits to implement the CCF's own data store while not taking advantages of the well-adopted data stores, especially when the performance differences are large. Replacing the CCF's data store with well-adopted, open-source databases or caches could further improve the CCF's data I/O and storage performance.

3.4. Configuration

Another pain point we've identified is that the configuration of the CCF requires knowledge that relates to Linux, software engineering, and networking. The prerequisite knowledge has blocked many users from setting up fresh CCF instances quickly without having to come to the developers of the CCF with questions. This situation also renders the CCF's setup documentations difficult to understand for users who are new to this field.

In addition, a typical set of the CCF's configuration files contain around 60 lines of configurations in JavaScript object notation (JSON) format, which further introduces work for the users to build a fresh CCF setup. Each of the decoupled microservices of the CCF (RA and WC) requires a separate configuration file, which adds difficulties and challenges to configuration and debugging for the users.

Considering the purpose and overall complexity of the CCF, the configuration and setup of it has been a major pain point since we shared the CCF with the industry and should be significantly simplified.

3.5. Deployment Challenges

The deployment challenges of the CCF are the collective outcome of the issues and pain points discussed in the previous sections. The performance of the CCF determines the size of the infrastructure that hosts CCF for large scale data collection. With CCF's performance being non-ideal, it's estimated to require a significant amount of computing resources for field data collection from millions of devices, which could add cost, maintenance work, and overall complexities to the data collection system.

In addition, the dependencies of the CCF are not compiled with the CCF's source code nor statically linked, making it complex to manage all of the CCF's dependencies in an internet-less deployment environment. Also, as the CCF is developed in Python and it depends on the NET-SNMP C library, the host of the CCF is required to run a complete Linux operating system which introduces overhead in CPU and memory usage for running the operating system (OS) and OS processes. When containerizing the CCF, this could result in large CCF images and heavy-weight containers.

Finally, the sophisticated CCF configuration process not only makes its instances difficult to set up, but also makes it harder to locate issues for the users to debug.

4. Resource Usage and Efficiency of Data Collection

As discussed in the previous sections, a high-performance, scalable, and easy to configure and deploy data collector is highly desirable to be the foundation of future network maintenance innovations and cost savings as the data collection demands continue to grow. The initial version of the CCF did not meet the requirements based on the pain points discussed in the previous sections. To build a data collector that could meet these requirements, we start from analyzing

the resource usage of common network data collection tasks, and then identify potential solutions. This analysis is discussed in the following sections.

4.1. Computation

First, understanding what computation tasks the data collector is responsible for during the data collection session is an important step to estimate how efficient the data collector could become.

During a network data collection session, the majority of computations happen remotely on the devices. For example, downstream RxMER per subcarrier data is measured, encoded, and uploaded by each CM. During the data collection session, the data collector is only responsible for facilitating and managing multithreaded tasks, sending requests to the devices, and waiting for the measurements to complete. An individual data collection task should use a negligible amount of CPU resource for most of the time during its lifecycle, which makes it promising that the data collector could become highly CPU efficient and could handle a very large number of data collection tasks at the same time if the concurrency is done efficiently.

Figure 3 shows the roughly measured procedure runtime of performing the DOCSIS 3.1 RxMER per subcarrier measurement and the DOCSIS 3.1 Spectrum Capture measurement.




Figure 3 – PNM procedure runtime measurement

We tracked how much time the measurement procedure actively spent for triggering the data collections and how much time it spent on idling and waiting for the collections to complete. In the results shown in Figure 3, the measurement procedure only spent negligible amount of time triggering the measurements compared to its idle time. During its idle time, the computing resources should be made available for other processes.

4.2. Networking

When collecting data from network devices, the data collector uses the server's network connection to query for data or trigger asynchronous measurements on the remote devices. For example, this procedure could involve sending user datagram protocol (UDP) packets (for SNMP) to CMs and CMTSSs. To provide a reference data point, the measured number of UDP (SNMP) packets sent to the Cable Modem during a downstream OFDM RxMER measurement session is shown in Table 1.

Table 1 – Downstream OFDM RxMER SNMP network load

Type	Number of Packets	Total Size	Direction
SNMP GET request	9	828 bytes	downstream
SNMP GET response	12	12,369 bytes	upstream
SNMP SET request	2	360 bytes	downstream

The SNMP GET responses transmitted the most amount of data because the data collector periodically checked the measurement status on the device by walking its measurement status MIBs. The entire measurement spent roughly 5 seconds to complete, which can add up quickly.

For RxMER measurement, the data collector performs SNMP SET on 6 MIBs; however, these SNMP SETs can be put into a single SNMP packet to reduce network loads. In the above measurement, the SNMP SETs were completed by sending 2 SNMP packets to the Cable Modem, resulting in 360 bytes of network usage.

If a data collection task uses longer time to complete, the number of SNMP GET request/response packets will increase as the data collector waits for longer durations while continuously sending SNMP GET requests to check the measurement status.

Based on the measured network load of downstream RxMER data collection from one CM, we can estimate the total network load for collecting downstream RxMER data for 1 million CMs at the same time, as shown in Table 2.

Table 2 – Network load estimations for concurrent data collection of OFDM RxMER from 1 million CMs (duration: 5 seconds)

Type	Traffic Rate	Direction
SNMP downstream network usage	1.9 Gbps	downstream
SNMP upstream network usage	19.79 Gbps	upstream

Both downstream and upstream traffic usage are high if the concurrent data collection for 1 million CMs is initiated from a single server.

In addition, packet per second (PPS) is another important statistic to estimate for network load as the servers and routers tend to be PPS bound instead of network throughput bound during large scale data collection. Table 3 shows the estimation of PPS loads.

Table 3 – PPS estimations for concurrent data collection of OFDM RxMER from 1 million CMs (duration: 5 seconds)

Type	PPS
SNMP downstream	2,200,000
SNMP upstream	2,400,000

The PPS numbers are also in the high range especially when we consider that most of the servers or VMs would be challenged to handle millions of packets per second. However, because the network usage is highly dependent on the protocols the data collector uses and the data collection procedures the devices implement, the data collector's network usage is primarily related to the deployment and scheduling scenarios. The data collector software may have limited room for improvement around network usage. Therefore, the above estimations are informative for the large-scale deployment of the data collector.

After the OFDM RxMER data collection from 1 million CMs completes, the CMs upload their measurement results to their designated TFTP servers. Assuming that the PNM files have an average size of 4 KB, 1 million CMs will upload roughly 4 GB of data to the TFTP server once the measurements complete. This would add additional network load to the system and transmit a considerably large number of packets per second through the network upstream as TFTP by default uses a packet size of 512 bytes. This calculation also leads to the analysis of data storage in the next section.

4.3. Data Collector's Storage

Data storage, or temporary data storage is another resource we should consider while building a large-scale data collection system. Assuming that a single instance of the data collector is responsible for data collection from 1 million devices, and each individual measurement data has an average size of 4 KB. This assumption results in 4 GB of storage usage per round and per measurement type during data collection.

For a high-performance data collector, an ideal design is to use dedicated long-term data stores for the collected data, and only cache the latest data collection and internal states within the data collector's temporary storage. This way, the performance of the data collector is less limited by I/O speeds as it would require less storage for data and could take advantage of in-memory caches for high-speed data access.

Ideal implementations of such caches could take advantage of open-source libraries such as BigCache for Golang, or take advantage of well-maintained and widely adopted caches such as Redis. Both are performant options and provide protections to the memory consumption through

configuring memory usage limits and aging off old data entries. And based on the above estimation of temporary data storage requirements, the data collector's cache can be fully implemented in memory while keeping the memory consumption within a reasonable range.

4.4. Data Collector's APIs

The data collector's APIs are responsible for the interactions between the data collector and northbound applications, schedulers, or the users etc. The APIs' performance is primarily determined by three factors:

- CPU processing speed
- Data I/O speed
- Network speed

The APIs' performance could be CPU bound if the API calls instantiate processing loads on the data collector which use significant amount of CPU resources. When the API calls require the data collector to communicate with the data store(s) very often, the performance of APIs could be I/O bound. And finally, when the data exchanges between the API callers and the data collector introduce significant network loads, the APIs' performance could be bound by the throughput of the network interface. The throughput of the network interface could be a limitation that's outside of the scope of the data collector's design considerations. However, minimizing the data payload sizes for API calls would be recommended. In addition, although the data collector could use high-speed in-memory caches as data stores, it's always recommended to reduce the number of direct data store hits from the API calls.

Ideally, for scalability, the data collector's APIs should not introduce high CPU loads and should focus on providing efficient connections to the data collector's data storage.

4.5. Parallelization

Highly optimized implementation of parallelization could drastically reduce the amount of resource the data collector uses for large-scale data collection tasks in deployment. In contrast, an inefficient implementation of parallelization could introduce a significant amount of memory overhead and CPU overhead. For example, the initial version of the CCF implements parallelization in Python using multiprocessing and event loops, which has the following shortcomings.

- High memory overhead introduced by multiprocessing in Python
- High CPU overhead introduced by increased Internal Procedure Calls (IPCs)
- Python as a programming language is not ideal in performance and efficiency

Fortunately, these shortcomings have already been addressed in modern programming languages such as Golang. With the Golang source code being compiled to native code and due to the goroutines and channels, highly efficient parallelization could be achieved for the data collector.

5. The Next-Generation CCF (ng-CCF)

Based on the analysis around resource usage and efficiency in the previous section, we now have clear objectives to develop the Next-Generation CCF (ng-CCF) [3] to resolve the pain points we've identified in the initial version of the CCF. The main objectives of the ng-CCF are to have

- High performance and efficient concurrency,
- High API performance,
- High data I/O performance,
- Efficient resource usage,
- Simplified installation and configurations,
- Simplified dependency management, and
- Implementation of all CCF's data collection functions and APIs for compatibility.

In addition to the main objectives, we also identify features that could be useful additions for ng-CCF:

- Having a built-in TFTP server
- Supporting the integration with remote/external TFTP servers
- Having a built-in data store
- Supporting the integration with external data stores
- Having a built-in SNMP client
- Supporting integration with gNMI targets/clients
- Cross-platform
- Horizontal and vertical scalability
- Small executables

Based on these objectives, we designed and developed the ng-CCF which has significant advantages and improvements compared to the initial versions of the CCF. The design, implementation, and performance analysis are discussed in the following sections.

5.1. Technology Stack

Referring to the analysis around resource usage and efficiency in the previous section, we decided to start the development of the Next-Generation CCF from completely rewriting the software in Golang as many objectives would be impossible to achieve if we build the ng-CCF based on the source code of the CCF which is in Python. This decision has a trade-off that the entire source code of the data collector needs to be rewritten, but it allows us to take advantage of Golang's ability to handle concurrency in a highly efficient way. This choice also allows the source of the ng-CCF to be compiled into a single statically linked executable which provides benefits to installation, configuration, and deployment of the data collector.

To implement the HTTP APIs for the ng-CCF, we chose to use Golang's Fiber [6] library instead of the built-in HTTP package because the Fiber library is built on top of Golang's Fasthttp [11]

package which provides superior API performance as shown in the benchmark results in Figure 4.




Figure 4 – gofiber’s API benchmark (requests per second)
<https://docs.gofiber.io/extra/benchmarks>

For data storage, we chose to define an interface for the ng-CCF’s data stores and implemented the interface with a built-in cache built on top of Golang’s BigCache library and a Redis client to communicate with the external Redis cache as 2 supported options. This interface can be conveniently implemented for the ng-CCF to integrate with any other types of data stores such as MongoDB or PostgreSQL. The built-in cache allows the users to start the ng-CCF without relying on external services, whereas the Redis cache allows multiple ng-CCF instances to share the same remote cache and make the cache accessible to be backed up, duplicated, or persisted. The Redis cache option is particularly useful in deployment because it largely reduces each ng-CCF instance’s memory usage and allows the users to host and manage the caches on dedicated servers. The built-in cache’s speed is bound by the time complexity of the in-memory data structure and the memory speed. And the Redis cache’s speed is bound by the speed of Redis, memory speed, Redis API call’s speed, and potentially network speed if the cache is remote.

For SNMP, we decided to not rely on OS dependencies such as NET-SNMP and chose an SNMP library gosnmp [7] which is fully implemented in Golang to make the executable portable. This SNMP library supports all SNMP functionalities the data collector needs, such as GET, SET, WALK, BULKWALK, and SNMPv3. Using this well-integrated library also allows the data collector to provide very detailed debugging messages for SNMP down to a per packet payload

level, which could be helpful for debugging the system when the devices don't respond as expected.

We also decided to give the ng-CCF a built-in TFTP server so that it becomes a completely self-contained software solution when there's need for a quick setup and trial. In the initial CCF, the TFTP server is an OS dependency and is integrated with the CCF through the Linux filesystem. In the ng-CCF, the TFTP server is implemented with Golang's tftp [8] library which allows the ng-CCF to process all uploaded files into memory without relying on system calls and using the slower hard drive.

In addition to the built-in TFTP server, the ng-CCF has TFTP, SFTP, and HTTP clients to handle different types of integrations with external TFTP, SFTP, and HTTP servers. For example, the ng-CCF can integrate with remote or external TFTP servers using its TFTP client. This ability could be particularly useful if there are already TFTP servers in deployment. Some CMTSs may implement an SFTP server for the applications to download PNM measurements instead of uploading the measurement results to a TFTP server. The SFTP client in the ng-CCF allows it to automatically switch between the TFTP client and the SFTP client based on the detected CMTS types. For integration with external HTTP services such as Prometheus, the ng-CCF driver can leverage its built-in HTTP client.

Finally, for concurrency, although Golang's goroutines are used in many submodules in the ng-CCF, we decided to employ Golang's "ant" package to build the primary data collection task pool to automatically manage task lifecycles and potentially achieve higher performance compared to using unlimited goroutines.

5.2. Architecture

The architecture of the ng-CCF is largely simplified compared to the initial version of the CCF. In the ng-CCF, there's no longer separate microservices that introduce communication overhead. To reduce the amount of computation, the architecture tries to reduce the amount of CPU processing and data store access when possible. The architecture drawing is shown in Figure 5.

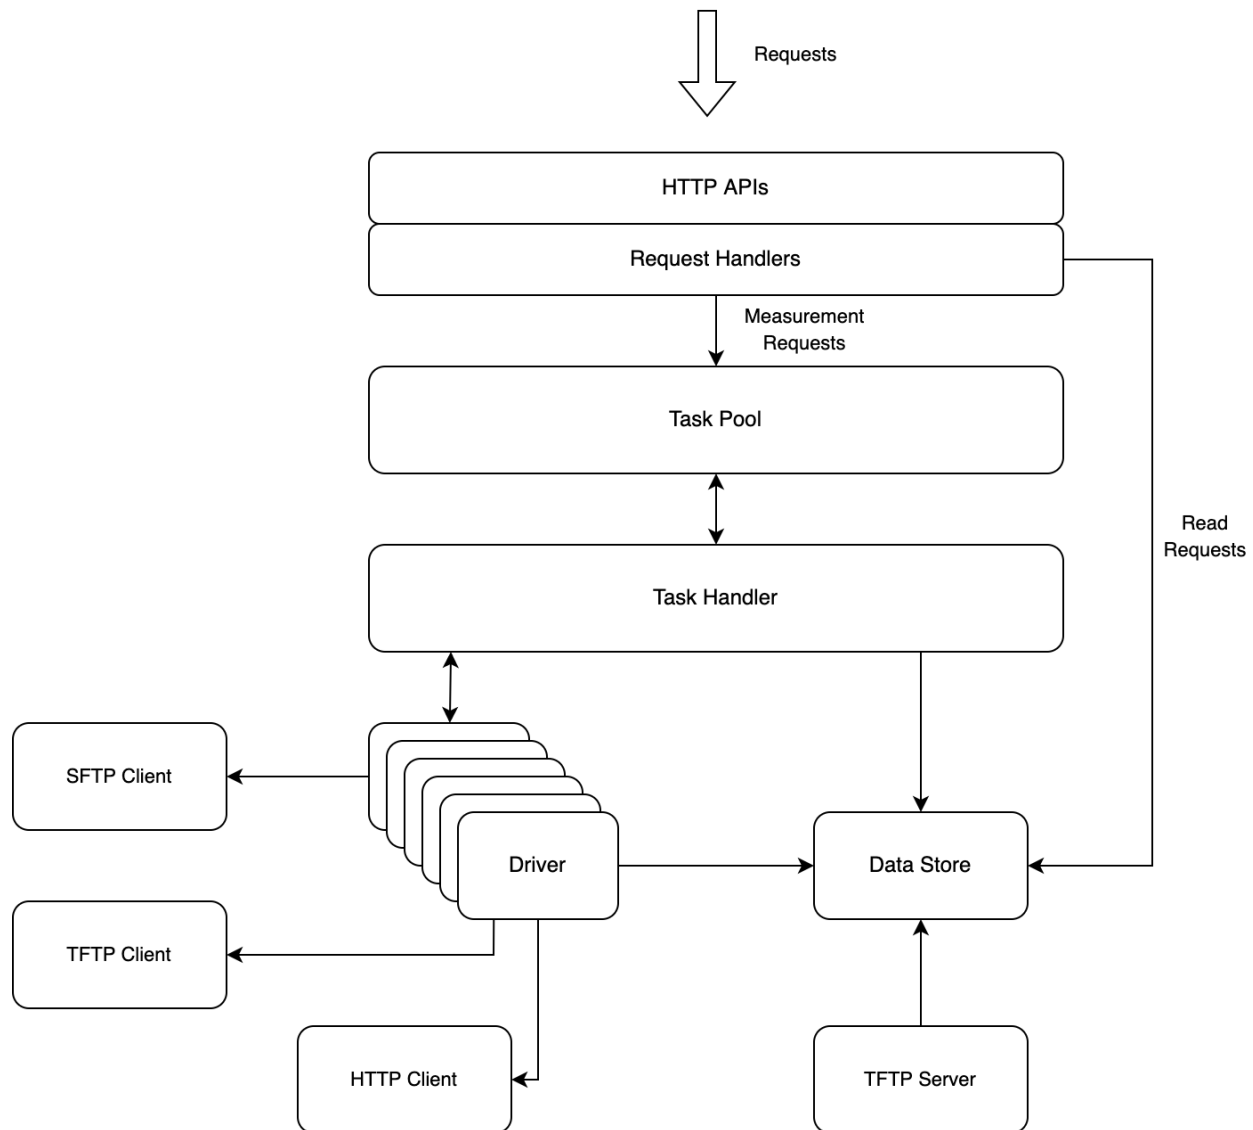


Figure 5 – ng-CCF's architecture

The ng-CCF's data store keeps the latest temporary data and serves as the state storage for data collection tasks and the storage for the TFTP server. The data store operations are thread-safe which makes the data store ideal for the goroutines to pass states and data blobs in addition to using Golang channels for local messaging between goroutines.

When a read request comes to the ng-CCF, the request handlers directly connect the request to the data store for the data request. Because the ng-CCF is designed to use in-memory caches, there's currently no need for an additional caching layer before the data store.

When a data collection and measurement request come to the ng-CCF, the request handlers create the description objects of the tasks according to the request and pass the task description objects to the task pool. The task pool then instantiates an individual, concurrent task handler for this specific task, and proceeds with executing the task and sub-tasks and managing their life

cycles. When the data collection tasks complete, the resulting data are collected from the drivers by the task handler and then stored into the data store along with the updates task states for future read requests.

In this architecture, the task pool handles the majority of the concurrency in the data collector, and the task handler manages sequential and parallel execution of individual sub-tasks and their life cycles, making the software easy to manage and maintain.

Concurrency can also be implemented within the drivers to boost data collection performance. For example, a driver can concurrently send SNMP GET and WALK requests to different MIBs for faster data collection. This is optional for the implementation of the ng-CCF drivers, and it is local to the drivers, which means it's isolated and modular and doesn't increase the overall complexity of ng-CCF's concurrency and internal messaging.

This architecture also makes it easy for testing each individual data collection drivers and testing the API handlers with mock drivers. The modular design makes the ng-CCF easy to maintain and update during long-term development and deployment.

5.3. Data Collection Functions

To make the ng-CCF compatible with existing applications that depend on the CCF and also improve its data collection capabilities, a wide variety of data collection drivers have been developed for the ng-CCF. The drivers provide functions to collect DOCSIS 3.1 specific data elements as well as DOCSIS 3.0 data elements. They also provide data collection functions from external services such as Prometheus, which can be used for integration with gNMI collectors. The current list of ng-CCF drivers is shown in Table 4.

Because the drivers are similar to plugins in the ng-CCF, new data collection functions can be conveniently added to the ng-CCF with concurrency, data storage, and API calls automatically handled.

Table 4 – ng-CCF data collection functions

Type	Description
CM OFDM downstream RxMER	A DOCSIS 3.1 specific data element that provides per subcarrier RxMER data of OFDM channels used by the CM
CM OFDM channel estimation coefficients	A DOCSIS 3.1 specific data element that provides per subcarrier channel estimation data of OFDM channels used by the CM
CM OFDM constellation diagram	A DOCSIS 3.1 specific data element that provides constellation diagram data of OFDM channels used by the CM
CM downstream histogram	A DOCSIS 3.1 specific data element that provides downstream power histogram data measured by the CM

Type	Description
CM OFDMA upstream pre-equalization	A DOCSIS 3.1 specific data element that provides the upstream OFDMA pre-equalization coefficients that the CM is using
CM OFDMA upstream pre-equalization last change	A DOCSIS 3.1 specific data element that provides the last adjustments to the OFDMA pre-equalization coefficients that the CM is using
CM downstream OFDM FEC summary	A DOCSIS 3.1 specific data element that provides a OFDM FEC summary for each individual modulation profile over a 10-minute or 24-hour time frame
CM upstream OFDMA RxMER	A DOCSIS 3.1 specific data element that provides per subcarrier RxMER data of OFDMA channels used by the CM
CM downstream spectrum capture	This measurement provides the full-band capture data of the CM's downstream spectrum
CM SC-QAM upstream pre-equalization	This measurement provides the upstream SC-QAM pre-equalization coefficients that the CM is using
CM capabilities	Collect and decode CM capability requests and responses per TLV 5 defined in the DOCSIS 4.0 MULPI specification
CM events	Collect CM device event logs, event times, and event IDs
CM OFDM channel topology	Discover the OFDM channel based logical topology
CM OFDMA channel topology	Discover the OFDMA channel based logical topology
Prometheus data	Collect any data from Prometheus APIs

5.4. Packaging

The source code of the ng-CCF is written in Golang compiled into a statically linked executable which contains all required dependencies, and no other software packaging process is required. This packaging allows the ng-CCF to run inside of a “scratch” docker container which is an empty container that has minimal storage overhead. The size of the ng-CCF’s “scratch” image is only negligibly larger than the executable’s size, making it extremely resource efficient in cloud deployment scenarios.

The original size of the ng-CCF executable is 15 MB. However, optionally, it’s possible to further reduce its size by using an executable packer such as Ultimate Packet for eXcutables (UPX). We used UPX to compress the ng-CCF’s executable, which doesn’t affect the requirements on the running system and doesn’t add any dependencies, yet it reduced the size of the ng-CCF’s executable from 15 MB to 3.7 MB which is only 24.67% of its original size.

Having such a compact executable could have benefits in deployment scenarios. Although it won't optimize memory consumption because the executable is auto-decompressed before running, it significantly reduces ng-CCF's storage footprint and potentially results in more efficient network usage during software updates, especially if we consider future development and extensions being applied to the ng-CCF where the executable size may continue to increase to hundreds of megabytes.

5.5. Configuration

Configuration challenges are a major pain point of the initial version of the CCF. Therefore, simplifying the ng-CCF's configuration is one of the high priority focuses of its development initiative.

The initial version of the CCF typically requires a certain level of knowledge in software engineering, Linux, and network engineering, and it requires 3 separate configuration files that consist of around 60 lines of JSON for the users to work through to get a minimal setup.

In the ng-CCF, because of the addition of a built-in TFTP server and a built-in data store, with default parameters, the user can start the ng-CCF using only one command and specifying only one command line parameter. The integration of ng-CCF and external TFTP servers and data stores is also largely simplified. For example, specifying an external Redis data store only requires three additional command line parameters, and replacing the built-in TFTP server with an external TFTP server only requires two additional fields in the API request payload.

From the trial experiences of the ng-CCF after its release, we've heard significantly less confusion regarding the setup and configuration of the data collector. The users have found it intuitive to start the data collector with one command, without having to work with the source code, installing dependencies and OS dependencies on offline machines, configuring microservices and debugging potential connectivity issues and package compatibility issues. The configuration improvement is an overall significant user experience improvement, and it largely reduces the friction of deployment of the ng-CCF as the data collector.

5.6. Scaling

Because the ng-CCF has a highly efficient concurrency implementation and one instance of the ng-CCF can fully utilize the computing resources on the host machine, it can be flexibly scaled vertically and horizontally depending on the use case needs.

Scaling the ng-CCF vertically requires computing resource upgrades on the host machine. Depending on the data collection requirements such as the number of devices, the number of measurement types, and the frequency of data collection, upgrading the computing resource on the host machine may not be viable if the data collection requirements exceed a limit. However, because of the ng-CCF's capabilities, this limit could be very high on a server with reasonable computing power. Therefore, the ng-CCF may have the potential to support large scale data collection using the computing power of a single server. The details and reference performance numbers are described in the following performance testing section.

Scaling the ng-CCF horizontally could be a more reasonable approach to consider in deployment scenarios. It's possible to balance the data collection loads by assigning dedicated ng-CCF instances to each CMTS's data collection needs. However, since the ng-CCF instances can share remote Redis caches for data storage and state storage, a better approach could be flexibly managing the number of running ng-CCF instances or containers based on the immediate data collection needs, and load balance by routing the ng-CCF's API calls to the ng-CCF instance pool. This approach draws an overall simpler and more flexible deployment system and can more efficiently utilize the computing resources as the granularity of load balancing becomes a single API call. Note that, in this approach, the Redis caches may need to be scaled to satisfy high volumes of access requests. An example drawing of such a deployment architecture is shown in Figure 6.

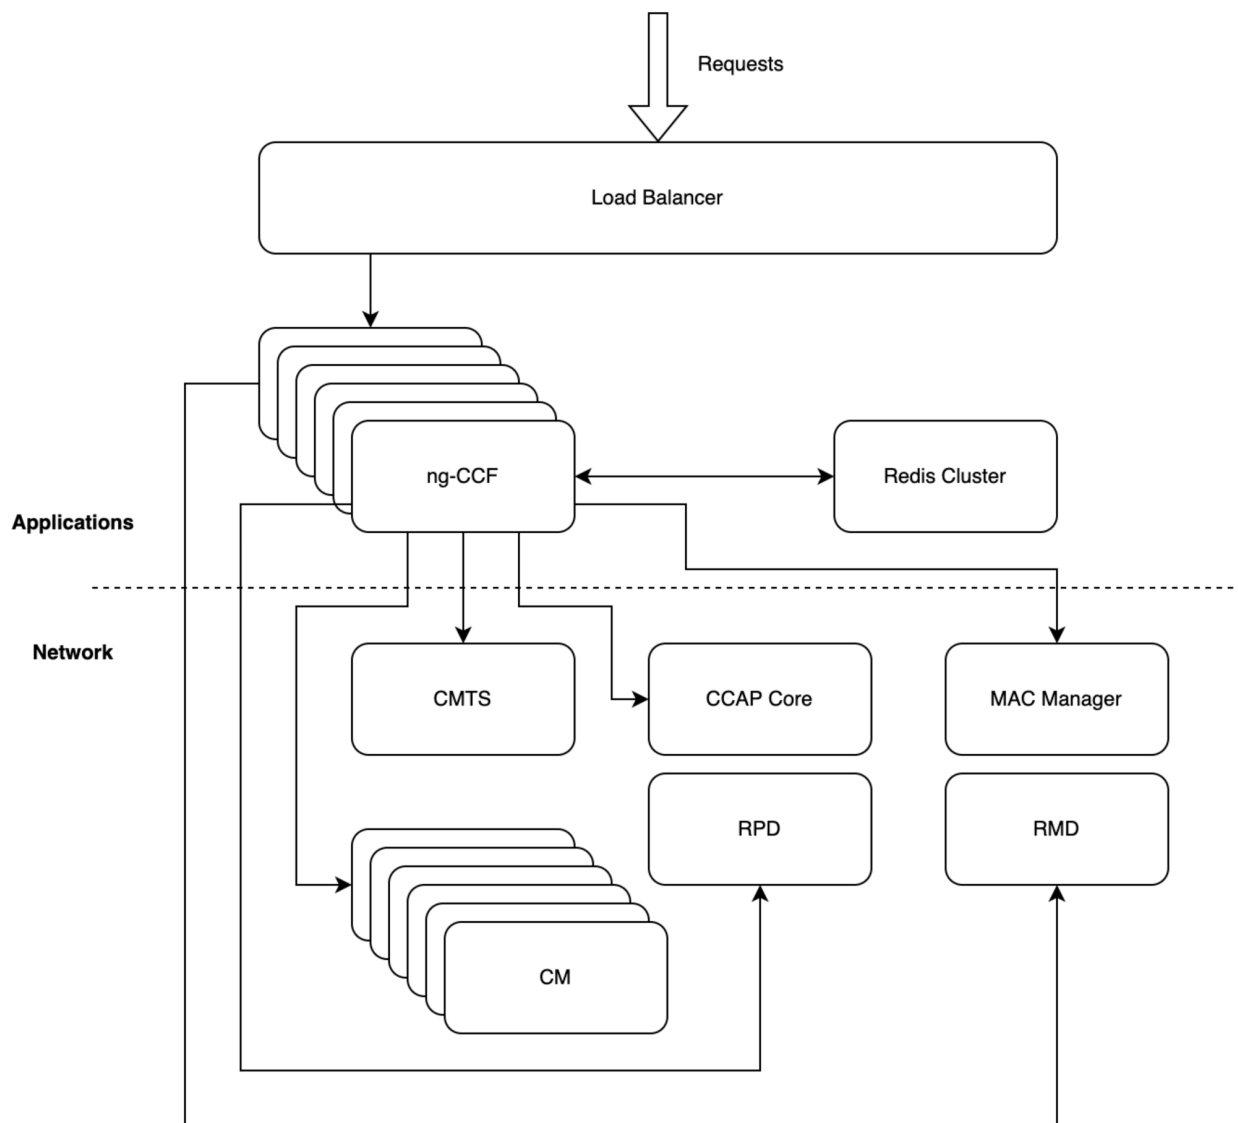


Figure 6 – ng-CCF deployment architecture (horizontal scaling)

5.7. Performance Testing

The performance testing of the ng-CCF focuses on three aspects:

- API calls handled per second,
- API call latency, and
- Mock data collection performance.

The number of API calls handled per second is an indicator of the API performance of the ng-CCF. And the API call latency is the indicator of the API responsiveness of the ng-CCF, which is another aspect of the ng-CCF's API performance. And finally, the mock data collection test uses a mock driver to simulate large scale data collection scenarios and measures the ng-CCF's data collection performance with a reasonable number of concurrent tasks and also pushes the ng-CCF to the limit to see how many concurrent data collection tasks a single ng-CCF instance could handle on a powerful server.

The tests were conducted on an Ubuntu 20.04 VM that's running on a 2021 Macbook Pro with 4 cores of the Apple M1 Pro processor and 8 GB of random-access memory (RAM). We used an open-source HTTP benchmarking tool, called wrk [5] (<https://github.com/wg/wrk>), running on a separate machine to start 12 threads with 400 concurrent connections to send API requests to both the CCF and the ng-CCF instances as fast as possible for 30 seconds. The measurement results are shown in Figure 7, Figure 8, and Figure 9.

12 threads, 400 concurrent connections, 30 second test duration




Figure 7 – ng-CCF and CCF's API performance (requests per second)

The initial version of the CCF only could handle 269.78 requests per second whereas the ng-CCF could handle more than 300,000 API requests per second, thanks to the Apple M1 Pro's high bandwidth memory. On an ordinary VM assigned with 4 heavily shared Xeon cores and a memory with lower bandwidth, the ng-CCF still was able to handle more than 110,000 API requests per second.

12 threads, 400 concurrent connections, 30 second test duration




Figure 8 – ng-CCF and CCF's API performance (request latency)

On the API latency measurement for the CCF and the ng-CCF, the difference is significant. The initial version of the CCF averaged more than 500 ms on API responses due to its usage of the Linux filesystem as its data store, and its API response latency peaked at almost 2 seconds. Whereas the ng-CCF averaged 7.4 ms of API response time and peaked at 137.46 ms which is possibly affected by garbage collection in Golang.

Mock CMs are designed to take 10 seconds to complete the mock PNM measurement. Hence the baseline of collection time is 10 seconds.

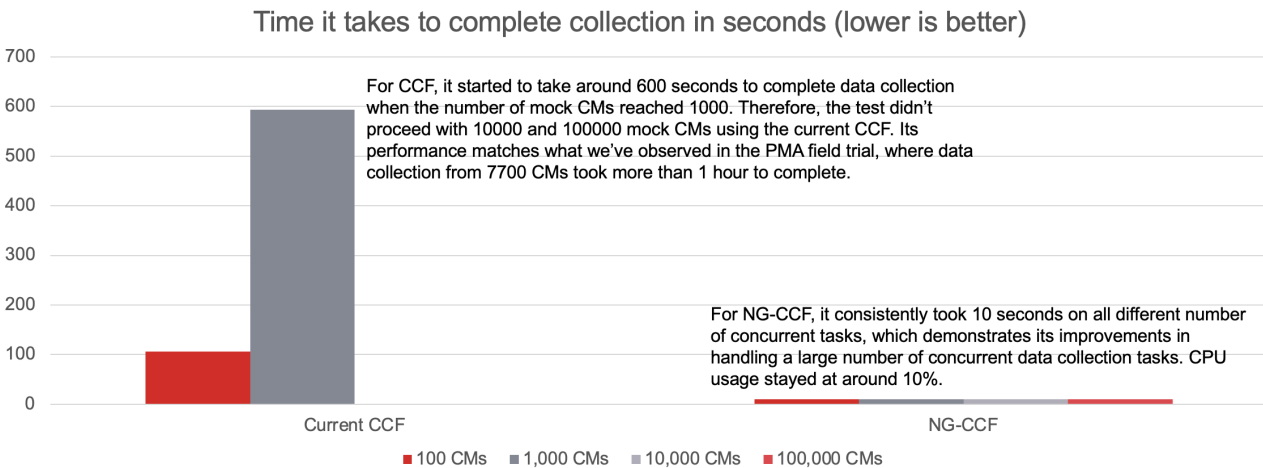


Figure 9 – ng-CCF and CCF's data collection performance

During the simulated data collection testing, we design each measurement to take 10 seconds to complete, most of the time used by the driver waiting for the device to complete the measurement. The initial version of the CCF matched its performance we observed during our first PMA field trial where it took more than an hour to complete measurement from about 7700 CMs. Meanwhile, the ng-CCF effortlessly handled 100,000 concurrent measurements while using only 10% of the 4 core CPU resource, and all tasks completed after the same 10 second wait time.

To find the limit of the ng-CCF for large scale data collection and estimate its requirements for vertical scaling, we conducted another test where one instance of the ng-CCF was hosted on a powerful workstation that has 256 GB of RAM and a 16-core Intel® Xeon® processor (Xeon® Gold 5218 CPU @ 2.30 GHz).

We let the ng-CCF run 0.5 million, 1 million, 2 million, 4 million, 8 million, and 10 million concurrent tasks during the test. Considering heap allocation bottlenecks, each task was designed to run 180 seconds to allow the ng-CCF to complete the allocation of large numbers of tasks.




Figure 10 – ng-CCF's CPU usage (16-core) with different numbers of tasks




Figure 11 – ng-CCF's memory consumption with different numbers of tasks

When the number of tasks surpassed 2 million, we started to observe slowdowns in the ng-CCF's ability to instantiate new tasks quickly. This is likely bound by heap allocation speed as the CPU usage was still low.

When the number of tasks reached 10 million, although the ng-CCF still didn't fully utilize the processing power of 16 CPU cores, the heap allocation delay became significant enough that 10

million was the maximum number of tasks the ng-CCF could allocate during a 180 second window.

When handling 1 million concurrent tasks, the ng-CCF used 9.75% of the 16-core CPU, and used 5.8 GB of RAM, making it promising to handle large scale data collection using only a few servers. When starting 10 million concurrent tasks on the ng-CCF, it used 66.41% of the 16-core CPU and used 68.8 GB of RAM at peak. The heap memory allocation speed could be a limiting factor for vertically scaling the ng-CCF any further. Therefore, it would be recommended to limit the number of concurrent tasks on each ng-CCF instance to less than 10 million.

6. Conclusion

In this paper, we shared the experience of how we identified the pain points and challenges in an existing data collector and analyzed the resource usage for data collection and potential approaches to improve the efficiency of a data collector. And we shared the details of how we designed and developed the Next-Generation Common Collection Framework (ng-CCF) to overcome the issues we've identified and demonstrated improvements in many different aspects such as improved performance and scalability, small and easily deployable executable, enhanced data collection functionalities, and significantly simplified configuration and setup process.

We hope to share our experience with the industry to help others target potential improvements that could be done in their data collectors. And we hope to share the ng-CCF's source code with the industry to help others take advantage of what we've developed. As the data collection continues to grow, a scalable, performant, and reliable data collector is highly desirable to be the foundation of future network maintenance innovations and cost savings.

Abbreviations

API	application programming interface
Bps	bits per second
CCAP	converged cable access platform
CCF	common collection framework
CM	cable modem
CMTS	cable modem termination system
FEC	forward error correction
Gbps	gigabits per second
gNMI	gRPC network management interface
gRPC	Google remote procedure call
HTTP	hypertext transfer protocol
Hz	hertz
JSON	JavaScript object notation
MER	modulation error ratio
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OS	operating system
PMA	profile management application

PNM	proactive network maintenance
RAM	random-access memory
REST	representational state transfer
RxMER	receive modulation error ratio
SCTE	Society of Cable Telecommunications Engineers
SFTP	SSH/secure file transfer protocol
SNMP	simple network management protocol
TFTP	trivial file transfer protocol
UDP	user datagram protocol
VM	virtual machine

Bibliography & References

- [1] CableLabs Proactive Network Maintenance Combined Common Collection Framework Architecture Technical Report, CL-TR-XCCF-PNM-V01-180814, August 14, 2018, Cable Television Laboratories, Inc.
- [2] Karthik Sundaresan, Jay Zhu, Mayank Mishra, and James Lin, “Practical Lessons from D3.1 Deployments and a Profile Management Application”, SCTE 2019
- [3] The Next-Generation Common Collection Framework (<https://code.cablelabs.com/CCF/ng-ccf>)
- [4] The Common Collection Framework (<https://code.cablelabs.com/CCF/dccf>)
- [5] Wrk: a modern HTTP benchmarking tool (<https://github.com/wg/wrk>)
- [6] Fiber: a web framework for Go (<https://github.com/gofiber/fiber>)
- [7] gosnmp: an SNMP library written in Go (<https://github.com/gosnmp/gosnmp>)
- [8] tftp: TFTP server and client library for Go (<https://github.com/pin/tftp>)
- [9] flask: the Python micro framework for building web applications (<https://github.com/pallets/flask>)
- [10] requests: a simple, yet elegant, HTTP library (<https://github.com/psf/requests>)
- [11] fasthttp: fast HTTP package for Go (<https://github.com/valyala/fasthttp>)

Matter

What It Is, How It Works and Why It Matters To The Cable Industry

Technical Paper prepared for SCTE by

Haefner, Kyle, Ph.D.

CableLabs

k.haefner@cablelabs.com

Haque, Asad

Comcast

asad_haque@comcast.com

Page, Jason

Charter Communications

jason.page@charter.com

Table of Contents

Title	Page Number
1. Introduction	4
1.1. General Overview of IoT	4
1.2. Lack of Interoperability - Lack of Open Standards	4
1.3. Consumer Branding/Labeling/Proximal Interoperability	5
1.4. Lack of Security/Uniform Standard for Security	5
1.5. Introduction to Matter	6
1.6. Impact for Cable Operators	7
2. Matter Architecture	7
2.1. Transports and Network Stack (Wi-Fi, Thread, BLE, IPv6, DNS-SD)	7
2.2. Fabrics	8
2.3. Data Model	9
2.4. Interaction Model	11
2.5. Bridging - Current Networks and Legacy Devices	11
3. Commissioning	12
5. Security	16
5.1. Secure Communications	16
5.2. Replay Prevention	17
5.3. Secure Group Communications	17
5.4. Access Control	18
5.5. Software Update	19
6. Administration	19
6.1. Operational PKI	19
6.2. Configuring Access Control	19
6.3. Fabric Management	20
6.4. Group Management	20
6.5. Software Update	20
6.6. Multi-Admin	20
7. Matter and Operators	21
Abbreviations	23
Bibliography & References	23

List of Figures

Title	Page Number
Figure 1: IoT Quilt	4
Figure 2: Matter Protocol Layers	7
Figure 3: Router Architecture	8
Figure 4: Matter Fabrics	9
Figure 5: Device Types and Clusters	10
Figure 6: Bridging Architecture	11
Figure 7: Matter Commissioning	12

Figure 8: DCL Architecture 14
 Figure 9: DCL Network 15
 Figure 10: Secure Group Communications..... 17
 Figure 11: Group Management..... 20
 Figure 12: Multi-Admin Sequence..... 21

1. Introduction

1.1. General Overview of IoT

A new phenomenon is sweeping our lives, involving both small and large devices communicating with each other, and providing rich information that enhances how we live, work, and play. We can tell from miles away if our home is at the right temperature if the porch light is on or if the front door is locked. This is possible because thermostats, bulbs, and door locks can communicate their status, receive commands and act on them. These devices employ a range of network protocols to communicate over the Internet. By using the Internet as the medium, devices can be placed in geographically dispersed locations and still provide near real-time status or environmental reading.

However, this type of convenience comes with a price i.e., how can IoT devices be adequately secured given their constrained nature? How can they fend for themselves in a hostile network environment? These devices are often inexpensive, easy to use, and quick to set up. However, they lack the proper hardware and software security needed to ensure data privacy and integrity.

1.2. Lack of Interoperability - Lack of Open Standards

Lack of standards is manifested in consumer frustration, manufacturer headaches and service provider anxiety. OEM's need to support multiple "stacks" and service providers need to build complex integration models to support thermostats from five different manufacturers. The high-level diagram below (Figure 1.) shows current silos and industry fragmentation in IoT landscape.

The lack of widely adopted IoT application standards is a significant barrier to mass adoption of IoT devices, especially by consumers. The current wall garden approach not only stymies innovation by developers but also acceptance by consumers. The lack of standard results in increased investment by developers because they must develop for multiple application stacks. Similarly, consumers have to invest in new devices every time they switch to a new IoT ecosystem because their current investment in devices will be incompatible.

The disparate protocols create an IoT quilt consisting of a patchwork of communication standards and device models leaving devices unable to interact with each other based on which part of the quilt they belong.

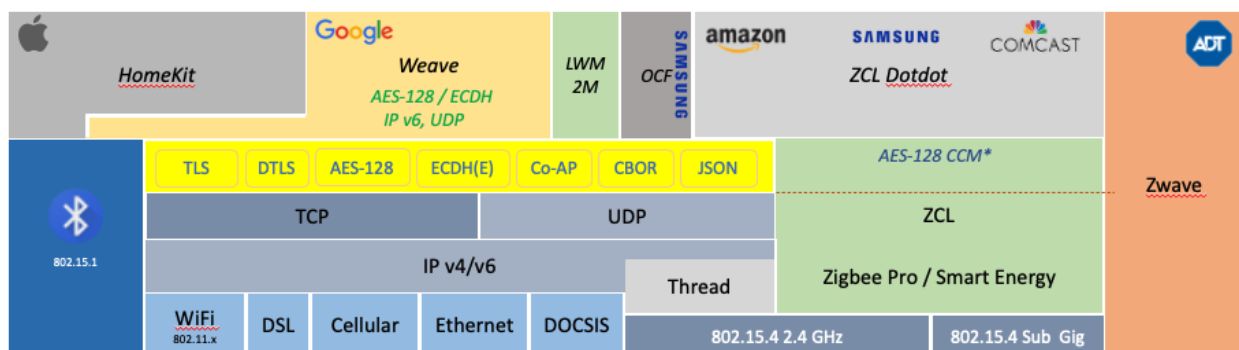


Figure 1: IoT Quilt

What industry needs is a cooperative effort that leads to a widely adopted standard with input from a wide swath of the IoT industry. This is exactly what Matter is. It grew from a need to create a new standard that

brings the best components from prevailing walled gardens in an open effort that guarantees participation from willing participants in drafting the new standard that will have wide acceptability from the launch.

1.3. Consumer Branding/Labeling/Proximal Interoperability

All the industry cooperation in the world will make little difference to the current smart home landscape if consumers do not literally buy in. The Connectivity Standards Alliance, CSA - the standards development organization (SDO) that maintains Matter - recognizes this and has developed marketing materials and a consumer-facing brand to help drive adoption. Starting this fall, consumers will be able to purchase products that contain the Matter logo on their packaging. This logo will signal to consumers that the device they purchase is guaranteed to work with their Matter applications. In the way that users currently associate the Wi-Fi logo with interoperable wireless connectivity, the Matter logo will be associated with guaranteed smart home interoperability.



Manufacturers can only display the Matter logo on their packaging if their devices have been certified. The certification process is run by CSA authorized test labs and is designed to guarantee conformance to the Matter specification. This process is analogous to devices being certified for Wi-Fi, Zigbee, Bluetooth, and many other protocols. With certification, a user can feel confident that the device they purchase will work seamlessly with their Matter application, access point, and other Matter devices. Certification is available to all CSA member companies.

1.4. Lack of Security/Uniform Standard for Security

IoT devices provide a fertile ground for hacking leading to either privacy violation, loss of service or dangerous situations. Two security researchers successfully commandeered a Jeep Cherokee using remote exploit (Greenberg, 2015). Incidents like these demonstrate that today's IoT needs to be empowered so that they can protect themselves in a hostile environment. In September 2016, the website of computer security consultant Brian Krebs was hit with 620 Gbps of traffic emanating from hijacked IoT devices, "many orders of magnitude more traffic than is typically needed to knock most sites offline" Koliass (2017, p.81). This attack exploited IoT devices such as cameras and DVRs that had default passwords to effectively turn thousands of these devices into a bot army. These devices were then used to launch a massive denial of service attack on DNS service provider Dyn, effectively causing an Internet blackout. This malware is dubbed "Mirai". One analysis of the attack showed that the Mirai malware had infected 49,657 unique devices, which included mostly IP based cameras, but also DVRs and routers (Herzberg, Bekerman, & Zeitman, 2016).

Such vulnerabilities are compounded by a lack of common IoT application standards that are created with security and privacy as one of the core principles. With disparate security implementations, IoT devices suffer from multiple possible attack vectors due to different security layers and their implementations.

1.5. Introduction to Matter

Matter is a next gen IoT application protocol designed from the ground up by IoT practitioners of the world. It is designed so that constrained IoT devices, controllers and Apps can interoperate, providing consumers with a rich experience that IoT devices offer. It also provides a standardized, compatible operating environment on which to build amazing experiences for developers. Similarly, retailers will benefit from simplified selling experience. Rather than invent underlying primitives, Matter uses proven technology and contributions from Amazon, Apple, Comcast, Google, Silicon Labs, Samsung, and others who have contributed resources in following key areas.

- Canonical Protocol Specifications
- Standardized SDK that implements the canonical Specification
- Test harness that checks SDK for compliance with the canonical Specification

Matter provides a unified out-of-box commissioning and pairing mechanism that ensures onboarding interoperability and compatibility between IoT devices made by a myriad of manufacturers. In addition to commissioning, included in the Matter stack are following features.

- Flexible device and service discovery
- A certificate-based device attestation process
- A secure end to end mutual authentication and encryption that provides security, privacy, and integrity of inter device communications.
- A unified and flexible interaction layer between data model and routing layer
- A unified data model with a wide range of device types defined

The following diagram depicts different layers of Matter protocol.

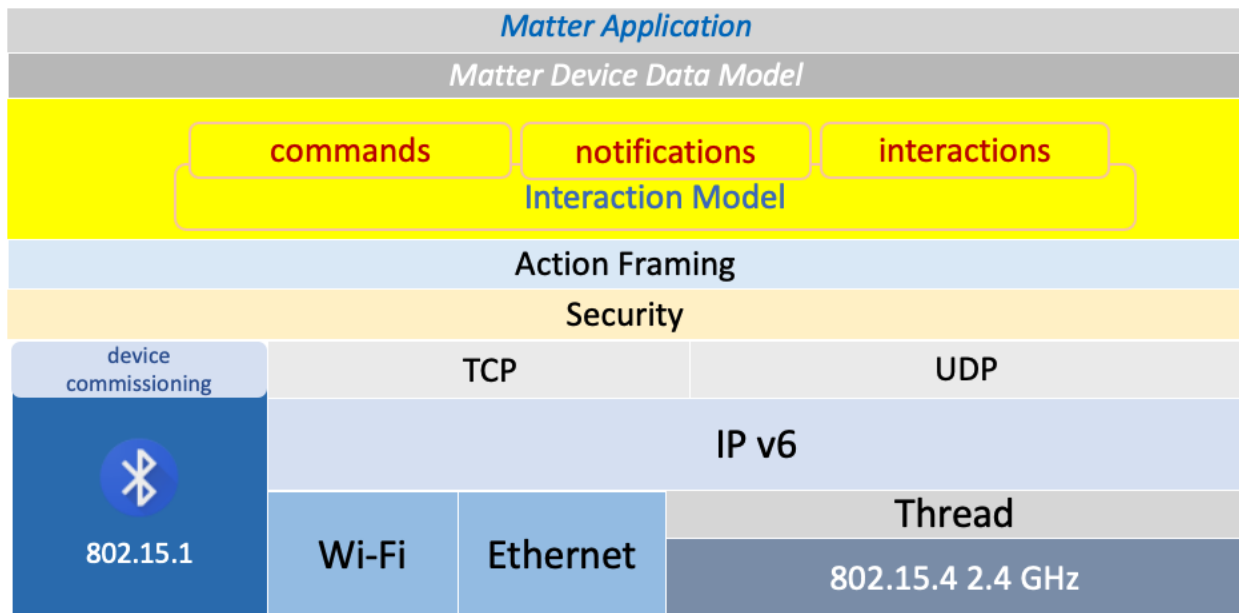


Figure 2: Matter Protocol Layers

1.6. Impact for Cable Operators.

Cable operators will be impacted by the emergence of Matter in several ways. Perhaps the most immediate impact is the introduction of a new wireless protocol, Thread, to the smart home. Over time, more and more low cost, low power, and low bandwidth devices will adopt this protocol and consumers will begin to expect support within their access point. Cable operators that fail to meet this expectation should expect increased call volume due to failed onboarding and lack of end device connectivity. Where Thread connectivity winds up being provided by some other device, such as a smart speaker, cable operators will lack any visibility into the state of the Thread network and have limited tools to support their customers. Matter will also increase the amount of broadcast and unicast traffic on a consumer's local area network. All this traffic will be encrypted leaving cable operators unable to effectively characterize it as malicious or safe. Participation in Matter will provide operators tools to better understand the devices on their customers networks and those devices' expected behavior. Matter also requires support for IPv6 and multicast DNS. Cable operators who do not currently support these technologies should expect increased call volume. Operators must support at least the basic requirements of Matter; Wi-Fi, Thread, IPv6, and mDNS. Failure to do so will risk their current managed access point offerings becoming obsolete, resulting in a lower take rate for this recurring source of revenue.

2. Matter Architecture

2.1. Transports and Network Stack (Wi-Fi, Thread, BLE, IPv6, DNS-SD)

At its core, Matter is a specification that defines a framework for interoperable, local network communication. The Matter specification allows for application layer messages to be transported over any IPv6 bearing network. Version 1.0 explicitly lists three compatible transports: Wi-Fi, Ethernet, and Thread. Wi-Fi and Ethernet are aimed at general purpose or high bandwidth applications and provide access to a local area network, LAN. Thread is aimed at low power and low bandwidth applications and

provides a mesh network topology known as a personal area network, PAN. A device known as a Thread Border Router allows for messages to be transported between the LAN and PAN. Cable companies that currently offer Wi-Fi routers should consider adding Thread support and enabling Border Router functionality. This will increase the range of devices they can currently support, provide access to useful information about the Thread PAN and network topology, and ensure their current connectivity solutions maintain pace with user expectations.

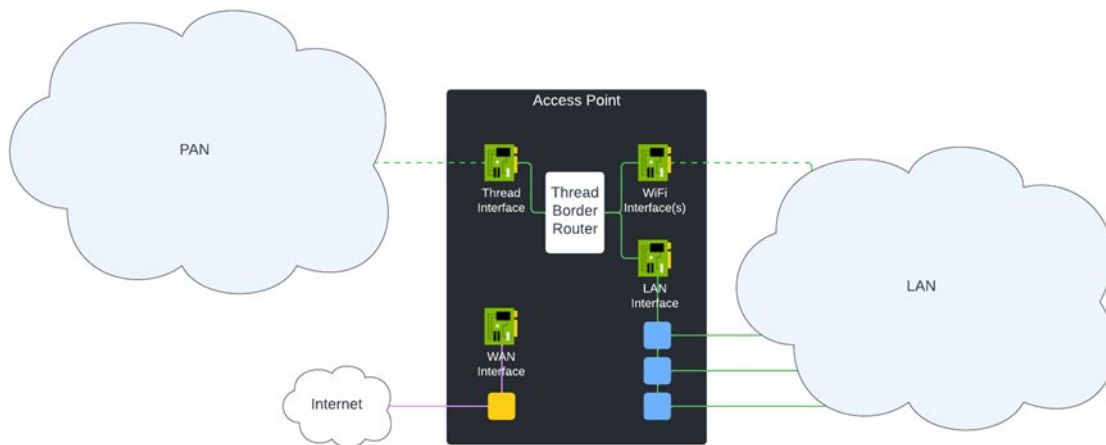


Figure 3: Router Architecture

In addition to IPv6 bearing transports, Matter also utilizes a few additional technologies to enable device discovery and network onboarding. One such technology is Bluetooth Low Energy, BLE. Matter utilizes BLE to enable discovery of commissionable devices and to facilitate a series of messages designed to verify the authenticity of a device, provision the device with Matter credentials, and to transmit credentials needed to connect to the user's Wi-Fi or Thread network. Devices that support BLE, this includes all Matter Thread devices, will initially advertise themselves over BLE when first powered on. Once they have been configured through the Matter commissioning process they will begin to operate over Wi-Fi or Thread.

Matter also makes extensive use of DNS service discovery, DNS-SD. This technology is utilized for device discovery and IP resolution. In the case of device discovery, devices that have already been onboarded to the user's LAN or PAN can publish themselves as commissionable for Matter commissioners to discover. In this situation the commissioning process would occur over IP on the respective transport and not utilize BLE. In the case of IP resolution, DNS-SD is utilized to resolve the IP address of commissioned devices, known in Matter as nodes, on the LAN or PAN. Utilizing DNS-SD to resolve a node's IP address allows for greater flexibility with IP assignment and does not require the node's IP to remain fixed over time. Nodes publish their presence with the `_matter._tcp` service and include an identifier consisting of their Fabric ID and Node ID. This allows the IP address of a Matter device to be resolved to its Fabric scoped Matter node identifier.

2.2. Fabrics

Matter groups collections of devices, Matter nodes, into what are known as Fabrics. Fabrics exist to provide: a common root of trust for this collection, enable encrypted communications between nodes, and

provide a set of scoped access controls. Fabrics exist only at the application layer and are completely decoupled from a user's LAN or PAN. In practice, most Matter users will operate multiple Fabrics, one per application they utilize to control their smart home.

Fabrics are created, or provisioned, by a Trusted Root Certificate Authority, TRCA. A Fabric ultimately consists of a unique 64-bit identifier and the public key certificate of the TRCA. A Matter device may belong to one or more Fabrics and would have a unique Node ID and Node Operational Certificate, NOC, for each Fabric it belongs to. This structure is what allows a Matter device to be controlled from multiple applications, provides users flexibility in granting applications access to Matter devices, and prevents ecosystem lock-in. Cable operators are well positioned to be TRCAs and provide supporting services to businesses that wish to participate in Matter but don't want to manage the required PKI.




Figure 4: Matter Fabrics

Nodes within a Fabric are permitted to communicate with each other based on access controls scoped to that Fabric. Every Matter node must maintain a set of access controls for each Fabric it belongs to. These access controls dictate permissions other nodes on the Fabric have to interact with functionality, exposed through Clusters, that a given node hosts. Communications between nodes are encrypted using the NOC they were provisioned during commissioning into a Fabric. This functionality allows for a secure channel to be established between nodes and is analogous to how communications are secured between modern Internet websites and browsers using HTTPS and PKI.

2.3. Data Model

In Matter, physical devices such as light bulbs, door locks, TVs, and more expose their functionality through the Zigbee Cluster Library, ZCL, data model. This data model enables interoperability between clients that issue commands and servers that accept commands. The highest order element in the Matter data model is a Device Type. Device Types can be things like dimmable light bulbs, thermostats, video players, and Wi-Fi access points. There are two high level Device Type categories in Matter, Utility Device Types and Application Device Types.

These Device Types are exposed through Endpoints that in many ways are akin to IP ports. A single physical device may expose multiple endpoints and therefore functionality of multiple Device Types. However, in practice most physical Matter devices will host two endpoints. One, endpoint 0 in all cases, that is used for administration and one that is used for actual application functionality. More information about Matter device types can be found in Connected Home over IP Device Library.

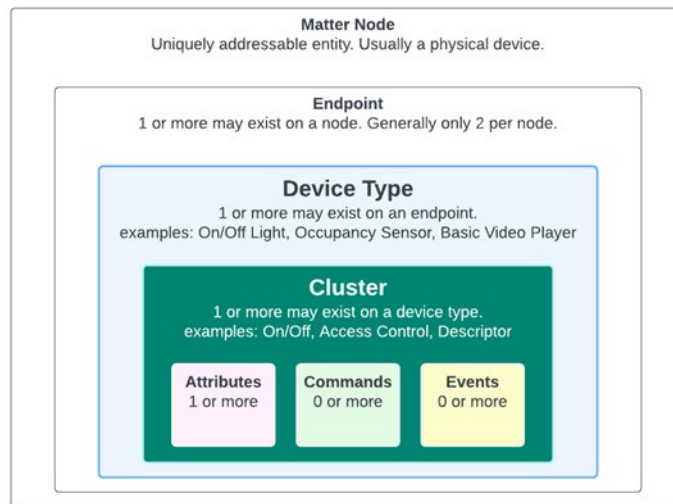


Figure 5: Device Types and Clusters

The functionality that a Matter Device Type exposes is done through implementing Clusters. Clusters are the foundational element of the ZCL data model and describe types of functionality such as the ability to turn something on or off, or adjust the level of something, like in a fan. Clusters are constructed from three lower-level elements: attributes, commands, and events. Attributes store state, such as the on or off state of a light bulb. Commands can be invoked to manipulate state, such as a toggle command to change the on or off state of a light bulb. Events can notify interested parties of updates to an attribute. An example of a complex interaction displaying all three Cluster conventions could involve an update to a thermostat that multiple occupants control from separate applications. In application 1, an occupant can invoke the 'SetpointRaiseLower' command hosted by the 'Thermostat Cluster' to adjust the 'OccupiedCoolingSetpoint' attribute. In a second application, another occupant that had subscribed to events on the 'Thermostat Cluster' could be notified through a Matter Cluster event of the update to the 'OccupiedCoolingSetpoint' attribute.

Clusters are defined as being either an application cluster or a utility cluster. Application clusters describe actual device functionality such as the ability to unlock a door lock. These clusters are generally hosted on endpoint 1 but may exist on multiple endpoints or have unconventional starting indices. Utility clusters are used to configure a Matter device or convey information about it. The functionality and information exposed on these clusters aids in commissioning, setting access controls, configuring bindings, reading logs, and more. In general, these clusters will be exposed through the Root Node device type that is always hosted on endpoint 0 for every Matter device. However, some utility Clusters, such as the Descriptor Cluster, exist on all endpoints that a device hosts. More information about application clusters can be found in Connected Home over IP Application Clusters. More information about utility clusters can be found in sections 9.5 – 9.14 and 11.1 – 11.19 of Connected Home over IP Specification.

2.4. Interaction Model

Matter defines four methods that a controller application can use to interact with the ZCL data model - read, write, invoke, and subscribe. These interactions ultimately allow an application to retrieve the current state of a device or to affect changes to its current state. Reads can be performed on cluster attributes to get their current state. An example would be reading the 'OnOff' attribute of a light bulb to determine if it was currently on or off. Writes can be performed on some cluster attributes to update their state. An example of this would be writing an entry to the 'NodeLabel' attribute to specify a user defined name for the Matter device. Invokes can be performed on cluster commands to affect the state of a cluster attribute. An example would be to invoke the 'Off' command to change the 'OnOff' attribute of a light bulb to off. Subscribes can be performed on many cluster attributes and events to receive notifications of changes to their state. An example would be to subscribe for updates to the 'CurrentLevel' attribute of a light bulb. With the subscription enabled, any time a user updates the level (brightness) of the bulb, all applications with active subscriptions would receive a message with the updated state.

2.5. Bridging - Current Networks and Legacy Devices

Bridges allow non-Matter IoT devices to participate in a Matter Fabric. This will allow consumers to continue to use some of their current (legacy) devices and will allow for a gradual transition to newer Matter-enabled devices. Bridges act as a translator between non-Matter protocols and Matter Fabrics. Similar to other Matter devices, a bridge can participate in many Matter Fabrics and must have at least one node on each Fabric. From that node the bridge will expose any number of endpoints through a specific device called an Aggregator. This device type has a cluster attribute that contains a list of all the bridged devices called a PartsList. Bridged endpoints are not full Matter nodes, they do not have their own Matter operational credential and all access control decisions for bridged devices will use the operational credential associated with the bridge. Access control entries can be specified using a set of targets that contain bridged device entries and bridged device types. For example, an ACL on a Matter lock could allow bridged device types of Window Sensor to read its state using the Node credential of the bridge.

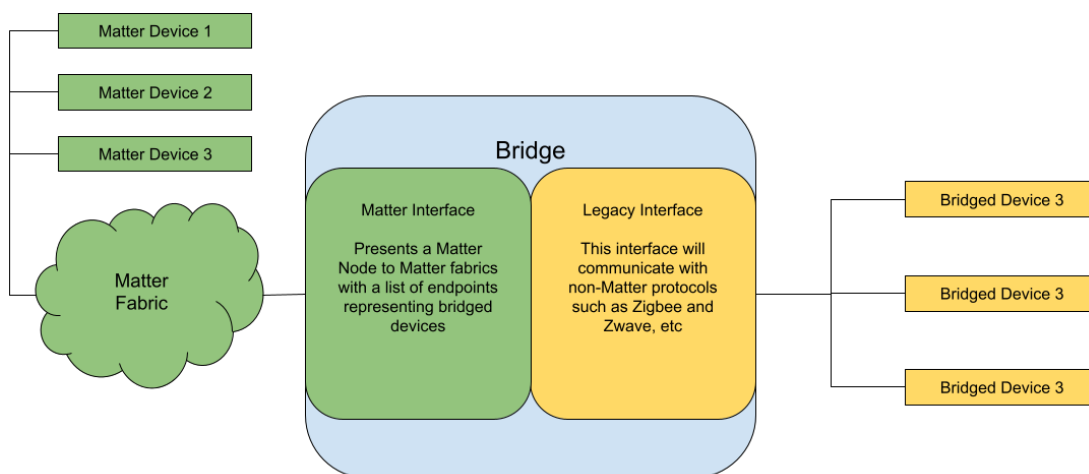


Figure 6: Bridging Architecture

3. Commissioning

The commissioning refers to a process of bootstrapping a network (aka a Fabric in Matter parlance) consisting of devices (including apps) that communicate with each other using Matter protocol. The commissioning process is designed to address following critical security principles.

- a. User Intent - user is explicit in the intention to commission the device
- b. Proof of Possession - the user is known to physically have and control the device
- c. Initial Secure Channel - all communications to commission and configure are secure and encrypted
- d. Device Authenticity - the user can be confident that the device is what it says it is




Figure 7: Matter Commissioning

In Figure X above the commissioning process ① is started by first scanning the QR code on the device or its packaging, this is usually done with an app on a smartphone referred to in Matter as a commissioner. The QR code contains among other things the product ID, the vendor ID and a passcode that is used to establish the initial secure connection with the device. This passcode can be manually entered by the user and even spoken into a smart speaker.

Next the commissioner will search for and connect to the device ②. This connection can be done over BLE, a temporary Wi-Fi connection to the device (SoftAP), or directly via IP. In this step the commissioner will check the device attestation certificate (DAC) and the certification declaration (CD) to establish that the device is authentic and has been tested and certified. The commissioner and the device will then establish a strong encrypted communication path with which to configure the device.

In ③ the commissioner configures the device with credentials for the Wi-Fi network as well as operational credentials that the device will use to connect to devices on the same fabric. Additionally, this step will often be used to set up access control lists so that the device will allow connections from other devices and applications.

Finally, in ④ the device will connect to the access point and begin normal operations.

The commissioning of a device can be done several times to add the device to new fabrics and to add administrators of the device.

4. Distributed Compliance Ledger

The Distributed Compliance Ledger (DCL) is a cryptographically secure, distributed network that allows device manufacturers (vendors), official test houses and Alliance certification centers to publish public information about a given device. Based on blockchain technology, it allows participants to update relevant device information that is cryptographically signed.

Connectivity Standards Alliance's (Alliance's) Distributed Compliance Ledger (DCL) is an industry-wide initiative to provide a cryptographically secure, distributed ledger of certified IoT devices and their roots of trust, without one company or an entity in charge of the ledger. By using an underlying permissioned blockchain framework, the DCL benefits from the following properties:

- Multi-node network that is run by the Alliance member companies
- Individually signed transactions using pre-approved keys
- Non-repudiation
- Distribution of data in different geographical locations
- Consensus protocol to ensure majority approval
- Public reads with available cryptographic proofs attached
- Transparency and auditability

4.1. Architecture

Conceptually, the DCL is a network of authorized "Validator Nodes" that comprise the consensus pool of the underlying blockchain framework. It works on the basis of a "Permissioned Ledger" where all write transactions are individually validated by the distributed validator nodes using the enrolled public keys. The CSA provides a public facing "Node" for full public access (Read Only). CSA members run their own nodes in their protected networks.

DCL 1.0 Topologies

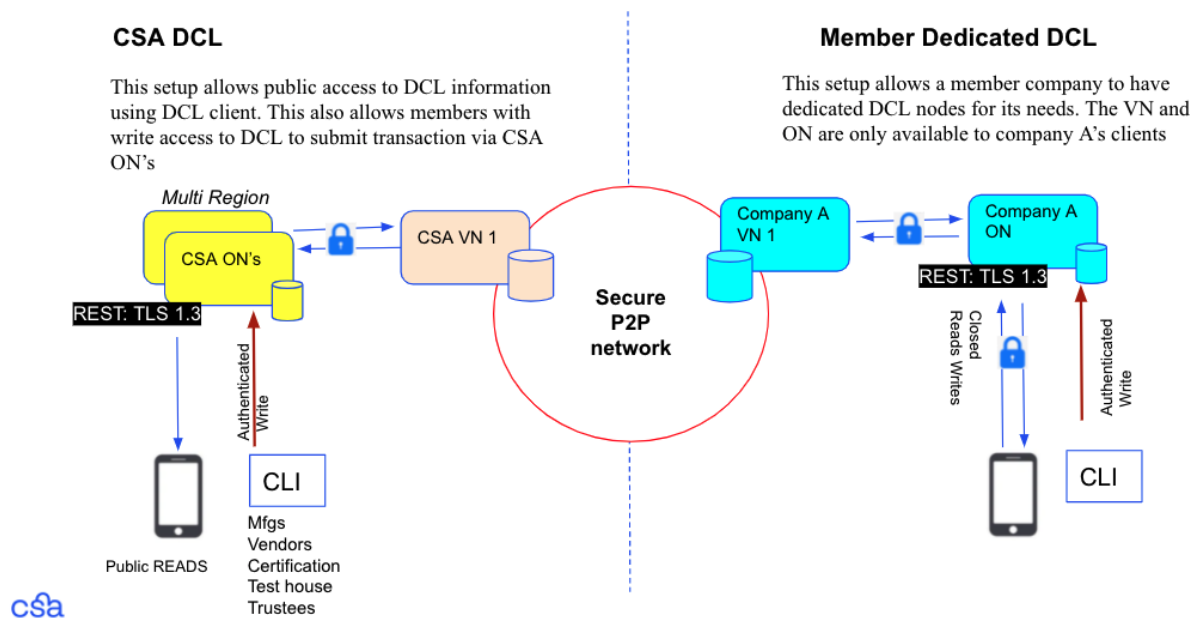


Figure 8: DCL Architecture

4.2. DCL Schema

DCL contains the following standard information, securely uploaded by individual members and validated using their public keys to provide nonrepudiation.

- **Vendor Info**, for example, company name, legal and preferred name, URL, etc.
- **Models and Model Software Versions**: models are identified by a unique combination of 'VID' (vendor ID) and 'PID' (product ID).
- **Model Certification**: DCL can be used as a source of information if certain Model Software versions are certified by a Certification Center. Certifications may be revoked.
- **PAA Certificates**: DCL can distribute authorized X509 Certificates (PAA and non-root).
- **Auxiliary Information**: DCL has an array of required auxiliary transactions:
 - **Validator Node** transactions define the current set of VNs participating in consensus.
 - **Account** transactions contain a role and a public key for every user wanting to write to the ledger and must be signed by the private key associated with the account.

- **Upgrade** transactions contain information about scheduled and completed updates of the DCL software.

4.3. DCL Nodes

The DCL network consists of a peer-to-peer network of Validator Nodes that communicate with each other using a secure and authenticated protocol.

- **Validator Node (VN)** is a full node that participates in the consensus protocol utilizing an authorized cryptographic key. A VN is responsible for validating new records. In order to run a VN, an organization must have an approved account with a Node Admin role - Access to VNs should be restricted by means of Sentry nodes and private networks. A better way to accept requests from the user is via Observer nodes.

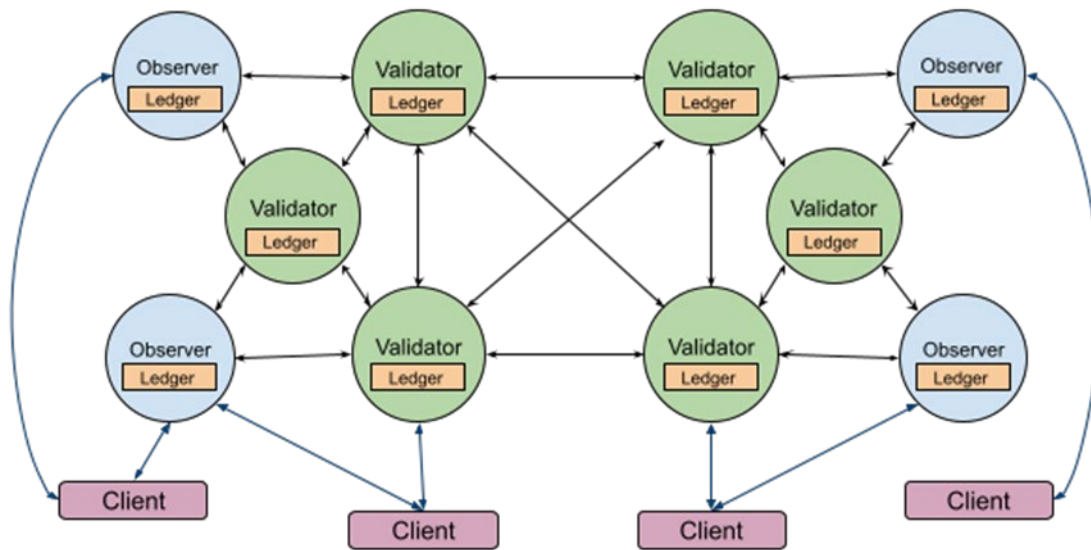


Figure 9: DCL Network

- **Observer Node (ON)** is a full node that does not participate in the consensus. ONs do not require an account with a Node Admin role to be created and member companies can govern their access to their ONs.

4.3.1.PKI Overview (PAA/PAI)

Matter specifies a three tier certificate chain for device attestation certificates to ensure device authenticity during commissioning. The top-level root is called “Product Attestation Authority” or PAA.

Each PAA can have one or more Intermediate known as ``Product Attestation Intermediate” or PAI. The PAI is responsible for issuing a unique device attestation certificate or DAC to each Matter device. The authorized pool of PAAs is entered in DCL so that they can be downloaded using the public DCL read interface to all Matter commissioners.

The Matter specification allows CSA members to operate their device PKI chain as long as the PAA and PAI adhere to the Matter Certificate Policy. Therefore, due care and best practice should be followed in setting up both a root and intermediate certificate authority and certain requirements must be met according to the Matter certificate policy for establishing and running a PAA/PAI. Several public PKI providers who are also CSA members can also provide PKI as a service by hosting PAA and PAI chains on behalf of Matter device manufacturers.

4.4. Running a DCL node

DCL consists of a pool of validator nodes that provide consensus based cryptographic integrity checks on the DCL entries. This provides blockchain level security and integrity for all the device data that is entered by authorized companies. Since DCL is a permissioned ledger, CSA and trustees of the DCL authorize each Validator node individually to become part of the DCL consensus pool.

CSA members can also run Observer Nodes which do not participate in the DCL consensus pool and therefore do not need to be authorized individually. CSA members can run Observer Nodes so that they have a full database of certified Matter devices in their network to enable individual company’s use cases.

5. Security

Matter takes a comprehensive approach to security that involves securely onboarding the device to the Matter Fabric, encrypted mutually authenticated connections between devices, fine-grained access control between devices and users, and a software update mechanism for issuing patches and new firmware to devices.

5.1. Secure Communications

The set of protocols used in Matter assures that all unicast communications are secured (encrypted and integrity checks), authenticated and provide builtin replay protection.

There are two methods of establishing a secure session in Matter and their use depends on if the device is being commissioned or is in normal operation. Both methods set up a shared symmetric session key for fast data encryption.

Passcode-Authenticated Session Establishment (PASE) is used when a device is being commissioned. PASE uses a password or passcode communicated through an out-of-band channel such as Bluetooth Low Energy (BLE), scanning a QR code, or Near Field Communication (NFC). Once both the device and the commissioner have the passcode they use it to calculate a common symmetric key on both sides. This key is then used to establish AES-CCM session keys to provide confidentiality and integrity of communications between devices. PASE is used to set up connections that establish operational credentials between devices.

Certificate-Authenticated Session Establishment (CASE) is used once a device is configured. It uses operational credentials in the form of certificates that are established at the time the device is

commissioned. These operational credentials are used to authenticate both ends of the communication and set up shared session keys. These shared symmetric session keys are used by the devices to encrypt traffic using the AES-CCM cipher suite.

5.2. Replay Prevention

Matter nodes use sessions as a way to quickly setup previously authenticated channels between two nodes. Sessions provide a mechanism to pause and resume a connection between devices as long as the session keys remain valid. Sessions also provide the Matter protocol a way to keep track of messages as they traverse two nodes. A randomly initialized message counter is established with each session that serves a two-fold purpose. First, it acts as an encryption nonce to ensure that every message is encrypted in a unique manner. Second, it acts as a replay and duplicate message detection mechanism.

Message duplication can be caused by network latency and errors where the sender did not receive an acknowledgement. Unhandled duplicate messages present problems if control messages are sent multiple times and can change the state of the device in unpredictable ways. Threat actors, while they are incapable of decrypting the message, can intercept encrypted messages and re-send them or “re-play” them. This can cause an unintended state in a device. A common example of this is a legitimate unlock command that is captured by an attacker and resent to a door lock later allowing the attacker to gain physical access. Matter prevents these issues by maintaining a history window of message counts from a sender to determine if the message is a duplicate. Duplicate messages are dropped by devices before they reach the application layer.

5.3. Secure Group Communications

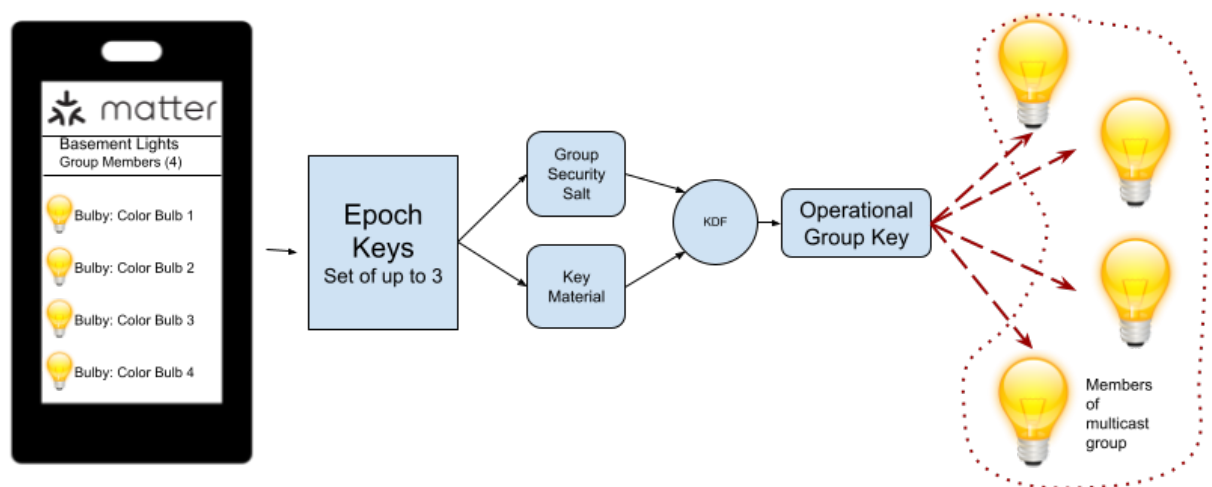


Figure 10: Secure Group Communications

Group communications in Matter take advantage of multicast networking. Multicast allows a single sender to send a message to a specific destination group and it is received by several devices simultaneously, relying on the network to replicate the message. An example use-case of this in the smart home is to send an “ON/OFF” command to a group of lights. Using multicast, a single command can be sent to a large group of devices instead of having to send individual unicast messages to each

device. This is both more efficient and faster than unicast as it requires fewer messages from the sender and can be done in parallel.

Because secure group communications is a one-to-many relationship i.e one sender and many recipients, securing this communication must be done in a way that provides three things:

1. All devices can prove that they are members of the associated group
2. All messages are encrypted and confidential.
3. Messages can only be decrypted by other members of the group

To accomplish the above goals Matter distributes through an Administrator Node a set of 1-3 keys to each device in the group called epoch keys. Devices use these epoch keys to derive an operational key that is used to encrypt and decrypt messages for the group.

5.4. Access Control

Access control provides fine grained access to devices. For example, your front door lock can be set up to allow one family member's phone (i.e., a parent) to manage who can change access settings on the lock (what other devices can view and perform actions on the lock), and another family member's phone (i.e., child) to operate (lock/unlock) the door.

Access control lists consist of the following attributes:

- Fabric Index: ID scoped to the associated Fabric
- Privilege Level: View, Proxy View, Operate, Manage, and Administer
- Targets: List of clusters (data model elements e.g DoorLock, temperature sensor)
- Subjects: List of sources of an action to which the ACL applies, often this will be the Node ID
- Authentication Mode: Type of secure channel (PASE, CASE)

An example of an ACL on a Matter device is shown below. The first entry gives "Administrator" privilege to a Matter node with ID "0xAABD65DF76230b54" over CASE authentication.

The second entry allows all devices on Fabric 1 to have "View" privilege on all target devices endpoint 1.

The last entry provides "Manage" privilege to subject 0x0000000000000001 on End Point 1 and Endpoint 3 of cluster ID 0x0000_0202.

```
ACL: [
  {FABRIC: 1, PRIVILEGE: ADMIN, AUTHMODE: CASE, SUBJECTS: [0xAABD65DF76230b54],
  TARGETS: []},
  {FABRIC: 1, PRIVILEGE: VIEW, AUTHMODE: CASE, SUBJECTS: [], TARGETS: [ENDPOINT: 1]},
  {FABRIC: 1, PRIVILEGE: MANAGE, AUTHMODE: GROUP, SUBJECTS: [0x0000000000000001],
  TARGETS: [{ENDPOINT: 1}, {ENDPOINT: 3, CLUSTER: 0x0000_0202}]}
]
```

5.5. Software Update

All devices will need to be updated at some point in their life cycle either to provide new functionality or to patch a vulnerability. Matter mandates that all Matter certified devices support over-the-air (OTA) software update. Software images must be signed by the vendor using a private key dedicated to signing images. Firmware images are transferred over a Matter specific Bulk Data Exchange (BDX) protocol that treats images as a collection of bytes with metadata. BDX is modeled after TFTP and can use either TCP or UDP as an underlying transport.

The software update process in Matter involves two nodes:

- OTA Requestor is any node that requires updating of software.
- OTA Provider is a Node that fulfills software update requests. The OTA Provider downloads the image from the vendor and stores a copy or alternatively can act as a proxy through which the OTA Requestor can download the firmware image.

Software updates must be checked for authenticity and integrity by the device prior to being installed. Additionally, devices will not install version numbers that are lower than their current running firmware version to prevent downgrade attacks.

6. Administration

6.1. Operational PKI

Devices are permitted access to a Matter Fabric through the provisioning of a Node Operational Certificate, NOC. The NOC uniquely identifies a particular node within the scope of a particular Fabric. This certificate is also used to establish secure communication channels with other nodes belonging to the same Fabric. NOCs may be directly issued by a Trusted Root Certificate Authority or an Intermediate Certificate Authority and are provisioned during Matter commissioning. A Matter device will have exactly one NOC for every Fabric that it belongs to.

6.2. Configuring Access Control

Matter provides a flexible fine grained access control list for participating devices. The rule-based ACL defines no implicit access by default. Therefore, it works on the principle of implicit “deny” unless an explicit rule defined on a device grants access to a requesting device, further fine grained to a given endpoint (service) on the target device.

An access list is created on a device during its commissioning so that when it is put on the network, it has proper security in place. During the commissioning phase, the access control module on a device grants implicit “Administrative” access to a Commissioner over PASE session which entails creating a secure channel using SPAKE2+ protocol between the device and commissioner using a passcode for that device. Further updates to Access Control lists can be made over the network on a CASE session as long as the source device has “Administrator privilege” defined in the target device ACL.

6.3. Fabric Management

There are two roles that are necessary to fully manage a Fabric. The commissioner role is used to provision new nodes into the Fabric and provide them with operational credentials that chain up to an operational root of trust. The administrator role is used to change access control lists within their respective Fabrics. Additionally, administrators can remove a Fabric that is not directly under their control. This is an important feature of Matter as it allows a device owner to remove a Fabric from the device that they may no longer control or a Fabric they no longer wish to control the device.

6.4. Group Management

It is expected that groups of devices will change. For example, you may wish to add or remove a light from a group. Adding a device is as simple as commissioning a new device and adding it to the group along with the required group operational key so that it can decrypt signals sent to the multicast address of the group. Removing a device from the group requires a re-key of the remaining devices in the group with new operational keys.

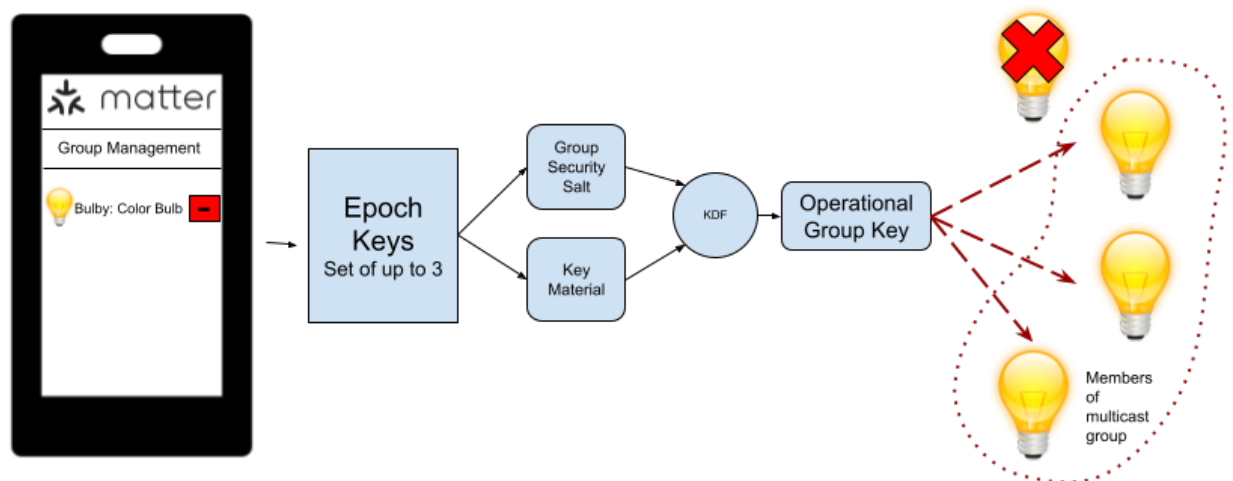


Figure 11: Group Management

6.5. Software Update

Software update announcements and notifications are automatic and happen in the background. To assure that the update does not disrupt the user, downloading and applying updates should be done after obtaining user consent. User consent can be stored for future use and is an optional process in Matter.

6.6. Multi-Admin

Matter has standardized simultaneous command and control of a device by multiple ecosystems. This key feature allows consumers to “commission” i.e. securely add control of their Matter devices by multiple ecosystems ensuring they are not locked into using a single vendor or ecosystem. For example, a door lock can be opened with multiple apps through multiple ecosystems if provided the initial owner intends to set it up in such a way.

The sequence diagram below depicts the steps initial Ecosystem (Eco A) takes to pair an already commissioned “bulb” to add it to the second ecosystem (Eco B). The process involves the user (Alice) instructing the device (Bulb) to go into commissioning mode using a new onboarding passcode and go into discovery mode. The commissioning window is open for a specified time in seconds. Alice conveys the new PIN to Bob who is connecting to the bulb using a different ecosystem. Bob discovers the bulb using “dns-sd” and completes commissioning using the new passcode.

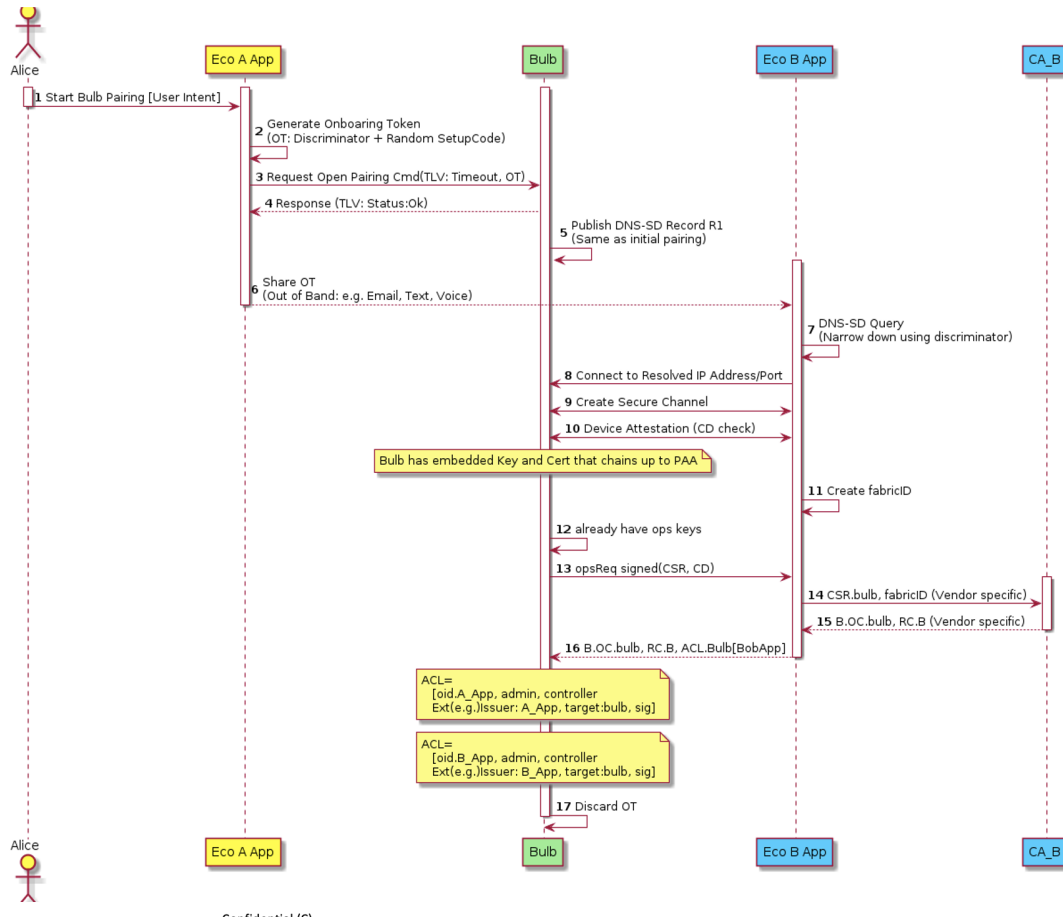


Figure 12: Multi-Admin Sequence

7. Matter and Operators

Matter represents a paradigm shift in traditional smart home management. This new specification simplifies the way applications and devices communicate by utilizing IP, the language of the Internet. It decouples the networking stack from the application stack, allowing for flexibility in use cases that require differing bandwidth and energy constraints. Matter is also open source and provides full access to specifications and reference SDKs. Given the unique relationship that Cable operators have as Internet and local connectivity providers to the smart home, they have a great opportunity to both add value to and derive value from Matter.

As wireless connectivity providers, Cable operators have a goal of ensuring their customer's devices seamlessly and consistently connect to their access points. Fundamentally, all that is needed to support the basic connectivity requirements of Matter is to provide Wi-Fi Access Points, Thread Border Routers, and IPv6 and multicast DNS support. Most operators already support many of these requirements and simply need to add Thread capable radios to their access points. Ensuring all Matter supported wireless transports, Wi-Fi and Thread, are present in customer's homes will reduce the potential for failed device connections and user confusion. This in turn will drive down customer service calls and improve user satisfaction. Cable operators should also consider certifying their access points as Matter certified Wi-Fi access points and Thread Border Routers. These additional device certifications are currently being defined and will be available in a future release of Matter.

As Internet Service Providers, Cable operators have a responsibility to help safeguard their customer's home networks and the broader Internet as a whole. The tools that allow this oversight often rely on the ability to uniquely identify network devices and to understand the applications they expose or interact with. In the evolving world of Internet security and privacy practices that encourage use of randomized MAC addresses and encrypted communication channels, the ability to acquire this information is becoming increasingly difficult. With Matter, Cable operators can directly query devices to uniquely identify them and understand what type of device they are. Additional helpful information such as a device's software version, hardware version, manufacturer, product ID, and serial number are also made available. Devices can be further vetted against the Matter Distributed Compliance Ledger to verify their authenticity and certification status. All of this represents a suite of tools that can be used to uniquely identify devices on a customer's network, understand the expected behavior of those devices, and provide a rich set of data to direct customers to appropriate support resources.

As customer premise equipment providers, Cable operators are well positioned to add value to Matter by proxying Matter communications from the Internet to a customer's local network. Matter is designed as a local network communication stack and does not define how the smart home may be interacted with when a user is not at home. Control of Matter devices while away from a user's local network will require routing messages through a locally connected device. Cable companies offer a range of devices including Wi-Fi routers, cable modems, and set-top-boxes that can act as proxies to relay these messages. Having a locally connected device that is always listening and addressable from the Internet provides Cable operators the ability to address this existing gap. Further value can be added by simplifying the device onboarding process. Cable's scale, access to the home, and ability to forge business to business relationships place it in a unique position to improve this area. Through arrangements with retailers, device onboarding information could be transferred to a customer's router at time of purchase so that the onboarding process automatically happens in the background when the device is first powered on. These are just a few of many potential value propositions the Cable industry can provide to Matter.

Matter is poised to transform the way that users, device manufacturers, and ecosystem providers approach the smart home market. Users will benefit from a recognizable brand that simplifies the purchasing process and ensures interoperability. Device manufacturers will be able to offer a reduced set of SKUs and take advantage of open-source SDKs. Ecosystem providers will be able to seamlessly connect to a range of devices with improved onboarding flows and enhanced security. With all these benefits in place, Cable operators must embrace Matter in order to meet their customer's evolving expectations and usher in the next generation of the smart home.

Abbreviations

AP	access point
BDX	Bulk Data Exchange
BLE	Bluetooth Low Energy
CASE	Certificate Authenticated Session Establishment
CPE	Customer Premise Equipment
CSA	Connectivity Standards Alliance
DAC	Device Attestation Certificate
DCL	Distributed Compliance Ledger
DNS	Domain Name System
DNS-SD	DNS - Service Discovery
IP	Internet Protocol
IPv6	Internet Protocol version 6
KDF	Key Derivation Function
LAN	local area network
NOC	Node Operational Certificate
PAN	personal area network
PASE	Passcode Authenticated Session Establishment
PID	Product Identifier
PKI	Public Key Infrastructure
ON	Observer Node
OTA	Over-the-Air
SDK	Software Development Kit
SKU	Stock Keeping Unit
TRCA	Trusted Root Certificate Authority
VID	Vendor Identifier
VN	Validation Node
ZCL	Zigbee Cluster Library

Bibliography & References

Greenberg, A. (2015, July 21). Hackers Remotely Kill a Jeep on the Highway—With Me in It. Wired. Retrieved from <https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/>

Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and Other Botnets. *Computer*, 50(7), 80–84. <https://doi.org/10.1109/MC.2017.201>

Herzberg, B., Bekerman, D., & Zeitman, I. (2016, October 26). Breaking Down Mirai: An IoT DDOS Botnet Analysis. In *Incapsula.com*. Retrieved from <https://www.incapsula.com/blog/malware-analysis-mirai-ddos-botnet.html>

Measuring IP Video Playback Quality of Experience

A Technical Paper prepared for SCTE by

Srinath V Ramaswamy
Principal Solution Architect
Comcast
1800 Arch Street, Philadelphia, PA 19103
(215) 301 5047
srinath_ramaswamy@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. IP Video Architecture	3
2.1. Playback Artifacts.....	3
2.2. Measuring IP Video Playback QoE.....	4
3. Video Viewing Quality(VVQ).....	4
3.1. VVQ Model.....	4
3.2. VVQ Formula.....	5
3.3. VVQ Coefficients Calibration.....	10
4. VVQ Implementation	11
4.1. VVQ in Analytics Engine Reference Architecture	11
4.2. Sample VVQ Scores	13
5. IP Video Playback QoE Management – Self healing.....	14
6. Conclusion.....	15
Abbreviations	15
Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1 – IP Video Architecture	3
Figure 2 – Scoring in Analytics Engine or IP Video Player for Video QoE Management	11
Figure 3 – VVQ Implementation in Analytics Engine	11
Figure 4 – VVQ Average Scores Linear.....	12
Figure 5 – VVQ Score Components Average values	13

List of Tables

Title	Page Number
Table 1 – Sample VVQ Scores	13

1. Introduction

To measure IP Video Playback Quality of Experience (QoE) there are several distinct metrics available that need to be reviewed individually to determine the overall playback quality. These may include rebuffering, startup latency, video bitrate, errors, and downshifts. This paper presents a novel approach to objectively predict the IP video playback quality as perceived by the end-user by calculating a single, final overall metric referred to as Video Viewing Quality (VVQ). VVQ factors in several distinct metrics such as rebuffering, playback errors, video quality, downshifts, startup latency, and recency. Inherent in some of these metrics are the impact from audio issues such as dropouts, missing audio. VVQ scores are currently being utilized to determine IP video playback QoE on various IP video devices and to troubleshoot playback issues. We will also discuss how this single metric, or score can further be leveraged to enhance IP Video playback experience by performing bandwidth and network optimizations. Measuring IP video playback quality is an important aspect of the transition to the delivery of video over IP from legacy QAM.

2. IP Video Architecture

At a high level, the IP video architecture is as shown in the Figure 1 below. In this figure, the compressed or uncompressed video stream is transcoded into several streams (variants), each at different fixed bitrate, to accommodate the varying network bandwidth available between IP video consumer devices and the content delivery network, and the various consumer devices. The transcoded streams might be encoded using one of the MPEG compression codecs (MPEG-2[4], MPEG-4/AVC [2], HEVC [3]). This is then packaged into one of the many streaming formats, for example MPEG-DASH [5] or HLS [6], where they are typically split into time-aligned segments of a few seconds duration. These are then placed on the Origin Server and IP CDN for delivery to IP video clients. This would apply to both linear and on-demand IP video. Consumer devices will typically try to download the highest bitrate video that they can support, with the intent of providing higher video quality to the viewer.




Figure 1 – IP Video Architecture

2.1. Playback Artifacts

During IP video playback several customer impacting issues could arise; some of them are listed below.

- Stalling of video playback due to rebuffering or errors in the stream
- Playback of low-quality video for extended durations

- Bad quality of encoded video/audio, causing video issues like macroblocking, blurry images, noisy, out of lip sync, flashing, video buffering errors (underflow, overflow), or missing audio
- Large startup latency
- Media playback failures (playback start fails or leads to ending playback)
- Playback control failures (for example, while pausing or exercising trick modes)
- Live latency
- Multiple and frequent occurrences of playback issues during a session

2.2. Measuring IP Video Playback QoE

Measuring IP Video playback QoE is important for numerous reasons including:

- To allow for the seamless transition to delivery of video over IP from legacy QAM; customers should not experience a degradation in video playback experience when they are switched to IP video
- Early detection of IP video infrastructure issues and enable self-healing
- Validation of new streaming technologies (for example content adaptive streaming, new players, codec)
- Facilitating IP video system component upgrades
- For the retention of video customers in a highly competitive marketplace
- To provide an entry point for identifying streaming video issues and assist in their resolution

3. Video Viewing Quality(VVQ)

As described earlier, there is a need to measure IP video playback QoE. There are several distinct metrics that can be reviewed individually to determine the overall IP video playback QoE including average video bitrate, startup latency, rebuffering duration, rebuffering count, top errors/failures, encoded video quality, and others. Viewing these distinct metrics one at a time to detect bad user experience is not feasible, especially with large-scale deployments. One option is to look at only a subset of them, but because customer experience can be impacted by any one of them, all need to be taken into consideration.

What is needed is a holistic metric for measuring IP Video Playback Quality of Experience (QoE) and the VVQ score meets that need. It objectively determines the IP video playback quality as perceived by the end-user and is expressed as a single, final overall score. VVQ would be applicable to any IP video playback use cases such as IP linear, VOD, and time-shifted streams.

3.1. VVQ Model

The VVQ model uses several metrics to compute the single final overall score. They are listed below.

- Video quality (as a result of lossy encoding)
In this model, variants/segments in the ABR ladder are classified into three video quality categories Low, Medium, High. There are several classification options that may include:
 - Video bitrates
 - Complex video analysis
 - Video quality metric from Peak Signal to Noise Ratio (PSNR), Structural Similarity Index (SSIM) based tools
 - Quantization values
- Quality switching (Downshift to Low/Medium video quality segments)

Considered in the quality switching are peak frequency value and number of downshifts (High to Medium, Medium to Low, High to Low)

- Low/medium video segments playback
Considered here are
 - The total playback duration of these lower quality segments during the entire measurement period
 - Continuous playback duration: duration of low/medium video playback before switching to optimal quality segment
- Rebuffering events
 - Total duration of rebuffering during the entire measurement period
 - Peak frequency value for these events
 - Continuous duration of rebuffering event
 - Number of rebuffering events
- Startup latency (Primacy)
- Last bad video quality impact (Recency)
Elapsed time since the last bad video quality event such as low-quality video or rebuffering
- Playback failures/errors
Considered here are failures and errors that disrupt video playback like frozen video, loss of audio, or error messages displayed to the user. During these types of errors/failures the following are determined
 - Total duration of error events during the entire measurement period
 - Peak Error frequency
 - Continuous duration of error event
 - Number of failures/error events
- End-user actions
Latency when executing trick mode commands like Pause, Seek, Fast Forward, Rewind
- Display dimensions

3.2. VVQ Formula

The formula to calculate VVQ score is shown next. It is an equation that factors in several IP video playback artifacts. From the customer point of view the main artifacts impacting them are quality of video and audio, rebuffering, startup latency, playback failures, quality switching, and time between bad video quality events.

Video Viewing Quality Score = 100 - Impact from Low Quality Video Playback - Impact from Medium Quality Video Playback - Rebuffering impact – Impact from the time interval between bad quality events - Playback startup times impact - Impact from Playback Errors - Impact from trick play latency

Expanding this to the actual terms used in the formula we would have,

VideoViewingQualityScore = 100 – (LowQualVideoImpact + ContinuousLowQualityImpact + DownshiftToLowQualImpact + MediumQualVideoImpact + ContinuousMediumQualityImpact + DownshiftToMediumQualImpact + ContinuousRebufferingImpact + RebufferingFrequencyImpact + OverallRebufferingImpact + LastBadQualityEventImpact + StartupTimeImpact + PlaybackFailureImpact + ContinuousErrorImpact + ErrorPeakFrequencyImpact + OverallErrorImpact + SeekLatencyImpact)

The terms in the above equation are further detailed below.

1. Low-quality video impact computation: This measures the overall impact from the playback of low-quality video at certain times or during the entire measurement period.

$$\text{LowQualVideoImpact} = A1 * \frac{m1}{n1}$$

$A1 = \text{Low Quality Video Impact Coefficient}$

$m1 = \text{Total Duration of low quality video playback}$

$n1 = \text{Total Duration of video playback}$

2. Medium-quality video impact computation: This measures the overall impact from the playback of medium-quality video at certain times or during the entire measurement period.

$$\text{MediumQualVideoImpact} = A2 * \frac{m2}{n1}$$

$A2 = \text{Medium Quality Video Impact Coefficient}$

$m2 = \text{Total Duration of medium quality video playback}$

3. Continuous low video quality impact computation: Computes impact based on how long the user client device is playing low-quality video before it switches to either medium or optimal quality video.

$$\text{ContinuousLowQualityImpact} = \sum_{l=0}^{m3} A3 * e^{A4 * t(l)}$$

$A3 = \text{Continuous Low Quality Impact Coefficient}$

$A4 = \text{Continuous Low Quality Impact Exponent Coefficient, } < 1$

$t(l) = \text{Low quality continuous duration,}$

$l = 0 \text{ to } m3. m3 - \text{Total number of continuous low video quality events}$

4. Continuous medium-video quality impact computation: Computes impact based on how long the user client device is playing medium-quality video before it switches to optimal or low-quality video.

$$\text{ContinuousMediumQualityImpact} = \sum_{l=0}^{m5} A5 * e^{A6 * t(m)}$$

$A5$ = Continuous Medium Quality Impact Coefficient

$A6$ = Continuous Medium Quality Impact Exponent Coefficient, < 1

$t(m)$ = Medium quality continuous duration

$l = 0$ to $m5.m5$ – Total number of continuous medium video quality events

5. Downshift to low video quality impact computation: Computes impact when user client device is downshifting to low-quality video during playback; uses peak downshift frequency and total number of downshift events

$$\text{DownshiftToLowQualImpact} = A7 * m7 + m4 * C1$$

$A7$ = Downshift to Low Quality Impact Coefficient

$m7$ = Peak frequency of downshift to low quality events

$m4$ = Total number of downshift to low quality events

$C1$ = Impact to score from downshift to low quality

6. Downshift to medium-video quality impact computation: Computes impact when user client device is downshifting to medium-quality video during playback; uses peak downshift frequency and total number of downshift events

$$\text{DownshiftToMediumQualImpact} = A8 * m8 + m6 * C2$$

$A8$ = Downshift to Medium Quality Impact Coefficient

$m8$ = Peak frequency of downshift to medium quality events

$m6$ = Total number of downshift to medium quality events

$C2$ = Impact to score from downshift to medium quality

7. Continuous rebuffering impact computation: Computes impact based on how long the user client device is rebuffering before it restarts video playout.

$$\text{ContinuousRebufferingImpact} = \sum_{p=0}^{m9} A9 * e^{A10*t(p)}$$

$A9$ = Continuous Rebuffering Impact Coefficient

$A10$ = Continuous Rebuffering Impact Exponent Coefficient, < 1

$t(p)$ = Continuous Rebuffering duration,

$p = 0$ to $m9$. $m9$ – Total number of rebuffering events

8. Rebuffering frequency impact computation: Computes impact based on how often the user client device is rebuffering.

$$\text{RebufferingFrequencyImpact} = A11 * m15$$

$A11$ = Rebuffering Frequency Coefficient

$m11$ = Peak frequency of Rebuffering events

9. Overall rebuffering impact computation: Computes the overall impact from rebuffering during the entire video playback.

$$\text{OverallRebufferingImpact} = A12 * \frac{m12}{n1}$$

$A12$ = Overall Rebuffering Impact Coefficient

$m12$ = Total Duration of video rebuffering

$n1$ = Total Duration of video playback

10. Last bad video quality event impact: A logarithmic curve is used to compute the impact to user viewing experience for the time intervals between bad video quality events such as a downshift to a low-quality video, rebuffering but then recovers after a certain time. Therefore, this impact is dependent on the time between the bad video quality events.

$$\text{LastBadQualityEventImpact} = \sum_{k=0}^{m10} C1 + A14 * \log_{m11} t(k)/A15$$

$C1$ = Last Bad Quality Event Impact Constant

$A14$ = Last Bad Quality Event Impact Coefficient

$m11 = \text{Last Bad Quality Event Impact Base} < 1$

$t(k) = \text{Time since last bad video quality event,}$

$k = 0 \text{ to } m10.m10 - \text{Total number of bad video quality events}$

$A15 = \text{Time scale factor}$

The time of occurrences of low-quality video, rebuffering, failure events are used to determine $t(k)$, the time since last bad video quality event.

11. Playback failure impact: This would be applicable to failures such as unable to playback video due to inability to retrieve video segments or unable to decode video or corrupt video without recovery.

$\text{PlaybackFailureImpact} = A16$

$A16 = \text{Playback Failure Constant}$

12. Startup time impact: Computes impact based on how long it takes for video playback to start after the user has initiated it. An exponential growth curve is used to compute the impact.

$\text{StartupTimeImpact} = A20 * e^{A21*t(z)}$

$A20 = \text{Startup time Impact Coefficient}$

$A21 = \text{Startup time Impact Exponent Coefficient,} < 1$

$t(z) = \text{Startup time}$

13. Seek latency impact: Computes impact based on how long it takes for video playback to start after the user has initiated a seek or trick mode operation like FFWD or RWD. An exponential growth curve is used to compute the impact.

$$\text{SeekLatencyImpact} = \sum_{p=0}^{m14} A22 * e^{A23*t(w)}$$

$A22 = \text{Seek Latency Impact Coefficient}$

$A23 = \text{Seek Latency Impact Exponent Coefficient,} < 1$

$t(w) = \text{Seek Latency}$

$p = 0 \text{ to } m14.m14 - \text{Total number of seek events}$

14. Continuous error impact computation: Computes impact based on how long the user client device is in error state before it restarts video playout.

$$\text{ContinuousErrorImpact} = \sum_{p=0}^{m15} A24 * e^{A25*t(p)}$$

A24 = Continuous Error Impact Coefficient

A25 = Continuous Error Impact Exponent Coefficient, < 1

t(p) = Error state duration,

p = 0 to m15. m15 – Total number of error events

15. Error frequency impact computation: Computes impact when the player runs into error events, uses peak error frequency and total number of error events.

$$\text{ErrorPeakFrequencyImpact} = A26 * m16 + m17 * C3$$

A26 = Error Peak Frequency Impact Coefficient

m16 = Peak frequency of error events

m17 = Total number of error events

C3 = Impact to score from all error events

16. Overall error impact computation: Computes the overall impact from errors during the entire video playback.

$$\text{OverallErrorImpact} = A27 * \frac{m18}{n1}$$

A27 = Overall Error Impact Coefficient

m18 = Total Duration of Error events

n1 = Total Duration of video playback

3.3. VVQ Coefficients Calibration

Calibration of VVQ coefficients in the above formula involved analyzing IP video playback data from production, publicly available Mean Opinion Score (MOS) data, and technical documents specifying the customer thresholds for various IP video playback issues.

In particular, the publicly available University of Waterloo QoE database test samples with MOS scores using 25 subjects was utilized [1].

4. VVQ Implementation

VVQ scoring implementation could be done either in IP Video Players or remote monitoring system/analytics engine as shown in the Figure 2. In the case of the VVQ scoring in analytics engine, the players are expected to emit playback data such as timestamps for bitrate changes, rebuffering start and stop times, and other events, to the analytics engine to permit VVQ scoring there. In the case of VVQ scoring in the IP video players, the calculated score is communicated to the remote monitoring system. The remote monitoring system analyzes the VVQ scores and determines the measures that need to be taken to address issues resulting in low scores.

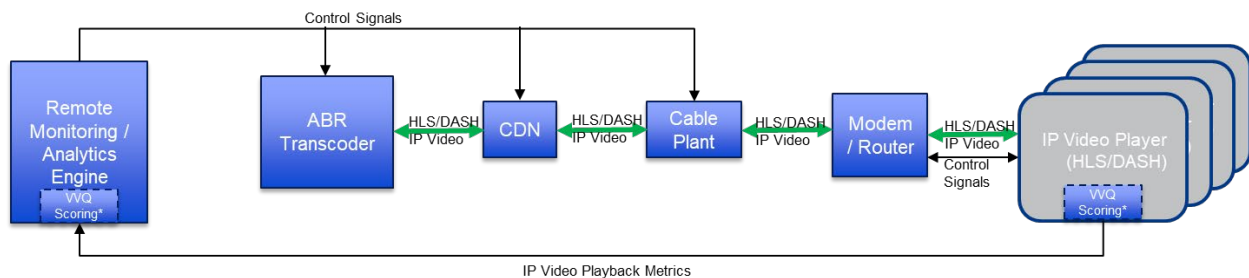


Figure 2 – Scoring in Analytics Engine or IP Video Player for Video QoE Management

4.1. VVQ in Analytics Engine Reference Architecture

Figure 3 below shows one implementation of VVQ in Analytics Engine. The IP video players in Set Top Boxes (STBs), Connected TVs, Mobile Devices and Desktops send player analytics data to Headwater 3 during IP video playback. Details on the Headwaters is available here [7]. Data in Headwaters is then stored in a Data Lake and ingested by ETL AWS Glue VVQ code to generate the scores. These scores are plotted using QuickSight user interface considering various parameters such as device type, application type, time of occurrence, and others.

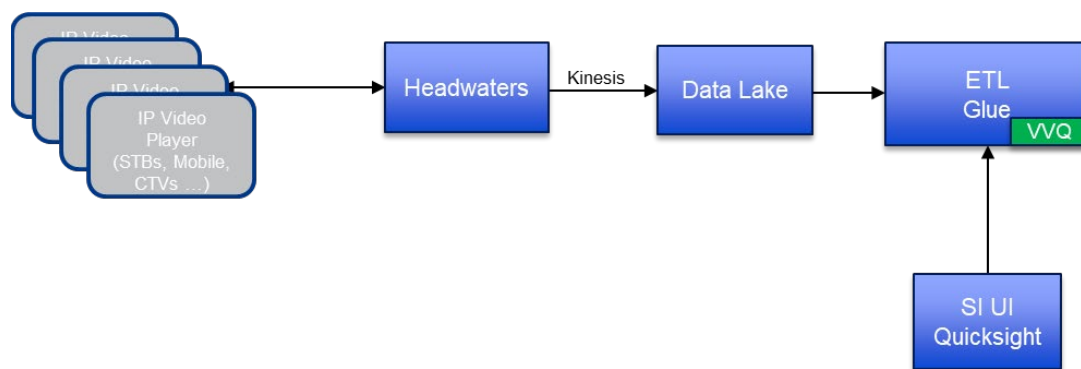


Figure 3 – VVQ Implementation in Analytics Engine

Variants/Segments in the ABR ladder are classified into three video quality categories - Low, Medium, High. The classification could be done via several approaches such as no reference or reference video encoding quality measurement tools, or simply based on video bitrates. The approach chosen here is using the video bitrates and their assignments to the three quality categories is predetermined.

The VVQ scores can be computed at various points in time during video playback. Some of the options are:

- At the end of video playback sessions, for example at the end of a movie playback, end of linear channel viewing session
- At predetermined fixed intervals during VOD or linear playback session
- Continuously as, and when, new playback data is available
- At the end of a program during linear channel viewing utilizing program metadata from Gracenote

In the architecture here the VVQ is computed at the end of video playback session.

Figure 4 below shows the screenshot of the VVQ scores generated from an implementation of this reference architecture for Linear IP Video Playback sessions.




Figure 4 – VVQ Average Scores Linear

Figure 5 below displays the VVQ score components average values for Linear IP Playback sessions. The components are the various terms detailed in the section 3.2 VVQ Formula.

Linear Sessions: Daily Average of VVQ Components




Figure 5 – VVQ Score Components Average values

4.2. Sample VVQ Scores

Listed in the table below are sample VVQ scores generated from IP video playback sessions across various devices.

Table 1 – Sample VVQ Scores

Device	VVQ Score	Description
STB	0	Total Linear playback duration: 88 mins 2 Rebuffering events, Max duration 7 mins 6 downshifts to bitrate 0.8Mbps (34s on low bitrate video*) 26 downshifts from 6 to 2Mbps (4mins on medium bitrate video*) In home issue, attributed to packet loss and Wi-Fi latency
STB	73.34	Total Recorded content playback duration: 51 mins 5 downshifts to 2.1Mbps Video (13mins on medium bitrate video*)
iPhone	0	Total VOD playback duration: 35 mins 9 mins on low bitrate video**
iPhone	82	Total VOD playback duration: 36 mins 3 Rebuffering events, Max duration 8s 2 downshifts to bitrate 0.7Mbps (31s on low bitrate video**)

Device	VVQ Score	Description
Desktop Edge	0	Total Linear playback duration: 48 mins 2 downshifts to low bitrate video (20mins on low bitrate video**) 3 downshifts to medium bitrate video (28mins on medium bitrate video**) 1 Rebuffering event, Max duration 3s
Android	20	Total Linear playback duration: 19 mins 17 downshifts to bitrate 0.5Mbps (3 mins on low bitrate video**) 56 downshifts to medium bitrate video 2Mbps (7mins on medium bitrate video**)

* Low bitrate video < 1.8Mbps, Medium bitrate video < 3Mbps, Video compression – AVC

** Low bitrate video < 0.75Mbps, Medium bitrate video < 1.5Mbps, Video compression - AVC

Please note that “We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws”.

5. IP Video Playback QoE Management – Self healing

One important objective of this scoring is to drive a self-healing feature of IP video playback issues. In the Figure 2 above the remote monitoring system based on the VVQ scores calculated interacts with ABR transcoder, CDN, cable plant, and with the client to control changes that will lead to better playback video experience (higher VVQ score). The client interacts with modem/router, for example, to switch to channels with less interference.

Possible self-healing scenarios are listed below,

- If the score is low due to in home Wi-Fi issues, automatically switch channels with less interference; also recommend upsell to better routers. This is made possible when the Wi-Fi Gateway and routers network analytics data is accessible by the IP Video playback device and is also controllable by it.
- If the in-home network bandwidth is maxed out causing low scores, automatically increase bandwidth and display upsell packages for higher BW product to customer
- If unable to fix occasional network issues via configuring in-home Wi-Fi equipment, display error messages for users to troubleshoot and identify possible sources of interference
- Detect low scores with new deployments and pull back releases/upgrades without manual intervention
- Detect low scores from all clients in a plant attributable to a plant DOCSIS channels outages and/or incorrect configuration and switch to redundant paths
- Detect CDN delivery issues (e.g., network connectivity failures) and allocate additional resources on the existing CDN or switch to alternate CDNs. If the issue is with an external CDN then potentially interface with their control interface to address these issues and indicate so to the customer.
- Switch to alternate transcoder delivery lanes when consistent media failures are detected, for example on linear channels.
- If most IP video players in a region are reporting low scores due to network BW capacity saturation, the remote monitoring system could instruct the transcoder to generate the highest profile stream in the ABR ladder at a lower bitrate that would reduce the saturation
- If the score is low due to out-of-home Wi-Fi/Carrier signal issues, provide scores to carriers for alternate solutions and recommend upgrades
- Proactively inform users of upcoming IP video outages

6. Conclusion

Measuring IP Video Playback Quality of Experience is important for several reasons including transition to all IP video, early detection of IP video infrastructure issues, enable self-healing, and video customer retention. VVQ meets these needs by presenting one final overall score that holistically measures IP Video Playback QoE rather than looking at several distinct metrics individually.

The VVQ model utilizes several playback metrics such as rebuffering, playback errors, video quality, downshifts, startup latency, recency, and others to compute the score and was validated with production data and publicly available MOS scores. VVQ is flexible and conducive to implementation either in analytics engine or IP Video Player. One such reference architecture implementation and its results are shown in this paper. VVQ is actively being utilized to gauge IP video playback QoE on various IP video devices and troubleshoot playback issues.

One another key benefit is self-healing of issues in the IP video eco system utilizing VVQ scores.

Author would like to thank several Comcast Video research and engineering teams that assisted with this work.

Abbreviations

ABR	Adaptive Bitrate
AVC	Advanced Video Coding
CDN	Content Delivery Network
DASH	Dynamic Adaptive Streaming over HTTP
HEVC	High Efficiency Video Coding
HLS	HTTP Live Streaming
IP	Internet Protocol
MPEG	Motion Picture Experts Group
PSNR	Peak Signal to Noise Ratio
QAM	Quadrature Amplitude Modulation
SSIM	Structural Similarity Index
VVQ	Video Viewing Quality

Bibliography & References

- [1] QoE Modeling for HTTP Adaptive Video Streaming-A Survey and Open Challenges. NABAJEET BARMAN, (Member, IEEE), AND MARIA G. MARTINI, (Senior Member, IEEE)
- [2] H.264, Coding of Audio-Visual Objects — Part 10: Advanced Video Coding (AVC), ISO/IEC 14496-10:2010.
- [3] H.265, High efficiency coding and media delivery in heterogeneous environments — Part 2: High efficiency video coding, ISO/IEC 23008-2:2013

[4] H.262, Generic Coding of Moving Pictures and Associated Audio Information (H.262), ISO/IEC 13818-2:2013.

[5] Dynamic Adaptive Streaming over HTTP (DASH), ISO/IEC 23009-1:2014

[6] R. P. Pantos, editor, HTTP Live Streaming, <https://datatracker.ietf.org/doc/html/rfc8216>

[7] Self-Service Dimensional Data Analytics, Francesco Dorigo, 2018 SCTE Cable-Tec Expo.
<https://www.nctatechnicalpapers.com/Paper/2018/2018-self-service-dimensional-data-analytics/download>

Modernizing Subscriber Management on the Road to 10G

A Technical Paper prepared for SCTE by

Sebnem Ozer, Ph.D.

Senior Principal Architect, CONNECT
COMCAST
1800 Arch St. Philadelphia PA
215 255 5730
sebnem_ozero@comcast.com

De Fu Li

Distinguished Engineer, CONNECT
COMCAST
defu_li@comcast.com

Jason Combs

Senior Principal Architect, CONNECT
COMCAST
jason_combs@comcast.com

Bob Gaydos

Fellow II, CONNECT
COMCAST
Robert_Gaydos@cable.comcast.com

Dan Rice

VP, CONNECT
COMCAST
Daniel_Rice4@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Current Service Fulfillment Architecture	3
2.1. Challenges	6
2.1.1. Scaling	6
2.1.2. Service Agility.....	7
2.1.3. Network Programmability.....	7
2.1.4. Operations.....	8
3. Proposed Solution.....	9
3.1. Device Management Application	9
3.2. Device Database	11
3.3. TFTP Proxy	11
3.4. Lightweight REST API.....	12
3.5. Deployment	12
4. Use Cases	13
4.1. Mid-Split Enablement	13
4.2. New Service Offerings	14
4.3. vBNG.....	15
5. Conclusion.....	17
Abbreviations	19
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1 - Cable Service Fulfillment Architecture	4
Figure 2 - Cable Modem Initialization and Provisioning.....	5
Figure 3 - PacketCable MultiMedia Architecture	6
Figure 4 - E2E System Model	10
Figure 5 - DMA Component View	11
Figure 6 - DMA Deployment Diagram.....	13
Figure 7 - New Service Use Cases	15
Figure 8 - DPoE-SP-MULPIv2.0 vCM Model	16
Figure 9 - vBNG System Components Utilizing DOCSIS Based NDP	17
Figure 10 - Data-driven and Knowledge-based Systems	18

1. Introduction

Network device provisioning (NDP) is a key component of the end-to-end system for service order fulfillment and activation. NDP is responsible for the configuration and service provisioning of subscriber's network devices. The DOCSIS network devices are the D2.0, D3.0, D3.1 and soon D4.0 modems, and their derivatives, such as set-top-box (STB) with embedded cable modem, media terminal adaptor with embedded cable modem (eMTA), eRouter, and more.

The number of device types, along with the number of available services applies a multiplication factor in the number of configuration files needed to provision. In the process of supporting tens of millions of individual devices and hundreds of services, the NDP of today is managing millions of DOCSIS configuration files. As network operators continue to invest in the roadmap to 10G, expanded network capabilities will add to the number of permutations of service and the number of device types involved. In addition, customers are expecting greater service agility in terms of faster development, deployment, and support for these innovative products and services offerings.

Although the DOCSIS subscriber management term is used extensively, currently the provisioning and management are done at the device level. This paper discusses an approach to modernizing today's device management and introduces new management systems with the goal of maintaining backward compatibility, employing network automation with dynamic service assurance, and increasing service activation velocity in support of the 10G technology roadmap. Use cases for current and possible future services are also covered towards a real subscriber management system.

2. Current Service Fulfillment Architecture

The service fulfillment architecture features a number of components that are vital to providing and implementing the customer's desired services (Figure 1).

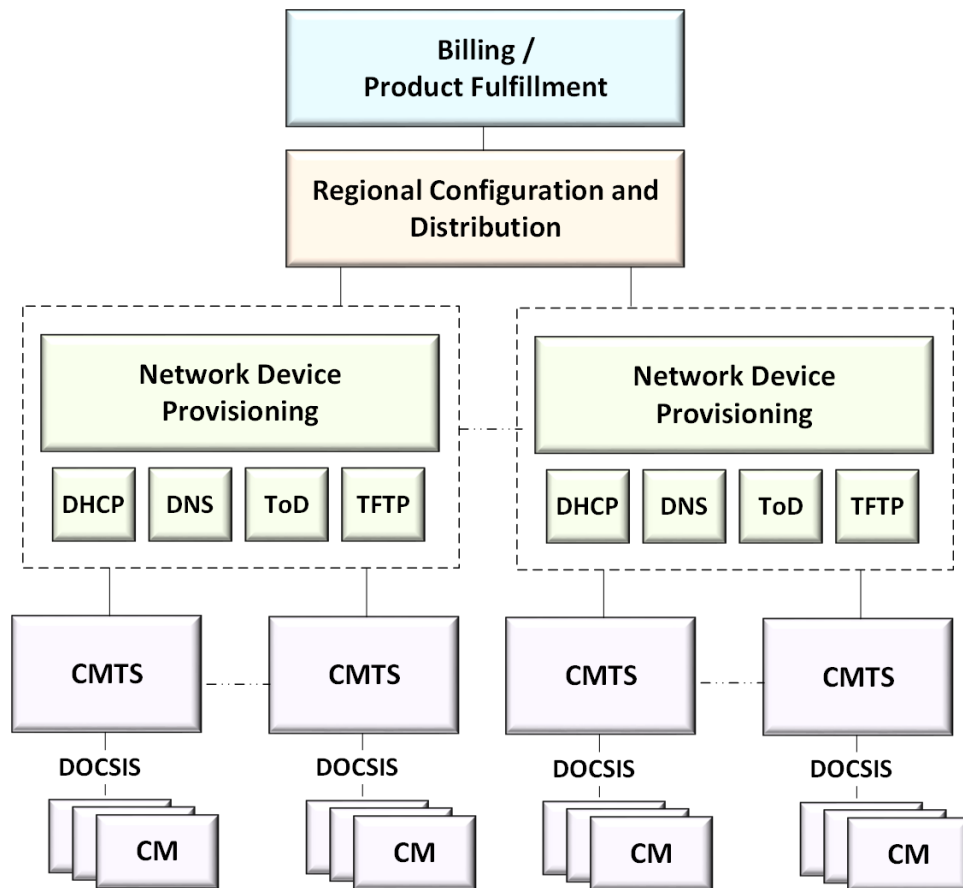


Figure 1 - Cable Service Fulfillment Architecture

As shown in Figure 1, a central provisioning and distribution system configures the servers and routers and distributes Dynamic Host Configuration Protocol (DHCP) instructions and device configuration files to the Device Provisioning Servers. These servers are integrated with the Network Registrar DHCP server to control the assignment of IP addresses for each device.

The devices supported by this architecture are the DOCSIS cable modems (CM) and embedded devices such as set-top-box (STB), media terminal adaptor (eMTA), eRouter, and more. The modem establishes DOCSIS data link layer connection to the cable modem termination system (CMTS), which in turn provides the L2 or L3 path to reach the NDP for its service provision and device configuration.

Figure 2 illustrates the main steps for CM initialization and registration. After ranging with CMTS and optional early authentication and encryption initialization, the IP initialization step takes place. The CM acquires an IP address in the cable operator address space, the current time-of-day, and a binary configuration file. The configuration files are delivered to the devices via the Trivial File Transfer Protocol (TFTP) server whose address is provided through DHCP information. During the registration step, the CMTS validates the configuration file contents sent by the CM, activates Medium Access Control (MAC) layer resources accordingly and sends MAC layer identities to the CM, as described in detail in CableLabs MULPI specifications. Services are defined statically in the configuration file except for dynamic services

that support eMTA voice service configuration file. Different service flows (SF) including High Speed Data, Voice Signal, Video data and signal, Community Wi-Fi and Business Services over DOCSIS (BSoD) services are defined with corresponding classifiers and QoS settings. Operators often use Service Class Names (SCN) that are defined in the CMTS and mapped to SFs in the configuration files but may include explicit service definition TLVs as well. A variety of other service attributes such as number of IP addresses, packet classifiers, and even diplexer settings may be included in the configuration files further expanding the unique sets of service attributes.

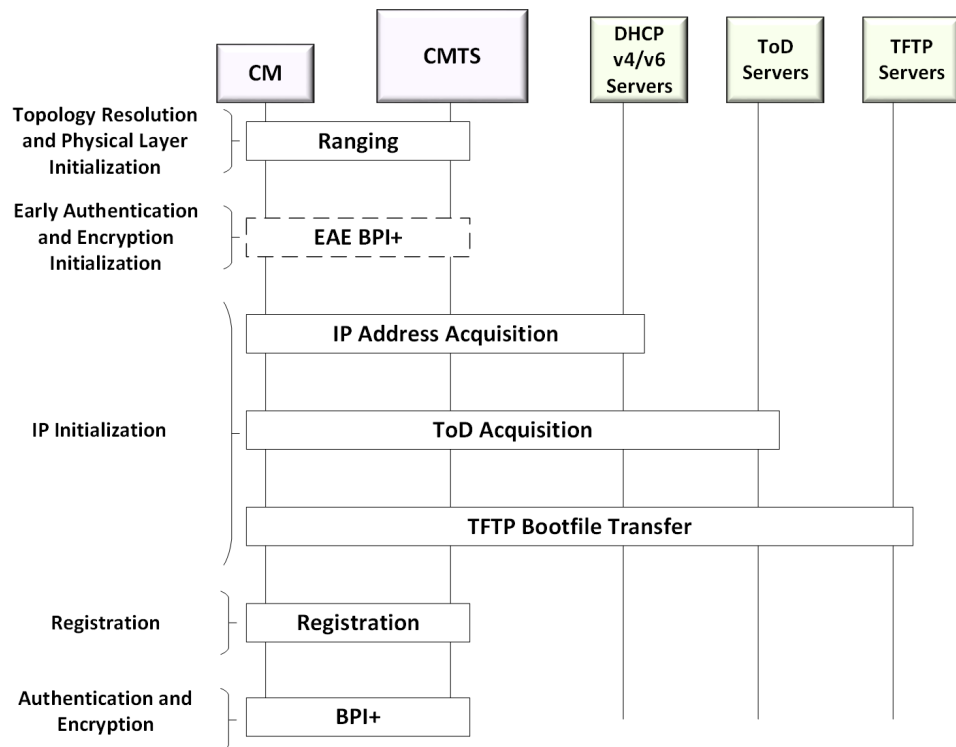


Figure 2 - Cable Modem Initialization and Provisioning

An SCN provides a set of Quality of Service (QoS) parameters for the SFs associated with it. At least two default SFs known as Primary Service Flows must be defined. As described in detail in the CableLabs OSSI specifications, the SCN also identifies the SF service characteristics to billing or customer service systems. Therefore, both SF identifier (SFID) and SCNs are used by the billing system. For most operators, the creation and integration of SCNs is a manual process.

Voice calls are established using PacketCable MultiMedia (PCMM) as shown in Figure 3. For eMTA voice calls, the SIP messages are delivered to P-CSCF that provides the reservation information to Policy Server (PS) for CMTS gate instructions. The Gate settings include classifiers and QoS settings defined for Upstream Unsolicited grant service (US UGS) and Downstream Best Effort (DS BE) calls. The PCMM architecture can also be used for non-voice SFs such as HSD SFs. All HSD dynamic SFs must be created using an SCN defined in the CMTS if the service is processed by the billing system with the current

architecture. The details on the PCMM architecture and OpenDayLight implementation of a Policy Server can be found in [6].

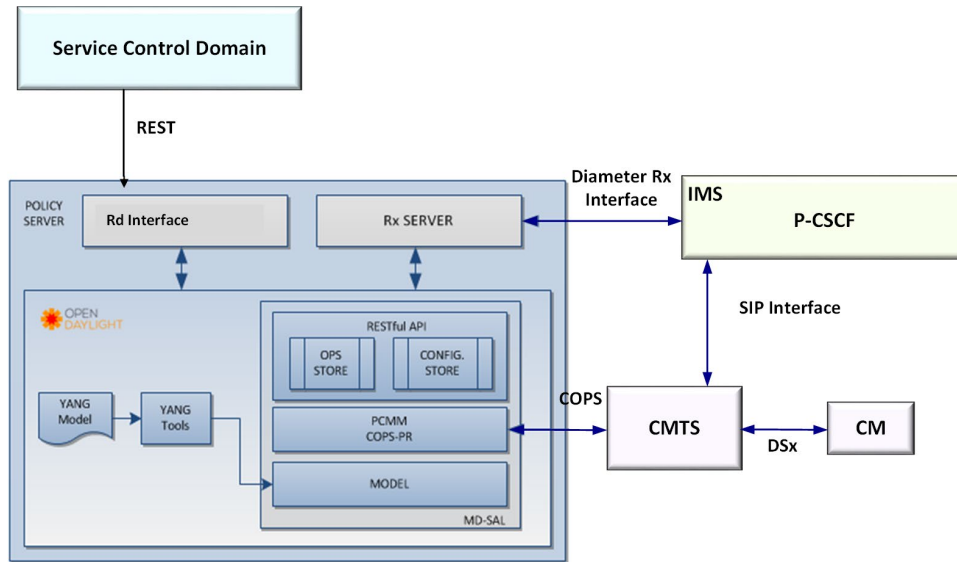


Figure 3 - PacketCable MultiMedia Architecture

2.1. Challenges

This section covers the challenges operators face with traditional network provisioning systems with the focus on DOCSIS CM and service provisioning.

2.1.1. Scaling

New configuration files are introduced for new speed tier rates and CM models and versions. Different number of CPE IP addresses and security attributes also require new configuration files. Furthermore EOL devices that are not swapped but required to support additional speed tier rates increase this number. If the configuration file management is not effective, e.g. not fully automated, many unused configuration files may exist within the operator's active database. In the process of supporting tens of millions of individual devices and hundreds of services, the NDP of today may be required to manage millions of DOCSIS configuration files.

Manual processing steps in the creation, update and deletion of a high number of configuration files may create issues in the existing deployments and new service launch.

2.1.2. Service Agility

Service agility for the 4As to have “Any content, Anywhere, Anytime and on Any device” is the ultimate goal for the ISPs. The key points of service agility are:

- Fast development of innovative products and services
- Configuring the services according to changing customer behaviors and requests
- Deploying new products and services without breaking existing systems and creating new issues
- Operating the services and underlying networks through proactive management techniques

A survey conducted by the TM Forum in October 2018 [1] shows that legacy Operations and Business Support Systems (OSS/BSS) remain the biggest challenge to network transformation causing slow progress. However, more recent TM Forum research [2] shows that 80% of ISPs have started to implement or benefit from digital transformation programs during the Covid-19 pandemic, which shifted the focus to customer experience and operations digitalization. Modular and configurable BSS are being adopted to reduce customization costs.

To achieve optimized service lifecycle management [3], some existing MSO systems, which have not been changed for more than two decades, must be replaced by a new design with automated service workflows and microservices. Today, the workflows for a new service and subscriber activation include manual steps for device and resource provisioning. The system is prone to human error, not suitable for fast activation and deployment and hard to operate and provide service assurance. These systems require customer equipment to be staged and configured initially in the warehouse and then shipped to the customer premises for installation. It also creates more need for dispatching technicians to activate services.

In these systems, initial estimations of the service group size and traffic load are used to project current and future scaling requirements. Therefore, service to resource mapping is not adaptable to more granular service and network changes. Traditional polling-based monitoring of data and manual troubleshooting create obstacles for automation and scaling.

To mitigate these issues, operators have started to automate the workflows gradually. A critical workflow that needs automation is the network infrastructure and device configuration.

2.1.3. Network Programmability

Service agility requires dynamic mapping of services to underlying network resources and automated provisioning based on traffic characteristics and service policies. The network resources may be physical or virtual function components and can be located in customer premises, the field, edge network, core network or in the cloud. The same network functionality may be distributed in different locations and scales.

Today, MSO networks consist of different access networks such as HFC, PON, Carrier Ethernet and Wireless. Redundant access networks and multi-path technologies are also being deployed with new service requirements. Although access networks have different resources and functionalities, the service to resource mapping and automated provisioning of workflows must be unified for effective service deployments and operations [3,4]. This can be achieved through microservices that can be reused by

different technologies. In this case, the services may be defined without requiring specific knowledge on the underlying technology. Such abstraction is crucial for both service agility and network programmability.

As an example of the current obstacles, today in most Cable operators' networks, configuration file TLV settings and SCN settings are provisioned separately. This creates an obstacle to change the network and device settings for customer needs without requiring a reset and interruption of the services. If a combination of SCN and explicit TLVs is used for modem provisioning, having a TLV value to be optimized based on an SCN parameter may not be feasible in this case. Optimization may require a vendor or the operator to introduce a new firmware code. This approach also makes network functionality upgrades impossible or very hard for early versions of the network components. For instance, a new optimization defined in the specifications can be applicable for both D3.0 and D3.1 modems but the operator may choose to apply it only for D3.1 since the code upgrades by the vendor for D3.0 may not be active.

Another example is the introduction of new speed tiers in the upstream direction. Mid-split and High-Split Enablement use cases involve CMs to range to OFDMA channels. This process has the following issues:

- It requires reconciliation among different tools and settings (SCNs and per CM Configuration file CoS)
- Backward compatible solutions must exist for iCMTS/iOLT but allow for more flexible solutions for vCMTS/vOLT
- Some of the new systems like FDX and PON do not use diplexers, requiring different provisioning
- Each new system must support different new modem models which require new configuration files

Not having programmable systems, e.g. requiring more configuration files also has direct scaling consequences.

One gradual step MSOs may take is to use PCMM based Dynamic Service Flow management for non-voice services. The current deployments are mostly for voice calls that are based on RSVP type resource allocation and management. This system may not be optimal for general dynamic service flow management and many ISPs have to pay license to a third party due to implementation and management complexities. Gradually, open source and APIs should be developed and adopted for all SF types.

2.1.4. Operations

Today manual and siloed workflows and resource management systems create obstacles for lighter operations and automated processes. Functionalities common to multiple access technologies such as subscriber management filters are implemented and managed differently depending on the access network. Therefore, multiple components with the same functionality are configured and operated for different implementations and APIs.

Services and network components are intertwined. Service definitions still include access network attributes, and they are not access technology agnostic. The billing system must know the SCNs defined in the access network provisioning system. The lack of abstraction is a challenge on the road to automated

service operations. Therefore, a unified service chaining over digitized network components cannot be deployed over multiple access technologies.

Having operational parameters of a CM provisioned using configuration files creates obstacles for a self-optimizing system without service disruption.

Therefore, there is a need for new network provisioning and subscriber management systems. The next section describes a solution that can be implemented as a first step towards completely agile service and automated network management.

3. Proposed Solution

3.1. Device Management Application

The proposed solution must be backward compatible, interoperable, and incremental because the existing back-end system is vast and impossible to replace as a single unit. Backward compatibility and interoperability are required for the front-end facing interfaces to cable modems and backend facing interfaces to the existing Billing-Production Fulfillments and NDP components.

The proposed solution is to introduce the device management application (DMA) component to proxy the interfaces between NDP, CMTS and modem devices. The DMA provides a light-weight API interface as an alternative, programmatic provisioning interface that enables a path toward modernizing the backend systems. Figure 4 illustrates the end-to-end (E2E) system model view.

The DMA has the following components (Figure 5):

- TFTP Client for downloading modem configuration file from NDP on behalf of modem devices,
- TFTP Server for modem devices to download their configuration file,
- Configuration File Manager (CFM) for managing the DOCSIS modem configuration file,
- API server for external orchestration application to programmatically managing modem devices provisioning,
- Device Database for modem devices configuration and provisioning attributes

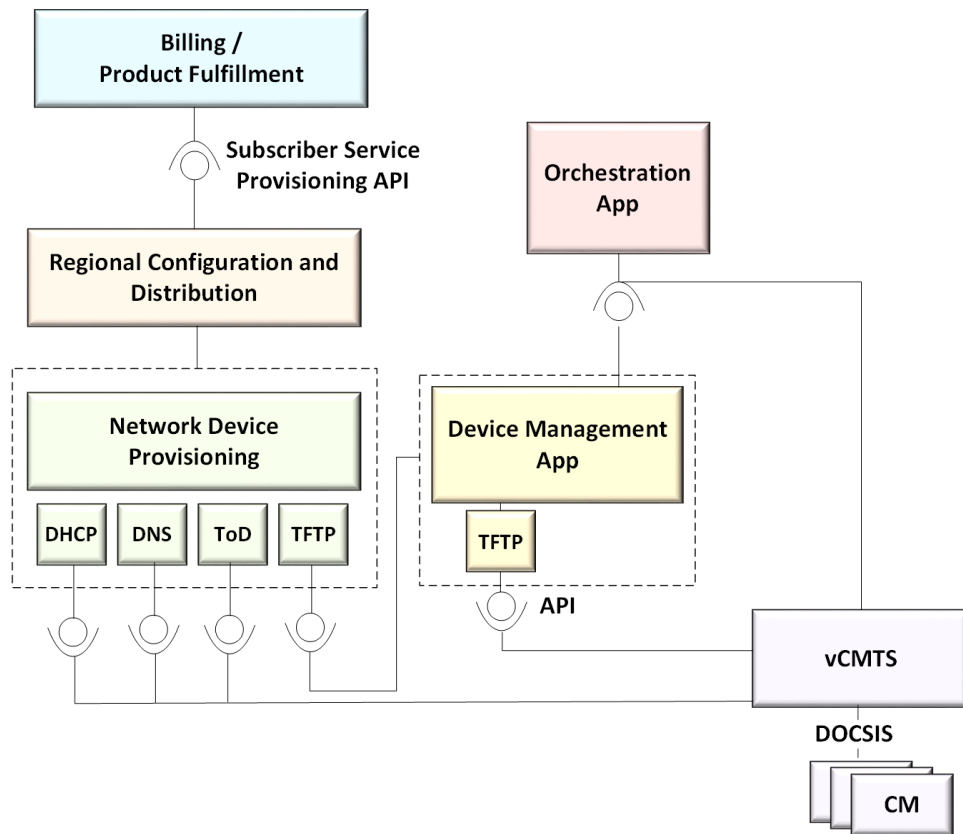


Figure 4 - E2E System Model

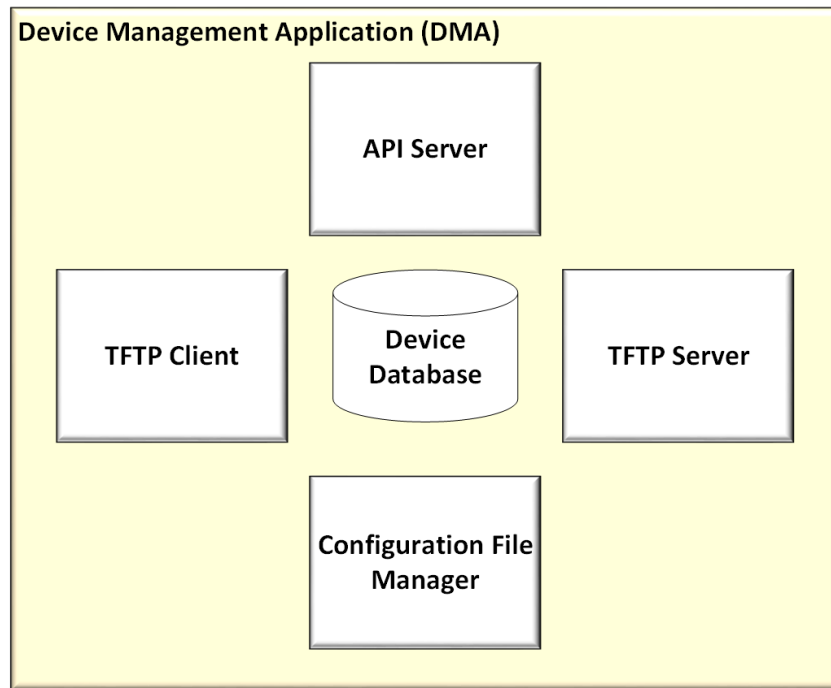


Figure 5 - DMA Component View

In the following subsections, the key functions and concepts are described.

3.2. Device Database

The device database must be persistent and highly available. Its database schema must support all objects and attributes in the DOCSIS modem configuration file. The schema must be extensible to support devices such as ONU for PON. The schema could support classifiers, policy rules and intent objects.

Automated database management without manual processes is needed to reduce errors and increase the speed of introducing new devices and services. Otherwise, increased number of device configurations may not be manageable and create scale and maintenance issues.

3.3. TFTP Proxy

DMA is the TFTP proxy server, consisting of TFTP client, TFTP server and configuration file manager (CFM) components. The configuration file manager downloads the configuration file on behalf of the cable modem. The configuration file is parsed, and the provisioned objects and their attributes are stored in the device database. At a high level, the sequence of events is:

- When modem sends DHCP request, CMTS DHCP relay agent would replace DHCP options 66 (TFTP server name) and 150 (TFTP server IP address) to DMA's internal TFTP server FQDN and IP address.
- CMTS then notifies the CFM.
- CFM initiates the TFTP client to download the modem configuration file from the NDP on behalf of the cable modems.
- CFM parses, validates, and updates the device database.
- CFM finally regenerates a new configuration file for the modem to TFTP download.

The DMA enables the combination of NDP provided configuration file and locally inserted device configuration attributes by regenerating a new configuration file and offering it to the cable modem.

3.4. Lightweight REST API

As an alternative to NDP, the API server provides REST interface for other orchestration applications to programmatically provision a modem. The API provides access to the provisioned objects and the attributes in the device database.

As described earlier, the device database schema is extensible to support objects such as:

- Policy, Rules, and Intents
- Device Classifiers
- New Service Offerings

These objects will be newly defined, beyond the scope of the existing DOCSIS configuration file. The objects are meant to be consumed by the CMTS or OLT at the core side for the active management of the device devices.

3.5. Deployment

The distributed access architecture (DAA) specification was introduced to evolve and modernize access network performance. This split allowed the core functions to run on a cloud computing platform. As such, virtualized CMTS (vCMTS) is a collection of software applications for the core functions. It is built upon a microservice architectural pattern and targeted for cloud computing platforms.

Here, vCMTS as platform is referring to the deployment of vCMTS core function software in MSO private cloud. In contrast, the integrated CMTS (iCMTS) represents the traditional custom hardware CMTS platform.

During the MSO roll out of the vCMTS platform, the production environment will have a mix of vCMTS and iCMTS components. It is essential for DMA to support new service offering across all iCMTS and vCMTS systems in the footprint. With that in mind, the DMA application can be deployed either in the MSO regional data center or public clouds.

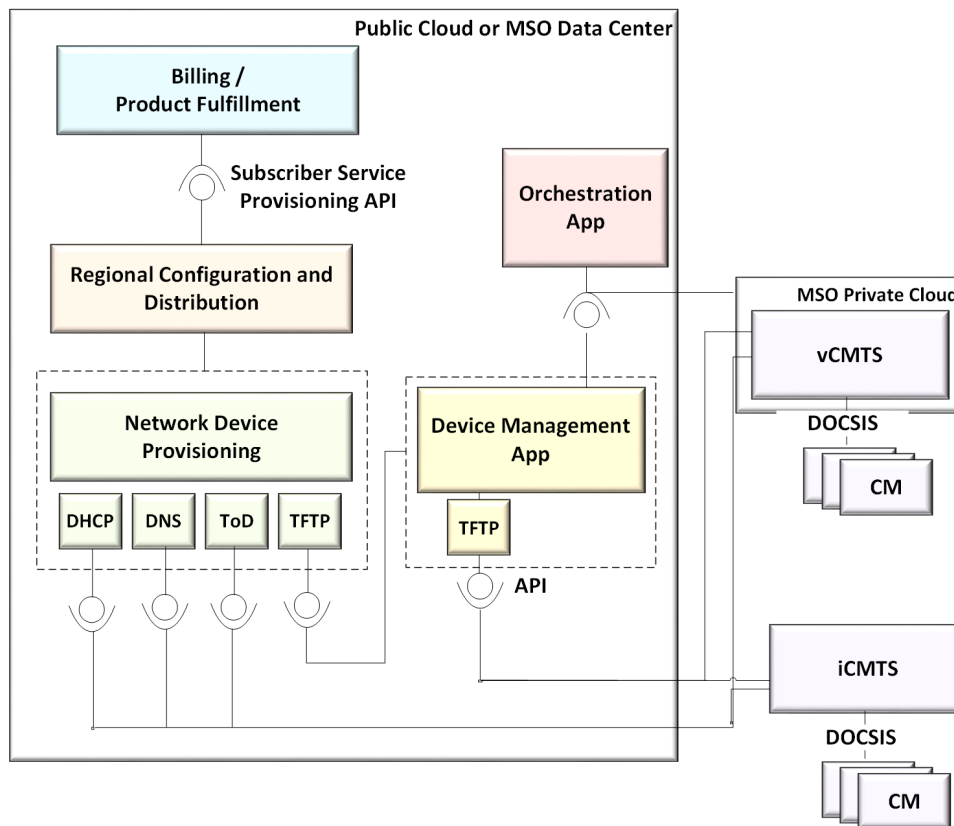


Figure 6 - DMA Deployment Diagram

4. Use Cases

4.1. Mid-Split Enablement

The benefits and challenges to Mid-split deployment are well documented in [5]. The obvious benefit is that a mid-split plant doubles the upstream bandwidth of the standard-split (aka low-split) plant. The challenges are:

- Coexistence of Mid-split (MS) and Standard-split (SS) CPE, where the MS-CPE and SS-CPE share the cable feed through a splitter. With imperfect isolation, the MS-CPE's upstream OFDMA signal may leak into the SS-CPE's downstream RF front end.
- Old infrastructure components such as drop-amps and splitter that makes MS-CPE inoperable on the OFDMA channel.

A detection and mitigation mechanism is proposed in [5] to address the above challenges. The mitigation mechanism is to steer the modem via the configuration file or DBC. Customizing modem configuration file at per device level is operationally challenging. The mechanism to DBC the modem requires external application to be developed, CMTS to expose its DBC API, and re-apply whenever the modem re-registers.

The management and mitigation mechanism can be simplified by utilizing the device management application. The interference detection tool would invoke the DMA API to update modem having SS-CPE interference or drop-amp. CMTS queries the DMA during modem registration. If the device database has

the interference attributes set, then the CMTS would enforce the OFDMA channel exclusion as part of the transmit channel set assignment process. Reports of modems having these types of interference can be easily obtained by querying the device database.

4.2. New Service Offerings

Operators are interested in providing personalized solutions for their subscribers without the need to have custom solutions for their network and service platforms. For example, services to support low latency applications, speed boost and continuation of speed rates over wireless (home and on-the go use cases) can be supported per subscriber's status and preference.

These services can be turned on and off by the subscriber or service options may be available to subscribers. In this case, an agile service platform can support runtime changes without interruption and resets. An example is provided in Figure 7 where Low Latency and Speed-Boost requests can be fulfilled dynamically using a PCMM based system. A REST API may provide the request to the PCMM that in turn opens or deletes gates for the vCMTS. As an alternative, PCMM gates may be integrated with Device Management App to facilitate future architectures (Figure 7).

Current PCMM specifications, at the time of publishing this paper, do not include Aggregate SF cases as defined in Low Latency DOCSIS specifications. However there is an increasing interest to extend PCMM SF definitions and have Policy Server implementations that are based on open source code and standard APIs.

This architecture can be further extended to use DMA only at the first initialization of subscriber devices and services or to remove it completely. In the former case, a basic common configuration file may be used to initialize the CM with primary SFs. Then, actual SFs can be dynamically provisioned using Orchestration app and PCMM. When backward compatibility is not needed and standard APIs may be used directly to configure the CMs, DMA may be replaced. This architecture is also a first step to have automated subscriber management by provisioning network and device components per subscriber service status and preferences.

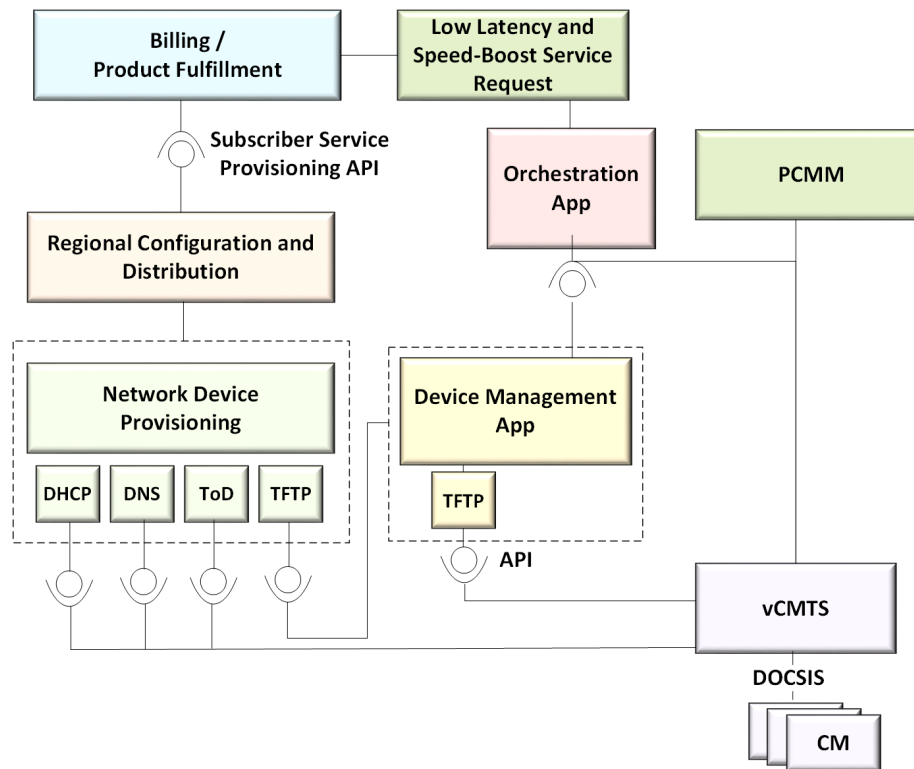


Figure 7 - New Service Use Cases

4.3. vBNG

CableLabs' DOCSIS Provisioning of EPON (DPoE) specification enables MSO to deploy EPON technology using the existing DOCSIS based backend systems. As shown in Figure 8 the ONU is emulated as virtual cable modem (vCM).

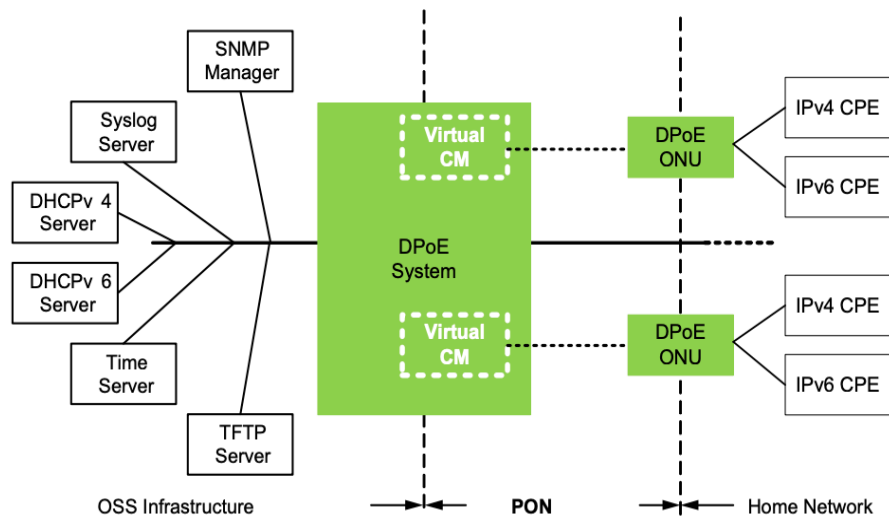


Figure 8 - DPoE-SP-MULPIv2.0 vCM Model

As MSOs are rolling out vCMTS platforms, the same private cloud can host virtual broadband network gateway (vBNG) [4] application to support EPON technology.

vBNG can be developed to utilize DOCSIS based NDP. In this case, vBNG invokes DMA provided API to initiate the download of the vCM configuration file and to query for the ONU device and service provisioning objects as shown in Figure 9.

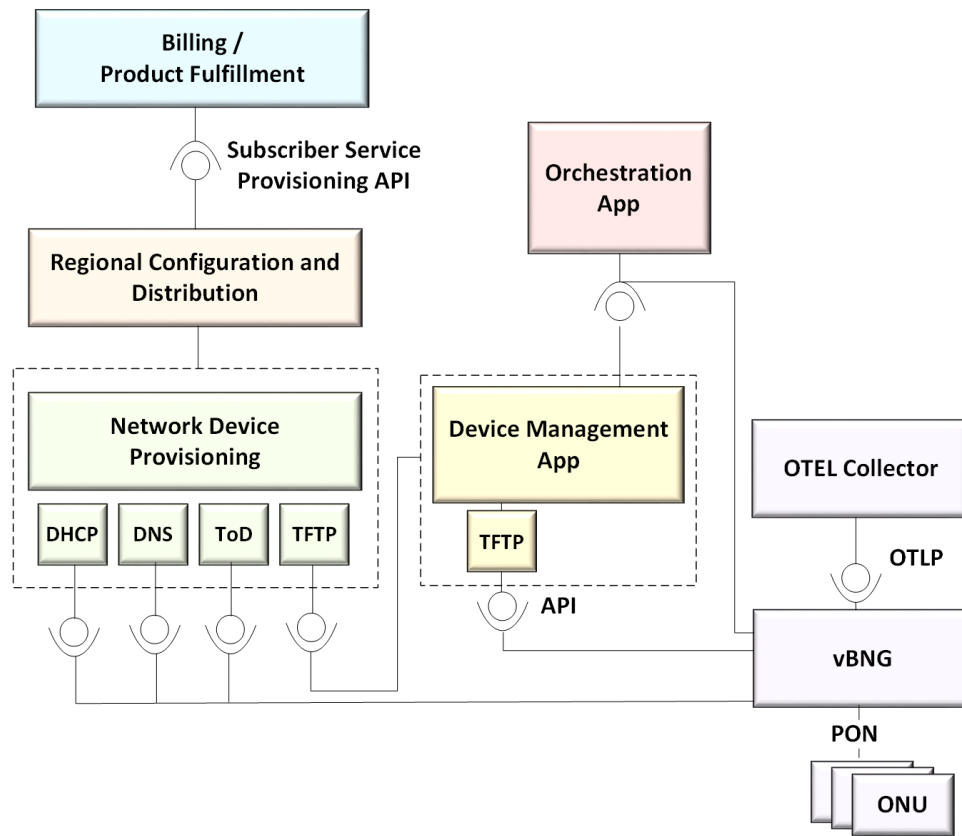


Figure 9 - vBNG System Components Utilizing DOCSIS Based NDP

The observability aspect of the operational support systems (OSS) for vBNG can be modernized by utilizing vendor-neutral open-source observability framework such as OpenTelemetry (OTEL).

5. Conclusion

The proposed DMA provides a path forward for modernizing the device management systems. It maintains backward compatibility and interoperability with the existing NDP systems. At the same time, it enables steps toward agile service deployments with programmable and automated networks. The DMA provides a programmatic interface to external applications to access the device database. The external application can be a CMTS, utilizing the DMA for the active management of subscriber services.

The proposed device management systems enable the operator to meet the goals of maintaining backward compatibility, employing network automation with dynamic service assurance, and increasing service activation velocity in support of the scalable 10G technology roadmap. The solution can be extended to unify management of access networks and introduce dynamic services with low latency and speed-boost support. These new services can be personalized based on the subscriber's status and preferences using dynamic SF provisioning. This architecture enables MSOs to integrate device, subscriber and service management as a unified system.

This is a crucial step to achieve optimized service lifecycle management. Operators started to transition to automated service workflows and microservices and have been deploying SDN, NFV and cloud enabled access networks [3]. The new system supports existing services while being upgraded for new services such as Gig Symmetric (Mid-Split, High-Split and FDX), MVNO and Low Latency services. As more services are being introduced at a faster pace and service personalization is becoming more crucial, a fully automated system as illustrated in Figure 10 is needed. The orchestrator and controllers use telemetry data and system knowledge to map services to resources, instantiate and manage the service and network microflows, and program and configure resource components. Once such a system is adopted, it can support dynamic changes to support customer runtime requests or network needs. Abstraction between services and resources can help to support multiple access technologies in a harmonized way, either as single or simultaneous technologies serving the same subscriber.

The following list describes the steps needed for modernizing device, service and subscriber management systems:

- Abstracted service and digitized resource management independent of underlying technologies
- Data and knowledge center based on cloud-enabled and push-based telemetry with standard APIs
- Data-driven orchestration for self-optimizing and scaling systems
- Microservices instead of monolithic functionalities and meshed service chains
- Containerized SW for faster initiation, efficient execution, and better isolation
- Automated inventory and topology management
- Zero touch installation and self-activation systems
- Customer centric personalization without requiring customized designs and operations

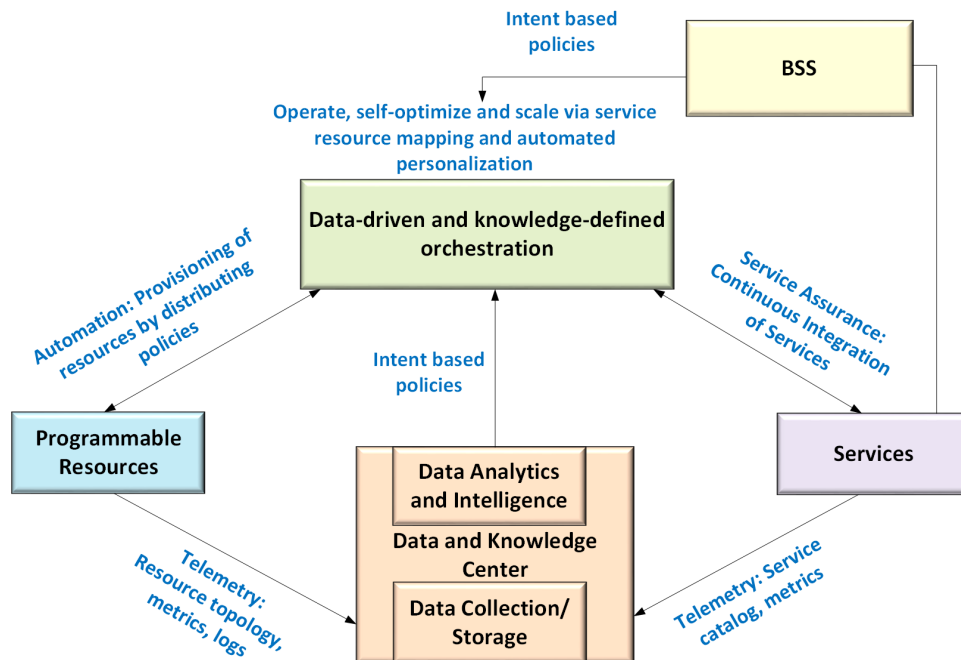


Figure 10 - Data-driven and Knowledge-based Systems

Abbreviations

API	Application Programming Interface
BSOD	Business Services over DOCSIS
BSS	Business Support System
CM	Cable Modem
CMTS	Cable Modem Termination System
DHCP	Dynamic Host Configuration Protocol
DMA	Device Management App
DNS	Domain Name System
DPoE	DOCSIS Provisioning over EPON
eMTA	Embedded Media Terminal Adaptor
HFC	Hybrid Fiber Coaxial
HSD	High Speed Data
MAC	Medium Access Control
MSO	Multiple System Operators
NDP	Network Device Provisioning
OSS	Operations Support System
OTEL	Open Telemetry
PCMM	PacketCable MultiMedia
PON	Passive Optical Network
QoS	Quality of Service
REST	Representational State Transfer
SCN	Service Class Name
SFID	Service Slow Identifier
STB	Set-top Box
TFTP	Trivial File Transfer Protocol
ToD	Time of Day
vBNG	Virtual Broadband Network Gateway

Bibliography & References

[1] *Digital Transformation Tracker 3: Why is network transformation so difficult?* TMForum Research report, by Mark Newman, October, 2018

[2] *Next Generation BSS: the case for a modular approach* by Ed Finegold, TMForum Knowledge, 2021

[3] *The Future of Operations: Building a Data-Driven Strategy* by Sebnem Ozer, Sinan Onder and Nagesh Nandiraju, SCTE Journal of Network Operations, 2019

[4] *A Virtual Broadband Network Gateway (vBNG) Approach for Cable Operators in a Distributed Access Environment*, by Jason Combs, SCTE Expo 2020

[5] *A Proactive Network Management Scheme for Mid-Split Deployment*, by Lei Zhou, Robert Thompson, Robert Howald, John Chrostowski & Daniel Rice, SCTE Expo 2020

[6] *SDN Ground Truth: Implementing a Massive Scale Programmable DOCSIS Network* by Sameer Patel, Mohammad Kabir Chowdhury, Jason Schnitzer, David Early, SCTE Expo 2016

Network Capacity Options on the Path to 10G

A Technical Paper prepared for SCTE by

Karthik Sundaresan
Distinguished Technologist
CableLabs
858 Coal Creek Cir, Louisville, CO, 80027
3036613895
k.sundaresan@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	4
1.1. DOCSIS Evolution.....	4
1.2. Technology Evolution.....	4
1.3. Spectrum Evolution	5
1.4. Modem peak speeds and service Tiers.	6
2. D3.1 Technology Options.....	6
2.1. Legacy Low Split with 750 MHz	6
2.1.1. Legacy Low split with 750 MHz, Adding OFDM, Adding 5 th SC-QAM	7
2.1.2. Legacy Low Split with 750 MHz, with OFDMA	7
2.1.3. D3.1 Low Split Upstream and 860Hz Downstream	8
2.2. D3.1 Mid-Split.....	9
2.2.1. Start with SC-QAMs	9
2.2.2. Add an OFDMA Channel	9
2.2.3. Replace with OFDMA	10
2.3. D3.1 High Split Deployments	10
2.3.1. SC-QAM and OFDMA.....	10
2.3.2. Full OFDMA	11
3. DOCSIS 4.0 ESD/FDX Technology Options.....	11
3.1. D4.0 ESD/FDD Technology Options.....	12
3.1.1. D3.1 High Split with D4.0 CM	12
3.1.2. D4.0 FDD 300 MHz Split.....	13
3.1.3. D4.0 FDD 396 MHz Split.....	13
3.1.4. D4.0 FDD 492 MHz Split.....	14
3.1.5. D4.0 FDD 684 MHz Split.....	15
3.2. D4.0 FDX Technology Options	15
3.2.1. D4.0 FDX 1.0 GHz with SCQAMs or OFDM.....	16
3.2.2. D4.0 FDX 1.2 GHz with SCQAMs.....	17
3.2.1. D4.0 FDX 1.2 GHz with OFDMA and 2 OFDMs	17
3.3. Comparison to FTTP/PON Technology Options.....	18
3.3.1. 10 G EPON	18
3.3.1. XGS PON.....	18
4. Conclusion.....	19
Abbreviations	21
Bibliography & References.....	22

List of Figures

Title	Page Number
Figure 1 – Legacy D3.0 Plant	6
Figure 2 – D3.1 Deployment on a Low Split Plant	7
Figure 3 – D3.1 Deployment on a Low Split Plant w OFDMA	8
Figure 4 – D3.1 Deployment on a Low Split Plant and 860 MHz.....	8
Figure 5 – D3.1 Deployment on a Mid-Split Plant.....	9
Figure 6 – D3.1 Deployment on a Mid-Split Plant with SC-QAM and OFDMA Channels	9
Figure 7 – D3.1 Deployment on a Mid-Split Plant with OFDMA Channel.....	10

Figure 8 – D3.1 Deployment on a High-Split Plant with SCQAM & OFDMA Channels.....	11
Figure 9 – D3.1 Deployment on a High-Split Plant with Full OFDMA Channels	11
Figure 10 – D4.0 FDD Spectrum Options	12
Figure 11 – D4.0 FDD CM on a High Split Network	13
Figure 12 – D4.0 FDD CM on UHS-300 MHz Split Plant.....	13
Figure 13 – D4.0 FDD CM on UHS-396 MHz Split Plant.....	14
Figure 14 – D4.0 FDD CM on UHS-492 MHz Plant.....	14
Figure 15 – D4.0 FDD CM on UHS-684 MHz Plant.....	15
Figure 16 – FDX Allocated Spectrum	15
Figure 17 – D4.0 FDX (1.0 GHz) Channel Allocations with SCQAMs	16
Figure 18 – D4.0 FDX (1.0 GHz) Channel Allocations with an OFDM	16
Figure 19 – D4.0 FDX 1.2 GHz Channel Allocations with SCQAMs and OFDM.....	17
Figure 20 – D4.0 FDX Channel Allocations with OFDM/A and 2 OFDMs	18
Figure 21 – DOCSIS CM DS/US capacity across technologies	20
Figure 22 – DOCSIS CM DS/US Usable spectrum usage	20

List of Tables

Title	Page Number
Table 1 – CM Spectrum Support	5
Table 2 – Summary of Peak Speeds on the Path to 10G	19

1. Introduction

Hybrid fiber cable (HFC) or data over cable service interface specifications (DOCSIS) networks are the most widely deployed technology for delivering Internet data services to the consumers. Cable operators today have variety of choices in front of them, in terms of upgrading and evolving their HFC networks. The upgrade choices with the current DOCSIS 3.1 technology includes Mid-split, High Split and distributed access architecture (DAA) and extending the downstream to 1.2 Gigahertz (GHz). With DOCSIS 4.0 technology there are even more plant upgrade options, from Full Duplex DOCSIS to the four new Ultra High Split upstream options (up to 684 Megahertz (MHz)) and the Extended spectrum to 1.8 GHz for the downstream. For each of these scenarios, operators are interested in knowing what are the data capacities that the system can realize from each of these plant upgrades. Also, for a consumer service, what are the potential service tiers that could be realized within each of these technology options. This paper gives a detailed analysis on each of these scenarios and layout the possible spectrum options for an operator given the different types of plant conditions. Understanding these options is a very important tool in figuring out what services can be reliably deployed. The cost benefit analysis for each of these options will help answer various for tactical and strategic network planning questions.

1.1. DOCSIS Evolution

DOCSIS 3.1 uses orthogonal frequency-division multiplexing (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 the upcoming DOCSIS 4.0 technology deployments in the plant. Within that pass band, digital television signals in 6 MHz channels are assumed present on the standard, as well as other narrowband and wideband digital signals.

The cable modem (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.

1.2. Technology Evolution

Allocating more spectrum to broadband is a priority for cable operators. Initially all of the downstream spectrum was allocated for carrying video programming. Over the past ~25 years, the downstream has transitioned by migrating the spectrum from video to broadband. In addition to transitioning the existing spectrum to broadband, there are several scenarios where additional downstream spectrum can be added and used for broadband service. Analyzing downstream DOCSIS capacity is different than analyzing upstream DOCSIS capacity. For the initial evolution from DOCSIS 1.0 to 2.0 to 3.0, the upstream spectrum allocation was fairly straightforward because the spectrum was mostly unused and could be allocated to broadband. In the last many years, upstream has become the bottleneck for cable networks, and is now driven by shifts in user behavior and symmetrical competition. Ultimately this became the primary driver for today's capacity upgrades, be it mid-split or DOCSIS 4.0. Downstream spectrum is much wider than upstream spectrum, and it also carries video signals. The table below shows the evolution in the spectrum supported by a cable modem for each DOCSIS version.

Table 1 – CM Spectrum Support

DOCSIS Version	Downstream Spectrum (usable)	Upstream Spectrum (usable)
DOCSIS 1.0	6 MHz	6.4 MHz
DOCSIS 1.1	6 MHz	6.4 MHz
DOCSIS 2.0	6 MHz	6.4 MHz
DOCSIS 3.0	24 to 192 MHz	25.6 (to 51.2) MHz
DOCSIS 3.1	576 MHz	192 MHz
DOCSIS 4.0	1152 MHz	Frequency Division Duplex (FDD) CMs 272/ 368/ 464/ 656 MHz Full Duplex DOCSIS (FDX) CMs 272/ 464/ 656 MHz

As shown by Table above the trend has been for each DOCSIS version to require more downstream and upstream spectrum support on the modem. The DOCSIS 3.0 specifications was a seminal moment when channel bonding was introduced. Early DOCSIS 3.0 modems would support up to 24 MHz of downstream spectrum, and the last DOCSIS 3.0 modems supported up to 192 MHz of downstream spectrum (256 MHz of downstream spectrum if Annex A quadrature amplitude modulation (QAMs) were used).

1.3. Spectrum Evolution

The cable modem termination system (CMTS) has more variability—it can support the minimum requirements as described above or it can support additional DOCSIS spectrum when there is spectrum available in the cable plant. With the move to both more broadband and Internet Protocol television (IPTV), and a recent trend is for a CMTS to support more downstream spectrum than a modem.

About 20 years ago, cable plants were carrying analog television (TV) signals. In North America, this equated to allocating 6 MHz of spectrum, one Consumer Technology Association (CTA) channel, to carry one analog TV channel in the National Television System Committee (NTSC) format. A sub-split 750-MHz system can carry about 115 CTA channels (each 6 MHz) for a total of 690 MHz of downstream spectrum.

Digital video technology was already on the rise to pack more television programming into a single CTA channel. MPEG-2 encoding allowed a single 6-MHz CTA channel to carry up to 10 standard-definition NTSC or 3 high-definition NTSC channels. MPEG-4 encoding allowed up to 17 standard-definition and about 9 high-definition NTSC channels to be put into a single 6-MHz channel.

Today, the coaxial cable is assumed to carry around 300 MHz of video programming. The amount of spectrum allocated to video programming varies widely, and 300 MHz is chosen as a "nominal" number. There are numerous solutions that can raise or lower that number, including the use of MPEG-4 encoding, switched digital video, and migration to IPTV. This number is assumed to go down as video switches to IPTV solutions.

Before 2008, the DOCSIS 1.0, 1.1, and 2.0 systems used only one downstream channel, a 6-MHz CTA channel. Since 2009, the DOCSIS 3.0 deployments allowed more downstream spectrum to be used for broadband, starting with four DOCSIS downstream channels (24 MHz) and, eventually, products that supported up to 32 downstream channels (192 MHz). The rise of digital video coincided with the availability of DOCSIS 3.0 technology, all around 2008, and spectrum that previously carried analog TV channels could now be reclaimed to carry both digital TV and broadband. As the internet bandwidth

consumption continued to rise, the need for additional spectrum for broadband became apparent. Around 2016, DOCSIS 3.1 technology was made available, resulting in the downstream needing additional spectrum for the new OFDM technology.

1.4. Modem peak speeds and service Tiers.

An old rule of thumb which multi system operators (MSOs) use is to offer a service tier which is about half of the CM's peak throughput capacity. (CM Capacity/Service Tier = 2/1) As service group bandwidth capacity increases with CMTSs providing more channels than the modem can handle, this ratio is slowly edging closer to the modem's peak capacity. We are now starting to see ratios of 1.5/1 or lower, as the higher network bandwidth capacity allows operators to take better advantage of the statistical multiplexing inherent in a medium like the DOCSIS network.

The throughput numbers described in this paper is peak broadband capacity per modem. As will be seen in the cases, the coaxial cable/CMTS can support more broadband capacity than a D3.1 CM or D4.0 CM can use, which can be used for load balancing and increasing service tiers. For purposes of this paper, the channel conditions are assumed to be good enough to get to the better/top modulation orders. All the capacity numbers are calculated as described in our previous paper [D3.1 Capacity]. Numbers will need to be adjusted if the actual channel conditions are lower than expected in certain deployments.

2. D3.1 Technology Options

This section talks about the various configurations which in operator can implement when deploying DOCSIS 3.1 technology and the capacities for each setting. An operator typically has one of the following types of networks in which they can deploy D3.1 technology. The first (and most common in the past), is a low split network with the upstream ranging from 5 to 42 MHz and the downstream ranging up to 750 or 860 MHz. The next is a mid-split network with the upstream ranging from 5 to 85 MHz and the downstream ranging up from 108MHz up to 1 GHz. The third is a high split network with the upstream ranging from 5 to 204 megahertz and the downstream ranging from 258Mhz to 1.2 GHz. Some operators may have networks with different combinations of the downstream with the upstream depending on when those plant upgrades were done.

2.1. Legacy Low Split with 750 MHz

Let's start with the legacy DOCSIS 3.0 plant with a low split for the upstream (42 MHz), and a 750 MHz plant for the downstream. The typical configuration used by MSOs has been 4 SC-QAM upstream channels and anywhere from 8 to 32 SC-QAM downstream channels. An operator would use 6.4 MHz channels upstream carriers and placing them directly adjacent to each other and using the modulation orders to 64 QAM for the upstream SC-QAM carriers.




Figure 1 – Legacy D3.0 Plant

With four 6.4-MHz-wide channels operating at 64 QAM, the aggregate throughput of the CMs is about **106 Mbps**. With 32 x 6 MHz-wide channels operating at 256 QAM, the aggregate throughput of this CM on the downstream is about **1.2 Gbps**.

2.1.1. Legacy Low split with 750 MHz, Adding OFDM, Adding 5th SC-QAM

Here we stay on to the legacy low split plant with the upstream still at 42 MHz, and a 750 MHz plant for the downstream. The downstream has the 32 SC-QAM downstream channels and with the DOCSIS 3.1 technology being introduced, we now have a new additional 96 MHz OFDM channel deployed, typically at the higher band edge of the downstream spectrum.

As a final step for the upstream in a sub-split network, operators could try adding one or more additional upstream SC-QAM carriers, up high or down low, as shown in Figure. The paper [Bandwidth Growth] draws the conclusion that even adding 10% additional upstream capacity can help alleviate upstream congestion. Operators have been successful in adding new carriers, which is a testament to maintaining the plant more diligently over the last decade. So, in many operator networks, the upstream configuration now adds a 5th SC-QAM channel.

The number of homes passed has decreased, which lowers the effect of noise funneling at low frequencies, and cascades have shortened, which lessens the impact of group delay close to the duplex filter cutoff frequency. Operators have been successfully running narrow carriers (typically with lower order modulation) both down to 10 MHz and closer to the duplex filter.

With 5 x 6.4-MHz-wide channels operating at 64 QAM, the aggregate upstream throughput of the CMs is about **133 Mbps**. With 32 x 6 MHz-wide channels operating at 256 QAM, and an additional 96 MHz OFDM channel (at 4096 QAM) the aggregate throughput of this CM on the downstream is ~ **2.1 Gbps**.




Figure 2 – D3.1 Deployment on a Low Split Plant

2.1.2. Legacy Low Split with 750 MHz, with OFDMA

Here we stay on the legacy low split plant with the upstream still at 42 MHz, and a 750 MHz plant for the downstream. For the upstream configuration, an operator could bring the number of SC-QAM upstream channels down to 1 and add a 25.6 MHz orthogonal frequency-division multiple access (OFDMA) channel. While this may not be a realistic scenario, due to the number of D3.0 CMs that may need to be supported, this is an interesting thought exercise on the capacity increase on the upstream for a low split plant. The downstream stays at 32 SC-QAM downstream channels and the 96 OFDM channel.

With a 25.6 MHz OFDMA channel operating at the max of 2048 QAM, along with the 1 SC-QAM channel, the aggregate upstream throughput of the CMs is about **267 Mbps**. With 32 x 6 MHz-wide channels operating at 256 QAM, and an additional 96 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **2.1 Gbps**.




Figure 3 – D3.1 Deployment on a Low Split Plant w OFDMA

2.1.3. D3.1 Low Split Upstream and 860Hz Downstream

If the sub-split deployment cannot be changed and there are enough DOCSIS 3.1 modems on the network, operators can consider replacing SC-QAM channels with an OFDMA channel. OFDMA technology makes better use of spectrum because it can operate nominally at 1024 QAM, whereas an upstream SC-QAM channel is limited to 64 QAM.

In this scenario, the legacy low split plant stays with the upstream at 42 MHz, and an 860MHz plant for the downstream. For the upstream configuration, an operator could bring the number of SC-QAM upstream channels to 3 and add a 12.8 MHz OFDMA channel. The downstream has 32 SC-QAM downstream channels and two 192 MHz OFDM channel. (Leaving about 226 MHz for about 37 CTA/video channels)

With a 12.8 MHz OFDMA channel operating at the max of 2048 QAM, along with the 3 SC-QAM channels, the aggregate upstream throughput of the CMs is about **200 Mbps**. With 32 SC-QAM channels operating at 256 QAM, and an additional 2 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **5 Gbps**.




Figure 4 – D3.1 Deployment on a Low Split Plant and 860 MHz

When there is a higher percentage of DOCSIS 3.0 modems on the network, as those modems cannot use the DOCSIS 3.1 spectrum, it may make sense to replace the second OFDM channel with 32 SC-QAM channels. When there is a higher percentage of DOCSIS 3.1 modems on the network, the 2 OFDM channels will provide more OFDM capacity for those modems to use. Because this case has more spectrum allocated to DOCSIS 3.1 technology, the aggregate throughput is 5 Gbps.

2.2. D3.1 Mid-Split

A mid-split HFC network has a return path of up to 85 MHz, or two times the spectrum of a sub-split network. With a mid-split, the forward path begins around 108 MHz; therefore, set top box (STB) based video services should be able to be maintained because the forward data channel can be up to 130 MHz. The additional upstream spectrum provided by a mid-split network can provide about 500 Mbps of capacity. Per the [PHYv4.0] specifications upstream SNR must increase from 22 dB for 64 QAM to 35 dB for 1024 QAM. An upgrade to mid-split offers the MSO an opportunity to access cleaner upstream spectrum to achieve this improvement. The SC-QAM channels depicted here may need to operate below 64 QAM if the mid-split upgrade does not also improve upstream SNR at the lowest frequencies. The downstream extends to a 1GHz plant but allows the same number of channels tunable by the CM as in the above scenarios. With 32 SC-QAM channels operating at 256 QAM, and an additional 2 x 192 MHz OFDM channel the aggregate throughput of these CMs on the downstream is about **5 Gbps**.

2.2.1. Start with SC-QAMs

In terms of additional upstream spectrum, a mid-split can be configured in several ways. As shown in the figure, a mid-split can fit 10 traditional upstream SC-QAM carriers of 6.4-MHz width (for a total of 64 MHz of upstream spectrum allocated to broadband).




Figure 5 – D3.1 Deployment on a Mid-Split Plant

This configuration can yield up to **265 Mbps** of aggregate upstream capacity. Note that DOCSIS technology allows SC-QAM to be modulated up to 256 QAM and no higher. That is, with SC-QAM, 512 QAM and 1024 QAM are not available. However, the newer DOCSIS 3.1 technology does allow higher order QAM modulation.

2.2.2. Add an OFDMA Channel

Figure shows another configuration of the traditional 4 upstream SC-QAM channels along with a 49.5 MHz OFDMA running at 2048 QAM up to 85 MHz.




Figure 6 – D3.1 Deployment on a Mid-Split Plant with SC-QAM and OFDMA Channels

This configuration can yield up to **572 Mbps** of aggregate upstream capacity and an increase of ~300 Mbps using the same spectrum because the OFDMA carrier can operate at a higher order of QAM modulation than single-carrier QAMs.

2.2.3. Replace with OFDMA

The next step for an operator could be to use 75 MHz of OFDMA, which can offer **708 Mbps** of upstream capacity, an increase of ~443 Mbps using the same spectrum. It may be necessary, however, to retain a single SC-QAM upstream channel for DOCSIS 3.0 modems.




Figure 7 – D3.1 Deployment on a Mid-Split Plant with OFDMA Channel

2.3. D3.1 High Split Deployments

A high-split HFC network has a return path of up to 204 MHz, or more than 4 times the spectrum of a sub-split network. With a high-split, the forward path begins around 258 MHz, so STB-based video services cannot be maintained because the forward data channel can only be moved as high as 130 MHz, as described in the SCTE-55 standards. However, the additional upstream spectrum provided by a high-split network can provide more than 1.8 Gbps of capacity.

An interesting observation is that a high-split 1.2-GHz system supports about the same amount of downstream spectrum as a sub-split 1-GHz system (12 MHz more downstream spectrum). The benefit of the high-split system is the additional upstream broadband capacity, though this additional 12 MHz of downstream spectrum can be used in a downstream OFDM channel. As a result, the approximate downstream throughputs for the 1.2-GHz system are only slightly higher than the sub-split 1-GHz system.

Though this system supports up to 1.2 GHz, many older CPE only support a lower top frequency. For example, early DOCSIS 3.0 modems were required to support only up to 870 MHz, though later DOCSIS 3.0 modems supported up to 1002 MHz. Similarly, different models of video CPE will also have different top ends. These constraints on CPE should be considered while laying out services on the coaxial cable.

2.3.1. SC-QAM and OFDMA

In the High-split plant, the upstream spectrum is up to 204 MHz, and a 1218 MHz plant for the downstream. The downstream configuration still remains about the same and can have 32 SC-QAM downstream channels and 2 x 192 MHz OFDM channels.

In the upstream spectrum, an operator can retain the 4 SC-QAM channels for DOCSIS 3.0 (and earlier) CMs and then expand the first OFDMA channel up to 108 MHz. Now the operator can add a second

OFDMA channel to the mix. This second OFDMA channel will start from 108 MHz and span up to 204 MHz.

This configuration can yield up to **1764 Mbps** of aggregate upstream capacity. With 32 SC-QAM channels operating at 256 QAM, and the 2 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **5 Gbps**.




Figure 8 – D3.1 Deployment on a High-Split Plant with SCQAM & OFDMA Channels

2.3.2. Full OFDMA

In the scenario where there are no more D3.0 CMs in the plant, an operator can remove the 4 SC-QAM channels for DOCSIS 3.0 (and earlier) CMs and then expand the OFDMA channels to span from 12 MHz to 204 MHz.




Figure 9 – D3.1 Deployment on a High-Split Plant with Full OFDMA Channels

This configuration can yield up to **1988 Mbps** of aggregate upstream capacity. With 32 SC-QAM channels operating at 256 QAM, and the 2 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream remains at about **5 Gbps**. If the Operator can free up more downstream bandwidth, the CMTS can source more downstream channels, but a D3.1 CM can only use the 32 SC-QAM and 2 OFDM channels worth of bandwidth, but these additional channels will allow an operator to load balance the different CMs on the network better across the available DOCSIS channels.

3. DOCSIS 4.0 ESD/FDX Technology Options

The DOCSIS 4.0 specifications builds upon the previous generations of DOCSIS specifications. It includes backward compatibility for the existing physical (PHY) layers in order to enable a seamless migration to the new technology. The DOCSIS 4.0 specifications introduces Full Duplex (FDX) DOCSIS PHY layer technology as an expansion of the OFDM/A PHY layer introduced in the DOCSIS 3.1 specification to increase upstream capacity without significant loss of downstream capacity within the same available spectrum. The DOCSIS 4.0 specification also builds upon DOCSIS 3.1 OFDM and OFDMA technology with an extended Frequency Division Duplex (FDD)

DOCSIS alternative. DOCSIS 4.0 FDD supports legacy high-split and also provides extended upstream splits up to 684 MHz in an operational band plan which is referred to as Ultra-high Split (UHS). DOCSIS 4.0 FDD also introduces expansion of usable downstream spectrum up to 1794 MHz (Extended spectrum DOCSIS). Both the FDX and FDD DOCSIS 4.0 alternatives based on the OFDM/A PHY carry the cable networks into the 10G space.

3.1. D4.0 ESD/FDD Technology Options

The 1.8-GHz systems are intended for DOCSIS 4.0 equipment where the CMTS is specifically designed to place more broadband spectrum on the coax, both upstream and downstream. However, these systems need to be backward compatible with DOCSIS 3.1, DOCSIS 3.0, and perhaps DOCSIS 2.0 modems. The DOCSIS 4.0 specification builds upon DOCSIS 3.1 OFDM and OFDMA technology with an extended frequency division duplex (FDD) DOCSIS alternative. DOCSIS 4.0 FDD supports both mid-split and high-split and provides extended upstream splits up to 300/396/492/684 MHz in an operational band plan referred to as ultra-high split (UHS).

DOCSIS 4.0 FDD also introduces the expansion of usable downstream spectrum up to 1794 MHz to support higher upstream splits. An FDD CM supports at least two or more of the following upstream upper band edges: 204 MHz; 300 MHz; 396 MHz; 492 MHz; and/or 684 MHz. The FDD CM supports a minimum of 5 independently configurable OFDM channels each occupying a spectrum of up to 192 MHz in the downstream, while the FDD node/CMTS supports 6 OFDM channels

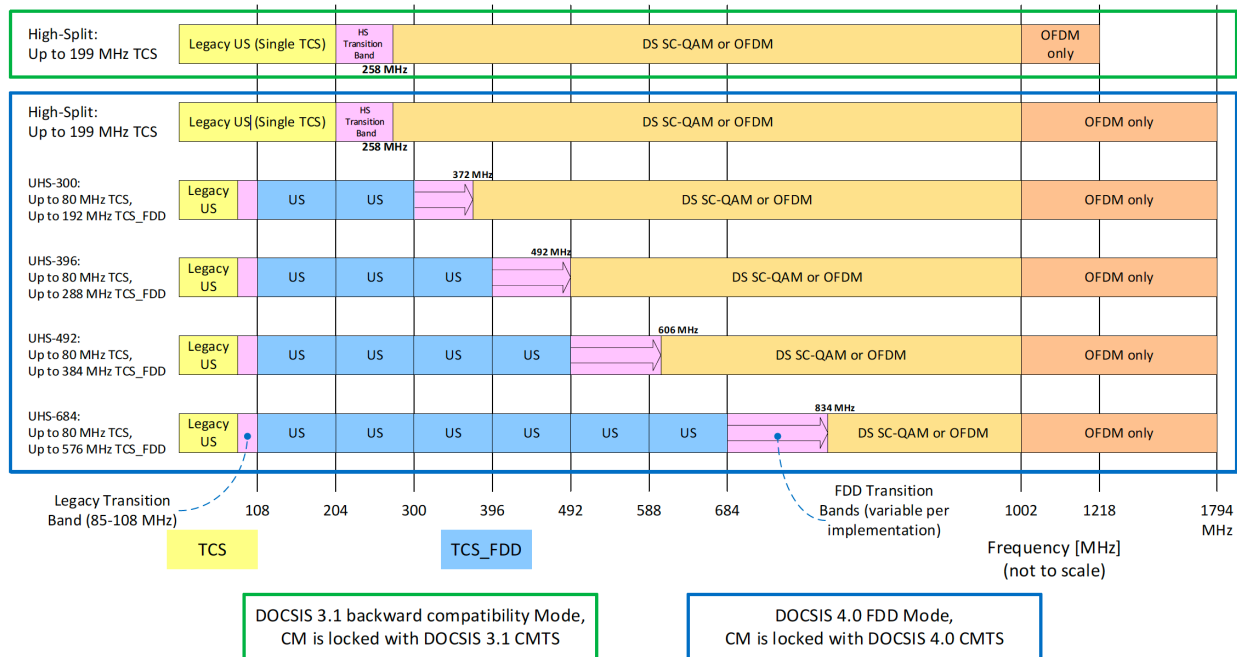


Figure 10 – D4.0 FDD Spectrum Options

Figure Source: [PHYv4.0] spec

3.1.1. D3.1 High Split with D4.0 CM

In the High split plant, the upstream spectrum is up until at 204 MHz, and the downstream up to 1218 Mhz. The downstream configuration can have 32-SC-QAM downstream channels and since the D4.0

CMs can support it, up to 4 x 192 MHz OFDM channels. In the upstream spectrum, an operator can retain the legacy 4 SC-QAM upstream channels and can have 2 OFDMA channels (70 and 96 MHz each)

This configuration can yield up to **1764 Mbps** of aggregate upstream capacity. The 32 SC-QAMs and the 4 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **8.8 Gbps**.




Figure 11 – D4.0 FDD CM on a High Split Network

3.1.2. D4.0 FDD 300 MHz Split

In the UHS-300 plant the upstream spectrum is up until at 300 MHz, and downstream extends up to 1794 Mhz. The downstream configuration can have 32 SC-QAM downstream channels and now 5 x 192 MHz OFDM channels. In the upstream spectrum, an operator can simply have 2 OFDMA channels from 108 to 300 MHz and one 75MHz OFDMA below 85 MHz. At this scenario and onwards, we are assuming the SC-QAM Channels are no longer needed, if they are the upstream numbers will need to be adjusted slightly lower, as seen in previous examples.

This configuration can yield up to **2696 Mbps** of aggregate upstream capacity. With 32 SC-QAM channels operating at 256 QAM, and the 5 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **10.7 Gbps**.




Figure 12 – D4.0 FDD CM on UHS-300 MHz Split Plant

3.1.3. D4.0 FDD 396 MHz Split

In the UHS-396 plant, the upstream spectrum is up until at 396 MHz, and downstream starts at 492 and extends up to 1794 Mhz. This makes a total of 1302 MHz of downstream spectrum, and the figure shows 1152 MHz for downstream DOCSIS channels. The downstream configuration can have 32 SC-QAM downstream channels and 5 x 192 MHz OFDM channels. In the upstream spectrum, an operator can simply have 3 OFDMA channels from 108 to 396 MHz and one 75MHz OFDMA below 85 MHz.

This configuration can yield up to **3690 Mbps** of aggregate upstream capacity. With 32 SC-QAM channels operating at 256 QAM, and the 5 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **10.7 Gbps**.




Figure 13 – D4.0 FDD CM on UHS-396 MHz Split Plant

3.1.4. D4.0 FDD 492 MHz Split

In the UHS-492 plant the upstream spectrum is up until at 492 MHz, and downstream starts at 606 and extends up to 1794 Mhz. This makes a total of 1188 MHz of downstream spectrum, and the figure shows 1152 MHz for downstream DOCSIS channels. The downstream configuration can have 32 SC-QAM downstream channels and 5 x 192 MHz OFDM channels. In the upstream spectrum, an operator can simply have 4 OFDMA channels from 108 to 492 MHz and one 75MHz OFDMA below 85 MHz.

This configuration can yield up to **4684 Mbps** of aggregate upstream capacity. With 32 SC-QAM channels and the 5 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **10.7 Gbps**.

This is an interesting case; in case an operator decides to keep video services and D3.0 CMs. With the downstream starting at 606 MHz, and if an operator decides to keep some amount of video spectrum, the DOCSIS 3.0 spectrum needs to be below 870 MHz (The top end for some early DOCSIS 3.0 modems). Also, no DOCSIS 3.0 modem is specified to use spectrum above 1002 MHz; hence, the DOCSIS 3.0 spectrum ends at 1002 MHz. Further, because video CPE possibly cannot tune above 1002 MHz either, the spectrum below 1002 MHz would need to be allocated such that both the video CPE and the DOCSIS 3.0 and earlier modems can best utilize the downstream according to service plans.




Figure 14 – D4.0 FDD CM on UHS-492 MHz Plant

3.1.5. D4.0 FDD 684 MHz Split

In the UHS-684 plant the upstream spectrum is up until at 684 MHz, and downstream starts at 834 MHz and extends up to 1794 Mhz. This makes a total of 960 MHz of downstream spectrum The downstream configuration can simply have 5 x 192 MHz OFDM channels. In the upstream spectrum, an operator can simply have 6 OFDMA channels from 108 to 684 MHz and one 75MHz OFDMA below 85 MHz.

This configuration can yield up to **6672 Mbps** of aggregate upstream capacity. the 5 x 192 MHz OFDM channel the aggregate throughput of this CM on the downstream is about **9.5 Gbps**.




Figure 15 – D4.0 FDD CM on UHS-684 MHz Plant

3.2. D4.0 FDX Technology Options

DOCSIS 4.0 FDX functionality significantly increases upstream capacity by using the spectrum currently used for downstream transmission for simultaneous upstream and downstream communications via full duplex communications. The FDX Allocated Spectrum is subdivided into FDX sub-bands each with a single FDX Downstream Channel and the associated FDX Upstream Channel(s) that can be assigned to modems according to system requirements. In the full FDX spectrum, the 3 resource blocks (each 192 MHz wide) will fit 6 FDX OFDMA upstream channels and 3 FDX OFDM downstream channels. The FDX node will simultaneously transmit and receive data on these channels with the Echo cancellation technology while the FDX CM either receives or transmits within a channel in a given sub-band. The CMTS assignment of FDX channels within the FDX band for Full Duplex DOCSIS operation can be done incrementally over time as a transition strategy, from existing DOCSIS networks to Full Duplex DOCSIS networks, as FDX-capable CMTSs and modems become available. The figure below shows the FDX allocated spectrum within the cable plant.




Figure 16 – FDX Allocated Spectrum

3.2.1. D4.0 FDX 1.0 GHz with SCQAMs or OFDM

In the initial FDX scenario, in a FDX plant the upstream-only spectrum is up until 85 MHz, and the downstream-only spectrum is from 804 MHz to 1002 MHz plant, and the FDX-Spectrum (simultaneous Upstream and Downstream) is from 108 MHz to 684 MHz. Staying at 1GHz downstream plant is a realistic initial scenario as it helps with not replacing all the taps at the start to 1.2 GHz.

In the spectrum between 684 and 804 is the as this is in the FDX transition band, to prevent any interference between the FDX bands and the downstream only region. An operator could choose to deploy video QAM channels here, and these channels can be used by other devices (STBs), but this spectrum cannot be used by the FDX CM itself

As regards a FDX CM's data capacity, in this type of a FDX plant, an operator can choose one of two options: to retain in the downstream-only spectrum (804 to 1002 MHz region) the 32 SC-QAMs or deploy a single 192 MHz OFDM channel and a single SC-QAM channel. In the upstream-only spectrum (5-85MHz), an operator can retain the 4 SC-QAM channels for DOCSIS 3.0 (and earlier) CMs or and cover the rest of the spectrum below 85 MHz with an OFDMA channel.




Figure 17 – D4.0 FDX (1.0 GHz) Channel Allocations with SCQAMs

In the FDX spectrum, the 3 resource blocks (each 192 MHz wide) will fit 6 FDX OFDMA upstream channels and 3 FDX OFDM downstream channels.

This configuration can yield up to **6537 Mbps** of peak aggregate upstream capacity. The peak aggregate throughput of this CM on the downstream, is **6937 Mbps**, if the 32 SCQAM option is chosen and **7666 Mbps**, if the 1 OFDM + 1 SC-QAM channel option is chosen.

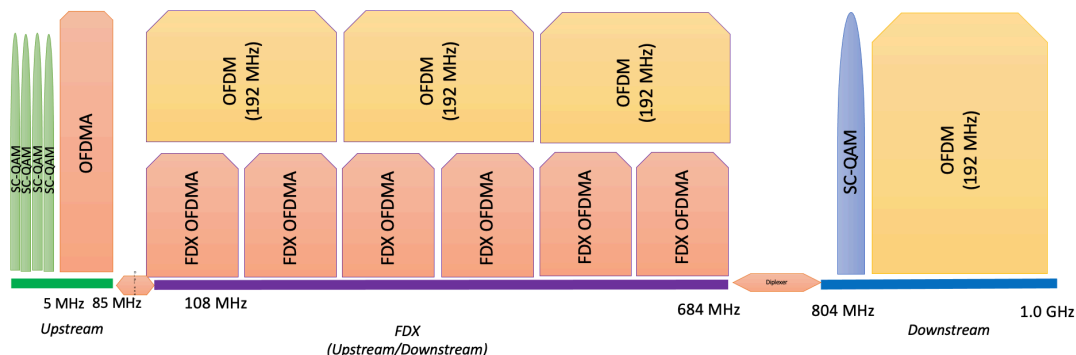


Figure 18 – D4.0 FDX (1.0 GHz) Channel Allocations with an OFDM

3.2.2. D4.0 FDX 1.2 GHz with SCQAMs

In the initial FDX scenario, in a FDX plant the upstream-only spectrum is up until 85 MHz, and the downstream-only spectrum is from 804 MHz to 1218 MHz plant, and the FDX-Spectrum (simultaneous Upstream and Downstream) is from 108 MHz to 684 MHz.

In this type of a FDX plant, an operator can choose to retain in the downstream-only spectrum (804 to 1218 MHz region), the 32 SC-QAMs and then deploy a 192 MHz OFDM channel and a second smaller 30 MHz OFDM channel. In the upstream-only spectrum (5-85MHz), an operator can retain the 4 SC-QAM channels for DOCSIS 3.0 (and earlier) CMs or and cover the rest of the spectrum below 85 MHz with an OFDMA channel.

The FDX spectrum, will fit 6 FDX OFDMA upstream channels and 3 FDX OFDM downstream channels.

This configuration can yield up to **6537 Mbps** of peak aggregate upstream capacity and **9141 Mbps** of the peak aggregate throughput of this CM on the downstream.




Figure 19 – D4.0 FDX 1.2 GHz Channel Allocations with SCQAMs and OFDM

3.2.1. D4.0 FDX 1.2 GHz with OFDMA and 2 OFDMs

Building on the FDX plant as described above, now the operator can remove the 4 SC-QAM channels for DOCSIS 3.0 (or earlier) CMs and expand the OFDMA channel to cover the spectrum below 85 MHz. Similarly for the downstream an operator could reduce the 32 SC-QAM channels to 4 and add an OFDM channel for more efficiency, for a total of 2 OFDM channels. The FDX band remains the same with 6 x 96MHz OFDMA upstream channels and 3 x 192 MHz OFDM downstream channels

This configuration can yield up to **6672 Mbps** of peak aggregate upstream capacity. The peak aggregate throughput of this CM on the downstream is about **9687 Mbps**.

One additional use case to consider would be FDX in a N+x configuration. While spectrum plans will likely remain quite similar, if we assume a lower RxMER (say 3-4 dB) or one modulation order lower for the FDX channels, the peak capacity will adjust lower to 6144 Mbps on the upstream and 9144 Mbps. Depending on the performance from the FDX amplifiers the modulation orders possible on the OFDM/A channels will change and the capacity calculation will also vary.

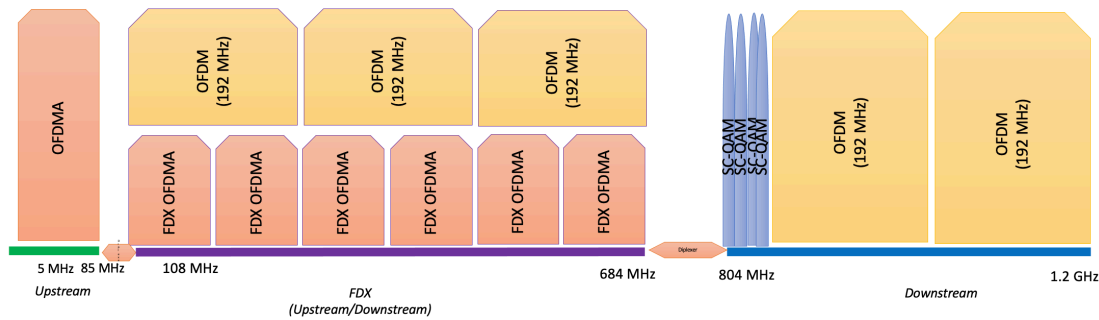


Figure 20 – D4.0 FDX Channel Allocations with OFDM/A and 2 OFDMs

3.3. Comparison to FTTP/PON Technology Options

Most FTTH networks are based on passive optical network architectures, simply as it is usually the lowest cost way to design a FTTH network. Many operators are deploying passive optical networks (PON) in many of their market segments

Per *[Broadband Pie]*, one operator's investigations and analysis has concluded that the overall cost of deploying FTTP to their residential footprint would be about 5 to 6 times the cost of rolling out DOCSIS 4.0 with DAA.

3.3.1. 10 G EPON

10G Ethernet PON (EPON) is the current and widely available generation of EPON technology, although the Nx25G EPON specifications have also recently been published. This technology is defined by Institute of Electrical and Electronics Engineers (IEEE) 802.3 standard and offers asymmetrical (10/1Gbps) and symmetrical (10/10 Gbps) rates for downstream and upstream.

The peak user throughput to an optical network unit (ONU) is around 8.6 Gbps upstream and 8.65 Gbps downstream

3.3.1. XGS PON

XGS-PON the current and widely available generation of gigabit PON (GPON) technology, although the 50G higher speed PON (G.HSP) specifications are now available. XGSPON is defined by the ITU-T G.9807.1. standard and offers asymmetrical (10/1Gbps) and symmetrical (10/10 Gbps) rates for downstream and upstream.

The peak user throughput to an ONU is around 8.55 Gbps upstream and 8.6 Gbps downstream

4. Conclusion

The table below offers a summary of the scenarios that we reviewed in this paper. It documents the number of downstream/upstream channels and the theoretical (best conditions) peak capacity that would be obtained in each of those configurations. (Do keep in mind that changing the underlying assumptions e.g., channel conditions, number of channels, spectrum allocated will change the numbers appropriately.)

Table 2 – Summary of Peak Speeds on the Path to 10G

Plant Config	Technology	DS Channels			US Channels				Theoretical Peak Capacity	
		SC QAM DS	DS-192-OFDM	FDX DS-192	SCQAM US	OFDMA-US-75	OFDMA-US-96	FDX US-96	Max DS Mbps	Max US Mbps
42 split / 750 MHz	D3.0	32			4				1216	106
42 split/ 860 MHz	D3.1	32	0.5		5	0	0		2170	133
42 split/ 860 MHz	D3.1	32	0.5		1	0.34	0		2170	267
42 split/ 860 MHz	D3.1	32	2		3	0.17	0		5030	200
85 split/ 1.2 GHz	D3.1	32	2		10	0	0		5030	265
85 split/ 1.2 GHz	D3.1	32	2		4	0.66	0		5030	572
85 split/ 1.2 GHz	D3.1	32	2		0	1	0		5030	708
204 split/ 1.2 GHz	D3.1	32	2		4	0.938	1		5030	1764
204 split/ 1.2 GHz	D3.1	32	2		0	0	2		5030	1988
204 split/ 1.2 GHz	D4.0 ESD	32	4		4	0.938	1		8844	1764
300 split/ 1.8 GHz	D4.0 ESD	32	5			1	2		10751	2696
396 split/ 1.8 GHz	D4.0 ESD	32	5			1	3		10751	3690
492 split/ 1.8 GHz	D4.0 ESD	32	5			1	4		10751	4684
684 split/ 1.8 GHz	D4.0 ESD	0	5			1	6		9535	6672
FDX/1.0 GHz	D4.0 FDX	32	0	3	4	0.66		6	6937	6537
FDX/1.0 GHz	D4.0 FDX	1	1	3	4	0.66		6	7666	6537
FDX/1.2 GHz	D4.0 FDX	32	1.15	3	4	0.66		6	9141	6537
FDX/1.2 GHz	D4.0 FDX	4	2	3	0	1		6	9687	6672
PON	10GEPON								8700	8650
PON	XGSPON								8660	8600

The figures below summarize the information in the table above in a graph, to get a feel for the differences in technology. The first figure shows the various downstream and upstream peak capacities, while the second figure shows the amount of actual usable spectrum (upstream versus downstream versus FDX) at the CM based on the number of channels in each of these configurations.




Figure 21 – DOCSIS CM DS/US capacity across technologies




Figure 22 – DOCSIS CM DS/US Usable spectrum usage

As seen in the table/figures above and in the example scenarios in this paper cable operators have many useful and cost-effective choices to increase the service speeds to their customers. Both DOCSIS 3.1 and DOCSIS 4.0 technologies give the operators plenty of options in both the upstream and downstream network design. The speeds provided by these technologies are more than capable to meet the competitive pressures in each of their respective markets. D3.1 can provide up to 5 Gbps downstream and almost 2 Gbps upstream peak speeds, ultimately enabling high service tier speeds to the customers. D4.0 FDD and FDX technologies take those peak downstream speeds to the 10 Gbps level and 6.6 Gbps in the upstream. This paper walks through the various channel planning scenarios that operator would need to go through to figure out which of these options they choose in each of their markets for scale, reliability, and cost for the benefit of speeds offered.

Abbreviations

CM	cable modem
CMTS	cable modem termination system
CTA	consumer technology association
DAA	distributed access architecture
DOCSIS	data over cable service interface specifications
EPON	ethernet PON
FDD	frequency division duplex
FDX	full duplex DOCSIS
GHz	gigahertz
HFC	hybrid fiber coax
IEEE	Institute of Electrical and Electronics Engineers
IPTV	internet protocol television
Mbps	megabits per second
MHz	megahertz
MSO	multiple system operator
NTSC	National Television System Committee
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
ONU	optical network unit
PHY	physical
PON	passive optical network
QAM	quadrature amplitude modulation
SC-QAM	single carrier QAM
SCTE	Society of Cable Telecommunications Engineers
STB	set top box
TV	television
UHS	ultra-high split

Bibliography & References

[DOCSIS MULPIv4.0] *DOCSIS 4.0, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv4.0-I05-220328*- March 28, 2022, Cable Television Laboratories, Inc.

[DOCSIS PHYv4.0] *DOCSIS 4.0, Physical Layer Specification, CM-SP-PHYv4.0-I05-220328*, March 28, 2022, Cable Television Laboratories

[D3.1 Capacity] *Accurately Estimating D3.1 Channel Capacity*, Karthik Sundaresan. SCTE 2017

[Broadband Pie] *DOCSIS 4.0 - A Key Ingredient of the 2030's Broadband Pie*, Zoran Maricevic et al. SCTE 2021

[Bandwidth Growth] *Managing the Coronavirus Bandwidth Surge: How to Cope with the Spikes and Long-term Growth*, John Ulm & Dr. Thomas Cloonan, CommScope, SCTE Expo 2020

Network Fingerprinting and Classification in Practice

An Operational Practice prepared for SCTE by

John Mansor

Vice President, Development Operations
OpenVault, LLC.

111 Town Square Place Suite 1180

Jersey City, NJ 07310

+1 (201) 677-8480

jmansor@openvault.com

Zach Simpson

Vice President, Engineering
OpenVault, LLC.

111 Town Square Place Suite 1180

Jersey City, NJ 07310

+1 (201) 677-8480

zsimpson@openvault.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Enrichment Templating	3
2.1. Augmenting Traffic Protocol Fields	3
2.2. Leveraging Additional Network Resources	4
2.3. Aligning Template Sources	5
2.4. Defining Enrichment Timing	6
2.5. Retaining Templated Sources.....	7
3. Applying Enrichment	7
3.1. Approaches to Traffic Protocol Enrichment	7
3.2. Traffic Flow Sampling and Enrichment Impact	8
3.3. Continuous Refinement and Update	9
3.4. Aggregation and Storage	10
4. Introducing AI Service Models	10
4.1. Fingerprinting Use Cases.....	11
5. Proactive Monitoring and Network Management.....	14
5.1. Monitoring, Automation and Grading	14
6. Conclusion.....	16
Abbreviations	16
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1 - Example flow record and enrichment source fields.....	4
Figure 2 - Example template source alignments.....	6
Figure 3 - Internal resource optimization via network system enrichment.....	11
Figure 4 - Capacity planning, analyzing past trends and known events.....	12
Figure 5 - Anomaly detection across selected classifications.....	13
Figure 6 - Identifying the potential for customer churn	14

List of Tables

Title	Page Number
Table 1 – Example enrichment source definitions	4
Table 2 – Template retention example	7
Table 3 – Traffic flow sampling scenario use case examples.....	9
Table 4 – Enrichment source automated refinement example	9
Table 5 – Aggregate storage options and timing example	10
Table 6 – Monitoring and automation use case examples	15
Table 7 – Deriving network resource scores from traffic intelligence	15

1. Introduction

Network fingerprinting is an emerging classification and filtering process that utilizes standard flow protocols to extract and enrich traffic records for analytics purposes. The process utilizes both public and private enrichment resources to create a modular, templated framework for use by automated machine learning (ML) and artificial intelligence (AI) systems. The goal, to create a predictive, proactive and forecast ready system for network traffic analysis, including the evolving diversity of traffic.

Through flexible templating, network fingerprinting enables a system to rapidly identify destination bottlenecks, detect anomalies within traffic flows and even recommend package adjustments. This approach has no deep packet inspection requirement and leverages flow record and packet metadata to store and enrich existing flow sources. The separation of enrichment from machine and AI techniques supports the use of homegrown solutions such as forecasting or monitoring while also allowing the use of additional open-source models for quick deployment and rapid time to value. This flexibility is designed to enable use cases across a variety of network, threat assessment and quality of service spaces and includes models to address proactive network management, self-healing actions (platform to network connections), anomaly detection, traffic monitoring, customer churn and capacity management.

The proceeding sections introduce the basic elements needed to achieve network fingerprinting and classification processes and demonstrate possible outcomes when leveraging those resources in traffic flow environments. The processes outlined focus on enrichment and augmentation, leveraging standard traffic flow protocols at a software layer without the need for specific network inspection hardware.

2. Enrichment Templating

Enrichment templating is essential for fingerprinting exercises and begins with defining the resources, business and technical requirements necessary to augment traffic flow data from a specific system or process. Defining enrichment sources in a modular and easily parsed format provides flexibility for automation, continuous updates and analytics reuse. Aligning traffic protocol fields to enrichment sources is key and ensures accurate data discovery and analysis.

2.1. Augmenting Traffic Protocol Fields

Before beginning the enrichment templating process, it is necessary to understand the traffic flow protocol and associated fields of the system being analyzed. While there are many common fields and properties between the various traffic flow exporters, identifying and understanding both the nuances and similarities of traffic protocols such as: NetFlow, Internet Protocol Flow Information Export (IPFIX), sampled flow (sFlow) etc. will greatly expedite the templating and alignment processes. Fields available to augment and enrich may vary slightly by traffic protocol so it is important to understand availability and alignment to any specific analysis use cases.

With the traffic protocol identified, assessment of the available fields and their key properties can begin. Developing valuable enrichment templates begins with understanding the consistency of fields and data within the traffic export a system will be receiving. This consistency will be critical to the application of enrichment via systematic lookups, joining of data via common fields or similar methods used by the traffic analysis system. Examples of common fields include private Internet Protocol (IP) addresses, public IP addresses, network type, network protocols, network ports and even traffic direction.

Flow Record Field Enrichment Source Fields derived from Client IP and Destination IP Addresses

client.ip	client.publicip	client.mac	client.package	client.monthusage	client.networkelement	org.name	application
10.0.0.1	192.168.0.1	00:00:00:00	GIGABIT PLUS - 1G	1.10 TB	CMTS-A01	MICROSOFT-CORP-MSN-AS-BLOCK	MS-Teams-Skype
10.0.0.2	192.168.0.1	00:00:00:00	GIGABIT PLUS - 1G	389.12 GB	CMTS-C01	MICROSOFT-CORP-MSN-AS-BLOCK	MS-Teams-Skype
10.0.0.3	192.168.0.1	00:00:00:00	750 MB PREMIER	753.24 GB	CMTS-G01	FASTLY	Disney Plus
10.0.0.4	192.168.0.1	00:00:00:00	500 MB Starter	1.34 TB	CMTS-A02	FASTLY	Disney Plus
10.0.0.5	192.168.0.1	00:00:00:00	GIGABIT PLUS - 1G	2.37 TB	CMTS-B01	FASTLY	Disney Plus

Figure 1 - Example flow record and enrichment source fields

Defining enrichment sources depends on analysis, modeling or automation requirements (see sections [4.2](#) and [5.1](#) for examples). Sources are normally qualified as either public or private. Public being sources or definitions broadly available or published openly by companies and individuals. While private sources normally represent internal network systems, purchased products or resources not publicly accessible.

The table below provides example definitions of several common enrichment sources, their classification type (public or private) and a brief description of their purpose in analysis systems.

Table 1 – Example enrichment source definitions

Enrichment Source	Type	Purpose
Autonomous Systems	Public	Align servers and destinations to known systems
Geographical IP Details	Public	Geographical data for public IP addresses
Application Identification	Private	IP address and protocol lookup for application enrichment
Cache Systems	Private	Internal network resource list for video or third-party service caching

When reviewing external enrichment sources, it is important to reference key fields and traffic flow protocol support to clearly decipher dependencies and alignment to protocol fields. Understanding how selected sources will impact and ultimately enrich the existing traffic flow data is crucial and plays a large role in defining not just data discovery but enrichment timing and performance. Identifying not just the fields but relevant aggregations required for analysis, can assist in preparation of the system and necessary classifications.

2.2. Leveraging Additional Network Resources

Enhancing a library of enrichment sources, begins by understanding the additional network resources capable of being leveraged for broader traffic intelligence. In some environments, traffic flow data may be the sole source of collection and in that case alternative public or private sources may be used for enrichment. However, in many network environments additional resources may be available or accessible through a variety of methods including Simple Network Management Protocol (SNMP) polling, Optical Line Termination (OLT) hardware or even Cable Modem Termination Systems (CMTS) via Internet Protocol Detail Records (IPDR). Additionally, Domain Name Systems (DNS) can provide insight into both local network destinations as well as public Internet resources, providing additional depth when analyzing traffic flow destinations.

Depending on policies and local regulations, another flexible and valuable resource for enrichment is billing system data. Even when anonymized or utilized only for subscriber package analysis, it can provide powerful insights into network health and ultimately customer experience within a network.

Billing systems may also include device and network resource details specific to subscribers, often simplifying the alignment process with traffic flow data.

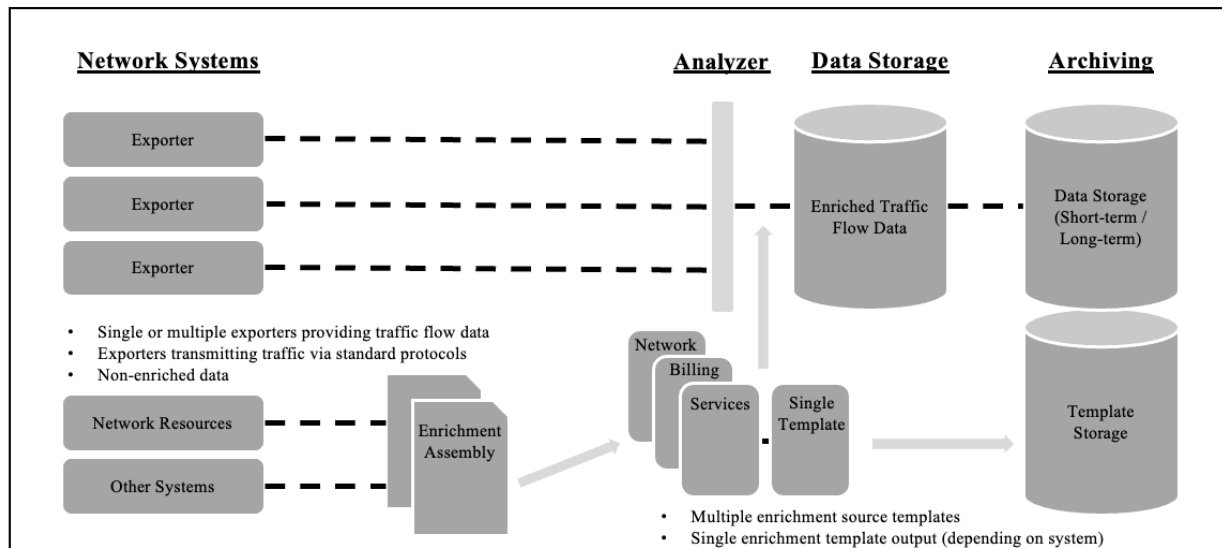
Successfully leveraging the various formats and outputs of network resources normally requires a series of crafted templates. Depending on the volume of data desired for analysis and presentation, templates can be either individually constructed or merged to a single network enrichment source. The latter normally only being preferred if continuous updating of resources is not anticipated.

2.3. Aligning Template Sources

Template alignment begins with identifying fields capable of being joined together through common keys or via the fields present within traffic records. The exact method of enrichment may be dictated by the system or application being utilized for traffic flow inspection and analysis. While many traffic analysis systems support enrichment from multiple sources, some require a single formatted file. The alignment processes described assume a multi-templated approach but with the option to join to a single file location if required.

When initially developing a solution for enrichment alignment, one commonality is IP address information. While originating IP addresses can exist in traffic records in either public or private form (depending on network settings, environment), it is normally one of the simplest and most consistent origin identifiers. When combined with subscriber enrichment sources it can provide the basis for numerous enrichment field additions. Before selecting a single origin field such as IP address, it is important to understand network behavior that can contribute to field consistency, including client/server protocol behavior and Dynamic Host Configuration Protocol (DHCP) settings within the network. Since not all enrichment sources may be leveraging IP address as a primary method of joining data, various combinations of ports, protocols and network service information can also be used. When facing more complex or non-IP address enrichment scenarios, defining unique combinations of enrichment source elements that are aligned to flow record data can assist in rapid lookups at time of application.

Aligning application specific enrichment sources within templates is sometimes a broader challenge, as the variation and combination of server or destination ports and IP addresses can require continuous enrichment from both public and private locations. This continuous enrichment is critical for the effectiveness and accuracy of application specific details (see section [3.3](#) for additional context). However, introducing the results of more complicated enrichment sources to traffic flow data remains a clearer task when adhering to described IP address or unique string combination solutions.



2.4. Defining Enrichment Timing

There a variety of traffic analysis systems and solutions available in the market and defining when to perform enrichment of traffic flow data may depend on the solution deployed. Timing selection may also depend on the use cases and type of analysis being constructed. While real-time performance is desired or required for many scenarios, such as threat analysis or capacity monitoring, it is not always necessary for point-in-time analytics or broader experience measurements. Depending on the flexibility of the solution this can lead to other alternative time considerations for analysis that may be focused on past outcomes, point-in-time comparisons or historical grading. The templating process is an ideal time to assess when enrichment should occur and whether it is necessary for enrichment updates to be applied only to newly collected (ingested) data or whether all historical datasets should be updated.

Many enrichment sources that include IP address, service or application focused details are most valuable remaining real-time. Simply meaning that at the moment enrichment sources are updated, all traffic record data from that point forward is enriched with the source details, until the source is updated again. Network fingerprinting and classification is often the most valuable when there is a clear history of events. Whether it is IP addresses dynamically changing and being associated to a new subscriber, a content delivery network (CDN) change occurring for a particular site or application and even location based internal resources such as cache systems becoming visible in the network.

This does not mean that there are no scenarios where historical collection updates should not be planned. Adjustments due to billing, subscriber, network resources or other field classification changes are often common to ensure that analysis done for historical or trending purposes, remains accurate. When assessing historical data adjustment, it is important to understand the timing and impact on the system, volume of data that is required to be queried and ultimately updated. For scenarios where an enrichment source may only be updated weekly or monthly, planning for historical reconciliation or adjustment on a set schedule will likely be the most practical option. For further timing considerations and tradeoffs when applying enrichment to traffic flow records, see section 3.1 for additional context.

2.5. Retaining Templated Sources

Enrichment sources can be a valuable commodity over time. Depending on the source, having historical references of enriched values can aid other analysis efforts, metrics or grading. Determining an optimal process for storage, access and use within additional data lake or analysis systems can depend on the enrichment source types. Leveraging long-term or cold storage options may make sense for enrichment snapshots that are continuously refreshed while more accessible storage may assist sources that will be more commonly reused. When deciding on retention strategies for enrichment templates it is often valuable to define how often each enrichment source is updated by underlying code or processes. Sources that are refreshed on a daily or more often basis may be more desirable for short-term storage options where the data can remain accessible if required.

Example retention timeframes for templated sources, including possible storage classes and conditions can be referenced in the table below.

Table 2 – Template retention example

Archive Description	Archive Timeframe	Storage Conditions	Storage Class	Details
Autonomous Systems	On Update	Past two versions	Standard / Short-term	Two versions of the complete source stored
Application Identification	Daily	7 days	Standard / Short-term	Saved daily and stored for 7 days
Billing System	On Update	Past two versions	Standard / Short-term	Two versions of the complete source stored

When designing an effective templating process, it is important to plan for the unexpected. In a real-time system, what happens if an enrichment source becomes populated with incorrect information? Having a plan to not only rollback to a previous enrichment source but also knowing the optimal method for updating historical data to a very specific point-in-time is invaluable during a production event. While the actual restoration and historical update process remains specific to the application or system deployed for traffic analysis, ensuring that a standard set of processes accompanies enrichment retention plans will ensure enrichment accuracy or adjustment is always a possibility.

3. Applying Enrichment

The art of enrichment application begins with determining the frequency for updating traffic records and ends with the aggregation and storage of the combined flow data (original record plus enrichment, if desired). Understanding the impact that sampling rates can have on specific use cases, enrichment actions and ultimately flow record storage will help ensure required modeling, analysis or presentation function as desired. The ability to successfully enrich traffic flow records accurately and at scale depends on the consistency and quality of enrichment sources and their field alignment. This makes previous planning and templating efforts key to a performant system, capable of supporting detailed network fingerprinting and classification.

3.1. Approaches to Traffic Protocol Enrichment

As previously noted, real-time traffic enrichment is not always a reality. Sometimes the volume of data or complexity of an enrichment source requires a point-in-time or scheduled update. When the time comes to apply enrichment sources to traffic flow data, each of these three: real-time, point-in-time and scheduled,

time-based scenarios is important to consider and apply within enrichment and flow record data collection.

Real-time is normally representative of enrichment that occurs as flow records are ingested and processed by the traffic analysis system. Relevant enrichment sources are leveraged to populate supplemental traffic record fields and augment the existing traffic flow data. In a real-time enrichment scenario, historical data is not updated. Depending on collection and fps rates, the actual moment an enrichment source is updated, collected data is supplemented. Real-time enrichment allows for supplemental traffic fields to be used for immediate visualization, presentation and modeling as records are ingested by the system.

Point-in-time enrichment occurs after flow records are ingested and processed by the traffic analysis system. Relevant enrichment sources are leveraged to populate supplemental traffic record fields and augment the existing traffic flow data. In a point-in-time enrichment scenario, data updates are made to historical traffic flow records only. Supplemental fields are added or updated after collection and only a dependency on system ability, performance is needed to make the updates. Point-in-time enrichment allows for supplemental traffic fields to be used for visualization, presentation and modeling between two time periods and after records are ingested by the system.

Scheduled enrichment is similar to point-in-time and occurs only after flow records are ingested and processed by the traffic analysis system. Relevant enrichment sources are leveraged to populate supplemental traffic record fields and augment the existing traffic flow data. In a scheduled enrichment scenario, data updates are made to historical traffic flow records only but are done so on a consistent frequency in order to align or aggregate key fields. Supplemental fields are added or updated after collection and only a dependency on system ability, performance is needed to make the updates. Scheduled enrichment allows for supplemental traffic fields to be used for visualization, presentation and modeling between two time periods and after records are ingested by the system.

Understanding tradeoffs that need to be made for the performant function of the traffic analysis system will save time and may help simplify the approaches selected for specific enrichment sources. When combined with sampling and continuous refinement processes, these approaches assist in the overall accuracy and classification capabilities of the system.

3.2. Traffic Flow Sampling and Enrichment Impact

Sampling rates of exported traffic flows can have an impact on the type of enrichment being performed. The level of sampling often pertains to specific use cases or automated actions being taken. When assessing traffic flow collection and enrichment, it is important to understand the volume of devices, flows per second (fps) and even internal or external traversal of the network element being collected. Each of these factors represent considerations and influence use cases for network fingerprinting and classification exercises. For example, in threat analysis and security spaces, no to extremely low sampling rates are often preferred to ensure that malicious traffic is not missed and can be cataloged and actioned effectively. While higher sampling rates can provide adequate intelligence for prime time and capacity planning use cases. While there is a wide variation in collection and enrichment use cases for traffic analysis systems, there are a few common scenarios that can help guide enrichment planning.

Determining an optimal sampling rate is often dependent on business, network and retention requirements. The network device itself may also dictate maximum sampling levels based on protocol or vendor defined parameters. The scenarios and use case recommendations provided in the table below may not apply to all customer environments. Threat analysis with low or no sampling scenarios should be reviewed on a case-by-case basis to ensure optimal experience and data retention.

Table 3 – Traffic flow sampling scenario use case examples

Use Case	Sampling
Capture and aggregate detailed traffic patterns across device base. Monitoring of key internal cache resources for capacity management, alerting and team response.	1:512
Detailed traffic flows for client/server, protocol and application analysis. Determine traffic locations and key server destinations for capacity planning and design.	1:2048
Develop overall network insights particularly during peak/primetime hours without retaining larger quantities of traffic data. Analyze and aggregate with network usage and performance metrics.	1:10000

Sampling rate is an incredibly important consideration when designing and setting up the actual platform where traffic flow record and enrichment will take place. It cannot only determine the effectiveness of the overall solution but also influence enrichment value, modeling and analysis.

3.3. Continuous Refinement and Update

Many enrichment sources require continuous care and feeding to ensure accurate results. As part of the templating process, enrichment sources should be clearly planned and defined. Once these sources are put to work enriching traffic flow records, keeping them up to date is critical to accurate fingerprinting of network resources. For certain enrichment templates, even a short lapse in updates can create data inaccuracy which may then require an adjustment to historical data. The complexity of continuous refinement within a traffic analysis system depends largely on the deployment and the context of the enrichment source. Enrichment sources focused on specific services classification, applications, CDNs and IP address ranges normally represent the highest risk for inaccuracy given change volatility. When finalizing enrichment application activities: documenting, alerting and implementing self-healing capabilities are often seen as best practices for these process critical sources.

Examples of refinement timing and associated enrichment sources are briefly described in the table below.

Table 4 – Enrichment source automated refinement example

Enrichment Source	Type	Timing	Details
Autonomous Systems	Public	Real-time	Align servers and destinations to known systems
Application Identification	Private	Real-time	IP address and protocol lookup for application enrichment
Billing System	Private	Weekly / Historical	Billing system reference and broadband package details

Leveraging built in functions of the traffic analysis system deployed or integrating surrounding scripts or functions to perform continuous enrichment on sources with the most risk, is always a recommended practice. Introducing more complex modeling, analysis and AI to enriched traffic flow data requires a high degree of refinement and accuracy.

3.4. Aggregation and Storage

A clear aggregation and storage strategy is necessary to ensure data remains accessible and retained for critical analysis and reuse. When considering aggregation and storage options it is often best to refer to enrichment template planning processes to guide decisions on what fields may require aggregation. Storing enriched traffic flow field data is not only valuable for historical purposes but often leads to an accelerated path for training AI models. Having aggregated data sets readily available to test or pilot various service techniques, can reduce time to value for analysis and modeling.

While simply storing all data, all the time, forever seems like the most ideal method for long-term historical reference, it is not always feasible for large traffic volumes. This means that for many traffic flow record scenarios, an approach that uses time-based aggregation can be considered a valuable alternative. Selecting the most ideal aggregate will depend on the use cases and traffic analysis system being utilized (support or performance of aggregation). There are many options for field aggregation in a traffic flow system, a time-based approach helps to ensure data can still be introduced to more advanced analysis models quickly. There any number of aggregate storage options based on roll-ups by minutes, hours, days etc. Choosing the right timeframe may also depend on audit scenarios, archive requirements and simply cost.

The table below introduces several timeframes and possible storage class options based on the type of aggregate data being stored.

Table 5 – Aggregate storage options and timing example

Archive Description	Data Timeframe	Storage Timeframe	Storage Class	Details
Enriched flow records	7 days	7 days	Standard / Short-term	Raw traffic and enriched records flow records available for 1 week
Raw flow records	7 days	7 days	Standard / Short-term	Saved for re-processing or comparison purposes
Enriched flow records aggregated	1 day	120 days	Cold / Long-term	Aggregated to a 24-hour period and stored
Enriched flow records aggregated	7 days	365 days	Cold / Long-term	Aggregated to a 168-hour period and stored

4. Introducing AI Service Models

Following an introduction to some of the techniques and processes needed to enrich and augment traffic flow records, it is time to leverage the available data for more advanced modeling and use cases. Introducing AI techniques to enriched traffic flow data enables network fingerprinting and associated classifications to take shape. The advent of numerous open data models as well as commercial solutions to assist with visualization and presentation can make detailed levels of classification a reality for almost any network. The success of each model and use case depends on the volume of data collected, availability of enriched fields and overall dependencies for the analysis being performed. Diligent and consistent enrichment processes help to ensure that the data is viable and more suitable for training and modeling processes. The process of data discovery and AI model selection normally begins with identifying key fields or aggregated points of interest within the enriched traffic flow data. These key data points will be used for model classification and training development activities. Assessing and preparing

enriched traffic flow data for modeling may require exporting or saving data sets for analysis and testing prior to production implementation.

4.1. Fingerprinting Use Cases

While the use cases for fingerprinting the network is ever widening and can even be system specific, there are now a variety of paths supported by readily available market and open solutions. Building and maintaining accurate fingerprints for the network normally means having a system and templating process that is adaptable and capable of growing with the network analysis and classification needs of the environment. This section provides several use cases and attempts to provide insight into specific enrichment templates, fields and techniques utilized to achieve the results. These use cases require various forms of network or traffic enrichment to be successful and are provided as examples. Detailed analysis and response as described throughout the scenarios, may depend on a variety of business and system configurations.

Scenario 1: Internal resource optimization via network system enrichment

In this scenario, the goal is to understand how external video service traffic is traversing the network over a specified period. During this primetime window, a series of internal video cache servers is being observed along with the primary external video service server destination. The video cache servers should be inclusive of the relevant markets being observed and therefore should carry most of the traffic from client systems. However, during the observed window there is still a large quantity of traffic, preferring the external destination vs. internal caches. Leveraging real-time traffic flow enrichment, the capacity and network systems adjust network configurations, re-shaping the traffic and ensuring cache usage. This scenario leverages several enrichment templates: autonomous systems, application identification and internal network resources (video cache systems).




Figure 3 - Internal resource optimization via network system enrichment

Scenario 2: Capacity planning, analyzing past trends and known events to predict future growth

In this scenario, the goal is to understand how streaming video service traffic is traversing the network and how that traffic may take shape based on historical references. During this primetime window, streaming services traffic is being monitored and a forecast view overlaid on the typical chart to indicate expected traffic patterns beyond the current time. To accomplish this both application specific enrichment along with processed, historical traffic data is combined to provide a forecast of usage for the upcoming hours. Leveraging a multi-week baseline of the exact time, day of the week and allowing for any specific events or conditions, the data is then constructed and appended to applicable charting metrics. The accessibility of processed, historical traffic data enables the system to make continual refinements and updates increasing the accuracy of forecasted usage over time. This scenario leverages several enrichment templates: autonomous systems, application identification and processed, historical traffic data.




Figure 4 - Capacity planning, analyzing past trends and known events

Scenario 3: Anomaly detection across selected classifications

In this scenario, the goal is to identify traffic anomalies early and enable rapid response through either manual or automated intervention. The visualization represents a series of connections across sets of classified traffic. This includes traffic that is categorized as gaming, CDNs, streaming services etc. The anomaly detection process is examining the enriched traffic classifications for deviation at either the upper or lower bounds. In a real-time system, as soon as enriched traffic flow data is made available to the detection process, anomalies are identified and can be actioned. In this case, the classification is a particular CDN that has been categorized by the system and is no longer receiving connections. This scenario leverages several enrichment templates: autonomous systems and application identification.




Figure 5 - Anomaly detection across selected classifications

Scenario 4: Identifying the potential for customer churn

In this scenario, the goal is to understand how a combination of traffic data, enrichment sources and network utilization information can help determine the potential for customer churn. This scenario provides a unique visualization, aggregating the various enriched data sources to plot client level data based on numerous activity conditions. These conditions include activity during peak and non-peak times, speed test usage and even references to supplemental customer support system resources. This provides a unique visualization of potential customers with activity that may represent the potential for churn. Drilling into the individual points of the data then provide inspection at the client level including all relevant enrichment detail available. This scenario leverages several enrichment templates: autonomous systems, application identification, internal network resources and customer support system data.



- **Big Consumers:** Power-users, routinely outside of assigned package
- **Baseline Consumers:** Consistent and expected activity based on assigned package
- **Limited Consumers:** Consistently below baseline activity of assigned package
- **Churn Potential:** Measurable decrease in activity, increased speed test usage and/or increased support interaction

Figure 6 - Identifying the potential for customer churn

5. Proactive Monitoring and Network Management

Proactive response remains one of the most powerful elements of network fingerprinting and classification. Leveraging enriched traffic flow data, modeling and analysis methods to enable proactive actions in many forms, should be the goal of any traffic analysis system. Whether those actions are designed to be fully automated, with intervention or in the form of customer experience metrics, there are any number of alerting and automation scenarios that can be developed.

5.1. Monitoring, Automation and Grading

Deploying monitors and automation within the network environment requires a functional traffic system with enrichment and/or deployed analysis use cases described in section [4.2](#). With an optimized traffic analysis system and with clear understanding of any sampling limits that may exist, a wide range of monitoring and automation use cases can be developed. Many of these use cases are made ever more powerful when combined with real-time monitoring and alerting functionality. This can include

automated dispatch of notifications, reports and information via team communication platforms, internal systems and devices. The ability to proactively tune and adjust devices based on analyzed data, either through internal processes or automated services, can assist in reducing manual intervention and drive self-healing abilities in the network. Example use cases for monitoring and automation range from capacity alerts to detailed threat assessments.

The table below provides several examples and includes common internal system dependencies.

Table 6 – Monitoring and automation use case examples

Use Case	Example Enrichment Dependency
Internal network capacity utilization (caches)	Internal network resources
Outage condition detection based on traffic and usage	Usage data, billing, IPDR
Outage detection for popular services / systems	Autonomous systems
Primetime or peak hour monitors	Applications
Congestion and network bottleneck visibility	Usage data, billing, IPDR
Top client / server monitors	Autonomous systems, Applications
Application and service optimization	Autonomous systems, Applications
Threat conditions and malicious traffic detection	Threat and security systems

In addition to monitoring and automation use cases, the availability of a broad range of enrichment templates can enable a variety of customer experience measurements. With network usage, subscriber or billing data availability it becomes possible to develop detailed experience and network grading components. Designing comprehensive experience monitoring and grading requires a deep understanding of not just the network system but relevant dependencies determining experience. This can include in-home devices, node, CMTS or optical layers, interoperability dependencies and even Over the Top (OTT) service performance. Establishing a grading process first requires assembling a list of available template fields that can be used either in real-time or as part of a point-in-time analysis. In its most simplistic form, scoring can be defined as a range of values equating to either a positive or negative network experience. Deciding on the fields to be leveraged may require an understanding of geographic, market, network or capacity constraints.

While the availability of enrichment sources and overall system capability may vary, the table below provides several example fields from previously discussed enrichment sources to attempt to aggregate a network experience score at a client level. Providing an overall score at the client level enables additional roll-up or aggregate opportunities and can provide additional granularity should it be required. In this example, a simple score of either a 0 or 1 is used per grading category or column based on pre-defined rules. This grading could be expanded to more sophisticated scoring and inclusive of many other factors and columns. These pre-defined rules can include evaluation of peak utilization of network elements and resources impacting a client over any number of days (30 days in the table below).

Table 7 – Deriving network resource scores from traffic intelligence

Client	Anomaly Impact	Anomaly Score	Peak Network Utilization	Network Score	Service Interactions	Service Score	Total
10.0.0.1	2.06 hours	0	93%	0	0	1	1
10.0.0.2	0 hours	1	71%	1	0	1	3
10.0.0.3	0 hours	1	68%	1	1	0	2

6. Conclusion

In summary, this practice has focused on enabling a structured process utilizing optimized enrichment templates and flow data analysis at the software layer. Network fingerprinting is an invaluable process for managing and enriching standard traffic flow protocol data. It is an inexpensive solution to stand up and does not require expensive deep packet inspection (DPI) appliances or additional hardware within the network. Most of the data provided by network routers are available from standard flow exporters and while the output can vary slightly by vendor, it is possible to rationalize those fields to get a full view of network traffic. While enrichment processes may need to run at varying frequency depending on source and desired outcomes, the combination of traffic flow data with additional network elements can help provide end-to-end network transparency. The inclusion of fiber and data over cable service interface specification (DOCSIS) termination systems can provide additional subscriber details that further augment network management potential. These classification and enrichment techniques help with fundamental understanding of traffic patterns, subscribers and can assist in improving quality of experience (QoE), by right-sizing internal vs external network utilization. The availability of collected data enables network, capacity and analyst personnel to better leverage traffic flow data through common AI and ML models. These models enable the visibility of anomalies and potential forecasting of customer churn. Proactive network management can be the culmination of data collection, enrichment, AI modeling and associated alerting. It enables continuous anticipation of networking issues and provides an additional mechanism for improving overall customer experience. Network fingerprinting as an emerging technique allows for multiple system operators (MSO) to anticipate issues, evolve network activities and ensure customers can continue to expand their broadband consumption.

Abbreviations

AI	artificial intelligence
ASN	autonomous system number
CDN	content delivery network
CMTS	cable modem termination system
DHCP	dynamic host configuration protocol
DOCSIS	data over cable service interface specification
DPI	deep packet inspection
fps	flows per second
IP	Internet protocol
IPDR	Internet protocol detail record
IPFIX	Internet protocol flow information export
ML	machine learning
MSO	multiple system operator
OLT	optical line termination
OTT	over the top
QoE	quality of experience
RDK	reference design kit
sFlow	sampled flow
SNMP	simple network management protocol

Bibliography & References

[Netflow] <https://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-netflow/index.html>

[sFlow] https://sflow.org/sflow_version_5.txt

[IPFIX] https://en.wikipedia.org/wiki/IP_Flow_Information_Export

[RDK] <https://wiki.rdkcentral.com/display/RDK/RDK+Documentation>

[PHYv3.1] DOCSIS 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I16-190121, January 21, 2019, Cable Television Laboratories, Inc

[PNM-3.1] Primer for PNM Best Practices in HFC Networks (DOCSIS 3.1), CM-GL-PNM-3.1-V03-220118, January 18, 2022, Cable Television Laboratories, Inc

[Anomaly Detection] <https://aws.amazon.com/blogs/machine-learning/anomaly-detection-with-amazon-lookout-for-metrics/>

[Customer Churn] <https://aws.amazon.com/blogs/machine-learning/predicting-customer-churn-with-no-code-machine-learning-using-amazon-sagemaker-canvas/>

All third-party trademarks, images, references and content are property of their respective owner(s).

Network Migration to 1.8 GHz

Operational “Spectral Analysis” Measured in nano-Hertz, a 30-Year Perspective

A Technical Paper prepared for SCTE by

Zoran Maricevic

Engineering Fellow

CommScope

+1 (203) 303 6547

zoran.maricevic@commscope.com

John Ulm

Engineering Fellow

CommScope

+1 (978) 609 6028

john.ulm@commscope.com

Craig Coogan

Senior Dir, Technology & Strategy

CommScope

+1 (630) 281 3143

craig.coogan@commscope.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Broadband Traffic Engineering and BW Growth Trends Overview	5
2.1.1. Broadband Traffic Engineering	5
2.1.2. Growth Rates for Broadband Peak Period Consumption	5
2.1.3. Extending Tavg Growth Trendlines to 2052	7
2.1.4. Projecting Service Group (SG) Sizes for a 30-year window	9
2.1.5. Network Capacity Modeling results for 1.2 GHz HFC Plant	9
2.1.6. Network Capacity Modeling results for 1.8 GHz ESD Plant	11
2.1.7. Network Capacity Modeling results for 10G PON Plant	12
3. Total Cost of Ownership for Various Network Upgrade Options	14
3.1. Network Upgrade Considerations	14
3.1.1. Network Evolution Example	14
3.1.2. Network Upgrade Options to consider	15
3.1.3. Time Value of Money	15
3.2. 30-year Total Cost of Ownership (TCO) Assumptions.....	16
3.2.1. Capital Expenditures (CAPEX)	16
3.2.1. Operational Expenditures (OPEX)	19
3.3. 30-year TCO Cash Flows – in Nominal \$	20
3.4. 30-year TCO Cash Flows – in 2023 \$.....	23
3.5. A Nano-Hertz Spectral Analysis.....	26
3.5.1. Network Upgrade Comparisons – D3.1 vs. ESD vs. FTTP	29
4. MSO Perspective – Current CAPEX vs. Upgrade TCO.....	30
5. Variations / sensitivity analysis.....	33
5.1.1. HFC Sensitivity analysis	33
5.1.2. R-PON Sensitivity Analysis.....	36
5.1.1. ESD to R-PON Sensitivity Analysis	39
6. Conclusion.....	41
7. Acknowledgements	42
Abbreviations	43
Bibliography & References.....	44

List of Figures

Title	Page Number
Figure 1 – DS Tavg Growth Projections for 2022 to 2037	6
Figure 2 – Max Subs per SG for Low, Moderate & High DS Tavg growth, 1794/396 MHz.....	6
Figure 3 – Downstream Tavg Trendline Predictions, 2022 – 52.....	8
Figure 4 – Upstream Tavg Trendline Predictions, 2022 – 52	8
Figure 5 – Max # of Typical Subs per SG – 258-1218 MHz HFC DS.....	10
Figure 6 – Max # of Typical Subs per SG – 204 MHz HFC US.....	10
Figure 7 – Max # of Typical Subs per SG – 492-1794 MHz ESD DS.....	11
Figure 8 – Max # of Typical Subs per SG – 396 MHz ESD US.....	12
Figure 9 – Max # of Typical Subs per SG – 10G PON DS	13
Figure 10 – Max # of Typical Subs per SG – 10G PON US	13

Figure 11: I-CCAP HFC network with head-end, field, & CPE network elements	14
Figure 12: Percentage of PV for an infinite cash flow stream.....	16
Figure 13: Initial CAPEX (\$ per HP) for D3.1 High-split; D4.0 ESD; & 10G R-PON upgrades	17
Figure 14: Initial CAPEX (\$ per HP) for D3.1 High-split & D4.0 ESD upgrades - detailed	17
Figure 15: Head-end changes for DAA D4.0 ESD upgrade	19
Figure 16: Field actives percentage fail heuristics curve	20
Figure 17: TCO of 1.2 GHz HS upgrade, with plant CAPEX + OPEX, and CPE CAPEX	21
Figure 18: TCO of 1.8 GHz ESD upgrade, with plant CAPEX + OPEX, and CPE CAPEX.....	21
Figure 19: TCO for 10G R-PON, 20% U/G, with plant CAPEX + OPEX, & CPE CAPEX	22
Figure 20: TCO for 10G R-PON, 80% U/G, with plant CAPEX + OPEX, & CPE CAPEX	22
Figure 21: TCO for ESD to R-PON upgrade with plant CAPEX + OPEX, & CPE CAPEX.....	23
Figure 22: TCO of 1.2 GHz high-split upgrade over time, in '23 dollars	24
Figure 23: TCO of 1.8 GHz ESD upgrade over time, in '23 dollars	24
Figure 24: TCO of ESD to R-PON upgrade in '23 dollars.....	24
Figure 25: TCO of 10G R- PON upgrade, 20% U/G, in '23 dollars.....	25
Figure 26: TCO of 10G R- PON upgrade, 80% U/G, in '23 dollars.....	25
Figure 27: 1.2 GHz high-split upgrade TCO in frequency domain, in '23 dollars	26
Figure 28: 1.8 GHz ESD upgrade TCO in frequency domain, in '23 dollars	26
Figure 29: ESD to R-PON upgrade TCO in frequency domain, in '23 dollars	27
Figure 30: 10G R-PON 20% U/G upgrade TCO in frequency domain, in '23 dollars	27
Figure 31: 10G R-PON 80% U/G upgrade TCO in frequency domain, in '23 dollars	28
Figure 32: Annual CAPEX for four USA MSOs: Comcast, Charter, Altice USA, Cable One	30
Figure 33: Annual CAPEX for four USA MSOs, Normalized per HP	31
Figure 34: CAPEX by category, from Charter's 1Q22 financial results	32
Figure 35: TCO Sensitivity for 1.2 GHz High-Split upgrade – 100,000 trials Monte-Carlo run	33
Figure 36: TCO Sensitivity for 1.8 GHz ESD upgrade – 100,000 trials Monte-Carlo run.....	34
Figure 37: Sensitivity Tornado chart – 1.2 GHz high-split	35
Figure 38: Sensitivity Tornado chart – 1.8 GHz ESD	36
Figure 39: TCO Sensitivity for R-PON 20% U/G upgrade – 100K trials Monte-Carlo runs	37
Figure 40: TCO Sensitivity for R-PON 80% U/G upgrade – 100K trials Monte-Carlo runs	37
Figure 41: Sensitivity Tornado chart for the 10G R-PON 20% U/G upgrade	38
Figure 42: Sensitivity Tornado chart for the 10G R-PON 80% U/G upgrade	39
Figure 43: TCO Sensitivity for ESD to R-PON upgrade – 100K trials Monte-Carlo runs	40
Figure 44: Sensitivity Tornado chart for the ESD to R-PON upgrade	41

List of Tables

Title	Page Number
Table 1: Initial upgrade CAPEX, per node area.....	18
Table 2: Network upgrade scenarios compared, in '23 dollars.....	29

1. Introduction

Frequency is the number of cycles per unit time. Communications often use signals occupying MHz and GHz spectrum. But when it comes to the network's operational health, it is the 1 to 10 nano-Hertz frequencies that matter the most!

How so? Those frequencies correspond to 30- and 3-year cycles, respectively. And most activities affecting the network's health (construction, upgrade, and maintenance cycles) fall within this range. This paper investigates the 5 W's of network upgrades around the pending Extended Spectrum DOCSIS 4.0 (ESD) rollout and its alternatives. It looks holistically at what makes economic sense over the full 30-year cycle, not just the next incremental step.

What network components are impacted? Over 30 years, a Hybrid Fiber-Coax (HFC) plant sees taps, amps and nodes going through 1 or more upgrade cycles, while Cable Modem Termination Systems (CMTS) and consumer premises equipment (CPE) upgrade even more often. Until now, these cycles were all independent. However, Distributed Access Architecture (DAA) and ESD now lock these together. Accounting for this, applying average per cycle costs and integrating over the "spectrum" is akin to performing a financial "spectral analysis." Both capital and operating expenditures may be captured this way, giving operators a long-term total cost of ownership (TCO) of the network.

When is the right time for various upgrades? Doing Fiber to the Premises (FTTP) now spends the bulk of the upgrade budget up front when 99% of its capacity goes unused. Can operators be wiser on when to invest in the network? When will the capacity be needed? Recent broadband bandwidth trends are reviewed and show a sharp decrease in subscriber consumption compounded annual growth rates (CAGR) over recent years. As an example, this might cause the network upgrade cycle to go from every 10 years to every 30 years with resultant economic impact. Upgrading too aggressively may be throwing away dollars in the near term.

Where do multi-system operators (MSOs) touch the network? Is it a complete overlay or just a surgical strike? Swap tap faceplates or go for 3 GHz tap housings? Keep existing amp cascade or push fiber deeper? Walking through the process forces operators to quantify materials, inventories, and project execution times.

Why choose a particular technology direction? Many upgrade decisions must be made soon and made well with the long term, full cycle consideration. Should operators upgrade amps to mid or high split or ESD?

Who will benefit from this? Network operators, small and large. From the long-term, big picture view this analysis gives, they'll be able to focus on just the near-term network operation and upgrade aspects, and plan for the budgets over the next 3-5 years with ease while keeping aligned with their 10- to 30-year cycle objectives. These are the 1-10 nano-Hertz frequencies that are our major concerns.

2. Broadband Traffic Engineering and BW Growth Trends Overview

2.1.1. Broadband Traffic Engineering

The CommScope (formerly ARRIS) team has led industry traffic engineering research for over a decade. [CLO_2014] introduced broadband Quality of Experience (QoE) using a simple formula with basic network capacity components. This evolved and [ULM_2019] gave an updated insight into calculating the service group (SG) capacity requirements:

Modified “COMMScope/CLOONAN’S CAPACITY EQUATION” Traffic Eng Formula:

$$C \geq (N_{sub} * T_{avg}) + (K-1) * T_{max_max} + T_{max_max} \quad (1)$$

The subtle change is that there are now three main components to the traffic engineering formula:

1. Peak Busy Period Average Consumption (i.e., $N_{sub} * T_{avg}$)
2. Peak Busy Period Ripple for managing QoE (i.e., $(K-1) * T_{max_max}$)
3. Headroom for maximum Service Tier Burst (i.e., $1 * T_{max_max}$)

While burst and ripple components manage a subscriber’s QoE, the consumption component is key to SG sizing. The T_{avg} growth rate has seen much research. ARRIS/CommScope has the most extensive broadband capacity monitoring history in the industry, collecting continuously since 2010 from the same MSOs. The 2022 data is in and downstream (DS) T_{avg} growth continues to slow.

The real multi-billion-dollar question is what’s the consumption growth for coming decades? For this paper, it is a three-decade window being considered. This growth drives our network investment strategies. Has T_{avg} growth slowed to a lower rate or is it no longer exponential? A companion paper by these authors [ULM_2022] investigated several possible growth trendlines including exponential, linear, Adoption S-curve and others. Our research measures how accurately each trendline matches last decade’s data. These bandwidth (BW) growth trajectories in [ULM_2022] were mapped out for 5/10/15 years. The resultant spaghetti plots in Figure 1 show a cone of uncertainty that grows over time, roughly doubling every 5 yrs.

2.1.2. Growth Rates for Broadband Peak Period Consumption

To understand the impact of these slowing growth rates, consider the following comparison to projections from just four years ago:

- 2018 DS Growth (43% CAGR) projection => DS T_{avg} = 100 Mbps/sub **by 2030**
- 2022 DS High Growth (21% CAGR) projection => DS T_{avg} = 100 Mbps/sub **by 2040**
- 2022 DS Low Growth (Linear) projection => DS T_{avg} = 100 Mbps/sub **in 200+ years**

[ULM_2022] implied that *the need for FTTP to all subscribers may be pushed back multiple decades*. This paper takes an economic view of different upgrade options for each of the low, medium, and high growth rate scenarios. From a network capacity planning perspective, [ULM_2022] conclusions on multiple growth trendline options were:

- the 5-year window provides a reasonably high confidence for near-term planning
- 10-yr window provides high, moderate, and slow growth ranges for longer term planning
- the 15-yr window shows too much variance and is more of an academic exercise.

This paper considers a 30-year window, so it is clearly noted that this falls under the academic exercise scenarios. However, it is still informative to see how this wide “cone of uncertainty” might impact our network migration economics.




Figure 1 – DS Tavg Growth Projections for 2022 to 2037




Figure 2 – Max Subs per SG for Low, Moderate & High DS Tavg growth, 1794/396 MHz

[ULM_2022] used the CommScope network capacity model to investigate several upgrade case studies. The 1794/396 MHz case study seen in Figure 2 showed that a node with 150+ subs can offer multi-gigabit service tiers; but the timing of additional node splits on the ESD plant is sensitive to which DS Tavg growth trendline it tracks. It shows there is no pressing need to push the HFC to very small (but inefficient!) Node + 0 (N+0) SG sizes in a 10-year window. This paper extends that analysis to 30-years to see how various upgrade options are impacted by the different growth trendlines.

2.1.3. Extending Tavg Growth Trendlines to 2052

The goal is to minimize up front investments while maintaining flexibility to increase network capacity and manage uncertainty risks. In this paper, there are several primary cases being considered in detail:

- Low growth scenario enabling 1.2 GHz DOCSIS 3.1 (D3.1) high-split to last 30 years
- Moderate growth scenario with a 1.8 GHz DOCSIS 4.0 ESD upgrade over 30 years
- High growth scenario with a 1.8 GHz DOCSIS 4.0 ESD upgrade in '23 and FTTP overlay in '44
- High growth scenario with a FTTP upgrade in '23

The impact of the other growth trendlines will also be investigated for each upgrade path to ascertain the operator's risk with each path.

The first step in projecting SG sizes out to 2052 is to extend the Tavg growth trendlines from [ULM_2022]. The DS Tavg growth trendlines are shown in Figure 3 and the upstream (US) Tavg growth trendlines in Figure 4. Because of the divergence of the various trendlines over the 30-year window, note that the Y-axis is now a log scale. Our studies will consider low, medium, and high growth scenarios.

The DS high growth scenario follows the 21% CAGR exponential trendline for the first 15 years. By the end of 30-years, the spread between the linear trendline and the 21% CAGR eventually becomes an almost absurd amount – 17 Mbps compared to 1,100 Mbps. For this study, the authors decided to drop the two extreme trendlines (upper and lower) for the 30-year point, maintaining that the probability of staying at the extreme for 30 consecutive years would be very low. The red ovals in Figures 3 and 4 indicate the trendlines used at 15-years and 30-years with a transition period between them.

For the 30-year period, the DS high growth then transitions to a more moderate 16% CAGR exponential trendline. The DS Tavg window being considered in 2052 still ranges from 22 Mbps to 356 Mbps. Note, even if DS Tavg stayed on the 21% CAGR for 30 straight years, the only impact on our analysis is that the 2052 numbers get pulled in 5-7 years earlier which is a minor impact to the overall financial analysis.

The US high growth scenario follows a similar methodology. It follows the 23% CAGR exponential trendline for the first 15 years, then transitions to a more moderate 18.5% CAGR exponential trendline. The US Tavg window being considered in 2052 ranges from 6 Mbps to 48 Mbps.




Figure 3 – Downstream Tavg Trendline Predictions, 2022 – 52




Figure 4 – Upstream Tavg Trendline Predictions, 2022 – 52

2.1.4. Projecting Service Group (SG) Sizes for a 30-year window

The ARRIS/CommScope Traffic Engineering formula shown earlier provides guidance on calculating SG sizes. For a given network (e.g., 1218/204 MHz HFC plant), the available capacity is known. Once T_{avg} and T_{max_max} are defined, then the maximum number of subscribers, N_{sub} , can be calculated. The DS and US T_{avg} for a given year are derived from the growth trendlines in the previous section. Low, medium, and high growth scenarios are considered.

The final piece to the puzzle is determining what to use for T_{max_max} . The authors decided to set the DS T_{max_max} to 5 Gbps. The rationale being that this can handle any known application today. If an application like virtual reality (VR) or augmented reality (AR) really takes off, then that will drive a higher T_{avg} growth trendline. Having a 5 Gbps service tier burst on top of that high T_{avg} should be more than sufficient. History has shown that it often takes a decade or more before new technologies become mainstream. This can be seen in how long it took High Definition (HD) video streams to become dominant. Ultra-HD 4K video stream has been around almost a decade and still is not dominant yet.

The max US service tier is defined by the particular type of network upgrade. The 204 MHz high split plant supports a 1 Gbps US tier. The 396 MHz ultra-high ESD split supports a 2.5 Gbps US tier. The 10G Passive Optical Network (PON) supports a symmetric 5 Gbps US tier. Again, there is no mainstream application in sight that would need more than a gigabit of burst speed.

Our analysis starts with a typical HFC SG with 200 subscribers (e.g., 400 homes passed (HP) @ 50% penetration). The CommScope network capacity model is then run for low, medium, and high scenarios to determine the maximum number of typical subs that can be support in each year through 2052. When that growth line passes through the 200 sub/SG limit, then the SG needs to be segmented (e.g., 1x1 Remote MACPHY Device (RMD) upgraded to 2x2 RMD). Similarly, once a growth line passes through the 100 sub/SG limit, then the SG needs to be segmented a second time (e.g., 2x2 RMD upgraded to 4x4 RMD). In reality, the number of homes passed per radio frequency (RF) leg is often unbalanced, so a 4x4 RMD may be of limited use and a node split is required instead. As will be seen in the upcoming results, the need for 4x4 segmentation is still 15+ years away for today's 200 sub service groups.

2.1.5. Network Capacity Modeling results for 1.2 GHz HFC Plant

The network capacity modeling results for a 258-1218 MHz HFC DS is shown in Figure 5. The low growth scenario can support 200 typical subs all the way until 2049 when the SG needs a 2x2 segmentation. The moderate growth scenario supports 200 subs per SG through 2033 when it needs a 2x2 segmentation. The network then supports 100 subs per SG until 2040 when a 4x4 segmentation is required. So, an MSO has a clear path to 2052 with a 1.2 GHz HFC DS for both the low and medium growth scenarios.

The high growth scenario is much more challenging for the 1.2 GHz HFC DS. The first 2x2 segmentation comes around 2030 while the next 4x4 segmentation is needed by 2035. The 4x4 segmentation then runs out of capacity around 2040. In reality, the operator will probably switch to either 1.8 GHz or 10G PON sometime in the 2035-40 timeframe.

The network capacity modeling results for a 204 MHz HFC US is shown in Figure 6. In the low growth scenario, US capacity is depleted sooner than the previous DS example. By 2037, the US needs to be segmented. This can be done with a 1x2 RMD replacing the 1x1 RMD. For medium growth, the US timeline is like the DS, although the 2x4/4x4 US does run out of capacity around 2045. This would then require further node splits or a switch to one of the other technologies. The US high growth scenario is very close timeline to the DS timeline.




Figure 5 – Max # of Typical Subs per SG – 258-1218 MHz HFC DS




Figure 6 – Max # of Typical Subs per SG – 204 MHz HFC US

2.1.6. Network Capacity Modeling results for 1.8 GHz ESD Plant

The network capacity modeling results for a 492-1794 MHz ESD DS is shown in Figure 7. The low growth scenario can support 200 typical subs through 2052 with no segmentation required. The moderate growth scenario supports 200 subs per SG through 2039 when it needs a 2x2 segmentation. The network then supports 100 subs per SG until 2047 when a 4x4 segmentation might be needed. Again, an MSO has a clear path to 2052 with a 1.8 GHz ESD DS for both the low and medium growth scenarios.

The high growth scenario is more interesting with the 1.8 GHz ESD DS. The first 2x2 segmentation is not until 2034. The next 4x4 segmentation is needed by 2038. The 4x4 segmentation capacity then holds out until 2044. At this point, the operator has at least three potential paths:

1. Continue to pull fiber deeper and split ESD nodes into smaller SG
2. Switch to FTTP for all customers
3. Do an FTTP overlay and selectively migrate heavy users to PON

The economics of the third option above is looked at in more detail in the upcoming sections.

The network capacity modeling results for a 396 MHz ESD US is shown in Figure 8. In general, all three US growth scenarios have very similar breakpoints as the DS scenarios. The only point that is different is with the medium growth scenario. The 2x2 US runs out of capacity around 2045 where a switch to 2x4 RMD might be needed.




Figure 7 – Max # of Typical Subs per SG – 492-1794 MHz ESD DS




Figure 8 – Max # of Typical Subs per SG – 396 MHz ESD US

2.1.7. Network Capacity Modeling results for 10G PON Plant

The 10G PON modeling assumes 128 HP per Remote Optical Line Terminal (R-OLT) port with 64 subs. The network capacity modeling results for a 10G PON DS is shown in Figure 9. The low growth scenario can support 64 typical subs through 2052 with no segmentation required. The moderate growth scenario supports 64 subs per SG through 2043 when it needs to drop to 64 HP and up to 49 subs through 2052.

The high growth DS scenario pushes the 10G PON capacity, just like it did ESD. Around 2036, it needs to drop to 64 HP and 32 subs. By 2042, the 10G PON needs to segment again to get to 32 HP and smaller SG sizes. This will probably be the time where an operator needs to start transferring heavy users to one of the next generations of PON technology (e.g., 25+ Gbps).

The 10G PON has plenty of US capacity and can handle low, medium, and high growth scenarios for 128 HP and 64 subs per SG through 2052, as is shown in Figure 10.




Figure 9 – Max # of Typical Subs per SG – 10G PON DS




Figure 10 – Max # of Typical Subs per SG – 10G PON US

3. Total Cost of Ownership for Various Network Upgrade Options

3.1. Network Upgrade Considerations

3.1.1. Network Evolution Example

Over the past ~30 years, HFC networks progressed from ragtag one-way community antenna video distribution networks to modern high-capacity video, voice, and data bi-directional networks of today. Figure 11 shows a high-level view of such a network. This example falls under the centralized architecture model, where all the sophisticated communication layer processing takes place in the Head-End. Here, signals are ‘packaged,’ and then ‘shipped’ over a more or less transparent physical network. The other side of processing takes place at the CPE / cable modem (CM).

The way the network was built is best described as an evolution: first the cables were installed; with the RF amplifiers and taps placed where required, and the same with the headend and customer premise gear. Once every ~30 years, the taps get a refresh – typically via faceplate upgrade – for example, 750 MHz taps would get upgraded to 1 GHz or 1.2 GHz taps. Amplifiers get renewed every ~15 years, either due to technological obsolescence or to reaching the end of reliable operation lifetime. Fiber nodes may get a refresh even more often than every 15 years. Node splitting, for example, is still an effective way to boost total capacity and service levels, in otherwise over-subscribed service groups. Finally, the headend and customer premise equipment refresh-cycles fall under just in time schedule – perhaps every 3-5 years, allowing for capacity boosts as needed, in the most cost-efficient way.

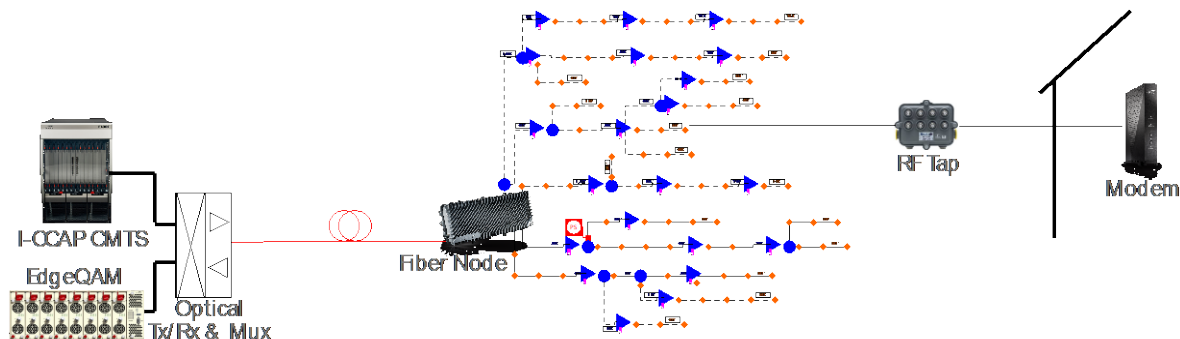


Figure 11: I-CCAP HFC network with head-end, field, & CPE network elements

The example network of Figure 11 may have 1 GHz taps installed 25 years ago, 750/870 or 1000 MHz RF amplifiers with 42/54 MHz RF split, updated 15 years ago, and perhaps 1 GHz nodes, with the same 42/54 MHz sub-split, updated 10 years ago. Headend and CPE may have been upgraded as recently as 5 years ago, with the latest wave of DOCSIS 3.1 (D3.1) deployments, correlated with introduction of 1 Gbps downstream data rates.

As learned in these COVID times, the networks delivered, and delivered marvelously, especially in the downstream. The 42 MHz upstream capacity was severely tested all day long with working and school from home. Some nodes struggled but were quickly upgraded with additional capacity. With FTTP/PON competition offering gigabit rates in the upstream - it is now imperative for MSOs to resolve the network's 42 MHz upstream capacity limitations. So, many operators are now considering their next network migration steps.

3.1.2. Network Upgrade Options to consider

CableLabs DOCSIS spec creators have envisioned these types of scenarios. Many paths exist to upgrade networks and boost its capacity in both downstream and upstream to extend the useful life of the network. Out of many options discussed in [Broadband Pie], the following upgrade scenarios of most interest are considered in this paper:

1. DOCSIS 3.1 Integrated Converged Cable Access Platform (I-CCAP) “high-split” (HS) upgrade
 - With 5-204 MHz upstream, 258 - 1,218 MHz downstream
2. DOCSIS 4.0 RMD ESD “ultra-high-split” (UHS) upgrade
 - With 5-396 MHz upstream, 492 - 1,794 MHz downstream, one of several ESD options
3. FTTP 10G R-PON (Remote OLT PON) upgrade
 - Effectively overbuilding the coaxial portion of the plant with fiber, and providing fiber drops to those homes that have signed up for the service

The FTTP upgrade provides lower operating costs (OPEX), in comparison to HFC networks [bbcmag FTTH OPEX] and [FTTH OPEX]. The more important question is: will those operating cost savings offset a much larger capital expenditure (CAPEX) required upfront to build an all-fiber network? Last year’s SCTE Cable-Tec Expo paper [Broadband Pie] considered this question, by looking into “total cost of ownership” (TCO) of nine various upgrade paths, including the three mentioned above. It used a fifteen-year period for which the total cost was calculated. This paper expands that analysis to consider a 30-year period, and in part to answer if the 15-year timeframe was too limiting.

3.1.3. Time Value of Money

To address any of these various duration questions, the “time value of money” (TVM) concept is one critical element to factor into the analysis. It is incorporated via a “discount rate” for the future years’ cash flows. For example, a discount rate at say 5% applied to \$100 received a year from now has a value today that would be ~\$95, or $1 / (1+5\%) = \$95.24$ exactly. The same applies to the expenses: one postponed by a year is ~5% less of an expense in today’s dollars. Similarly, today’s \$100 is worth \$105 a year from now.

Economists employ a dividend discount model (DDM) [Gordon] to estimate the present value of an infinite-series of future cash flows. A perpetual stream of annual \$100 expenses would be valued at \$2000, \$1333, and \$1000 in today’s dollars, with 5%, 7.5% and 10% discount rates, respectively. A difficult question is what exact discount rate to use – short and long-term interest rates, overall economy growth rate, and inflation rate are contributing factors, and the discount rate chosen affects the present value of future cash flows a lot!

Highlighted points in Figure 12 illustrate what percentages of that present value (PV) is still achieved if the perpetual flow were to cease flowing 15 or 30 years from now – also shown as a function of the discount rate. This concept, applied to TCO analysis, explains that working with a 15-year period considers only between 52% and 76% of the OPEX contributions, provided the discount rates are between 5% and 10%. Moving to 30 years, however, improves those numbers to between 77% and 94%. Given all the uncertainties involved, and that the authors will be thrilled if the overall model accuracy reaches +/-20 %, one may conclude, especially at 7.5% discount rate, to get to the 89% of the total value of OPEX is more than “good enough.” The 7.5% discount rate is chosen as a reasonable middle point between 5% (considered low) and 10% (considered high) discount rates.




Figure 12: Percentage of PV for an infinite cash flow stream

The 30-year window is also an important consideration for CAPEX too. Components such as amplifiers might need replacing every 15 years. With a 7.5% discount rate, the replacement costs for an amplifier 15-years from now adds an additional ~34% in today's dollars to the amplifier's CAPEX costs. Another replacement after 30-years adds another ~12% to the CAPEX. Looking at the finances from a spectral perspective provides a better insight into these recurring costs for some of the options.

3.2. 30-year Total Cost of Ownership (TCO) Assumptions

To get the total cost of ownership, both CAPEX and OPEX are included over a 30-year period.

3.2.1. Capital Expenditures (CAPEX)

CAPEX is determined by adding cost of materials, plus the cost of labor necessary to install the materials. Headend, field, and customer premise equipment are all accounted for. Breakout of what the initial upgrade CAPEX for the three outlined options is shown in Figure 13.

The network shown in Figure 11 is the starting point for each upgrade option. Its characteristics:

- I-CCAP topology, with 5-42 / 54-860 MHz upstream and downstream
- 21,120 feet of hardline coax plant
- 400 HP with a 50% take rate (i.e., 200 subscribers)
- One fiber node, 7 bridger & 14 line-extender RF amplifiers and 100 RF taps

Statistics for this node area work out to:

- 100 homes-passed per mile
- 5.25 RF amplifiers per mile
- 19 homes-passed per amplifier
- ~5 taps per RF amplifier and ~4 homes-passed per tap

This is very representative of a typical suburban North American HFC plant.




Figure 13: Initial CAPEX (\$ per HP) for D3.1 High-split; D4.0 ESD; & 10G R-PON upgrades




Figure 14: Initial CAPEX (\$ per HP) for D3.1 High-split & D4.0 ESD upgrades - detailed

Table 1: Initial upgrade CAPEX, per node area

	I-CCAP 1218/204 MHz High Split	DAA 1794/396 MHz UHS ESD	FTTP 10G R-PON 20% Underground	FTTP 10G R-PON 80% Underground
CMTS / Video license	\$3,250	\$800	\$800	\$800
CIN / Ethernet Switches		\$871	\$3,485	\$3,485
Head-End Optics	\$3,210	\$1,000	\$4,000	\$4,000
Node Hardware	\$4,100	\$7,000	\$11,875	\$11,875
Field Hardware	\$10,500	\$17,640	\$500	\$500
Field labor, incl node	\$6,050	\$6,800	\$13,563	\$13,563
Fiber, material + labor		\$2,112	\$84,480	\$211,200
Taps/splitters, mtrl + lbr	\$4,000	\$8,000	\$7,500	\$7,500
Drops, material + labor			\$40,000	\$40,000
CPE		\$240	\$24,000	\$24,000
Total, per SG or node	\$31,110	\$44,463	\$190,203	\$316,923
Total, per HP	\$78	\$111	\$476	\$792

**** Disclaimer: Price points discussed and shown in this document are meant to provide indicative general trends for these architectures, and as such should not be construed as an offer for selling any products at any of the price points shown. ****

Material and labor per node area are best-effort estimates, shown in Figure 14 and detailed in Table 1.

For I-CCAP high-split option, CMTS license covers additional DS and US enabled D3.1 spectrum. Furthermore, complete replacement of head-end optics, node and RF amplifiers is assumed, with digital return for the 5-204 MHz upstream spectrum. Field labor includes \$500 for the node replacement, \$250 per bridger and \$200 per line extender, plus \$1,000 for documentation update for the area. Only the tap faceplate is upgraded, at \$40 per tap.

For ESD 396/492 MHz ultra-high-split option, the head-end side of Figure 11 gets upgraded to the one shown in Figure 15 – the Cable Modem Termination System (CMTS), digital video Quadrature Amplitude Modulation (QAM) generators and analog headend optics are replaced by converged interconnect network (CIN), interfacing the Ethernet and video core with the node-located Remote PHY device (RPD) and/or RMD devices. Pluggable Dense Wavelength Division Multiplexing (DWDM) enhanced Small Form-factor Pluggable (SFP+) modules are shown at \$1,000 each, with one on each end of the digital optical link interfacing CIN and the node. 1.8 GHz Node and RF amplifier hardware are new products and assumed to be at 40% premium over those for 1.2 GHz. The 1.8 GHz ESD upgrade has some additions in comparison to the 1.2 GHz case, including a provision for 5% of aerial plant cable replacement, and an increase in the number of actives, from 7 & 14 to 10 & 14 bridgers & line extenders, respectively. Furthermore, a complete tap housing upgrade to 1.8 GHz is included at \$80 per tap, including material & labor.

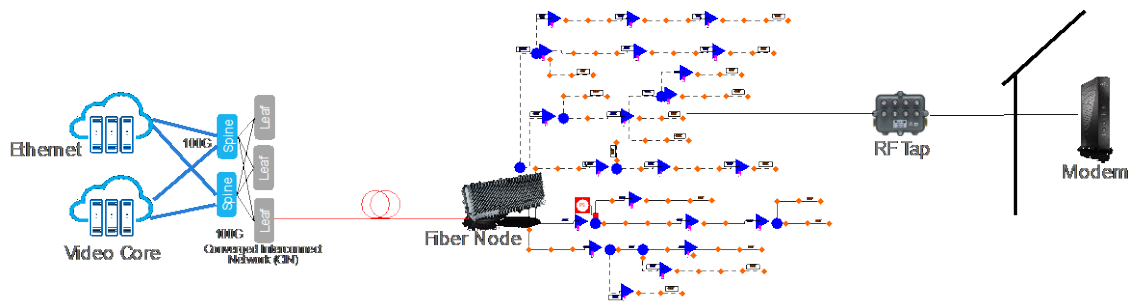


Figure 15: Head-end changes for DAA D4.0 ESD upgrade

For the 10G R-PON upgrade, Figure 15 topology serves as a blueprint, except that the complete hardline coax plant is overbuilt with fiber. The first PON option assumes an 80/20 percent mix of aerial to underground plant, with \$2/foot for aerial, and \$12/foot for underground - which comes to a blended cost of \$4/foot – for both material & labor to install. The second PON option assumes a 20/80 percent mix of aerial/underground plant with a blended cost of \$10/foot. The same CIN headend network of Figure 15 feeds node located R-PON remote optical line terminals (R-OLTs), using SFP+ modules at both ends, just like the ESD case. A quantity 4 of optical wavelengths are required, given the 1x 128 splitting ratio presumed for the 10G R-PON case. Field optical splitters costs are shown in taps row; drops, at \$200 each, and CPE, at \$120 each, are allotted for subscribed premises (50% of homes passed) only.

As can be seen by the results at the bottom of Table 1, the initial PON upgrade costs dwarf the initial HFC upgrade costs, by up to 10x more. The billion-dollar question is whether the PON networks save the operator enough over 30-years in OPEX and additional CAPEX savings to make this investment worthwhile.

3.2.1. Operational Expenditures (OPEX)

The OPEX includes headend and field power consumption, plant cable and drop cable repair and maintenance costs, and field active maintenance (material and labor to address equipment failures).

Field active maintenance cost in time is based on the heuristics curve from [Broadband Pie] and is repeated in Figure 16. As can be seen, there is a significant rise in failures after year 10.




Figure 16: Field actives percentage fail heuristics curve

I-CCAP 1.2 GHz headend powering needs are estimated as 124 Watts (W) per node area, including 50% cooling provision overhead for the buildings Heating, Ventilation and Air Conditioning (HVAC). Field line power supplies are presumed to operate with 85% efficiency, feeding 90W, 45W, and 25W to nodes, bridgers, and line extenders, respectively, assuming 3% in-coax ohmic loss.

For the 1.8 GHz DAA, the headend side consumption drops to 18W, while the field component assumptions change to 150W, 60W, and 35W, for nodes, bridgers, and line extenders, respectively.

The R-PON R-OLTs are modeled at 18W and 22W per 10G OLT port, on the headend and node side, respectively.

With \$0.12 per kilowatt-hour (kWh) assumption, the powering cost conveniently comes out to “kilo-dollar per kWh” for the whole year:

- 24 hours x 365.25 days = 8,766 hours in a year, times \$0.12 / kWh = \$1,052 per kWh per year

For the HFC upgrades, an annual upkeep is expected to be required for 1% of the hardline plant, as well as 1% of the drop-coax. This has also been built into OPEX. For the PON case, however, only 0.35% upkeep assumption of the cables, and 0.5% for the drop fibers has been considered.

3.3. 30-year TCO Cash Flows – in Nominal \$

Our economic models take in all the above CAPEX and OPEX assumptions and spit out TCO cash flows, over the 30-year period considered. Figures 17-21 display these cashflows, with plant OPEX, plant CAPEX, and CPE CAPEX color coded. Note the difference in scales. These are in nominal dollars and do not reflect the time value of money mentioned earlier. For each case, the initial upgrade CAPEX shows in year 2023. HFC plant OPEX grows slightly in time, driven by the increase of field actives failures, per Figure 16. Year 2038 shows the next big CAPEX investment, to address the aging of the HFC plant actives. HFC CPE CAPEX assumes 20% of all CPEs getting replaced every 3 years, in the D3.1 and D4.0 cases; and 25% of all CPEs replaced every 5 years in the case of the 10G PON.

Different Tav_g growth rates were used with different scenarios to show the potential range of costs. The 1.2 GHz high-split upgrade scenario assumes a low growth rate. This shows a best-case scenario from a cost perspective. The 1.8 GHz ESD upgrade uses the medium growth rate to represent a middle of the road scenario. Both R-PON scenarios and the ESD to R-PON migration assume a high growth rate.

The 2049 CAPEX upgrade in Figure 17 reflects the need to split the 1.2 GHz I-CCAP service group in order to keep up with the low-CAGR growth assumption. Similarly, the 2046 and 2042 CAPEX upgrades in Figures 18-20 reflect the need to split 1.8 GHz ESD and 10G PON service groups, in order to keep up with the moderate and high-CAGR scenarios, respectively.




Figure 17: TCO of 1.2 GHz HS upgrade, with plant CAPEX + OPEX, and CPE CAPEX




Figure 18: TCO of 1.8 GHz ESD upgrade, with plant CAPEX + OPEX, and CPE CAPEX




Figure 19: TCO for 10G R-PON, 20% U/G, with plant CAPEX + OPEX, & CPE CAPEX




Figure 20: TCO for 10G R-PON, 80% U/G, with plant CAPEX + OPEX, & CPE CAPEX




Figure 21: TCO for ESD to R-PON upgrade with plant CAPEX + OPEX, & CPE CAPEX

Figure 21 represents an interesting scenario where the operator invests in an ESD upgrade in 2023. However, if a high growth scenario is followed for 20+ years, then the ESD plant SG sizes start to become challenged in the '40s decade. This case study assumes a FTTP overlay that starts in 2044 and assumes 20% underground plant. The “heavy” DOCSIS users are migrated to FTTP. For our analysis, the 20/80 rule was followed where the top 20% move to FTTP while the lower 80% remain on ESD. With the 20/80 rule, the top 20% represent 80% of the BW usage (i.e., FTTP) while the lower 80% represents only 20% of total BW consumption. A blended approach like this can keep the majority of subscribers on HFC for many, many decades.

3.4. 30-year TCO Cash Flows – in 2023 \$

CAPEX and OPEX cash flows over time, as shown in Figures 17-21 above, provide lots of information, however, are hard to compare to each other given the different timing of different expenses. One way to deal with this issue is to bring valuation of all the future flows back to the present time – or, as shown in Figures 22-26 - to bring the valuations to 2023 dollars. As stated previously, a 7.5% annual discount rate has been used to perform this “cash travel in time.” In particular, note the decrease in the cost components that are 15-30 years in the future. This is perhaps most dramatic in figure 24 for the ESD to R-PON migration scenario. Compare Figure 24 to Figure 21 to see the effect of TVM.




Figure 22: TCO of 1.2 GHz high-split upgrade over time, in '23 dollars




Figure 23: TCO of 1.8 GHz ESD upgrade over time, in '23 dollars




Figure 24: TCO of ESD to R-PON upgrade in '23 dollars




Figure 25: TCO of 10G R- PON upgrade, 20% U/G, in '23 dollars




Figure 26: TCO of 10G R- PON upgrade, 80% U/G, in '23 dollars

3.5. A Nano-Hertz Spectral Analysis

Looking at these ‘cash flows over time’ graphs, a clear pattern emerges: some CAPEX takes place only once in 30 years (fiber plant build, HFC tap upgrade, demand-driven service group splits), every 15 years (aged actives upgrade), every 5 years (PON ONUs), 3 years (HFC CPEs), and some annually (OPEX). “Your mileage may vary” adage applies here, however; the above frequencies accurately reflect the assumptions made.

This is where the signal analysis time domain / frequency domain analogy came from: what would these expenses look like if viewed in the ‘frequency domain’? Figures 27-31 provide the answer.




Figure 27: 1.2 GHz high-split upgrade TCO in frequency domain, in ‘23 dollars




Figure 28: 1.8 GHz ESD upgrade TCO in frequency domain, in ‘23 dollars




Figure 29: ESD to R-PON upgrade TCO in frequency domain, in '23 dollars




Figure 30: 10G R-PON 20% U/G upgrade TCO in frequency domain, in '23 dollars




Figure 31: 10G R-PON 80% U/G upgrade TCO in frequency domain, in '23 dollars

Frequency of 'once every year' can be calculated in Hz, which is 1/s, by counting how many seconds there are in an average year (31,557,600), and then taking the inverse; to get 31.7×10^{-9} or 31.7 nano Hertz (nHz). That's why the annual OPEX in Figures 27-31 is positioned where it is. Similarly, 30-year periodicity is at ~1 nHz, 15-year at ~2 nHz, 5-year at ~6 nHz, 3-year at ~10 nHz, and so on.

The ~2 nHz dominant CAPEX peaks in Figures 27 and 28 reflect the 15-year cycle of active network elements upgrades (nodes, RF amps, headend optics), while ~1 nHz peaks comprise the 30-year cycle of RF tap replacements and occasional once in 30-year node splits.

The ~2nHz CAPEX peak in Figure 29, ESD to PON upgrade, is an almost exact replica of the same peak of Figure 28, ESD only upgrade, while the ~1 nHz peak captures the once in 30 years fiber overbuild. For the two PON upgrades in Figures 30 and 31, the ~2 nHz peak reflects 15-year cycle of R-OLT refreshes, while the left-most peak of ~1 nHz captures the rest: fiber overbuild, splitting network, drops buildout – investments that are made just once over the observed period. For all the five cases in Figures 27 -31, CPE capex reflects 3-year cycle for HFC, and the 5-year cycle for PON, while the OPEX shows at the annual-cycle frequency of ~31 nHz and is ~50% lower for the PON upgrades, in comparison to the HFC ones.

3.5.1. Network Upgrade Comparisons – D3.1 vs. ESD vs. FTTP

Comparison among various upgrade approaches seems easier in the frequency domain – in good part because this domain represents ‘integral over time,’ and in our case expressed in 2023 dollars. Table 2 compares the considered scenarios, by adding up NPV of plant CAPEX, plant OPEX and CPE CAPEX, already shown in the frequency domain plots, to get the comparison in NPV TCO.

Table 2: Network upgrade scenarios compared, in '23 dollars

Upgrade Scenario	1.2 GHz D3.1 HS	1.8 GHz D4.0 ESD	1.8 GHz ESD (w PON overlay ⁴⁴)	10G R-PON 20% Underground	10G R-PON 80% Underground
Tavg Growth Scenario	Low	Medium	High	High	High
NPV (Plant CAPEX)	\$103	\$146	\$209	\$468	\$785
NPV (Plant OPEX)	\$100	\$126	\$131	\$44	\$58
NPV (CPE CAPEX)	\$31	\$36	\$38	\$89	\$89
NPV TCO (7.5% TVM)	\$235	\$308	\$380	\$601	\$932
Total cash outlays over time	\$496	\$643	\$971	\$826	\$1,176

While the initial 10G PON CAPEX towers at up to ~10x the cost for the HFC upgrades, the TCO ratios for the 30-year period with OPEX included have dropped to 2x to 3x, for PON upgrade compared to 1.8G ESD upgrade; 2.5x to 4x compared to 1.2G HS I-CCAP upgrade. Time value of money discounting improves the HFC upgrades, given many expenses that are delayed in time. The PON upgrades, however, are heavily front-weighted, even though PON benefits from 50-60% lower OPEX, as compared to the two HFC upgrades.

Perhaps the most interesting result is the middle column – the ESD to R-PON migration. Traditionally, operators may ask why should they invest in ESD now if they must jump to FTTP in the future? Looking at the total cash outlays, in nominal dollars, in the last row of table 2, notice that this option comes in at \$971/HP while the equivalent 10G R-PON is at \$826/HP. However, factoring in the time value of money results in the ESD to PON migration scenario TCO of only \$380/HP in '23 dollars, while the R-PON

option is \$601/HP. That is almost 60% more cost for R-PON only in '23 dollars. So, if an operator chooses the ESD path in '23, the downside risk, if consumption were to follow the high growth rate, is very manageable and would costs significantly less than R-PON.

In summary, the PON upgrades still cost the operator double, triple, or even quadruple the amount of '23 dollars invested in the plant. But is this difference significant or immaterial to operators' budgets today? To answer this question, the next section compares these cost outlays to operators' existing business-as-usual CAPEX rates.

4. MSO Perspective – Current CAPEX vs. Upgrade TCO

Stepping back for a second, let's look at some key MSO's current capital expenditures to see how these various upgrades fit into their business-as-usual spending. Four USA cable operators: Comcast [CMCSA – cable segment only], Charter [CHTR], Altice USA [ATUS], and Cable One [CABO] are publicly traded companies and provide a wealth of information in their annual report filings and quarterly earning updates. Figure 32 shows annual amount of capital expenditures of each operator for the years 2019 – 2022. (The amounts for 2019-2021 are actually spent, the 2022 number is based on operators' full-year guidance numbers declared at their 2022 Q1 earnings calls). Comcast Cable and Charter, each with ~60 and ~55 million homes and businesses passed, are shown separately from the two smaller ones: Altice USA and Cable One, with ~9.3 and ~2.7 million home passings, respectively.

So, what, some may say? Well, these numbers on their own maybe don't say much. That is why they're often expressed as a percentage of revenue – and typically > 10% of the revenue, because cable is one capital intense business. Nevertheless, those same annual CAPEX numbers normalized by operators' number of homes and business passed (i.e., passings), are much more informative: Figure 33 shows such a normalization, with a weighted average line also added. Even without the addition of the weighted average line, most of the annual data points fall in the \$110-\$140 per home or business passed range, with the weighted average coming closely to the \$130/year value.

Of the four operators, ATUS has been the most vocal about migrating its customers to FTTP. Note how much higher their CAPEX jumped in 2022 compared to the other three operators.




Figure 32: Annual CAPEX for four USA MSOs: Comcast, Charter, Altice USA, Cable One




Figure 33: Annual CAPEX for four USA MSOs, Normalized per HP

But what percentage of annual CAPEX is directly related to the network of Figure 11? Charter [CHTR] 1Q22 investors presentation offers a clue – during the year 2021, about 27% of the CAPEX went to ‘CPE/Install,’ with the other 33% to the network side – for ‘Line Extensions’ (~23%) and ‘Upgrade/Rebuild’ (~10%) categories, as shown in Figure 34.


One thus may extrapolate that about $27\% + 23\% + 10\% = 60\%$ of overall operators CAPEX goes into the network, which, multiplied with the values of ‘weighted average’ from Figure 33 gives ~\$77 per-year per-home-passed, of which a slightly larger part (~\$42) applies to the plant and the rest (~\$35) to the CPE CAPEX. Note that a large portion of the CPE CAPEX applies to set top boxes (STB) and digital video recorders (DVR) for video service inside subscriber’s homes. As operators migrate to (Internet Protocol (IP) Video, there will be some reduction in the CPE CAPEX that could be applied to the network CAPEX.

Capital Investment

Capital Expenditures by NCTA Category

(In Millions)

■ CPE/Install ■ Scalable Infrastr. ■ Line Ext. ■ Upgrade/Rebuild ■ Support



Capital Expenditures

(In Millions)

			LTM	
	1Q21	1Q22	1Q21	1Q22
Cable	\$ 1,709	\$1,783	\$ 7,242	\$7,227
Mobile	112	74	533	444
Total	\$ 1,821	\$ 1,857	\$ 7,775	\$ 7,671
<i>Of which: Commercial</i>	\$ 333	\$ 365	\$ 1,397	\$ 1,477
<i>Of which: Rural Construction Initiative</i>	n/a	\$ 232	n/a	\$ 232

Highlights

- 1Q22 capex of \$1.9B comprised of \$1.8B cable and \$74M mobile
- Cable capex includes \$232M of rural construction initiative spend, most of which is included in line extensions capital
- Y/Y increase in line extensions of \$143M due to the rural construction initiative
- Y/Y decrease in support of \$48M primarily due to lower mobile capex and timing of vehicle spend
- Y/Y decrease in scalable infrastructure of \$40M primarily due to timing of spend
- Mobile capital expenditures of \$74M primarily for back office systems, most of which is included in support capital

Figure 34: CAPEX by category, from Charter's 1Q22 financial results

As a first sanity check, let's look at \$~50 per HP per year for network CAPEX, expanded every 30 years, with a 7.5% discount rate. This results in \$590 NPV TCO values expressed in 2023 dollars. Thus, compared to the values in Table 2, if CAPEX were to continue at the rates shown above, any of the HFC upgrade paths could be afforded. A fiber upgrade may be possible if the plant is mostly above ground. Otherwise, a significant increase in CAPEX over 30-years is needed if there is significant underground plant.

Now consider that the operator has roughly \$25/HP to \$30/HP per year available for these major network upgrades. It turns out that the operator could upgrade all systems to 1.2 GHz over three years; or upgrade all systems to ESD over a 4- to 5-year window. This seems very reasonable. A hybrid of mainly HFC upgrades, plus a careful mix of the FTTP PON ones, where necessary, is another possibility. The necessity is primarily driven by competitors' actions.

If the operators stay at \$25/HP to \$30/HP per year for the FTTP upgrades, these will need to get spread over 20 to 30-year window, and this would likely not fix the BW problems in a timely manner. If the operator bumps the network CAPEX investments up to \$80/HP to \$100/HP range (e.g., similar to what ATUS may have done), then FTTP upgrade path would take 5- to 6-years for mostly aerial plant and 8- to 10- years for mostly underground plant. And this doesn't include any HFC upgrades in the interim to remain competitive and provide required capacity to maintain existing QoE.

Caveat Emptor (buyer beware) warning is in order here – above statement are valid, provided all of the assumptions made above are valid.

5. Variations / sensitivity analysis

“One should not make predictions, especially about the future” is a quip variously attributed to Samuel Goldwyn of Metro-Goldwyn-Mayer fame, [Goldwyn], Yogi Berra of baseball-playing philosopher fame [Yogi], and to Niels Bohr, Nobel-prize-winning quantum physicist [Bohr]. Taking these esteemed gentlemen’s advice to heart, this section is more about the range of possibilities rather than some precise foretelling of how the networks shall evolve – because only time will tell.

5.1.1. HFC Sensitivity analysis

Thus, rather than provide predictions, Monte-Carlo analysis of Figure 35 and Figure 36 show a range of possible outcomes for TCO of the two HFC upgrades. The graphs show a range and probability of outcomes, based on a certain set of assumptions specified. These ‘frequency views’ are formed after a run of 100,000 trials is completed.




Figure 35: TCO Sensitivity for 1.2 GHz High-Split upgrade – 100,000 trials Monte-Carlo run




Figure 36: TCO Sensitivity for 1.8 GHz ESD upgrade – 100,000 trials Monte-Carlo run

Rather than discuss a single value under certain assumptions, as done in the previous section, these charts offer a range of outcomes, given certain ranges of assumed variables. Thus, a 95% confidence interval for 1.2 GHz high-split I-CCAP upgrade is \$193-\$310, vs. \$254-\$405 for the 1.8 GHz ultra-high-split ESD upgrade. Note that the high end of both 95% confidence intervals is ~30% higher than the base case.

To better understand the model's sensitivity to various assumptions, Figure 37 and Figure 38 display "Tornado charts", for the two HFC upgrade cases: 1.2 GHz high-split I-CCAP and 1.8 GHz ultra-high-split ESD. The Tornado chart ranks each variable's impact from most on the top to least on the bottom. The HFC upgrades are the most sensitive to the number of passings per node, followed closely by the discount rate assumed. Other variables had noticeably less impact and ranked as: cost of kWh of power, number of amplifiers in the network, and so on.




Figure 37: Sensitivity Tornado chart – 1.2 GHz high-split




Figure 38: Sensitivity Tornado chart – 1.8 GHz ESD

5.1.2. R-PON Sensitivity Analysis

10G PON cases have charts of their own: Figure 39 displays the distribution and range for the TCO of 10G R-PON upgrade with 20% underground plant. Its 95% confidence interval spans \$512-\$817; under various assumption ranges, as shown in the sensitivity Tornado chart of Figure 41. To no surprise, the % of underground plant had the biggest impact on sensitivity. The high end of the 95% confidence interval is ~36% higher than the baseline. This shows that FTTP upgrades have more sensitivity in their cost analysis than their HFC counterparts.

Figure 40 displays the distribution and range for the TCO of 10G R-PON upgrade with 80% underground plant. Its 95% confidence interval spans \$782-\$1,370 under various assumption ranges, as shown in the sensitivity Tornado chart of figure 42Figure 41. Because the % of underground plant is already very high,

the cost of fiber material & labor had the biggest impact on its sensitivity. The high end of the 95% confidence interval is almost 50% higher than the baseline, showing even more variability. This shows that the higher the % of underground plant, then the higher FTTP sensitivity becomes.




Figure 39: TCO Sensitivity for R-PON 20% U/G upgrade – 100K trials Monte-Carlo runs




Figure 40: TCO Sensitivity for R-PON 80% U/G upgrade – 100K trials Monte-Carlo runs

Order of variables affecting the TCO of PON upgrade differ from those of the HFC upgrades – understandably so – the cost of running new fiber is highly sensitive to the % of the plant that's underground, as opposed to aerial. Plant length is closely behind, as is fiber installation cost. Interestingly, the discount rate does not affect the outcome as prominently here as it does for the HFC upgrades. This can be explained by HFC annual costs fairly evenly distributed in time, and thus

benefiting or not from a high/low discount rate, while the PON costs are mainly upfront and as such don't get much of a benefit if the discount rate is high.




Figure 41: Sensitivity Tornado chart for the 10G R-PON 20% U/G upgrade




Figure 42: Sensitivity Tornado chart for the 10G R-PON 80% U/G upgrade

5.1.1. ESD to R-PON Sensitivity Analysis

Perhaps the most interesting upgrade path under consideration for the high growth projection is starting with an ESD upgrade in '23 and then to add an R-PON overlay in '44. This had a baseline NPV TCO of \$380/HP. The Monte Carlo sensitivity analysis for this scenario is shown in Figure 43. The 95% confidence interval is \$307/HP to \$533/HP. The top end of this blended upgrade scenario is ~40% higher than the base line.

The Tornado chart is shown in Figure 44. For this case, the discount rate is the most impactful variable followed by HP per parent node. The % of underground plant is a distant third to these inputs, with variation of ~\$37 per 20% increase in underground runs. This means that even an 80% underground plant

would add just \$110 to the NPV TCO of the ESD-to-R-PON upgrade, as compared to adding \$331 if PON upgrade were to be done in 2023.




Figure 43: TCO Sensitivity for ESD to R-PON upgrade – 100K trials Monte-Carlo runs




Figure 44: Sensitivity Tornado chart for the ESD to R-PON upgrade

6. Conclusion

Damn the torpedoes – the humankind’s inability to predict the future ought not to stop us from trying to envision what a possible range of network upgrade outcomes may look like – that’s what this paper is about. Common to all scenarios was the starting point of 400HP, 200 subs per node, with the top tier of 5/1 Gbps DS/US as the common goal.

The 1.8 GHz ESD plant upgrade handled the low to moderate DS Tav_g growth projections just fine, through the 30-year window. The moderate case needed two node splits – one in 2038 and the other in 2046. Its TCO/HP came to \$308 in '23 dollars. If DS Tav_g follows the high growth projection for 20+ years, then a R-PON overlay might be needed to migrate ‘heavy’ customers to FTTP. This increased the NPV TCO/HP up to \$380. This effectively provides a ceiling for the potential NPV TCO based on low to high growth scenarios.

For the low growth DS Tav_g projections, the 1.2 GHz high split I-CCAP upgrade held up surprisingly well, with just one node split required in 2049. Its TCO/HP came to \$235. For non-competitive markets

and/or tight capital budgets, this remains a solid option. However, if DS Tavg follows the moderate to high growth projections, then a switch to ESD or FTTP will be needed during the next decade.

The FTTP PON upgrades are often associated with the high growth scenarios, but cost 2x higher than ESD NPV TCO/HP (i.e., \$308 vs. \$601) when overbuilding 80% aerial plant; and 3x higher (i.e., \$308 vs. \$932) for 20% aerial plant (and the rest underground). Along the way, the 10G R-PON original 1:128 splitting group had to be halved twice: once in 2035, and second time in 2042. And if the DS Tavg growth follows the low or moderate projections, then the operator may be spending a lot of excess money it might not need to.

Pushing big spending decisions into the future does provide significant value in today's dollars, as seen in "1.8 GHz ESD upgrade now, overbuild with fiber in 2044" scenario. This scenario benefits from the lower NPV TCO of the ESD route, as well as from the future high capacity of PON, but only implemented if needed. The blended ESD/PON path NPV TCO comes in 37% saving compared to doing PON now, and only at a 23% premium, in 2023 dollars, as compared to staying with ESD for 30 years.

Based on the track record of the publicly traded USA MSOs, the CAPEX dollars are likely to stay available for the lower cost upgrade approaches, provided operators' profitability stays on part. Another option is to use a hybrid approach: a lower-cost HFC upgrades where the market drivers call for it, and the higher cost/higher capacity PON in competitive markets. To conclude, if in doubt, let the market show you the way.

7. Acknowledgements

David Bowler, for expert guidance on finer aspects of 1.8 GHz ESD plant upgrade.

Abbreviations

10G	10 gigabits per second
AR	augmented reality
BW	bandwidth
CAGR	compound annual growth rate
CAPEX	capital expenditures
CATV	community antenna television
CIN	converged interconnect network
CM	cable modem
CMTS	cable modem termination system
COVID	coronavirus disease
CPE	consumer premises equipment
D3.1	DOCSIS 3.1
D4.0	DOCSIS 4.0
DAA	distributed access architecture
DDM	dividend discount model
DFN	distribution fiber network
DOCSIS	data over cable service interface specification
DS	downstream
DWDM	dense wavelength-division multiplexing
ESD	extended spectrum DOCSIS
FTTP	fiber-to-the-premises
HEO	head-end optics
HFC network	hybrid fiber-optic and coaxial cable network
HP	homes-passed
HS	high-split
I-CCAP	integrated converged cable access platform
IP	internet protocol
IPTV	internet protocol television
MAC	media access control
Mbps	megabits per second
MDU	multi dwelling unit
MSO	multiple system operator
nHz	nano Hertz
NPV	net present value
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OLT	optical line terminal
ONT	optical network terminal
ONU	optical network unit
OPEX	operating expenditures
PHY	physical layer
PNM	proactive network maintenance
PON	passive optical network
PV	present value

QAM	quadrature amplitude modulation
QoE	quality of experience
RF	radio frequency
RPD	remote PHY device
RMD	remote MAC/PHY device
ROI	return on investment
R-PON	remote PON
SC-QAM	single carrier quadrature amplitude modulation
SDV	switched digital video
SFP	small form factor pluggable
SFP+	enhanced small form-factor pluggable
SG	service group
SLA	service level agreement
TCO	total cost of ownership
TVM	time value of money
UHS	ultra-high-split
US	upstream
VR	virtual reality
W	Watt

Bibliography & References

[ATUS] Altice USA corporate investors page:

<https://investors.alticeusa.com/investors/overview/default.aspx>

[bbcmag FTTH OPEX] <https://www.bbcmag.com/broadband-applications/to-reduce-network-operating-expenses-choose-ftth>

[Bohr] The Economist magazine, letters to the editor in response to “The Perils of Prediction” article
<https://www.economist.com/letters-to-the-editor-the-inbox/2007/07/15/the-perils-of-prediction-june-2nd>

[Broadband Pie] Z. Maricevic, J. Andis, T. Cloonan and J. Ulm: “DOCSIS 4.0 - A Key Ingredient of the 2030's Broadband Pie”, 2021 Cable-Tec EXPO

[CABO] Cable One corporate investors page: <https://ir.cableone.net/corporate-profile/default.aspx>

[CHTR] Charter corporate investors page: <https://ir.charter.com/investor-overview>

[CLO_2014] T. J. Cloonan et. al., “Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements,” SCTE Cable-Tec 2014, SCTE

[CMCSA] Comcast corporate investors page: <https://www.cmcsa.com/investors>

[FTTH OPEX] Fiber Broadband Association: “Operational Expenses for All-Fiber Networks are Far Lower Than for Other Access Networks”, June 2020, <https://www.fiberbroadband.org/page/fiber-research>

[Goldwyn] The Economist magazine: “Biology’s Big Bang”
<https://www.economist.com/leaders/2007/06/14/biologys-big-bang>

[Gordon] Dividend Discount Model: <https://www.investopedia.com/terms/d/ddm.asp>

[HOWALD_2022] “Collision-Free Hyper-Speeds on the Bi-Directional FDX Highway”, Dr. Robert Howald, John Ulm, Saif Rahman, Dr. Zoran Maricevic; SCTE Cable-Tec 2022, SCTE

[ULM_2022] Ulm, John and Zoran Maricevic, Ram Ranganathan; “Broadband Capacity Growth models – will the End of Exponential Growth eliminate the need for DOCSIS 4.0?”, SCTE Cable-Tec 2022, Sept 19-22, Philadelphia, PA.

[ULM_2019] “The Broadband Network Evolution continues – How do we get to Cable 10G?”, John Ulm, Dr. Tom Cloonan, SCTE Cable-Tec 2019, SCTE

[Yogi] The Economist magazine: “The Perils of Prediction” <https://www.economist.com/books-and-arts/2007/05/31/the-perils-of-prediction>

Open Source, a Mindset and How it Has Transformed Common Platform Enumeration Stack Software Development using Reference Design Kit

A Technical Paper prepared for SCTE by

Khem Raj
Fellow
Comcast
1050 Enterprise Way, Sunnyvale, CA 94089
(408) 768-2010
Khem_raj@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Open Source Development.....	3
2.1. Community	3
2.2. Reliability and Security.....	4
2.3. Transparency	4
2.4. Support and Maintenance	4
2.5. Working with Open Source Communities	5
2.6. Open Source Licensing.....	6
2.7. Supporting Key Open Source Projects	7
2.8. Seeding New Open Source Communities	7
2.9. Open Source Program Office (OSPO).....	8
2.10. Attracting Best Talent.....	9
3. RDK and Open Source	9
3.1. RDK Infrastructure	10
3.2. RDK development and workflow	10
3.3. RDK Collaboration	11
3.4. Key Takeaways from Open Collaboration	11
4. Conclusion.....	12
Abbreviations	12
Bibliography & References.....	12

1. Introduction

“Software is eating the world” quipped Mark Andreessen in 2011 [4]. Today it is apt to say that **“Open Source is eating the software world.”** From toasters to a helicopter on Mars, systems are using a robust amount of open source software in their stack. According to the 2021 Open Source Security and Risk Analysis Report by Synopsys [1], 75% of all codebases were comprised of open source code alone.

However, it is still not always easy to work with open source projects 30 years after open source came into being. Organizations and individuals experience challenges in gauging the health of a project: is it stable, secure, sustainable, and will it be there if I depend on it, are common concerns. While the acquisition of open source is free, there is significant effort needed to comply with the hundreds of open source licenses used. To be successful, it is important for organizations that consume open source to contribute back, and to engage with the project. Projects come in all shapes and sizes, with different processes, tools, and funding models; knowing how to navigate open source is a learned skill. As open source companies mature, they also need to learn to release code that they have created, and learn to build communities of users and contributors around it.

2. Open Source Development

Adopting open source development models is flourishing as more businesses discover the advantage of open source projects. Those advantages over proprietary solutions, and the mindset around open source, continue to change as it becomes a more prevalent solution pool. Open source is a mindset which can be applied to solve different problems; it is effectively thriving in software development, but we are seeing an open source development model being used for new hardware development as well, such as Reduced Instruction Set Architecture (RISCV), which involves open source being adopted across the central processing unit (CPU) industry. The recent COVID-19 pandemic also saw open source approaches used for contact tracing and collaboration across the medical industry. While a traditional approach followed a series of steps including requirement, design, implementation, testing and quality assurance (QA), system integration, and deployment all performed serially, in open source the work is done in smaller increments and a “release early release often” approach is adopted. Many of the steps described above are performed in parallel and are iterated more often which allows end users to access the code early in development cycles. This approach is quite successful as it allows faster development and testing, transparency and openness, rapid innovation, enhanced collaboration, and high-quality peer reviewed code.

2.1. Community

All successful open source projects build upon a healthy, growing community of users, developers, maintainers, and promoters. They all share a common goal which is served by the project. The communities are global in nature, comprised of people from diverse backgrounds, which empowers the project with new ways of solving problems with novel ideas. This helps project members to develop new features and capabilities at a faster rate than isolated internal teams could. A collective approach of the community consisting of talented people across the

world can deliver solutions faster and more rapidly innovate. This approach establishes an environment of efficient contributions and community team collaboration. An open source community has a strong desire to engage in meaningful productive work, aspiring to develop a sense of belonging, and make it easy for new members to feel at home and connected with others. Open source communities differ in many aspects and each has a unique culture. It cannot be assumed that what is learned as part of one community will be directly applicable in another community. Openness in community fosters authenticity, and members tend to respect it. Therefore, most thriving development communities are those that conduct their activities in open.

2.2. Reliability and Security

Open source projects host the source code in open repository infrastructures such as GitHub or GitLab, making the code more accessible to a wider audience of developers and individuals. This approach results in more individuals looking into the source code regularly. Some users will test newly developed code and provide feedback to developers. Because the code is reviewed in an open environment, healthy reviews are the norm, making the code more robust prior to acceptance[7]. Since code is subjected to much more thorough review and testing, it also tends to be more secure [6], and any security issues that do arise afterwards are handled briskly. Updates are frequent which provides a smaller window for exploiting any weaknesses in code. A major motivation in decision-making is to select the best possible solution. In some cases, this means selecting best technical solution instead of being influenced by other factors which might provide a quick but technically inferior solution. This merit-based approach is key reason for open source projects being more reliable and secure[5]. It is good to use security response time for patching security related bugs of a given project as part of selection criteria for an open source project.

2.3. Transparency

In an open-source community, code is available openly, and all discussions happen on open channels such as internet relay chat (IRC), web forums, or through mailing lists. All decisions regarding new features and defects also follow open channels. This information is very valuable for users as they can see the technical merits or limitations of the solution and avoid any unwanted surprises later in deployments. Open channels also encourage users to discuss the code issues they encounter with the community and forge collaborative solutions. This also avoids vendor lock-in which might be a consequence of using closed solutions.

2.4. Support and Maintainance

Any software will need support, and open source projects are no exception. While proprietary solutions are extensively supported by vendors providing a single point of contact, this model may not translate directly to open source, as the software is created by a number of developers covering different parts of the project. Therefore, it is important to look into the support aspect of projects and find third-party dedicated support or use the online community-provided support if that is sufficient. There may be no upfront cost to using an open source project, though there could be additional cost due to integration or maintenance of the project needed in the business solution. After some time, an open source project may stop supporting a prior version that was deployed in products, which could mean additional cost to engage third-party software support. It is beneficial to understand these hidden costs beforehand for better estimation of support costs

for the products. It is worth stating that the costs for support and maintenance will be reduced compared to a traditional approach, but it will not be zero, as some might proclaim or assume since the code is available for free.

2.5. Working with Open Source Communities

Each open source community is unique and has subtle characteristics which need to be well understood for one to effectively participate in the community. It is important to understand that communities are united by a common interest; finding this interest and understanding the goals of a given project is of utmost importance. Communities have laid out certain rules of engagement, though not all of them may be documented. These can be learned by reading through community channels, forums, and mailing lists, or by interacting with experienced community members.

Usually, projects have a README file which is a good place to get acquainted with the project. There are other common files which may be provided by projects including:

- CONTRIBUTING file, which would list contributing procedures and tools required.
- MAINTAINERS file, which is a good resource to learn about key developers and maintainers who can be valuable points of contact as you start engaging with the community.
- LICENSE or COPYING file, which lists the open source license governing the content of the source provided by the project. This is crucial information which should be well understood. It may require review by your organization's legal team to provide guidance on the mode of engagement that can be pursued.

Some projects with a large codebase and community create governing bodies to take care of project administration and future technical directions. These governing bodies are instrumental in shepherding the project and it is important to understand the role and functions of such committees if they exist.

There are frequently multiple projects available for a given requirement in a product. Therefore, it becomes important to select the best fit. The maturity of a project is an important mark; looking at the number of commits and rate of commits can give a good hint of project maturity. Some infrastructures such as GitHub also provides stars for a project which are awarded by users of the project; more stars means that more users found the project useful for their needs. This could be another parameter to check. Regular release cycles can be used to deduce how actively code is maintained. If a project does not have a deterministic release cadence it becomes difficult to engage with such projects because it would be hard to align with your project roadmap. In other instances, the number of open defects and the rate at which the issues are handled can serve as a good mark of project development. A larger number of contributors is a good sign as well, since it would mean that it would be easier to find help when needed, while diversity in committers could also be indication of good quality within a project. There are other markers that can indicate the strength of a project including Continuous Integration (CI), good documentation, and examples and lists of other projects using or depending upon the project. Open source project space is dynamic, and projects can be started and soon become stale or dead. The risk that arises due to this unpredictability can be minimized by taking precautions as mentioned above.

Contributing to a project is important and is the lifeline of any open source project. If you depend on an open source project, you should consider contributing to the project. Doing so would benefit your use case as it also supports the overall project. When we speak of contributing, it is often equated with contribution made in the form of code. Even though it is a significant contribution category, this is not the only way to contribute to a project. There are other aspects of projects that can benefit from contributions as well. Documentation is one key aspect of any project and does not strictly involve code right away. Open source projects interact via online forums or mailing lists, where many questions are asked, so active members can contribute by answering these questions. Website designs, infrastructure setups, providing computers for CI, defect triaging, organizing events, connecting community members, promoting and marketing the project, and mentoring new community members are additional ways to contribute to a project in a significant manner.

As one becomes a regular contributor, they can begin providing inputs for project directions along with new features and design decisions. These are deeper engagements which could require a significant time investment. Doing this, however, would be impactful and help steer the project and promote you into an influencer or leadership role in a project or community.

2.6. Open Source Licensing

Most open source software projects have licenses. If there is no license mentioned explicitly, then you should assume the software is not being licensed and therefore may not be copied, distributed, modified or otherwise used.

Some of the common open source licenses are GPL family, Apache 2.0, MIT, Mozilla Public license, and BSD family. Broadly, these licenses are classified as either copyleft or permissive. Copyleft licenses typically demand reciprocity and require that any derivative work is also licensed under the same obligations. The GPL v3 license is an example of a copyleft license. Some licenses are permissive, which grants license to use and re-license. BSD, MIT, and Apache are common examples of permissive licenses. Proprietary licenses are more restrictive than open source licenses on granting rights; these are typically non-free licenses. It is very important to learn and understand the implications of the license of a specific open source project. It will help in understanding the contribution and distribution mechanisms of the source code and its derivatives.

There are obligations to fulfill as a consumer of the project; a different set of obligations may be required to be fulfilled when distributing the project to your end-users and consumers. Generally, Open Source Program Office (OSPO) departments which help organizations in formulating the open source policies are the best resources to provide a detailed understanding of the open source licenses to developers and technical teams.

When open source components are used as part of a product, there may be more than one license which is in play. Therefore, it is absolutely important to be aware of every open source license used in products and create a license manifest. There are commercial and open source tools which can come to aid in building the open source license manifests. It is important to note that

these tools are not the final authority on signing off these manifests; assigned auditors should review the licenses and ensure that they are listed as required. These manifests are living documents, which can change as soon as the component list or dependencies change. As such, it is beneficial to generate these manifests along with production builds and have these manifests correspond to each exact source that is being used in each release.

For scalability, it is good to define an intake process in which teams are provided with the correct information and guidelines to choose a component and assess its license. This approach helps in making appropriate decisions early and avoids license-related showstoppers at later stages of a project, which can result in lost time, and in many cases has caused projects to be cancelled. A vetted process to get approval for using a particular license is good to have, with checks and balances to ensure that these kinds of problems are avoided. It is good practice to require that developers regularly take license-related training to stay current on open source licensing trends.

When creating open source projects and seeding them, it is also important to license the code appropriately. There could be many factors influencing this choice, including dependencies of the project, perhaps the inclusion of another set of open source components, or the situation in which code is licensed under a license which may be incompatible with the license you are trying to use for the new project. It is also important to align the license with project goals and community involvement. Some licenses may require more effort to fulfill compliance within the conditions in which the project will be used. Therefore, software architecture including static or shared linking may also play a part in determining which licenses are more suitable than others. It is important to consult with knowledgeable teams, especially the legal team supporting your group, to come to a decision on choosing a license. Choosing the “wrong” license can have serious consequences, so it is important to choose the best license for each project.

2.7. Supporting Key Open Source Projects

Open source dependencies for products can run deep, and in many cases the product may have a critical dependency on a potentially dead open source project. A typical response is to keep using dead code, as switching to an alternative could be quite disruptive if it would come at a time which was not planned. It is important to find your open source dependencies and evaluate projects closely for their health and long-term viability. Doing so provides a good overview of the strength of a product’s foundation. An open source project may be in need of computers for hosting testing infrastructure, there may be a downward trend in contributors, or the project may be maintained by a single person in their free time, making it important to find out how you may be called on to help a project. You may be able to support a project by donating funds, contributing bug fixes, making improvements, or introducing features. Some projects require conference sponsorship or event hosting. Creating a comprehensive support plan can be useful in maintaining continuity and helps avoid surprises when something in the community breaks down.

2.8. Seeding New Open Source Communities

Open source engagement evolves from being a user to becoming a participant, then further to being a contributor, influencer, or maintainer. These are helpful experiences to gain before starting a new community or seeding a project. Starting a new open source project involves many

activities besides making code available. There are millions of open source projects, and it is likely that a similar implementation is already available; how you differentiate your project is a key aspect to consider. An open source project would not go far without a vibrant and healthy community around it. When starting a new project, it is a big challenge to build a community of users, developers, and promoters. As such, one needs to define clear goals and a meaningful value for the audience that will drive its consumption. Doing so will help with community growth; there must be clear directions for someone new to get started with the project, and contribution guidelines and steps should be crisply defined. As communities are comprised of humans, it is important that members see a clear, tangible value in being part of community. Communities develop culture. It is not something that will be pre-defined; and, as it is essential for a community to endure, the proper care and feeding of community is necessary. Engagement models will define how new members are on-boarded, how friction is removed for regular members, and how core contributors are incentivized. This documentation will provide a good process to facilitate learning for members.

2.9. Open Source Program Office (OSPO)

As open source adoption has become prevalent, scaling is a challenge, particularly in large organizations. Some teams may be open source savvy and forge an open source-based solution which could be quite optimized and best in class, though it is not easy to replicate that level of success across different products and teams without effort. OSPOs are established to tackle such problems. OSPOs act as an interface between open source communities and the organization. OSPOs also helps connect various internal teams wanting to benefit from open source. They can be instrumental in defining open source policies for the company. Open source communities are quite dynamic in nature, and it is often required to have someone to actively look into the ongoing developments and make meaning of them from an organization's point of view. As noted above, projects periodically change their licensing terms. If these projects are used in a company's products, the OSPO can help in interpreting these changes and monitor relevant changes. OSPOs help in optimizing open source consumption, collaboration, and contributions. They keep track of statistics and data and monitor the progress on a regular basis. This type of data is important in measuring productivity relative to open source goals that the organization might have set forth. There are various events and conferences related to different open source technologies, and as projects are organized across different geographies, OSPOs track event calendars and bring vital information to the company teams. The OSPO also connects internal team members who may be participating in the same events. They can also help with internal initiatives to drive the open source mindset inside the organization, to promote inter-team collaboration and code sharing, resulting in similar benefits for proprietary code development as well. To assist with career growth in engineering, developers might be interested in speaking opportunities at open source events. While the call for presentations for these events happens months ahead of the event itself, the OSPO keeps this information current and keeps the teams informed about the deadlines. OSPOs can also help in adopting various processes for open source compliance and recommending scalable tools across the organization. Given these important functions, OSPOs are becoming an integral part of modern organizations.

2.10. Attracting Best Talent

Due to their open nature, ease of access to source code, and other related tools and infrastructure, open source projects provide fertile ground for training new developers. Many developers are enthusiastic about working in an open source project or environment where an “open first” approach is used for software development. If an organization is involved in using, collaborating, and contributing to open source projects, it can be a big incentive for top class developers to seek positions there. If an organization employs well-known open source developers, it motivates other developers to join the company to work and learn from the best developers and collaborators. To attract talented developers to your organization, showcasing your community participation, involvement, and contribution can go a long way. Doing this creates a large pool of developers and engineers with experience in open source technology, which in turn ensures that there are enough people available to work on future critical open source projects.

3. RDK and Open Source

Reference Design Kit (RDK) provides a set of software components which are used to build software stacks for video clients, video gateways, broadband routers and gateways, smart camera solutions, wireless access points, and more devices used in the home. RDK is used to quickly build efficient stacks. RDK has transformed cable software stacks by providing a robust platform backed by open source solutions and technologies. The initial profiles were video client and gateway devices, which soon expanded into broadband profiles, building routers, access points, and providing Data Over Cable Service Interface Specification (DOCSIS[®]), digital subscriber line (DSL), and fiber wide area network (WAN) based devices. In the future we will see 5G modems and routers, as well as smart TVs being built using the RDK software platform.

RDK started as a foundation with RDK Management being the custodian of the code and infrastructure. RDK has adopted open source best practices and improved upon its own processes iteratively over time, establishing clear contribution guidelines and processes which are readily available for community. There are also webinars and How-To documents made available for new and seasoned contributors. The community growth is measured in terms of new members and contributors joining the project. It also measures lines of code growth and number of changesets submitted and accepted on a quarterly basis. These measurements offer key insights into community engagement and growth. RDK started with most of the code being under RDK license but has since migrated to an Apache 2.0 open source license for all new code and also relicensed most of the original code under Apache 2.0. All new components are released under an Apache 2.0 license and made open source from day one. In some cases, components are being incepted on GitHub and developed directly in the open with community collaboration.

Such best practices have been adopted and tuned over a period of time as the project grew its userbase and the needs of the community for open collaboration grew beyond the initial set of members.

RDK has also regularly hosted community events including the annual RDK Conference and RDK Tech Summit. Attendance and topics discussed in these events have grown multifold. Where in the first conference topics focused mainly on video projects, topics now include broadband technologies, wireless connectivity, augmented reality, smart speakers, smart

cameras, and new open source code and features that are being created by the RDK community. These events also serve as key collaboration and networking opportunities for the community of developers and users. There are additional events like coding challenges which are organized throughout the year.

RDK has also embraced popular open source platforms for ease of development and to broaden community outreach. Projects are mirrored on GitHub as well. Starting with just a few projects, RDK today contributes directly to many upstream projects including Yocto, Webkit, and Gstreamer. At the same time, there are RDK initiated projects which are also being used beyond RDK in other projects. This fact demonstrates how RDK started as an open source user/consumer and slowly graduated to a participant, contributing to key projects, before finally producing and seeding projects and growing the community. RDK's commitment to open source has led to leadership and influence in the cable industry in particular. RDK now has templates for open sourcing projects. These templates include key files that an open source project should have. Members can utilize these templates to quickly learn and become proficient in open source processes. It now also helps new members to become good at using RDK components by offering trainings, webinars, and dedicated community support. There is a clear path for members to become contributors and maintainers of components.

As contributions grow, conflicts will happen where a decision to adopt a given implementation or another must be made. RDK has established Technical Advisory Boards for different profiles. These committees of maintainers and contributors broker a decision based on the technical merits of contributions and the RDK roadmap.

The RDK open source community also keeps a keen eye on the latest trends and developments in other open source communities and adopts tools or best practices that were successful for other communities such as adopting Slack for instant messaging and collaboration.

From custom licensing requirements to an “open first” development mindset, RDK has transformed itself and keeps learning from open source development and processes, incorporating these learnings to address its own community challenges.

3.1. RDK Infrastructure

RDK is hosted on <https://rdkcentral.com>, which includes instructions on how to join the community, documentation, and the code itself. RDK uses self-hosted gerrit and GitHub cloud infrastructure to serve the community. Jira is used for workflow management, defect tracking, and technical support. There is extensive wiki documentation available for technical know-how, webinars, FAQs, and more. Additionally, regular meetings are held for roadmap sharing, and general community synchronization and updates. Special interest groups (SIGs) are formed as needed to forge new features or address large scale problems such as security vulnerabilities.

3.2. RDK development and workflow

All new components for RDK are developed in open source as community members and developers contribute code on a regular basis. Contributions go through an established intake process. This process involves scanning code for license compliance, build verification, runtime

sanity testing, and code review. In addition, some members also validate the patches in their internal setup and provide the testing results back as community contributions. This approach establishes a robust intake process which provides sufficient testing during code development and helps avoid regressions that may be introduced by newly added code.

3.3. RDK Collaboration

The RDK ecosystem consists of different types of members. Hardware manufacturers, typically consisting of original equipment manufacturers (OEMs), silicon on chip (SOC) providers, and specialized hardware add-ons, build hardware for specific products. Independent software vendors provide value on top of RDK platforms. Application developers build and port their applications to the platform. Service providers use these products as vehicles to deliver the services to customers. Students and universities work with the RDK platform on experimenting with technologies, learning, and research. As we can see, there is a wide array of users and contributors. Each could be working on enhancing a given segment of the RDK, such as a SOC vendor interested in unleashing the power of hardware capabilities of CPUs and other co-processing units. They would require working toolchains and infrastructure to build upon. Similarly, application developers would require a working stack on production hardware or reference systems. RDK enables these key elements by providing infrastructure, tools, and techniques upon which everyone can build. A common set of workflows for code development and acceptance criteria means that contributors are adding value for each other, and an established code quality is maintained. A common set of references are maintained as hardware platforms which are easily accessible to the community and are close enough to product hardware which users might be working on. This result provides a quick on-ramp for porting and increasing feature velocity.

3.4. Key Takeaways from Open Collaboration

The open platform that RDK has built provides many additional benefits. It has improved code reuse; since code is openly available, there is less duplication. A new stack uses common building blocks which are the same across all profiles, and scale horizontally. This stack also includes common build infrastructure and CI systems. It has resulted in better use of resources and enabled the community to avoid redundant work. Promoting re-use means that a modular software design has emerged. A departure from a monolithic mindset to a component-based mindset has been an on-going transformation. As components take shape, it has improved code ownership and maintenance, as expert teams are able to function independently while developing components maintained by them. It is an engineering community with collaboration mainly around source code. This construct has provided opportunities for any new developer, student, or other entity to submit pull requests and contribute. The feedback loop is essential for learning, and it has helped individual members and member organizations develop in-house expertise rapidly. Open source RDK has also helped training organizations to develop good programs and they can easily keep these modules up to date as the code changes, providing fresh courses and training events. These programs serve as a critical part in growing the developer base. Application developers can build their applications once and deploy many times on different RDK-based systems, streamlining the efforts to improve the application experience. There is a rapid expansion of device profiles which can be attributed to the key points discussed above.

4. Conclusion

The use of open source development methods and processes is on the rise, offering significant advantages over traditional closed software development methods. Major pieces of many complex software stacks today consist of open source software, and its share is on the rise as well. It is important to become a wise consumer and thoughtful producer of open source software. Doing this requires a deep understanding of open source communities, open source licenses, and various projects involved in building products. Strategically aiding and sustaining the communities supporting and building critical software is essential. Scaling to use open source software across larger organizations is a challenge which is being addressed by dedicated Open Source Program Offices which serve as liaisons between internal development teams and reach out to open source communities. The “open first” mindset is also on the rise, where open source solutions are sought first and developed and customized.

Abbreviations

CPU	central processing unit
DOCSIS	Data Over Cable Service Interface Specification
DSL	digital subscriber line
IRC	internet relay chat
OEM	original equipment manufacturer
OSPO	open source program office
QA	quality assurance
RISCV	reduced instruction set computer five
RDKit	reference design kit
SIG	special interest group
SOC	silicon on chip
WAN	wide area network

Bibliography & References

- [1] 2022 Open Source security and risk analysis report - <https://www.synopsys.com/content/dam/synopsys/sig-assets/reports/rep-ossra-2022.pdf>
- [2] People Powered – How communities can supercharge your business, brand and teams – Jono Bacon
- [3] 4 Stages of Open Source Consciousness – Khem Raj
- [4] Why Software Is Eating the World – <https://a16z.com/2011/08/20/why-software-is-eating-the-world/>
- [5] Reliability Issues in Open Source Software - <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.259.434&rep=rep1&type=pdf>
- [6] Is Open Source Software More Secure than proprietary Products - <https://www.govtech.com/security/is-open-source-software-more-secure.html>

[7] Open Source Case for Business - https://opensource.org/advocacy/case_for_business.php

Optimizing Wi-Fi Channel Selection in a Dense Neighborhood

A Technical Paper prepared for SCTE by

Yonatan Vaizman

Lead Researcher, Machine Learning
Applied AI & Discovery, Comcast
1325 G Street NW, Washington, DC 20005
202-524-5077
Yonatan_Vaizman@comcast.com

Hongcheng Wang

Distinguished Engineer, Machine Learning
Applied AI & Discovery, Comcast
1325 G Street NW, Washington, DC 20005
332-301-5055
Hongcheng_Wang@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Wi-Fi Pain Metric.....	3
3. Optimization Problem and Solvers.....	5
3.1. MIQP Problem Solver	6
3.2. Neural Network Gradient Descent	6
4. Preliminary Experiments	8
4.1. Estimating Potential Pain	8
4.2. Optimization Details	10
4.3. Results	10
5. Conclusion.....	10
5.1. Future Directions.....	11
6. Acknowledgements	11
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 – Channel Allocation for a Dense Area	5
Figure 2 – Neural Network Approach with Gradient Decent (GD) and Back Propagation	7
Figure 3 – Estimating the Co-Usage, Sensing, and Potential Pain Matrices.....	9

List of Tables

Title	Page Number
Table 1 – Experimental Results	10

1. Introduction

In dense neighborhoods, there are often dozens of homes in close proximity. This can either be a tight city-block with many single-family homes (SFHs), or a multiple dwelling units (MDU) complex (like a big apartment building or condominium). Each home in such a neighborhood (either a SFH or a single unit in a MDU complex) has its own Wi-Fi access point (AP). Because there are few (typically 2 or 3) non-overlapping radio channels for Wi-Fi, neighboring homes may find themselves sharing a channel and competing over airtime, which may cause bad experience of slow internet (long latency, buffering while streaming movies, *etc.*). Existing APs sometimes have smart channel selection features, but because they work independently (the APs do not coordinate), this can cause a cascade of neighboring APs constantly switching channels, which is disruptive to the connectivity of the homes. Wi-Fi optimization over all the APs in a dense neighborhood is highly desired to provide the best user experience.

We present a method for Wi-Fi channel selection in a *centralized* way for all the APs in a dense neighborhood. We describe how to use recent observations to estimate the potential-pain matrix: for each pair of APs, how much Wi-Fi-pain would they cause each other if they were on the same channel. We formulate an optimization problem – finding a channel allocation (which channel each home should use) that minimizes the total Wi-Fi-pain in the neighborhood. We design an optimization algorithm that uses gradient descent over a neural network to solve the optimization problem. We describe initial results from offline experiments comparing our optimization solver to an off-the-shelf Mixed-Integer-Programming solver. In our experiments we show that the off-the-shelf solver manages to find a better (lower total pain) solution on the train data (from the recent days), but our neural-network solver *generalizes* better – it finds a solution that achieves lower total pain for the test data (“tomorrow”).

2. Wi-Fi Pain Metric

To measure the “pain” caused to the users in a dense Wi-Fi space, we define a new Wi-Fi Pain Metric. The main cause for Wi-Fi density pain is when a home’s neighbors are using the same radio channel and occupying much of its airtime: when my home’s AP senses high interference (“others” are using the channel), my home’s devices (including my AP) will have to wait longer times before they can send their packets over the radio channel, and this will cause the experience of slowness.

However, if my home barely has internet traffic during the night, while my neighbors use the same Wi-Fi channel heavily, the interference at night doesn’t cause me any pain. The pain comes when my neighbors use the channel heavily *while* my home tries to use the same channel. In addition, my home may have a lot of internet traffic at the same time as another home in my apartment building, but because there are five floors separating the two homes, our Wi-Fi signals never interfere with each other (the homes cannot “sense” each other – we will define this more formally later).

To simplify, we notice that in a dense neighborhood, homes cause each other Wi-Fi pain when three conditions are met: the homes can sense each other, they tend to have a lot of internet traffic at the same times, and they use the same radio channel. The first two are regarded as given conditions of the neighborhood (we can measure or estimate them, but we cannot control them)

and the third is the aspect that we can control – which channel does each home use. We treat these three components as independent. Let's now formalize the overall pain mathematically with these three components, for a neighborhood with n homes and n_c Wi-Fi channels:

- The **(binary) sensing** matrix, $S^b \in \{0,1\}^{n \times n}$. $S^b_{i,j}$ is 1 iff home i can sense (and be interfered by) home j .
- The **co-usage matrix** $U \in \mathbb{R}_+^{n \times n}$. This describes how much homes tend to have internet traffic at the same time. Notice, it doesn't matter which channel each home is using, and it doesn't matter if the homes can "sense" each other. This component only cares about the behavior patterns of the homes' residents and devices (specifically, the internet-activity patterns).
- The **channel allocation matrix**: $C \in \{0,1\}^{n \times n_c}$. For each home (row) which channel is assigned to it – exactly one channel (out of the n_c options) has a value of 1. Typically, n_c is 2 or 3.

The pain that home j causes to home i depends on the three conditions we mentioned – this is expressed with multiplication:

$$\sum_{c=1}^{n_c} S^b_{i,j} U_{i,j} C_{i,c} C_{j,c}.$$

Notice, that we use matrix C twice in the formula and inside a sum over the possible channels (c) – this is to capture if the two homes are using the same channel: if the two homes are not on the same channel, the whole sum will be 0, but if they are on the same channel, the sum will have a single non-zero element $S^b_{i,j} U_{i,j}$. Similarly, if the two homes don't even sense each other ($S^b_{i,j} = 0$), the whole sum will be 0 (even if they are using the same channel) – this can describe two homes that are physically far away from each other in the neighborhood, or have many walls between them, so the radio signal doesn't travel from one to the other. We assume additivity: the pain that home i senses from the neighborhood is the sum of the pain that it senses from all the neighborhood's homes:

$$pain_i = \sum_{c=1}^{n_c} \sum_{j=1}^n S^b_{i,j} U_{i,j} C_{i,c} C_{j,c}.$$

To simplify the formula, we combine the two components that we cannot control and define the **potential-pain matrix** $P = S^b \circ U$ (elementwise multiplication). $P_{i,j} = S^b_{i,j} U_{i,j}$ describes the pain that home j would add to home i if they were using the same channel. The total pain in the neighborhood is a sum over the homes:

$$pain^{total} = \sum_{c=1}^{n_c} \sum_{i,j=1}^n S^b_{i,j} U_{i,j} C_{i,c} C_{j,c} = \sum_{c=1}^{n_c} \sum_{i,j=1}^n P_{i,j} C_{i,c} C_{j,c}$$

And we can express it in matrix form:

$$pain^{total} = \sum_{c=1}^{n_c} [C^T P C]_{c,c} = Tr(C^T P C)$$




Figure 1 – Channel Allocation for a Dense Area

Figure 1 illustrates part of a made-up dense neighborhood (right image) – a floor plan with 8 apartments in an apartment building, and the potential-pain matrix for the 8 homes (left image), where darker shades of gray represent higher potential-pain value. The floor plan in the figure has two colors to the APs in the homes, representing a possible channel-allocation to two channels (blue and green).

Homes 101 and 104 are far away from each other (see the floor plan), so their APs never sense each other – this explains why they have a blank (0) value in the matrix – they have 0 potential to cause each other pain. This also explains why a smart channel allocation may allocate the same channel (green) to these two homes.

Home 106 represents a heavy internet user (most of the day has a lot of traffic), so it has the potential to cause much pain (darker shade in the matrix) to the homes that can sense it and typically have internet traffic at the same times (103, 104, 105). Homes 101 and 102 can sense home 106, but they may have internet traffic at different times of the day than home 106, so they have lower potential pain from 106 (medium gray shade). It makes sense to put home 106 on the blue channel and isolate it from homes 103, 104, and 105 (allocated the green channel).

3. Optimization Problem and Solvers

We can now define the main optimization problem as follows:

$$C^* = \arg \min_{C \in \{0,1\}^{n \times n_c}} Tr(C^T P C)$$

$$s. t.$$

$$\forall i \in \{1 \dots n\}: \sum_{c=1}^{n_c} C_{i,c} = 1$$

This problem assumes we know (or estimate from recent data) the potential-pain matrix P – it is the conditions of the neighborhood, the potential of homes to cause Wi-Fi pain to one another.

The task of the optimization is to select a good combination of per-home channels, to minimize the overall pain that the homes cause each other. One of the reasons for this *centralized* channel selection approach is to avoid too many channel changes – frequent changes can be disruptive to the users’ connectivity experience. So, a typical use would be to solve this optimization problem, set the selected channels to all the neighborhood’s APs, and keep the channels fixed for a while (e.g., a whole day, a whole week).

3.1. MIQP Problem Solver

We note that our optimization problem is a mixed-integer quadratic programming (MIQP) problem: the search parameter C appears in the objective function (the formula for total pain) in a quadratic form, and its values are constrained to be integers. This is a non-convex problem, and we don’t have an algorithm that can guarantee finding the global optimum (the very best combination of per-home channels) in reasonable time.

There are commercially available solvers, like Gurobi (Gurobi Optimization, 2022), that use a branch-and-bound approach to solve mixed integer programming problems (including the quadratic type). These methods iteratively try to rule out parts of the parameter-space and narrow down where we can find the global optimum, as well as narrow down the gap between lower and upper bounds for the optimal objective value. These tools often manage to reach the global optimum and they employ various heuristics to try to speed up the process.

3.2. Neural Network Gradient Descent

We propose an alternative method to solve the optimization problem. We construct a neural network model to calculate a soft-approximation of the neighborhood’s total pain, given any combination of channel allocation, and use gradient descent with back-propagation to change the underlying parameters until the pain reduces to a local minimum. The model is illustrated in Figure 2.

The model’s parameters are represented as a matrix $W \in \mathbb{R}^{n \times n_c}$. The input to the model is a dummy scalar variable $\beta \in \mathbb{R}_+$. Using W and β , the model calculates a “soft” version of channel allocation $C^{\beta, W} \in [0, 1]^{n \times n_c}$ by using the softmax operation on each row of βW :

$$C_{i,c}^{\beta, W} = \frac{e^{\beta W_{i,c}}}{\sum_{d=1}^{n_c} e^{\beta W_{i,d}}}.$$

The resulting matrix $C^{\beta, W}$ has each row (for home i) describing a probability distribution over the n_c optional channels. This is not a valid channel allocation (in practice each AP only uses a single channel at a time), but this is a soft approximation of a valid channel allocation.

The model then incorporates the potential pain matrix P as a fixed given input and uses $C^{\beta, W}$ to calculate a soft approximation of the total pain:

$$\text{pain}^{\beta, W} = \text{Tr}(C^{\beta, W^T} P C^{\beta, W}).$$

Notice, that the input variable β controls the order of the approximation: with a small value, like $\beta = 0.1$ the soft channel allocations in $C^{\beta,W}$ will be closer to a uniform distribution over the n_c channels. With a higher value, like $\beta = 100$, the soft channel allocations better approximate a valid channel allocation – where for each home only a single channel gets a value close to 1 and the other channels get a value close to 0.

To solve the optimization problem, we start by randomly initializing the parameters W (e.g., using an *i.i.d.* standard normal distribution), and then use gradient descent (with back-propagation) to reduce the approximated total pain $pain^{\beta,W}$. In addition, we start by using a small value of β as input, and slowly increase it. This helps the algorithm first find a good global area and only later fine tune the parameters to a local minimum. After this procedure converges to a local minimum, and the parameters are tuned to values W^{end} , we can get the solution (the chosen channel allocation) by looking at the approximated channel allocations (for large β) and thresholding their values:

$$C_{i,c}^{end} = 1 [C_{i,c}^{1000,W^{end}} > 0.5].$$

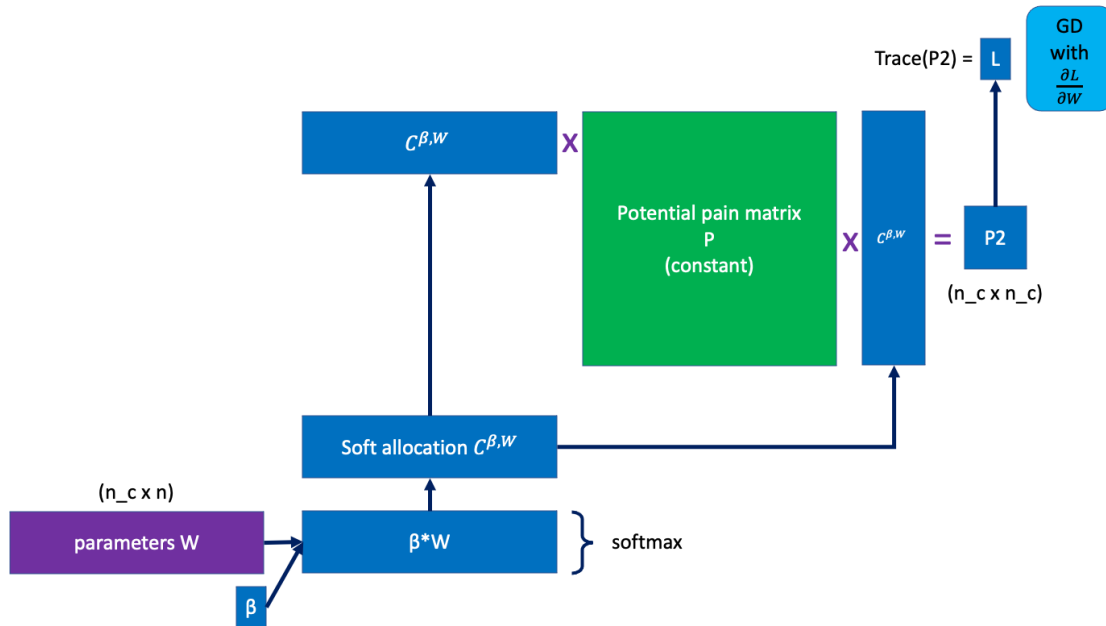


Figure 2 – Neural Network Approach with Gradient Decent (GD) and Back Propagation

While this gradient descent approach does not presume to find a better (lower) optimum than off-the-shelf solvers, we want to highlight a few advantages it has:

- This approach doesn't assume that the potential pain matrix P is symmetric, while other methods may rely on convex relaxations of the optimization problem, requiring them to have a symmetric matrix for the quadratic form.

- This approach can be modified to solve a different optimization problem that tries to minimize the *worst-home-pain* instead of the total, or *average-home-pain*. By making slight changes to the neural network, it can approximate the pain of the worst suffering home, and the optimization will try to minimize that value.
- This approach runs efficiently, quickly reaching a local minimum.
- This approach does not “try too much” to get to the global minimum. We want to generalize to near-future data, so we should avoid overfitting to the most recent days’ data.

4. Preliminary Experiments

During summer 2021, we conducted a few offline experiments with data from a big apartment building. We had data from 66 homes in the building, so we treated them as “the neighborhood’s homes” for the experiment. We tried various combinations of different aspects, and we share here some of our preliminary experiments and results. In these experiments, we simulated running the optimization on a reference date, to select the channel allocation for the following day. We collected data from the homes in the neighborhood from the recent days up to (and including) the reference day (the “train days”), calculated the potential pain matrix, and solved the channel allocation problem. We did a similar calculation to get the potential pain matrix for the day following the reference day (the “test day”). We evaluated the total pain on both the train days and the test day, given the chosen channel allocation, keeping in mind that the real goal is to improve (minimize) the pain on the *test day*.

4.1. Estimating Potential Pain

We estimated the two components of the potential pain separately: the (binary) sensing matrix S^b and the internet co-usage matrix U . Figure 3 illustrates this process: the colors in the matrices represent the cell values, ranging from 0 (dark blue) to high values (bright yellow). Each matrix has a different range (see the color-bar to the right of each image).

The co-usage can be defined as some version of multiplying two home’s internet-traffic time-series ($u_{t,i}$ represents home i ’s usage at time t). In this paper we use $U_{i,j} = \log(1 + \sum_t u_{t,i} u_{t,j})$, but we can have many variations: sum each home’s time-series first and then multiply, use a different non-linearity than logarithmic, apply non-linearity on $u_{t,i}$ alone to produce a non-symmetric version, *etc.* To estimate the co-usage matrix U , we used periodic measurements that each AP takes every 15 minutes. Specifically, we used a measurement of percentage of airtime that the AP occupied the channel to *transmit* data to the home’s devices (the “download” direction, assumed to occupy the majority of airtime in a typical home). We smoothed the quarter-hourly measurements to hourly quantities. We experimented with both measurements from whole-days (all hours of the day) and evening-time (only using measurements from 7pm-10pm local time), but here we focus our results on evening-time. For estimation with the recent n_d days, this results in a time-series (vector) of $3n_d$ hourly values for each home. We then calculated the cross-correlation between homes (the dot product of two homes’ time-series) and took the $\log(1 + x)$ of these values.

Figure 3 top row illustrates the process of estimating the co-usage matrix: starting with a Wi-Fi usage time-series for each home (top left). The image shows 10 homes and airtime-percentage values from 96 time points. This narrow matrix is multiplied by its transpose to produce the usage correlation matrix (for each pair of homes the value is the dot product of their two time-series). These correlation values can be extremely large (notice the color-bar reaching values of 200k), so we then apply logarithmic compression to form the co-usage matrix U .

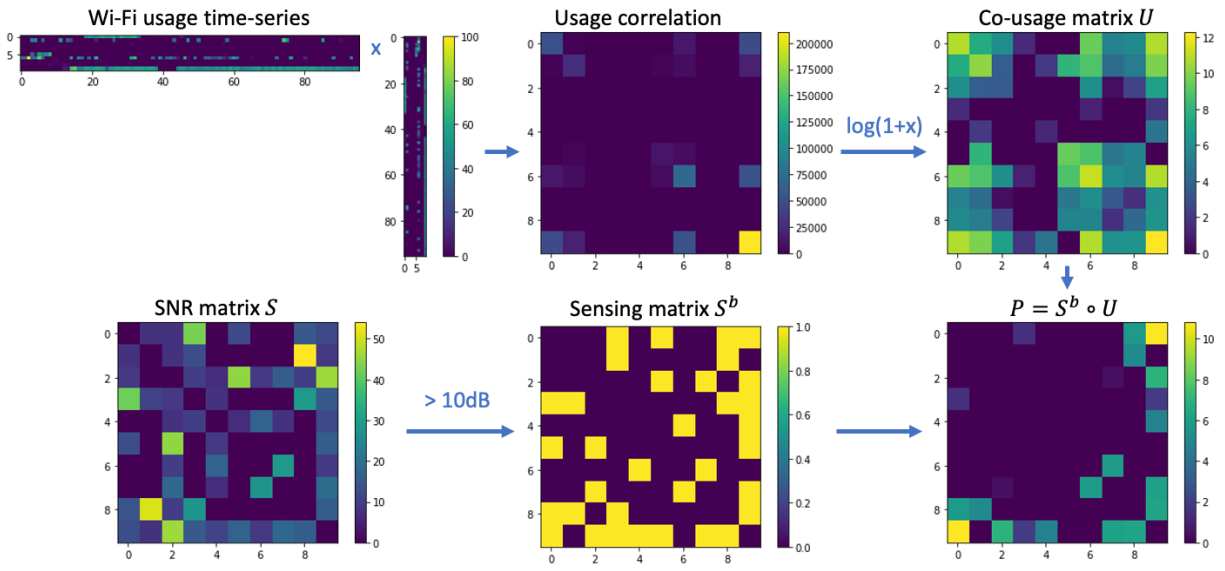


Figure 3 – Estimating the Co-Usage, Sensing, and Potential Pain Matrices

For estimating the sensing matrix S^b , we used radio-scan reports from the APs in the neighborhood: each AP performs a scan multiple times a day to look for Wi-Fi beacons in the air. The AP records the media access control (MAC) address of every other AP that it senses, and the signal to noise ratio (SNR) of the sensed beacon. We mapped sensed Wi-Fi MAC addresses to the familiar APs that are part of the neighborhood. The scans reported additional sensed entities that came from external APs (which we don't know and cannot control). For each pair of homes $\langle i, j \rangle$ in the neighborhood, we averaged the SNR values (over a period, like a week) of how strongly home i 's AP senses home j 's AP. We can call these variables the SNR matrix S (typically having non-negative real values), illustrated in Figure 3 bottom left image for 10 homes. In our experiments, we chose to symmetrize the sensing matrix: $S \leftarrow 0.5(S + S^T)$. We applied a threshold of 10dB to produce the binary sensing matrix S^b (Figure 3, bottom middle image). Notice that since an AP never "sensed itself" in the radio scans, we naturally get zeros in the diagonal. This fits our formulation, because we wish to only model the pain that homes cause other homes, not themselves.

We multiplied (elementwise) these two estimated matrices U and S^b to form the potential pain matrix P (bottom right image in Figure 3).

For the test day, we calculated the co-usage matrix U from the test day's usage measurements. However, we used the same SNR matrix S as we did for the train days. This is because we didn't

have sufficient scan measurements from every day, and because we assumed that “who can sense whom” stayed stationary over a longer time (~a month).

4.2. Optimization Details

We used the Gurobi package (Gurobi Optimization, 2022) as a MIQP solver. For our neural network algorithm, we implemented the network using TensorFlow (Martín Abadi, 2015) and Keras (Chollet, 2015). Every update step had just a “single example” input into the network. We increased the value of the input variable β in phases (running 6,400 update steps in each phase) with values: 1, 10, 100, 1000. We used ADAM optimizer with learning rate 0.001.

4.3. Results

Table 1 – Experimental Results

	Train days	Algorithm	Total pain – per train day	Total pain – test day
1.	1 (Aug 24)	Gurobi	58.2	194.0
2.	1 (Aug 24)	Neural Network	58.2	184.9
3.	4 (Aug 21-24)	Gurobi	64.5	166.3
4.	4 (Aug 21-24)	Neural Network	69.9	143.8

We show in Table 1 results from a few of our offline experiments with a single neighborhood. These were all done with train days up to (and including) August 24th and testing on usage data from August 25th. In these experiments, we used usage (and scan information) from the 2.4GHz frequency and we simulated solving the channel allocation for $n_c = 2$ channels. Rows 1-2 show experiments where there was only a single training day, compared to 4 training days in rows 3-4 (the table reports the average total pain per train day). The results show that when training with data from more days, we could achieve a worse (higher) total pain on the train data, but a better (lower) total pain on the test day, which is what we want to achieve. As expected, our neural network solver did not beat Gurobi’s solution on the train days. However, the neural network solver’s solution generalized better to the test day – it achieved a lower pain than Gurobi’s solution (in both the 1-train-day and 4-train-days scenarios).

5. Conclusion

We have discussed the problem of Wi-Fi airtime competition in a dense neighborhood and the need for a centralized channel selection solution. We defined a Wi-Fi pain objective, based on the co-occurrence of close neighbors having a lot of internet traffic at the same time on the same radio channel. We formulated the pain such that all the relevant information is captured in a single square matrix P , indicating for each pair of homes how much pain would one add to the other if they were using the same channel. We formulated an optimization problem and offered two alternative solvers for it: an off-the-shelf MIQP problem solver and a tailored neural network solver. We conducted preliminary offline experiments with data from a real neighborhood and demonstrated how we can achieve better generalization (lower pain for “tomorrow”) with more training days and by using our neural network solver.

5.1. Future Directions

There are still many more directions to research, including various flavors of Wi-Fi pain (non-symmetric definitions of potential-pain, weighting different days of the week, *etc.*), and adjustments to the optimization algorithm (e.g., regularization on the parameters W , schedule of changing β). An interesting direction is minimizing the worst-home pain and seeing how it influences the average-home pain. We will conduct more offline experiments with many more neighborhoods. Additionally, actual trials will reveal more reliably how helpful is centralized channel selection and A/B tests can help demonstrate which methods are better. In actual channel-selection experiments, we can more directly measure the sensed interference that every AP experiences from its environment. More importantly, we'll have to assess the effect on the residents' subjective experience of slow internet and Wi-Fi "pain".

6. Acknowledgements

We thank Zhehao (Kenny) Zhang for running the offline experiments, and Shelley Leung for processing telemetry data. We thank Obi Asinugo, Zekun (Katherine) Yang, Ali Mohammadi, Teddy ElRashidy, and Colleen Szymanik for great discussions and guidance about this interesting Wi-Fi problem.

Abbreviations

AP	access point
DFS	dynamic frequency selection
Hz	Hertz
SCTE	Society of Cable Telecommunications Engineers
MIQP	mixed integer quadratic programming
MAC	media access control
SFH	single-family home
MDU	multiple dwelling unit
SNR	signal to noise ratio

Bibliography & References

- Chollet, F. a. (2015). *Keras*. Retrieved from <https://keras.io>
- Gurobi Optimization, L. (2022). *Gurobi Optimizer Reference Manual*. Retrieved from <https://www.gurobi.com>
- Martín Abadi, A. A. (2015). *TensorFlow: Large-scale machine learning on heterogeneous systems*. Retrieved from <https://www.tensorflow.org/>

Peas In a Pod: What Makes Them Green?

A Technical Paper prepared for SCTE by

Defu Li

Distinguished Engineer
Comcast Cable
Massachusetts
+1 (267) 586-7680
Defu_Li@comcast.com

Richard Grivalsky

Senior Energy Engineer
Comcast Cable
+1 (802) 316-6553
Richard_Grivalsky@comcast.com

Robert Gaydos, Comcast Cable

Ashok Ramasamy, Comcast Cable

Eric Stonfer, Comcast Cable

Gianni DiGregorio, Comcast Cable

Table of Contents

Title	Page Number
1. Introduction.....	3
1.1. vCMTS and DAA.....	3
1.1. vBNG.....	3
1.2. PPODs	3
1.3. Carbon Neutrality and Purpose.....	4
2. Observability.....	4
2.1. Framework	4
2.2. Collector and Monitor	5
2.2.1. Collectd	5
2.2.2. IPMI Plugin.....	6
2.2.3. Turbostat Plugin	6
2.2.4. Grafana Dashboard	7
3. Hot Standby & Power Saving Mode.....	8
3.1. Hot Standby.....	8
3.2. Intel CPU Power Saving Mode.....	11
3.3. Future Considerations.....	12
4. Conclusion.....	14
Abbreviations	14
Bibliography & References.....	15

List of Figures

Title	Page Number
Figure 1 - MHA v2 [Source: CM-SP-R-PHY Specification].....	3
Figure 2 - Logical View of Access Network and Compute Nodes in a Leaf/Spine Architecture.....	4
Figure 3 - PPOD Observability Framework.....	5
Figure 4 - Snippets of Prometheus Configmap and Collectd Pod Spec ¹	6
Figure 5 - IPMI Metrics.....	6
Figure 6 - Screenshot of Turbostats Output ¹	7
Figure 7 - Power Consumption Dashboard ¹	7
Figure 8 - Dual-Redundant vs. Hot Standby Line Drawing.....	8
Figure 9 – Server Rear Elevation Dual-Redundant vs. Hot Standby Watts Consumed.....	10
Figure 10 - Power Consumption for Host with C-State Disabled ¹	11
Figure 11 - Power Consumption of a Host with C-State Enabled ¹	12
Figure 12 - Workload Utilization Percentage for a PPOD for One Week ¹	13
Figure 13 - Stacked Workload Utilization Percentage for a PPOD for One Week ¹	13

List of Tables

Title	Page Number
Table 1 - Hot Standby Enabled Then Disabled Data Segment	9
Table 2 - B-Side Breakers Closed	10
Table 3 - B-Side Breakers Open	10

1. Introduction

1.1. vCMTS and DAA

The distributed access architecture (DAA) specification, or modular head-end architecture version 2 (MHAv2), was introduced to address cable headend space and power limitations. The traditional integrated CMTS (iCMTS) or cable converged access platform (CCAP) functions were split into two: the physical (PHY) function, and the core function. The remote PHY device (RPD) provides the PHY function, while the core functions consist of CMTS and CCAP operating on the MAC or IP layers.

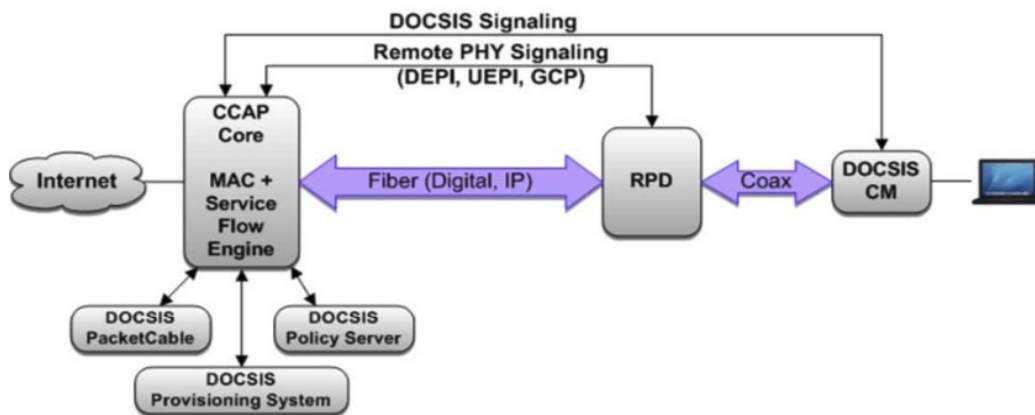


Figure 1 - MHAv2 [Source: CM-SP-R-PHY Specification]

The split allows the core functions to run on a cloud computing platform. The virtualized CMTS/CCAP (vCMTS/vCCAP) is a collection of software applications, built upon the microservice architecture pattern and targeted for cloud computing platforms. Comcast Cable has built its own private cloud in order to host these vCMTS software applications.

1.1. vBNG

CableLabs' DOCSIS Provisioning of EPON (DPoE) specification enables an operator to deploy EPON technology using the existing DOCSIS based backend systems. This specification allows an optical network unit (ONU) to be emulated as virtual cable modem (vCM).

The Comcast Private Cloud can host the virtual broadband gateway (vBNG) application which supports EPON technology. Like DPoE, vBNG emulates an ONU in order to utilize DOCSIS based network device provisioning backend systems.

1.2. PPODs

A typical private cloud is likely to consist of many server racks. Like peas in a pod (PPOD), these server racks are built identically providing operational efficiency and easy scalability. There are thousands of 'peas' (servers) spread across hundreds of PODs in an operator's network.

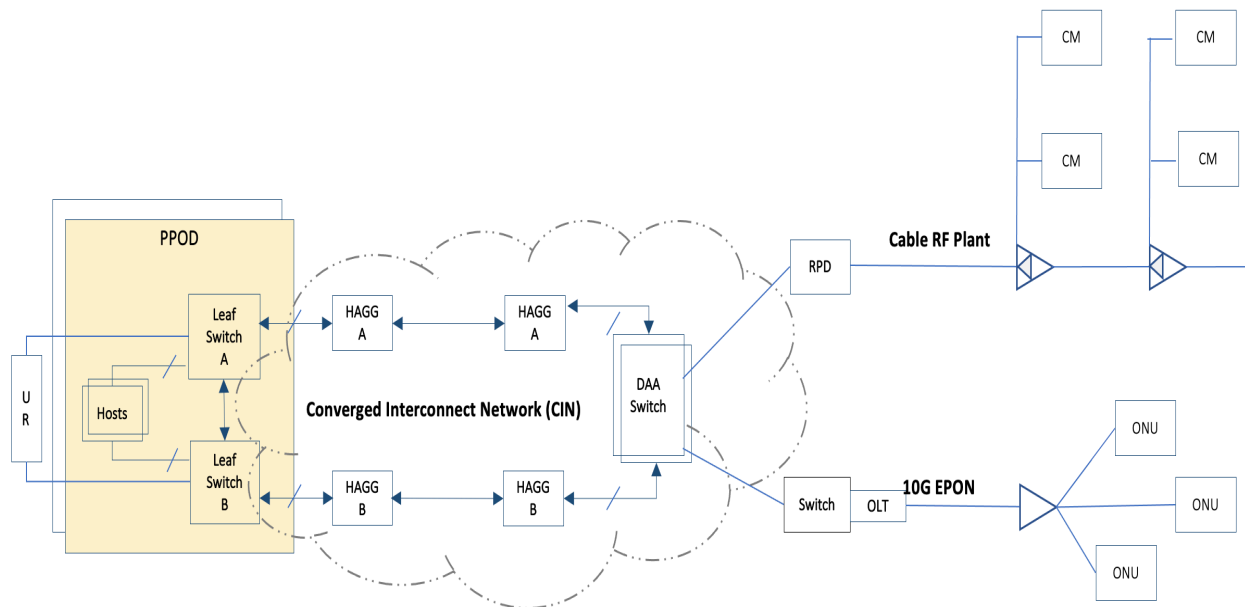


Figure 2 - Logical View of Access Network and Compute Nodes in a Leaf/Spine Architecture

The PPODs are deployed across hundreds of sites nationwide. Each PPOD contains several servers which form a compute cluster, each server has dual ethernet ports connecting to a pair of leaf switches. vCMTS and vBNG workloads are deployed and replicated on these PPODs OLT and RPD traffic is tunneled to or from the vBNG and vCMTS via the converged interconnect network (CIN) via a leaf-spine switch fabric. Upstream traffic to or from the internet is routed via upstream routers (URs)

1.3. Carbon Neutrality and Purpose

Looking ahead to Comcast's commitment to being carbon neutral by 2035, the question becomes, what can we do to “green-up” our PPODs, make them more energy efficient, and in the process reduce our operational expenditures?

In this paper, we will discuss what is involved to provide energy consumption observability for the Comcast Private Cloud. We will discuss the techniques that we employ to provide immediate energy saving as well as more advance techniques based on load and demand characteristics of our containerized network function (CNF) workloads. This paper will conclude with the lessons learned and future strategy for energy efficiency looking beyond the PPODs, in the wider Comcast ecosystem.

2. Observability

2.1. Framework

Since the start real time observability has been crucial for the Comcast Cable Private Cloud, as such we have built a stack based upon on Elasticsearch, Logstash, and Kibana (ELK) stack and Prometheus, a time series database (TSDB).

At a high level, the Prometheus server periodically scrapes metrics from targets in the PPOD. The long-term metrics are pushed to S3 ThanosStore, with Grafana dashboards providing a human friendly interface. ThanosQuery provides a distributed query engine for short-term locally cached metrics.

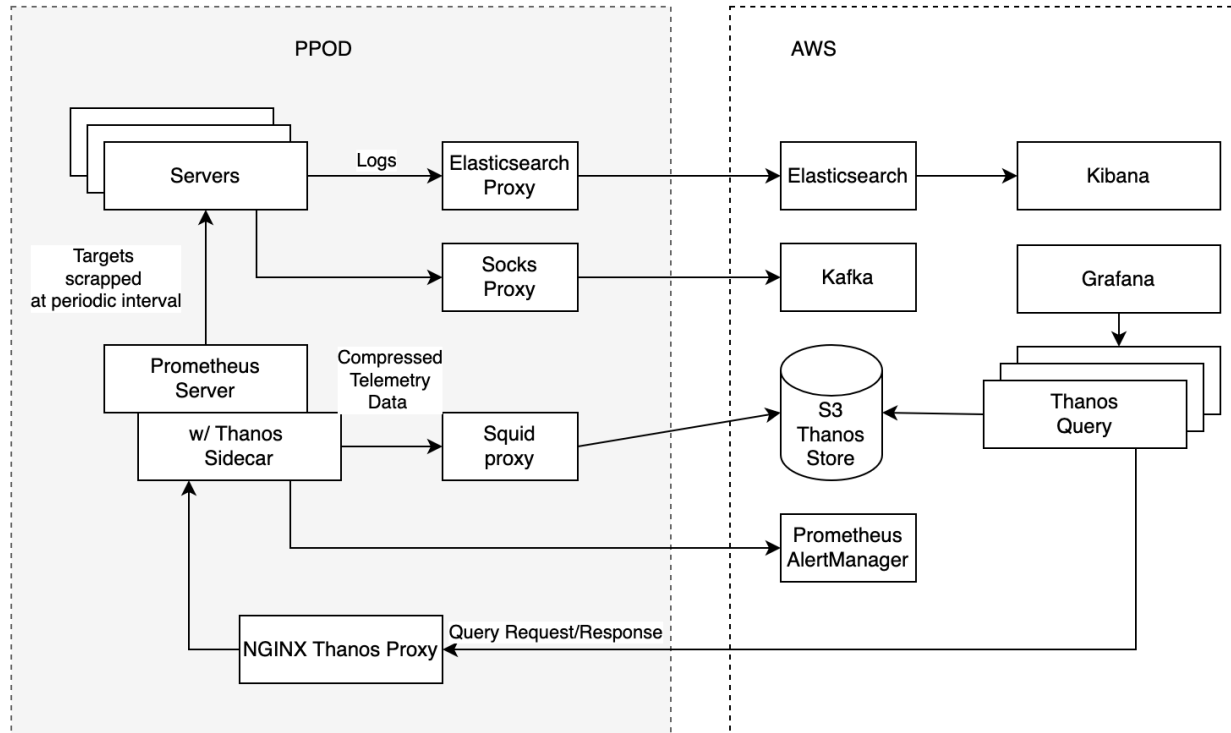


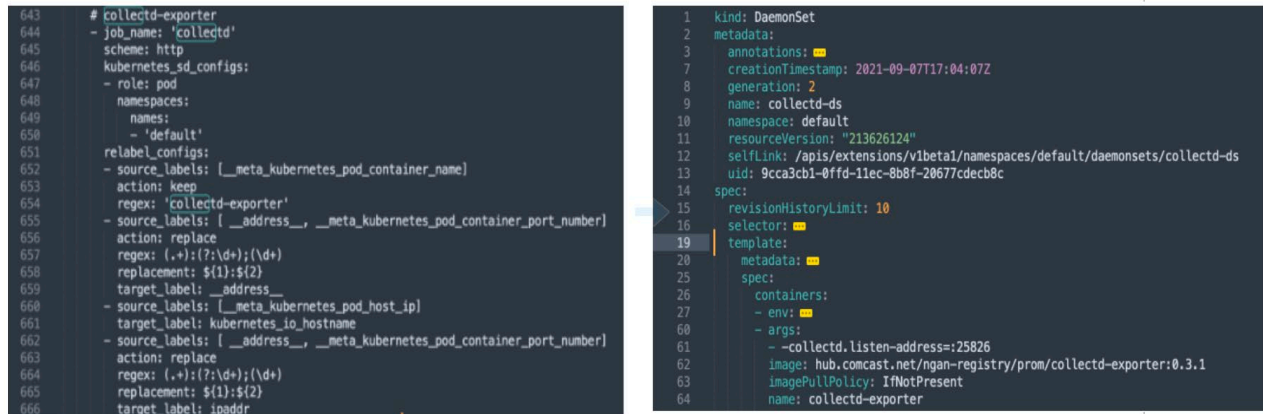
Figure 3 - PPOD Observability Framework

In the subsections which follow, we describe the components which provide power consumption metrics.

2.2. Collector and Monitor

2.2.1. Collectd

Collectd is a Linux daemon that collects, stores and transfers performance metrics on a per host level. Collectd is deployed as a DaemonSet for all hosts, in each PPOD. Prometheus in turn scrapes the metrics provided by Collectd.



```

643 # collectd-exporter
644 - job_name: 'collectd'
645   scheme: http
646   kubernetes_sd_configs:
647   - role: pod
648     namespaces:
649     - 'default'
650   relabel_configs:
651   - source_labels: [__meta_kubernetes_pod_container_name]
652     action: keep
653   - source_labels: [__address__, __meta_kubernetes_pod_container_port_number]
654     action: replace
655     regex: '(.+):(?:\d+);(\d+)'
656     replacement: ${1}:${2}
657     target_label: __address__
658   - source_labels: [__meta_kubernetes_pod_host_ip]
659     target_label: kubernetes_io_hostname
660   - source_labels: [__address__, __meta_kubernetes_pod_container_port_number]
661     action: replace
662     regex: '(.+):(?:\d+);(\d+)'
663     replacement: ${1}:${2}
664     target_label: ipaddr
  
```

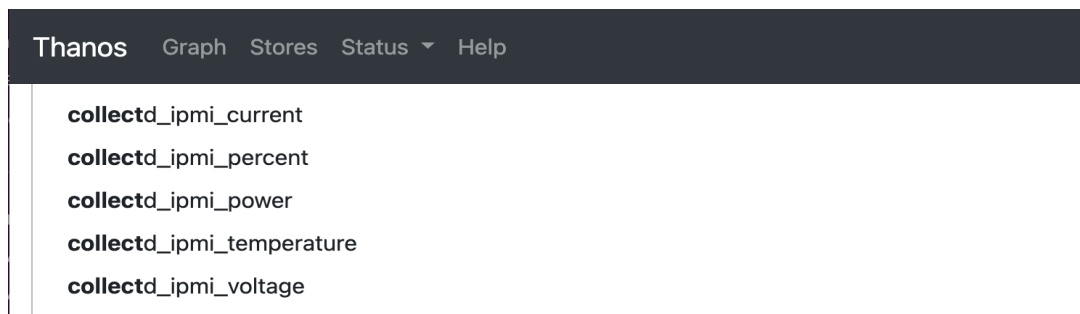
```

1 kind: DaemonSet
2 metadata:
3   annotations:
4     creationTimestamp: 2021-09-07T17:04:07Z
5     generation: 2
6   name: collectd-ds
7   namespace: default
8   resourceVersion: "213626124"
9   selfLink: /apis/extensions/v1beta1/namespaces/default/daemonsets/collectd-ds
10  uid: 9cca3cb1-0ffd-11ec-8b8f-20677cdec8c
11 spec:
12   revisionHistoryLimit: 10
13   selector:
14     template:
15     metadata:
16     spec:
17     containers:
18     - env:
19     - args:
20     - --collectd.listen-address=:25826
21     image: hub.comcast.net/ngan-registry/prom/collectd-exporter:0.3.1
22     imagePullPolicy: IfNotPresent
23     name: collectd-exporter
  
```

Figure 4 - Snippets of Prometheus Configmap and Collectd Pod Spec¹

2.2.2. IPMI Plugin

Collectd supports numerous loadable plugins. The Intelligent Platform Management Interface (IPMI) plugin uses the OpenIPMI library to read hardware sensors on the host in order to provide power consumption metrics as shown in the figure below.



Thanos Graph Stores Status ▾ Help

- collectd_ipmi_current
- collectd_ipmi_percent
- collectd_ipmi_power
- collectd_ipmi_temperature
- collectd_ipmi_voltage

Figure 5 - IPMI Metrics

2.2.3. Turbostat Plugin

Turbostat is a Linux tool that reports processor frequency and statistics. The Turbostats Plugin utilizes Turbostats for reporting processor performance metrics.

Package	Core	CPU	Avg_MHz	Busy%	Bzy_MHz	TSC_MHz	IRQ	SMT	POLL	C1	C1E	C6	POLL%	C1%	C1E%	C6%	CPU/c1	CPU/c6	CoreTemp	PkgTemp	PkgNot		
t	-	-	RAWMott	PKG_%	RAW_%																		
00	0	0	0	0.00	63	3.02	2100	2894	108432	0	2280	123551	0	0	0.00	96.95	0.00	0.00	96.98	0.00	53	53	61.650
0	0	0	0	0.00	69	3.30	2100	2895	2969	0	177	3622	0	0	0.01	96.74	0.00	0.00	96.70	0.00	50	52	32.100
0	0	0	0	0.00	16	53	2.55	2100	2895	3551	0	20	4062	0	0	0.00	97.48	0.00	0.00	97.45			
0	1	1	67	3.18	2100	2895	3286	0	253	3972	0	0	0.01	96.86	0.00	0.00	96.82	0.00	49				
0	1	17	59	2.82	2100	2895	3983	0	3	4377	0	0	0.00	97.23	0.00	0.00	97.18						
0	2	2	43	2.04	2100	2895	3330	0	12	3583	0	0	0.00	98.00	0.00	0.00	97.96	0.00	49				
0	2	18	146	6.96	2100	2895	4254	0	7	4615	0	0	0.00	93.08	0.00	0.00	93.04						
0	3	3	75	3.58	2100	2895	3167	0	17	3538	0	0	0.00	96.46	0.00	0.00	96.42	0.00	49				
0	3	19	56	2.69	2100	2895	5258	0	883	5577	0	0	0.03	97.36	0.00	0.00	97.31						
0	4	4	76	3.62	2100	2895	3110	0	2	3456	0	0	0.00	96.42	0.00	0.00	96.38	0.00	50				
0	4	20	66	3.15	2100	2895	4401	0	236	4802	0	0	0.01	96.90	0.00	0.00	96.85						
0	5	5	153	7.29	2100	2895	2910	0	3	3236	0	0	0.00	92.75	0.00	0.00	92.71	0.00	49				
0	5	21	65	3.09	2100	2895	2869	0	24	3445	0	0	0.00	96.94	0.00	0.00	96.91						
0	6	6	42	2.01	2100	2895	3797	0	119	4170	0	0	0.01	98.04	0.00	0.00	97.99	0.00	49				
0	6	22	49	2.32	2100	2895	3112	0	213	3944	0	0	0.01	97.72	0.00	0.00	97.68						
0	7	7	98	4.67	2100	2895	3233	0	118	3634	0	0	0.01	95.37	0.00	0.00	95.33	0.00	50				
0	7	23	67	3.21	2100	2895	3271	0	9	3838	0	0	0.00	96.83	0.00	0.00	96.79						
1	0	8	45	2.16	2100	2895	2884	0	9	3151	0	0	0.00	97.88	0.00	0.00	97.84	0.00	52	53	29.620		
00	0	0	0	0.00	0.00	0.00	0.00	0	0	0	0	0	0.00	96.32	0.00	0.00	96.28						
1	0	24	78	3.72	2100	2895	3619	0	13	3754	0	0	0.00	96.32	0.00	0.00	96.28						
1	1	9	58	2.79	2100	2895	3326	0	27	3740	0	0	0.00	97.25	0.00	0.00	97.21	0.00	52				
1	1	25	74	3.52	2100	2895	2801	0	3	3271	0	0	0.00	96.51	0.00	0.00	96.48						
1	2	10	54	2.56	2100	2895	2889	0	16	3980	0	0	0.00	97.48	0.00	0.00	97.44	0.00	52				
1	2	26	46	2.18	2100	2895	3611	0	9	3887	0	0	0.00	97.86	0.00	0.00	97.82						
1	3	11	67	3.21	2100	2895	4087	0	13	4848	0	0	0.00	96.84	0.00	0.00	96.79	0.00	51				
1	3	27	39	1.87	2100	2895	2992	0	0	3317	0	0	0.00	98.16	0.00	0.00	98.13						
1	4	12	46	2.19	2100	2895	3071	0	3	3543	0	0	0.00	97.85	0.00	0.00	97.81	0.00	52				
1	4	28	45	2.13	2100	2895	2817	0	0	3140	0	0	0.00	97.90	0.00	0.00	97.87						
1	5	13	43	2.03	2100	2895	3905	0	66	4185	0	0	0.01	98.01	0.00	0.00	97.97	0.00	53				
1	5	29	50	2.37	2100	2895	3043	0	10	3774	0	0	0.00	97.67	0.00	0.00	97.63						
1	6	14	42	1.99	2100	2895	3941	0	11	4276	0	0	0.00	98.06	0.00	0.00	98.01	0.00	51				
1	6	30	44	2.12	2100	2895	2588	0	0	3553	0	0	0.00	97.92	0.00	0.00	97.88						
1	7	15	44	2.08	2100	2895	3168	0	1	3434	0	0	0.00	97.96	0.00	0.00	97.92	0.00	52				
1	7	31	71	3.40	2100	2895	3189	0	3	3827	0	0	0.00	96.64	0.00	0.00	96.60						
Package	Core	CPU	Avg_MHz	Busy%	Bzy_MHz	TSC_MHz	IRQ	SMT	POLL	C1	C1E	C6	POLL%	C1%	C1E%	C6%	CPU/c1	CPU/c6	CoreTemp	PkgTemp	PkgNot		

Figure 6 - Screenshot of Turbostats Output¹

2.2.4. Grafana Dashboard

Power consumption data visualization can be easily created and customized via a Grafana Dashboard. The figure below shows the power consumption by host and by PPOD. The charts on the left show the total power consumption by host and by PPOD. The charts on the right show the breakdown by the power supply unit by host and by PPOD.

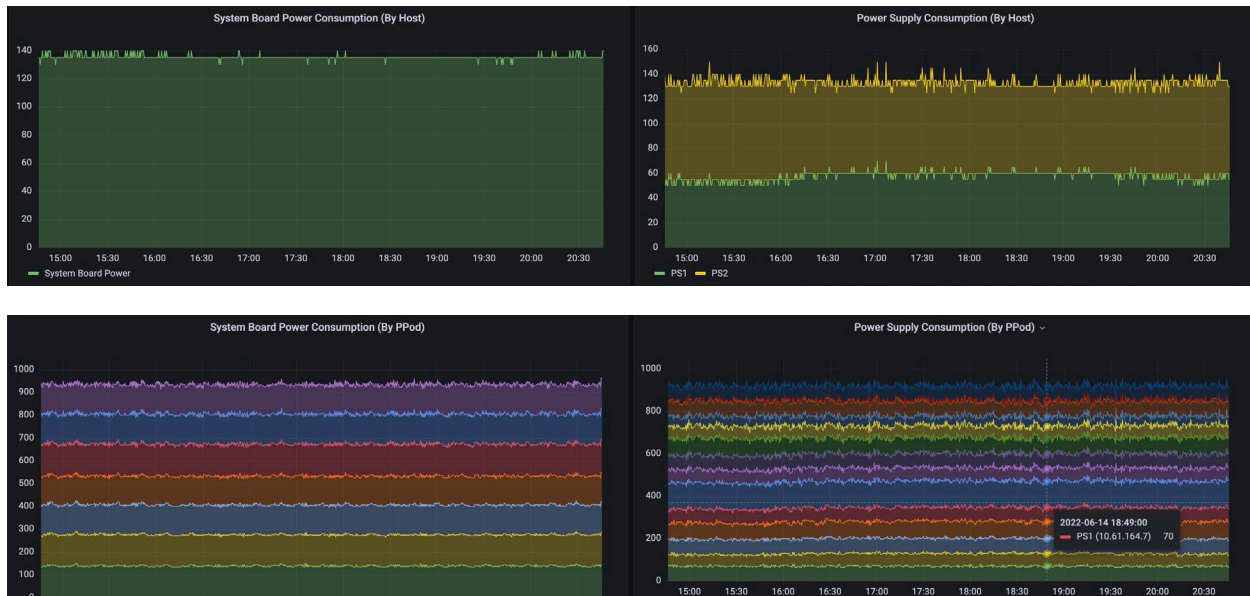


Figure 7 - Power Consumption Dashboard¹

3. Hot Standby & Power Saving Mode

3.1. Hot Standby

Power supply unit (PSU) hot standby, also referred to as hot sparing, is the ability for a single power supply to transform input to platform required voltage while keeping an idle power supply in reserve, as seen in Figure 8. The platform stages PSUs on or off dependent upon the platform's throughput and required power. The power supply in standby configuration does not transform input voltage to platform required voltages but does maintain telemetry, connection to the common buss, and is instantaneously available to support higher energy demand or in support of loss to the active-primary power supply.

It is through this idle state operators can realize a reduction in energy consumption. In dual redundant mode, the platform's required load will be split $\approx 50\%$ on each PSU. If the PSU has an 800-watt capacity and is only loaded with 100-watts (12.5%), it may not be optimally loaded. This creates transformation through two PSUs and impacts energy dissipated to transform voltage; total dissipation is dependent on the efficiency curve of the PSU and the output load. If we optimize the load, by enabling hot standby, we can improve the efficiency and reduce the number of locations voltage is transformed. By moving from dual-redundant operation at 100-watts (12.5%) of PSU0 and PSU1's capacity utilized, hot standby operation loads PSU0 at 200-watts (25% capacity).




Figure 8 - Dual-Redundant vs. Hot Standby Line Drawing

Several original equipment manufacturers (OEMs) offer this platform setting today. In initial trials, a $\sim 4.5\%$ reduction in energy consumption for the deployed solution was documented.

Load-side power distribution measurements were captured with dual-redundant power supplies active for a minimum of (7) days prior to enabling of hot standby for the platforms tested as seen in Table 1. The platforms were then set to hot standby mode and measurements captured over time (30 days) before returning to dual-redundant mode for additional measurement.

During the hot standby trial time two distinct events were observed in which platform power exceeded the threshold for a single supply, the hot standby was brought into operational state, and then returned to hot standby with no service impact observed.

Table 1 - Hot Standby Enabled Then Disabled Data Segment

Date	Bus A Volt	Bus B Volt	Circ A03 Current	Circ B03 Current	Circuit 3 Total Watts	Circ A04 Current	Circ B04 Current	Circuit 4 Total Watts
11/28/2021 20:00	53.5	53.64341	8.389999	0.21	460.1300626	7.94	0.2	435.518682
11/29/2021 0:00	53.5	53.63208	8.41	0.21	461.1977368	7.98	0.2	437.656416
11/29/2021 4:00	53.5	53.654583	8.26	0.21	453.1774624	7.91	0.21	434.4524624
11/29/2021 8:00	53.5	53.665543	8.21	0.21	450.504764	7.86	0.21	431.779764
11/29/2021 12:00	53.5	53.650021	8.25	0.21	452.6415044	7.86	0.21	431.7765044
11/29/2021 20:00	53.599998	53.500065	4.81	4.12	478.2362582	5.57	2.89	453.1671767
11/30/2021 0:00	53.599998	53.500057	4.78	4.14	477.6982264	5.59	2.94	456.9141564
11/30/2021 4:00	53.599998	53.50074	4.76	4.12	475.5590393	5.57	2.93	455.3091571

For platforms which do not currently offer hot standby configuration, measurements were captured in dual-redundant state by aggregating channel loads as seen in Table 2. Channel 12, for example, was drawing a total of two amperes in dual-redundant mode. Power was removed from the B-side load distribution by way of opening breakers in order to test the specific platform load and its effect on PSU efficiency curve as seen in Table 3. Channel 12 was now only drawing 1.7 amperes. Energy avoidance was calculated with the nominal consumption of a PSU in hot standby factored for those devices by adding 13-watts of load, as seen in Figure 9. This process was performed on two separate PPODs of varying compute load and tested for 24-72 hours before restoring B-Side power.

Table 2 - B-Side Breakers Closed

Module 2A					
SNMP Chan Mapping	Channel	Load	Ampacity	Inventory	Brkr Status
22	12	1.2A	20A	YES	ON
23	13	2.1A	20A	YES	ON
24	14	2.2A	30A	YES	ON
25	15	2.3A	30A	YES	ON
Module 2B					
SNMP Chan Mapping	Channel	Load	Ampacity	Inventory	Brkr Status
32	12	0.8A	20A	YES	ON
33	13	2.2A	30A	YES	ON
34	14	1.8A	30A	YES	ON
35	15	2.2A	30A	YES	ON

Table 3 - B-Side Breakers Open

Module 2A					
SNMP Chan Mapping	Channel	Load	Ampacity	Inventory	Brkr Status
22	12	1.7A	20A	YES	ON
23	13	4.0A	20A	YES	ON
24	14	3.7A	30A	YES	ON
25	15	4.3A	30A	YES	ON
Module 2B					
SNMP Chan Mapping	Channel	Load	Ampacity	Inventory	Brkr Status
32	12	0.0A	20A	YES	OFF
33	13	0.0A	30A	YES	OFF
34	14	0.0A	30A	YES	OFF
35	15	0.0A	30A	YES	OFF




Figure 9 – Server Rear Elevation Dual-Redundant vs. Hot Standby Watts Consumed

As DAA expands the proliferation of server-based rack architectures, hot standby presents itself as a low-impact, reliable, and sustainable practice to aid in the drive toward carbon neutrality. Given the change and impact to installation and operational practices, processes must be built to ensure load is equally distributed across AC & DC plant circuits.

3.2. Intel CPU Power Saving Mode

All servers in a PPOD are Intel CPU based. Intel processors can be controlled by the following:

- Per core C-State
- Per core P-State

The C-State is an idle power state in which the processor is not executing instruction. The P-State is for various voltage or frequency levels in which the processor is still executing instructions.

For configuring and controlling C-states, on most modern Linux platforms C-states are automatically enabled, this is done via a combination of basic input/output system (BIOS) settings and the intel_idle driver. In order to dynamically force the system to a lower C-state (more power intensive) one can open the file `/dev/cpu_dma_latency`, and write a low value, usually (5) or under to this file. The value found in `/dev/cpu_dma_latency` represents the amount of latency in microseconds allowed for C-state transitions, by forcing this to a low value this should limit the CPU to C0 during active workloads and C1 during idle. For as long as this file remains open the C-states will be forced to these lower states.

Our first step was to understand C-state, workloads, and how they influence the power consumption of a host. This is accomplished by scheduling the workloads on a host, then measuring the power consumption data with C-state enabled and disabled.




Figure 10 - Power Consumption for Host with C-State Disabled¹

Our test result shows that with C-state saving mode disabled, the power consumption level remains constant for any number of CNF workloads scheduled on to a given host.

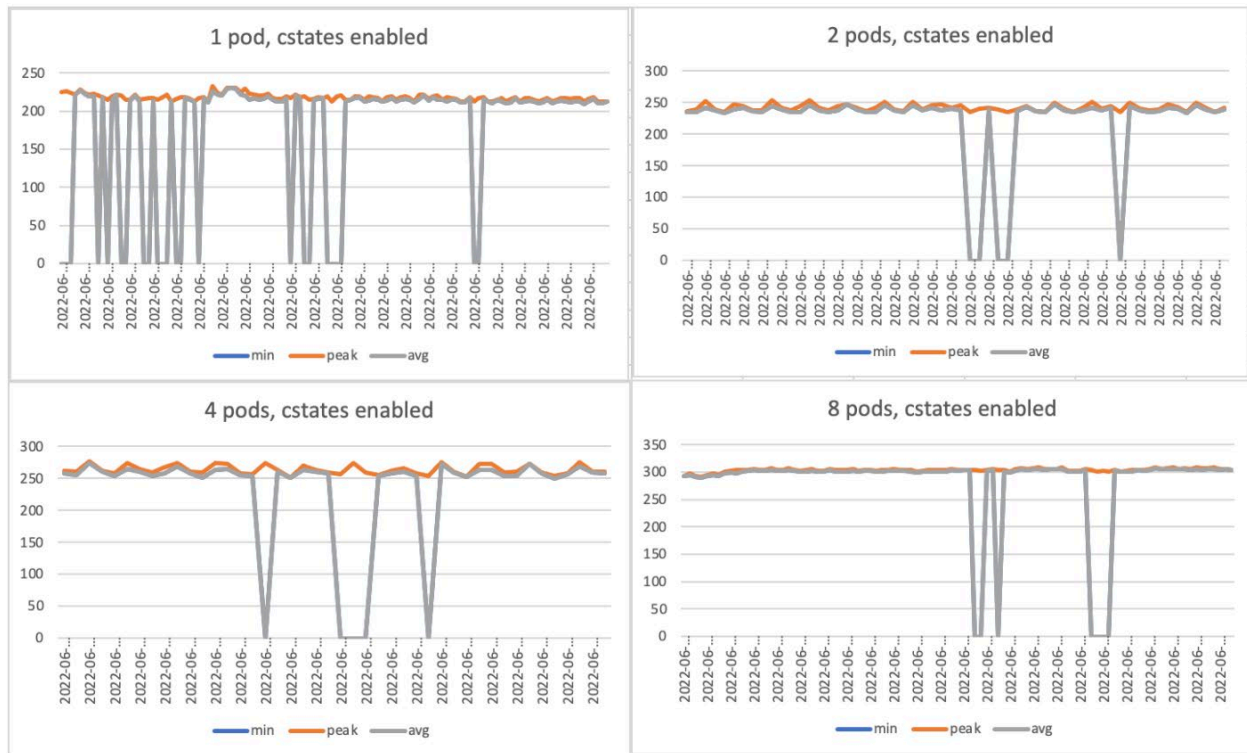


Figure 11 - Power Consumption of a Host with C-State Enabled¹

With C-state enabled, as the number of CNF workloads deployed on a given host increases, the power consumption level rises. This is shown in Figure 11.

3.3. Future Considerations

The aggregate usage can be generalized as similar demand curve each day. An example of subscriber usage aggregated across all workloads within a PPOD is shown in Figure 12 and Figure 13. The Y-axis represent the workload usage percentage normalized by the access technology capacity service by the CNF. The X-axis is day of the month.

Figure 12 plots several diverse, individual workloads usage in a PPOD. Figure 13 is a stacked version, which provides a better view of the aggregate usage curve. The pattern is similar for all PPODs across all sites.




Figure 12 - Workload Utilization Percentage for a PPOD for One Week¹




Figure 13 - Stacked Workload Utilization Percentage for a PPOD for One Week¹

A system is elastic and can adapt to workload changes by provisioning and deprovisioning resources, in order to meet demand. On the compute cluster, significant CPU core resources are isolated and dedicated to the DPDK CNF workloads. The required CPU core resources for the CNF workloads are very much traffic or network IO bandwidth driven.

In the future, we will be exploring the following approaches to match the CPU resources to the traffic demand:

- A. CPU P-State control by software application to match the short-term traffic demand
- B. Bin packing of workloads across hosts to match the longer-term traffic demand
- C. Combination of the above

Beyond the PPOD, we could also explore the overall capacity planning process in terms of spectrum activation. For example, an average 50-subscriber service group having low bandwidth demand requires far less spectrum and compute resources activated, as compared to that of an average 500-subscriber service group having much higher bandwidth demands. Can the capacity planning process be automatic, just-in-time, and elastic?

¹We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

4. Conclusion

We started with energy consumption observability for the Comcast Private Cloud. We quantified the power savings of ~4.5% by reconfiguring the PSUs to hot standby; this will be operationalized for PPODs being put into production in late-2022/early-2023. As hot standby is platform specific, integrated software-based, user-defined configuration, we will be working with OEMs to explore how many existing platforms can be integrated through software upgrades. All hosts in PPODs are already provisioned with C-State power saving mode enabled.

The measured power consumption metrics for host with CPU C-State setting and various CNF workloads provide us the insight into the potential savings. The future strategy for energy efficiency is very much aligned with our cloud native architecture evolution; meaning it is just-in-time and elastic to workload changes by automatically provisioning and deprovisioning resources, such that the available resources match the demand.

Abbreviations

AC	Alternating current
BIOS	Basic input/output system
CCAP	Cable converged access platform
CIN	Converged Interconnect Network
CNF	Containerized network function
CPU	Central processing unit
CM	Cable modem
DAA	Distributed access architecture
DC	Direct current
DOCSIS	Data over cable interface specification
DPDK	Data plane development kit
DPoE	DOCSIS provisioning of EPON
EPON	Ethernet passive optical network
ELK	Elasticsearch, Logstash, and Kibana stack
iCMTS	Integrated cable modem termination system
HAGG	Headend aggregation switch
IP	Internet protocol
IPMI	Intelligent platform management interface
MAC	Media access control
MHAv2	Modular headend architecture version 2
OEM	Original equipment manufacturer
OLT	Optical line termination
ONU	Optical network unit
PHY	Physical
PPOD	Physical pod
PSU	Power supply unit
RF	Radio Frequency
RPD	Remote physical device
TSDB	Time Series Database
UR	Upstream router

vBNG	Virtual broadband gateway
vCM	Virtual cable modem
vCMTS	Virtual cable modem termination system

Bibliography & References

Dell PowerEdge Manuals

HPE ProLiant Manuals

Cable Labs R-PHY Specification

Perceptual Video Coding Optimization Techniques: Most Recent Trends and Future Directions

A Technical Paper prepared for SCTE by

Dan Grois, PhD

Principal Researcher

Comcast

Comcast Center 1701 JFK Blvd. Philadelphia, PA 19103

1-215-286-1700

dan_grois@comcast.com

Alex Giladi

Fellow

Comcast

Comcast Center 1701 JFK Blvd. Philadelphia, PA 19103

1-215-286-1700

alex_giladi@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background: Human Visual System.....	4
3. Perceptual Video Quantization Framework	5
3.1. Background: Contrast Sensitivity Function	5
3.2. Perceptual Quantization Matrices for UltraHD Resolution Displays	7
3.3. Perceptual Quantization Matrices for Mobile Device Displays	10
4. Perceptual Video Masking Framework	11
4.1. Background: Visual Masking.....	12
4.2. Forward and Backward Masking Encoding Scheme	12
4.2.1. Experimental Results and Brief Discussion	13
5. Future Directions for Perceptual Video Coding Optimizations.....	15
6. Conclusion.....	17
Abbreviations	18
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1 – The schematic block diagram of the H.265/MPEG-HEVC encoder.	3
Figure 2 – The default perceptual quantization matrices defined in the HEVC standard.....	5
Figure 3 – Upsampling the default 8×8 HEVC matrix for obtaining default matrices for 16×16 and 32×32 transform block sizes.	6
Figure 4 – A sample frames from the tested sequences: (a) “Lucy”; (b) “Everest”; (c) “Warcraft”; (d) “Regatta”.	8
Figure 5 – A schematic illustration of the proposed joint backward and forward temporal masking framework.....	13
Figure 6 – The schematic block diagram of the H.266/MPEG-VVC encoder.....	15
Figure 7 – A multi-resolution encoding framework for enabling efficient sharing of analysis information across representations.	17

List of Tables

Title	Page Number
Table 1 - HEVC Transform Block Type/Size-Dependent Quantization Matrices.....	6
Table 2 - HDR UltraHD test video sequences.	8
Table 3 - BD-BR PSNR and SSIMPlus bit-rate savings for the HEVC encoding.	9
Table 4 - BD-BR PSNR and SSIMPlus bit-rate savings for the HEVC encoding.	9
Table 5 - SSIMPlus scores for encoding the Regatta video sequence.	10
Table 6 - BD-BR SSIMPlus bit-rate savings for the HEVC encoding.	10
Table 7 - SSIMPlus scores for encoding the Regatta video sequence.	11
Table 8 – Test Sequences.	13
Table 9 – Bitrate Savings in Terms of BD-BR.....	14

1. Introduction

There is currently a strong demand for high resolution video content, particularly for the high-definition (HD) and UltraHD video content to be displayed on a variety of devices, ranging from Smart TVs and laptops to mobile devices and smartwatches. There is a continuous need to decrease video transmission bit-rate, especially for delivery over wired and/or wireless/cellular networks without reducing visual presentation quality [1]-[4], [5].

In addition, the HDR UltraHD video content is recently attracting a lot of attention due the relatively high luminance levels and fine shadow details, which extend much beyond conventional Standard Dynamic Range (SDR) content. The HDR technology makes it possible to present highly bright signals along with very dark signals on the same video frame, thereby providing a high contrast ratio within the same image. In addition, the HDR video content is usually combined with a Wide Color Gamut (WCG), such as BT.2020, thereby enabling to present video with a significantly extended color spectrum. Particularly, HDR has gained its popularity after the development and approval of the High Efficiency Video Coding (HEVC) standard, i.e. H.265/MPEG-HEVC, in 2013 [6].

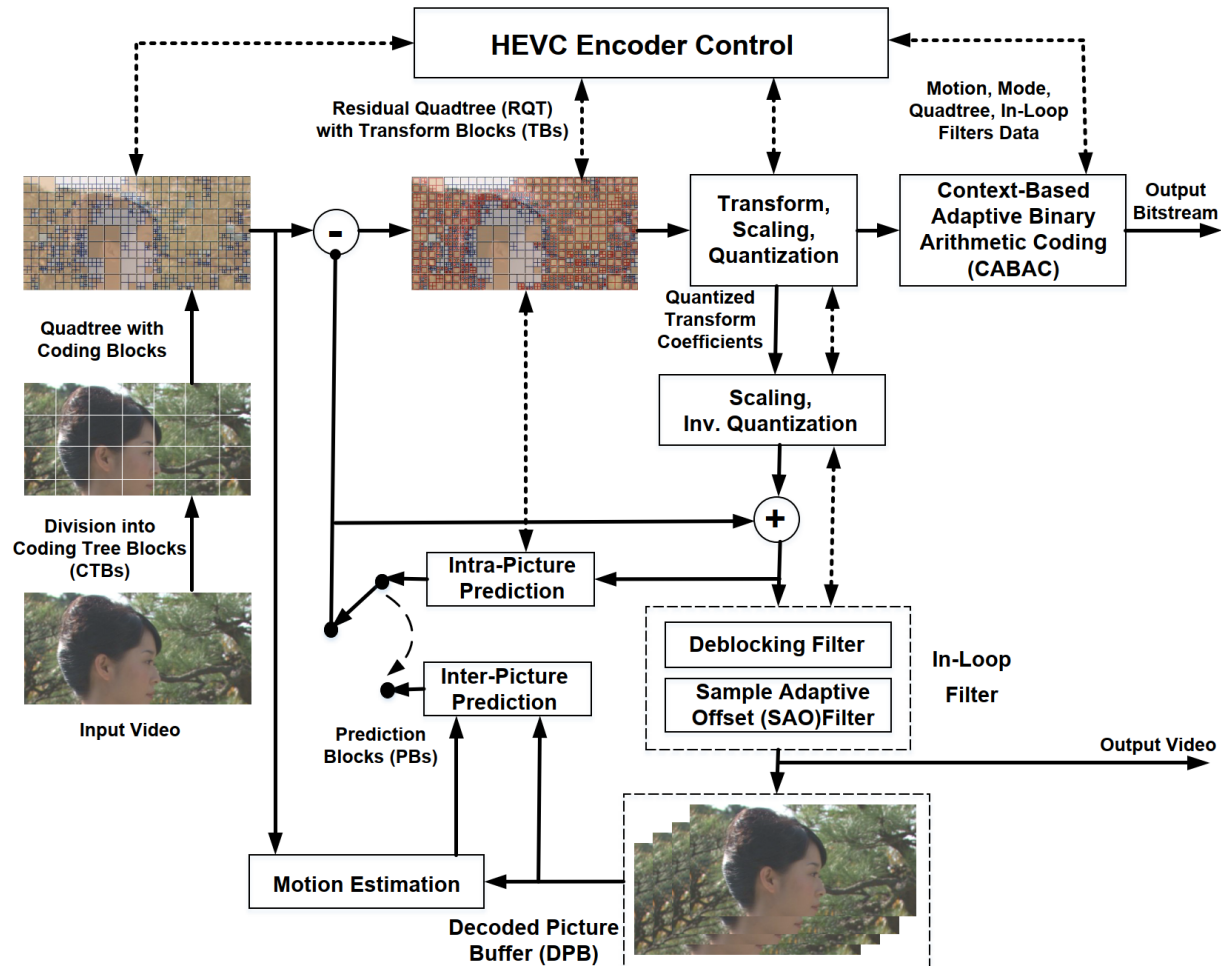


Figure 1 – The schematic block diagram of the H.265/MPEG-HEVC encoder.

The development of the first version of HEVC by the Joint Collaborative Team on Video Coding (JCT-VC) of ITU-T Video Coding Experts Group (VCEG) and ISO/IEC Moving Pictures Expert Group (MPEG) was officially finalized in January 2013 [6]. *Figure 1* illustrates a block diagram of the H.265/MPEG-HEVC encoder. After that, the final aligned HEVC specification was approved by ITU-T as Recommendation H.265 and by ISO/IEC as MPEG-H, Part 2. About one year later, the 2nd HEVC version was finalized, incorporating the Range Extensions (RExt) as well as the Scalable and Multi-view Extensions (SHVC and MV-HEVC, respectively) [7]. In turn, the 3rd and 4th HEVC edition were issued in 2015 and 2016, further containing the 3D Video Coding Extensions (3D-HEVC) and the Screen Content Coding Extension (HEVC-SCC), respectively [8], [9]. When developing the H.265/MPEG-HEVC standard, high-resolution video coding was considered as one of its main potential application scenarios, while keeping it applicable to almost all existing use cases that were already targeted by H.264/MPEG-AVC. The development process of H.265/MPEG-HEVC was driven by the most recent scientific and technological achievements in the video coding field. As a result, when compared to its predecessor - H.264/MPEG-AVC, H.265/MPEG-HEVC is able to achieve a bitrate reduction of roughly 50% for substantially the same visual quality [1]-[4].

Video applications continue to gain a lot of traction and to have an enormous demand. A very significant increase in the bandwidth requirements is expected by 2023, particularly due to the increase in the resolution supported by devices. It is expected that 66% of the connected flat-panel TV sets will have the support for the Ultra-High Definition (UltraHD) resolution compared to only 33% in 2018 (note that “UltraHD” in this paper refers to the 3840x2160 resolution, also known as 4K or 2160p). The typical bitrate for a 60fps 4K HDR10 video is between 15 to 24 Mbps, nearly four times the typical High-Definition (HD) video bitrate [5].

As a result, there is a continuous strong need to further decrease video transmission bitrate, especially for the UltraHD content, substantially without reducing the perceptual visual quality.

2. Background: Human Visual System

One of the most popular approaches for improving video quality is related to considering *spatial frequency sensitivity* of Human Visual System (HVS). As known, the HVS system is a part of the central nervous system, which enables processing of visual details and generating non-image photo response functions by obtaining and processing visible information. Thus, for example, during the coding process, the values of the Discrete Cosine Transform (DCT) frequency coefficients can be attenuated by applying quantization matrices: i.e. lower spatial frequencies are usually quantized with smaller quantization parameters (QPs), while higher spatial frequencies - with larger QPs [10], [11].

However, the improvements in visual quality of the related state-of-the-art approaches are relatively small, and more efficient solutions are desirable. In addition, most of the state-of-the-art pre-processing methods are designed for the relatively low-resolution SDR video content, and as a result, these methods found to be mostly inefficient for HDR UltraHD.

In turn, this is also true for the coding schemes that aim to remove fine details below a predefined visibility threshold, which is referred as Just Noticeable Difference (JND). As a result, the state-of-the-art JND-based schemes do not provide sufficient video quality improvement for the HDR UltraHD video content as well.

In this paper, two important perceptual video coding optimizations techniques are presented and discussed in details: *Section 3* describes a novel perceptual video optimization framework based on authors' work presented in [10], [11], and *Section 4* describes a joint backward and forward masking encoding scheme, based on authors work presented in [12]. Then, future directions for perceptual video coding optimizations are provided in *Section 5*, while this paper is concluded in *Section 6*.

3. Perceptual Video Quantization Framework

The human visual system (HVS) is considered to be a very complex system, while a level of contrast that is required to generate a response perceived by HVS is known as a contrast threshold of a sinusoidal luminance pattern. In turn, an inverse of this threshold is called “contrast sensitivity”, and it varies as a function of a spatial frequency.

3.1. Background: Contrast Sensitivity Function

The relationship between the spatial frequency and contrast sensitivity is known as a contrast sensitivity function (CSF) that differs for achromatic and chromatic scenes. The term Contrast Sensitivity (CS) often relates to visual acuity, thereby being able to differentiate between the object and the background [10].

In turn, CSF generally defines the sensitivity of the observer to various frequencies of visual stimuli, e.g., sensitivity to vertical black and white strips grating as a function of spatial frequencies [13],[14]. In case, the above frequencies are higher than a threshold predefined by the Human Visual System (HVS), the human observers are not able to differentiate between the strips. Also, generally, the HVS sensitivity to luminance significantly differs from the HVS sensitivity to chrominance.

The HVS is more sensitive to low spatial frequencies than to high spatial frequencies [15]-[18], and by assuming that HVS is isotropic, it can be modeled as a nonlinear point transformation that is followed by a Modulation Transfer Function (MTF) [19].

Later, this approach was practically used in developing a HVS-based quantization table for the JPEG still image compression standard [20], [21]. Authors of [21] derived this table by incorporating a HVS model developed by Daly [15]-[18] with an uniform quantizer, and further claiming that by replacing the JPEG quantization table with their HVS-based quantization table, obvious perceptual quality improvements are achieved. More specifically, the authors of [21] applied a 1st order low-contrast MTF of the HVS model proposed by Daly for generating a HVS-based quantization table for the baseline JPEG image compression standard, as follows below.

16 16 16 16 17 18 20 24		16 16 16 16 17 18 21 24
16 16 16 17 18 20 24 25		16 16 16 16 17 19 22 25
16 16 17 18 20 24 25 28		16 16 17 18 20 22 25 29
16 17 18 20 24 25 28 33		16 16 18 21 24 27 31 36
17 18 20 24 25 28 33 41	16 16 16 16	17 17 20 24 30 35 41 47
18 20 24 25 28 33 41 54	16 16 16 16	18 19 22 27 35 44 54 65
20 24 25 28 33 41 54 71	16 16 16 16	21 22 25 31 41 54 70 88
24 25 28 33 41 54 71 91	16 16 16 16	24 25 29 36 47 65 88 115
<i>Default Quantization Matrix of 8x8 size for Inter-prediction unit (including the Luma, Cb and Cr components)</i>	<i>Default Quantization Matrix of 4x4 size for Inter-/Intra-prediction unit (including the Luma, Cb and Cr components)</i>	<i>Default Quantization Matrix of 8x8 size for Intra-prediction unit (including the Luma, Cb and Cr components)</i>

Figure 2 – The default perceptual quantization matrices defined in the HEVC standard.

The HEVC video coding standard allows usage of perceptually-tuned frequency-dependent quantization matrices, instead of applying a constant quantization parameter (QP) on each coding block. These matrices better suit the HVS characteristics by allowing to quantize higher frequencies in a stronger manner, while their sizes vary from 4x4 to 32x32. However, the specification of the HEVC standard [6]-[9] only defines default quantization matrices for 4x4 and 8x8 transform blocks (see *Figure 2*).

The rest of matrices, i.e. for transform block sizes of 16×16 and 32×32 , are obtained by upsampling the original 8×8 perceptual quantization matrix respectively. More specifically, the original 8×8 matrix is replicated: each block in the 8×8 matrix is replicated to the 2×2 area of the 16×16 transform block and to the 4×4 area of the 32×32 transform block, as shown in *Figure 3*.




Figure 3 – Upsampling the default 8×8 HEVC matrix for obtaining default matrices for 16×16 and 32×32 transform block sizes.

Depending on the transform block type (i.e. used for *Intra* or *Inter*-picture prediction) and transform block size (i.e. 4×4 , 8×8 , 16×16 or 32×32), the HEVC standard employs twenty quantization matrices: 8 matrices for *Y (Luma)* component and 6 matrices for each of *Cb* and *Cr (Chroma)* components, as specified in *Table 1* below.

Table 1 - HEVC Transform Block Type/Size-Dependent Quantization Matrices.

Block Component	Type/size-Dependent Quantization Matrices
Y (<i>Luma</i>)	Intra 4×4 , Intra 8×8 , Intra 16×16 , Intra 32×32 ; Inter 4×4 , Inter 8×8 , Inter 16×16 , Inter 32×32 .
Cb (<i>Chroma</i>)	Intra 4×4 , Intra 8×8 , Intra 16×16 ; Inter 4×4 , Inter 8×8 , Inter 16×16 .
Cr (<i>Chroma</i>)	Intra 4×4 , Intra 8×8 , Intra 16×16 ; Inter 4×4 , Inter 8×8 , Inter 16×16 .

In addition, HEVC allows to use other quantization matrix values (i.e. *customized* quantization matrix values) besides the default values. For that, the above-mentioned customized quantization matrix values can be transmitted within the HEVC bitstream Sequence Parameter Set (SPS) or Picture Parameter Set (PPS), while coding these customized values by using so called Differential Pulse Code Modulation or in short DPCM. Similarly, the 16×16 and 32×32 quantization matrices are obtained by upsampling corresponding 4×4 and 8×8 quantization matrices (see *Figure 3*).

In spite of the fact that the HEVC default perceptual quantization matrices of *Figure 3* are based on HVS, they were initially developed and tested on low-resolution JPEG images, such as 512×512 pixels. Therefore, they almost didn't provide any benefits for UltraHD video content, that has the 3840x2160 resolution in terms of luma samples, which is the most popular resolution nowadays. As a result, this is currently also a reason for the relatively low popularity of these default perceptual quantization matrices, which most often are not used at all.

In the following section, the design and development of novel perceptual quantitation matrices for encoding UltraHD HDR video content is presented, further being inspired by investigating CSF of a human visual system. The novel perceptual quantitation matrices significantly improve perceived video quality without a need for pre-processing and without an increase in coding computational complexity.

3.2. Perceptual Quantization Matrices for UltraHD Resolution Displays

The contrast sensitivity of the human eyes, and more generally – of the human visual system as the whole, is one of the main factors how humans perceive achromatic or chromatic images. Therefore, when developing HVS-based models, it is especially important to determine the Contrast Sensitivity Function (CSF) as accurate as possible. For that, it is important to consider substantially all known HVS characteristics that have any impact of the CSF [22].

With this regard, in addition to the HVS-based models developed by Mannos&Sakrison and Daly at the end of the 20th century, Barten in his paper from 2004 proposes a more accurate HVS-based physical model/formula for the contrast sensitivity of the human eye. Particularly, in his work, Barten considers a plurality of HVS parameters, such as photon noise, neural noise, external noise, lateral inhibition, eye pupil diameter, eye pupil size, angular size of the object, luminance conditions, etc.

As a result, Barten's HVS-based CSF model is considered to be the most accurate for representing the HVS contrast sensitivity, and considered to be the best CSF model to date [23]-[25]. However, Barten's model is very complex, and its usage for an accurate determining of efficient perceptual quantization matrices, either in HEVC or in other emerging video coding standards, is found to be very challenging.

Therefore, at the 1st step of authors work [10], perceptual quantization matrices to be employed during the video coding loop have been designed by fitting the Daly's HVS-based model into the Barten's HVS-based CSF model.

In turn, at the 2nd step, the CSF-tuned human visual coefficients are empirically optimized by gradually attenuating high frequencies in a much stronger manner than low frequencies, and further giving priority to luminance (Luma) over chrominance (Chroma) due to the fact that human eye is more sensitive to Luma changes than that of Chroma [10].

For obtaining experimental results presented in this section, a special emphasis was made on video sequences having a 10-bit sample representation and UltraHD spatial resolution (particularly, the 4K resolution – i.e. 2160p, or more specifically, the 3840x2160 resolution in terms of luma samples), as presented in *Table 2* below.

Table 2 - HDR UltraHD test video sequences.

Tested Video Sequences	No. of Frames	Frame Rate per Second	Resolution	Dynamic Range
Lucy (provided by NBCUniversal [®])	8425	24	3840x2160	HDR
Everest (provided by NBCUniversal [®])	7202	23.98	3840x2160	HDR
Warcraft (provided by NBCUniversal [®])	8177	23.98	3840x2160	HDR
Regatta (provided by UltraHD forum [®])	5841	59.94	3840x2160	HDR

In *Figure 4* below, sample frames of the above-mentioned tested sequences are presented.



Figure 4 – A sample frames from the tested sequences: (a) “Lucy”; (b) “Everest”; (c) “Warcraft”; (d) “Regatta”.

These tested video sequences can be generally characterized as follows:

- “Lucy” – includes many action scenes, many fast motion scenes, mixed content [26];
- “Everest” – includes mountains views, many snow scenes, mostly slow motion scenes [26];
- “Warcraft” – includes various computer-generated content, mostly fast motion scenes [26];
- “Regatta” – includes many water scenes, many fast motion scenes [27].

The x265 open source HEVC-based encoder [28] was selected for implementing the proposed perceptual quantization framework due its ubiquity in the industry and flexibility in configuration, as well as due to its

good coding performance. When encoding the video sequences of *Table 2* with the target bit-rates of 8Mb, 10Mb, 12Mb and 14Mb, the BD-BR coding gains are significantly large in terms of both PSNR [29] and SSIMPlus [30], [31]. Specifically, the BD-BR PSNR and BD-BR SSIMPlus bit-rate savings for the HEVC encoding with the proposed QMs versus HEVC encoding with the constant QP reach coding gain of 15.5%, for encoding the “Lucy” video sequence, as is shown in *Table 3* below.

Table 3 - BD-BR PSNR and SSIMPlus bit-rate savings for the HEVC encoding.

Tested Video Sequences	BD-BR SSIMPlus Proposed QMs vs. Default HEVC QMs	BD-BR SSIMPlus Proposed QMs vs. no QMs
Lucy	-13.5%	-15.5%
Everest	-3.8%	-5.9%
Warcraft	-6.1%	-6.0%
Regatta	-10.6%	-11.9%

As can be clearly seen from *Table 3*, by employing the proposed perceptual QMs, significant coding gains of up to about 16% are achieved.

Especially, these coding gains are significant in terms of the SSIMPlus metrics, while for the “Regatta” video sequence the coding gain is the most significant - the “Regatta” video content is considered to be hard to encode, since it contains many water scenes, and the proposed perceptual QMs perform much better for such content.

In turn, *Table 4* below presents BD-BR PSNR and BD-BR SSIMPlus bit-rate savings for the HEVC encoding with the proposed perceptual QMs versus HEVC encoding with the constant QP per Common Test Conditions (CTC) defined in [32] – i.e. without employing the default HEVC QMs.

Table 4 - BD-BR PSNR and SSIMPlus bit-rate savings for the HEVC encoding.

Tested Video Sequences	BD-BR PSNR Proposed QMs vs. no QMs	BD-BR SSIMPlus Proposed QMs vs. no QMs
Lucy	-2.5%	-6.3%
Everest	-0.9%	-5.8%
Warcraft	-1.0%	-7.9%
Regatta	-2.4%	-11.3%

In this case, as seen from *Table 4*, the coding gain as a result of employing the proposed QMs is even larger and is up to 11.3%.

Below, as an example, the breakdown of the “Regatta” video sequence is presented, thereby showing the SSIMPlus score in a range between 0 and 100, while the larger the number - the better the video quality is (100 is the best possible quality). As is clearly seen from *Table 5*, when the proposed perceptual QMs are employed, the SSIMPlus score [30], [31] is significantly higher – i.e. it is more than 1 point - compared to encoding with a constant QP per CTC [32] (i.e. marked as “no QMs” in the above table). Similarly, the encoding with the default HEVC QMs provides a little improved visual quality compared to the above-mentioned constant QP encoding, but still much worse quality compared to the encoding with the proposed perceptual QMs. In turn, the minimal SSIMPlus score is increased by 1 point for the bit-rates of 8Mb, 10Mb, 14Mb, and by even 2 points for the bit-rates of 6Mb, 12Mb, which is visually noticeable.

Table 5 - SSIMPlus scores for encoding the Regatta video sequence.

Target Bit Rate (Kb)	SSIMPlus (no QMs)	SSIMPlus (Default HEVC QMs)	SSIMPlus (Proposed QMs)	Minimal SSIMPlus (no QMs)	Minimal SSIMPlus (Default HEVC QMs)	Minimal SSIMPlus (Proposed QMs)
6,000	83.09	83.14	84.16	68	68	70
8,000	86.31	86.42	87.49	74	74	75
10,000	88.61	88.74	89.80	78	78	79
12,000	90.28	90.42	91.40	81	81	83
14,000	91.49	91.65	92.56	84	84	85

Perceptual quantization matrices for mobile devices, such as tablets and smartphones, which have much smaller display sizes, thereby allowing a significant reduction in the overall video transmission bit-rate by removing non-perceivable details from each video frame, are discussed at the next section below.

3.3. Perceptual Quantization Matrices for Mobile Device Displays

For mobile devices, which have smaller display sizes, there is a need to develop dedicated perceptual quantization matrices for coding High Dynamic Range (HDR) mobile device-based video content. The perceptual quantization matrices proposed at [11] are based on Human Visual System (HVS) and utilized for reducing video transmission bit-rate and for optimizing perceived visual quality of video content to be displayed on mobile devices, such as tablets and smartphones.

According to video coding scheme proposed at [11], visual quality of the HDR UltraHD video content is significantly improved, for substantially the same bit-rate, in terms of the popular objective quality metric SSIMPlus. On the other hand, the video transmission bit-rate is significantly reduced by up to about 25%, while keeping visual quality of the video content, to be displayed on a mobile device screen, substantially at the same level.

Similarly to what is explained at the previous section, the x265 open source HEVC-based encoder [28] was selected for implementing the proposed perceptual quantization framework due its ubiquity in the industry and flexibility in configuration, as well as due to its good coding performance.

Table 6 below presents, in its right column, the BD-BR SSIMPlus bit-rate savings for the HEVC encoding with the proposed perceptual quantization matrices (QMs) that are optimized for *mobile devices* versus HEVC encoding with the default QMs, as defined in the HEVC specification. In addition, in the middle column, are presented the BD-BR SSIMPlus bit-rate savings for the HEVC encoding with the proposed perceptual QMs that are optimized for *mobile devices* versus HEVC encoding with the constant QP as defined in CTC [32]– i.e. without employing the default HEVC QMs. As can be clearly seen from Table 6, by employing the proposed perceptual QMs, significant coding gains of up to about 25% are achieved. It should be noted that for the “Regatta” video sequence the coding gain is the most significant - the “Regatta” video content is considered to be hard to encode, since it contains many water scenes, and the proposed perceptual QMs perform much better for such content. Below, as an example, the quality scores for encoding the “Regatta” video sequence with target bit rates varying between 2Mb and 5Mb are presented, thereby showing the SSIMPlus score in a range between 0 and 100, while the larger the number - the better the video quality is (100 is the best possible quality).

Table 6 - BD-BR SSIMPlus bit-rate savings for the HEVC encoding.

	BD-BR SSIMPlus	BD-BR SSIMPlus
--	----------------	----------------

Tested Video Sequences	Proposed QMs vs. Default HEVC QMs	Proposed QMs vs. no QMs
Lucy	-15.5%	-16.8%
Everest	-21.4%	-22.2%
Warcraft	-5.9%	-7.8%
Regatta	-22.2%	-23.9%

Also, as it is clearly seen from *Table 7*, when the proposed perceptual QMs that are optimized for *mobile devices* are employed, the SSIMPlus score is significantly higher – i.e. it is up to about *2 points* compared to encoding with a constant QP according to CTC [32], i.e. marked as “no QMs”. Similarly, the encoding with the default HEVC QMs provides a little improved visual quality compared to the above-mentioned constant QP encoding, but still much worse quality compared to the encoding with the proposed perceptual QMs. In addition, the minimal SSIMPlus score is increased by a very significant number of up to *7 points* for the bit-rate of 3Mb, which is visually clearly noticeable.

Table 7 - SSIMPlus scores for encoding the Regatta video sequence.

Target Bit Rate (Kb)	SSIMPlus (no QMs)	SSIMPlus (Default HEVC QMs)	SSIMPlus (Proposed QMs)	Minimal SSIMPlus (no QMs)	Minimal SSIMPlus (Default HEVC QMs)	Minimal SSIMPlus (Proposed QMs)
2,000	76.82	76.85	78.61	48	48	54
3,000	82.06	82.10	83.48	58	58	65
4,000	84.82	84.87	86.10	65	66	72
5,000	87.20	87.28	88.48	73	73	77

In the next section, another promising approach for increasing video coding gain is discussed. In this approach, “visual masking” is applied [12], which is based on the human visual system (HVS) characteristics.

4. Perceptual Video Masking Framework

As known, HEVC was especially designed for coding of HD and UltraHD video content with a much larger coding gain compared to its predecessor H.264/MPEG-AVC, thereby reducing both spatial and temporal video content redundancies in a much more efficient way, which in turn significantly assisted in compression of the HDR UltraHD video content [1]-[5]. However, coding of the HDR video content still remains challenging due to users’ demands for high visual quality, which in turn requires allocating more bits and increasing a video coding depth (e.g., from 8 bits to 10 bits). In addition, the transmission bandwidth is normally limited due to a typical limitation of the existing network infrastructure, especially in case of the transmission over wireless/cellular networks. As a result, in order to stay within the transmission bandwidth limits, the high-resolution HDR video content is often compressed with visually perceived coding artifacts. Moreover, encoding of the HDR content normally consumes significant computational resources due to a requirement to preserve fine details within the HDR video. Therefore, there is further a strong demand to improve perceived visual quality of the compressed HDR video substantially without increasing its bit-rate [10], [11].

4.1. Background: Visual Masking

One of the promising approaches for increasing video coding gain is applying “visual masking”, which is based on a very interesting phenomenon observed in the human visual system (HVS) [12]. According to this phenomenon, two or more stimuli are presented sequentially to a viewer, with one stimulus acts as a target which has to be detected and described, while other stimuli are used to mask the visibility of that target. With this regard, a good amount of research has been carried out in the video compression field, such as [33],[34] for example, which exploits the above-mentioned phenomenon by providing a psycho-visual algorithm that has been implemented in the x264 encoder [35]. In turn, more advanced studies of are further presented and discussed in [36]. In addition, in the most recent work, such as [37], it is proposed to mask temporal activities that are unnoticeable by human visual system by using a masking coefficient. Further, [38] presents a video Just Noticeable Difference (JND) scheme by employing compound spatial and structure-based temporal masking, further measuring a JND threshold for each transform coefficient of a color video. Also, [39] proposes an improved transform-based JND estimation model considering multiple masking effects.

However, all surveyed existing visual masking approaches, the most interesting of which are indicated above, lead to relatively low bitrate savings. As a result, these approaches have not been adopted in the video streaming/coding industry to date. In addition, computational complexity of existing visual masking schemes is relatively high due to the utilization of relatively complex quantization models [19].

4.2. Forward and Backward Masking Encoding Scheme

In this section, a masking technique for videos is exploited and discussed in detail. Extensive experiments have been carried out for the unidirectional (either forward or backward) temporal masking, but due to the lower coding gains when compared to the bidirectional (i.e. *joint* forward and backward) temporal masking, and to keep the presentation of the experimental results of this work in a clear and simple manner, this paper is focused on the bidirectional temporal masking only. The x265 open source HEVC-based encoder [28] was selected for implementing the proposed joint forward and backward temporal masking framework due its ubiquity in the industry and flexibility in configuration, as well as due to its good coding performance.

With this regard, *Figure 5* presents a schematic illustration of the proposed framework, which includes three forward sub-windows 1 to 3, and three backward sub-windows 4 to 6. Each window can have a different length and for each window, a set of different quantization parameters (QPs) can be assigned by adding the following QP offsets: ΔRef_1 , $\Delta Non-Ref_1$; ΔRef_2 , $\Delta Non-Ref_2$; ΔRef_3 , $\Delta Non-Ref_3$; ΔRef_4 , $\Delta Non-Ref_4$; ΔRef_5 , $\Delta Non-Ref_5$; ΔRef_6 , $\Delta Non-Ref_6$. Also, different QP offsets can be assigned separately to reference (e.g., *P-frames*, *B-frames*) frames and to non-reference frames (e.g., *b-frames*) present inside each masking window. The above-mentioned QP offsets are predefined in the x265 code [28] for reference *B-frames* and for non-reference *b-frames*, the offsets for *P-frames* are automatically reduced by 30%, thereby applying only 70% of the ΔRef offset value, to improve their quality and to increase a coding gain. In addition, no QP offsets are applied to *I frames*, regardless of the fact whether the *I-frame* is a scenecut or not. In case when an *I-frame* is present inside a masking window, the masking is avoided on all frames after this *I-frame* in a given masking direction (either forward or backward). The values of QP offsets can be customized using the x265 command line [28].




Figure 5 – A schematic illustration of the proposed joint backward and forward temporal masking framework.

For evaluating the proposed framework, the authors selected a wide range of cinematic content, mostly in 10-bit UltraHD resolution, which can be characterized as follows: (a) “El Fuente” – includes mixed content, with both fast and slow motion; (b) “Lucy” – includes many action scenes, many fast motion scenes, mixed content; (c) “Warcraft” – includes various computer-generated content, mostly fast motion scenes; (d) “Everest” – includes mountains views, snow scenes, mostly slow motion scenes; (e) “Regatta” – includes water scenes, many fast motion scenes. Main technical parameters of the above-mentioned cinematic content are presented in *Table 8* below:

Table 8 – Test Sequences.

No	Sequence name	Resolution	Frame count	Frame rate	Duration (sec.)	Bit depth
1	El Fuente	3840x2160	1500	60	25	8
2	Lucy	3840x2160	480	24	20	10
3	Warcraft	3840x2160	495	23.98	~20	10
4	Everest	3840x2160	480	23.98	~20	10
5	Regatta	3840x2160	1199	59.94	~20	10

The test environment, lighting, and the rest of requirements for obtaining optimal viewing conditions were set according to [40].

4.2.1. Experimental Results and Brief Discussion

The experiments were conducted according to [40],[41]. Due to the COVID-19 restrictions, viewing sessions were done remotely, and strict instructions were provided to all participants in accordance with [40],[41]. The video playback was done on high-end consumer 4K TV displays capable of playing HEVC-encoded video content with a minimal screen size of 55” (OLED TV). High-quality 32” professional SDI reference/grading displays were used as well.

Table 9 – Bitrate Savings in Terms of BD-BR.

Sequence Name	CRF	Without Masking (Reference)		With Masking (Tested)		BD-BR Savings
		Bitrate	MOS	Bitrate	MOS	
El-Fuente	20	23960.64	88	18266.86	89	-10.7%
	24	14701.49	85	11378.59	82	
	28	9013.59	80	7115.96	79	
	32	5629.3	73	4560.42	72	
	36	3616.78	64	3054.72	64	
Lucy	20	14074.81	89	11390.49	88	-5.6%
	24	8042.85	85	6677.24	84	
	28	5003.31	74	4210.82	75	
	32	3211.48	69	2737.33	62	
	36	2113.21	63	1832.32	55	
Warcraft	20	7092.24	88	6655.14	88	-2.3%
	24	4082.99	85	3852.15	85	
	28	2572.63	82	2436.43	83	
	32	1705.72	73	1621.98	71	
	36	1171.32	66	1121.69	58	
Everest	20	13420.24	85	11418.12	85	-14.8%
	24	5738.53	83	4943.96	82	
	28	2516.55	76	2224.67	78	
	32	1480.22	69	1335.23	70	
	36	1007.73	65	922.89	65	
Regatta	20	34769.29	83	26042.09	84	-26.3%
	24	21264.06	80	15836.46	82	
	28	13164.48	82	9791.08	79	
	32	8342.09	76	6249.4	72	
	36	5349.47	66	4192.14	61	
Average				-11.9%		

As seen from *Table 9*, the largest bitrate savings of more than 26% are for “Regatta” sequence, which has the highest bitrate, on average.

On the other hand, the smallest bitrate savings of 2.3% are for the “Warcraft” sequence, which has the smallest bitrate, on average. Further, there is a substantial decrease in the overall computational complexity in terms of encoding times (for simplicity, the results are not presented).

One of the *important findings* is that the proposed joint backward and forward temporal masking framework tends to perform better for higher bitrates and frame rates, as well as for content that includes textures, such as water and snow.

In this paper, two approaches have been presented and discussed in detail: perceptual quantization matrices and visual masking. Future directions with this regard are discussed at the next section.

5. Future Directions for Perceptual Video Coding Optimizations

In spite of the fact that the HEVC standard was especially designed for the HD and UltraHD video content, more efficient video compression techniques are still desired, especially for streaming UltraHD video content as well as Panorama video content (so called 360° video content) from concerts, shows, sport events, etc. Therefore, in order to fulfill this demand, the exploration phase for future video coding technologies beyond HEVC (ITU-T H.265 | ISO/IEC 23008-2) started in October 2015 by establishing a Joint Video Exploration Team (JVET) on Future Video Coding of ITU-T VCEG and ISO/IEC MPEG. In turn, these future technologies were integrated into the Joint Exploration Test Model (JEM), and the official standardization activities for the next-generation video coding standard officially started in April, 2018 - right upon publishing results of the Call for Proposals (CfP) for future video coding technologies (after completing the Call for Evidence (CfE) in 2017).

The emerging video codec under the JVET development was titled “Versatile Video Coding”, or in short, VVC [42].

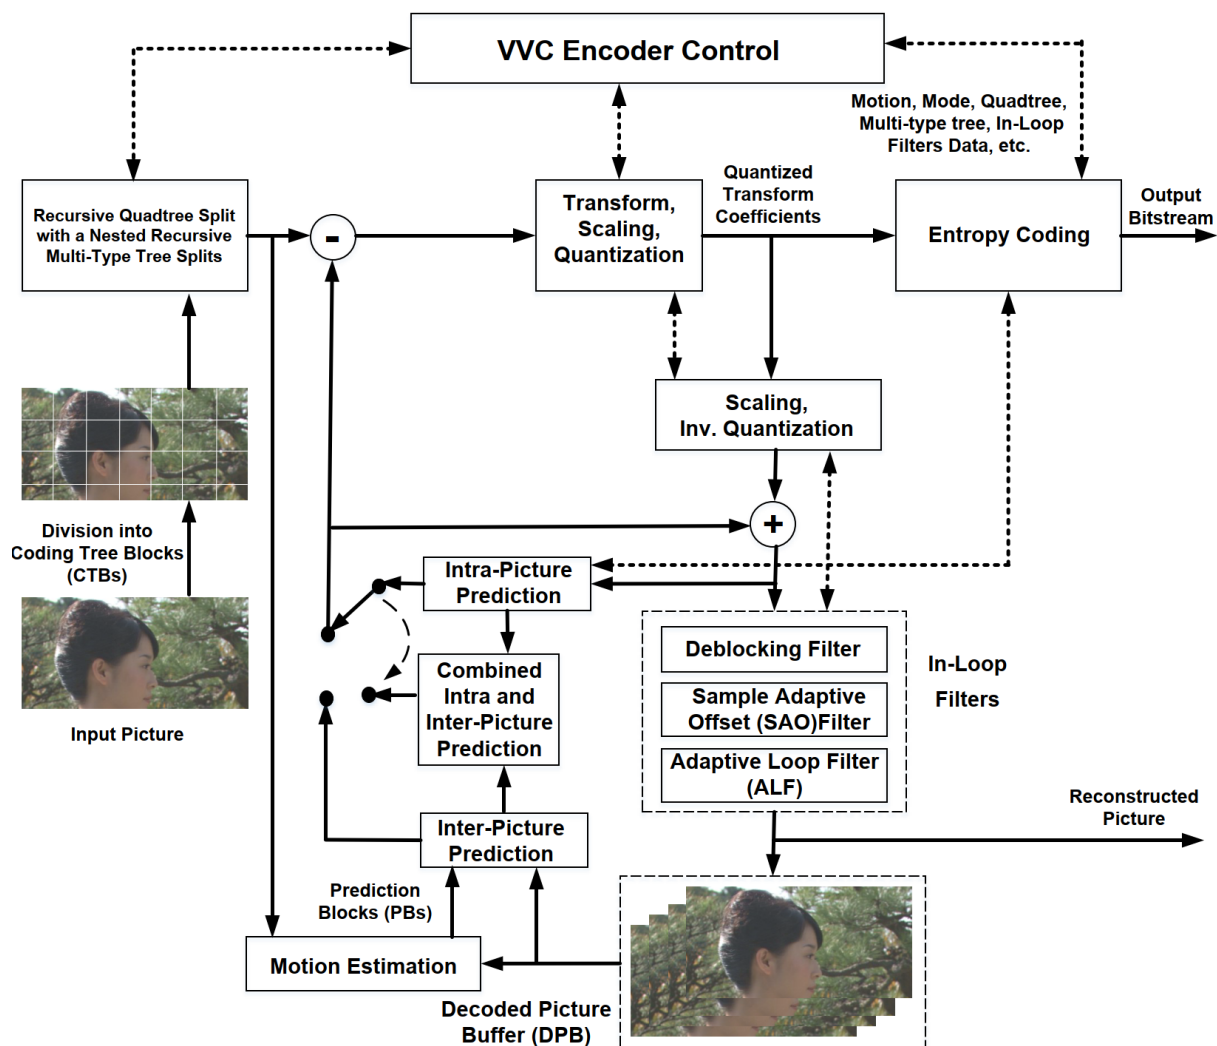


Figure 6 – The schematic block diagram of the H.266/MPEG-VVC encoder.

In turn, the first VVC draft along with the VVC Test Model 1 (VTM1) was published right after the April, 2018 meeting. As one of the main tools that provides a significant coding gain, VVC includes a quadtree with nested multi-type tree by using a coding block structure of binary and ternary splits. Further, additional tools and features include: Intra-mode coding with 67 Intra-picture prediction modes; Intra Block Copy (IBC); Bi-Directional Optical Flow (BDOF); Adaptive Motion Vector Resolution (AMVR); Geometric Partitioning Mode (GPM); Combined Inter and Intra Prediction (CIIP); Adaptive Loop Filter (ALF); *and many others*. The first version of the VVC standard (i.e., VVC v1) was officially finalized during the 19th JVET meeting, which took place between June 22 and July 1, 2020, and the VVC codec is currently starting to be widely deployed worldwide [42]. The schematic block diagram of the H.266/MPEG-VVC encoder is presented in *Figure 7*.

Therefore, the perceptual quantization matrices and masking encoding schemes described in this paper, can be applied during the VVC encoding process as well, thereby leading to even larger coding gains.

In addition, the presented perceptual video coding optimizations can be applied to the framework of [43], as described below. According to this framework, the x265-based multi-resolution encoding architecture is leveraged to significantly speed-up the encodes through sharing the analysis information from lower to higher resolutions. For example, the encodes that generate representations at the 960x540 resolution (i.e., 540p) can share information to those that generate representations at the 1920x1080 resolution (i.e., 1080p, a dyadic multiple of 540p), which in-turn can speed up the 3840x2160 resolution encode (i.e., 2160p, a dyadic multiple of 1080p). In addition, the encodes that generate representations for the 1280x720 resolution can share information to the encode that generates a representation at the 2560x1440 resolution.

Therefore, a framework can run efficient encodes, which generate both the reference and dependent representations in parallel, while also efficiently handling dependencies with the sharing analysis information among these presentations. In turn, the Adaptive Bit-Rate (ABR) ladder allows the more efficient sharing of the aforementioned analysis information among the representations at different resolutions.

Further, *Figure 7* schematically presents a high-level block-diagram of the proposed framework for the efficient usage of the x265-based [28] adaptive multi-resolution encoding architecture, thereby enabling the efficient sharing of analysis information across representations. As shown in *Figure 7*, the proposed framework is being used for converting an input video source, which has the $W \times H$ resolution, to n representations, each representation having a scaling factor of L_i (*with $i = 0$ to n*) and bitrate B_i . This framework uses a configuration file in order to define the encoding graph and the degree of encoder decision reuse between a dependent and a reference representation. In addition, it is possible to add representation-specific parameters to the configuration file, such as limiting the Coding Tree Unit (CTU) size for lower resolutions.

For the more detailed description regarding the syntax of the x265 configuration file, the reader is referred to [28].

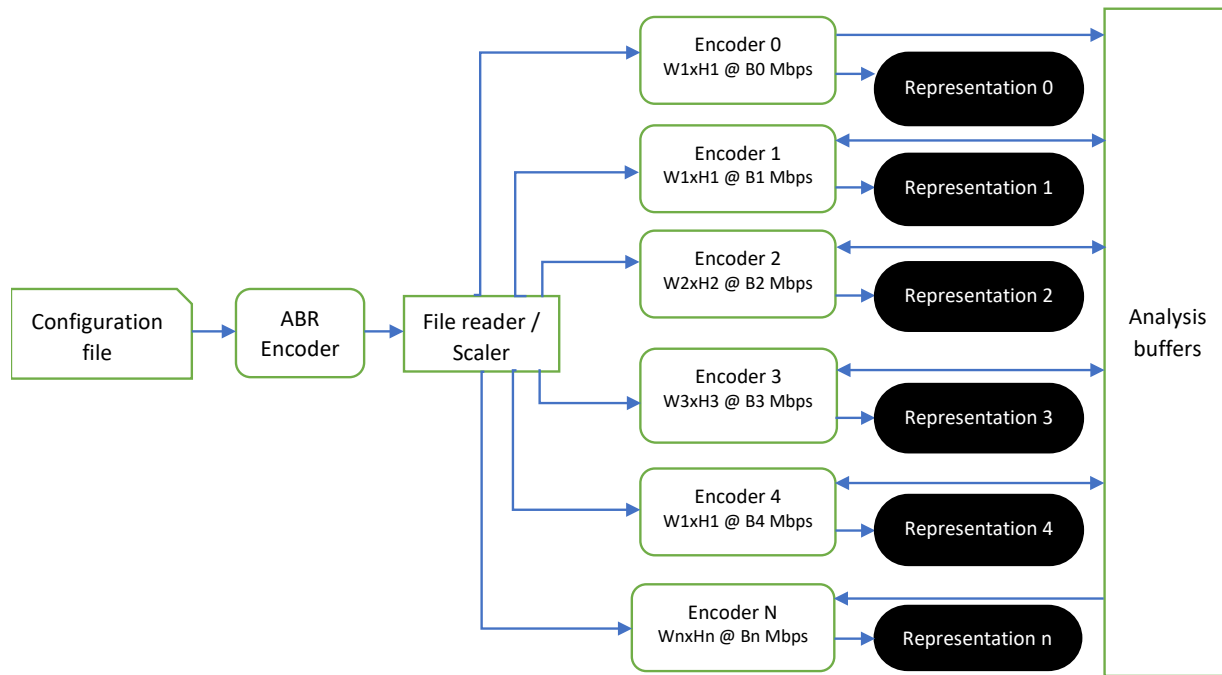


Figure 7 – A multi-resolution encoding framework for enabling efficient sharing of analysis information across representations.

By such a way, i.e. by utilizing perceptual quantization matrices and masking encoding schemes with VVC and with the above-mentioned multi-resolution encoding framework, further significant bit-rate savings can be achieved.

6. Conclusion

In this paper, perceptual video coding optimization techniques, including most recent trends and future directions, have been discussed in details. A special emphasis was made on the UltraHD resolution, such as the 3840x2160 (4K) resolution in terms of luma samples, and on the most recent, and currently the most advanced, H.265/MPEG-HEVC video coding standard. The development of the perceptual quantization matrices has been motivated by the Daly HVS-based perceptual model, which was further fitted into the more advanced and more complex Barten model (that incorporates a variety of HVS parameters) for much more accurate generation of these matrices. As a result, the video transmission bit-rate for UltraHD displays was reduced up to 11.3% in terms of SSIMPlus, while keeping the visual quality at substantially the same level. On the other hand, for smaller size mobile device displays, the video transmission bit-rate was reduced up to about 25% in terms of SSIMPlus objective quality metric, while keeping the visual quality substantially at the same level. In addition, a joint backward and forward temporal masking framework was presented, which considers temporal distances between frames and the closest scenecuts. This framework has been implemented in the popular x265 HEVC-based encoder. Based on the extensive subjective quality assessments, significant bitrate savings of up to about 26% are achieved for substantially the same perceived visual quality. The future direction mostly refer to implementing the presented framework with the most recent VVC video coding standard, further getting benefit from the presented x265-based multi-resolution encoding architecture.

Abbreviations

ABR	Adaptive Bit-Rate
ALF	Adaptive Loop Filter
AMVR	Adaptive Motion Vector Resolution
AVC	Advanced Video Coding
BDOF	Bi-Directional Optical Flow
CfE	Call for Evidence
CfP	Call for Proposals
CIIP	Combined Inter and Intra Prediction
CSF	Contrast Sensitivity Function
CTC	Common Test Conditions
CTU	Coding Tree Unit
DCT	Discrete Cosine Transform
DPCM	Differential Pulse Code Modulation
fps	frame per second
GOP	Group of Pictures
HDR	High Dynamic Range
HEVC	High Efficiency Video Coding
HVS	Human Visual System
IBC	Intra Block Copy
JCT-VC	Joint Collaborative Team on Video Coding
JND	Just Noticeable Difference
JVET	Joint Video Exploration Team
MPEG	Moving Pictures Experts Group
PPS	Picture Parameter Set
QM	Quantization Matrix
QP	Quantization Parameter
SDR	Standard Dynamic Range
SPS	Sequence Parameter Set
VCEG	Video Coding Experts Group
VTM	VVC Test Model
VVC	Versatile Video Coding

Bibliography & References

- [1] D. Grois, D. Marpe, A. Mulyoff, B. Itzhaky, and O. Hadar, "Performance comparison of H.265/MPEG-HEVC, VP9, and H.264/MPEG-AVC encoders," *Picture Coding Symposium (PCS) 2013*, pp.394-397, 8-11 Dec. 2013.
- [2] D. Grois, D. Marpe, T. Nguyen, and O. Hadar, "Comparative Assessment of H.265/MPEG-HEVC, VP9, and H.264/MPEG-AVC Encoders for Low-Delay Video Applications", Proc. SPIE 9217, *Applications of Digital Image Processing XXXVII*, 92170Q, Sept. 2014.
- [3] D. Grois, T. Nguyen, and D. Marpe, "Coding Efficiency Comparison of AV1/VP9, H.265/MPEG-HEVC, and H.264/MPEG-AVC Encoders," *Picture Coding Symposium (PCS)*, Dec. 2016.
- [4] D. Grois, T. Nguyen, and D. Marpe, "Performance Comparison of AV1, JEM, VP9, and HEVC Encoders", Proc. SPIE 10396, *Applications of Digital Image Processing XL*, 103960L, 7-10 Aug., 2017.
- [5] D. Grois *et al.*, "Performance Comparison of Emerging EVC and VVC Video Coding Standards with HEVC and AV1," in *SMPTE Motion Imaging Journal*, vol. 130, no. 4, pp. 1-12, May 2021.
- [6] ITU-T, Recommendation H.265 (04/13), Series H: Audiovisual and Multimedia Systems, Infrastructure of audiovisual services – Coding of Moving Video, High Efficiency Video Coding.
- [7] ITU-T, Recommendation H.265 (10/14), Series H: Audiovisual and Multimedia Systems, Infrastructure of audiovisual services – Coding of Moving Video, High Efficiency Video Coding.
- [8] ITU-T, Recommendation H.265 (04/15), Series H: Audiovisual and Multimedia Systems, Infrastructure of audiovisual services – Coding of Moving Video, High Efficiency Video Coding.
- [9] ITU-T, Recommendation H.265 (12/16), Series H: Audiovisual and Multimedia Systems, Infrastructure of audiovisual services – Coding of Moving Video, High Efficiency Video Coding.
- [10] D. Grois, and A. Giladi, "Perceptual quantization matrices for high dynamic range H.265/MPEG-HEVC video coding", Proc. SPIE 11137, *Applications of Digital Image Processing XLII*, 111370O, 2020.
- [11] D. Grois, and A. Giladi, "HVS-Based Perceptual Quantization Matrices For HDR HRVC Video Coding for Mobile Devices", pp. 1-14, IBC 2020.
- [12] D. Grois, A. Giladi, P. K. Karadugattu, and N. Balasubramanian, "Novel temporal masking framework for perceptually optimized video coding", In *Proceedings of the 1st Mile-High Video Conference (MHV '22)*. Association for Computing Machinery, New York, NY, USA, 119–120.
- [13] G. M. Johnson and M. D. Fairchild, "On Contrast Sensitivity in an Image Difference Model", In *Proceedings of the IS&T PICS Conference*, pp. 18- 23, Portland, OR, 2002.
- [14] B.A. Wandell, "Foundation of Vision," Sinear Associates, Sunderland, MA, 1995.
- [15] S. Daly, "Subroutine for the generation of a two dimensional human visual contrast sensitivity function," Technical Report 233203Y, Eastman Kodak, 1987.
- [16] S. Daly, "The visible differences predictor: An algorithm for the assessment of image fidelity," In A. B. Watson, editor, *Digital Images and Human Vision*, pp. 179–206, 1993.
- [17] S. Daly, "A visual model for optimizing the design of image processing algorithms," *International Conference on Image Processing (ICIP)*, vol. 2, pp. 16–20, Nov.1994.
- [18] S. Daly, T. Kunkel, X. Sun, S. Farrell, and P. Crum, "Preference limits of the visual dynamic range for ultra high quality and aesthetic conveyance," *Proceedings of the SPIE*, 8651:86510J–86510J–11, 2013.
- [19] J. L. Mannos, and D. J. Sakrison, "The Effects of a Visual Fidelity Criterion on the Encoding of Images," *IEEE Trans. on Info. Theory*, Vol. IT-20, No. 4, July 1974.
- [20] R. Rosenholtz and A. B. Watson, "Perceptual adaptive JPEG coding," In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, vol. 1, pp. 901–904, Sep. 1996.
- [21] L.W. Chang, C.Y. Wang and S.M. Lee, "Designing JPEG quantization tables based on human visual system," *Proceedings 1999 International Conference on Image Processing (Cat. 99CH36348)*, Kobe, 1999, pp. 376-380 vol.2.
- [22] M. Nezamabadi, S. Miller, S. Daly, and R. Atkins, "Color signal encoding for high dynamic range and wide color gamut based on human perception," *Proceedings of the SPIE*, 9015:90150C–90150C–12, 2014.
- [23] P. G. J. Barten. Physical model for the contrast sensitivity of the human eye. *Proceedings of the SPIE*, 1666:57–72, 1992.
- [24] P. G. J. Barten. "Contrast sensitivity of the human eye and its effects on image quality," volume 72, *SPIE press*, USA, Dec. 1999.
- [25] P. G. J. Barten, "Formula for the contrast sensitivity of the human eye," *Proceedings of the SPIE*, 5294:231–238, 2003.

- [26] IMDB content database, Online: <https://www.imdb.com/title/tt2872732>.
- [27] UltraHD Forum, Online: <https://ultrahdforum.org>.
- [28] Projects from VideoLAN, x265 software library and application, Online: <https://www.videolan.org/developers/x265.html>.
- [29] G. Bjøntegaard, "Calculation of average PSNR differences between RD-curves", ITU-T Q.6/SG16 VCEG 13th Meeting, Document VCEG-M33, Austin, USA, Apr. 2001.
- [30] Eurofins[®], "Which is the best objective video quality measure and why use SSIMPLUS", Eurofins[®], Online: <https://cdnmedia.eurofins.com/digitaltesting/media/116613/qoe-why-ssimplus.pdf>.
- [31] SSIMWAVE[®], "SSIMPLUS Outperforms VMAF", 2017.
- [32] F. Bossen, "Common HM test conditions and software reference configurations," document JCTVC-L1100 of JCT-VC, Geneva, CH, Jan. 2013.
- [33] V. Adzic, H. S. Hock, and H. Kalva, "Visually lossless coding based on temporal masking in human vision," Proc. SPIE 9014, Human Vision and Electronic Imaging XIX, 90141C (25 February 2014)
- [34] V. Adzic, H. Kalva and B. Furht, "Exploring visual temporal masking for video compression," *2013 IEEE International Conference on Consumer Electronics (ICCE)*, 2013, pp. 590-591.
- [35] Projects from VideoLAN, x264 software library and application, Online: <https://www.videolan.org/developers/x264.html>.
- [36] V. Adzic, H. Kalva and B. Furht, "Temporal visual masking for HEVC/H.265 perceptual optimization," *2013 Picture Coding Symposium (PCS)*, 2013, pp. 430-433.
- [37] Siddique AA, Qadr MT, Mohy-Ud-Din Z. Masking of temporal activity for video quality control, measurement and assessment. *Measurement and Control*. 2020;53(9-10):1817-1824.
- [38] K.-C. Liu, "Color Video JND Model Using Compound Spatial Masking and Structure-Based Temporal Masking," in *IEEE Access*, vol. 8, pp. 136760-136768, 2020.
- [39] H. Wang, L. Yu, H. Yin, T. Li, and S. Wang, "An improved DCT-based JND estimation model considering multiple masking effects," *Journal of Visual Communication and Image Representation*, vol. 71, 102850, 2020.
- [40] ITU-R, Recommendation ITU-R BT.500-14 (10/2019), Methodologies for the subjective assessment of the quality of television images.
- [41] M. Wien, and V. Baroncini, "Status Report on SDR HD Verification Test Preparation", Doc. JVET-T0043, Teleconference, 7-16 Oct. 2020.
- [42] ITU-T, Recommendation H.266 (08/2020), Series H: Audiovisual and Multimedia Systems, Infrastructure of audiovisual services – Coding of Moving Video, Versatile Video Coding.
- [43] A. Mathesawaran, P. K. Karadugattu, P. Ramachandran, A. Giladi, D. Grois, P. Venkatesan, and A. Balk, "Open source framework for reduced-complexity multi-rate HEVC encoding," *Proc. SPIE 11510, Applications of Digital Image Processing XLIII*, 115101Y (25 August 2020);

Photon Avatars in the Comcast Cosmos: An End-to-End View of Comcast Core, Metro and Access Networks

A Technical Paper prepared for SCTE by

Venk Mutalik

Fellow

Comcast

1800 Arch Street, Philadelphia, PA 19103

+1 (860) 262-4479

Venk_Mutalik@Comcast.com

Steve Rupp, Senior Principal Engineer

Fred Bartholf, Senior Principal Engineer

Bob Gaydos, Fellow

Steve Surdam, Vice President

Amarildo Vieira, Principal Optical Engineer

Dan Rice, Vice President

Table of Contents

Title	Page Number
1. Introduction.....	4
2. The Big Picture.....	4
2.1. Backbone, Metro and Access	4
3. The Backbone	6
3.1. Comcast Backbone History.....	6
3.2. Optical Innovations in the Backbone.....	7
3.3. Capacity of the Backbone	8
3.4. Future of the Backbone.....	8
3.5. Purpose Built Networks.....	9
3.6. Power Consumption.....	9
3.7. IPoDWDM and Alien Waves	9
3.8. Automation	9
4. The Metro Network.....	10
4.1. The 28 CRANs of Comcast.....	11
4.1.1. In the beginning. Video on demand (VOD).....	12
4.1.2. Next up! High-speed internet (HSI).....	12
4.1.3. 10 x the bandwidth and then some!.....	13
4.2. Optimizing further - Capacity vs. Connectivity	13
4.2.1. Internet protocol over dense wave division multiplexing (IPoDWDM).....	14
4.3. Leaveraging access and metro together to support commercial services.....	15
5. The Access Network	17
5.1. Converging the Access Network.....	19
5.1. Dual Laser Bidirectional Coherent Transmission.....	21
5.2. Convergence Continuous Pervasive Monitoring.....	24
6. Conclusion.....	25
Abbreviations	25
Bibliography & References.....	27

List of Figures

Title	Page Number
Figure 1 – Comcast Backbone and Access Networks.....	5
Figure 2 – Comcast Architecture	5
Figure 3 – Comcast Backbone Architecture - 1	6
Figure 4 – Comcast Backbone Architecture - 2	7
Figure 5 - The beginning of the Metro network. 3 independent CATV providers	10
Figure 6 - Metro network today. One MSO. Flex grid and ROADMs. The Swiss Army Knife	11
Figure 7 - U Ring Topology.....	12
Figure 8 - ROADM impact on 400GbE ZR+ optics.....	15
Figure 9- Metro and Access used together to support commercial customers	16

Figure 10 – Illustrating the DAA Network of Comcast 17
 Figure 11 – Typical Secondary to Node distances in the Access Plant 18
 Figure 12 – Illustrating the Converged Comcast Access Network..... 20
 Figure 13 – Converging Direct detect and Coherent bi-Directional Wavelengths 20
 Figure 14 – Illustrating Bi-Directional Coherent systems 22
 Figure 15 – Hybrid Loop testing the Coherent Optics in the Access Plant..... 23
 Figure 16 – Illustrating Continuous Pervasive Monitoring 24

1. Introduction

As the largest broadband company in the US, Comcast serves millions of customers and businesses with a reach that stretches coast to coast. All of this is the result of a large optical network that spans core, metro and access layers with multiple intersecting points all intended to increase capacity, reduce latency, and enhance reliability.

In this paper, we describe for the first time an end-to-end view of our optical network including the core, metro, and access layers. At the core and metro, we increase capacity with a move towards flexible 400G connections and reduce latency and enhance reliability with an infrastructure that meshes color-less, direction-less, and contention-less reconfigurable multiplexers thru to each of our headends. At the access layer that connects these headends to customers and businesses, we discuss capacity increases with our move to all-digital fiber links and the distributed access architecture paradigm. Of note is a cost effective environmentally hardened dual laser bidirectional coherent 100G system and the converging all bidirectional access transmission formats on one single optical fiber. Reliability enhancements accrue with real-time continuous and pervasive optical monitoring of all these access assets. We then briefly describe the infrastructure that helps provision, visualize and event these layers. Finally, we will venture into the future of optical technology at Comcast and its positive impact on network robustness and enhancing the customer experience.

2. The Big Picture

Photons flood into the Comcast backbone network from giant Internet routers and reach the various metro routers. At the metro center, photons reincarnate and course thru the highly meshed Converged Regional Area Network infrastructure that terminates in the thousands of our headends. At the various headends, photons reincarnate again, traverse access fibers and light up the many homes, businesses and fiber nodes eventually completing their journey in the downstream. A similar process ensues in the upstream where photons transmigrate thru the access system to the various headends and mesh metro circuits before making their way back thru to the internet.

2.1. Backbone, Metro and Access

Presented below is a picture that illustrates the reach of our backbone and access systems in Comcast [1].

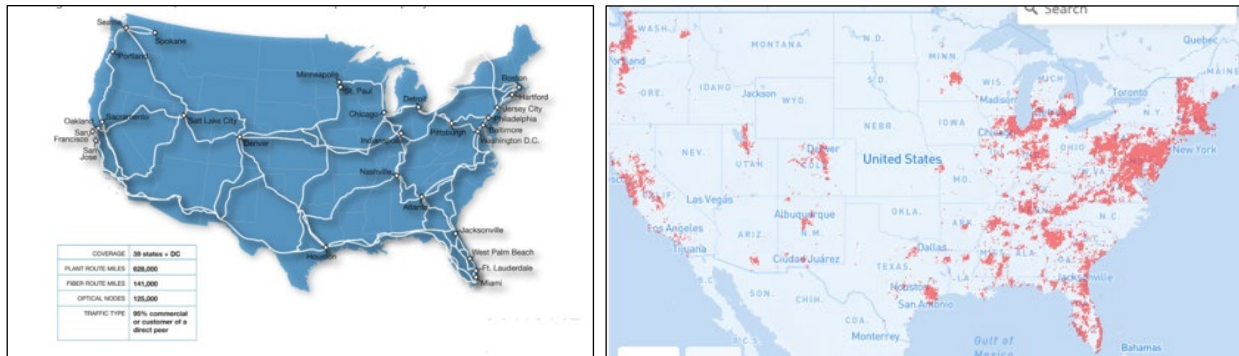


Figure 1 – Comcast Backbone and Access Networks

With the breadth that reaches coast to coast connecting important Internet Data sites and the depth of reaching 60M homes and businesses, the Comcast communications system comprising the Backbone, Metro and Access sub-systems today is now all under one single group in Comcast.

Internet sites located at Chicago, Atlanta, Ashburn and others are connected to the Backbone and via route redundancy reach all our Metro sites that interconnect hubs and headends within individual cities. At the extremities of the Metro network, residential access network interconnect with Metro and via vCMTS (virtualized Cable Modem Termination System) connect to Remote PHY based digital fiber Nodes. Commercial/Business services which were always closely connected to Metro are now converging with Residential access networks and are sharing common connectivity and fiber assets.

The residential access network is orchestrated by Comcast own provisioning, monitoring and fix agents and provide an end-to-end view. The Metro and Backbone orchestrations for provisioning, monitoring and fix are more closely related to the vendors of choice.




Figure 2 – Comcast Architecture

With this understanding in the next few sections, we describe a more detailed view of each of the subsystems along with a brief history evolving to the current state and potential innovations that we are looking into for meeting capacity and performance and transform Comcast into the leading communications provider in the country.

3. The Backbone

The early days of the Comcast network began as 28 separate networks or islands connected using Transit connectivity and leveraging local access interconnects. Comcast's network needs dramatically changed in the early 2000's. The company was looking to centralize video downlinks in order to distribute them terrestrially, which coincided with the initial roll out and testing of Broadband cable modems. Between the needs of nationally distributed video, transit for high-speed internet access, and the birth of data centers for VOD / back-office systems, there was an obvious need for a national network. The decision was made to build a National Backbone that enabled more control over the products and services Comcast would deliver by acquiring the Indefeasible Right of Use to a 2-fiber national footprint. The footprint touched all 28 Comcast metro networks and had a presence in the majority of all the tier 1 cities that enabled easy connectivity to our transit provider(s). Since the fiber passed close to Comcast buildings, but not through Comcast buildings, short fiber laterals would need to be constructed to attach metros to the Backbone.

3.1. Comcast Backbone History

Each of the 28 metros chose 2 locations for Backbone connectivity, and core locations were also strategically selected in major cities. Each core location was given a core router for Comcast "internal video" and a core router for "external" or "internet" services. Although multiple "networks" were built from a routing perspective, fiber availability, and cost, drove the need for a single transport platform to carry all traffic and lines of business. The original Backbone was designed and built as a greenfield taking advantage of new MEMS based reconfigurable add, drop multiplexers (ROADM) technology. ROADMs were placed at each metro aggregation points, core location, and fiber branch. This created an 88 channel, fixed 50GHz, meshed network, that allowed channels to be turned up between ANY add/drop locations. A large portion of the footprint contained amplifier sites that didn't serve any add/drop purpose. These sites were built as amplifier only sites, unless optical regeneration was needed in which the site was built as an add/drop.

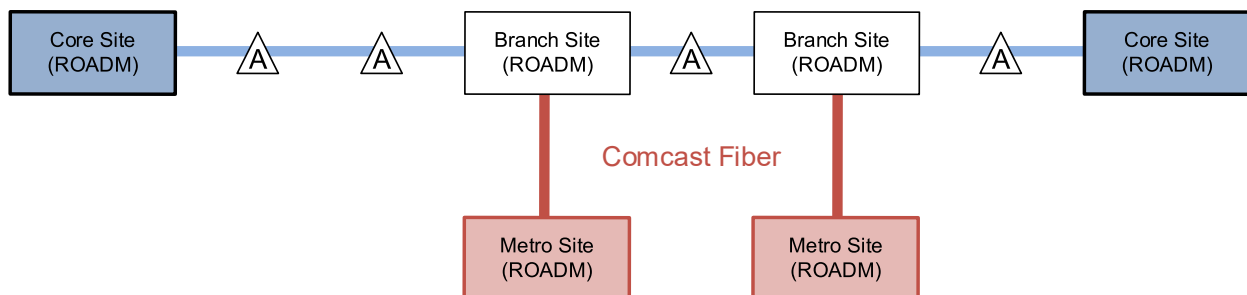


Figure 3 – Comcast Backbone Architecture - 1

Traffic was then turned up using 10G DWDM transponders. Initial calculations arrived at each metro requiring a 10G channel for “internal” Comcast video traffic, and a 10G channel needed for “external” data/internet traffic. To support this in the core, all core routers were also connected via 10G channels per supported service.

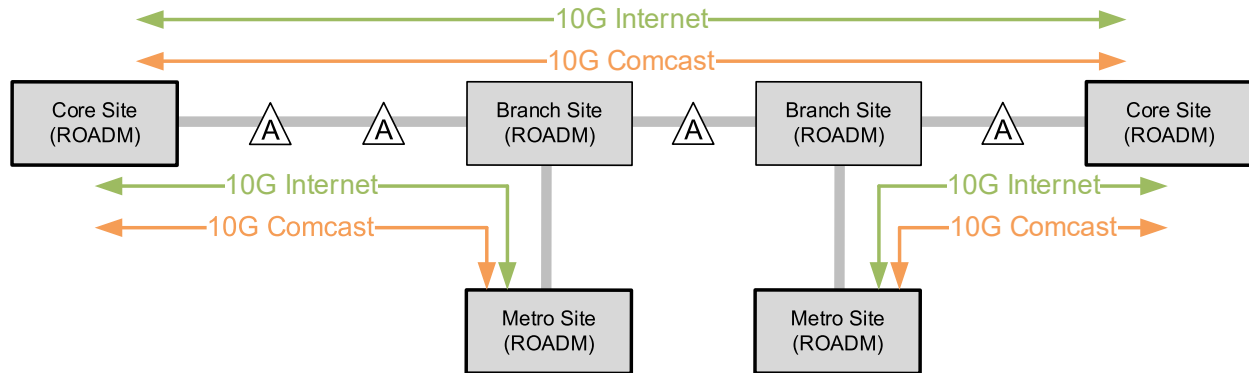


Figure 4 – Comcast Backbone Architecture - 2

Before the entire greenfield was even completed, it became VERY obvious that the exponential growth of the internet would quickly drive 2 changes.

First, Comcast needed to transition from leveraging a transit provider, to becoming a full-service provider. Core sites needed to move to locations that were more conducive to Settlement Free Interconnects. “Carrier Hotel” sites were chosen, and “Edge” router ports now faced offnet Settlement Free Interconnects and customers.

Second, the Compounded Annual Growth Rate (CAGR) growth far exceeded Comcast’s ability to deploy technology such as 10G transponders. We quickly realized that in order to support our growth that we would need to drive Optical Technology roadmaps to support higher Optical speeds. The move to 40G (2006) and then to 100G (2011) was required to manage our optical spectrum efficiently to avoid exhaustion on the original line system.

By 2015, the original fixed grid system was full and unable to utilize the new 400G transponders that required flex grid. Comcast chose to acquire more fiber and construct a 2nd line system instead of trying to upgrade the existing system. The 2nd backbone optical system was completed in 2018 and was almost exclusively 400G from the start with built in capabilities to expand to 800G and beyond.

3.2. Optical Innovations in the Backbone

The tremendous CAGR growth realized over the last 15+ years, drove many optical industry innovations in an effort to keep pace. In addition, these innovations had to be economically viable driving Watts per Bit efficiencies.

The primary obstacle that separates backbones from other optical networks is simply the fiber distance. Fiber Distance drove the initial 10G and 40G NRZ deployments to install over 40 optical regen sites across the country. The first big break in technology was the arrival of the

40G coherent modem. It not only quadrupled speeds per channel, but just as importantly, it more than doubled the reach. Cost per bit dropped, and with regens reductions dropping into the teens, network costs were reduced significantly. Within a few years, Comcast was then able to augment longer spans with Raman amplifiers, move to the industry's first 100G coherent modem, and reduce Optical regen sites to 4.

The second backbone built in 2017-2018 was built utilizing colorless direct attach (CDA), flex grid ROADMS, with a flexible add/drop structure. Combining low loss fiber and 2nd Gen Raman amplifiers, Comcast is now able to cross every link without regen at a minimum of 200G and maximum of 800G. Additionally, all shelves were deployed with L-band amplifiers, so L-band channels can be added in the future without any service interruption.

Finally, the integration of an OTDR into the optical line system must be mentioned. When your network spans 18,000 route miles, knowing exactly where an outside plant problem occurs is invaluable with the hopes of reducing the Mean Time to Mitigate. (MTTM)

3.3. Capacity of the Backbone

The most valuable commodity in National Optical Networks is long haul fiber and the efficient use of optical spectrum management. Fiber longevity within the access and metro networks is also important, but typically more feasible and economical than acquiring and deploying national fiber where such options are very limited. Any fiber, if available, usually comes at an exorbitant cost.

Comcast has historically been a proponent of partnering with the Optical Industry to drive key technology innovations and continues to collaborate on next generation Optical capacity and efficiencies innovations.

Table 1: Fiber Capacity

Year	Line System	Modem	Fiber Capacity
2005	50GHz Fixed	10G	880GB
2009	50GHz Fixed	40G Coherent	3.5TB
2013	50GHz Fixed	100G Coherent	8.8TB
2018	CDA Flex	400G Flex	12.8 - 25.6TB*
2021	CDA Flex	800G Flex	16.8 - 33.6TB*

* C-band only (all fibers are L-band capable which should double capacity when modems are available)

3.4. Future of the Backbone

There is no denying the need for long haul networks to interconnect metro networks and peers, both in footprint and offnet. Even with the overall CAGR slowdowns after the Pandemic, the need to scale efficiently and economically will not disappear anytime in near the future. Moving forward, it is critical to drive software, tools, automation and integration in addition to the speeds and feeds of the typical hardware innovations. It is important that as an Industry we continue to partner and drive these Optical technology evolutions in support of future network products and services.

3.5. Purpose Built Networks

Comcast traffic contains 2 basic types. The majority (>95%) is very predictable internet destined traffic between the metro aggregation shelf and the edge. The smaller portion is very unpredictable traffic that has varied source and destination points. The large portion emphasizes the need for lowest cost per bit, while the other demands flexibility over cost. A point-to-point network of a flexible CDA line system with state-of-the-art modems would really push capacity up, and costs down, but with almost no flexibility. Where a CDC network can move traffic between any location, but at a lower capacity, and at a higher initial cost. Since the original Backbone was full and reaching end of life, a replacement was needed. Comcast utilized the luxury of having more than 1 fiber pair and decided to build 2 purpose built networks instead of one Swiss Army knife of a network. This allows the cost-effective network to grow quickly, while preserving the life of the costlier one.

3.6. Power Consumption

Power consumption is an industry wide concern. However, this is amplified in peering locations (BB Core sites) specifically. Peering locations are needed to interconnect with other service providers, and they are independently owned by facility managers. This forces Comcast and other service providers to rent space and power. Network growth has forced these locations to fill and make space and power very scarce. If a space fills, a migration to a new facility is needed. To avoid (or delay) the high cost and complexity of moving a core peering location, driving power consumption down is a must. Comcast has seen watts per optical bit fall from 7.3 in 2005 to below 0.5 today. Unfortunately, growth has outpaced the reduction in power consumption, so this continues to be a hot topic.

3.7. IPoDWDM and Alien Waves

IP over DWDM, primarily led by ZR+ plugs, is getting serious testing in Comcast. Although they are not as applicable in the Backbone as the metros, there is still some key use cases. The Comcast Backbone network tried this in 2006-2007 with a 40G optic being placed in our core routers. Engineering limitations were understood and overcome, but operational limitations ultimately ended the experiment. The same operational shortcomings were identical to alien wave issues. Comcast continues to seek operational solutions both internal and external in hopes of taking advantage of the possible cost reductions of alien transponders and plugs.

3.8. Automation

Turn up of new bandwidth and services across large national networks can seem like it happens at a snail's pace. Equipment and resources tend to be needed in several locations, often very far apart. Planning, deploying, testing, and turn up requires a lot of time. This time is stretched even longer as they sit on long waiting lists seen in very large networks. Comcast has been able to streamline and automate some of the predictable, business as usual (BAU), bandwidth growth, but much more is needed. Commercial and wave services will, due to their unique and custom designs, amplify this need. Comcast continues to drive automation from planning to turn up and utilize more flexible equipment to drive down equipment installation times.

4. The Metro Network

Designed to support multiple lines of business (LOB) with a variety of bandwidth needs and traffic patterns, however, to say designed is a bit of a misnomer. This would imply the metro was conceived and deployed from day one in a greenfield manner to meet all the needs of today. What had actually happened is that over the last 25 years these networks came to be through mergers of smaller cable television (CATV) providers into multiple system operators (MSO). There were technology evolutions from passive transport systems with amplifiers (AMP) and multiplexers (MUX) to what is present today with reconfigurable add, drop multiplexers (ROADM) colorless add / drops, and flex grid technology. The type of data and rate has also dramatically changed. The earliest deployments were to support video to areas that were beyond the reach of long coaxial (COAX) trunk runs, followed by hybrid fiber-coaxial (HFC) and 1 gigabit ethernet (GbE) commercial customers. Today 100G to 800G wavelengths are the norm! The more than 60 metro networks at Comcast have now either been upgraded or planned to go to colorless, dispersion less designs with flex grid ROADMs optimized for coherent wavelengths. Additional paths have been added between sites (referred to as degrees) for resiliency and latency reduction as well. The Metro networks are unique in that they are intentionally not purpose built to support a single demand but are LOB agnostic. The Swiss Army Knife of transport at Comcast!




Figure 5 - The beginning of the Metro network. 3 independent CATV providers




Figure 6 - Metro network today. One MSO. Flex grid and ROADMs. The Swiss Army Knife

4.1. The 28 CRANs of Comcast

The largest consumer of metro transport bandwidth are the Comcast regional area networks (CRAN). These networks are the core of the internet protocol (IP) infrastructure consisting of a pair of large aggregate routers (AR), residential U routers (RUR) and commercial super U routers (SUR). Each edge site is U ring with a pair of RURs with one router connecting to one market AR and the other connecting on diverse path to the other AR. If there are commercial services at the site, the same topology is applicable but with SURs in addition to the RURs. Each CRAN has two ARs which are connected to two backbone connections on diverse routes. This topology ensures failures of equipment or fiber will not impact customers. They are called U rings because the shape of the letter U describes the topology. The top of each leg of the U is an AR location and the bottom being the RUR / SUR site.

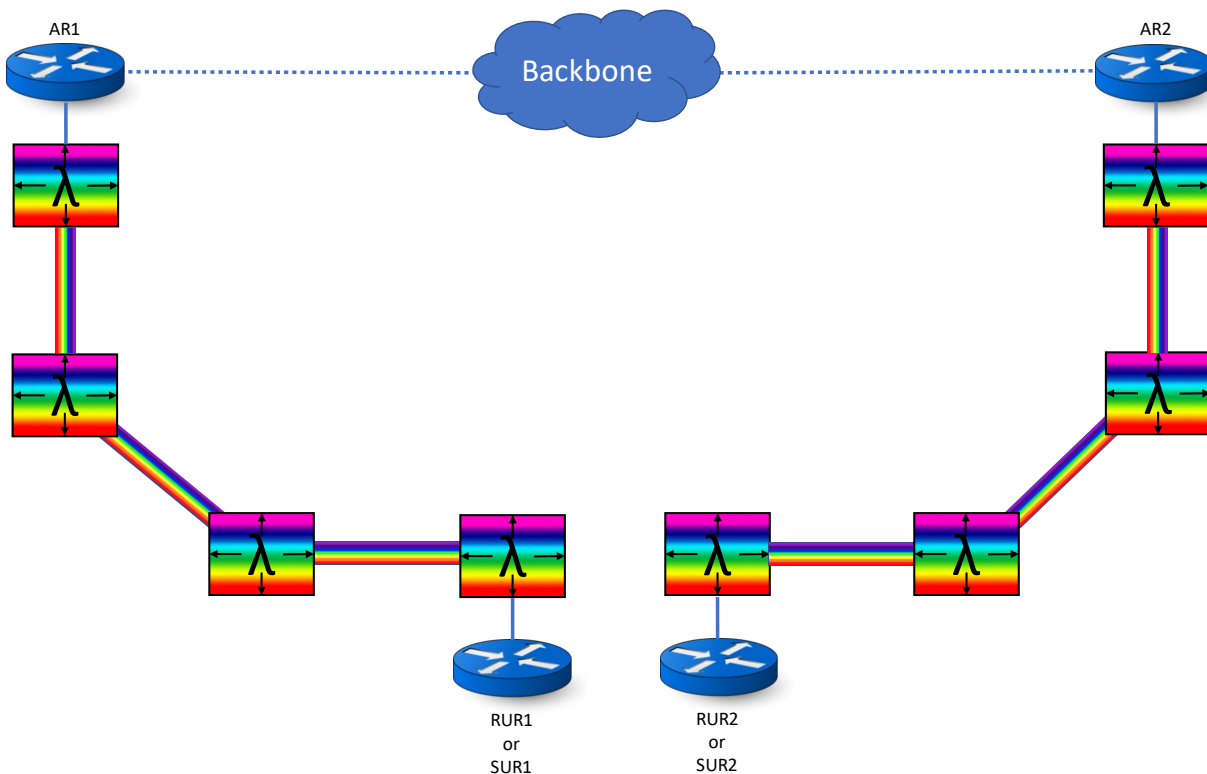


Figure 7 - U Ring Topology

4.1.1. In the beginning. Video on demand (VOD)

The Metro Transport Network tends to evolve to meet the needs of the routed network. The original iteration was dense wave division multiplexing (DWDM) and coarse wave division multiplexing (CWDM) uni-directional 10GbE optics placed directly in routers with MUXs and AMPs as needed. This was to support VOD and was also the birth of the U ring at Comcast. The routers we aptly named U routers (UR). A major difference in the topology compared to today is that that these URs to support VOD were configured hop by hop. Each router was connected in a chain. Since the circuit never had to go more than one span this allowed the use of point to point amplified optical links referred to as passive transport

4.1.2. Next up! High-speed internet (HSI).

If all routers need the same amount of capacity the hop-by-hop VOD architecture worked fine. With the introduction of high-speed internet (HSI) not only did the traffic need to become bi-directional but also the capacity needs of the routers began to become more fluid based on consumption of the product. For example, if there were five routers in the leg of a U all at 10GbE of demand and the fourth router in the chain grew and needed 20GbE then all the spans between the AR and the fourth router need to be augmented to 20GbE. This applied to both sides of the U. This was not a scalable model to throw away bandwidth along the way to reach a downstream router.

This led to the installation of active transport with ROADMs which allowed each router to get a direct connection to the ARs bypassing sites along the way optically. This first generation of transport had dispersion compensation, fixed grid wave selective switches (WSS), fixed grid MUXs and was optimized for 10GbE. Depending on the era this equipment could be either 50 gigahertz (GHz) or 100GHz spaced. This spacing is not critical for the 10G but would play a role in future evolutions.

As ROADMs were being installed and traffic could be steered at site optically it made sense to start breaking up large rings into smaller sub rings by building degrees between sites and including more remote sites into the network. These new shorter paths began to open the door for commercial services like cell backhaul (CBH).

4.1.3. 10 x the bandwidth and then some!

The next substantial change was to 100GbE interfaces on routers and the 100G coherent wavelengths required on transport. This was a huge leap! Fortunately, the metro transport was able to carry this traffic without upgrading the line system. A typical 100G coherent wavelength is ~37.5GHz wide. With most of the technology at this time being deployed with 50GHz spaced ROADMs and MUXs the 100G wavelength fit without issue. In fact, up to a 200G wavelength can fit in these networks.

At this point the transport was one step ahead of the routers. With 400G line rates right around the corner and the router interfaces to follow the transport needed to maintain this lead and be ready for 400G as soon as possible.

In an interesting twist some of the oldest ROADM transport networks were the first ones capable of supporting 400G. As mentioned above both 50GHz and 100GHz systems were deployed with the 100GHz typically the first generation. With a 400G wavelength being ~75GHz wide it fit in the oldest technology deployed allowing another entire evolution of bandwidth increases to be deployed without an upgrade to the photonics. This is however the end of the road. 800G is already being deployed and with its ~112GHz wide wavelength it can only be deployed using flex ROADMs and a colorless add / drop structure. Given the fixed 50GHz systems are not capable of 400G and the networks that can are the oldest, the standard has been set to only deploy new networks as colorless and Flex grid to ensure any size future wavelength would be supported. Having also learned that there can never be enough available capacity a decision has been made to design all new networks to be C+ L band capable as well. Adding the L band doubles the capacity of the fiber! A final benefit to deploying all colorless networks going forward is the availability of Layer 0 (L0) control plane. This sets the stage for further resiliency by offering protection and restoration at the optical level.

4.2. Optimizing further - Capacity vs. Connectivity

Now the network is well prepared for bandwidth expansion by being able to support any width wavelength and the addition of L band has doubled the capacity. The construction of new degrees and creation of sub rings improves resiliency, latency and provides new routes for commercial customers. This architecture is the new performance baseline. Can it be optimized further though?

Of the traffic on the CRANs the largest and most predictable is residential. RUR U rings are high capacity with many being over 1 terabit (Tb). Although the physical flow of the traffic is passing through several ROADMs the logical circuit looks like point-to-point. With the large and predictable growth of these U rings does it make sense to go back to the original VOD architecture and build large, simple point to point networks? The answer might be yes, but not exactly.

4.2.1. Internet protocol over dense wave division multiplexing (IPoDWDM)

Much of the focus on technology advancements has been on creating larger wavelengths that consume less spectrum, however the transport is now outpacing the IP platforms with up to 1.2TB line rates while the routers are just starting to adopt 400GbE. To best utilize and align the IP and optical platforms the answer could be IPoDWDM. Simply put, this is making the source of the DWDM wavelength a pluggable optic and placing it directly in the router, thus eliminating the transponder. Sound familiar? This is exactly how the first generation of transport to support VOD was deployed as described previously. This may make sense again but with several important caveats.

The optic that has been developed for this application is called a ZR+. It has a 400GbE rate and like other 400G transport interfaces has a spectral width of ~75GHz. It does have roadblocks to overcome before being deployed at scale.

4.2.1.1. Form factor and transmit power

The first two issues are related. They are form factor and optical transmit power. There are two current versions of the optic. C form factor pluggable (CFP) and quad small form factor pluggable (QSFP). The form factor is directly related to optical transmit power. The CFP is larger and can house the components required to achieve ~0 decibel-milliwatts (dBm) optical transmit power commonly seen in transponders. However, IP host platforms that use CFP are no longer common at Comcast. They have been replaced by devices that use the much smaller QSFP optic. The issue is given the smaller size of the QSFP optic it can only produce ~ -10dBm transmit power with the current technology. This is as much as 14dBm lower than transponders! To allow the high-power transponder and low power optic to coexist in the same ROADM they need to have a similar input. There are two options. Either turn down the transponder and greatly reduce its performance or amplify the low power ZR+ optic. Neither of the solutions are acceptable in a brownfield environment. Fortunately, both problems are remedied by high power ZR+ optics in a QSFP form factor becoming available in late 2022. In fact, Comcast is collaborating with an optical partner to be one of the first to get this optic in the lab in advance of the general availability release.

4.2.1.2. Reach and ROADM limitations

The third issue is the reach of the optic. The ZR+ optic not only has distance limitations but every ROADM it passes through adds penalties and decreases the reach. In the end all the ROADMs that were added to increase flexibility in steering traffic and reducing latency now could be preventing the use of these optics. One solution being explored is to build ROADM bypass express paths to decrease the number of ROADMs and increase the reach.




Figure 8 - ROADM impact on 400GbE ZR+ optics

4.2.1.3. Operational support

Finally, the largest issue is the operational and engineering support of these optics. By disaggregating the DWDM source from the optical line system the topology connection between the DWDM wavelength source and add / drop structure is no longer known. It is not possible to know what the pluggable is connected to. This splitting of the optical network into two domains essentially makes the link an alien. This was the case also in the VOD era however with only a handful of links to support per market it was manageable with static documentation. Today with thousands of links this is no longer possible. Comcast engineers and software developers and the industry in general are currently working on solutions to operationalize aliens which would make the use of ZR+ optics possible.

If these challenges can be overcome ZR+ optics do not replace transponders everywhere nor do the ROADM bypass express paths displace the multiple path options of many meshed degrees. Instead, it becomes yet another layer, optimizes further and solidifies the Metro Optical Networks role of being the do it all, Swiss Army Knife at Comcast.

4.3. Leaveraging access and metro together to support commercial services

Where the metro ends the access starts. Unlike other networks that are completely independent and do not converge the metro and access are different. They can and do and have a symbiotic relationship.

Comcast combines the metro optical network, optical transport network (OTN) tails over the access fiber, small optical shelves at customer sites called network terminating equipment (NTE) and dedicated commercial shelves at hub sites called wave integration shelves (WIS) to support multiple types of commercial business.

As seen in the above network evolution diagrams the metro transport has become more meshed and better able to support an any to any traffic pattern. The access has also gone through evolutions but what has stayed the same is the fact that it stops at the hub site. What has been found is that most commercial customers can connect to the access and be brought back to a hub site where metro is present, but how can they be connected to each other if their fiber terminates at different hubs? Or in the case of ethernet dedicated internet (EDI) how do they get to the router that provides internet access? This is where the access and metro can work together by using OTN tails.

OTN is taking an Ethernet (or other protocol) payload and encapsulating it in a wrapper to be carried across an optical transport network. It is completely transparent, and a standard governed by International Telecommunication Union (ITU) G 7.09. An OTN tail is transmitting that wrapped payload over dark fiber or a channel on a MUX over fiber. Using this tail and transponders (or Muxponders) in the NTE and WIS facing each other can extend the metro core over the access without using ROADMs or even AMPs.

This tail is then transmitted over the metro by a second card in the WIS with an optical transport unit (OTU) client to client connection to the card facing the NTE. This back-to-back configuration allows a DWDM OTN signal to face the NTE and the add / drop structure of the metro. The client-to-client connections have been made in the past as ethernet where a similar design was used for regens. In this case, keeping them OTN so that the circuit remains transparent from end to end has been selected as the standard. There is no conversion to ethernet and back to OTN. One advantage of this is that a trail trace identifier (TTI) remains intact from end to end even if the circuit crosses to another network or even to another vendor. You can transmit a trace on the A end and read it on the Z end. This is particularly useful for trouble shooting and topology verification.

This Unified Optical Architecture to Support Wavelength and High Bandwidth Ethernet Services is now a common way the metro and access work together at Comcast and is fully detailed in another paper being presented at the SCTE Cable-TEC Expo 22. [1]

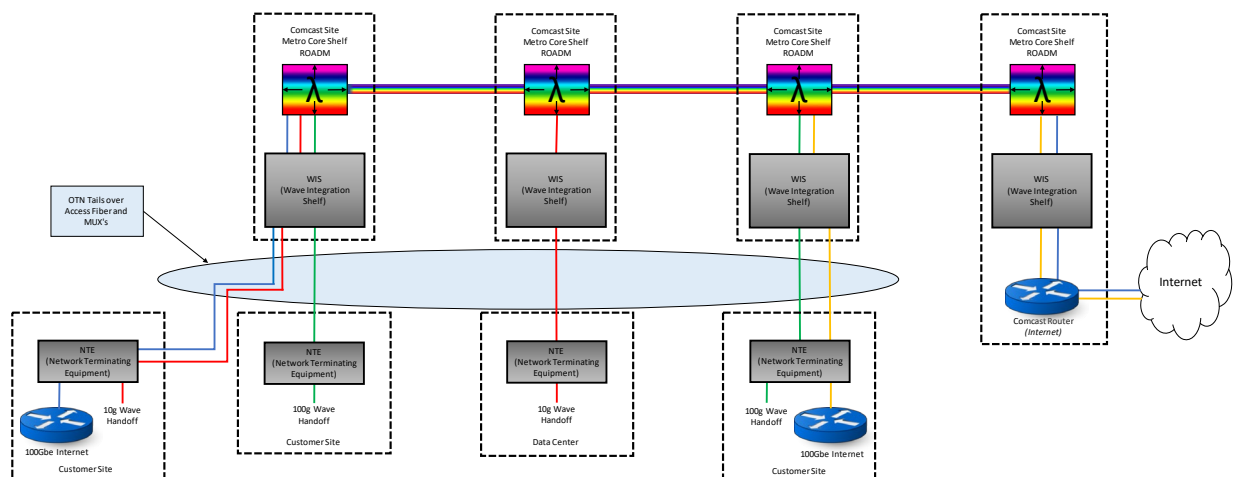


Figure 9- Metro and Access used together to support commercial customers

5. The Access Network

The Comcast Access network consists of a myriad set of architectures some organic, some acquired and many augmentations over the years. Some of these architectures are the traditional 1310nm/1550nm analog and digital/analog return nodes connected to traditional CMTS. Others are DWDM QAM Overlay solutions and yet others are the latest distributed access architecture (DAA) based nodes.

Comcast has pioneered the DAA that has virtualized the CMTS functionality and separated the PHY layer from the CMTS and distributed it out into nodes in the field. These are now called the Remote PHY Device (RPD) nodes. In Comcast virtually all of the traditional nodes and CMTSs will be replaced by the DAA in the next couple of years. For this reason, this paper will concentrate of the DAA system but will address other systems along the way as needed.

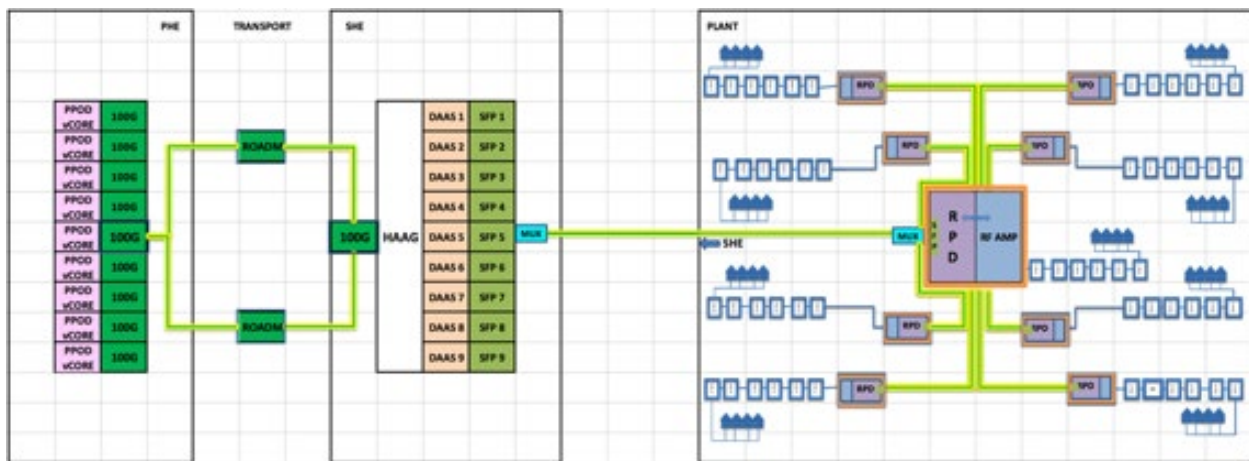


Figure 10 – Illustrating the DAA Network of Comcast

Shown above is a simplified version of the DAA system. vCMTS cores from the Primary headend are connected to DAA switches (DAAS) in the secondaries. Note here that these connections are typically done via the Comcast's own Metro networks described in the above sections. While the Primary to Secondary connections are based on 100G or higher capacity coherent optics, the connections from the DAAS on wards in the access domain are always 10Gbps DWDM based signals. The aggregating switches called HAAGs typically combine US signals and de-combine DS signals to the respective DAAS ports.

Once out of the DAAS, both the US and DS 10G wavelengths are multiplexed on a single strand of fiber and sent over into the field, where they are demultiplexed at an OSP enclosure location. For there, dual strands of fiber carry the 10Gbps signals over to the various RPDs in the immediate location. The distance from the Secondary to the field OSP mux is called the trunk fiber, while the fiber from the OSP Mux location to the various RPDs is called the distribution fiber. On the trunk fiber Comcast has a 100GHz ITU WL Plan standard that has adjacent odd and even pairs of ITU channels serve each node beginning at ITU 61/60 pair to ITU 15/14 pair. Thus, all in all, 24 pairs of wavelengths are available on each trunk fiber capable of serving 24 different remote devices.

While distribution fibers are new fibers sometimes put in explicitly for newer nodes, the trunk fibers are in extremely short supply as they were installed decades back and augmented over the many years to serve nodes that sprung up as population centers sprung up. It is for this reason that all DAA trunk fibers are bi-directional as that relieves pressure on the trunk fibers and promotes fiber efficiency across the organization.




Figure 11 – Typical Secondary to Node distances in the Access Plant

Generally, primary headends also supply node signals directly from their location to the local nodes as well. It is for this reason that the distance between a headend and the node is around 10km, but a fair number of nodes extend up to 40km. This is illustrated in a Comcast survey of optical links from a couple of year back [2]. In general, distribution links are on average 0.5km and as discussed construction crews may build this fiber for some of the newer nodes which arise due to node splits or fiber extensions from the OSP Mux. For this reason, there are multiple fibers connecting OSP Mux to fiber nodes. It must be mentioned here that most of access field infrastructure requires Industrial-Temp operation, this means that all optics actives and passives must be qualified up to 85C. Therefore, the field trunk fibers are standardized to 100GHz spacing utilizing thin-film based optical filtering and all the SFPs used in remote devices are Industrial-Temp qualified tunable SFPs.

Our fiber to the home (FTTH) offering comprises initial RFoG deployments but that is expected to move over to 10G EPON offerings for all the newer builds. The 10G EPON DS is at 1577nm while its US is at 1270nm. With the advent of SFP based OLT and ONUs, it is expected that current RPD nodes that have the capability to received 10G signals will receive 10G signals from their respective DAAS ports and convert via the SFP based OLT to the 10G EPON standard and serve the FTTH customer base. A lot of attention has been paid to enable the convergence of conventional RPD devices that can serve residences via standard HFC plant AND the FTTH customers thru the remote OLT (also called virtualized broadband network gateway or vBNG).

The vBNG line puts out 1577nm in the DS and accepts 1270nm in the US one single fiber and thus enables centralized or decentralized splitting as it serves homes via single fiber stands.

5.1. Converging the Access Network

The previous has been a simple description of Comcast network, but while the direction of digitization and distributed access is clear Comcast still has various older architectures comprising analog DS wavelengths and potentially analog or digital US wavelengths and they must all be accommodated as the DAA transition takes place. In addition, one interesting aspect of access plant is that most of the businesses that require separate optical wavelengths are also interspersed among the residential customers. The ability to have analog and digital WLs coexist on the same fiber has been described in earlier papers and is summarized here. Briefly, full spectrum analog WLs in the C-Band require a comprehensive WL plan to minimize optical non-linearities such as 4WM, XPM, SRS and overcome dispersion and imperfections of the optical passives. For this reason, they are not on a uniform spacing, but rather non-uniformly spaced. In Comcast, we have carved out Analog and 10G WLs in DS and, 10G and DRT WLs on the US on the same single strand of fiber.

Our current business base is a mix of single and dual fiber CWDM and DWDM connections that reflects the years of organic and acquired growth. Many of our current business customers that require fiber to their premises require 1G to 10G services (discussed in the Metro section of this paper) and could be in the vicinity of existing OSP Mux enclosures. For a go forward plan, since these are digital signals, these customers can be served by the same strand of trunk fiber as would connect analog and digital nodes. This is an important way to relieve stress on our fiber assets by enabling a more effective use of our fiber assets. In addition, it enables a common way to continuously and pervasively monitor optical assets which is described in a subsequent section.

But what is interesting in the figure below is that we have now been able to incorporate Coherent optics in the access plant in much the same way as direct detect has been deployed. This type of convergence requires game changing technology of bidirectional dual laser pluggable optics conceived and implemented in Comcast in recent years and is described in this section.

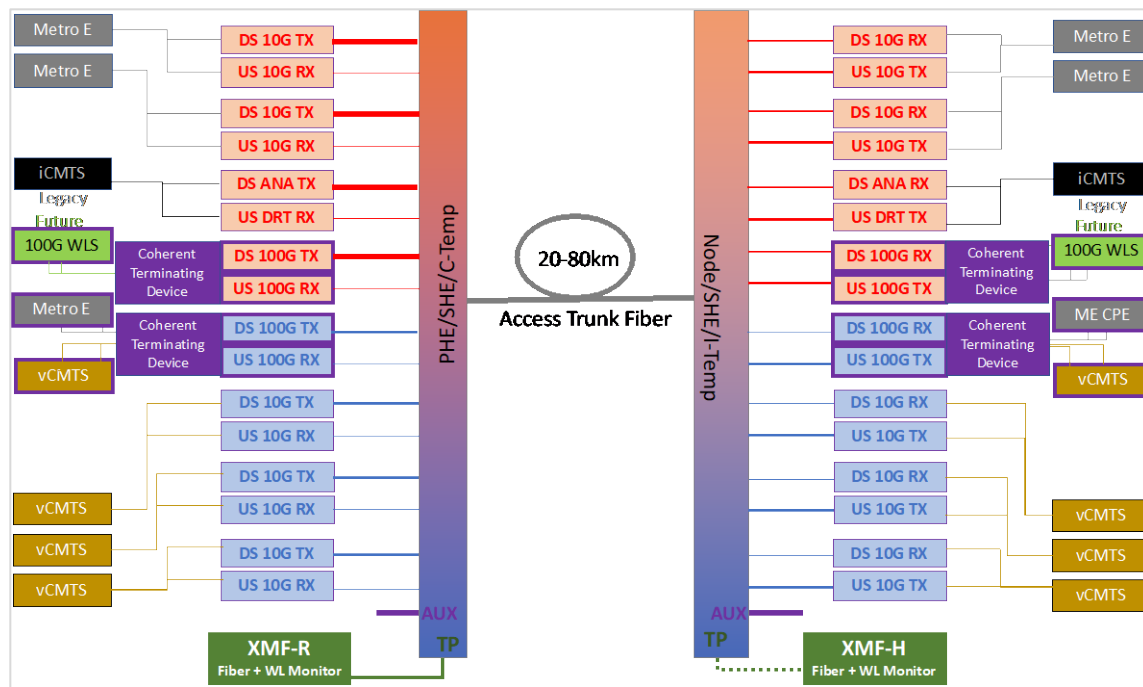


Figure 12 – Illustrating the Converged Comcast Access Network

Since ROADM cascades are not used in our DAA access networks due to modest fiber lengths, the design of high-capacity (100G – 800G) solutions could be elegant since there is no degradation due to these devices. Furthermore, these links can support bidirectional transmission quite easily since there are no non-reciprocal elements in the fiber path. Indeed, the relaxed relatively modest fiber reach required in DAA enables us to innovate on optical devices as well as reduce cost, oftentimes by innovating on the trailing edge of technology, which is discussed next.




Figure 13 – Converging Direct detect and Coherent bi-Directional Wavelengths

Presented above is a measured plot of the various direct detect and coherent optical signals overlaid with the optical passband of typical filters being deployed in Comcast [3]. It is seen here that 10Gbps direct detect as well as 100Gbps QPSK Coherent all the way to 400Gbps 16QAM Coherent optical signals fit within the optical passband of deploying optical filters. Although there is little concern for data rates up to 400Gbps, we still need to further investigate the case for 800Gbps where, even with higher QAM modulation orders and with higher OSNR requirements, it is expected that the occupied bandwidth will start to encroach into the WDM passband and impose power penalties that need to be considered, particularly over temperature. More generally, a skillful use of optical passives express/upgrade ports with a contiguous C-Band can easily be dedicated for higher capacity/bandwidth modulations or for flex grid type applications.

5.1. Dual Laser Bidirectional Coherent Transmission

Coherent optics relies on having a laser at the transmitting end send out phase and amplitude information in the form of QAM constellations to a receiver. At the receiver, the incoming signal is ‘beat’ up against a laser that has in effect the exact wavelength of the transmitting laser and it teases out the phase and amplitude information out. Since information is coded in phase and amplitude, as opposed to only amplitude in direct detect systems, the information carrying capacity of the coherent system is vastly superior to that of direct detect systems for the same amount of bandwidth available.

To reduce cost and also to ensure that the wavelengths are identical, most of the plugs available in the market use just one laser in each plug and split its light to do double duty, in effect to act as the encoding transmitters and the decoding receiver. This is illustrated in the top half of the picture below.

The industry has gotten quite good at this type of design and has successfully used it to provide 100s of Tbps over 1000s of kms in core and metro networks. Unfortunately, this approach always requires two separate fibers, one for downstream and another for upstream information. When a fiber can be fully loaded with data, say for 100s of Tbps, this approach seems reasonable. But when the fibers are unable to be fully loaded, like what happens in access networks, an approach of this kind is very inefficient with fibers and can cause severe fiber exhaustion.

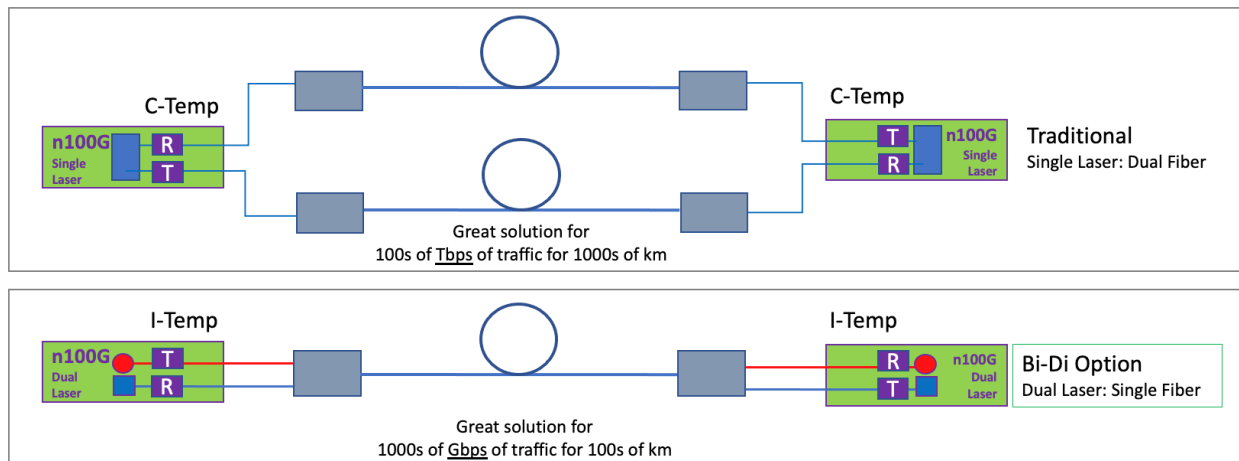


Figure 14 – Illustrating Bi-Directional Coherent systems

It is for this reason that a new approach had to be thought of for accommodating Coherent optics in the access domain and enable high-capacity wavelengths to feed our residential DAA systems and as standalone high-capacity wavelengths for our business customers.

We note here that the use of optical circulators is a valuable option and could potentially allow the same wavelength to traverse in opposite directions, but that approach can be prone to reflections in the fiber plant and also to fiber Rayleigh backscattering that could limit the link budget. A more robust solution is needed for wide bi-directional deployments that envision the use of existing fibers with other bi-directional signals already running on it.

Such a robust solution to this problem is to have a separate laser for the transmitter and one for the receiver at each end. Doing this would enable the two wavelengths to be multiplexed (or combined) together and transmit bidirectionally (Bi-Di) on the same single strand of fiber. In fact, many Coherent optical wavelengths could be multiplexed on the same fiber along with direct detect wavelengths thus leading to a converged system that was described in the previous section.

Interestingly, low earth orbit (LEO) satellite communications systems have also come up with the same requirement of dual laser bi-directional transmission. In free space, coherent optical modules communicate with adjacent satellite and hop signals from satellite to satellite to minimize the need for ground stations. But the vast speeds of the LEO satellites and their relative speeds to other satellites create the familiar doppler shift in frequency sufficient to thwart the use of a single laser design. So, in this effort of defining a dual laser plug, we have had the unusual conjoining of space requirements as well.

With this in mind, Comcast have specified dual laser bidirectional fully C-Band tunable industrial temperature CFP2 plugs for use in access networks much in line with the CableLabs point to point Coherent specification [4]. Current systems support 100Gbps rates with QPSK modulation at 32GBaud with a 25dB link budget, enough to support point-to-point links up to 100kms. As mentioned previously there are no ROADMs in the access plant and there are no

non-reciprocal elements in the system, this allows for relaxation of certain specifications and needs for chromatic dispersion compensation that enable these new innovative devices to be cost effective as well.

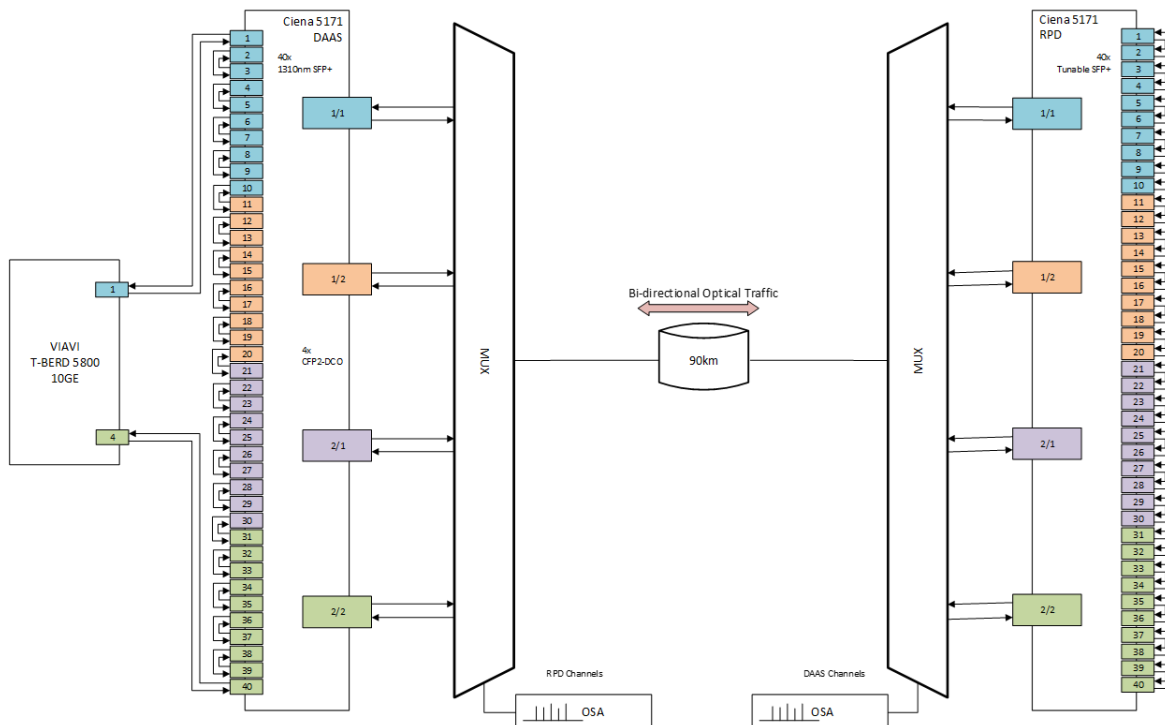


Figure 15 – Hybrid Loop testing the Coherent Optics in the Access Plant

While 100G coherent systems can exist on their own and support individual wavelength services for our commercial customers in the converged access architecture described in the previous section, a particularly important application of the dual laser bi-di system is to support DAA systems. Often these DAA systems have to operate out of smaller OTN or secondary cabinets that make direct metro transport system connections difficult due to limited critical infrastructure. And adding fiber to the trunk line is an expensive and time-consuming effort. In such cases, we have developed a system that takes in 10 of the 10Gbps streams and electronically converts to 100Gbps stream that modulates the CFP2 which is then multiplexed to a single fiber and sent over to the secondary. At the secondary the 100Gbps stream is split back to its constituent 10Gbps streams and serves DAA nodes from there.

To test this system for its robustness, we set up a hybrid loop test where 10Gbps signal was injected into one of the 10Gbps port that then entered the 100Gbps CFP2 which then traversed 90km of SMF fiber and was converted to its constituent 10Gbps stream and looped back via the CFP2 and traversed the 90km. In this way the traffic traversed $90\text{km} \times 2 \times 40 = 7200\text{km}$ before closing the loop. In our tests this was done with zero errors and with latency that was predominantly dictated by the optical fiber time of flight.

The system designed and shown above can support up to 40 10Gbps streams each way that modulate 4 of the 100Gbps streams in each of the two locations. In effect, 6 such systems can be optically multiplexed together and serve up to 240 10Gbps Nodes all on a single strand of fiber. As we have said before these our efforts at innovations will be to extend to higher capacity systems cost effectively over the coming years.

5.2. Convergence Continuous Pervasive Monitoring

Our current business base is a mix of single and dual fiber CWDM and DWDM connections that reflects the years of organic and acquired growth. Many of our current business customers that require fiber to their premises require 1G to 10G services (discussed in the Metro section of this paper) and could be in the vicinity of existing OSP Mux enclosures. For a go forward plan, since these are digital signals, these customers can be served by the same strand of trunk fiber as would connect analog and digital nodes. This is an important way to relieve stress on our fiber assets by enabling a more effective use of our fiber assets. In addition, it enables a common way to continuously and pervasively monitor optical assets which is described in a subsequent paragraph.

In the figure below we have shown a converged system that is a combination of residential and commercial services, incorporating coherent and direct detect systems, all capable of running on a single fiber. Since all services can be multiplexed on a single fiber, the current series of optical passives in Comcast have consolidated test ports that provide access to view forward and return wavelengths of live links. Furthermore, the test ports also allow unrestricted 1611nm access for use of OTDRs to shoot over live links and provide fiber impairment and cut information.

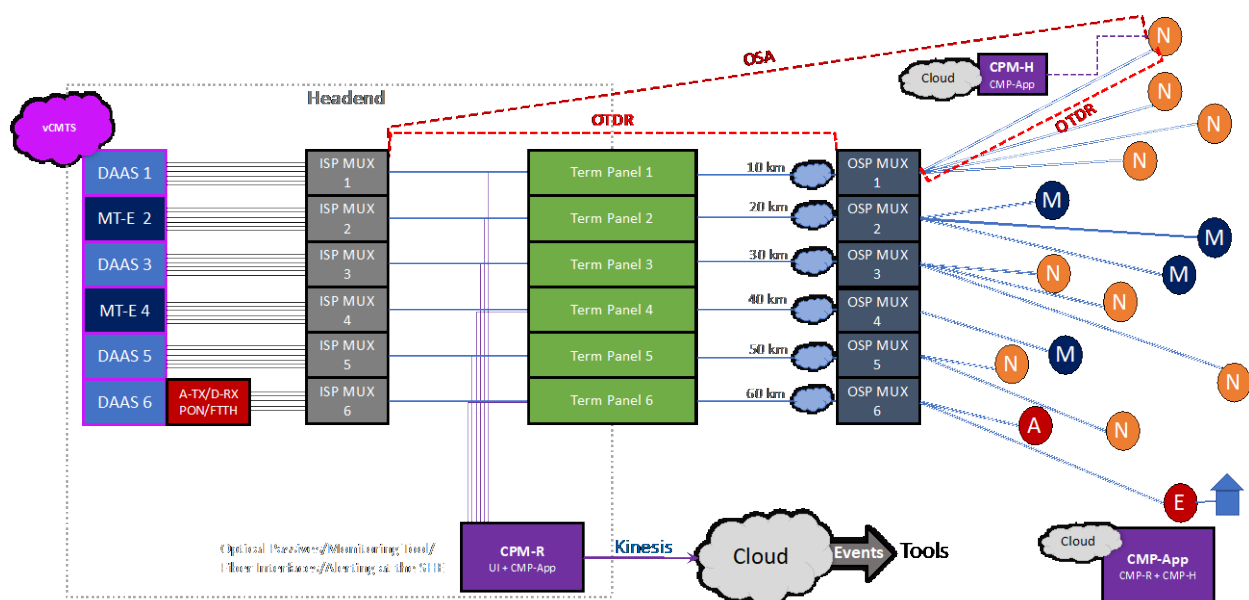


Figure 16 – Illustrating Continuous Pervasive Monitoring

Detailed description of the continuous and pervasive monitoring paradigm in Comcast was described in earlier SCTE papers. By way of summary, all optical passives are connected to a continuous monitor comprising an OSA and OTDR and an optical switch. This arrangement

continuously monitors all connected fibers in a round robin fashion and generates fault information comprising either fiber impairments or individual wavelength impairments on average within 90 seconds. The entire information set is sent into the cloud and alarms and events are sent to fix agents.

To promote a common language between headend and field technicians, a version of the monitor is available as a handheld device that has the same cloud connectivity as the headend version. Connecting the headend unit or the hand-held unit or both to the cloud now enables the entire region, division or company tie in and view all events in real time and strive to achieve closure on important customers impacting issues. This is a real game changer with the net effect of providing continuous monitoring on residential, commercial, direct detect and coherent optical signals across the access network.

6. Conclusion

In this paper, we described for the first time an end-to-end view of our optical network including the core, metro, and access layers. A brief history of the organic and acquired properties and technologies at play in Comcast was described. At the core and metro, we have increased capacity with a move towards flexible 400G connections and continue to reduce latency and enhance reliability with an infrastructure that meshes color-less, direction-less, and contention-less reconfigurable multiplexers thru to each of our headends. At the access layer that connects these headends to customers and businesses, we discussed capacity increases with our move to all-digital fiber links and the distributed access architecture paradigm. The vision shared here is part of a larger commitment that leverages Comcast's technical and business expertise to deliver cutting edge and reliable services to our residential and commercial customers combining a robust and scalable core and metro network with a converging range of access networks ranging from wavelength services to traditional HFC to FTTH solutions.

Abbreviations

4WM	Four Wave Mixing
AMP	Amplifier
AP	access point
AR	Aggregate Router
bps	bits per second
FEC	forward error correction
HCF	Hollow core fiber
Hz	hertz
K	kelvin
SCTE	Society of Cable Telecommunications Engineers
AMP	amplifier
AR	aggregate router
CATV	cable television
CDA	Cororless Direct Attach
CBH	cell backhaul

CFP	C form-factor pluggable
CMTS	Cable Modem Termination System
COAX	coaxial
CRAN	Comcast regional area network
CWDM	course wave division multiplexing
DAA	Distributed Access Architecture
DAAS	Distributed Access Architecture Switch
dBm	decibel-milliwatts
DRT	Digital Return Transmitter
DS	Downstream
DWDM	dense wave division multiplexing
EDI	ethernet dedicated internet
FTTH	Fiber To The Home
G	gig
GbE	gigabit ethernet
GHz	gigahertz
HCF	Hollow Core Fiber
HFC	hybrid fiber-coaxial
HSI	high speed internet
IP	internet protocol
IPoDWDM	Internet protocol over dense wave division multiplexing
ITU	International Telecommunication Union
L0	layer 0
LOB	line of business
MEMS	micro-electromechanical system
MTTM	Mean Time To Mitigate
MSO	multiple system operators
MUX	multiplexer
NRZ	Non-Return-to-Zero
NTE	network terminating equipment
OLT	Optical Line Termination
OSP	OutSide Plant
OTDR	Optical Time Domain Reflectometry
OTN	optical transport network
OTU	optical transport unit
PHY	Physical layer
QAM	Quadrature Amplitude Modulation
QSFP	quad small form factor pluggable
RFoG	Radio Frequency (RF) over Glass
ROADM	reconfigurable optical add drop multiplexer

RPD	Remote PHY Device
RUR	residential U router
SRS	Stimulated Raman Scattering
SUR	super U router
TB	terabit
TTI	trail trace identifier
UR	U router
US	Upstream
vBNG	virtualized Broadband Network Gateway
vCMTS	Virtualized Cable Modem Termination System
VOD	video on demand
WIS	wave integration shelf
XPM	Cross Phase Modulation
WL	Wavelength
WSS	wave selective switch

Bibliography & References

- [1] It is 10PM, Do You Know Where Your Wavelengths Are – Continuous and Pervasive Monitoring of Optical Assets in the Access Domain, Venk Mutalik, Dan Rice, Rick Spanbauer, Simone Capuano, Rob Gonsalves, Bob Gaydos, SCTE 2020
- [2] Accelerating the Virtualization: Introducing Hybrid Fiber Shelf into the Mix, Venk Mutalik, Bob Gaydos, Dan Rice, Jorge Salinger, SCTE 2020
- [3] Approaching the Universal Speed Limit: Introducing Hollow Core Fibers for Low Latency and High Capacity, Venk Mutalik, Amarildo Vieira, Bob Gaydos, Elad Nafshi
- [4] Cable Labs point to point Coherent Optics specification

Planned Maintenance Tool (PMT): A Data-Driven Approach to Recommending the Best Time for Planned-Maintenance

A Technical Paper prepared for SCTE by

Pete Quesada

Sr. Principal Engineer
Comcast Cable
4100 E Dry Creek Rd, Centennial, CO 80122
720.692.8683
Pete_Quesada@Comcast.com

Julianne Heinzmann

Software Developer & Engineer
Comcast Cable
1800 Arch Street, Philadelphia PA, 19103
267.586.7928
Julianne_Heinzmann@Comcast.com

Nishesh Shukla

Software Developer & Engineer
Comcast Cable
4100 E Dry Creek Rd, Centennial, CO 80122
303.263.8121
Nishesh_Shukla@Comcast.com

Resmi Vijayan

Software Developer & Engineer
Comcast Cable
1800 Arch Street, Philadelphia PA, 19103
817.908.7353
Resmi_Vijayan@Comcast.com

Mike O'Dell

Director Network Maintenance
Comcast Cable
Virtual Location
412-417-0481
Michael_Odell@cable.comcast.com

May Merkle-Tan

Lead Machine Learning Researcher
Comcast Cable
4100 E Dry Creek Rd, Centennial, CO 80122
Heng-RuMay_Tan@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
1.1. Problem Statement	4
1.2. Proposed Solution	4
2. PMT Core Components	5
2.1. Data Sources.....	5
2.1.1. Customer Usage Information (CUI)	5
2.1.2. Device to Network Mapping (DNM)	5
2.1.3. Plant Topology Information (PTI)	5
2.1.4. Service Interruption Events (SIE).....	5
2.1.5. Customer Interactions (CI).....	5
2.2. Data Processing	6
2.2.1. CUI data filtering for test region	6
2.2.2. Selected Accounts for Recommendations	6
2.2.3. SIE Data Processing for Assessment.....	7
2.3. Recommendation Algorithm.....	7
2.3.1. Data Analysis	7
2.3.2. Ranking-Based Algorithm	9
3. Assessment.....	13
3.1. Aim	13
3.2. Method	13
3.3. Findings.....	15
3.3.1. Comparison of REC and DNR delta_ratios	15
3.3.2. Time-of-day Effects	23
3.4. Assessment Summary	24
4. Application.....	25
4.1 Architecture	25
4.1.1. Overview	25
4.1.2. Process Flow.....	27
4.1.3. Grouping Service	28
4.1.4. User Interface.....	29
5. Discussion	30
5.1 Learnings from the field and insights from the assessment.....	30
5.1.1. Trial	30
5.1.2. Learnings and Observations	30
5.1.3. Considerations	31
6. Conclusion.....	31
Abbreviations	32
Acknowledgments	32

List of Figures

Title	Page Number
Figure 1 - Data Processing Workflow for the PMT Recommendations Process	6
Figure 2 - How CUI data are aggregated to derive proportional usage for account and selected accounts	8
Figure 3 - Scoring of best and worst hours	9
Figure 4 - PMT Algorithm Formula.....	13

Figure 5 - Three different CI baselines were considered in our assessments.....	14
Figure 6 - CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with prior 24hrs as a baseline	16
Figure 7 - Non-self-service CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with prior 24hrs as a baseline.	17
Figure 8 - Cost-incurring CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with prior 24hrs as a baseline	18
Figure 9 - CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with SIE_ImpactStartHr in prior 7 or 14 days as a baseline.....	19
Figure 10 - Non-self-service CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with SIE_ImpactStartHr in prior 7 or 14 days as a baseline.	20
Figure 11 - Cost-incurring CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with SIE_ImpactStartHr in prior 7 or 14 days as a baseline.	21
Figure 12 - Distribution and relative proportion of SIEs for each PMT category.	24
Figure 13 - Overview of PMT Application and its different components.....	26
Figure 14 - High-Level Overview of PMT Recommendation Service	27
Figure 15 - Example groupings determined by Grouping Service.....	28
Figure 16 - Example PMT application recommendations	29

List of Examples

Title	Page Number
Example 1 - Calculating Proportional Usage per Account.....	8
Example 2 - Determining Best and Worst Hours according to CUI	10
Example 3 - Proportional Network Usage According to Group.....	10
Example 4 - Calculating Proportional Users	10
Example 5 - Calculating Account Level Prediction Weights	11
Example 6 - Calculating Prediction Weights	11
Example 7 - Final Recommendation Values.....	12

List of Tables

Title	Page Number
Table 1 - Derivation of delta_ratios for comparison across SIE categories using 24hrs prior CIs as a baseline.	22
Table 2 - Derivation of delta_ratios for comparison across SIE categories using CIs in the prior 7 days at the same impact start hour as a baseline	22
Table 3 - Derivation of delta_ratios for comparison across SIE categories using CIs in the prior 14 days at the same impact start hour as a baseline	23

1. Introduction

1.1. Problem Statement

Planned maintenance is a daily activity for any number of complex systems, including cable plants. It is important to think of a cable plant as a living, breathing organism that requires care and feeding, involving replacing parts that are continuously exposed to the elements. Repairs often include identifying problems, making repairs, and replacing parts while temporarily interrupting a customer's service. A service interruption event (SIE) averages between five to ten minutes. In an ideal world, all SIEs would be performed during the evening maintenance window, but in practice, most short-duration SIEs must be performed outside of this maintenance window. Currently, SIEs are scheduled without the benefit of knowing which hours would have the least or the highest amount of subscriber impact. This information would be invaluable in optimizing the ideal time to perform an SIE.

We set out to find if a data-driven system could be developed to determine the best time to conduct SIEs. Performing SIEs during times when they will have the least impact on subscribers would not only provide a better subscriber experience but also could potentially cut the expenses incurred by responding to customer interactions (CI), such as calls to our care agents, unnecessary truck rolls, chat sessions, and other triaging events.

1.2. Proposed Solution

Our research involved identifying data that would have sufficient signal to indicate the least and the most impactful times to perform an SIE for the set of subscribers that each SIE would impact; we chose to calculate recommendations on each subscriber's hourly high speed data usage, which we will refer to, here, as customer usage information (CUI).

The main objective of our research was to develop a Planned Maintenance Tool (PMT) to assist with field operations. The algorithm that drives the PMT evaluates the hourly customer usage information (CUI) for each set of customer accounts that an SIE will impact. Then the PMT returns the hours when the SIE will be the least impactful toward those customers. We assessed the validity of our algorithm with historical SIEs and corresponding CI data from a geographic area that we will refer to as the 'test region.'

We have found that CIs typically increase when there is an unexpected SIE. We theorized that if we could create an algorithm to identify the best time(s) to perform an SIE, we would see a less severe CI increase around the hour the SIE is performed. Our assessment, though limited in breadth, appeared to follow this expected trend, and the findings pinpointed subsets of CIs we could monitor and assess periodically for financial and customer impact.

In the following sections, we discuss the constraints that motivated the initial version of the PMT user application and the refinements we think would be necessary further to improve the user experience and reliability of the PMT.

2. PMT Core Components

At its core, the PMT tool comprises two components that provide the data-science backbone of the system. They are the data processing and the recommendation algorithm, which are described in detail in this section.

2.1. Data Sources

It was posited that the best time to perform SIEs with the least impact on customers would be the hours when the least amount of data was consumed. We aggregated data from the following sources to collect a good source of data consumption for a group of homes.

2.1.1. Customer Usage Information (CUI)¹

Customer usage information is aggregated, in bytes, for upstream and downstream traffic for each DOCSIS (Data Over Cable Service Interface Specification) -capable device at hourly intervals.

2.1.2. Device to Network Mapping (DNM)

We needed to map individual MAC addresses to the network elements, including nodes, regions, and CMTS (Cable Modem Termination System). DOCSIS-capable network infrastructure allows the implementation of a system that polls all devices six times a day to check for any impairments, noise, and other factors. This data also records all devices connected and active in the network, along with their mapping to customer account numbers, location, and elements in the network.

2.1.3. Plant Topology Information (PTI)

The plant topology information includes data from various source systems to provide a hierarchical view of multiple elements in the network, such as head-end devices, CMTSs, RF cables, power supplies, taps, nodes, drop cables, and customer devices in the network. Although not used in the initial assessment, plant topology will be used by the developed application.

2.1.4. Service Interruption Events (SIE)

Service interruption events are initiated when an interruption in the plant is needed to correct RF system impairments, if the plant needs to be disconnected to replace a network component, or if periodic maintenance needs to be performed. Data about service interruption events is recorded and includes the time they occurred, the length of time that service was interrupted and a list of accounts that were affected.

2.1.5. Customer Interactions (CI)

Customer interaction events are identified as indicators of the impact SIEs could have on customers. Customer interaction events consist of logs of activity in a customer's timeline. They include billing, communications, customer chats, customer request tickets, inbound calls, tech appointments, speed tests performed, and equipment orders. The events selected for this analysis were limited to inbound customer

¹ We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

2.2. Data Processing

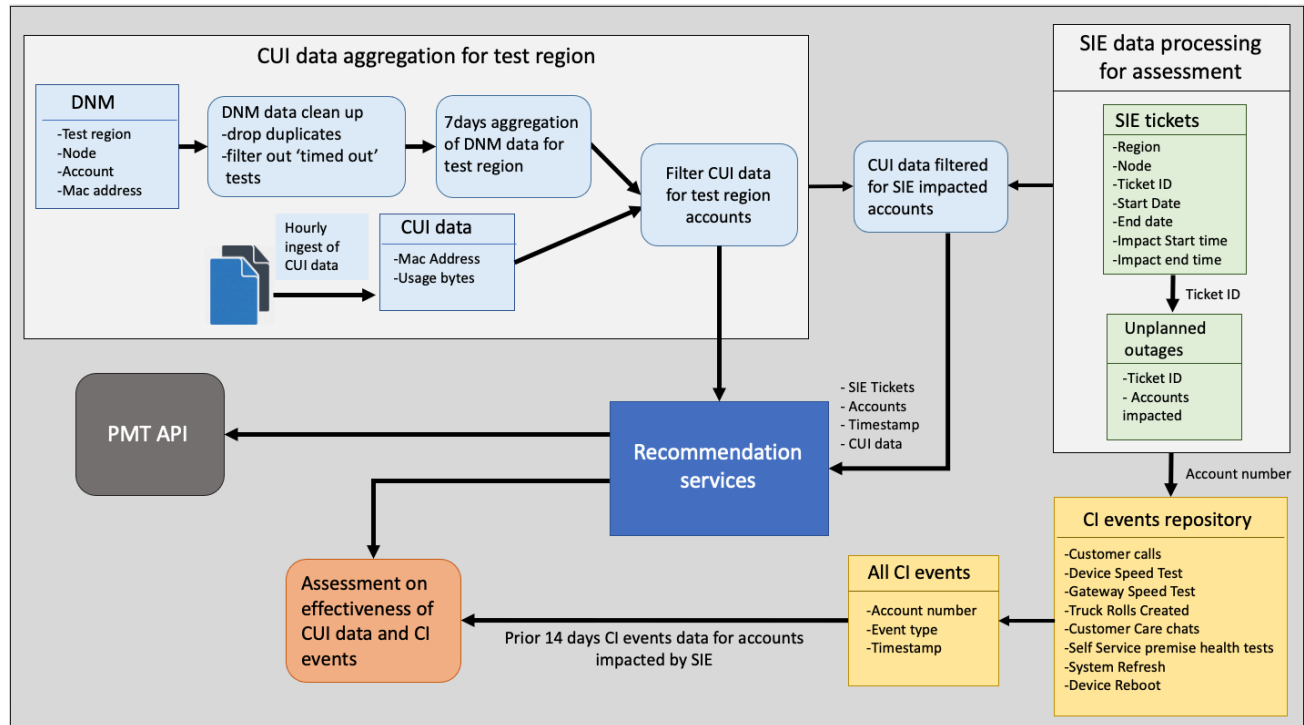


Figure 1 - Data Processing Workflow for the PMT Recommendations Process

2.2.1. CUI data filtering for test region

Daily polls of the DNM from the test region are collected and cleaned to filter out ‘timed out’ and inactive devices. A 7-day aggregate of DNM data is performed to accommodate account additions or deactivations changes. This seven-day aggregated DNM is then used to filter the stream of hourly CUI. We derive an hourly aggregate of CUI at account levels for the total 24 hours daily, which serves as the input into the PMT algorithms for any accounts (households) likely to be impacted by a network SIE.

2.2.2. Selected Accounts for Recommendations

Based on CMTS and node segment ID from PTI data, the geolocation latitude/longitude for customer service addresses are derived. The latitude/longitude information is passed to a Grouping Service (see section 4.3.1) to find all field topology information such as cables, taps, buildings, and addresses inside the node boundary. Accounts are then formed into multiple groups based upon the topology information by performing a fuzzy comparison of street addresses. Additionally, it was found useful to resolve some addresses by comparing their latitudes/longitudes using proximity. The account groupings² are further used in generating recommendations (See section 2.3.1).

² In Section 3 where we detail the assessment of the CUI and PMT algorithm, these “account groupings” are simply the set of SIE-impacted accounts, derived directly from SIE logs.

2.2.3. SIE Data Processing for Assessment

We collected SIE data for the test region for this assessment. The accounts impacted by each SIE are derived from an ‘unplanned outages’ dataset. For each account affected by an SIE, we collected customer interaction events from 14 days before the onset of a SIE, which we used to derive 3 different CI baselines. The SIE ticket information and CI events were then used to create the assessment detailed in Section 3.

CI event types include:

- Customer calls
- Virtual assistant chats
- Device speed tests
- Gateway speed tests
- Truck rolls created
- Customer care chats
- Self-service premise health tests performed through a web application
- System refreshes
- Device reboots

We summarize the various data processing pipelines in Figure 1.

2.3. Recommendation Algorithm

The PMT Tool uses existing network usage data to create a historical picture of how customers interact with our services within their homes. This section details the two main components of our recommendations: Data Analysis and the Ranking-Based Algorithm.

2.3.1. Data Analysis

CUI data is ingested and filtered down to Account ID, Timestamp, and Total usage on a given Hour (upstream bits + downstream bits = total). Then, Data Aggregations are performed on the filtered dataset. The important fields to note are those that are shown in Figure 2.

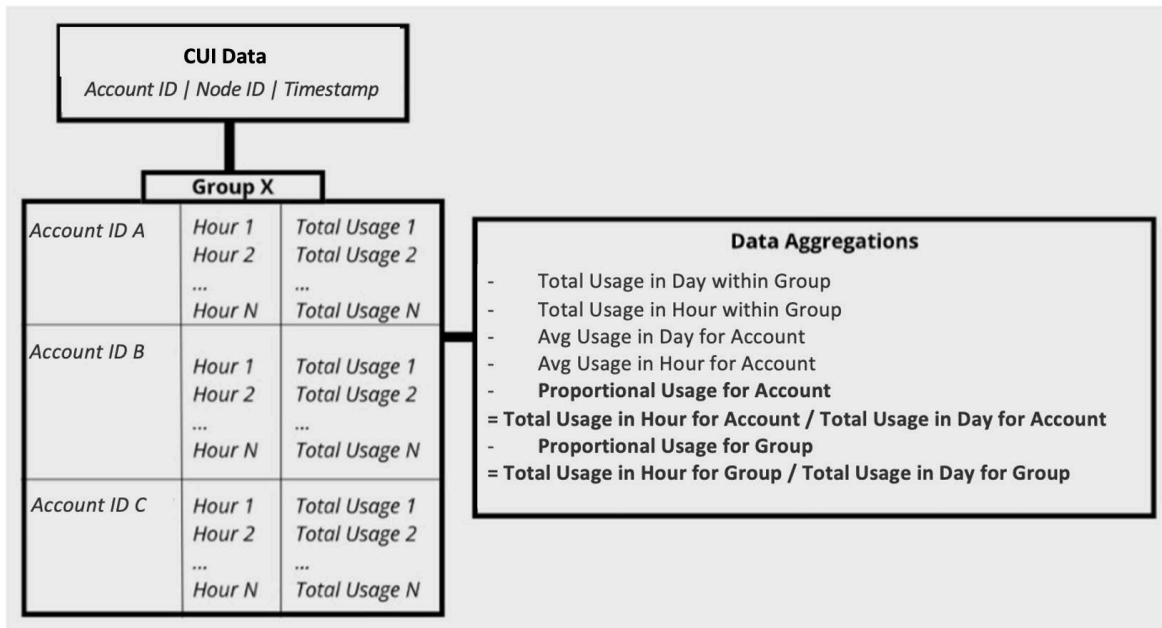


Figure 2 - How CUI data are aggregated to derive proportional usage for account and selected accounts

Proportional Usage for Account: We take the total usage seen in an hour for each Account ID and divide it by the sum of the total usage for the Account ID over the entire day (or however many hours are being considered in the comparison).

Proportional Usage for Group: We collect the total usage seen in an hour for the selected accounts (abbreviated as ‘Acct’ in the following Examples); this selected set of accounts is also referred to as a “Group”. We then divide it by the sum of the total usage for the Group over the entire day (or the specific hours considered in the comparison).

Consider Example 1 for an applied example of calculating the Proportional Usage per Account:

Acct A (Proportional Usage)	Acct B (Proportional Usage)
<i>Total Usage for Day = 99</i>	<i>Total Usage for Day = 102</i>
Hour 1 = 10 / 99 = 0.101	Hour 1 = 9 / 102 = 0.088
Hour 2 = 40 / 99 = 0.404	Hour 2 = 6 / 102 = 0.059
Hour 3 = 3 / 99 = 0.03	Hour 3 = 5 / 102 = 0.049
Hour 4 = 15 / 99 = 0.151	Hour 4 = 12 / 102 = 0.118
Hour 5 = 20 / 99 = 0.202	Hour 5 = 10 / 102 = 0.098
Hour 6 = 5 / 99 = 0.051	Hour 6 = 20 / 102 = 0.196
Hour 7 = 6 / 99 = 0.061	Hour 7 = 40 / 102 = 0.392

Example 1 - Calculating Proportional Usage per Account

With Proportional Usage (PU) per Account derived, we can move on to describe how we ranked an account's PU relative to others within the selected set of accounts ("Group" and "Acct" are used interchangeably).

2.3.2. Ranking-Based Algorithm

Scoring $N = \# \text{ Hours in Group} // 2$	
Best Hours Hours ~ N LOWEST Total Usage / Acct % of Acct for each Hour Returned	Worst Hours Hours ~ N HIGHEST Total Usage / Acct % of Acct for each Hour Returned

Figure 3 - Scoring of best and worst hours

Figure 3 outlines the theory behind making recommendations based on CUI usage data. All recommendations are made according to the number of hours to be considered before creating the recommendation. The defined hour arrangements are as follows:

Morning Hours: 6 am-5 pm (inclusive)

Extended Working Hours: 6 pm-11 pm (inclusive)

All Day: 6am-11pm (inclusive)

The number of hours being compared in each arrangement is divided by 2 (dropping any remainder) to get N .

Morning Hours: $N = 5$

Extended Working Hours: $N = 2$

All Day: $N = 8$

N is used to determine the best and the worst hours to perform maintenance for each Account ID. For example, looking at a 7-hour window of time, N would be equal to $7 // 2 = 3$. We apply this in *Example 2*.

Acct A (Total Usage)	Acct B (Total Usage)
Hour 1 = 10	Hour 1 = 9
Hour 2 = 40	Hour 2 = 6
Hour 3 = 3	Hour 3 = 5
Hour 4 = 15	Hour 4 = 12
Hour 5 = 20	Hour 5 = 10
Hour 6 = 5	Hour 6 = 20
Hour 7 = 6	Hour 7 = 40
N Best Hours Acct A = Hours 3, 6, 7	N Best Hours Acct B = Hours 3, 2, 1
N Worst Hours Acct A = Hours 2, 5, 4	N Worst Hours Acct B = Hours 7, 6, 4

Example 2 - Determining Best and Worst Hours according to CUI

Next, the results are accumulated to get the percentage of each Account ID returned for each hour within a group. *Example 3* is expanded to demonstrate this:

Group AB (Acct A and Acct B)	
Best Hours Group AB:	Worst Hours Group AB:
Hour 1 = 0.5	Hour 1 = 0.0
Hour 2 = 0.5	Hour 2 = 0.5
Hour 3 = 1.0	Hour 3 = 0.0
Hour 4 = 0.0	Hour 4 = 1.0
Hour 5 = 0.0	Hour 5 = 0.5
Hour 6 = 0.5	Hour 6 = 0.5
Hour 7 = 0.5	Hour 7 = 0.5

Example 3 - Proportional Network Usage According to Group

We now take our example's Proportional Users (PU), where $PU = \% \text{ Accts at Best Hour} - \% \text{ accounts at the worst hour}$. *Example 4* breaks this down:

Best Hours Group AB:		Worst Hours Group AB:	(PU):
Hour 1 = 0.5	-	Hour 1 = 0.0	= 0.5
Hour 2 = 0.5	-	Hour 2 = 0.5	= 0.0
Hour 3 = 1.0	-	Hour 3 = 0.0	= 1.0
Hour 4 = 0.0	-	Hour 4 = 1.0	= -1.0
Hour 5 = 0.0	-	Hour 5 = 0.5	= -0.5
Hour 6 = 0.5	-	Hour 6 = 0.5	= 0.0
Hour 7 = 0.5	-	Hour 7 = 0.5	= 0.0

Example 4 - Calculating Proportional Users

We define the Weight of our Prediction based upon the summation of the Proportional Usage for the accounts over each of the Hours in Best Hours and Worst Hours, as briefly explained in *Example 1*. We then combine that with the results from *Example 2* to get the Proportional Usage of Accts across the Best Hours and the Proportional Usage of Accts across the Worst Hours. This is written out in *Example 5*. The Best Hour results corresponding to each Account are marked with a “*”, while the Worst Hours are denoted with a “-” minus sign.

* Best Hours for Mac - Worst Hours for Mac			
Mac A	Mac B	B-Level	W-Level
*Hour 1 = 0.101	*Hour 1 = 0.088	$(0.088 + 0.101) = 0.189$	$(0) = 0$
-Hour 2 = 0.404	*Hour 2 = 0.059	$(0.059) = 0.059$	$(0.404) = 0.404$
*Hour 3 = 0.03	*Hour 3 = 0.049	$(0.03 + 0.049) = 0.079$	$(0) = 0$
-Hour 4 = 0.151	-Hour 4 = 0.118	$(0) = 0$	$(0.151 + 0.118) = 0.269$
-Hour 5 = 0.202	-Hour 5 = 0.098	$(0) = 0$	$(0.202 + 0.098) = 0.3$
*Hour 6 = 0.051	-Hour 6 = 0.196	$(0.051) = 0.051$	$(0.196) = 0.196$
*Hour 7 = 0.061	-Hour 7 = 0.392	$(0.061) = 0.061$	$(0.392) = 0.392$

Example 5 - Calculating Account Level Prediction Weights

Finally, we get the Weight of Our Prediction (WP) by taking W-Level – B-Level, shown in Example 6.

Hour	W-Level		B-Level	(WP)
1	0	-	0.189	= -0.189
2	0.404	-	0.059	= 0.345
3	0	-	0.079	= 0.079
4	0.269	-	0	= 0.269
5	0.3	-	0	= 0.3
6	0.196	-	0.051	= 0.145
7	0.392	-	0.061	= 0.331

Example 6 - Calculating Prediction Weights

Our final step is to take our PU from *Example 4* and subtract the WP to get our overall recommendation. This final step is calculated in *Example 7*.

<i>Recommendation = (PU) – (WP)</i>			
Hour	(PU)	(WP)	Recommendation
1	0.5	-0.189	= 0.689
2	0.0	0.345	= -0.345
3	1.0	0.079	= 0.921
4	-1.0	0.269	= -1.269
5	-0.5	0.3	= -0.8
6	0.0	0.145	= -0.145
7	0.0	0.331	= -0.331

Example 7 - Final Recommendation Values

The recommendation value is then divided into three separate categories: Do-Not-Recommend (DNR), Caution (CAU), and Recommend (REC). To be categorized as DNR, the recommendation value will be ≤ -0.1 . Conversely, a REC result will be ≥ 0.1 . This leaves “CAU” to be between -0.1 and 0.1.

Figure 4 summarizes the above Examples as our PMT algorithm:

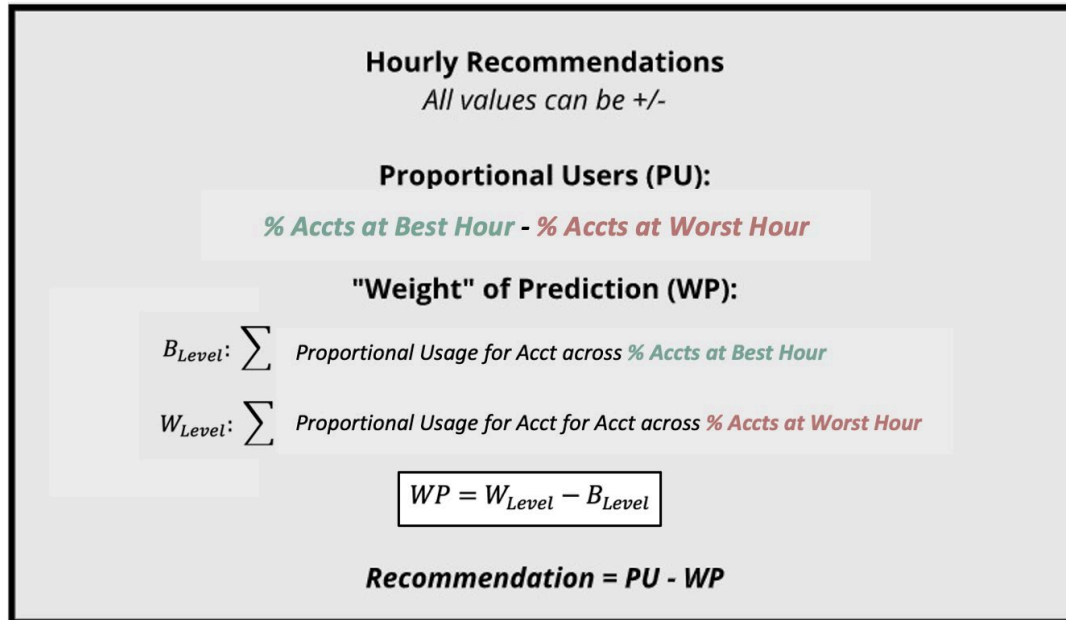


Figure 4 - PMT Algorithm Formula

3. Assessment

The PMT algorithm leverages CUI to find the optimal times when customers served by the same branch of a network have the lowest usage pattern levels relative to other hours of the day. This algorithm assumes that customers will likely experience less service impact during periods of proportionally lower data usage relative to their respective total use across the day. Identifying a collective lower usage pattern helps drive the recommendation towards the most optimal hours for a group of customers whose network branch requires planned maintenance.

3.1. Aim

Since the PMT algorithm is dependent on the CUI, we must determine if the identified CUI has sufficient signal to allow the PMT to make meaningful recommendations.

Specifically, we want to ensure that PMT suggestions of REC, CAU, or DNR for performing an SIE show variation in customer impact (approximated by the volume of CIs). We anticipate that if CUI has sufficient signal, SIEs occurring during PMT's "recommended" hours will show less relative CI increase than SIEs performed during PMT "do-not-recommend" hours. Otherwise, we would expect that the difference in these relative SIE increases across different PMT recommendations is negligible.

3.2. Method

We took advantage of the availability of historical CUI, CI, and SIE data to perform the assessment. This approach mimics an "idealized" scenario because we already know i) the SIE occurred, ii) the exact set of households impacted by the specific SIE, and iii) have access to the household's corresponding CUI for the same day as the SIE, which is used to derive the PMT recommendation for the hour each SIE occurred. Given that we were planning to trial an application based on the PMT algorithm in one of the

regions of our service footprint (aka the ‘test region’), we focused our assessment on the same region of interest from February 01 to March 11, and April 14, 2022. This time period avoids date ranges that could potentially be affected by daylight savings in 2022 (13th March (USA) and 27th March (UK)) as well as 2021 winter holiday seasonal effects³.

A total of 3341 SIEs in the test region occurred during our collected data sample's assessment period of interest. We derived PMT recommendations for each SIE and its set of corresponding households' CUIs for each hour from 6 am to 11 pm ET on the day each SIE took place. The hourly recommendations for the same day as the SIE allowed us to determine if the start of each SIE occurred at an hour the PMT yielded i) REC, ii) CAU, or iii) DNR result.

Categorizing the time of SIEs onset by PMT recommendations allowed us to compare the average volume of CI associated with each set of affected households during the start hour of SIE relative to their corresponding mean CI baseline across these PMT categories of SIEs. We refer to this metric of relative change in the mean volume of customer interactions as our ‘delta ratio’ ($dRatio_{CI}$):

$$the\ dRatio_{CI} = \frac{CI_{SIE_{startHr}}}{CI_{baseline}}$$

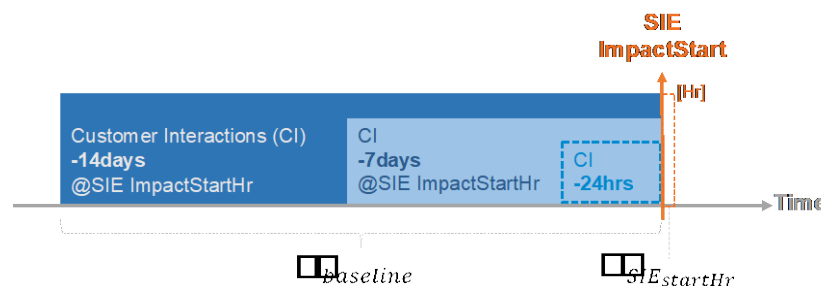


Figure 5 - Three different CI baselines were considered in our assessments.

For each SIE, the start of SIE impact (denoted as SIE ImpactStart) is used as a reference to derive the average CIs from the prior 24 hours, the average CIs from the same first hour since the start of SIE impact (denoted as a dotted orange bar) in the prior 7 or 14 days.

We derived the *delta ratio* metrics using three different CI baselines (see Fig 5): 24hrs prior, seven days prior, and 14 days before the same SIEs. It is helpful to get a sense of the fluctuations in CI in the 24hrs before the SIE, as it can show the potential time-of-day effects of customer interactions (e.g., customers may tend to interact at certain hours of the day). This comparison also prompted us to consider the volume of CI in the last 7 or 14 days during the same period as the hour following the start of the service impact event (@SIE). At the same time, this does not fully account for potential day-of-the-week effects (for which we needed a sample of historical CI data going back for more weeks than we were able to

³ There is a +5hrs (before 13th March) +4hrs (13-26th March) +5hrs (27th March onwards) UTC – EST/EDT difference. While there is potentially no differences in terms of Service Interruption Event (SIE) impactHr timestamps, the associated Customer Interactions (CI) events would include some days with +4/+5 ET depending on number of days or hours (7 or 14days | 24hrs baselines) one looks back relative to the CI data time zone. As such, the date-range we work with is primarily to avoid dealing with daylight savings conversions.

sample at the time of writing). We at least established a prior seven or 14-day baseline for the same period as the hour following each SIE.

3.3. Findings

3.3.1. Comparison of REC and DNR delta_ratios

We are particularly interested in the difference between DNR and REC delta_ratios. Specifically, if a DNR delta_ratio is more extensive compared to REC delta_ratio, we can infer that the SIE at a PMT REC hour is a less customer “impacting” time.

Figure 6 illustrates how delta_ratio(s) are derived for the three PMT categories of SIEs for the assessment using 24hrs before baseline with all customer interaction types considered. The analytical approach is similarly applied for the evaluations performed using the same SIE for each day in the previous 7 or 14 days as baselines (Figs 9—11).

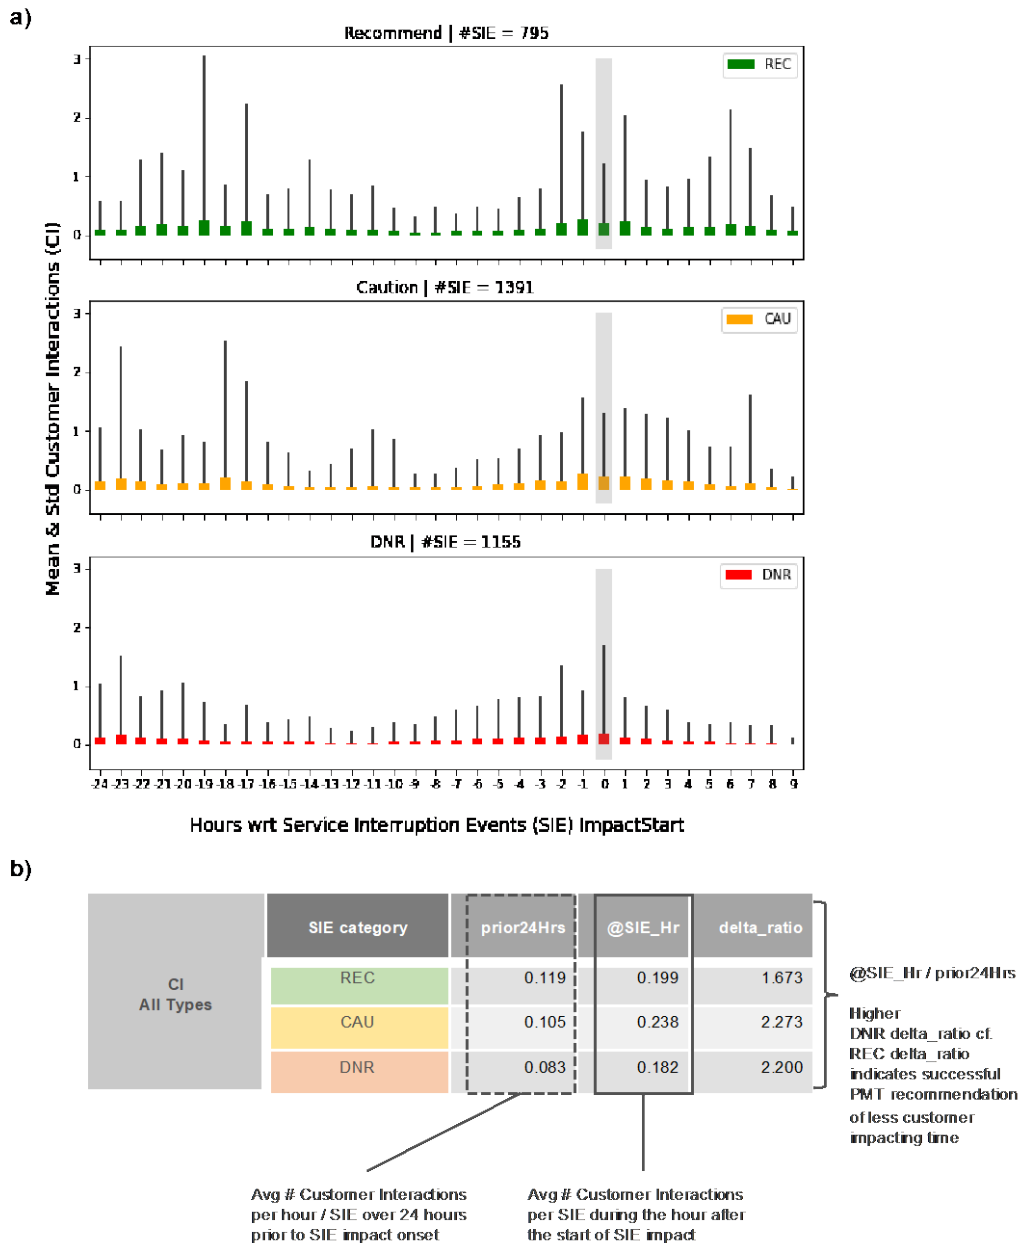


Figure 6 - CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with prior 24hrs as a baseline

a) Hourly mean and standard deviation of CI across all SIEs concerning impact start hour (i.e., 1—24 hours prior and 1—9 hours after). All types of CIs are considered. The prior 24 hrs as a baseline is denoted with SIE categorical shading, relative to SIE start hour, as indicated by the grey vertical bar; b) Summary of average CI i) across all hourly means of prior 24 hours (prior24Hrs), ii) for an hour after SIE impact start (@SIE_Hr), as well as the respective delta_ratio for each SIE category.

Figures 7 and 8 provide the summary of customer interactions 24 hours before the SIE impact start hour for assessments with sub-types of customer interactions.

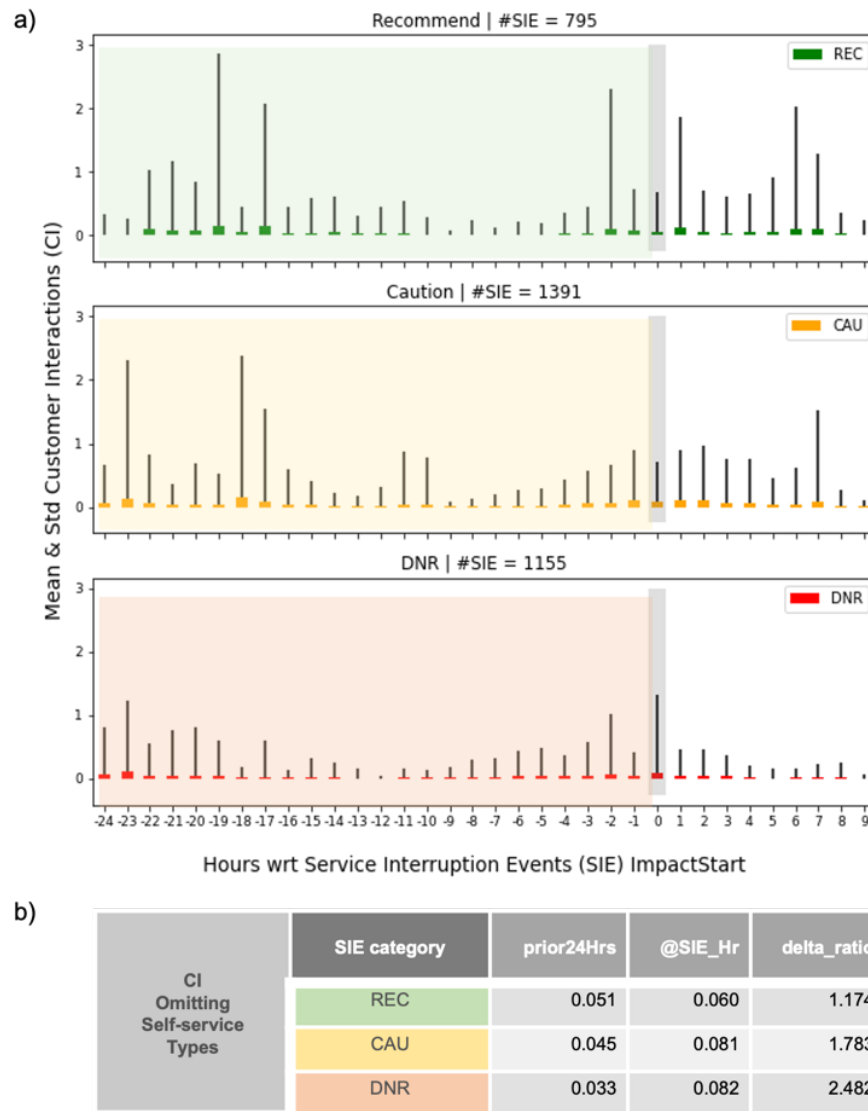


Figure 7 - Non-self-service CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with prior 24hrs as a baseline.

a) Hourly mean and standard deviation of CIs across all SIEs concerning SIE impact start hour (i.e., 1—24 hours prior and 1—9 hours after). CIs without self-service event types are considered. The prior 24hrs as a baseline is denoted with SIE categorical shading, relative to SIE start hour, as indicated by the grey vertical bar; b) Summary of average CI i) across all hourly means of prior 24 hours (prior24Hrs), ii) for an hour after SIE impact start (@SIE_Hr), as well as the respective delta_ratio for each SIE category.

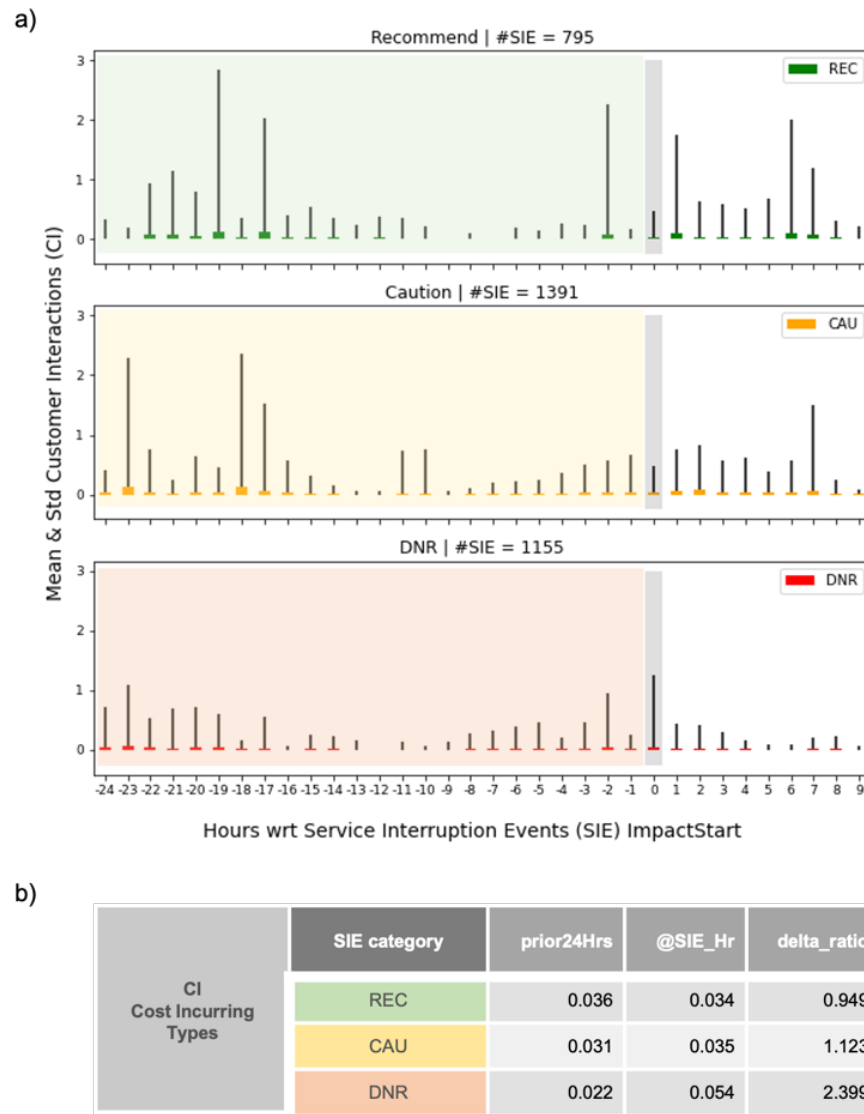


Figure 8 - Cost-incurring CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with prior 24hrs as a baseline

a) Hourly mean and standard deviation of CI across all SIEs concerning SIE impact start hour (i.e., 1—24 hours prior and 1—9 hours after). Cost-incurring CI types (e.g., technician visit scheduling; repair call/chats) considered. The prior 24hrs as a baseline is denoted with SIE categorical shading, relative to SIE start hour, as indicated by the grey vertical bar; b) Summary of average CI i) across all hourly means of prior 24 hours (prior24Hrs), ii) for an hour after SIE impact start (@SIE_Hr), as well as the respective delta_ratio for each SIE category.

Similarly, the baselines of assessments performed using the same SIE in the prior seven days and 14 days are summarized in Figures 9—11.

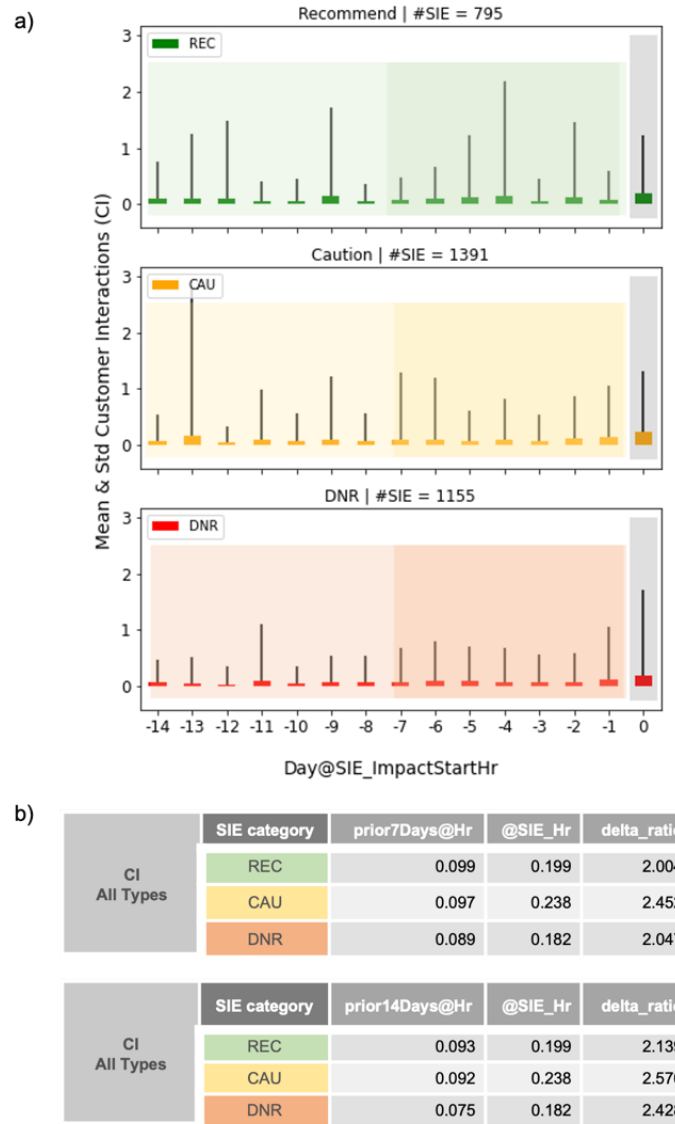


Figure 9 - CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with SIE_ImpactStartHr in prior 7 or 14 days as a baseline

a) Daily mean and standard deviation of CI (at SIE impact start hour across all SIEs. All types of CIs are considered. The prior 7 or 14 days CI at SIE impact start hour as a baseline is denoted with SIE categorical shading, relative to SIE start hour, as indicated by the grey vertical bar; b) Summary of average CI i) across all hourly means of prior 7 or 14 days, ii) for an hour after SIE impact start (@SIE_Hr), as well as the respective delta_ratio for each SIE category.

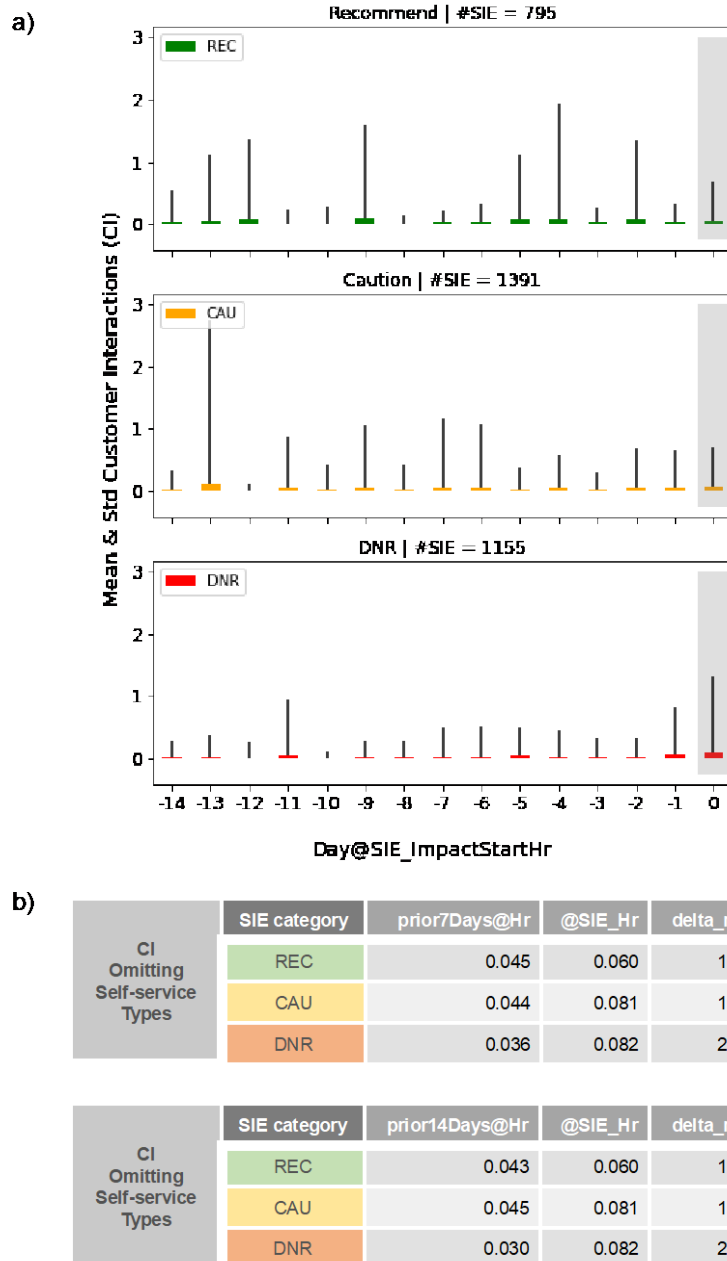


Figure 10 - Non-self-service CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with SIE_ImpactStartHr in prior 7 or 14 days as a baseline.

a) Daily mean and standard deviation of CI at SIE impact start hour across all SIEs. CIs without self-service event types are considered. The prior 7 or 14 days CI at SIE impact start hour as a baseline is denoted with SIE categorical shading, relative to SIE start hour, as indicated by the grey vertical bar; b) Summary of average CI i) across all hourly means of prior 7 or 14 days, ii) for an hour after SIE impact start (@SIE_Hr), as well as the respective delta_ratio for each SIE category.




Figure 11 - Cost-incurring CIs relative to SIE_ImpactStartHr and derived delta_ratios for the PMT categories of SIEs with SIE_ImpactStartHr in prior 7 or 14 days as a baseline.

a) Daily mean and standard deviation of CI at SIE impact start hour across all SIEs. Cost-incurring CI types (e.g., technician visit scheduling; repair call/chats) considered. The prior 7 or 14 days CI at SIE impact start hour as a baseline is denoted with SIE categorical shading, relative to SIE start hour, as indicated by the grey vertical bar; b) Summary of average CI i) across all hourly means of prior 7 or 14 days, ii) for an hour after SIE impact start (@SIE_Hr), as well as the respective delta_ratio for each SIE category

We observed that our assessments across the different baselines (24hrs prior; 7 days prior; 14 days prior) and the various combinations of customer interaction types (e.g., all types; excluding self-service trouble-

shooting; inclusion of only cost-incurring types) all show a consistent trend: Service Interruption Events (SIEs) performed during PMT “recommended” hours show less relative customer interaction (CI) delta_ratio increase compared to SIEs performed during PMT “do-not-recommend” hours (see Tables 1—3).

Table 1 - Derivation of delta_ratios for comparison across SIE categories using 24hrs prior CIs as a baseline.

a)	CI All Types	SIE category	prior24Hrs	@SIE_Hr	delta_ratio
		REC	0.119	0.199	1.673
		CAU	0.105	0.238	2.273
		DNR	0.083	0.182	2.200
b)	CI Omitting Self-service Types	SIE category	prior24Hrs	@SIE_Hr	delta_ratio
		REC	0.051	0.060	1.174
		CAU	0.045	0.081	1.783
		DNR	0.033	0.082	2.482
c)	CI Cost Incurring Types	SIE category	prior24Hrs	@SIE_Hr	delta_ratio
		REC	0.036	0.034	0.949
		CAU	0.031	0.035	1.123
		DNR	0.022	0.054	2.399

Table 2 - Derivation of delta_ratios for comparison across SIE categories using CIs in the prior 7 days at the same impact start hour as a baseline

a)	CI All Types	SIE category	prior7Days@Hr	@SIE_Hr	delta_ratio
		REC	0.099	0.199	2.004
		CAU	0.097	0.238	2.452
		DNR	0.089	0.182	2.047
b)	CI Omitting Self-service Types	SIE category	prior7Days@Hr	@SIE_Hr	delta_ratio
		REC	0.045	0.060	1.333
		CAU	0.044	0.081	1.828
		DNR	0.036	0.082	2.293
c)	CI Cost Incurring Types	SIE category	prior7Days@Hr	@SIE_Hr	delta_ratio
		REC	0.036	0.034	0.955
		CAU	0.031	0.035	1.140
		DNR	0.024	0.054	2.249

Table 3 - Derivation of delta_ratios for comparison across SIE categories using CIs in the prior 14 days at the same impact start hour as a baseline

a)	CI All Types	SIE category	prior14Days@Hr	@SIE_Hr	delta_ratio
		REC	0.093	0.199	2.139
		CAU	0.092	0.238	2.576
		DNR	0.075	0.182	2.428
b)	CI Omitting Self-service Types	SIE category	prior14Days@Hr	@SIE_Hr	delta_ratio
		REC	0.043	0.060	1.415
		CAU	0.045	0.081	1.798
		DNR	0.030	0.082	2.777
c)	CI Cost Incurring Types	SIE category	prior14Days@Hr	@SIE_Hr	delta_ratio
		REC	0.034	0.034	1.005
		CAU	0.035	0.035	1.012
		DNR	0.020	0.054	2.704

What is particularly interesting to note is that when omitting self-service CI types or considering only cost-incurring CIs – such as scheduling a technician visit and calls and chats with a customer agent – we observed relatively larger differences in DNR—REC delta_ratios. Although the differences in delta-ratios observed were not statistically significant ($p > 0.05$ ⁴; likely due to data sample sizes), these findings are encouraging because they point to specific CIs that we could potentially monitor and assess for the financial and customer impact on an ongoing periodic basis.

3.3.2. Time-of-day Effects

In addition to deriving and assessing the DNR—REC delta_ratios, categorizing the time of SIEs onset by PMT recommendations also allowed us a view into when different categories of PMT recommended SIEs tended to occur over the period of 06:00 hrs—23:00 hrs in the test region. Figure 12(a) shows the distributions of SIEs for each PMT recommendation category over time.

Additionally, our assessments highlighted a time-of-day effect: hours that the algorithm would recommend tended to occur earlier in the day (06:00-14:00hrs), while the caution hours shifted to later (08:00-15:00hrs), and for hours that PMT yielded a do-not-recommend, we saw the latest (09:00-19:00hrs).

Combining the separate categorical plots in Fig 12(b) as a relative proportion of all SIEs performed during the onset hour of service interruption across the period of 06:00 hrs—23:00 hrs, we can begin to appreciate the time-of-day effects together with the previously described trend in delta-ratios. Specifically, we observe that performing an SIE after 15:00 hr almost always is associated with high customer impact compared to performing an SIE before 08:00 hr when it shows low impact for customers.

⁴ Our t-tests were performed with probability (p) significance test against the threshold $\alpha = .05$; we assume and allow for a 5% chance level of how extreme our observed results must be to reject the null hypothesis of no delta_ratio difference. At the set threshold of $p < .05$, an observed test probability p below 0.05, would indicate the alternative hypothesis of a delta_ratio difference is ‘statistically significant’ and that we could reject the null-hypothesis. In our case we observe that the t-test performed on the delta_ratios derived from our limited sample exceeds the set threshold of acceptable chance level, and we conclude that the observed trend is not statistically significant.




Figure 12 - Distribution and relative proportion of SIEs for each PMT category.

- a) Hourly distribution of SIEs for each PMT recommendation category – REC; CAU; DNR – over time;
b) Relative proportion of all SIEs in data sample during onset hour of service interruption across the period of EST 06:00 hrs—23:00 hrs.

3.4. Assessment Summary

Overall, the assessments for the test region over the date range highlighted the following:

- 1) CUI has sufficient signal to drive the PMT algorithm;
- 2) Different categories of PMT recommendation are associated with a consistent trend: SIEs performed during PMT “recommended” hours show less relative customer interaction (CI) increase (as measured by delta_ratio) compared to SIEs served during PMT “do-not-recommend” hours;
- 3) Operationally, a simple and effective way to improve customer experience is by retroactively measuring CI for SIEs in defined geographic regions and using the results to provide guidance. PMT recommendation is most effective if it is run:
 - i. Only when a planned maintenance job needs to be scheduled (i.e., beforehand, before being in location)
 - ii. Derived with CUI of the exact SIE-affected accounts

We reiterate that our data-driven assessment approach mimics an idealized scenario, which may not be easily achieved in the field without a robust data ingest and computational platform. To approximate the idealized system, future customer usage information must be forecasted. This relies on the stability of historical data, which may be affected by seasonal events – an area of research we hope to pursue.

Notwithstanding, the findings discussed provide valuable insights and recommendations for how the PMT could be applied in the field. Importantly, given that the data sampled was over a short period in the current year and for only the selected test region, periodic assessments would be needed to monitor if our observed delta_ratios and time-of-day trends hold across seasons, time, and indeed, if there may be regional differences in such movements.

Next, we discuss practical considerations of applications in the field and how our findings could be best incorporated into deploying an early version of the PMT.

4. Application

An application was developed, for use by technicians, to determine the best time to perform SIEs, specifically for ‘planned maintenance’ events to be performed during regular and extended working hours. For better or worse, the application was developed with scalability and performance, which means that its function does not match that of the Assessment described above in Section 3. Although much further testing would be required to measure its effectiveness, it helped expose the challenges faced by developing a field tool that would rely on data generated by various corporate systems. A description of the application follows.

4.1 Architecture

4.1.1. Overview

The application is divided into three major systems, including the same Recommendation Service (Section 2.3) developed for the Assessment (Section 3). The additional components include a Grouping Service (Section 4.1.3), designed to help with performance, and an Application Programming Interface (API), where aggregated data and field tech requests were processed. Figure 13 provides an overview of the PMT Application and its different components.

As described earlier, the Recommendation Service uses CUI metrics and groups of selected accounts to calculate recommendations. However, in the PMT application, we created a ‘Grouping Service’ that pre-determined groups of up to 40 accounts for which recommendations would be calculated. This approach allows the recommendations to be calculated before the technicians need them, thus minimizing query latency, unlike systems that must perform data pulls and calculations on-demand. We provide a summary of this in Figure 14 and Section 4.1.2.

The Grouping Service (Section 4.1.3) consumes data about the structure of the cable plant (see description of PTI in Section 2.1.3) from each CMTS down to each account. It then creates groups of up to 40 accounts on the same network branch. This is then stored in the MySQL database via the API, which is used to fetch this data by the Recommendation Service. These pre-created groups are approximations of areas that could be affected by a representative SIE event based on their connectivity to the cable plant.

The application is built with a RESTful web service that takes requests from the Field Tech’s Laptop through a user interface, which helps them generate a list of accounts that would be involved in an SIE. This list of accounts is then sent to the API, which determines the group it best represents, looks up the pre-calculated recommendations for that group, and returns the result to the user interface on the Field Tech’s laptop, as illustrated in Figure 13.

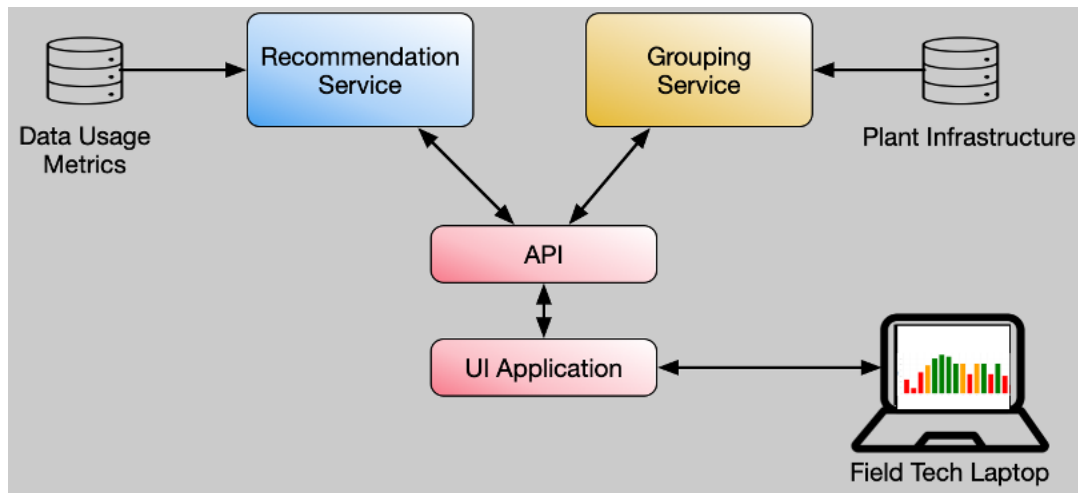


Figure 13 - Overview of PMT Application and its different components.

Detailed descriptions of each of these components are now discussed, except for the Recommendation Algorithm (Section 2.3), which has already been described.

4.1.2. Process Flow

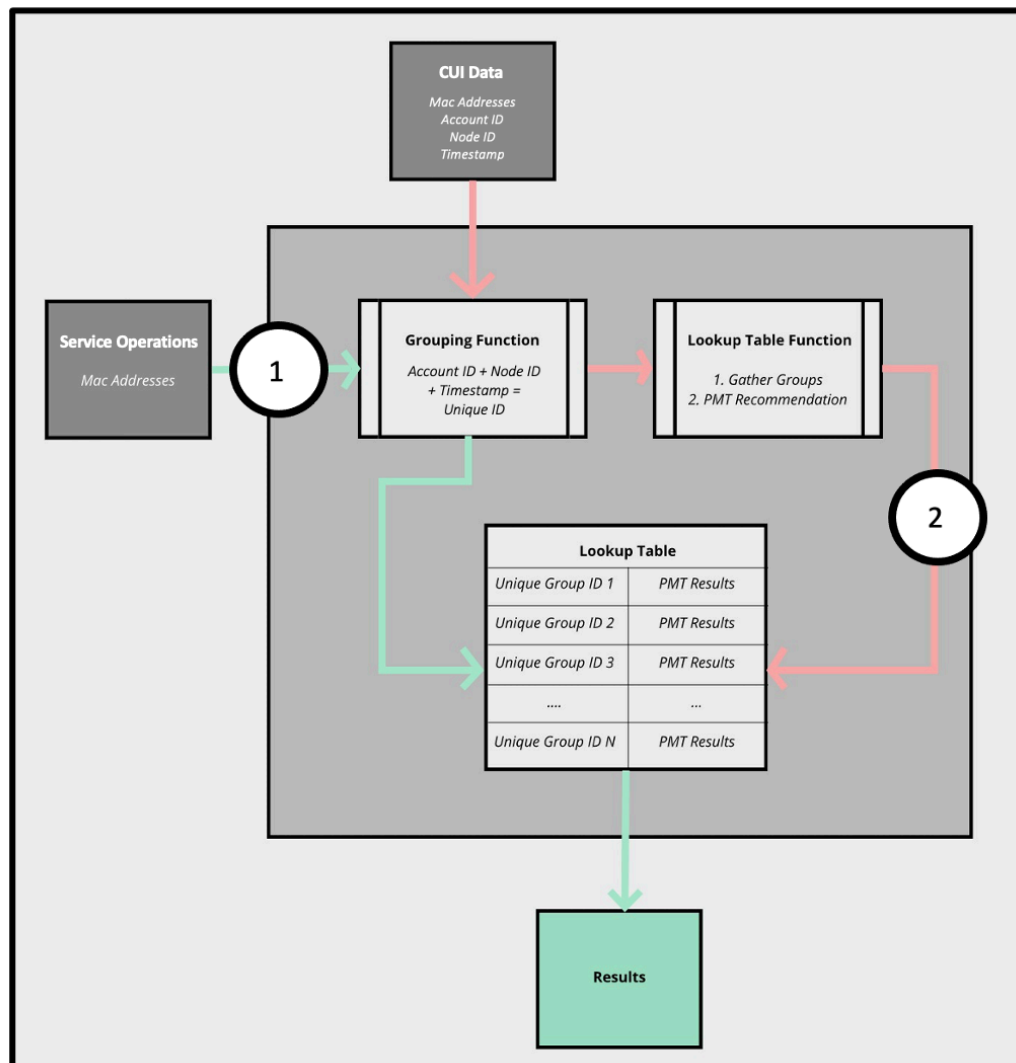


Figure 14 - High-Level Overview of PMT Recommendation Service

Flow 1, described by the **green** arrows, represents a technician's process when scheduling a job. It takes Mac Addressees provided by our Service Operations Tool and then matches them to Account and RF Node IDs. This data then gets passed through the Grouping function to create a Unique Identifier for the group. That Unique ID goes to the Lookup Table, and the PMT Result tied to the Unique ID is returned.

Flow 2, described by the **red** arrows, represents the backend flow where data is processed and prepared for a technician to query. It takes in CUI Data, comprised of Upstream and Downstream rates, and creates a unique ID using the same Grouping function as Flow 1. Then, this ID and CUI Data are passed to the Lookup Table Function, making a PMT Recommendation. The results are then populated into the Lookup Table to await a technician to access the results.

4.1.3. Grouping Service

An assumption was made that given a neighborhood where our services are provided, customers within closer proximity and/or customers who share a pedestal or an amplifier may experience similar plant activity. For example, fiber connects the CMTS to the RF Node in a neighborhood; there would be actives and amplifiers downstream. From there on, there would be taps with up to 8 ports that supply our services for up to 8 households in a residential neighborhood. If the amplifier is faulty, every household connected downstream from that amplifier would be affected. This led to creating groupings of 5-40 accounts.

An internal geographical topology tool, which is built on top of the Geographic Information System Framework (GIS), provides geographical and physical connection information regarding the plant's infrastructure: CMTS, Nodes, Actives, Passives, Amplifiers, Taps, and houses that have been set up to receive our services. Once information regarding a particular CMTS is retrieved, PMT's Grouping Service creates the groups based on common ancestors (such as taps or amplifiers) and geographical proximity within the infrastructure tree. It saves it into a graph database using a graph framework (See Figure 15 below). The last step in the Grouping Service process is to send lists of accounts for all the identified groups and underlying information, including account numbers, MAC addresses, and physical addresses, over to the MySQL database via the API. This information is then pulled into the Recommendation Service via an automated pull.




Figure 15 - Example groupings determined by Grouping Service.

Two groups being calculated based on their geographical proximity by the Grouping Service

There are tradeoffs in assuming that the data usage patterns are similar between customers sharing a common ancestor or geographical proximity. It would be more accurate for the PMT application to evaluate each account's statistically considered data consumption when calculating a recommendation for a SIE.

4.1.4. User Interface

Due to expediency and the fact that technicians use a smartphone or their in-truck laptop, we decided to use a browser-based solution. In the PMT application, the UI was created with ReactJS, and the backend was developed via NodeJS, Express API, and MySQL as a database. This was done to provide a visualization of the recommendations to a technician that would be concise and relatively quick to read.

As mentioned in section 4.1.3, the input for the area affected by the SIE is acquired via an internal GIS-based tool in which the accounts within the impacted area are identified and provided to the PMT application. Within the UI, the technician selects future dates (as far as two days in the future) in which they desire to view the recommendations for that area. These inputs go to the API, which determines which accounts belong to which groups via a “majority rules” filter to select which pre-determined group has the most accounts for an assigned area. Using a graph, PMT then fetches the recommendations for the identified group from the MySQL database and displays them on the ReactJS frontend.




Figure 16 - Example PMT application recommendations

After the accounts are input by the user, PMT would run its microservices and reveal the hourly recommendation based on the hours of the day and the date selected.

Figure 16 shows the displayed recommendations in three color categories: green for REC, yellow for CAU, and red for DNR. In addition, the bar graph’s magnitude is determined by the values calculated by the Recommendation Service. Additionally, a user can look at today’s, tomorrow’s, and the day after tomorrow’s data to find an ideal time to perform an SIE.

5. Discussion

5.1 Learnings from the field and insights from the assessment

We discuss the learnings from our small-scale field trial and the insights derived from assessing our algorithm with historical SIE and CI data.

5.1.1. Trial

The PMT application was trialed in a limited area (0.15% of total nodes in the region), or approximately 30 nodes in a test region with tens of thousands of nodes. While the trial is ongoing, we have already discovered that deployment of the PMT on a fraction of the regional network plant does not generate sufficient data to concretely determine if the use of the developed PMT application improves customer experience.

We note that the assessment described in Section 3 was performed on a sample of a few thousand SIEs across the whole test region. As just mentioned, the application trial covered a small fraction of the test region by comparison. Therefore, to match the integrity of the assessment, the trial would need to be expanded to include the entire region. However, such an expansion would be equivalent to creating a large-scale production-level application deployment, which would require committing resources to an unverified design.

As of this writing, we determined it would be best to extend our assessment approach within the same region, as well as to the other areas, to verify that findings in the test region: scale up in volume of historical SIEs and CIs, and check if they are consistent across the geography of different regions and over time.

5.1.2. Learnings and Observations

We acknowledge that our assessment of derived PMT recommendations, based on historical data, simulated an *idealized* scenario wherein customer usage information was extracted for the set of SIE-affected accounts and all hours on the same day (that is still unfolding) when the SIE occurred. In a real-world scenario, this application would have to reliably and efficiently estimate future customer data used to calculate recommendations for any set of accounts and any day. Building such an application will require:

- Reliable forecasting of future customer data usage from historical CUI;
- Adequate computing and data ingest resources to continuously store, fetch and compute the stream of available CUI data;
- Adequate computing resources to calculate and return recommendations, at scale, in a performant fashion when a field technician makes a request;
- A system to continuously monitor system effectiveness to provide feedback for improvements.

A noteworthy finding is the time-of-day pattern of SIEs and their associated PMT recommendation category observed in our assessment: SIEs with PMT REC hours tended to occur earlier in the day. In contrast, an SIE during hours that PMT would yield CAU or DNR occurred at higher frequencies later in the afternoon. Field operations groups can immediately use this time-of-day insight to guide them to make sensible choices and possibly even schedule directives on when to perform investigations and preventative maintenance that results in a SIE.

5.1.3. Considerations

In the future, it may be prudent to store data about SIE's impact on customers for an entire year to understand annual patterns. Correlations could be based on locality, especially given the different schedules that universities, K-12 schools, and local holidays, among other events, will affect data usage. For example, university towns have transient populations that drop significantly between terms. K-12 schools have schedules imposed by local governments and are typically published before each school year. Knowing when children are not going to be at school means knowing that they're probably going to be at home sharing bandwidth with parents who may be working from home.

Moreover, events like the COVID pandemic have significantly impacted the use of residential gateways, making knowledge of data usage patterns even more critical as livelihoods have been, and continue to be, made from home offices. Although epidemics and pandemics do not have a known schedule, local health directives and infection rates could be studied and used for predicting increases in residential data usage.

6. Conclusion

Insights from our assessments can be used to recommend how a PMT user application could be architected, deployed, and used in real-time to guide technicians to the optimal time to schedule an SIE. While it should be cautioned that the data used in our assessments were limited in sample size, within one northeastern region in our national footprint, and over a limited time range (February 01 – March 11 and April 16th onwards), the findings we observed suggest that customer usage information could be a good indicator of potential customer impact during SIEs.

Observing trends using the assessment method described in section 3 can help simplify the deployment of the PMT application across our service footprint. It offers a way to standardize its use and potentially roll out randomized field trials across multiple regions for a more robust assessment of the value of using the PMT application in the field.

We also know that the recommendation algorithm may not always yield the expected trend during the year. This could be due to many factors, including:

- seasonal customer data usage patterns
- customer experience campaigns rolled out in parallel
- changes in data usage patterns due to pandemics or natural disasters

Overall, we learned that it helps to employ “hindsight,” using retrospective analyses as we did in assessing our algorithm with historical SIE and CI data. We suggest this as good practice before building a field application for trial. The insights derived from a data-driven analytical approach can help guide how such an application would be best used and deployed. Since there is no way to predict which nodes in our network will require the most SIEs in the future, we determined retroactive data was the most efficient way, albeit idealized, to validate our assumption that customer usage data at particular times during the day would be indicative of customer impact.

In conclusion, we entered this trial with the premise that the best time to interrupt the network with a SIE is when nobody is using it. This premise was assessed using the methodologies described in this paper. The measurable reduction of CIs during SIEs that occurred during times recommended by the PMT indicates that a data-driven model to predict the best times to conduct invasive maintenance is possible and deserves further development.

Abbreviations

API	Application Programming Interface
CAU	Caution
CI	Customer interactions
CMTS	Cable Modem Termination System
CUI	Customer Usage information
DNM	Device to Network Mapping
DNR	Do Not Recommend
DOCSIS	Data Over Cable Service Interface Specification
ERSI	Environmental Systems Research Institute
GIS	Geographic Information System Framework
PMT	Planned Maintenance Tool
PTI	Plant Topology Information
PU	Proportional Users
REC	Recommend
SIE	Service interruption events
WP	Weight of Our Prediction

Acknowledgments

Many thanks for the support and contributions made to this project by our esteemed Comcast colleagues:

James Cortese, who joined billing data to plant topology using a combination of character matching and geolocation methods used in the Grouping Service.

Mai Lam, who provided project management support by driving project schedules, authoring status reports, and coordinating meetings for a project and team that changed course many times during its journey.

Susan Owens, for editing services on her first project with our team.

Peter Lester, who helped to refine the PMT assessments and asked critical questions.

PMA Improvements – Strategies Employed for Faster Mitigation, Increased Capacity, and Cost Savings

A Technical Paper prepared for SCTE by

Jonathan Leech

Principal Engineer II
Comcast

183 Inverness Dr West, Englewood, CO 80112
(720) 289-1561
jonathan_leech@cable.comcast.com

Andy Martushev

Principal Engineer II
Comcast

183 Inverness Dr West, Englewood, CO 80112
720-284-2277
andy_martushev@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction	4
2. US D3.0 Profile Management	4
2.1. Data Granularity	4
2.2. Collection and Recommendation Interval	4
2.3. Recommendations	5
2.4. Downgrade Profile	5
2.5. Upgrade Profile	5
2.6. Capacity Gain	5
2.7. Dashboards	5
3. DS D3.1 Profile Management	7
3.1. Model	7
3.2. Clustering Algorithm	8
3.3. Performance	8
3.4. Capacity Gain	9
3.5. Dashboards	9
4. US D3.1 Profile Management	12
5. Future Work	14
6. Conclusion	15
Abbreviations	15
Bibliography & References	15

List of Figures

Title	Page Number
Figure 1 – Overall US D3.0 PMA Stats	5
Figure 2 - Example US D3.0 PMA Stats for Mountain West Region	6
Figure 3 - Example US D3.0 PMA Monitoring Dashboard	6
Figure 4 – DS D3.1 PMA Profile Stats	9
Figure 5 – Example DS D3.1 OFDM Utilization Dashboard	10
Figure 6 – Example DS D3.1 AWS Dashboard	10
Figure 7 – Example profile 1 modem MER histograms, percentiles, and MER threshold	11
Figure 8 – Example profile 2 modem MER histograms, percentiles, and MER threshold	11
Figure 9 – US D3.1 PMA IUC Stats	12
Figure 10 – Example IUC 5 modem MER percentiles, histograms, and threshold	13
Figure 11 – Example IUC 6 modem MER percentiles, histograms, and threshold	13
Figure 12 - Example IUC 9 modem MER percentiles, histograms, and threshold	14

List of Tables

Title	Page Number
Table 1 – OFDM Profile MER Thresholds, 5 dB more aggressive than DOCSIS 3.1 spec	8

Table 2 – OFDMA Profile MER Thresholds from DOCSIS3.1 spec 12

1. Introduction

A Profile Management Application (PMA) is a critical component of DOCSIS downstream and upstream for both speed and reliability. This is especially true with increased bandwidth demands in recent years. As such, it is critical to react quickly to issues in order to provide the best customer experience. Faster mitigation of network issues reduces customer impact and call volumes. The profile recommendation interval was lowered from 6 hours in the previous system (Harb, 2020) to 5 minutes in the DOCSIS 3.0 (D3.0) upstream (US), and from 3.5 days to 1 day in the DOCSIS 3.1 (D3.1) US and downstream (DS), while reducing operational costs and improving capacity.

Ingesting and analyzing large amounts of data at a high rate creates high demand for both storage and CPU. The technology stack was refactored and costs were lowered by eliminating redundancy and leveraging streaming, batching, cloud computing, and parallel processing. Aligning the polling and PMA processing using Simple Storage Service (S3), Simple Notification Service (SNS), and Simple Queueing Service (SQS) allows for processing of a single batch of related data immediately after polling. Storage demand was reduced by moving components from a relational database to S3 with a large batch size. CPU demand was reduced by moving the analysis logic from a large Apache Spark cluster to a smaller Elastic Kubernetes Service (EKS) cluster. CPU demand was further reduced by refactoring the clustering algorithm to use Single Instruction Multiple Data (SIMD) parallel processing.

Making PMA recommendations more often improves capacity to an extent, but larger capacity gains were made by making changes to the profile selection and clustering algorithms. For D3.1, the Modulation Error Ratio (MER) data model was improved by using histograms and a time decay function. Better utilization and capacity estimates were created by using the model and the added time dimension. Optimal percentiles and corresponding weights are generated for each modem and used as inputs to the clustering algorithm. The result was capacity gains of greater than 8% and 5 Tb/s.

2. US D3.0 Profile Management

Capacity gains were increased primarily by decreasing the recommendation interval from 6 hours to 5 minutes. To do this, the data granularity was also changed, and the operating model changed from batch to streaming. Additionally, some changes were made to the profile recommendation business logic. Finally, costs were reduced by eliminating redundancy in the polling and data ingest pipeline.

2.1. Data Granularity

Per-modem SNR and Forward Error Correction (FEC) statistics were replaced with interface-level aggregates in the data ingest pipeline. Interface-level data is orders of magnitude smaller than the equivalent per modem data, which allows for more frequent collection and the incorporation of more historical data for making profile recommendations. Data ingested includes interface level FEC, MER, Profile, and Utilization metrics.

2.2. Collection and Recommendation Interval

Interface-level aggregate SNR and FEC is ingested in 5-minute intervals. The incoming data is retrieved from AWS S3 storage and an SNS notification informs of the availability of new data. The data is written to a time series database. There is minimal lag between the poll and the ingest, and a new set of profile recommendations is made immediately. This allows the system to respond quickly to a change in the plant conditions and results in a better customer experience. As a result, a 60% reduction in major alarms and a 50% reduction in minor alarms was achieved.

2.3. Recommendations

US profiles for D3.0 channels are comprised primarily of a modulation order from Quadrature Phase Shift Keying (QPSK) to 64-Quadrature Amplitude Modulation (QAM), and data and parity lengths for short and long grants. A recommendation is then made and implemented for each upstream channel on a Cable Modem Termination System (CMTS). Recommendations are made by choosing the most appropriate profile for the channel from a fixed set of available profiles. The chosen profile minimizes FEC error rates while maximizing throughput. Additional consideration is given to interfaces with high utilization, e.g., a profile change won't be made if it would increase the utilization above 80%.

To reduce profile flapping (repeated profile changes downgrading and upgrading the same channel), the algorithm will lower the profile more readily than it will raise it. It also takes into consideration past profiles and their corresponding MER values.

2.4. Downgrade Profile

The decision to downgrade a profile is made when the channel uncorrectable FEC rate is impaired (> 1%) for 20 minutes.

2.5. Upgrade Profile

The profile is upgraded after 30 minutes of clean uncorrectable and correctable FEC rates (< 0.0001%) and MER > 36 dB, or 75 minutes of Uncorrectable FEC < 0.1%, and a +2 decibel (dB) improvement in MER and a 50% reduction in correctable FEC.

2.6. Capacity Gain

On average, the capacity gain is measured at 16.6 %. This is about a 1.5% capacity improvement vs. the previous version of US D3.0 PMA. Gains in various markets vary from 14% to 22%. US D3.0 PMA adds about 4.2 Tb/s in total capacity to the upstream, which represents an average gain of 3.1 Mb/s per channel. System performance and capacity gains are continuously monitored via dashboards and automated alerting.

Overall Profile Stats				
cmts_count	gainpct	interface_count	rawgain	avgintgain
3756	16.6	1304932	4.18 Tb/s	3.05 Mb/s

Figure 1 – Overall US D3.0 PMA Stats

2.7. Dashboards

The dashboards provide insight into the performance and health of the system. Figure 2 shows a dashboard view of the performance of the US3.0 PMA system, showing statistics for the Mountain West region, including profile distribution, capacity gain, interface degradation, and the number of changes applied by the PMA System.

Figure 3 shows a dashboard view of the health of the US3.0 PMA system. Metrics shown include processing time and volume, message lag, and quantities of profile changes.




Figure 2 - Example US D3.0 PMA Stats for Mountain West Region




Figure 3 - Example US D3.0 PMA Monitoring Dashboard

3. DS D3.1 Profile Management

Capacity was increased while lowering cloud computing costs, by improving the data model, improving the clustering algorithm, and making performance improvements to the clustering algorithm. Additionally, the processing model was changed from batch-oriented to streaming. The performance improvements allowed the recommendations to be made more frequently as compared to the prior version, lowering the recommendation interval to 1 day from 3.5 days. As a result, approximately 8% more OFDM capacity was gained, more than 5 Tb/s in total.

3.1. Model

The MER data model was changed to use a two-dimensional histogram, or heatmap, incorporating an exponential decay function. Figures 7 and 8 below show representative histogram rendered using a color scale. This allows for the quick computation of any percentile value between 1 and 100 without sorting. The purpose of the decay function is to age out older data over time and introduce a recency bias. A half-life of 4 days was chosen for the decay function. An entire CMTS worth of compressed JSON data is stored in a single S3 file, organized by interface. On average, the data for a modem (with a single 96 MHz wide Orthogonal Frequency Division Modulation (OFDM) channel, 25 kHz subcarrier spacing) takes up only 75 kilobytes of space, or about 20 bytes per subcarrier. This space savings is the key to cost savings in cloud storage, as the full dataset is both read and written hourly.

Modem level per-subcarrier MER data is ingested on an hourly basis. The polling system produces a single S3 file per CMTS and sends an SNS notification when the file is ready. The large batch size minimizes costs for reading the incoming data. Due to the streaming nature of the system, there is a minimal delay between the poll completion and the processing.

The previous system (Harb, 2020) used a static MER percentile and a weight of 1 for each modem. Using a static 10th percentile of MER of each modem as an input to the clustering algorithm was insufficient to achieve optimal gains. This is based on the observation that more than 90% of the time the modem MER is higher than the 10th percentile MER and could thus be using a higher profile. On the other side, up to 10% of the time the MER is lower than the 10th percentile MER, and the modem might be using profile 0 (see below). As a result, the system under-estimated capacity gain.

The changes to the model allow for a more accurate estimation of capacity gain, which in turn allow for greater capacity gain. A fixed profile 0 at 64 or 256-QAM is used, and PMA generates 3 additional dynamic profiles. Depending on the MER distribution, a modem may spend a percentage of time on a single profile or a subset of all the profiles. Using the time percentages as weights, the estimated capacity for a modem is calculated as the weighted average of the profile bit-loading. The estimated capacity for an OFDM interface is then calculated as the average of the per modem estimates. Further increases of estimation accuracy at the OFDM interface level could be obtained by weighting the average by per-modem utilization. This would result in larger capacity gains and would be useful for comparing the predicted capacity against actual capacity but could unfairly punish lower utilization users. This is because modem utilization is not equal. For example, a modem with a high utilization on a low profile would skew the actual capacity downwards as compared to the estimate.

A single, optimal profile for a modem can be computed by iterating over the percentiles 1 – 100, estimating the resulting capacity as above, and choosing the percentile that maximizes capacity in conjunction with profile 0. Similarly, a set of N optimal profiles can be computed for a modem by choosing N percentiles of increasing value such that the chosen percentiles maximize capacity.

3.2. Clustering Algorithm

The clustering algorithm as described in (Harb, 2020) already minimizes capacity loss. The inputs were a profile per-modem (computed from the 10th percentile of MER values), with a weight of 1 for each profile. The clustering algorithm reduces the input to the desired number of profiles, e.g., 3, by merging the two profiles together that result in the lowest capacity loss. The improvements to capacity gain were achieved primarily by improvements to the inputs. Versus the previous inputs of a single profile for each modem, 2 optimal profiles are instead computed for each modem (as described above), as well as an additional input representing profile 0. Rather than a weight of 1.0 for each input, varying weights are used. The weight used for each input represents the estimate of the percentage of time spent on the profile, as mentioned above. The weight for profile 0 is the sum of the percentages of time each modem is estimated to use profile 0. Including its contribution to the profile 0 weight, the sum of the weights for each modem's inputs equals 1.

Rather than excluding modems with low MER prior to the clustering, the clustering algorithm merges those inputs as appropriate into the profile 0 cluster. As such, depending on the distribution of MER values, some modems may be ultimately assigned 0, 1 or 2 profiles.

Table 1 – OFDM Profile MER Thresholds, 5 dB more aggressive than DOCSIS 3.1 spec

Bits	Modulation	MER Threshold (dB)
2	QPSK	6.0
4	16-QAM	12.0
6	64-QAM	18.0
7	128-QAM	21.0
8	256-QAM	24.0
9	512-QAM	27.5
10	1024-QAM	31.0
11	2048-QAM	34.0
12	4096-QAM	38.0

3.3. Performance

The core clustering algorithm was made 25 times faster. The biggest speed increases were due to rewriting the distance function using SIMD parallel processing via the Advanced Vector Extensions 512 (AVX-512) instructions on Intel CPUs and in Java via the `jdk.incubator.vector` package. This allows for 16 integer operations to run in parallel on a single core. The next largest boosts came from switching the internal representation from double to int, and from immutable Scala Array types to primitive arrays, which was also necessary to take advantage of SIMD. Finally, upgrades to versions of the Java Virtual Machine (JVM), Scala, and Spark software were made.

Some of the performance gains were offset due to doubling the number of inputs. As the algorithm has an $O(N^2)$ complexity, doubling the inputs results in 4 times slower performance. The net speed increase of more than 6 times faster directly results in lower cloud computing costs. Rather than a single large batch, the work is spread throughout the day. This lowers cloud computing costs by keeping a smaller number of cores continuously busy. The cycle time was reduced from twice a week to daily, and is expected to be reduced to hourly on an as-needed basis, matching the MER ingest rate. As it was also observed in the

D3.0 US, faster response to changes in the plant result in greater capacity gain / and or less capacity loss. However, and an important difference: because the model and clustering account for MER changes over time, and the CMTS assigns profiles periodically throughout the day, a daily cycle is sufficient in most cases.

3.4. Capacity Gain

Greater gains are possible on OFDM than on D3.0 channels, due to the higher modulation orders available. The average capacity gain on OFDM is 43%. Actual capacity gain is measured vs a static 8-bit (256-QAM) modulation profile. As the maximum modulation is 12-bit (4096-QAM), the maximum gain possible on a perfect network would be 50%. This is about an 8% improvement to capacity vs the previous version. This represents an average bits / symbol across all OFDM channels of 11.47, compared to 10.83 on the prior version. In terms of raw capacity, the new DS D3.1 PMA is adding 27.1 Tb/s of total DS capacity, about 5 Tb/s more than the previous version.

Overall							
cmts_count	ds_ofdm_count	avg_speed	prof0speed	prof0bitload	prof0pct	gain	rawgain
2172	114596	801 Mb/s	476 Mb/s	6.93	1.26%	43.0%	27.1 Tb/s

by CMTS							
cmts +	ds_ofdm_count	avg_speed	prof0speed	prof0bitload	prof0pct	gain	rawgain
acr01.a1atlanta.ga.atlanta.comcast.net	64	828 Mb/s	561 Mb/s	8	0.547%	27.2%	9.78 Gb/s
acr01.a3atlanta.ga.atlanta.comcast.net	56	822 Mb/s	561 Mb/s	8	0.974%	32.8%	10.3 Gb/s
acr01.a4atlanta.ga.atlanta.comcast.net	48	820 Mb/s	561 Mb/s	8	1.99%	46.4%	12.5 Gb/s
acr01.a5atlanta.ga.atlanta.comcast.net	48	823 Mb/s	561 Mb/s	8	2.18%	46.6%	12.5 Gb/s
acr01.a6atlanta.ga.atlanta.comcast.net	48	825 Mb/s	561 Mb/s	8	2.24%	45.8%	12.3 Gb/s
acr01.abercomist.ga.savannah.comcast.net	40	822 Mb/s	561 Mb/s	8	1.52%	46.8%	10.5 Gb/s
acr01.abingdon.va.knox.comcast.net	40	820 Mb/s	561 Mb/s	8	2.90%	43.7%	9.81 Gb/s

Lowest 100 Ranked Interfaces							
cmts	if_name	avg_speed	prof0speed	prof0bitload	prof0pct +	gain	rawgain
acr06.northave.il.chicago.comcast.net	cable-ds-ofdm.12/4/48	553 Mb/s	561 Mb/s	8	98.8%	9.422%	2.35 Mb/s
acr04.spokane.in.indiana.comcast.net	cable-ds-ofdm.12/2/48	425 Mb/s	420 Mb/s	5	98.1%	-24%	-135 Mb/s
acr02.pembroke.fl.pompano.comcast.net	cable-ds-ofdm.9/15/48	575 Mb/s	561 Mb/s	8	90.9%	2.54%	14.8 Mb/s
acr05.hallandale.fl.pompano.comcast.net	cable-ds-ofdm.8/13/48	577 Mb/s	561 Mb/s	8	90.9%	2.81%	15.8 Mb/s
acr05.hallandale.fl.pompano.comcast.net	cable-ds-ofdm.8/14/48	581 Mb/s	561 Mb/s	8	90.9%	3.68%	20.6 Mb/s
acr02.pembroke.fl.pompano.comcast.net	cable-ds-ofdm.9/13/48	575 Mb/s	561 Mb/s	8	90.9%	2.54%	14.8 Mb/s
acr05.hallandale.fl.pompano.comcast.net	cable-ds-ofdm.8/15/48	577 Mb/s	561 Mb/s	8	90.9%	2.85%	16.0 Mb/s

Figure 4 – DS D3.1 PMA Profile Stats


3.5. Dashboards

The dashboards for DS D3.1 PMA show the system performance and health. Figure 4 shows the overall performance of the system, including the capacity gain, as well as a break-down by CMTS, and a view of the poorest performing interfaces.

Figure 5 shows a dashboard view of the performance of a specific OFDM interface. Metrics shown include traffic distribution by profile, channel speed by profile, interface statistics, etc.

Figure 6 shows a dashboard view of the health of the DS D3.1 PMA system, focused on AWS resources. Metrics shown include CPU utilization of the components, and metrics for the various notification topics and queues used by the system.

Figures 7 and 8 show example modem MER histograms, the optimal percentiles chosen as inputs to the clustering algorithm, and the resulting segmented profiles.



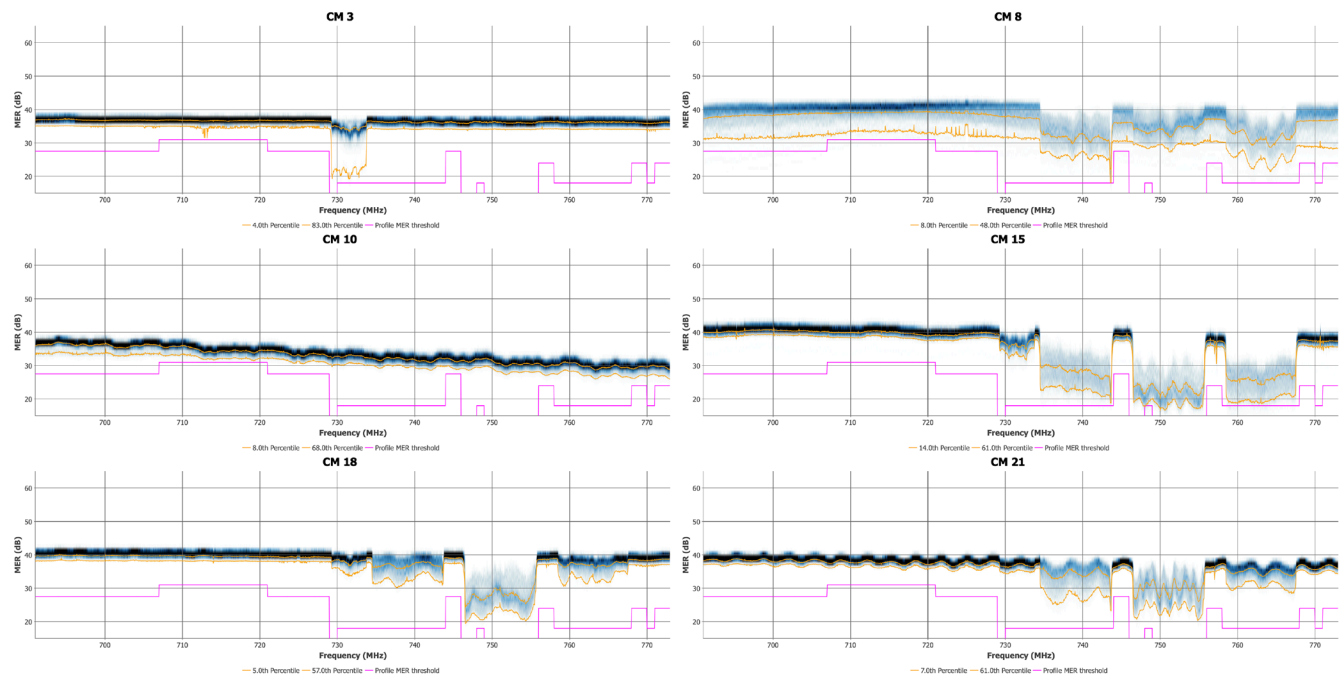


Figure 7 – Example profile 1 modem MER histograms, percentiles, and MER threshold

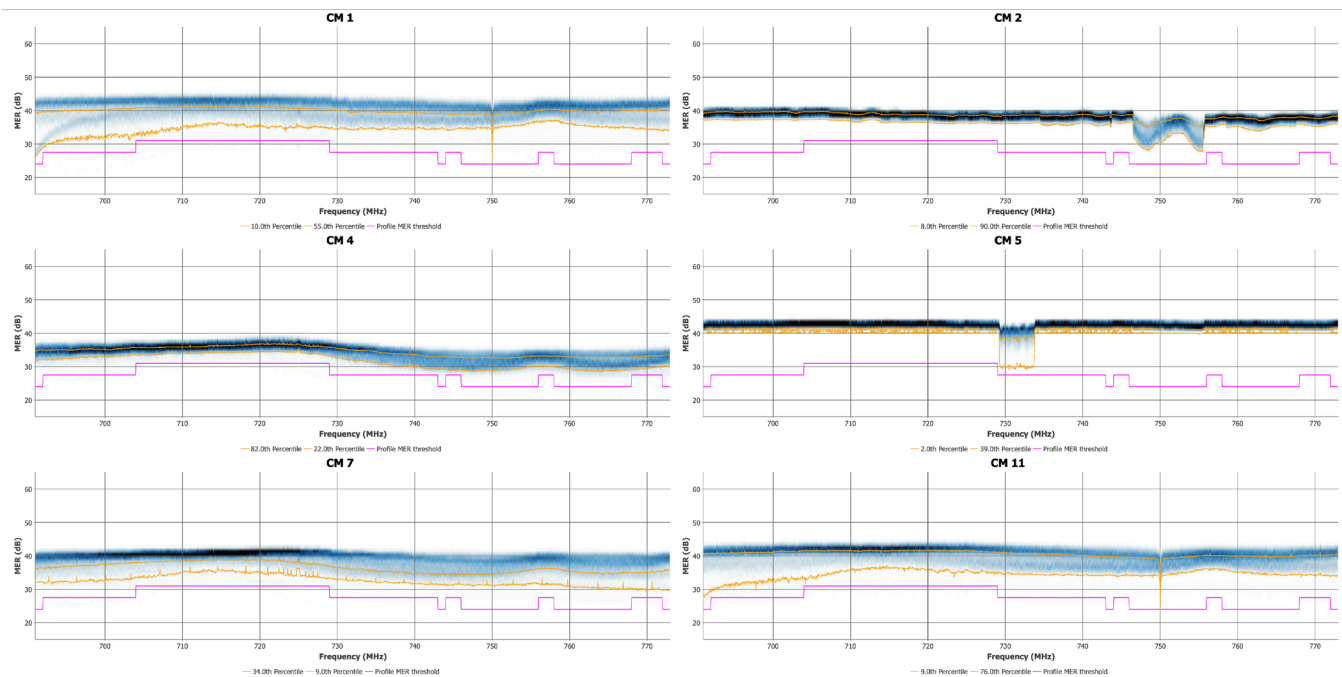


Figure 8 – Example profile 2 modem MER histograms, percentiles, and MER threshold

4. US D3.1 Profile Management

Profile Management for Orthogonal Frequency Division Multiple Access (OFDMA) on US D3.1 is similar to OFDM with a few notable differences. The MER values are per-minislot (400 kHz) instead of per subcarrier. The data is ingested at a 5-minute interval instead of hourly. The increased rate of collection helps offset the loss in granularity. The upstream uses IUCs instead of profiles, and more of them (7 vs 4). A single, static 10th percentile of MER per modem is used as the input to the clustering algorithm. This appears to be sufficient as there is less volatility in the MER values within a single modem as well as across the population of modems on an OFDMA interface. Having less volatility is a factor of the averaging of subcarrier MER into minislots, the upstream funnel effect, the ability of modems to adjust transmit power levels to achieve a desired receive power, and the effectiveness of the pre-equalization process. As of the time this writing, OFDMA PMA is not in full production, but based on trials, capacity gains similar to OFDM PMA are expected.

Table 2 – OFDMA Profile MER Thresholds from DOCSIS3.1 spec

Bits	Modulation	MER Threshold (dB)
2	QPSK	11.0
3	8-QAM	14.0
4	16-QAM	17.0
5	32-QAM	20.0
6	64-QAM	23.0
7	128-QAM	26.0
8	256-QAM	29.0
9	512-QAM	32.5
10	1024-QAM	33.5
11	2048-QAM	39.0
12	4096-QAM	43.0

Overall							by PP0D						
ppod_count	ofdma_count	avg_bit_rate	luc13bitrate	luc13pct	gain		ppod	ofdma_count	avg_bit_rate	luc13bitrate	luc13pct	gain	
35	1333	7.71	5	3.98	54.2		argapp104	2	7.85	5	0.219	57.0	
							cabapp107	1	7.86	5	0.0696	57.1	
							conopp103	19	7.73	5	2.90	54.6	
							fl06pp101	2	7.78	5	2.88	55.5	
							flarpp101	170	7.82	5	0.539	56.3	
							flarpp102	39	7.80	5	0.223	56.0	
Lowest 100 Ranked Interfaces													
ppod	rid	if_name	avg_bit_rate *	luc13bitrate	luc13pct	gain							
qampone104	6AWSD0930A	0a80.5/0/0	5.17	5	90.0%	3.41%							
flneec102	FLMRD00501	0a72.0/0/0	5.33	5	72.1%	6.58%							
qampone104	6AWSD0930A	0a9.6/0/0	5.42	5	82.5%	8.40%							
flneec104	FLQPD84801	0a106.6/0/0	5.53	5	78.9%	10.5%							
qampone102	A68TD16004	0a3.3/0/0	5.55	5	62.2%	11.0%							
qampone108	GAMPD09201	0a34.0/0/0	5.72	5	72.4%	14.4%							

Figure 9 – US D3.1 PMA IUC Stats




Figure 10 – Example IUC 5 modem MER percentiles, histograms, and threshold




Figure 11 – Example IUC 6 modem MER percentiles, histograms, and threshold




Figure 12 - Example IUC 9 modem MER percentiles, histograms, and threshold

5. Future Work

Increase the profile recommendation rate for Downstream and Upstream DOCSIS 3.1 from daily to hourly on as-needed basis. Faster reaction to changing conditions would result in even greater capacity gains, or less capacity loss.

Make further performance optimizations and improvements to the clustering algorithms. A better result can be achieved by adjusting the weights accordingly as the clusters are merged, and optimizations can be made to reduce the algorithmic complexity.

Investigate generating more profiles and using profile 0 to a greater effect by modifying it from a flat 64-QAM modulation to a segmented profile. In many cases profile 0 could be made to be more robust, have higher throughput, and automatically handle cases otherwise handled by exclusion zones.

The MER thresholds used in the clustering algorithms to determine modulation are estimates, and appropriate values for the thresholds have been chosen, albeit generally conservative ones. There are cases where more aggressive thresholds can be used without driving FEC rates and using them would result in greater capacity gains. There may also be cases where more conservative thresholds are necessary in order to lower uncorrectable FEC rates. In these cases, the capacity estimates that the PMA algorithm makes would be high, and modems would be using profile 0 at a greater percentage than expected. Lowering the thresholds would lower estimated capacity but increase actual capacity, by shifting traffic to higher profiles. By comparing estimated profile utilization with actual profile utilization, the MER thresholds may be refined, creating realistic thresholds per CMTS, interface, or modem.

6. Conclusion

By making strategic changes to the US D3.0, DS D3.1, and US D3.1 PMA software, the system is able to mitigate issues faster, increase network capacity, and lower cloud computing costs. This was done by eliminating redundancy and leveraging a combination of streaming and batching, cloud computing, and vector processing. Improvements were also made to the underlying data model, business logic, and clustering algorithms. The net result was a reduction in the US D3.0 recommendation interval from 6 hours to 5 minutes, and from 3.5 days to 1 day for DS 3.1 and US 3.1. Capacity was increased by 1.5% for US D3.0 and by 8% for DS D3.1.

Abbreviations

AVX-512	advanced vector extensions, 512-bit SIMD instruction set on Intel
AWS	Amazon web services, a cloud computing platform
CMTS	cable modem termination system
CPU	central processing unit
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
dB	decibel
DOCSIS	data over cable service interface specification
DS	downstream
EKS	elastic Kubernetes service
FEC	forward error correction
Hz	hertz
IUC	interval usage code
JSON	JavaScript object notation
JVM	Java virtual machine
MER	modulation error ratio
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PMA	profile management application
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying, a 2-bit modulation form
S3	simple storage service, object storage on AWS
SIMD	single instruction multiple data, a type of parallel processing
SNR	signal to noise ratio
SNS	simple notification service, message delivery on AWS
SQS	simple queue service, message queuing on AWS
US	upstream

Bibliography & References

1. “Full Scale Deployment of PMA”, SCTE Cable-Tec Expo 2020,
https://www.scte.org/documents/3132/1781_Harb_3030_paper.pdf, M. Harb, B. Santangelo, D. Rice, J. Ferreira.
2. “A Machine Learning Pipeline for D3.1 Profile Management”, SCTE Cable-Tec Expo 2019,
https://www.scte.org/documents/3816/799_Harb_A_Machine_Learning_Pipeline_for_D3.1_Profile_Management_paper.pdf, M. Harb, J. Ferreira, D. Rice, B. Santangelo, R. Spanbauer.
3. “Practical Implementation of Profile Management Application (PMA) to Improve Data Throughput in the Presence of Impairments”, SCTE Cable-Tec Expo 2021,
https://www.scte.org/documents/4700/2121_Volpe_3479_paper.pdf, B. Volpe.
4. “Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)”, SCTE Cable-Tec Expo 2019,
https://www.scte.org/documents/3867/806_Sundaresan_Practical_Lessons_D3.1_Deployments_Profile_Management_Application_PMA_p.pdf, K. Sundaresan, J. Zhu, M. Mishra, J. Lin
5. Cablelabs, Data-Over-Cable Service Interface Specifications DOCSIS[®] 3.1, Physical Layer Specification, CM-SP-PHYv3.1-I02-EC, 2014.
6. “AVX-512”, Wikipedia, <https://en.wikipedia.org/wiki/AVX-512>

Preparations for Deploying & Lessons Learned from Deploying High Split (204 MHz) on I-CCAP, R-PHY, & R-MACPHY

High Split as a Steppingstone Towards DOCSIS[®] 4.0

A Technical Paper prepared for SCTE by

Craig Coogan

Vice President, Product Management
CommScope, Inc.
Lisle, IL
+1 (630) 281 3143
craig.coogan@commscope.com

Jamie Brown

Director, Next Generation Access Network Technology
Shaw Communications, Inc.
Calgary, AB, Canada
+1 (403) 750 4591
jamie.brown@srjb.ca

Zoran Maricevic, Ph. D.

Engineering Fellow
CommScope, Inc.
Wallingford, CT
+1 (203) 303 6547
zoran.maricevic@commscope.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. High Split & Cable Access Architecture Overview.....	4
2.1. High Split Overview – What is It?.....	4
2.2. Market Drivers for High Split.....	7
2.3. Historic & Projected Service Group Capacity of DOCSIS Versus PON.....	8
2.4. What Does High Split Buy the Cable Operator?.....	11
3. Architectural Considerations for High Split.....	13
3.1. CMTS Architectures: I-CCAP, R-PHY, & R-MACPHY.....	13
3.2. Spectral Considerations.....	16
3.3. Leakage Detection.....	17
3.4. Customer Premise Equipment (CPE) Considerations.....	18
3.4.1. Potential High-Split CPE Interference on Legacy CPE.....	19
4. Operational Considerations for High Split.....	21
4.1. Upstream Level Considerations.....	21
4.1.1. Levels Out of Cable Modem (CM).....	21
4.1.2. Levels into RF Amplifier Upstream Ports.....	24
4.2. Field Components Upgrade Considerations – “Rip and Replace” or “Modules Only?”.....	27
5. Shaw Communications – High-Split Lessons Learned.....	29
5.1. Why the Need to Modernize the Network Beyond Mid Split?.....	29
5.2. Key High-Split Readiness Activities (with a Mind Towards DOCSIS 4.0).....	29
5.2.1. Network Architecture Readiness.....	29
5.2.2. Set-Top Boxes with Legacy Out-of-Band Signaling.....	30
5.2.3. On-Premises Architecture.....	30
5.3. A Potential Evolution Path to DOCSIS 4.0.....	32
5.4. Summary of Lessons Learned.....	35
6. Conclusion.....	36
Abbreviations.....	38
Bibliography & References.....	41

List of Figures

Title	Page Number
Figure 1 – Example Spectra & Overview of Low/Sub Split (42 MHz), Mid Split (85 MHz), & High Split (204 MHz).....	7
Figure 2 – DOCSIS Service Group Bandwidth Over the Years Utilizing DOCSIS 3.1 Technology.....	8
Figure 3 – DOCSIS Service Group Bandwidth Over the Years, Including Future Expansion Using DOCSIS 4.0.....	10
Figure 4 – Max Subs per SG for Low, Moderate & High DS Tavq growth, 1,218/204 MHz.....	11
Figure 5 – I-CCAP Centralized Access Architecture.....	13
Figure 6 – Remote PHY Distributed Access Architecture.....	14
Figure 7 – Remote MACPHY Distributed Access Architecture.....	14
Figure 8 – DOCSIS 4.0 Frequency Split Roadmap vs. FCC-governed Aeronautical Bands.....	18
Figure 9 – Two Potential High-Split CPE Interference Scenarios.....	19
Figure 10 – D3.1 PHY Spec Snippets, Requiring TCP of 65 dBmV Out of the CM.....	23
Figure 11 – Max HS CM Levels (Green) with Various Ultra-High-Split Options of D4.0 FDD.....	23

Figure 12 – Dynamic Range Window (DRW) Considerations from D3.1 PHY	24
Figure 13 – An Example Section of HFC Plant in Between Two RF Amplifiers	25
Figure 14 – Actual Per-channel DS Levels, at an RF Amplifier Output and at Customer Premises	25
Figure 15 – CM Upstream Levels, and the Resulting Input-into-RF Amp Levels	26
Figure 16 – Conditioned Taps Example for an Even Further Optimized Plant Section.....	27
Figure 17 – How to Replace a Node – Initiation and Replumbing Part of the Process	28
Figure 18 – Setup & Wrap-up Part of the Node Replacement Process	28
Figure 19 – Step 1: D3.1 High Split (204 MHz / 1 GHz)	33
Figure 20 – Step 2: D3.1 High Split with 1.8 GHz Taps (in 3.0 GHz Housings)	34
Figure 21 – Step 3: D3.1 High Split with 1.8 GHz Taps (in 3.0 GHz Housings) and 1.8 GHz Amplifiers...	34
Figure 22 – Step 4: DOCSIS 4.0 High Split	35

List of Tables

Title	Page Number
Table 1 – Raw Throughput Examples for Low/Sub Split, Mid Split, & High Split	6
Table 2 – Upstream Cable Modem Levels from D2.0 RFI Specification	21
Table 3 – Assumed Upstream RF Channel Characteristics from D2.0 RFI	22
Table 4 – D3.0 per Channel Power Levels Out of Cable Modem, with Highlighted Max Points	22
Table 5 – Node Serviceability Flags for Back-office Software in Support of High Split.....	32

1. Introduction

Vendors who make equipment for hybrid fiber-optic & coaxial cable (HFC) networks and cable service providers (also known as “cable operators”) have been discussing high split (HS)—moving the upstream (US) frequency split to 204 MHz—for quite some time now. Over the last few years, vendors have made available the various components required for high-split operation, and they continue to add features and capabilities to simplify the introduction of high split by operators into HFC networks.

The time for talk is over! Some operators have recently taken the brave plunge in deploying high split, while others are busy preparing to migrate to high split in the near-term. This paper will explore the drivers for pursuing high split (204 MHz) on an HFC network. The paper will explore what it takes to deploy high split, including the consideration of architectural issues that need addressing to support high-speed data (HSD), video, and voice over internet protocol (VoIP) services on high-split HFC networks for both residential and commercial services. Furthermore, the paper will explore the considerations for each of the services and the features that are required in the products that comprise the end-to-end solution that enables high split.

The paper will present a unique cross-industry view of a cable operator who is actively deploying high split as well as a vendor who is providing most of the elements to enable this type of migration. The paper will explore the considerations for three of the most popular HFC access architectures today, including Integrated CCAP (I-CCAP) as a Centralized Access Architecture (CAA) as well as the two leading Distributed Access Architectures (DAAs): Remote PHY (R-PHY) and Remote MACPHY (R-MACPHY). There is value in deploying high split in all three of these architectures, and the technology is available today. The paper will explore the main issues which must be addressed in preparation for deployment, as well as key lessons learned from both the preparations and deployments. This paper will position high split as an important steppingstone to DOCSIS 4.0 and explore an HFC network’s evolution through this lens.

2. High Split & Cable Access Architecture Overview

This section will provide an introduction and overview on what exactly high split is, what are the market drivers for high split, and what does high split gain for cable operators in order to lay the groundwork for reviewing the architectural and operational considerations for implementing high split.

2.1. High Split Overview – What is It?

The current return spectrum in HFC networks, typically allocated between 5 MHz and 42 MHz (65 MHz for some geographic regions), has been a critical resource over the years to support the enablement and growth of interactive services. The Data Over Cable Service Interface Specification (DOCSIS) specification versions that define how to provide standards-based high-speed data service up to and including DOCSIS 3.1 have traditionally used frequency division duplex (FDD) to split the upstream (US) and downstream (DS) bands and their associated traffic with a guard band in between the two to prevent interference.

The technology behind the HFC upstream path has grown in complexity and efficiency over time. More recently, DOCSIS 3.1 brought orthogonal frequency domain multiple access (OFDMA) digital modulation technology to the return path along with a higher 204 MHz split. OFDMA improved resiliency to return path noise and impairments, which has resulted in support for even higher speeds and capacities without the need for physical expansion of the return bandwidth spectrum.

However, the 42/65 MHz upstream has become cable's Achilles Heel. With gigabit downstream tiers, it becomes very difficult to pair a complimentary upstream tier with it. This creates the need for upstream bandwidth augmentations which is achieved through implementation of higher frequency splits. DOCSIS 3.1 technology introduced an upstream spectrum of 5-204 MHz, also referred to as a high-split (HS) configuration, with downstream spectrum starting at 258 MHz. The latest DOCSIS 4.0 specification provides an extended FDD option to enable FDD upstream spectrum up to 684 MHz while also providing for incremental downstream bandwidth capacity to offset the expansion of the upstream spectrum.

The vast majority of HFC networks today were built to either 750 MHz or 860 MHz as the maximum DS frequency with sub-split returns in the 5-42 MHz range (5-65 MHz range for some other regions). Expanding services are driving consideration for network expansions to address growing capacity needs in both the downstream and upstream directions. Past expansions were focused on expanding the downstream, while leaving upstream spectrum mostly untouched. As upstream bandwidth expansions are planned, particularly for sub-split and mid-split architectures evolving the return to 204 MHz, a phased approach should be considered that not only considers the impact on downstream bandwidth and services, but also considers the impact of future DOCSIS 4.0 technology options. This will be further explored later in the paper.

In comparison to the upstream spectrum, the downstream spectrum has stretched from 54 MHz (or approximately 87 MHz in other regions) up to 550 MHz, 760 MHz, 860 MHz or, more recently, 1,002 MHz or 1,218 MHz. Whichever is the maximum DS frequency, the DS spectrum clearly has much wider spectrum and supported bandwidth and capacity than the US has. It has been clear for some time that cable operators would need to eventually expand the utilized US spectrum in order to meet the needs of their subscribers as well as to compete in the broadband access market. The drivers for expanding US spectrum will be further described in Section 2.2.

Over time, both the DOCSIS standards and the equipment that implements these standards enabled an ability to expand the upstream spectrum. Up through the DOCSIS 3.1 standard, cable operators can expand the US spectrum to two additional levels:

- (1) Mid split (MS): US spectrum of 5-85 MHz with DS spectrum starting at 108 MHz
- (2) High split (HS): US spectrum of 5-204 MHz with DS spectrum starting at 258 MHz.

Eventually, as DOCSIS 4.0 technologies are widely adopted, additional Ultra-High Splits (UHS) will open the door to even greater upstream bandwidths and capacities with splits that can reach 300 MHz, 396 MHz, 492 MHz, or 684 MHz. However, regardless of which upstream expansion option is selected, operators must recognize certain technical challenges and be prepared to address them as will be explored in the rest of this paper.

An example of the raw total throughput that a cable operator could send in the downstream and receive in the upstream for each service group (SG) is instructive of the relative throughputs that each of the splits outlined above could potentially provide. Table 1 shows three separate implementations, one for each of low split, mid split, and high split. For this example, a mixture of DOCSIS 3.0 single carrier QAM (SC-QAM) and DOCSIS 3.1 orthogonal frequency-division multiplex (OFDM) channels was assumed in the DS, and likewise, a mixture of DOCSIS 3.0 advanced time division multiple access (ATDMA) SC-QAM and DOCSIS 3.1 orthogonal frequency-division multiple access (OFDMA) channels was assumed in the US, as outlined in Table 1.

Table 1 – Raw Throughput Examples for Low/Sub Split, Mid Split, & High Split

US Frequency Split	US Spectrum Usage (MHz)	DS Spectrum Usage (MHz)	US Raw Throughput (Mbps)	DS Raw Throughput (Mbps)
Low/Sub Split	Total: 5-42 ATDMA: 18-42 OFDMA: N/A	Total: 54-860 SC-QAM: 54-588 OFDM: 588-860	87	5380
Mid Split	Total: 5-85 ATDMA: 18-42 OFDMA: 42-85	Total: 108-1,218 SC-QAM: 108-642 OFDM: 641-1,218	575	7812
High Split	Total: 5-204 ATDMA: 18-42 OFDMA: 42-204	Total: 258-1,218 SC-QAM: 258-642 OFDM: 642-1,218	1302	6912

In this example, the maximum DS frequency when operating in low/sub split is assumed to be 860 MHz, a starting point for many cable operators today, and the example also assumes that the operator is utilizing DOCSIS 3.0 SC-QAMs only at this stage in the US, representing yet again another common starting point for cable operators. When moving to high split, the DS loses its lower 204 MHz of spectrum, and a move to 1,218 MHz in the DS is assumed to maximize total spectrum allowed when using DOCSIS 3.1-capable equipment. The example in Table 1 just assumes raw throughput to and from the SG, independent of whether the channel is used for MPEG or DOCSIS QAM.

For the US in this example, usable bandwidth was assumed to be starting at 18 MHz due to noise that typically exists in the 5-18 MHz range caused by noise funneling from the cable plant. However, operators can and have harvested this lower bandwidth utilizing the capabilities of the newer DOCSIS 3.1 OFDMA channels, but that is not explored in this example. Figure 1 shows an overview of the low split, mid split, and high split spectrum and the estimated raw throughput from Table 1. The migration to services that are closer to symmetric is a key driver for cable operators to compete with operators who offer services over a fiber-to-the-home (FTTH), fiber-to-the-premise (FTTP), fiber-to-the-building (FTTB), or generally, a fiber-to-the-X (FTTx) architecture. The competitive drivers for a migration to high split will be covered in Section 2.2.

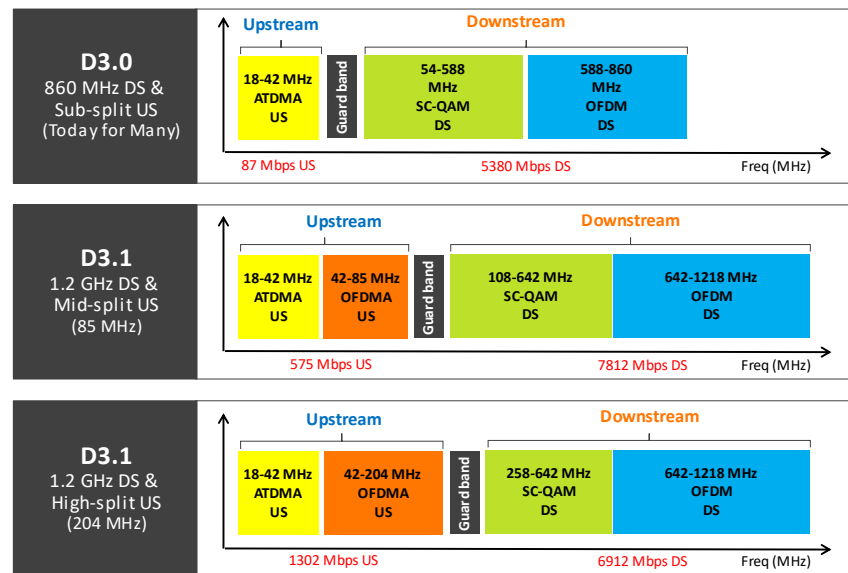


Figure 1 – Example Spectra & Overview of Low/Sub Split (42 MHz), Mid Split (85 MHz), & High Split (204 MHz)

2.2. Market Drivers for High Split

Cable operators face different levels of competition that vary by market. Historically, cable operators had competed against the legacy telecommunication providers (also known as the “telcos”), who initially offered high-speed data utilizing digital subscriber line (DSL) service over twisted-pair copper lines. Many telcos and other alternative operators have opted to implement a fiber passive optical network (PON) for FTTx services. Service providers who are utilizing PON are offering service tiers up all the way up to 5 to 8 Gbps, with many offerings currently in the range of 2 to 3 Gbps as of the writing of this paper. Deploying some versions of PON for FTTx service offerings enables symmetric services, and this distinguishes these offerings from service offerings over HFC.

While cable operators do face tough competition from service providers that utilize DSL, satellite, and fixed wireless access (FWA), the services offered on these technologies generally compete with only specific segments of the broadband internet market either based on location or service availability as well as services that generally fit the lower end of the bandwidth speeds offered, which generally also align with lower cost. Cable operators face much tougher competition from service providers who offer services over FTTx PON due to the higher speeds, potential for symmetrical services, and lower latencies that this technology provides. Therefore, the existence of service providers using FTTx technologies overbuilt on an HFC network and/or the threat of a service provider overbuilding an FTTx network on an HFC network are two of the main market drivers for operators moving their DOCSIS deployments to high split.

Due to the actual or potential competition from FTTx, cable operators need to increase their DS SLA speeds, and when they do this, they also need to increase the US SLA speeds for two main reasons: (1) to have enough bandwidth in the US to support the required Transmission Control Protocol (TCP) Acknowledgements (ACKs) to sustain the offered DS speed and (2) to offer US speeds that may be required or desired by a subset of their subscribers, most notably business/commercial subscribers as well as gamers and video uploaders.

Deploying FTTx services typically comes at high cost whether it is done by cable operators, telcos, or other service providers. Cable operators generally have a key advantage, especially against the potential for a new fiber optic overbuilder: the advantage is the cable network they already have in place! One of the hallmarks of DOCSIS has been the ability to enable cable operators to expand services in an evolutionary manner with incremental investments while maintaining backward compatibility for existing services, as is explored in a separate paper entitled “Network Migration to 1.8 GHz – Operational “Spectral Analysis” Measured in nano-Hertz, a 30-Year Perspective” [Maricevic_2022]. This ability afforded by DOCSIS is a key reason why cable operators are now at the point of implementing or seriously considering a move to high split (204 MHz) in the US and 1.2 GHz in the DS. The decision for an operator to deploy PON or expand with DOCSIS is not an “either/or” situation but can many times be a coexistence story where the operator deploys PON in strategic portions of its network where it makes business sense in addition to high split with 1,218 MHz on HFC. Each operator needs to evaluate its own economic, competitive, cost, strategic, etc. situation to decide which path to go relative to expanding capacity with a pure HFC network implementation, or with a hybrid HFC / PON type deployment.

2.3. Historic & Projected Service Group Capacity of DOCSIS Versus PON

Since FTTx PON competition is a key market driver for cable operators to implement high split, it would be instructive to show how DOCSIS has evolved service group bandwidth over the years and to project how it might grow. This section explores a couple of potential evolutions of an HFC network SG with both spectral changes and channel lineup changes and how that impacts the raw bandwidth per service group (SG) in the DS and US.

Figure 2 shows an example of how SG bandwidth capacity has changed over the years, illustrative of the general DS and US SG bandwidth trends over time. The example only shows an expansion up to DOCSIS 3.1 technology and maximizing out at 1.2 GHz DS and 204 MHz US.




Figure 2 – DOCSIS Service Group Bandwidth Over the Years Utilizing DOCSIS 3.1 Technology

In this example, by 2020, the operator had grown to utilize 32 SC-QAM channels and a single 192 MHz OFDM channel in the DS while the upstream remained at low split (42 MHz). The operator in this example utilized a single 3.2 MHz and four 6.4 MHz ATDMA SC-QAM channels in the US. This yielded just over 3 Gbps per SG DS, which was roughly 25 times what could be achieved in the US.

In 2022, this example added another 192 MHz block of OFDM in the DS and opted to go to mid split (85 MHz). The DS SG throughput grew to nearly 5 Gbps, while adding OFDMA to fill up the remaining 85 MHz of spectrum grew the US SG throughput to approximately 540 Mbps. This reduces the DS:US SG bandwidth ratio to about 10x. Had the operator instead opted to move to high split (204 MHz) and fill the remaining spectrum up to 204 MHz with OFDMA channel capacity, they would have increased the US SG throughput to approximately 1.5 Gbps, further dropping the DS:US SG bandwidth ratio to about 3x. While moving to high split brought the operator closer to symmetrical in DS versus US, the move did not appreciably increase the DS throughput nor efficiently utilize the spectrum that was available.

The final step in this sample HFC SG evolution has the operator moving to 1.2 GHz of spectrum in the DS and 204 MHz in the US, yielding approximately 8.6 Gbps throughput in the DS and 1.5 Gbps in the US with the ratio of DS:US at approximately 5.5x.

By way of comparison, PON offers both asymmetrical and symmetrical offerings, depending on the version of PON in use. Gigabit passive optical networking (GPON) provides 2.5 Gbps DS with 1.2 Gbps US for a 2:1 ratio. The Institute of Electrical and Electronics Engineers (IEEE) specifies a 10 gigabit per second Ethernet passive optical networking (10G EPON) downstream that can be deployed with either a 1 Gbps or a 10 Gbps upstream for either a 10:1 or a 1:1 ratio. The International Telecommunication Union (ITU) Telecommunications Specifications Sector (ITU-T) specifies 10 Gbps PON (10G-PON), which is also known as XG-PON, that pairs 10 Gbps DS with 2.5 Gbps US at a 4:1 ratio and 10 Gbps symmetrical PON (XGS-PON) that provides symmetric 10 Gbps in both DS and US for a 1:1 ratio.

When including the impact of forward error correction (FEC) on throughput, XGS-PON (and other 10G PON technologies) net capacity is ~8.5 Gbps to the SG. Considering traffic engineering, average throughput per user, and a SG size of 64 subscribers, the maximum service level agreement (SLA) that XGS-PON technology can offer at a reasonable subscriber quality of experience is ~7.5 Gbps. Therefore, the 7.5 Gbps potential SLA from XGS-PON sets the bar for the near-term future evolution of DOCSIS and HFC networks to strive for.

Figure 3 below shows an alternative SG bandwidth evolutionary path along with providing a longer future projection of the HFC network SG bandwidths. The example includes those from Figure 2 above through year 2020. In 2022 in the new example, the operator has chosen to jump straight from an 860 MHz DS and low split (42 MHz) cable plant to a 1.2 GHz DS and high split (204 MHz) cable plant, taking a larger leap in an earlier timeframe and skipping the work associated with implementing a mid-split cable plant. Starting from 2025 and later, the operator has chosen to implement D4.0 FDD, which is also known as Extended Spectrum DOCSIS (ESD), to further increase the DS and US SG bandwidth.




Figure 3 – DOCSIS Service Group Bandwidth Over the Years, Including Future Expansion Using DOCSIS 4.0

When the operator moves to D4.0 FDD, even though the maximum DS frequency can be moved up to as high as 1.8 GHz, the DS bandwidth in this example does not grow as much relative to the high split implementation from 2022 because some of what had been DS spectrum is now allocated to the US. The example shows that as the operator moves the UHS to higher frequencies, the US bandwidth grows while the DS bandwidth shrinks as the operator allocates more spectrum to the US at the expense of the DS. In this example, the operator moves through three different UHS settings: 300, 396, and 492 MHz. The example in Figure 3 shows the following:

1. DOCSIS 4.0 can achieve greater than 8.5 Gbps throughput per DS SG, more than what XGS-PON can provide
2. DOCSIS 4.0 can enable the HFC network to offer services that are much closer to symmetrical than prior implementations of DOCSIS
3. Implementing 1.2 GHz DS and high split (204 MHz) US is a clear steppingstone on the path to DOCSIS 4.0 and provides a significant amount of value with equipment that is available today.

The example in Figure 3 focuses the future steps to DOCSIS 4.0 as following the extended spectrum FDD path. A similar story could be laid out for operators who choose to implement DOCSIS 4.0 Full Duplex DOCSIS (FDX). One of the key technical differences between D4.0 FDD and D4.0 FDX is that D4.0 FDX utilizes some portions of the spectrum for both DS and US transmission while keeping the maximum DS frequency set at 1.2 GHz. An operator could chart a bandwidth evolution course to high split and then on to either D4.0 FDD or D4.0 FDX and, in theory, achieve similar results. The factors for

deciding between D4.0 FDD and D4.0 FDX are beyond the scope of this paper, but regardless of that future path, making the step to high split now puts the operator in a better position to compete and set up for the next phase of its network evolution.

At the January, 2019 Consumer Electronics Show (CES), the Internet & Television Association (NCTA), CableLabs[®], and Cable Europe announced an industry initiative generally called 10GTM, which was defined as “the cable industry’s vision for delivering 10 gigabit networks” and “a powerful, capital-efficient technology platform that will ramp up from the 1 gigabit offerings of today to speeds of 10 gigabits per second and beyond – to consumers in the United States and across the globe in the coming years” [NCTA 10G]. DOCSIS 4.0 is clearly one of the key technologies that enables 10GTM, and as high split is a logical steppingstone to DOCSIS 4.0, it is also clearly also one of the smartest steps an operator can take on the Path to 10GTM.

2.4. What Does High Split Buy the Cable Operator?

CommScope has been modeling HFC network operator bandwidth trends for well over a decade now and has utilized this data to project how operators can evolve their plants to meet the projected bandwidth demands of their end subscribers. A separate paper entitled “Broadband Capacity Growth Models – Will the end of Exponential Growth eliminate the need for DOCSIS 4.0?” by John Ulm, Dr. Zoran Maricevic, and Ram Ranganathan [ULM_2022] provides some excellent insights about what implementing high split actually buys the cable operator that will be summarized here.

The completed study performed modeling of a high (21% CAGR), moderate (16% CAGR), and low (linear) DS average busy hour user throughput (Tavg) growth rates on a 1,218/204 MHz plant offering a 5 Gbps DS, 1 Gbps US SLA. It analyzed SG size and the maximum number of subs supported over a 10-year window. The results are shown in Figure 4.




Figure 4 – Max Subs per SG for Low, Moderate & High DS Tavg growth, 1,218/204 MHz

For comparison purposes, Figure 4 also includes the 2018 projections of 43% CAGR “uber growth” – CAGR much higher than the more recent tapering of DS bandwidth growth. The 2018 projections show the max subs per SG supported would drop to 28 subs by 2030. This drove many people to think that FTTx would be required by 2030. However, the reality is that a 1,218/204 MHz plant supporting 5G x 1G tiers can easily last into the next decade, maybe even further if the slower growth projections hold.

In this case, the operator moves to high split in 2024 while simultaneously completing its migration to IPTV in the same year, reducing the legacy MPEG QAM video spectrum to zero. In 2025, the 5G DS tier is introduced and fills up most of the available 1,218 MHz of spectrum. Note that the SG size is still at 400 subs. After 2025, the SG size is reduced as needed to keep within the allotted 1,218 MHz.

Perhaps the key point of this 1,218/204 MHz case study is that a node with 150+ subs can be upgraded to 1,218/204 MHz and support a service tier of 5 Gbps x 1 Gbps for the next decade and beyond. There is no pressing near-term need to push the HFC to very small (and inefficient!) SG sizes, that could be, for example, achieved in N+0 systems. Even if the throughput follows high growth rates, the cable operator is not stuck and can node split further to a smaller SG size to meet the needs of their subscribers and/or migrate to DOCSIS 4.0.

3. Architectural Considerations for High Split

3.1. CMTS Architectures: I-CCAP, R-PHY, & R-MACPHY

There are three predominate Cable Modem Termination System (CMTS) / Converged Cable Access Platform (CCAP) architectures in the market today:

1. **Integrated CCAP (I-CCAP)**
2. **Remote PHY (R-PHY)**
3. **Remote MACPHY (R-MACPHY)**

I-CCAP is a Centralized Access Architecture (CAA) whereby an I-CCAP provides DOCSIS Media Access Control (MAC) and physical layer (PHY) functionality in a single, integrated, highly available chassis that provides high-speed data, voice, and video services on HFC networks. Figure 5 shows an example of an I-CCAP deployment in an HFC cable network.




Figure 5 – I-CCAP Centralized Access Architecture

R-PHY is a DAA whereby the PHY component of the CCAP is moved from a centralized location (in an I-CCAP) and out to a Remote PHY Device (RPD), which is housed in a shelf or a fiber node. Figure 6 shows an example R-PHY architecture implementation. The RPD handles the RF generation and reception of the signals that traverse the cable plant. The CCAP MAC functionality is provided by either a physical CCAP Core, which is typically an I-CCAP that has been evolved to support R-PHY operation, or a virtualized CCAP Core that runs on common off-the-shelf (COTS) servers in the headend. From headend to RPD, the signals typically traverse an Ethernet network and over digital optics. The migration from analog / amplitude modulated optics as utilized in an I-CCAP CAA architecture to digital optics provides benefits by reducing the noise introduced on the HFC plant by the analog optics. Given the reduction in noise on the HFC network and the pushing of the RF signal generation to the edge of the network, the channels can be received typically with higher modulation error ratio (MER) than in an I-CCAP architecture, and this results in the ability to utilize higher modulation orders and a more efficient use of the RF spectrum.

Furthermore, moving the PHY function to a fiber node can reduce the cable operator's rack space, power, and cooling requirements for the cable headends and operator facilities. In this architecture, the RPDs are aggregated by a Converged Interconnect Network (CIN), which is comprised of switches/routers that connect the RPDs to the cores, software systems, and the internet to round out the network.




Figure 6 – Remote PHY Distributed Access Architecture

R-MACPHY is another DAA whereby both the MAC and PHY components of the CCAP are located in a Remote MACPHY Device (RMD), which can also reside in a shelf or a fiber node. Figure 7 shows an example R-MACPHY architecture implementation.

Like R-PHY, R-MACPHY also utilizes Ethernet and digital optics for the fiber optic distribution network and has similar benefits as R-PHY. With R-MACPHY, both the MAC and PHY functionality are moved to the remote location—typically in a fiber node—and this colocation has some additional benefits including better latency performance and, depending on the application and implementation, the potential to remove the need for the IEEE 1588 Precision Timing Protocol (PTP) grandmaster clock and associated timing network, further simplifying the network.

Similar to R-PHY, RMDs are also aggregated via a CIN. In the cable operator headend and similar to R-PHY, the RMDs connect to OSS/BSS, a video core, and, in the newer Flexible MAC Architecture (FMA) as defined by CableLabs, an FMA Core and/or MAC Manager. RMD is the main access component, which is an intelligent device performing the following functions, among others:

- Support DOCSIS MAC functionality, including DOCSIS signaling functions
- Provide all PHY-related circuitry such as downstream QAM & OFDM modulators and upstream QAM & OFDMA demodulators
- Convert downstream MPEG video received from a video core and downstream legacy out-of-band signals received over a digital transport link, such as Ethernet, into analog for transmission over RF
- Convert upstream legacy analog out-of-band signals received over RF into digital for transmission over a digital transport link, such as Ethernet.

R-MACPHY deployments arguably consume less total system power than either the I-CCAP or R-PHY alternatives.




Figure 7 – Remote MACPHY Distributed Access Architecture

The great news for cable operators is that a migration to high split is both possible and can make sense for any of these architectures, and high split deployments are currently in the field on all three architectures using today's technology!

When an operator upgrades the HFC network to high split, the main components that get upgraded in all of these architectures are the actives in the cable plant, including the fiber node and RF amplifiers (line extenders (LEs) and bridger amps (MBs)). Whether deploying a CAA or DAA, the RF tray in the node needs to be upgraded to support high split (204 MHz) operation. Operators may typically want to leverage the fiber node visit to also move the downstream to support up to 1.2 GHz. With a CAA, the amplitude modulated optics can be implemented as analog in the return path, but the use of digital optics in the return path as offered by some vendors provides significantly improved performance. Generally, analog return path is better suited to shorter fiber links (approximately 10 miles or less) while digital return provides even more pronounced benefits at longer fiber links. When implementing a DAA, then the fiber optic link is handled by Ethernet links and digital optics.

Independent of implementing a CAA or DAA, the considerations for the upgrades to the actives and passives in the cable plant between the fiber node and the subscribers are the same. Generally, the RF amplifiers (actives) need to be upgraded to support 204 MHz operation in the upstream path, and most of the options already include support for 1.2 GHz operation in the downstream path as well. These upgrades to the amplifiers typically allow for the housing to remain in place while the active part of the product is swapped out. This saves time and cost and reduces the potential for problems to arise from cutting out and replacing products. Regarding passives, if taps currently deployed are limited to 1 GHz operation, the operator may choose to upgrade these to support 1.2 GHz via a tap faceplate swap while working on the cable plant. When operators implement the upgrade to high split, they will generally need to upgrade the node and subtended amplifiers at the same time or avoid using a portion of the spectrum while making the switch. If an operator is using conditioned taps, then the conditioning in place—whether via attenuation, equalizers (EQs), or cable simulators—may need to be adjusted to handle the new frequencies in use.

There are some considerations for implementing high split when looking at these three main CMTS/CCAP architectures. When embarking on an HFC network evolutionary step like high split, it is helpful to look out a couple potential steps ahead to leverage the work that needs to happen to better prepare the operator for a step beyond the high split. A DOCSIS 4.0 system, whether FDD or FDX, is generally predicated on implementing a DAA. For D4.0 FDD, a DAA is required to be in place because the HFC access technologies vendors are not currently planning on implementing analog / amplitude modulated optics that can support beyond 1,218 MHz. For D4.0 FDX, the spectrum that is shared between downstream and upstream can only really function well if the PHY is located near the edge of the network. Therefore, given that the operator will likely need to move to a DAA as the next evolutionary step to DOCSIS 4.0, it may make sense to leverage the planned node visit for doing the high-split upgrade to also implement a DAA at the same time. However, the great news is that moving to a DAA is definitely not required. Moving to a DAA requires additional coordination and work by the operator, so the operator may choose to implement high split with I-CCAP CAA first to get the benefits for some portions of the network and then subsequently upgrade to a DAA later in time. Products in the market today provide the flexibility to support multiple evolutionary steps that include a migration to high split!

One of the most common locations for RMDs and RPDs is in the optical fiber node. An RMD will contain network interfaces, such as Ethernet, and perform DOCSIS MAC and PHY functions, QAM video PHY functions, and RF functions including upstream analog signal receive and processing, all contained within a single module. As such, operators will want to ensure that legacy HFC node housings

currently deployed can be initially upgraded to a high split configuration and are later upgradeable to FDD-, FDX-, and/or FMA-compliant nodes via installation of RMD or RPD modules.

Production hardware capable of 1.8 GHz D4.0 FDD or 1.2 GHz FDX operation is not yet available as of the writing of this paper. However, many operators have either already deployed or are trialing D3.1 RMD and RPD solutions and planning for eventual production deployments. Therefore, initial high split deployments must be part of larger initiatives that lay the groundwork for higher frequencies and bandwidth in both downstream and upstream directions. This strategy ensures support for future migrations to DAA solutions that will lead to improved RF signal quality and greater efficiencies in speed, reliability, latency and security.

Operators may need to support more symmetrical gigabit upstream speeds on their I-CCAP infrastructure. The initial phase of a transition to a high-split architecture includes the RF amplifier upgrades and may continue to leverage traditional analog optics for downstream signal transport if the lasers are already 1.2 GHz-capable. For the upstream, digital return optics based on sampling and digitization of return analog TDMA and OFDMA carriers at the node transport those signals back to the I-CCAP at the hub or headend facility. An upgrade to 5-204 MHz modules is necessary, both for the node transmitters and the headend digital receivers. Legacy analog return optical links, typically capable of supporting expanded return bandwidths, can continue to be used but care must be taken to ensure their set operating points are adjusted for the expanded high-split bandwidth load to maintain desired signal-to-noise ratio (SNR) and MER performance targets.

As DAA architectures further mature, deployment of DAA (RMD and/or RPD) technologies can start in select areas as the next rollout phase supporting the migration to a high-split HFC network. At that point, new bi-directional digital Ethernet links will replace traditional analog optics and legacy digital optics as nodes are converted to DAA. Legacy nodes upgraded with RMDs or RPDs remain in place supporting DOCSIS 3.1 services over upgraded high-split coaxial networks, enabling operators to benefit from significant improvements in RF signal quality and ready to support a future transition to D4.0 FDD or FDX operation.

3.2. Spectral Considerations

When operators implement high split (204 MHz) in the upstream, operators need to clear out the spectrum from 54 MHz to 258 Mhz that had typically been used in the downstream and will now be repurposed. A simultaneous migration to 1,218 MHz can help by increasing the overall spectrum in use on the cable plant and enabling operators to shift the downstream spectrum up in frequency. However, this may not be enough to free up the desired spectrum. Therefore, the operator may need to implement additional strategies to free up spectrum, some of which include the following:

- Harvesting / removal of video channels that have a low subscriber watch rate
- Implementation of Switched Digital Video (SDV) for MPEG QAM video
- Migration of QAM video from MPEG-2 to MPEG-4 for greater compression
- Evolution from MPEG QAM video to IP video delivery.

When moving to high split, if an operator has legacy MPEG QAM set-top boxes (STBs) deployed, then one of the legacy STB out-of-band (OOB) is typically utilized, either SCTE 55-1 [SCTE 55-1] or SCTE 55-2 [SCTE 55-2]. While these protocols have served cable operators well in providing interactive video services to subscribers for years, the signals that are used in the downstream for these implementations typically fall between 85 MHz and 204 MHz and are thus now in the upstream direction once a cable

plant is migrated to high split. Two popular methods of dealing with this issue when evolving to high split include the following:

1. Migrate to STBs with embedded cable modems (CMs), potentially also compliant with the DOCSIS Set-Top Gateway (DSG) specification [DSG]. Note that some legacy STBs can be field upgraded from supporting legacy STB OOB to supporting DSG.
2. Migrate to all IP video delivery with IP STBs.

When employing one of the two strategies above, operators should take care that any STB upgrades that take place truly remove the need for any legacy STB OOB signaling as some STB implementations still required legacy STB OOB to either boot-up or for specific functions that required subsequent STB software upgrades to disable.

Cable operators may have some additional signals on the plant that need to be either disabled, relocated, or specially handled, especially when moving to a DAA. For example, cable operators in Europe who have previously operated frequency modulation (FM) radio application on the cable plant will likely need to shut this service down in order to support high split. When moving to a DAA, some of the plant alignment tones can be locally generated by the RPD or RMD. For other narrowband signals that may need to be preserved, the operator can implement Narrowband Digital Forward / Narrowband Digital Return (NDF/NDR) [R-OOB].

In summary, the cable operator needs to explore the entire spectrum in use today and ensure that all services are accounted for in some way when migrating to high split and/or a DAA. Products and features are available to help ensure that nearly all of the services can be available after the migration, but the operator may need to make some decisions on specific services to impact in order to get the greater benefit of improved upstream and downstream bandwidth.

3.3. Leakage Detection

Migration to high split and ultra-high split systems will require moving the upstream/downstream split further up into and potentially above the aeronautical band. Figure 8 below illustrates how the migration to a high-split architecture today—and migrations to even higher splits in the future—will incrementally turn legacy downstream spectrum currently within the aeronautical bands into new upstream bandwidth. This will require new approaches from operators to continue to monitor and repair signal leakage in compliance with existing United States Federal Communications Commission (FCC) and other international governing body rules. Some of the approaches being explored include:

1. Leveraging OFDMA Upstream Data Profile (OUDP) test bursts
2. Implementation of exclusion zones
3. Upgrading legacy leakage meters to detect downstream OFDM pilots.

With OUDP, upstream OFDMA signals will be generated by all cable modems that are provisioned for HS mode to allow new test equipment to detect egress from coaxial plant and accurately measure it for enterprise reporting. The concept involves cable modems that are instructed to generate OUDP test bursts in open timeslots and using specific CMTS configurations and OFDMA parameters that are considered optimum for accurately detecting leakage when driving either slowly or quickly past an RF leak. Specific parameters include subcarrier spacing, cyclic-prefix, roll-off-period, symbols-per-frame, data IUC modulation, pilot pattern, transmit burst gap between CMs, transmit duration and the frequency transmit location(s). Each modem will be instructed to repeat transmissions in a round-robin fashion, so RF Leakage detectors are not reliant on customer upstream data traffic. The primary benefit of using signals from cable modems is that their RF level is the same amplitude as the OFDMA signal which allows for a

much higher capture resolution and accuracy when a vehicle is in motion. A further exploration of this strategy can be found in a prior SCTE paper, “Leakage In A High Split World – Detecting and Measuring Upstream Leakage Levels in a One Gbps Symmetrical High Split Hybrid Fiber Coax Network” by John Chrostowski, et. al. [Chrostowski 2020].

Implementation of exclusion zones carved out of downstream OFDM carriers would allow the insertion of tag signals which would allow the continuous use of legacy leakage detection field meters.

Upgraded meters would allow detection of OFDM pilots within a large bandwidth range.

Interoperability testing to refine and evaluate the effectiveness of these proposed approaches, and their variations, is currently under way. Results today are encouraging, and options to support leakage monitoring obligations are within reach.




Figure 8 – DOCSIS 4.0 Frequency Split Roadmap vs. FCC-governed Aeronautical Bands

3.4. Customer Premise Equipment (CPE) Considerations

When operators decide to move to a high-split implementation, one of the first steps they need to take is to start pre-seeding the market with customer premise equipment (CPE) / CMs that have a selectable diplex filter that is remotely switchable to support the current split—low split (42 MHz or 65 MHz) or mid split (85 MHz)—in addition to high split (204 MHz). Having a large set of modems already in the field that are capable of high split will make it easy to take full advantage of the benefits of high split once the rest of the network elements have been evolved to make it a reality. Therefore, it is imperative for cable operators to start deploying cable modems that have the desired high-split diplexer capability early in time so that the market has enough critical mass of modems in use to justify and take advantage of the upgrade to high split.

After the operator has deployed enough high-split-capable CPE and enabled high split in the network, remotely forcing the CPE to use the new US and DS spectral ranges is critical. There are two main methods for this. One method involves configuring the MAC Domain in the I-CCAP, CCAP Core (for R-PHY), or RMD to tell the modems to start using the high split diplexer setting. This is then communicated down to the cable modems through the MAC Domain Descriptor (MDD) as type-length-value (TLV) 21 as defined in the D3.1 MAC and Upper Layer Protocols Interface Specification [DOCSIS 3.1 MULPI]. This implementation is relatively easy since it is down at the MAC domain level and not at the individual CM level. An alternate method is to set the diplexer setting via the CM config file, and this utilizes TLV 84 in the config file as defined in the same specification and which supersedes any TLV 21 setting in the

MDD [DOCSIS 3.1 MULPI]. While more difficult to manage down to the CM level, TLV 84 provides operators with more granular control of CMs that utilize high split and may become useful. With these two methods, operators have the ability to more broadly implement or to specifically pinpoint high split on the installed CPE as required for the given implementation.

When implementing high split, cable operators may need to address the legacy video CPE that are on the network due to the legacy STB OOB signals that would traditionally be in the downstream spectrum that is now utilized in the upstream. This topic has been explored in more detail in Section 3.2 and will also be explored in Section 5.2.2 but is listed in this section for completeness.

3.4.1. *Potential High-Split CPE Interference on Legacy CPE*

When implementing high split on CPEs, it is also possible that the upstream transmissions from the CMs may interfere with other legacy CPE, such as video STBs and legacy D2.0 and D3.0 CMs, that may not be high-split capable. In short, the transmissions in the upstream band from 54 MHz to 204 MHz may overlap with the downstream spectral window for the legacy CPE, and given that the high-split CM and legacy CPE are close to each other on the cable plant, the transmissions will be relatively high power and may cause issues.

There are two typical potential scenarios that have been considered relative to the potential interference of US transmit signals from high-split CPE on legacy video STBs and legacy D2.0 and D3.0 CMs: single home interference and neighbor-to-neighbor interference, as shown in Figure 9. Testing conducted by CableLabs, vendors, and multiple cable operators across a large set of legacy STBs has shown that STBs generally cannot handle an adjacent D3.1 transmit interference level known as the Carrier-to-Adjacent Carrier Interference Ratio (CACIR) where the interfering US transmission signal is 20 dB higher than the downstream video QAM receive level. Above this level of interference leads to MPEG video data corruption and resulting in an impacted video signal, including tiling and other potential artifacts. D2.0 or D3.0 legacy modems may also experience packet loss. The impacts are generally caused because the automatic gain control (AGC) functional blocks in the legacy CPE devices may be overdriven due to the relatively high-power US transmissions coming from the nearby D3.0 high-split CPE.




Figure 9 – Two Potential High-Split CPE Interference Scenarios

Luckily for both cable operators and end subscribers, interference that has a CACIR of 20 dB or higher is generally very uncommonly seen by operators in the field. Studies have shown that the neighbor-to-neighbor interference would require an extremely rare set of conditions for both neighboring premises to have. Typically, the scenarios that have impacts are with single home interference for a small set of

subscribers who have significant RF losses either in the drop to the premise or in the premise. For these scenarios, the DS signal level into the video STBs will already be quite low while the D3.1 modems running with high split enabled will be operating near the top of their RF output power range to overcome the losses within the in-home cable network. These subscribers have typically already been highlighted as needing improvements by typical cable plant monitoring and maintenance efforts. Additionally, operators can use the monitoring of STB receive power levels and D3.1 high-split CPE transmit levels to identify premises that might be at risk of experiencing such interference and then take steps to correct these scenarios.

Cable operators can mitigate cases of potential or actual interference through several potential solutions, including the following:

- Improving the in-premise cable network and/or drop to the premise to decrease the signal attenuation
- Reconfiguration of the in-premise cable network to better isolate the D3.1 high-split CPE by not placing them on the same RF splitter as legacy CPE that may be impacted by interference
- Replacing in-premise RF splitters with RF splitters that have higher isolation
- Installing band-stop RF filters to filter out the 54 MHz to 204 MHz spectral range for the legacy CPEs. Note that this solution should be used sparingly as the band-stop RF filters need to be very strategically placed within the subscriber premise, and the subscriber potentially moving such filters could cause issue with the high-split CPE deployments.

4. Operational Considerations for High Split

4.1. Upstream Level Considerations

4.1.1. Levels Out of Cable Modem (CM)

In DOCSIS 2.0 (D2.0), quadrature phase shift keying (QPSK) modulation remained as one of the constellations for the return path, and was assigned higher maximum power than the max power for the higher-level constellations, such as 16QAM and 64QAM, as shown in Table 2 from the DOCSIS 2.0 Radio Frequency Interface (RFI) specification [DOCSIS 2.0 RFI]. At first glance, this specification may appear as an error because the higher-level constellations do require a higher signal to noise and as such, would benefit from higher power transmission – provided, that the underlying noise is thermal / white Gaussian in nature.

Table 2 – Upstream Cable Modem Levels from D2.0 RFI Specification

Table 6-6 Constellation Gains and Power Limits

Constellation	Constellation Gain G_{const} Relative to 64QAM (dB)	P_{min} (dBmV)	P_{max} (dBmV) TDMA	P_{max} (dBmV) S-CDMA	$P_{min} - G_{const}$ (dBmV)	$P_{max} - G_{const}$ (dBmV) TDMA	$P_{max} - G_{const}$ (dBmV) S-CDMA
QPSK	-1.18	8	58	53	9.18	59.18	54.18
8QAM	-0.21	8	55	53	8.21	55.21	53.21
16QAM	-0.21	8	55	53	8.21	55.21	53.21
32QAM	0.00	8	54	53	8.00	54.00	53.00
64QAM	0.00	8	54	53	8.00	54.00	53.00
128QAM	0.05	8	N/A	53	7.95	N/A	52.95

But there was no error in Table 2, because QPSK, even though not very bit-per-second-per-hertz efficient, remained to deal with the presence of sizable ingress interference. Table 3 gives further clarification – QPSK was needed for those situations when carrier-to-interference plus ingress was in the mid-20 dB range, and as pointed to in Note 2 from the table, those time-varying ingress bursts could be as high as 10 dB below the signal. Therefore, there was a desire and need to keep QPSK as one of the options and to keep its max power high.

Table 3 – Assumed Upstream RF Channel Characteristics from D2.0 RFI

Table 4-2 Assumed Upstream RF Channel Transmission Characteristics (see note 1)

Parameter	Value
Frequency range	5 to 42 MHz edge to edge
Transit delay from the most distant CM to the nearest CM or CMTS	<= 0.800 msec (typically much less)
Carrier-to-interference plus ingress (the sum of noise, distortion, common-path distortion and cross-modulation and the sum of discrete and broadband ingress signals, impulse noise excluded) ratio	Not less than 25 dB (Note 2)
Carrier hum modulation	Not greater than -23 dBc (7.0%)
Burst noise	Not longer than 10 μ sec at a 1 kHz average rate for most cases (Notes 3 and 4)
Amplitude ripple 5-42 MHz:	0.5 dB/MHz
Group delay ripple 5-42 MHz:	200 ns/MHz
Micro-reflections – single echo	-10 dBc @ <= 0.5 μ sec -20 dBc @ <= 1.0 μ sec -30 dBc @ > 1.0 μ sec
Seasonal and diurnal reverse gain (loss) variation	Not greater than 14 dB min to max

2. Ingress avoidance or tolerance techniques may be used to ensure operation in the presence of time-varying discrete ingress signals that could be as high as 10 dBc. The ratios are guaranteed only within the digital carrier channels.

How high? As high as 58 dBmV per single QPSK channel of D2.0 as sourced from a cable modem. Then, the DOCSIS 3.0 Physical Layer Specification [DOCSIS 3.0 PHY] takes it a step higher, as shown in Table 4: a D3.0 cable modem must be able to generate a QPSK channel with a minimum of 61 dBmV if using a single channel and the same total power if distributed across 2 or 4 QPSK channels.

Table 4 – D3.0 per Channel Power Levels Out of Cable Modem, with Highlighted Max Points

Table 6-21 - Electrical Output from CM⁴⁶

Parameter	Value
Level range per channel (Multiple Transmit Channel mode disabled, or only Multiple Transmit Channel mode enabled with one channel in the TCS)	TDMA: P _{min} to +57 dBmV (32-QAM, 64-QAM) P _{min} to +58 dBmV (8-QAM, 16-QAM) P _{min} to +61 dBmV (QPSK) S-CDMA: P _{min} to +56 dBmV (all modulations)
Level range per channel (two channels in the TCS)	TDMA: P _{min} to +54 dBmV (32-QAM, 64-QAM) P _{min} to +55 dBmV (8-QAM, 16-QAM) P _{min} to +58 dBmV (QPSK)
Level range per channel (three or four channels in the TCS)	TDMA: P _{min} to +51 dBmV (32-QAM, 64-QAM) P _{min} to +52 dBmV (8-QAM, 16-QAM) P _{min} to +55 dBmV (QPSK)

With the introduction of D3.1 [DOCSIS 3.1 PHY], no per-channel max power increases take place; however, the total composite power is nudged up, from 58 dBmV for D2.0, and 61 dBmV for D3.0, to 65 dBmV, as shown in Figure 10. D3.1 also increased the maximum US spectrum to 204 MHz. The latest

DOCSIS spec, D4.0, stays with the same 65 dBmV total composite power (TCP) limit, provided the whole D4.0 upstream spectrum, up to 684 MHz, is utilized [DOCSIS 4.0 PHY].

7.4.12.2 Transmit Power Requirements

The transmit power requirements are a function of the number and occupied bandwidth of the OFDMA and legacy channels in the TCS. The minimum highest value of the total power output of the CM P_{\max} is 65 dBmV, although higher values are allowed. The total maximum power is distributed among the channels in the TCS, based on equal power spectral density (PSD) when the OFDMA and legacy channels are fully granted to the CM. Channels can then be reduced in power from their max power that was possible based on equal PSD allocated (with limits on the reduction). This ensures that each channel can be set to a power range (within the DRW) between its maximum power, $P_{1.6\text{hi}}$, and minimum power, $P_{1.6\text{low}}$, and that any possible transmit grant combination can be accommodated without exceeding the transmit power capability of the CM.

Maximum equivalent channel power ($P_{1.6\text{hi}}$) is calculated as $P_{1.6\text{hi}} = P_{\max} \text{ dBmV} - 10\log_{10}(N_{\text{eq}})$.

For a CM operating with a DOCSIS 3.1 CMTS, even on a SC-QAM channel, the CMTS MUST limit the commanded $P_{1.6\text{hi}}$ to no more than 53.2 dBmV+ ($P_{\max} - 65$) if the bandwidth of the modulated spectrum is ≤ 24 MHz. This enforces a maximum power spectral density of P_{\max} dBmV per 24 MHz. This limit on power spectral density does not apply for a CM operating with a DOCSIS 3.0 CMTS, where the fidelity requirements are the DOCSIS 3.0 fidelity requirements and not the DOCSIS 3.1 fidelity requirements.

SC-QAM channels that are 6.4 MHz in BW have a power of $P_{1.6\text{r}_n} + 6$ dB.

Figure 10 – D3.1 PHY Spec Snippets, Requiring TCP of 65 dBmV Out of the CM

Figure 11 shows how the above TCP-focused specs translate to per-channel levels with the upper left-most green line showing D3.1 high-split levels and the other lines/colors showing various levels of D4.0 FDD ultra-high splits. Please note that the vertical y-axis shows power per 1.6 MHz-wide channel, the method used in the PHY specs. For power per 6.4 MHz-wide channel, one should add 6 dB to all the per-channel powers.




Figure 11 – Max HS CM Levels (Green) with Various Ultra-High-Split Options of D4.0 FDD

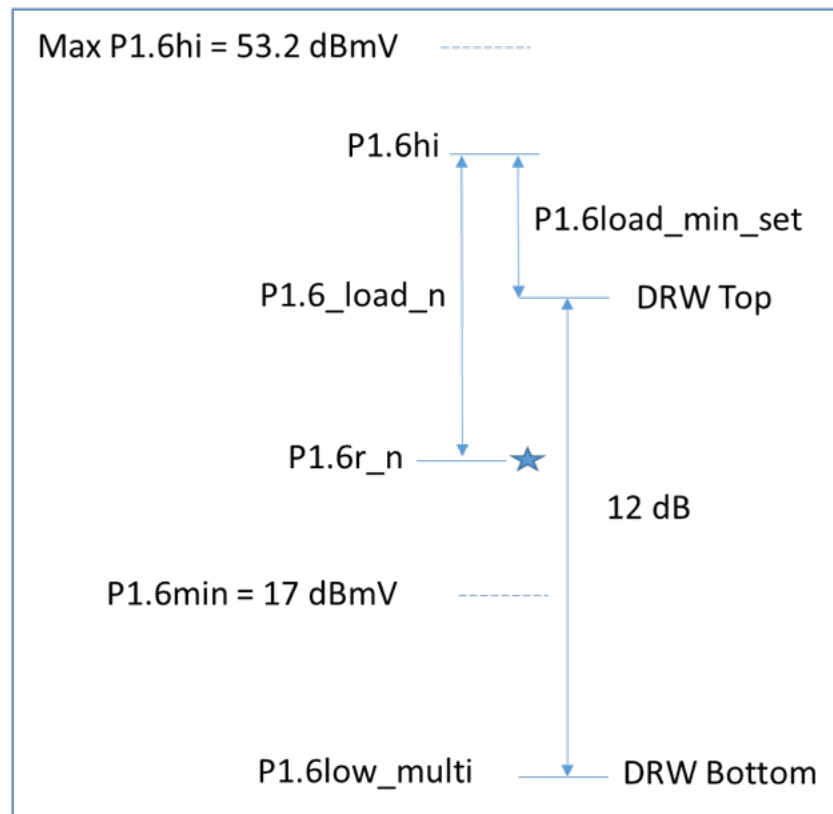


Figure 3.2.7-4. Transmit power at the CM

Figure 12 – Dynamic Range Window (DRW) Considerations from D3.1 PHY

Thus, in comparison to US levels specified in previous DOCIS versions, both TCP and per-channel D3.1 levels are high enough – which gives plenty of room to overcome whatever spurious noises the upstream plant may present. In comparison with the D4.0 levels, the maximum D3.1 high-split levels are also higher—about 6 dB higher and range from ~48 to ~52 dBmV per 6.4 MHz-wide channel, as can be seen in Figure 11. Operators, however, should exhibit caution in approaching the maximum levels, for several reasons as outlined below.

The first reason is that cable modem termination systems (CMTSs) “normally administers dynamic range window (DRW) of 12 dB.” Exact dynamic range window (DRW) details per D3.1 PHY spec are shown in Figure 12. One way to think of this aspect is that HS CM shall not be commanded a TCP higher than 59 dBmV, which corresponds to ~44 dBmV/6.4 MHz channel. If reached, this TCP range is often denoted as a “red zone” by the operators.

The second reason is outlined in Section 4.1.2 below.

4.1.2. Levels into RF Amplifier Upstream Ports

The second reason operators should exhibit caution in approaching the maximum modem transmit levels when operating high split is that the upstream power into the return RF amplifiers should fall into a “Goldilocks” (also known as “just right”) range recommended by the amplifier manufacturers: not too low so as to get affected by thermal noise and the amplifier’s noise figure, and not too high so as to get distorted in the upstream gain stage. Per CommScope’s MB120 amplifier data sheet [MB120 data sheet],

as an example, the high-split configuration upstream distortion specs are shown for 33 upstream channels at 5 dBmV per 6.4 MHz at the amplifier's upstream input port.

Figure 13 shows an in-between RF amplifiers section of HFC plant, as a backdrop for various levels discussed. For completeness, the tap values were selected in order to set the customer-premise downstream levels to fit within -6 to +8 dBmV per 6 MHz-wide downstream channel range, as shown in Figure 14.

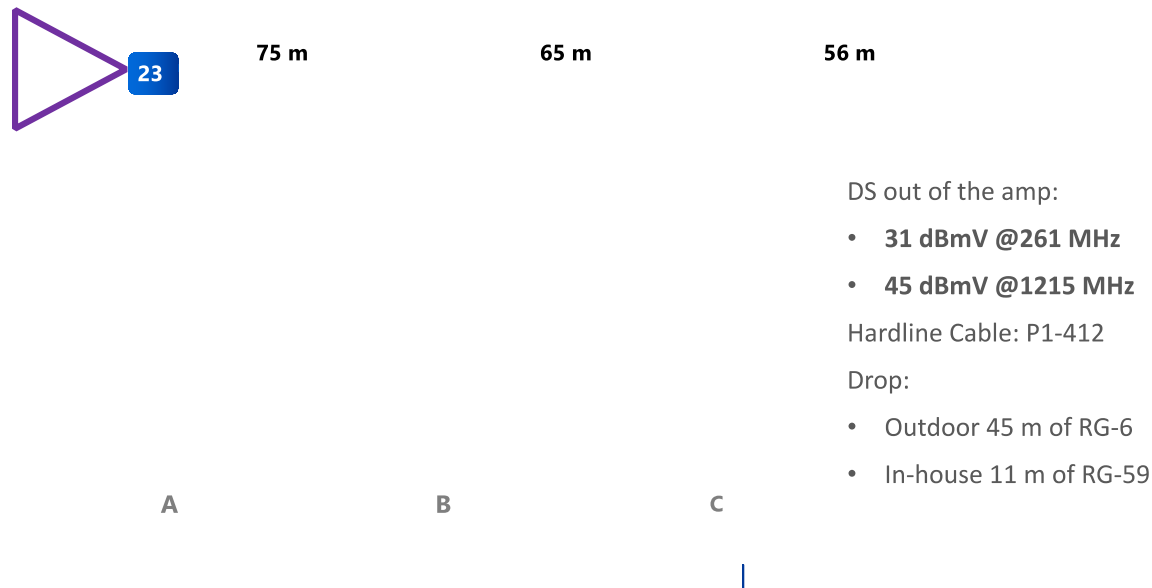


Figure 13 – An Example Section of HFC Plant in Between Two RF Amplifiers




Figure 14 – Actual Per-channel DS Levels, at an RF Amplifier Output and at Customer Premises




Figure 15 – CM Upstream Levels, and the Resulting Input-into-RF Amp Levels

The upstream levels into the RF amplifier, driven by the CM's 31/35, 27/35, and 24/34 dBmV at lowest/highest upstream channel at premises A, B, and C, respectively, are shown in Figure 15. The “smile” shape in the 3 “@Amp” curves is due to the cable loss signature, driven by the selected linearly-up-sloped CM output. The CM TCP comes to 48, 46, and 44 dBmV at premises A, B, and C, respectively.

A more exacting method of DS/US alignment can be carried if “conditioned taps” are used in place of the regular ones. In Figure 16, the first tap, with the original 23 dB value, is replaced by a 17 dB tap with an internal 6 dB cable simulator plugin. Similarly, the third tap with the original 11 dB value, is replaced by a 4 dB tap with an internal 6 dB cable equalizer plugin. As a result, the downstream levels at the conditioned A and C drops will match those of B; and similarly, as shown in Figure 16, the same CM levels out of A, B and C cable modems will produce an identical level into the upstream RF amplifier input port.




Figure 16 – Conditioned Taps Example for an Even Further Optimized Plant Section

The effort required to perform this level of balancing does not appear to produce enough improvement in the high-split plant to pay for itself, and is hereby mentioned as an option, although rarely employed operationally. That may, however, change with introduction of D4.0 and ultra-high splits.

4.2. Field Components Upgrade Considerations – “Rip and Replace” or “Modules Only?”

At minimum, high-split upgrades affect the fiber node and RF amplifiers. In some cases, the RF taps and passives may get a refresh as well. One of the key decisions to consider when upgrading nodes and amps is whether to “rip and replace” the whole housings or to just upgrade the modules inside.

The modules-only approach has several advantages, such as:

- Faster to complete – which leads to less down time
- Crews are more efficient – get more nodes done in a shift
- Fewer opportunities for the things to go wrong (such as broken fiber pigtails, cracked coaxial cable, bent connectors) – thus avoiding “opening a can of worms”
- Lower cost of labor encountered – by eliminating re-connectorization of coaxial cable and/or fiber splicing.

There are, however, several reasons where it makes sense to “rip and replace”:

- If the new device offers future-proof options not available in the old one
- Is the old housing in acceptable shape? If not, replacement is the only option
- Is the old device type an unsupported product? If yes, cut out and use the “standard” product.
- If standardization driven – even if the old device is a supported product, it may get replaced because of standardization.

Standardization does matter as it improves inventory management and thus reduces overall inventory needs. Fewer spares are carried on a technicians’ trucks, and tech training is simplified, which further minimizes opportunities for human error.

Figure 17 and Figure 18 show details of the node replacement process, in support of and as an illustration of how involved it may get, especially if executing the “rip-and-replace” strategy.

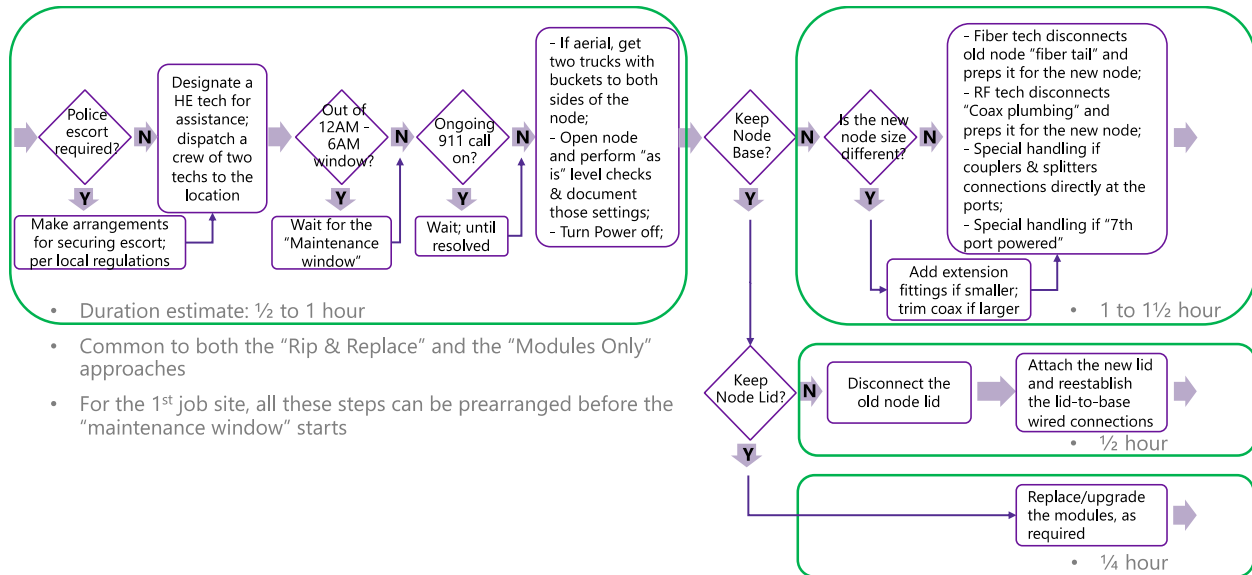


Figure 17 – How to Replace a Node – Initiation and Replumbing Part of the Process

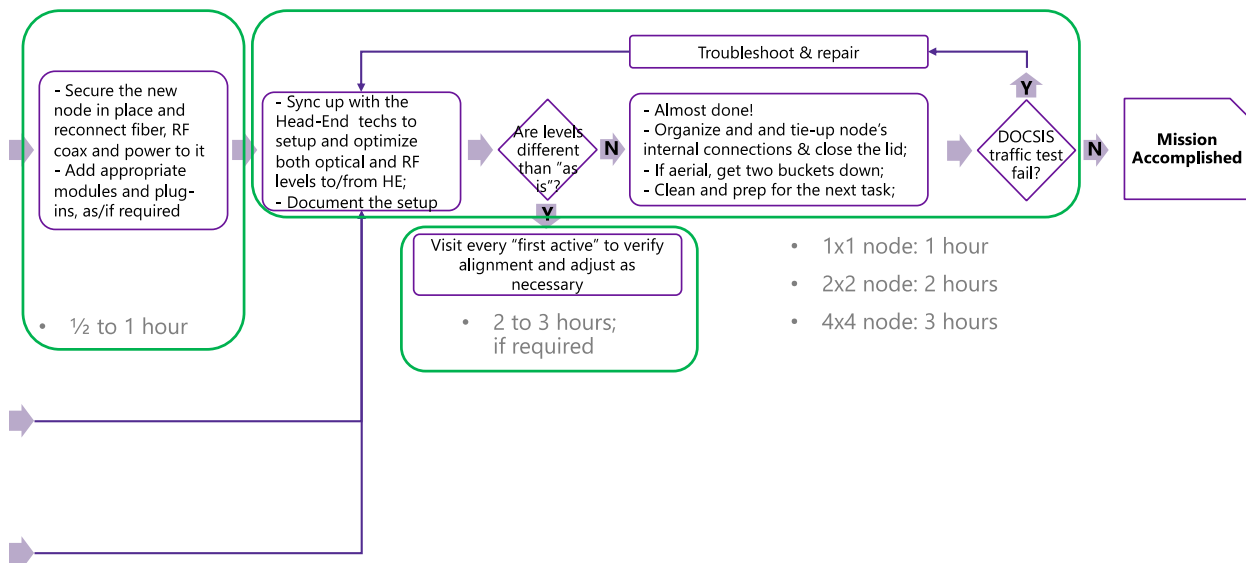


Figure 18 – Setup & Wrap-up Part of the Node Replacement Process

5. Shaw Communications – High-Split Lessons Learned

Shaw Communications has embarked on deploying both mid split (85 MHz) and high split (204 MHz) along with 1,218 MHz downstream in their HFC network. This section describes thoughts behind the upgrade, key considerations, potential future migrations, and lessons learned.

5.1. Why the Need to Modernize the Network Beyond Mid Split?

Mid-split deployments are becoming more and more common for great reasons. With the sudden rise in upstream traffic during the COVID-19 pandemic, mid-split deployment plans for most operators have increased in priority and scale. Mid-split upgrades provide a rapid, cost-effective method for immediate congestion relief. However, for most operators the business case for mid-split upgrades has been primarily focused on congestion mitigation while the benefits of broadband tier enablement are usually discounted. As a result, the broadband speeds offered in the marketplace are typically defined solely by an operator's congestion management strategy. Competitive FTTx deployments will force many operators to shift this thought process.

In many markets around the world, cable operators are facing more competition from well-funded players (typically quadruple-play telco operators) who are aggressively deploying FTTx networks at scale. Even operators with significant mid-split deployments (such as Shaw) face substantial market pressures from the aggressive advertising on the purported advantages of a pure fiber network, along with very aggressive speeds tiers in market. In the Canadian market, Shaw is already seeing competitive 2.5 Gbps symmetrical tiers become broadly available as of the writing of this paper.

While broadband tiers enabled by a mid-split network have been very successful in competing against GPON offerings, these tiers will not be as competitive once FTTx deployments shift to 10 Gbps-capable PON technologies, such as 10G EPON and XGS-PON. As a result, it is paramount that cable operators think beyond mid split and look towards future technologies such as high split and D4.0 to ensure they remain relevant in a very competitive marketplace.

5.2. Key High-Split Readiness Activities (with a Mind Towards DOCSIS 4.0)

Operators who have D4.0 FDD deployment ambitions can utilize high-split deployments as an opportunity to set themselves up for success by dealing with some of the foundational readiness activities early. The following is a summary of key activities an operator can undertake within their high-split programs to assist a future D4.0 transition.

5.2.1. Network Architecture Readiness

As noted, Section 3.1, there are multiple architecture options that can enable high-split deployments. If an operator has D4.0 deployment ambitions, starting down the DAA path is critical. It is recommended that operators include the following architectural considerations when defining their high split upgrade program:

- DAA to implement: R-PHY or R-MACPHY
- Node fiber availability to support analog-to-digital optics conversion
- Proactive cascade reductions
- HFC power supply upgrades
- Elimination of 3x3 and 4x4 (DS-SG x US-SG) node configurations (to ensure compatibility with DAA upgrades given the capabilities of the RPDs and RMDs in the market today).

5.2.2. Set-Top Boxes with Legacy Out-of-Band Signaling

Most operators have set-top boxes deployed in their networks that require SCTE 55-1 or 55-2 out-of-band (OOB) signals to operate. As noted previously in this paper maintaining these signals requires some extra efforts in a high-split network.

Shaw elected to fully reclaim these legacy set top boxes and only support DOCSIS Set-Top Gateway (DSG) or IPTV video set-tops in Shaw's high split network deployments. Note that DSG STBs utilize an embedded cable modem and thus preclude the need for the legacy SCTE 55-1 or 55-2 STB out-of-band signaling. This approach operationally simplifies the deployments while ensuring Shaw is effectively managing the lifecycle of legacy equipment in the network.

It should be noted that should an operator elect to leverage DSG set-top boxes with high split, it is critical that they test this equipment to ensure the capability of operating in a mode where the legacy OOB signals are not present. Some of the DSG STB equipment requires the legacy OOB to come online to function correctly. In this situation, these DSG STBs are not compatible with a high split deployment and will require software updates to eliminate the dependency on the OOB carrier. The potential issue is solvable by working with the STB vendor, but it is good to plan ahead so the operator is not caught off-guard waiting for a vendor STB software update that was unexpected.

5.2.3. On-Premises Architecture

When looking at the on-premises network architecture, operators need to consider the following when defining their go-to-market strategy for implementing high split:

5.2.3.1. On-Premises Interference Risk

Many operators deploy multiple coaxial cable-based customer premises equipment (CPE) in the same premises network to support high-speed data, voice, and video services. There is a risk when high-split and sub/mid-split CPE deployed in the same in-home network that a high transmit level on high split CPE can overload the tuner of the sub/mid split CPE in the same home. If this were to occur services supported by the sub/mid split CPE could be intermittent.

5.2.3.2. Wi-Fi & Ethernet Distribution

With the advancement of next generation Wi-Fi technologies (Wi-Fi 6, 6E) and the proliferation of structured Ethernet wiring (Cat 5E/6), the home IP network is becoming more capable and easier to use than ever. In addition, the market scale of these technologies will drive future capabilities and cost optimizations that are not easily replicated on the coaxial networks in the premise.

5.2.3.3. Future Access Network Considerations

To minimize future re-work, it is important to align HFC install practices with future FTTx, and D4.0 considerations. In these use cases, coaxial cable-based in-home networks are not ideal, and customers are best supported by Wi-Fi and wired Ethernet systems.

Operators should strongly consider moving to a point-of-entry style modem install for their high-split customer activations. This type of installation eliminates the use of on-premises coaxial networking and leverages a converged services gateway (broadband, voice, & video) with Wi-Fi and Ethernet for on-premises signal distribution. This will provide the operator the following benefits:

- Leverages future-proof in-home networking technology (Wi-Fi and Ethernet)
- Simplifies on premises networks and reduces operational support costs
- Improved customer experience
- Supports easy future FTTH and D4.0 upgrades.

5.2.3.4. Spectrum Readiness

The upgrade of existing low-split or mid-split plant to high split does require the reallocation of downstream spectrum to use in the upstream. Shaw had to identify 25 downstream carriers to be reclaimed in its mid-split nodes to enable a seamless transition to high split. This reclaim was enabled primarily through optimizations to the legacy QAM video package. These optimizations included the following:

- MPEG-2 to MPEG-4 conversions (high definition (HD) and standard definition (SD))
- Legacy VOD QAM carrier reclaims
- Improved closed-loop stat multiplexing of linear video services.

To enable the radio frequency (RF) spectrum for high-split deployments, Shaw developed the capability to build a distinct channel lineup that is used for high-split nodes deployments. This distinct channel lineup is delivered by the DAA node via an auxiliary video core. This configuration allows an operator to target any spectrum reclamation activities to the node level and minimize the operational impact of large-scale video STB swaps while eliminating the cost and complexity of managing multiple channel lineups through analog head ends and hub sites.

5.2.3.5. Modem Pre-Seeding

Deployment of modems with capabilities to support future plant configurations is a great way for operators to set themselves up for future success. Shaw was able to leverage a very high penetration of mid-split-capable modems to provide immediate congestion relief and large-scale broadband tier upgrades as nodes were upgraded to mid split.

When looking at high-split upgrades, operators should include the pre-seeding of high split-capable CPE into their overall network upgrade programs. With high-split-capable CPE available today, this is a “no regrets” investment!

High-split modem pre-seeding benefits include the following:

- The full-band capture capability of the modem can be used to validate performance of newly activated downstream spectrum (1-1.2 GHz)
- Investments in high-split CPE are fully leveraged in D4.0 FDD (and also D4.0 FDX) deployments
- Immediate traffic offload to newly activated upstream and downstream spectrum once high-split upgrades have been completed
- Future broadband tier upgrades can be done without a modem swap.

5.2.3.6. Back-office Software Readiness

Back-office software (BSS) updates are often overlooked and undervalued when operators consider network upgrades such as high split. However, the definition of node-level serviceability flags is a key enabler to a successful upgrade. These types of flags tell the billing system what the service capabilities of

the node are and what equipment variants are to be allowed on the node. The following are examples of node service availability flags that an operator could implement in support of high-split investigations:

Table 5 – Node Serviceability Flags for Back-office Software in Support of High Split

Flag	Function
Legacy STB Reclaim Underway (Hybrid)	Informs the customer service representative that a legacy STB reclaim is underway and warns them that activation of new legacy OOB video CPE should not be done on the customers node.
Legacy STB Reclaim Underway (All IP)	Informs the customer service representative that a legacy STB reclaim is underway and warns them that activation of new legacy QAM video CPE should not be done on the customers node.
Legacy STB Reclaim Completed (Hybrid)	Informs the customer service representative that the customer's node has completed all required legacy OOB video CPE reclaims and implements a provisioning rule to deny all activations of non-DSG and non-IPTV hardware.
Legacy STB Reclaim (All IPTV)	Informs the customer service representative that the customer's node has completed all required legacy video CPE reclaims and implements a provisioning rule to deny all activations of non-IPTV hardware.
High Split Completed (Hybrid)	<p>Informs the customer service representative that the customer's node has completed a high-split upgrade and informs them of the new high-split-enabled broadband service catalog.</p> <p>In addition, informs the access network configuration systems of which RF channel lineup to apply to the DAA node (i.e., high split legacy QAM video lineup + high split DOCSIS configuration).</p>
High Split Completed (All IP)	<p>Informs the customer service representative that the customer's node has completed a high-split upgrade and informs them of the new high-split enabled broadband service catalog.</p> <p>In addition, informs the access network configuration systems of which RF channel lineup to apply to the DAA node (i.e., full spectrum high split DOCSIS configuration).</p>

5.3. A Potential Evolution Path to DOCSIS 4.0

For a cable operator to be successful in a market where there is an aggressive FTTP builder, they will need to have the ability to rapidly deploy competitive speeds to 10G PON technologies, while at the same time increasing focus on service reliability, service experience, and customer service. D4.0 FDD technology gives operators the necessary tools to compete in this type of marketplace. D4.0 FDX technology also gives operators the necessary tools to compete and may be preferred by some operators. Focusing in on D4.0 FDD, it provides operators the following benefits:

- Highly economical upgrade
- Fastest path to compete with 10 Gbps-capable PON technologies
- Enough capacity to reduce focus on speed (i.e., take speed "out of the conversation")

- Improved network reliability
- Ability to leverage previous high-split investments.

Shaw has developed their network upgrade strategy with a view on the most efficient path to D4.0 FDD. In this context, high split is an important and logical steppingstone in setting the stage for future D4.0 FDD deployments.

The following is a potential four step process that operators who are currently deploying mid split could implement to enable a seamless transition through high split and to D4.0.

Step 1: D3.1 High Split (204 MHz / 1 GHz)

In this phase as shown in Figure 19, the operator begins deploying high split using an architecture that can evolve gracefully to D4.0 FDD deployments. This configuration will allow operators to compete in highly competitive markets with gigabit symmetric services while setting the stage for future upgrades.

The following elements are included in this initial step:

- Reclaim of all legacy OOB STBs
- Upgrade any taps and passives with less than 1 GHz capabilities to 1.2 GHz while keeping 1 GHz-capable taps in place
- Resolution of any plant architecture issues (e.g., plant powering upgrades, elimination of 3x3 (DS-SG x US-SG) and 4x4 node configurations, proactive cascade reductions)
- DAA node deployment (Remote PHY or Remote MACPHY) coincident with an upgrade to 204 MHz / 1.2 GHz
- Reclaim legacy QAM video channels to support the high split conversion.



Figure 19 – Step 1: D3.1 High Split (204 MHz / 1 GHz)

Step 2: D3.1 High Split with 1.8 GHz Taps

In this phase as shown in Figure 20, the operator will include a full upgrade of taps and passives into the high-split upgrade efforts. In this phase a sweep of the plant will be completed to upgrade all taps and passives housings to be 3 GHz-capable, and 1.8 GHz-capable faceplates.

Adding a full tap and passive swap enables the operator to achieve the following benefits:

- Minimizes future rework to enable full D4.0 FDD capacities
- Allows the operator to gain experience and optimize full scale tap and passive swaps
- Enables additional plant hardening by allowing the inspection and resolution of craft issues that can cause ingress and RF performance issues
- Improves plant stability issues that are caused by seizure screw related issues

- Improves DAA node uptime by ensuring power bypasses are included in all taps and passives
- Enables 1.2 GHz plant operation and allows the deployment of an additional 200 MHz of downstream capacity.



Figure 20 – Step 2: D3.1 High Split with 1.8 GHz Taps (in 3.0 GHz Housings)

Step 3: D3.1 High Split with 1.8 GHz Taps and 1.8 GHz Amplifiers

In this phase as shown in Figure 21, the operator will include a full upgrade of the HFC amplifiers to be 1.8 GHz-capable equipment (compliant with [SCTE 279 2022]).

Adding a full amplifier upgrade enables the operator to achieve the following benefits:

- Gaining experience with 1.8 GHz amplifiers in production environments
- Improved compliance of amplifier level setup by leveraging auto-setup capabilities of 1.8 GHz amplifiers
- Minimizing operational complexity by using soft-selectable pads and EQs
- Minimizes future rework to enable full D4.0 FDD capacities.




Figure 21 – Step 3: D3.1 High Split with 1.8 GHz Taps (in 3.0 GHz Housings) and 1.8 GHz Amplifiers

Step 4: DOCSIS 4.0 High Split

In this final phase as shown in Figure 22, the operator will include the deployment of a D4.0-capable node and shift to a RMD to enable the full D4.0 FDD capacities.

Adding a D4.0 RMD enables the operator to achieve the following key benefits:

- Enables full D4.0 FDD capacities
- Improves network resiliency by fully separating the data and control planes on the DAA node
- Reduction of technical facility power and cooling requirements.



Figure 22 – Step 4: DOCSIS 4.0 High Split

5.4. Summary of Lessons Learned

As part of Shaw’s high-split deployments and Shaw’s overall D4.0 readiness activities, the following are some key lessons learned:

1. The Frequency Modulation (FM) and Very High Frequency (VHF) Bands are unproven for upstream usage

These portions of the spectrum have historically been used for downstream signals; their use for upstream has not been proven out. Over the air broadcasts can significantly impact the quality of these channels. It is important that operators have a plan on how to validate the impacts of ingress and how much DOCSIS capacity this spectrum can provide.

2. The performance of DAA deployments is amazing!

DAA deployment will be required for DOCSIS 4.0; however, there are significant performance benefits that operators can achieve by including DAA deployments in their high-split programs. Shaw has seen an average downstream modulation error ratio (MER) improvement of up to 8.2 dB and an upstream MER improvement of 5.2 dB.

3. Pre-seeding modems is an effective way to prepare for the future

Having modems pre-deployed in the network that are immediately able to utilize the new spectrum created by deploying high split is a key element to the high-split business case. These modems enable immediate congestion mitigation, tier upgrades without modem swaps, as well as enabling the operator to use the full-band capture capability to test and validate the new spectrum before customer use.

4. Avoid “big bang” changes

A successful D4.0 deployment involves many changes that can be highly impactful to operational teams. It is important to layer these changes in over time to allow for effective change management. High split represents an opportunity for operators to begin their D4.0 journey today and prove out and harden technology changes one at a time.

5. Plan for measuring and minimizing customer impact when deploying high-split change events

As has been highlighted by the COVID-19 pandemic, the reliability of the home’s broadband connection has never been so important. Due to the sheer number of changes required to

deploy high split and ultimately D4.0, it is critical that the operator has a plan on how to measure and optimize the deployments to ensure customer impacts are minimized.

6. Conclusion

Subscribers' growing bandwidth demand and increased competitive pressures, especially from internet service providers deploying 10G PON technologies, are driving cable operators to expand upstream (and downstream) capacity. In addition to these market factors, operators may upgrade plant actives (RF amplifiers) due to reaching the end of their typical 10-20-year lifetime—an opportune time to make the change to 1,218/204 MHz actives.

High-split (204 MHz) operation in the upstream—especially when paired with an upgrade to 1,218 MHz in the downstream—can provide quite the improvement in the HFC network service group capacity of up to 8.6 Gbps capacity in the downstream and 1.5 Gbps capacity in the upstream. Some projections show high split and 1,218 MHz providing enough bandwidth for a service group of at least 150 subscribers with 5 Gbps down x 1 Gbps up speed tiers / SLAs for the next 10 years.

Moving to high split is not only a steppingstone on the path to 10G and to DOCSIS 4.0, whether moving to FDD or FDX, but a huge step in terms of benefits to both operators and end subscribers. Luckily for cable operators, technologies and products in the market—deployable today—enable operators to confidently move forward with high split in the three leading CMTS/CCAP architectures: I-CCAP, R-PHY, and/or R-MACPHY. A node visit, required for the high-split upgrade, is also the opportune time to upgrade to a DAA, whether R-PHY or R-MACPHY. However, a simpler and faster high-split upgrade with I-CCAP, is still a strong option to consider, especially when the bandwidth and market drivers demand an earlier move to high split—before the operator can migrate to a DAA. High split on I-CCAP can have even better performance when utilizing digital return in place of analog return on the fiber optics link.

In planning to implement high split, several potential architectural considerations must be taken into account, such as strategies on how to free up the high-split spectrum, like implementing SDV, reducing the video channel line-up, migrating from MPEG-2 to MPEG-4, or moving to IPTV. Tools and solutions exist that support leakage detection, which in the high-split world now must detect signals in the reverse path. Operators must develop a CPE strategy that aligns with high split, including pre-seeding the market with high-split-capable CPE, migrating legacy video STBs to STBs with embedded cable modems (DSG-capable or otherwise), and dealing with potential interference from high-split-capable CPE with other legacy CPE in the home or in close neighbors. Operators need to consider the cable modem RF levels and the inputs to the RF amplifiers when making the move to high split. Operators also need to consider whether to take a “rip-and-replace” or a “module-only” upgrade path for the RF components, with each one having benefits and drawbacks.

Shaw Communications shared some considerations and lessons learned from embarking on a high-split upgrade. Shaw explained the Canadian market dynamics that drove them to deploy mid split and then high split to compete against FTTx service providers. Shaw explored some of the key architectural considerations for moving to high split, including network architecture readiness, strategies for handling legacy STB out-of-band signals, preparing for the spectral shift, pre-seeding high-split-capable modems, and back-office software readiness, among others. Shaw provided a step-by-step potential evolutionary path to high split and then to DOCSIS 4.0 FDD-capable DAA devices in preparation for the next step of enabling D4.0 FDD. Finally, Shaw shared some of their key lessons learned, including the unproven nature of the FM and VHF bands for upstream usage, the amazing performance of DAA (average MER improvements of 8.2 dB down and 5.2 dB up), the importance of pre-seeding modems with high split

support, avoiding making major changes at the same time, and the importance of minimizing subscriber impacts during the upgrade through proper monitoring.

In short, the time for talk is over, and there is no better time like the present to embark on the journey to implement high split as a key step on the path to 10G!

Abbreviations

10G EPON	10 Gbps Ethernet passive optical networking
10G-PON	10 Gbps passive optical networking (also known as XG-PON)
ACK	acknowledgement (typically from Transmission Control Protocol)
AGC	automatic gain control
BSS	business support system
BW	Bandwidth
CAA	centralized access architecture
CACIR	carrier to adjacent carrier interference ratio
CAGR	compounded annual growth rate
CAPEX	capital expense
CCAP	Converged Cable Access Platform
CDF	cumulative distribution function
CES	Consumer Electronics Show
CM	cable modem
CMTS	Cable Modem Termination System
COTS	common off-the-shelf
CPE	customer premises equipment
D2.0	Data Over Cable Service Interface Specification 2.0
D3.0	Data Over Cable Service Interface Specification 3.0
D3.1	Data Over Cable Service Interface Specification 3.1
D4.0	Data Over Cable Service Interface Specification 4.0
DAA	Distributed Access Architecture
dB	Decibel
dBmV	decibel-millivolts (decibels relative to 1 millivolt)
DCA	Distributed Cable Architecture
DOCSIS	Data Over Cable Service Interface Specification
DRW	dynamic range window
DS	Downstream
DSG	DOCSIS Set-top Gateway
DSL	digital subscriber line
DTA	digital terminal adapter
EOL	end of line
EPON	Ethernet Passive Optical Network (aka GE-PON)
EQ	Equalizer
ESD	extended spectrum DOCSIS
FCC	Federal Communications Commission
FDD	frequency division duplex
FDX	full duplex DOCSIS
FM	frequency modulation
FMA	Flexible MAC Architecture
FTTB	fiber to the building
FTTH	fiber to the home
FTTP	fiber to the premise

FTTx	fiber to the 'x' where 'x' can be any of multiple option for subscriber locations, including home, premise, building
FWA	fixed wireless access
Gbps	gigabit per second
GHz	gigaHertz
gNMI	Google Network Management Interface
HD	high definition
HEO	head end optics
HFC	hybrid fiber-coax
HP	homes passed
HS	high split (204 MHz)
HSD	high-speed data
HW	hardware
I-CCAP	Integrated Converged Cable Access Platform
IEEE	Institute of Electrical and Electronics Engineers
IP	internet protocol
ITU	International Telecommunication Union
ITU-T	International Telecommunication Union Telecommunications Specification Sector
kbits	kilobits per second
LDPC	low density parity check (FEC code)
LE	line extender (amplifier)
LS	low split (42 MHz or 65 MHz, depending on region)
MAC	media access control
MB	multi-port bridger (amplifier)
Mbps	megabit per second
MDD	MAC Domain Descriptor
MDU	multiple dwelling unit
MER	modulation error ratio
MHA	Modular Headend Architecture
MHz	megaHertz
MPEG	Moving Picture Experts Group
MPTS	multi-program transport stream
MS	mid split (85 MHz)
MSO	multiple system operator
MULPI	MAC and upper layer protocols interface
N+0	node+0 actives
NCTA	The Internet & Television Association (formerly National Cable & Telecommunications Association)
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OPEX	operating expense
OSP	outside plant
OSS	operational support system
PHY	physical interface/layer
PON	passive optical network
PSD	power spectral density
PTP	(IEEE 1588) Precision Timing Protocol

QAM	quadrature amplitude modulation
QoE	quality of experience
QPSK	quadrature phase shift keying
R-MACPHY	remote MACPHY
R-PHY	remote PHY
RF	radio frequency
RFI	radio frequency interface
RMD	remote MACPHY device
RPD	remote PHY device
SCTE	Society of Cable Telecommunications Engineers
SDV	switched digital video
SG	service group
SLA	service-level agreement (also known as speed tier)
SNR	signal-to-noise ratio
SS	sub split (42 MHz or 65 MHz, depending on region)
STB	set top box
Tavg	average bandwidth per subscriber
TCP	total composite power
TCP	Transmission Control Protocol
Tmax	maximum sustained traffic rate – DOCSIS Service Flow parameter
TX	Transmit
UHS	ultra-high splits (300 MHz, 396 MHz, 492 MHz, or 684 MHz)
US	upstream
VOD	video on demand
VoIP	voice over internet protocol
VHF	very high frequency
XG-PON	10 Gbps passive optical networking (also known as 10G-PON)
XGS-PON	10 Gbps symmetrical passive optical networking

Bibliography & References

[Maricevic_2022] “Network Migration to 1.8 GHz – Operational “Spectral Analysis” Measured in nano-Hertz, a 30-Year Perspective”, Dr. Zoran Maricevic, John Ulm, Craig Coogan, 2022

[NCTA 10G] <https://www.ncta.com/media/media-room/introducing-10g>, NCTA, CableLabs[®], & Cable Europe

[DOCSIS 2.0 RFI] “DOCSIS 2.0 Radio Frequency Interface Specification”, CM-SP-RFIv2.0-C02-090422, CableLabs, 2009

[DOCSIS 3.0 PHY] “DOCSIS 3.0 Physical Layer Specification”, CM-SP-PHYv3.0-C01-171207, CableLabs, 2017

[ULM_2022] “Broadband Capacity Growth Models – Will the end of Exponential Growth eliminate the need for DOCSIS 4.0?”, John Ulm, Dr. Zoran Maricevic, Ram Ranganathan, SCTE Cable-Tec Expo 2022, SCTE

[SCTE 55-1] “ANSI/SCTE 55-1 2019 – Digital Broadband Delivery System: Out of Band Transport Part 1: Mode A”, ANSI/SCTE, 2019

[SCTE 55-2] “ANSI/SCTE 55-2 2019 – Digital Broadband Delivery System: Out of Band Transport Part 2: Mode B”, ANSI/SCTE, 2019

[DSG] “DOCSIS Set-top Gateway (DSG) Interface Specification”, CM-SP-DSG-I25-170906, 2017

[R-OOB] “DOCSIS DCA - MHA v2 Remote Out-of-Band Specification”, CM-SP-R-OOB-I13-220531, CableLabs, 2022

[Chrostowski 2020] “Leakage In A High Split World – Detecting and Measuring Upstream Leakage Levels in a One Gbps Symmetrical High Split Hybrid Fiber Coax Network”, John Chrostowski, Greg Tresness, Dan Rice, & Benny Lewandowski, 2020

[DOCSIS 3.1 MULPI] “DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification”, CM-SP-MULPIv3.1-I23-220328, CableLabs, 2022

[DOCSIS 4.0 PHY] “DOCSIS 4.0 Physical Layer Specification”, CM-SP-PHYv4.0-D01-190628, CableLabs, 2019

[MB120 Data Sheet] CommScope STARLINE MB120 1.2 GHz MiniBridger Amplifier Data Sheet, CommScope

[SCTE 279 2022] “SCTE 279 2022 – 1.8 GHz Broadband Radio Frequency Hardline Amplifiers for Cable Systems”, SCTE, 2022



Creating Infinite
Possibilities.

Preparing the Outside Plant Network for the road to 10G

Mike Spaulding

V.P. of Network Maintenance
Comcast Corp.

Introduction

- Multiple paths to the same destination
- If amplifiers are involved, there could be some shared “pot-holes”
- Lessons learned from the activation of expanded spectrum



Methodologies

Higher or Wider?

FDX or FDD


Higher Modulation Orders

Wider Bandwidth

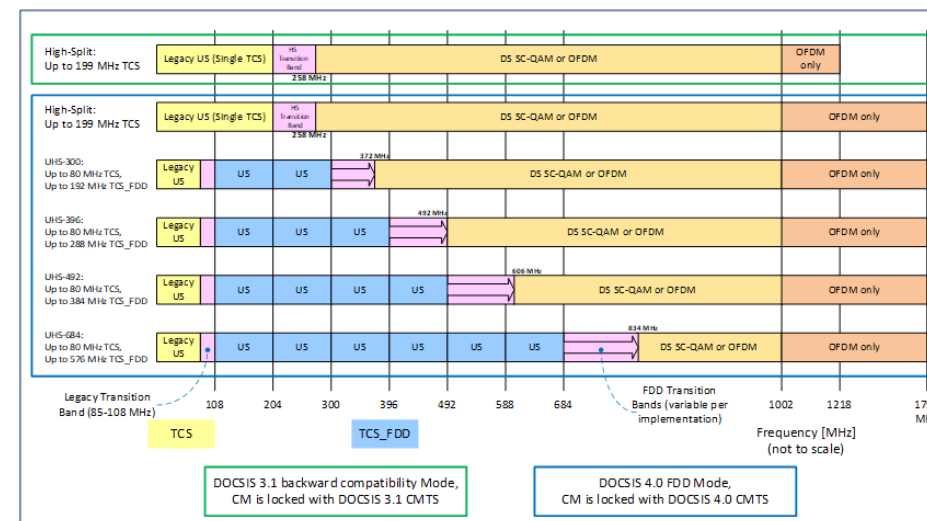
Both?

	Modulation Order	OFDM	OFDMA
Constellation minimum SNR/MER Performance of OFDM and OFDMA carriers	64-QAM	N/A	23dB
	128-QAM	N/A	26dB
	256-QAM	27dB	29dB
	512-QAM	30.5dB	32.5dB
	1024-QAM	34dB	35.5dB
	2048-QAM	37dB	39dB
	4096-QAM	41dB	43dB

Configurable FDX Allocated Spectrum Bandwidths




Configurable FDD Allocated Spectrum Bandwidths




Optical Architectures and Capabilities

- Analog or digital optics vs.
- R-Phy or MAC-Phy

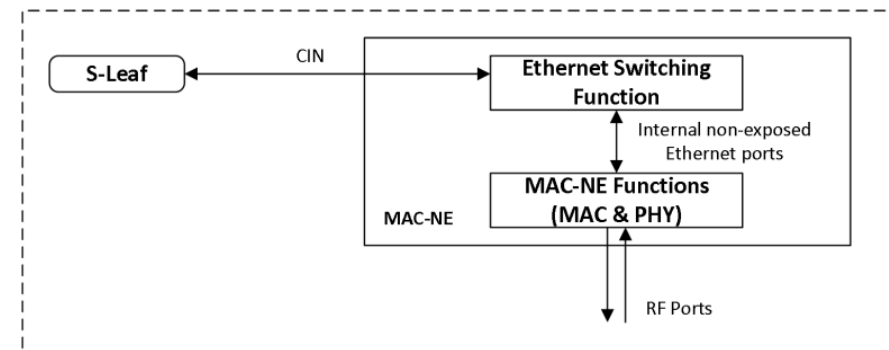
Analog DFB Laser NPR Curve



R-PHY Internal and External Connections






MAC – Network Element Internal and External Connections



RF Architecture Characterization

- Amplifier and BW selection
- Amplifier Characterization
- Legacy Plant Evaluation




Evaluating Legacy Plant

Active and Passive Obstacles


Active Devices

- “Temporary” amplifier placements
 - May not be in design files or maps
- Sub-Split Components in Mid or High Split amplifiers



Passive Devices


- Line Equalizers
- Conditioned Taps
 - Both may contain duplexers for upstream conditioning



Evaluating Legacy Plant




Coaxial Cable and Connectors

- Fused Disc cables with “mold spike” beyond 1 GHz
- Hardline connectors may not be “tuned” to higher frequencies





Craft Sensitivity

- Coaxial Cable Management
- Component Frequency Response
 - Extended Frequency Tuning
 - Fused Disc Cable
- Frequency Specific Testing



PNM Tools

- OFDMA/OFDM analysis
- FBC
- Thresholds



Device Health															
Registration State	6 (Online)														
Down Rx Power	5.2	5.2	5.2	5.4	5.5	5.7	5.8	6	6	6.1	6.2	6.1	6.1	6	6
Downstream SNR	35.5	35.4	35.5	35.5	35.5	35.5	35.6	35.6	35.7	35.3	35.8	35.8	35.8	35.7	35.7



Additional Considerations

Parallel Efforts

- Behavioral Identification
- Training
- Test Equipment





Creating Infinite Possibilities.

Thank You!

Mike Spaulding

V.P. Plant Maintenance

michael_spaulding@cable.comcast.com

Privacy Posture 2022

The Intersection of Technology, Policy, Standards, and Security

A Technical Paper prepared for SCTE by

Brian Scriber

Distinguished Technologist & Vice President Security and Privacy Technologies
CableLabs

858 Coal Creek Circle, Louisville, CO 80027

b.scriber@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Security and Privacy	3
3. Privacy Enabling Technology.....	4
3.1. Privacy Compliance	4
3.2. Technology Solutions.....	4
3.3. Intersection of Tools Supporting Technology and Compliance	6
4. Risks and Threats to Protected Data	6
4.1. Risks Resulting in Exposure	6
4.2. Risks of Data Exposure.....	7
4.2.1. Malicious Actors	8
4.2.2. Regulatory, Reputational and Legal Response	8
5. Technology Policy and Privacy	8
5.1. United States.....	9
5.2. International.....	10
6. Privacy Standards Community.....	11
7. Conclusion.....	12
Abbreviations	13
Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 North America Data Privacy Software Market Size (Fortune Business Insights, 2022).....	4
Figure 2 Real-Time Bidding Broadcasts Per Day (ICCL, 2022)	7
Figure 3 Growth of State Privacy Regulation.....	9
Figure 4 US State Privacy Legislation Tracker 2022	10

List of Tables

Title	Page Number
Table 1 Standard Development Organizations and Initiatives.....	11

1. Introduction

Privacy Engineering: The Intersection of Technology, Policy, Standards, and Security

While the term “security” is often used interchangeably with “privacy,” these two disciplines require different skillsets and process. Privacy focusses on access control and data usage. The toolset for the nascent privacy engineering field requires awareness of the development in the technology sector, awareness of risks and threats to protected data, privacy law in relevant jurisdictions, and different privacy standards, as these standards help set the bar as to what constitutes appropriate privacy protections. Security is focused on data access, integrity, and confidentiality. Organizations holding data need to be aware of the business and regulatory risks associated with this data and shape their internal privacy engineering and privacy policy compliance teams to mitigate these risks. This paper addresses the data security and privacy landscape for 2022 to provide some assistance in assessing your organizations’ security and privacy posture.

2. Security and Privacy

Security and privacy get grouped together, it’s easy to think of them as the same thing, they are not. Privacy is about the decisions that surround competing claims for access to, modification of, deletion of, and altering the disposition of information (Bambauer, 2013), as well as the legal right, uses, and potential ownership of data. Security, by the other token, is about the confidentiality, the integrity, and the availability of the data being accessible to only those people or system with appropriate identity and credentials.

In the cable ecosystem, a network operator occupies an interesting position, the operator helps to secure the data in transit through technologies like the CableLabs[®] DOCSIS Security Specification, protecting the confidentiality and integrity of the data through encryption, hashing, and message authentication. At higher levels of the OSI stack, over 92% of US web traffic transiting the internet today uses HTTPS (Google, 2022) and encrypting from endpoint to endpoint.

In many cases, the tools used to deliver security can be used to protect privacy, an example of this is how confidentiality in the above two examples protects from unauthorized users seeing message contents; recall, however, that privacy is about the decisions and competing claims to access information. An encrypted network does not help protect a user from endpoints and services that are going to sell user behavior and identification. The endpoint operator may claim that those observations and data were collected using their hardware, software systems, and algorithms, and how they handle the disposition of those data is their prerogative. The user, on the other hand, may claim that the observations collected about them, and their identification data is personal, and that the expectation was that the interactions made were not reasonably, or lawfully, expected to be shared with others or were used in ways that were not in the interests of the user.

Those differing claims create a new privacy discipline, new technologies and new tools unique to privacy engineering that are not exclusively within the security discipline. This work addresses the current state of how we capture, classify, and protect these aspects and nuance of data.

3. Privacy Enabling Technology

Privacy Engineering, as a discipline, is nascent, as is the supporting ecosystem of products and tooling. Some of the tools in this area begin to show how their application and outcomes drive advancement in data protection and privacy compliance.

3.1. Privacy Compliance

Compliance can depend on the legal jurisdiction, the domain data is held within, and the status of who holds the data (IAPP 2022). A hospital in Tennessee running field trials of new treatments will have different compliance hurdles from a Californian network operator or a direct-to-consumer retailer in Germany, but all of these do have compliance considerations. To address the multitude of combinations of the above, there are new tools being developed and expanded in the privacy sector which are disparate in many ways from the security sector. The Data Privacy Software market size in 2022 is \$US3.26B in 2022 but has a CAGR of 40.8% leading to a projected size of \$US25.85B by 2029 (Fortune Business Insights, 2022). North America alone will grow from \$US682.9M to \$US9.3B by 2029 (See Figure 1).




Figure 1 North America Data Privacy Software Market Size (Fortune Business Insights, 2022)

3.2. Technology Solutions

Several solutions exist in this growing Privacy Enabling Technologies (PET) space, many of these can be grouped together, but custom offerings are appearing in the market and in conferences that cannot always be categorized early.

Traditional data masking techniques include obfuscation (signal-to-noise ratio reductions to hide data), anonymizers (hiding identity or limiting cross-session tracking), pseudonymization (replacing protected data with different identifiers), and data minimization to reduce the overall amount of work performed by remaining PETs.

Self-sovereign identity solutions take blockchains, smart contracts, and an ecosystem approach to exchange of information and verifiable rules for how that data is accessed, used, stored, deleted, and shared. Protected data is stored either off-chain, or encrypted on the chain, and the sharing of that data between participants in the ecosystem is managed through the smart contracts of the chain, recording agreements, limiting use, and setting up the rules for data disposition.

Homomorphic encryption and Zero-Knowledge Proofs (ZKP) are mathematical concepts related to encrypted data whereby an operation such as a query can be performed on the ciphertext without decrypting it. In the hospital example from above, consider a study participant who is being evaluated for treatment of a new concern that arose during the study. While the patient's research status is protected (whether they are in the control group or the treatment group), the treating physician may need to consider a pharmaceutical treatment for the new concern. They query the system to find out the answer to a query that includes the patient's research status, and if they have allergies to the potential medication, and if there are drug-interactions with the research study course of treatment. With zero-knowledge proofs and homomorphic encryption, an authorized inquiry can get an answer for situations like the above without exposing the underlying data, but if the query isn't carefully constructed to be only what's necessary, data can be exfiltrated from multiple queries or overly broad queries.

Differential privacy is another tool based on mathematics. In hospital examples, like the one above, the concern is that medical data will be revealed, but if any given datum or small collection of data were to have a sufficient probability of being false, or untrue, then the release of information would not be damaging to an individual and the aggregate dataset would still be useful for use cases like that of the hospital field trial. For example, modifying data with an acceptable amount of noise to reduce the probability of identifying an individual but still providing a sufficiently reliable result to a query. For example, a query made with an exact age may be used to identify an individual. Adding noise to the data so an age range is created lessens the probability of identifying an individual but still returns clinically relevant information, such as adding sufficient noise to the age data so the query returns results for individuals aged from 34 to 38.

Federated Learning, also Privacy Enhanced Federated Learning (PEFL) is a subset of machine learning and artificial intelligence for using multiple datasets without sharing data (M. Hao, 2020). The data are distributed across independent local hosts and algorithms are trained for each host or dataset and for contributing to aggregate observation. This approach allows for data to be neither centralized, nor homogenous in its structure or distribution. What is shared are the resultant weights and biases for neural networks which can then be used in aggregate without risk of sharing private data used to obtain the trained network.

Secure Multi-Party Computation (SMPC or MPC) is based on the idea of several parties or systems who work together to solve problems using disparate stores of data, not unlike Federated

Learning, but in SMPC the objectives, rules, and model parameters are shared using finite field cryptographic tooling.

3.3. Intersection of Tools Supporting Technology and Compliance

There exist several different hybridized versions of technical solutions and approaches from Technology Solutions, some integrating cryptography in various steps for additional protection from accidental exposure or malicious actors. There also exist several tools for enterprise compliance with privacy regulation, these help to automate the auditing of datastores, databases, files, as well as system interfaces and communication tools for the storage or transit of data that is likely to be protected. The identification of that likelihood can range from simple rules to machine learning based on general training or custom training for specific jurisdictions and industries. The newest entrants into the tooling market are development operations tools (DevOps and PrivOps) where engineers using build tools for custom software development have steps where fields likely to have protected data are identified, in part to help train the privacy engineers as well as identify potential areas inside the software for additional protection.

4. Risks and Threats to Protected Data

Risks resulting in data exposure are different from the risks of exposure. Several risks can result in exposure, these range from accidental, to intentional, and even established markets for this protected information.

4.1. Risks Resulting in Exposure

Risks resulting in exposure include accidental exposure (e.g., misplaced documents), endpoint sale of data, online tracking (e.g., browser fingerprinting, cookies, etc.) intentional aggregators and profile building, linked data sets building identifiable collections (unintentional or otherwise), active fraud (these often have colloquial terms like “social engineering”, “phishing” and “spear-phishing”), real-time markets for data (whether sold or shared), insider threats, data loss or deletion (leading to misidentification or misclassification which can result in as many concerns as incorrect data), intentional exfiltration, technical vectors (insufficient security, unpatched systems, trojans, compromised websites, mobile and personal devices, removable media, poor configuration management, and access to local or cloud servers), as well as several other security compromises from various actors with different intentions (hactivism, cybercriminal networks, disgruntled employees, etc.).

Some of these risks may be more likely than others, and some risks may have smaller or more targeted datasets. The granularity and compartmentalization of the protected assets can limit the scope of a breach and each of the above risks has its own set of potential mitigation steps. This is an area where there is a significant overlap with security, however, some of the risks to exposure, above exist outside that intersectionality.

Risks where the entity or system on the other side of the endpoint has access to protected data (as defined by the laws of the prevailing jurisdiction) don’t necessarily have a security solution;

these include the intentional collection and sharing of data by the operators of that endpoint. Browser fingerprinting, activity monitoring (e.g., how quickly or slowly the user scrolls past different content, or when and where a mouse hovers, or what videos are played and when they are stopped), markers from previous websites, these are all aspects of how even unintentional data is collected. When that data is then shared with larger networks, when linked (or collections of unlinked data) are sold in real-time bidding markets (ICCL, 2022), or when it is retained and associated with future transactional records, individuals and corporations lose control over data that may be sensitive or protected.

Empirical evidence on the public understanding of the scale and scope of what is tracked seems to indicate a significant underestimation of how much data on individuals is collected each day. The Irish Council on Civil Liberties (ICCL) published a breakdown by European and US locality showing this data being sold in the greater than \$US117B Real-Time Bidding (RTB) network market, shown in Figure 2.




Figure 2 Real-Time Bidding Broadcasts Per Day (ICCL, 2022)

4.2. Risks of Data Exposure

Risks of exposure, what happens after the data is in the hands of another party, focuses on two primary aspects, the use of that data by malicious actors, and the regulatory and legal response.

Operational risk is not addressed here, because this paper focuses upon data exposure, not necessarily data obfuscation, modification, or deletion as is common in ransomware attacks.

4.2.1. Malicious Actors

Buyers in the markets described in Risks Resulting in Exposure include sales and marketing organizations, but the buyers also include malicious actors who seek to monetize the data collected. This can be accomplished through resale (there are underground marketplaces specializing in the sale of personal information), through direct ransom for the data (pay to have it returned and presumably not used or resold), and through extortion and the threat of public release of the data.

Ransomware used to be primarily about encryption and extortion payments to return to normal operations, but ransomware attacks that exfiltrated data increased from 22% of cases in Q2 2020 to a massive 81% in Q2 2021 (~270% YOY increase) (Schein 2022). These actors have recognized that threats of exposure can be part of criminal extortion, and some of these actors, when stymied by corporate victims that are reluctant to pay, have turned to threats (and action) of direct engagement with the individual people whose data is the subject of the breach.

4.2.2. Regulatory, Reputational and Legal Response

Those organizations who find that data in their stores has been compromised also find that there may now be both legal responsibilities and regulatory notifications required. Some jurisdictions require notification of real or potential breaches of protected data; this can be based upon the size of the company, the size of the breach, the content of the breach, and the industry within which the company may play. The legal impact may also involve public or private civil responses from those who suffered harm through the breach; depending on several factors, including the jurisdiction, this could involve significant damages.

In addition to the legal requirements on notification, there may be regulatory responses that could involve significant fines. In the European Union, significant violations of the General Data Protection Regulation (GDPR) can result in fines up to 4% of the corporation's gross annual revenue for each offense. Notifications also carry with them the reputational risk to the organization, potentially signaling to the market that they may be ineffective or careless with protecting this sensitive data. Starting with the probability of breach based on controls and tooling; then comparing that to the legal, regulatory, and reputational price organizations may pay in the event of a breach, an organization can balance risk and plan for mitigation investments.

5. Technology Policy and Privacy

The disparate technology policy and regulation across multiple jurisdictions makes this discipline difficult to standardize. These policies have an impact on the technology and are critical to understand for businesses who are engaged in commerce in these jurisdictions. This

paper is not intended as legal advice, please check with your own legal counsel as this may be out of date or incomplete.

5.1. United States

The United States does not have comprehensive national preemptive consumer privacy policy, and although talks on this front continue as of this paper's submission, rapid changes in that status are not anticipated. This said, there are several privacy legislative initiatives, some of which are laws, that cover specific sectors and actors. Examples of areas where laws exist in the US are in health privacy, finance, and protecting children. Other actions have an impact on privacy including cybersecurity, trade, and restrictions on governmental actions. Existing state legislation enacted over the last couple years can make for difficult corporate navigation of privacy rulemaking.

The growth of privacy initiatives around legislative action in the United States does show tremendous interest and inertia; this topic appears to be growing, but the resultant proposals are not consistent, and due to that non-standard deployment, privacy compliance is also growing in expense and potential for errors across jurisdictions. The growth of state privacy legislation can be seen in **Error! Reference source not found.** (IAPP 2022).




Figure 3 Growth of State Privacy Regulation

Currently, five states have enacted laws addressing private data and enterprise responsibilities, California (California Consumer Privacy Act & California Privacy Rights Act), Colorado (Colorado Privacy Act), Virginia (Virginia Consumer Data Protection Act), Utah (Utah Consumer Privacy Act), and Connecticut (Connecticut Data Privacy Act). Additional bills were considered during this 2022 legislative session, visible in Figure 4 (IAPP 2022).

The five states that have laws on the books, have some commonalities, but some dramatic differences. The consumer right to access, rectify, delete, and restrict records exists in all five, as do the business responsibilities for opt-in as the default, transparency, and limits on processing based on purpose of the data. Additionally, enterprises cannot discriminate against consumers who are exercising their rights.

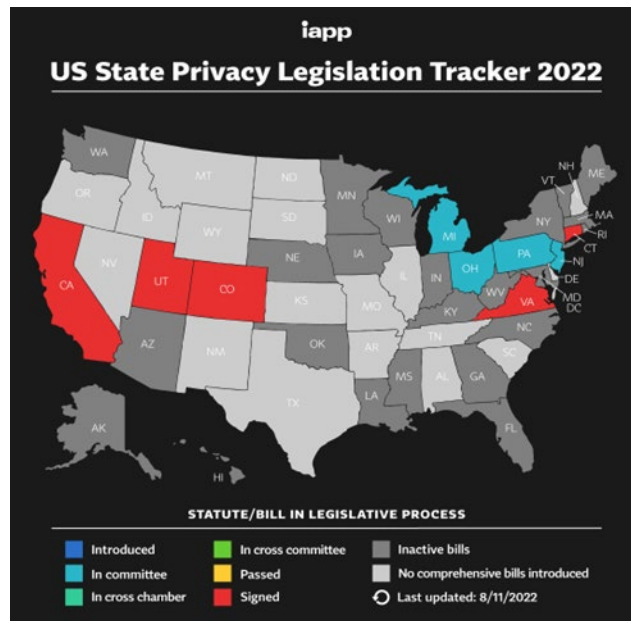


Figure 4 US State Privacy Legislation Tracker 2022

Utah does not have restrictions on automated decision-making based on protected data while the other four states do restrict this (California added this in the CPRA). In Colorado, Connecticut, and Virginia, additional rules are in place for risk assessments and those three states rely upon the Attorney General to prosecute cases. California is currently the only state in the United States with a private right of action for privacy violations. Detailed rulemaking is still taking place in these states, so the fine-grained details of these regulations are still subject to some level of change.

5.2. International

The European Union and eighteen other countries (excluding the United States) have comprehensive consumer privacy legislation and regulation, these include Argentina, Armenia, Australia, Benin Republic, Brazil, Canada's Personal Information Protection and Electronic Documents Act (PIPEDA), China, Colombia, European Union, Hong Kong, Israel, Kenya, New Zealand, Nigeria, Philippines, Singapore, South Africa, South Korea, and Turkey.

As in the United States the consumer right to access, rectify, delete, export, and age-based opt-in options are predominantly consistent across these. The default opt-in or opt-out as well as the consumer right not to be subject to automated decision-making seem to be inconsistently applied or available. Business obligations in these markets often include transparency, purpose limitations for data retention, data minimization, security requirements (somewhat varied), record keeping and breach notification requirements. Requirements for data protection officers, international data transfer restrictions, preemption, and sector-specific regulation are inconsistently across these jurisdictions. The EU, Australia, and a few others require privacy by design raising questions about enforcement interpretation with older systems and processes.

6. Privacy Standards Community

Tool development for inconsistent requirements like those listed in Section 5.2 is difficult due to the lack of economies of scale across ecosystems, in some cases this is made even more difficult by conflicting expectations (e.g., opt-in versus opt-out). The lack of mature, consistent, international standards and due to the variability in technology options, network operators and enterprises are left buying privacy compliance tools that cannot interoperate effectively with other tools in the space, where one tool can't complete all the tasks, and where incumbent vendors and contracting rules inhibit innovation.

For enterprises and operators who are looking for standards that support operations and enable innovation, Table 1 is meant to serve as an incomplete guide:

Table 1 Standard Development Organizations and Initiatives

Organization	Document/Specification/Initiative	Notes
ISO/IEC	Privacy Information Management Systems Scheme (PIMS Scheme) 27701 & 27702 (ISO 2019)	
National Institute of Standards and Technology (NIST)	Privacy Framework	1.0 (Jan 16, 2020) United States Dept. of Commerce
International Association of Privacy Professionals (IAPP)	www.iapp.org	Policy-neutral information privacy organization. Certification and professional credentialing
World Wide Web Consortium (W3C)	https://www.w3.org/Privacy/IG/	Public-interest non-profit web-focused privacy standards group.
3 rd Generation Partnership Project (3GPP)	https://www.3gpp.org/	Mobile broadband standards development organization. Consumer data privacy is a working group topic for R18 in 2022.
Connectivity Standards Alliance (CSA/Matter)	https://csa-iot.org/all-solutions/matter/	Consumer IoT standards development organization.
Institute for Electrical and Electronic Engineers (IEEE)	https://digitalprivacy.ieee.org/	Privacy collaboration, policy, and research for individual private needs online

Future Directions Digital Privacy Initiative		
Wireless Broadband Alliance (WBA)	https://wballiance.com/wi-fi-imsi-privacy-protection/	Permanent IMSI privacy protection initiative
Multi-Party Computation Alliance (MPC Alliance)	https://www.mpcalliance.org/	Standards and advocacy group focused on the adoption of MPC technology

7. Conclusion

Privacy technology is advancing on several fronts, technology, policy, standards, and the discipline's overlap with security. Entirely new disciplines like Privacy Engineering are being developed (Carnegie Mellon 2022), nascent tools are being brought to the market, standards are catching up, and legal, compliance and regulatory frameworks are being established and renovated. The objective of this work was to provide the reader the tools necessary to begin their own assessment of the posture of the reader's own organization, to understand the different aspects of the space, better evaluate the nuance and differences between privacy and security and begin to learn to discern where and how privacy technologies should be applied within their influence.

Abbreviations

CAGR	Compound Annual Growth Rate
CCPA	California Consumer Privacy Act
CDPA	Connecticut Data Privacy Act
CPA	Colorado Privacy Act
CPRA	California Privacy Rights Act
DevOps	Development Operations
DNS	Domain Name System
DOCSIS	Data Over Cable Service Interface Specification
GDPR	General Data Protection Regulation
HTTPS	HyperText Transfer Protocol Secure
ICCL	Irish Council for Civil Liberties
MPC	Multi-Party Computation
PIMS Scheme	Privacy Information Management Systems Scheme
PIPEDA	Personal Information Protection and Electronic Documents Act
PrivOps	Privacy Operations
RTB	Real-Time Bidding
SCTE	Society of Cable Telecommunications Engineers
SMPC	Secure Multi-Party Computation
UCPA	Utah Consumer Privacy Act
VCDPA	Virginia Consumer Data Protection Act
W3C	World Wide Web Consortium
ZKP	Zero-Knowledge Proofs

Bibliography & References

Bambauer, Derek E. "Privacy versus security." *J. Crim. L. & Criminology* 103 (2013): 667.

Google, 2022: Transparency Report on Network Traffic, HTTPS encryption on the web:
<https://transparencyreport.google.com/https/overview?hl=en>

M. Walker, B. Scriber, K. Shockey, D. Slagle "Have Your Privacy Cake and Eat It Too: How New Technologies Look to Protect Consumer Privacy While Promoting Innovation." *TPRC47: The 47th Research Conference on Communication, Information and Internet Policy*. 2019.

Fortune Business Insights, Report ID: FBI105420, Data Privacy Software Market Size, Share & COVID-19 Impact Analysis, By Deployment (On-premises and Cloud), By Application (Compliance Management, Risk Management, Reporting & Analytics, and Others), By Organization Size (Small & Medium Enterprises (SMEs) and Large Enterprises), By Industry (BFSI, IT and Telecommunication, Government, Manufacturing, Retail, Healthcare, and Others), and Regional Forecast, 2022-2029

Dilmegani, Cem, Information Security Privacy Enhancing Technologies and Use Cases,
<https://research.aimultiple.com/privacy-enhancing-technologies/> *AI Multiple* (2022)

M. Hao, H. Li, X. Luo, G. Xu, H. Yang and S. Liu, "Efficient and Privacy-Enhanced Federated Learning for Industrial Artificial Intelligence," in *IEEE Transactions on Industrial Informatics*, vol. 16, no. 10, pp. 6532-6542, Oct. 2020, doi: 10.1109/TII.2019.2945367.

O. Goldreich, S. Micali, A. Wigderson "How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority". STOC 1987: 218-229

ICCL (Irish Council for Civil Liberties), *Note on scale of Real-Time Bidding data broadcasts*, <https://www.iccl.ie/wp-content/uploads/2022/05/Mass-data-breach-of-Europe-and-US-data-1.pdf>, 2022

M. Schein, 16 Mar 2022, Marsh McLennan, National Co-Chair Cyber Center of Excellence

IAPP International Association of Privacy Professionals), 2022 Resources: <https://iapp.org/resources/article/us-state-privacy-legislation-tracker/> , <https://iapp.org/resources/article/the-growth-of-state-privacy-legislation-infographic/> , https://iapp.org/media/pdf/resource_center/global_comprehensive_privacy_law_mapping.pdf

ISO, 2019, <https://www.iso.org/news/ref2419.html>

NIST, Privacy Framework 1.0, <https://www.nist.gov/privacy-framework/privacy-framework>

Carnegie Mellon, Privacy Engineering, 2022, <https://privacy.cs.cmu.edu/>

Proactive Network Maintenance (PNM) Paves the Way for More Upstream Bandwidth

A Technical Paper prepared for SCTE by

Takashi Hayakawa

Engineer 5
Comcast Corporation
4100 E Dry Creek Rd Centennial, CO 80122
720- 938-4393
Takashi_Kayakawa@cable.comcast.com

Mike O'Dell

Director, Network Maintenance
Comcast Corporation
215 E. North St. New Castle, PA 16101
412-417-0481
Michael_ODell@cable.comcast.com

Paul Schauer

Distinguished Engineer
Comcast Corporation
183 Inverness Drive West, Englewood, CO 80112
720-276-8308
Paul_Schauer@cable.comcast.com

Larry Wolcott

Comcast Fellow
Comcast Corporation
183 Inverness Drive West, Englewood, CO 80112
303-726-1596
Larry_Wolcott@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Evolution of Pre-Equalization Analysis and Proactive Network Management.....	4
2.1. Differences Between SC-QAM and OFDMA Pre-EQ	6
3. Coefficient Reporting.....	7
3.1. Raw Coefficients	7
3.2. Pre-equalizer Metrics	8
3.3. Coefficient Overflows and Underflows	9
3.4. Monitoring a Large Group of Nodes and Modems	9
4. Learnings From the Field	13
4.1. Legacy plant.....	13
4.2. Partial bonding	15
4.3. Performance limitations – how to detect?	16
4.4. Customer performance aspects – speed tests	16
4.5. Implications of PMA	19
5. Conclusion.....	19
Abbreviations	20
Bibliography & References.....	20

List of Figures

Title	Page Number
Figure 1 – Example of SC-QAM equalized spectrum	6
Figure 2 – Example of OFDMA equalized spectrum	7
Figure 3 – Example of OFDMA pre-equalizer coefficients and metrics	8
Figure 4 – Multiple modems' pre-equalizer amplitudes on a node	10
Figure 5 – Two modems under impairment with different signatures	10
Figure 6 – Multiple modems under heavy impairment with the same signature (red)	11
Figure 7 – Impaired modems mapped on the same network branch	12
Figure 8 – Same 12 modems' frequency response after repair	12
Figure 9 – Example of passive plug-in components	13
Figure 10 – Conditioned line equalizer	14
Figure 11 – Sub-split amplifier installed in a mid-split system	14
Figure 12 – Sub-split return equalizers installed in mid-split amplifiers	15
Figure 13 – Frequency response of OFDMA channel through a sub-split equalizer	15
Figure 14 – Field meter speed test results in high-split system.....	17
Figure 15 – Upstream signal analysis of expanded upstream using a field meter	18

List of Tables

Title	Page Number
Table 1 – PNM Features compared	5
Table 2 – ICFR and coefficient metrics	9
Table 3 – Line equalizer properties showing 5 MHz to 42 MHz operation	13

1. Introduction

Cable operators are moving towards wider upstream spectrum for improved speed and bandwidth. Activating new upstream spectrum can come with a number of unforeseen challenges. Fortunately, over a dozen years of Proactive Network Maintenance (PNM) experience can provide an excellent foundation for detecting, characterizing and repairing many of the common obstacles. The result is faster, smoother orthogonal frequency division multiple access (OFDMA) deployments and finding and fixing network problems before our customers are impacted.

In this paper, the authors, Hayakawa, O'Dell, Schauer and Wolcott present the latest results of mid-split and OFDMA from the field. They review a number of unforeseen issues including technical and operational solutions. From frequency response issues caused by long-forgotten network equipment and filters to partial-bonding caused by in-home amplifiers, there are plenty of opportunities to get ahead of. This paper provides valuable insight for cable operators embarking on mid-split and high-split transformations in conjunction with OFDMA technology.

2. Evolution of Pre-Equalization Analysis and Proactive Network Management

CableLabs has evolved three generations of DOCSIS PNM. Starting in DOCSIS 1.1, evolving through DOCSIS 2.0, and reaching wide deployment with the advent of DOCSIS 3.0, cable operators have a rich body of data to identify and correct impairments in the RF plant. While there are numerous data sets available in the DOCSIS and PNM requirements, the pre-equalization coefficients are among the most useful to isolate subtle degradations in RF plant performance.

Processing equalization data involves manipulating complex coefficients in the frequency and time domains utilizing Fourier transformations and related mathematics. The scope of these mathematics is outside of this presentation. They may be summarized as the dynamic manipulation of RF energy across frequency bands to compensate for non-ideal coaxial cable plant parameters.

Collection of pre-equalization data also depends on the cable modem SNMP Set and TFTP file return mechanism described in the DOCSIS 3.0 PNM specification. This process is outside the scope of this presentation as well; operators are encouraged to develop their infrastructure to support these new data sets.

In 1999, the first-generation pre-equalization work in DOCSIS 1.1 codified publishing complex equalization coefficients and dynamic RF management to the cable plant. It also experienced growing pains. The original eight coefficients, the accuracy of the measurements, and the implementation of the specification required significant revision as it moved from the laboratory to full production.

The second-generation pre-equalization work implemented in DOCSIS 2.0 improved the pre-equalization modeling. The number of coefficients expanded from 8 to the now familiar 24 as well as improvements to the algorithms.

Pre-equalization analyses advanced to a current best practice with the third generation of PNM in DOCSIS 3.0. CableLabs released a formal PNM specification and built richer data sets into the base DOCSIS MIBs as well.

With DOCSIS 3.1 OFDMA, CableLabs brings the fourth generation of PNM to cable plants. Changing from classic 6.4 MHz wide Single Carrier QAM channels to OFDMA channels that are up to 96 MHz wide with thousands of 25 or 50 kHz subcarriers brings new capabilities and analytical needs.

The DOCSIS 3.0 PNM capabilities have been covered in depth in previous SCTE papers. They are summarized and compared to the capabilities that DOCSIS 3.1 add to the analyses. These reflect major changes that operators need to embrace and will be detailed later in this paper.

Table 1 – PNM Features compared

US Measurement	DOCSIS 3.0 SC-QAM	DOCSIS 3.1 OFDMA
Channel Ranging Status	Exposed for each SC-QAM US channel on the CM and vital for detecting partial bonding.	Exposed for each OFDMA channel on the CM and include stronger indications of partial service.
Speed test	Coarse reductions in expected speed due to US channel impairments. Losing one out of four bonded US channels will result in up to a 25% reduction in measured US speed.	Subtle reductions in expected speeds due to OFDMA US channel impairments. Partial service capabilities, improved error correction, and larger bandwidth will ameliorate but not eliminate reduction in US speed
Spectrum analysis at CMTS	Collected from CMTS spectrum analysis table with limits on narrow or wide channels	New and expanded CMTS spectrum analysis table. New US triggered spectrum capture PNM file available for all US types
Rx power, MER, FEC at CMTS	Collected from CMTS and analyzed along D1.1, D2.0, and D3.0 evolutionary path.	New channel status table entry with expanded power, MER, and OFDMA performance data. New profile status entry reflecting total, correctable, and uncorrectable codewords per US OFDMA profile. New PNM US OFDMA Rx Power and Rx MER files reflecting OFDMA specific
Spectrum analysis at CM	Available for DS spectrum with limited support for US capture and extended upstream bands.	Expanded to include mid-split and high-split frequencies.
Tx power at CM	Long established 53dBmV (with some extensions)	Higher overall due to OFDMA: must be capable of transmitting a total average output power of 65dBmV
Upstream pre-EQ coefficients at CM	Collected from CM and analyzed along D1.1, D2.0, and D3.0 evolutionary path.	New PNM file. Revised analyses describe in the next sections.

2.1. Differences Between SC-QAM and OFDMA Pre-EQ

There are important changes in Data Over Cable Service Interface Specification (DOCSIS) 3.1 and OFDMA that provide improved troubleshooting capabilities over the DOCSIS 3.0 predecessor, single channel quadrature amplitude modulation (SC-QAM). The first is the larger DOCSIS 3.1 OFDMA channel width of up to 96 MHz, contrasted to the fixed SC-QAM channel widths of either 3.2 MHz or 6.4 MHz. With the wider channels, there is greater time resolution, which provides more precise distance calculations than narrower channels that also contain guard bands and roll off. The time resolution is calculated as the reciprocal of the total equalizer bandwidth.

For example, if an OFDMA channel is configured to operate with 50-kHz subcarrier spacing. In the frequency domain, each equalizer coefficient represents 50 kHz. For SC-QAM using 24-tap equalization coefficients of 6.4 MHz-width, each coefficient represents 233.33 kHz. This is over 4.6 times improvement in the frequency resolution of the equalized channel bandwidth. However, the resolution improvement is even greater when considering the additional inefficiencies inherent with SC-QAM channel shaping and adaptive pre-equalization. First, if the SC-QAM channels have spacing or guard bands, which the unoccupied spectrum will not be subject to the equalizer bandwidth used for troubleshooting. Also, because DOCSIS uses a root raised cosine shaping filter, the SC-QAM channels have a roll off factor of 0.2, leaving only 0.8 of the channel used for the symbol bandwidth. In the case of 6.4 MHz wide channels, only 5.12 Msym/s are visible to the equalizer. Likewise with the 3.2 MHz wide channels, the channel alpha is 2.56 Msym/s.




Figure 1 – Example of SC-QAM equalized spectrum

Figure 1 shows example comparison of 4 SC-QAM channels configured at 6.4 MHz channel bandwidth. While these channels occupy 25.6 MHz of spectrum, only 20.48 MHz of spectrum is available for troubleshooting. Also note that 5.12 MHz of that spectrum is spread across 5 spans of non-contiguous frequency spectrum, creating gaps that cannot be used reliably for troubleshooting.




Figure 2 – Example of OFDMA equalized spectrum

In comparison, Figure 2 shows an equivalent OFDMA channel configured at the same frequency spectrum of 25.6 MHz wide at 50-kHz subcarrier spacing. The entire 25.6 MHz of spectrum is equalized at 50 kHz frequency resolution, providing a total of 512 contiguous sample points, compared to nearly 88 sample points in the previous example (Figure 1).

3. Coefficient Reporting

3.1. Raw Coefficients

When a DOCSIS 3.1 modem has an operational OFDMA channel and its pre-equalizer is enabled, the modem can report the pre-equalizer coefficients. It will upload the coefficients on a specified trivial file transfer protocol (TFTP) server upon a request via simple network management protocol (SNMP) when the coefficients are available. The request typically takes several seconds involving a few SNMP commands and polling, sometimes ending up with a timeout and a failure.

The pre-equalizer covers all the subcarriers between the lowest and highest active subcarriers. Each coefficient is a complex number for a subcarrier and specifies an amplitude and phase adjustment to the subcarrier's transmission. The complex number is encoded with a pair of 16-bit fixed point numbers for the real and imaginary parts. Presenting the coefficients' magnitude values in the log scale gives the familiar amplitude of frequency response. Taking the differentials of coefficients' phase values gives the group delay.




Figure 3 – Example of OFDMA pre-equalizer coefficients and metrics

3.2. Pre-equalizer Metrics

The raw pre-equalizer coefficients are thousands of complex numbers. In order to assess the plant health, we need simpler metrics (Figure 3). The most straightforward is the in-channel frequency response (ICFR) which is the difference between the maximum and minimum of the subcarriers' amplitudes in dB.

Meanwhile, the modem performs the linear fit on the amplitude curve and publishes the following metrics on the SNMP table **docsPnmCmUsPreEqTable**. They are cheaper to retrieve than the coefficients as it takes only a few SNMP queries. The modem does the same calculation on the group delay, too.

- Linear fit
 - Mean (docsPnmCmUsPreEqAmpMean)
 - Slope (docsPnmCmUsPreEqAmpSlope)
- Nonlinear component (ripple)
 - Peak-to-peak (docsPnmCmUsPreEqAmpRipplePkToPk)
 - RMS (docsPnmCmUsPreEqAmpRippleRms)

The modem will have a low in-channel frequency response (ICFR) as well as low amplitude slope and ripple metrics when it is on a clean leg of the hybrid fiber-coaxial (HFC) node. We looked at thousands

of modems on a 44.4 MHz-wide channel among a few hundred nodes and observed the following relationship between the ICFR and the linear fit metrics.

Table 2 – ICFR and coefficient metrics

ICFR Range (dB)	Modem Count		Average Slope (dB/MHz)	Average Ripple Pk-to-pk (dB)	Average Ripple RMS (dB)
0..2	342	7.7%	0.015	1.148	0.204
2..3	1027	23.1%	0.037	1.278	0.235
3..5	1672	37.6%	0.072	1.505	0.289
5..7	798	17.9%	0.111	1.752	0.345
7..10	337	7.6%	0.159	2.353	0.498
10..15	216	4.9%	0.238	3.400	0.703
15..20	24	0.5%	0.204	9.876	2.084
20 or more	33	0.7%	0.158	21.812	3.082

3.3. Coefficient Overflows and Underflows

Each of real and imaginary coefficient parts can take a value between -4 and 4 with a 1/8192 increment. The practical maximum magnitude of a coefficient is 4 in linear (+12 dB) because the magnitudes over 4 will be clipped when the phase is at or near a multiple of 90 degrees.

The magnitude shouldn't go very low either because the coefficient will lose the resolution. In an extreme case where the magnitude is the minimum (1/8192), the phase resolution is 90 degrees. The amplitude resolution is 0.004 dB at the 1/4 magnitude (-12 dB) whereas it is 0.03 dB at the 1/32 magnitude (-30 dB). In most cases the amplitude stays above -15 dB. In very extreme cases we have seen some coefficients collapsed down to zero (0+0j).

Coefficient amplitudes above +12 dB or below -15 dB including the 0+0j value are a good indication of anomaly.

3.4. Monitoring a Large Group of Nodes and Modems

Retrieving the raw coefficients is slow and could fail after all. When monitoring thousands and millions of modems, it is economical to ask for the raw coefficients only when the modem can report them and its HFC condition looks interesting. There are several steps to consider before asking for the coefficients. The step 1-3 can be done by querying just the cable modem termination system (CMTS).

1. Does the HFC node have an OFDMA channel?
2. Is the modem online?
3. Has the modem succeeded in ranging the OFDMA channel?
4. Is the modem's preequalizer enabled?
5. Does the modem have the coefficient metrics at docsPnmCmUsPreEqTable?
6. Are the coefficient metrics interesting or bad enough? Such as the amplitude ripple peak-to-peak is greater than 5 dB or the RMS is greater than 3 dB.

When a modem with bad coefficient metrics is found, it makes sense to retrieve the raw coefficients from it as well as the modems nearby or of the entire node altogether to see whether the modem's impairment

is isolated or prevalent and to narrow down the possible location of impairment cause. If some modems exhibit a similar impairment and they are in the same neighborhood on the HFC topology, it is likely that a single cause is affecting those modems. Overlaying the modems' pre-equalizer amplitude curves gives a good insight.




Figure 4 – Multiple modems' pre-equalizer amplitudes on a node

Figure 4 is an example of overlay plot for a node. The modems on this node and most of the others had relatively low ICFRs. Some of them had many modems sharing the same minor microreflections whose cause would be near the CMTS or remote PHY device (RPD).




Figure 5 – Two modems under impairment with different signatures

On this node, two modems saw a high ICFR with very different curve signatures (Figure 5). The causes of impairment for the two were certainly different and likely between the tap and house. Indeed, they were on very different legs of the node. Several more nodes had multiple modems with isolated impairments.




Figure 6 – Multiple modems under heavy impairment with the same signature (red)

The more interesting situation, where multiple modems reported very similar frequency responses, was found in several other nodes. On one of such nodes, twelve out of 67 OFDMA-enabled modems showed a very similar pair of amplitude spikes whereas others didn't (Figure 6). The spike and dip at 48 MHz were anomalies by themselves as they exceeded the +12 dB ceiling and the -15 dB floor. The twelve were so similar that they were likely to be nearby each other under the same impairment. It turned out that they were all on the same branch of the node and the branch had no unimpaired modem, seen when mapped in Figure 7.




Figure 7 – Impaired modems mapped on the same network branch

It was an easy guess that the impairment came from where the bad branch began. It allowed the field technician to spot an old line equalizer easily. Once the equalizer was removed, the spikes went away in Figure 8.




Figure 8 – Same 12 modems' frequency reponse after repair

4. Learnings From the Field

4.1. Legacy plant

Every bandwidth expansion in the Hybrid Fiber Coaxial network has been met with unforeseen obstacles, both in the downstream and upstream. The flexibility of coaxial cable itself has allowed operators to upgrade electronics and optics and leverage the same coaxial cables successfully for many years. As the designs have evolved, so have the active and passive components deployed in the network.

When we consider the “return” portion of the spectrum, today understood to be 5 MHz – 42 MHz, even this has undergone several iterations. From 30 MHz as the upper band edge, to 40 MHz and ultimately 42 MHz, each of those seemingly minor changes have introduced complexity to the activation and reliable operation of the newly activated portion. Line passives, and return equalizers are a couple of the common passive components that can be disruptive to a bandwidth expansion, or network upgrade.

Line passives like equalizers, and conditioned taps have allowed designers and engineers to build networks that deliver very consistent levels to a large number of subscribers within the service area. Using passive plug-in components (Figure 9) like cable equalizers, cable simulators, return attenuators, or even high pass filters in the taps themselves, it is possible to condition the output signals at every tap port in the network to be within a small target window. They can also allow for the expansion of the feeder or distribution systems to a greater service area, resulting in fewer nodes or amplifiers, and reducing upgrade and operating costs.

However, some of these components do have frequency specific duplexers that would need to be replaced when expanding the bandwidth, and particularly in the upstream portion of the spectrum.




Figure 9 – Example of passive plug-in components

Some line equalizers, in addition to adding equalization to the forward bandpass to account for normal cable losses over long coaxial cable spans, can have a diplexed return path portion to introduce attenuation to the upstream path. These line equalizers would prohibit any expanded bandwidth above the diplex frequency from being enabled at the customer premises southbound of the passive device, examples in Figure 10 and Table 3.

Table 3 – Line equalizer properties showing 5 MHz to 42 MHz operation

Parameter	Frequency (MHz)	Equalization Mode	
		Typical	Maximum
Bandwidth Split (MHz)	---	42/51	
Cable Equalization (dB)	5-42	0	
	51-750	8	



Figure 10 – Conditioned line equalizer

In many cases a walk out is done prior to a system upgrade. This is useful in identifying the location of the previously installed amplifier locations, any traffic or access issues, or verification of passive devices that may need replaced. However, on occasion amplifiers are placed and the locations of these active devices are not updated on the system maps, or design documents. Instances such as: disaster recovery, temporary or special events, or temporary amplifier placements due to coaxial cable deterioration. While these are intended to be temporary, they can be overlooked and left in the system unnoticed. In these instances, they can cause unexplained failures in the recently expanded portions of the spectrum and may require a full physical walk-out of the network to locate.




Figure 11 – Sub-split amplifier installed in a mid-split system

In addition to “old” amplifiers, another potential obstacle to the activation of expanded spectrum is “old” passive plug-in components, example in Figure 9. Several amplifier platforms have remained largely unchanged through several generations, and the form factors for the passive plug-ins in the amplifiers themselves have also remained unchanged. Forward and Return equalizers in particular are a concern. Often, then new nodes and amplifiers are placed, the new portions of spectrum are not immediately activated with signals. The legacy channel line-ups are often retained for a period of time after splicing is complete. In these instances, it is difficult to immediately identify if a legacy plug-in component has been installed in an amplifier using PNM tools. Legacy verification practices such as the use of a signal generator and a spectrum analyzer would be required to ensure that the expanded bandwidth was usable, and there were no legacy components improperly installed in either an active or passive device.



Figure 12 – Sub-split return equalizers installed in mid-split amplifiers

As with the figure above we recognize that it can be difficult to identify the bandpass frequency of the return equalizers, unless you have a key to the letter code which denotes the upper band edge. In the example of Figure 12, the “S” on the return equalizer denotes an upper band edge of 40 MHz. The form factor however is the same as the mid-split return equalizer and can easily be installed in error. If the construction or maintenance teams use legacy sub-split channels for amplifier alignment or device recovery verification, then all indications will be that the system is operating normally and without error.

However, when the additional spectrum is activated, these devices can cause excessive tilt, due to the loss cause by the sub-split passive component, seen in Figure 13.



Figure 13 – Frequency response of OFDMA channel through a sub-split equalizer

4.2. Partial bonding

When the expanded portion of spectrum, like OFDMA is activated, the long loop automatic level control (ALC) function of DOCSIS devices located in the customer premises will attempt to use that portion of spectrum (if capable) and attempt to adjust the transmit power of the device at that frequency to accommodate the CMTS range request. This can result in a significant difference in transmit levels at the device when compared to the traditional sub-split upstream channels. If the transmit delta or tilt are severe enough to violate the dynamic range window of the device, the device will ultimately be forced into a partially bonded state.

There can also be legacy devices left in the network or drop system that can result in failure to activate the expanded portion of the spectrum. In-home amplifiers are commonly used to increase the amplitude of

signals in the home. There have been many different generations of in-home amplifiers deployed, and many, if not most of them, were designed and built for CATV systems that operated sub-split return paths. These sub-split diplexed devices can result in a failure to activate any additional bandwidth above 40 or 42 MHz. Other considerations for the subscriber drop system, are legacy filters and traps. These devices have been used for many years to manage service and channel access, or for frequency specific bandwidth management like, high pass or low pass filters to manage ingress in the premise. The location and placement of these filters may result in some DOCSIS devices to be unable to acquire a clear and unobstructed path back to the node. These instances would result in capable devices being considered as being in a partially bonded state.

4.3. Performance limitations – how to detect?

Depending on the phase of construction, and the business unit involved, there can be different methods to detect obstacles to the activation of “new” spectrum. As discussed in the Section 4.1, there are plenty of legacy devices that can remain in the trunk and distribution system that can affect the signal flow through the expanded spectrum areas. We have also identified several potential obstacles in the premise wiring system. While these effectively result in the same thing, there may be different methodologies to identifying where the obstacle exists.

During the construction phase of the upgrade, once the node, and all the amplifiers have been spliced in and activated, there are several legacy practices that can still be used for spectrum testing and validation. For many years, active sweep transmitters installed in the headend allowed technicians to prove connectivity, frequency response and level validation of a newly constructed network. The introduction of technologies like Remote-PHY and Remote MAC-PHY have rendered this practice obsolete, as the signal generation point is at the node, and no longer in the headend or hub. There are still however, hand-held signal generators and receivers that can successfully be used to insert a carrier at an insertion point, like a terminating tap, and receivers that can detect that inserted signal at an endpoint, like the optical node return input test point. This method is still perfectly valid; however, it requires a lot of time and mobility to validate all the potential end points in a node. Some test equipment can generate QAM signals in the return band to be received on a spectrum analyzer and be evaluated for both amplitude and quality.

The use of intelligence tools is a more efficient way to validate the upstream signal path, however it is also not without challenges. At the earliest stages of the construction and activation process, the downstream and upstream channel plans may not be fully evolved to activate capable devices in the premises. Further, the intelligence tools themselves may not yet be capable of collecting the additional telemetries available from the CMTS or v-CMTS and localizing them to a point in the network which could be causing the failure. It is incumbent on the operators to engage the tooling teams early in the engineering process of an upgrade to begin to develop the application programming interfaces (APIs) and endpoints between systems to be able to integrate newly activated spectrum telemetry into existing tool suites. The use of field strength meters with embedded D3.1 modems can be useful spectrum validation tools as well.

4.4. Customer performance aspects – speed tests

After all the construction, activation and validation steps have been completed, and the capable devices have been configured with the appropriate firmware to successfully use the newly activated spectrum, it is

finally time to see the fruits of all that labor. Just how fast can we go? There are many different methods to evaluating throughput, or speed, and different locations in the network where it can be measured.

Wireless device tests, hardwired computer tests, and field test equipment used by the technical staff can all garner slightly different results. Which one is right? The answer may be that all of them are. As with many aspects of the telecommunications network, the answer may be a nuanced one. Speed tests are a snapshot in time, of the performance of the path between the requesting device and a speed test server located somewhere on the network. There are many factors that affect the results of a speed test, such as: the type of device being used to initiate the test, the methodology it uses to connect to the network (i.e. wired or wireless), the location of the speed test server, relative to the initiating device, the volume of traffic on the network, and whether the test is customer initiated, or a background test executed by the provider.

There have been many studies and publications created on the different types, and the benefits of the various methodologies, but the focus of this paper is not to evaluate speed test methodologies, rather to provide some experiential suggestions on how best to use PNM tools and practices to successfully activate and validate mid or high split networks. To that end, here are a couple of simple suggestions. The first and possibly most gratifying way is to use a capable field strength meter at a point in the network to validate the use of all the available spectrum. The ability to use a field meter with a DOCSIS 3.1 embedded modem and configured to hit a speed test server can quickly confirm that the upstream path is capable of transporting signals in the expanded portion of the spectrum.

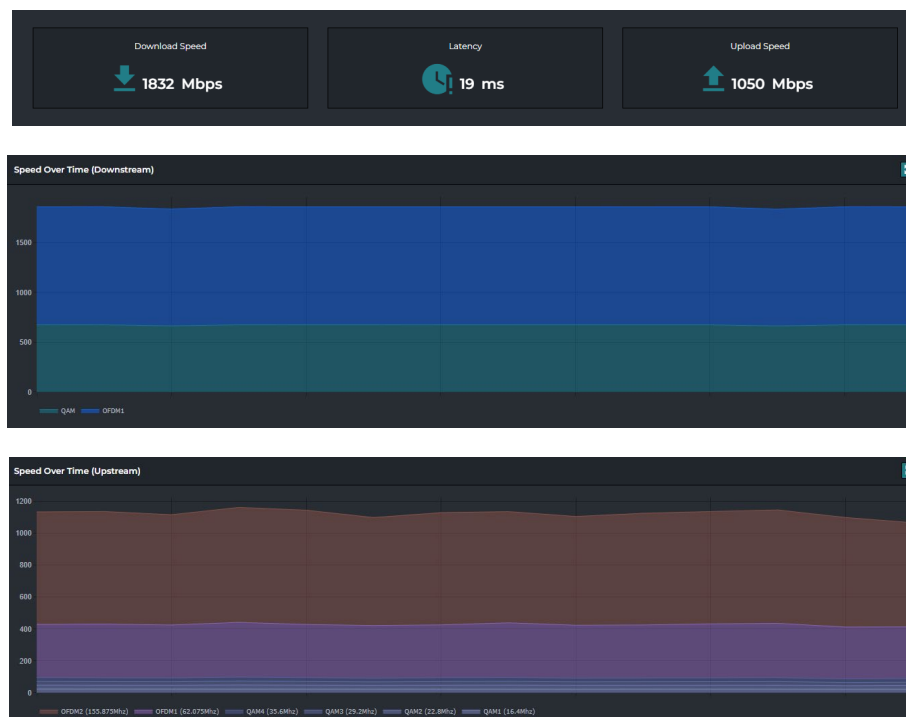


Figure 14 – Field meter speed test results in high-split system

This simple test validates a few things. First that the modem in our test equipment was successfully able to connect and transmit data through the expanded spectrum, and that we were able to achieve a reasonably high data rate through the network. Understanding that this is only one test from a specific

point in the network, and a point in time, we can replicate this at various terminations in the network to validate that there aren't any significant obstacles between the point of the test, and the node location. Devices like the one used for the speed test illustrated in Figure 14 can also provide additional data regarding the quality of the network as it pertains to any impairments that may exist upstream from the test point.

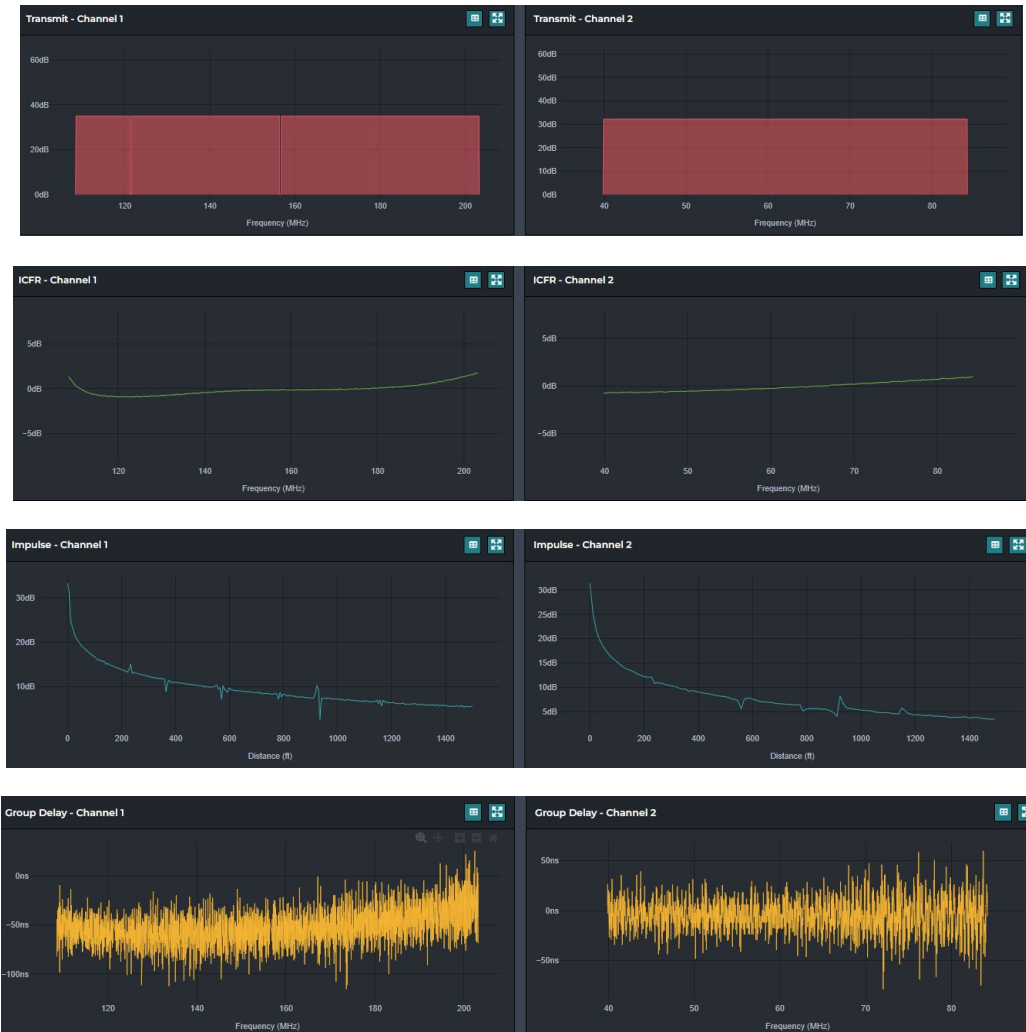


Figure 15 – Upstream signal analysis of expanded upstream using a field meter

Having data points like these are analogous to PNM tool results once the spectrum is activated, and devices are enabled to begin using this spectrum. They can be reliably used to indicate impairments in the transmission medium (coaxial cables), or of passive or active devices that may still be present in the network and have gone unidentified.

There is obvious value in using field test equipment for validation of speeds, and the additional performance data that it can provide. It is, however, time consuming and labor intensive. The other alternative is to do bulk “background” speed testing once devices are activated and online. There are pros and cons to this method also, however. On the plus side, the volume of devices located in different parts of the network can be very useful in localizing potential plant impairments that influence speed. This is also useful in validating those devices are successfully using the expanded spectrums based on their speed

results. While they may not be extremely precise on the actual network speed the way that a subscriber-initiated test could be, they are at least directional with respect to the functional operation of the outside plant portion of the network. The bulk speed tests may not fully stress the network in the way that a subscriber-initiated test does either or transmit as much data to calculate a more precise throughput or speed result, but it can be a good indicator of relative health and functionality of the network. Another potential downside to the bulk testing, is that it doesn't provide any meaningful intelligence on the quality of the coaxial plant or any potential impairments, provided that the subscriber modems are still able to transmit data through the expanded spectrum at some level.

4.5. Implications of PMA

Profile Management Applications (PMA) aren't uniquely associated to the introduction of OFDM or OFDMA signals in expanded spectrums. CMTS's and cable modems have been able to be configured for some level of adaptability to changing plant conditions for many years. These applications are extremely useful in managing data packet loss in the presence of impairments in the network. As the reliability and quality of the customer experience on broadband networks is becoming increasingly valued, their deployments are becoming more commonplace. These applications can adjust the modulation profiles of A-TDMA transmissions as well as OFDMA transmissions. While this is important to the experience of the customers in the degraded service area, it can add complexity to the operational priorities. Prior to the enablement of PMA, an impairment in the network that resulted in measurable packet loss would be identified using packet counters, or other intelligence tools, and would trigger a prioritized repair ticket for a system technician in a relatively short period of time. With the advent of these dynamic management applications, there are some operational considerations that need to be understood. For instance: 1) The poller timing of a monitoring tool(s). 2) The frequency of the PMA analysis and profile change recommendation. 3) Ticket creation and dispatch timing for a packet loss event.

These timing components need to be understood and synchronized to prevent scenarios where a technician is dispatched to an event before the profile management application has had an opportunity to mitigate the impairment. It is understood that the impairment still exists and may need to be identified and mitigated by a network technician, however, the synchronization of the event, mitigation and ticket creation will allow the business to properly prioritize the workforce to address the most impactful events first. The cable plant is a dynamic and ever-changing organism. Many events caused by transient noise in the upstream path, are moderate in severity and can be easily managed by the profile management application without ever sending a technician to affect a repair. There could be instances where the profile management application will intentionally not be engaged to downgrade profiles. Instances where the upstream interface has high utilization, or the profiles downgrade options have already been exhausted, are occasions where it is recommended to dispatch a repair technician to the node as a high priority repair.

5. Conclusion

As cable operators begin to use wider upstream spectrum on OFDMA channels and improve speed and bandwidth, spectrum activation can have challenges. In this first OFDMA PNM attempt, we revisited and applied some of our lessons learned with PNM and SC-QAM adaptive pre-equalization and were able to quickly find several nodes with multiple modems severely impaired and to locate some of the impairment causes. As we deploy mid-split and high-split to more nodes and customers, we will explore more PNM techniques, measurements, and channel configurations, such as wider channels, excluded subcarriers, cataloging ICFR signatures and impairment causes, and analysis along other measurements (Tx power, Rx MER, etc.), to serve us well for troubleshooting and repairing this new service.

Abbreviations

ALC	automatic level control
APIs	application programming interfaces
CM	cable modem
CMTS	cable modem termination system
dB	decibel
DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
HFC	hybrid fiber-coaxial
ICFR	In-channel frequency response
kHz	kilohertz
MHz	megahertz
Msym/s	megasymbols per second
OFDMA	orthogonal frequency division multiple access
PNM	proactive network maintenance
SC-QAM	single channel quadrature amplitude modulation
RF	radio frequency
RPD	remote PHY device
RMS	root mean square
SNMP	simple network management protocol
TFTP	trivial file transfer protocol
US	upstream

Bibliography & References

- Data-Over-Cable Service Interface Specifications DOCSIS 3.0 Operations Support System Interface Specification CM-SP-OSSIv3.0-I20-121113 (Cable Television Laboratories) [Note: Full band capture was first introduced in this version of the OSSI specification]
- Primer for PNM Best Practices in HFC Networks (DOCSIS[®] 3.1), CM-GL-PNM-3.1-V02-210114
- Larry Wolcott, John Heslip, Bryan Thomas, Robert Gonsalves, “A Comprehensive Case Study of Proactive Network Maintenance”, Cable-Tec Expo 2016, Philadelphia PA, September 26-29, 2016
- Alberto Campos, Eduardo Cardona, Lakshmi Raman, CableLabs Pre-Equalization Based Pro-Active Network Maintenance Model Version 3, September 5th, 2007

Redesigning the Voice Relevance Engine for Scale and Reliability

A Technical Paper prepared for SCTE by

Eileen Bengston

Principal Solutions Architect
Comcast
1700 JFK Boulevard, Philadelphia PA
(267)260-3370
eileen_bengston@cable.comcast.com

Ferhan Ture

Senior Director Machine Learning
Comcast
1325 G Street NW, Washington DC
(202)524-5060
ferhan_ture@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Current Voice Relevancy System	4
3. Redesigning Voice Relevancy System	5
3.1. External Interfaces	5
3.1.1. Queries.....	6
3.1.2. Responses	7
3.2. NLU and AR	8
3.3. Annotations, Synonyms, Ontology, and Testing.....	10
4. Conclusion.....	11
Abbreviations	11
Bibliography & References.....	11

List of Figures

Title	Page Number
Figure 1 - Basic NLP Process.....	3
Figure 2 - VREX Voice Transaction (Before).....	4
Figure 3 - VREX Voice Transaction (After).....	6
Figure 4 - Example Intent Pseudocode.....	7
Figure 5 - Core vs Experience Layers (Before)	9
Figure 6 - Core vs Experience Layers (After)	9

1. Introduction

At SCTE in 2012, a group from Comcast presented the Voice Relevance Engine for Xfinity (VREX). Its goal was to understand human utterances for video search and discovery and to take an appropriate action. Since then, Comcast's voice system expanded to cover numerous platforms, won a technical Emmy, and was used for more than 10 billion voice commands last year alone. After running at scale for 10 years, the voice system is being redesigned to suit the new and growing demands for voice control across a range of products.

The original system as described at SCTE in 2012 [1] was divided into three major areas: automated speech recognition (ASR), natural language processing/natural language understanding (NLP/NLU), and action resolution (AR). These services and the division of labor therein include: first determine what the user said, second understand what the user meant, lastly decide what the system should do to resolve the request.

Generic Voice Transaction

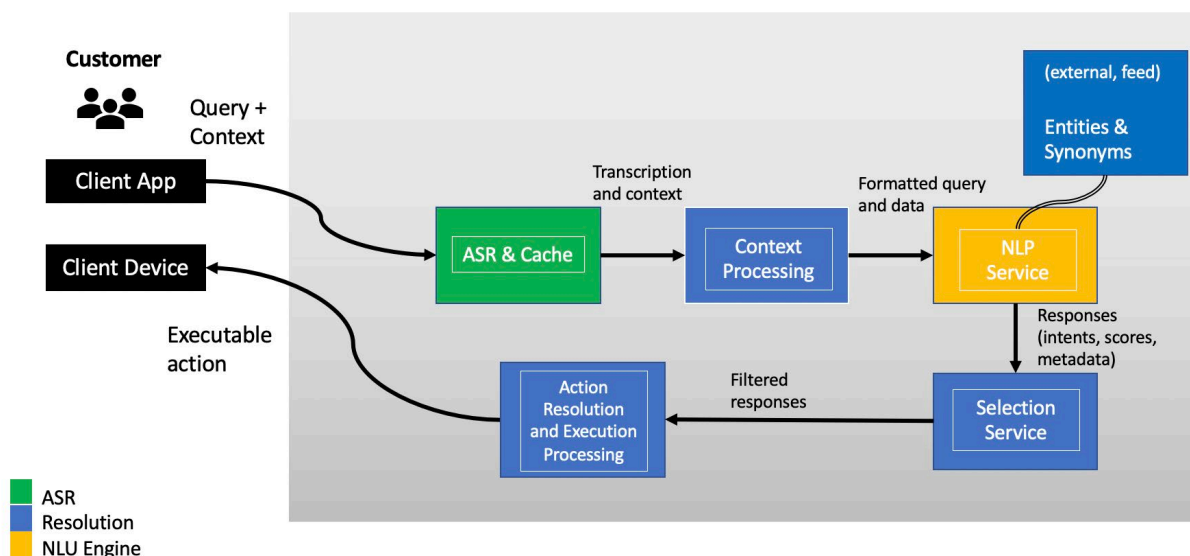


Figure 1 - Basic NLP Process

For the example “show me Adam Sandler movies” the ASR will transcribe the utterance to a text string “show me Adam Sandler movies”. Then, the NLP will attempt to ascribe meaning to the transcription. In this case it will canonically break up the string into <show me> an action, <Adam Sandler> a person, and <movies> a content type. In a simple system, Action Resolution involves the system choosing to provide a list of movies starring actor Adam Sandler. The information is provided in a way that can be displayed on the target device.

The above model is simple and linear. But what happens when there are several domains to cover? (A domain is an area of knowledge like video, customer experience, or home.) More information is needed to determine the appropriate resolution. Is the user interacting with a television? Are they walking in the front door? Are they talking to their car? Context as to when and where the utterance occurred affects

what the domain is for processing the request and what the action resolution will be. There may be several domains that could be applicable, so there can be multiple responses. Selection then must occur based on some quantification of how likely it is that the result is the expected solution. The Action Resolution phase increases in complexity as well because there is more to resolving the action. If the user is in their car speaking to the GPS, the Adam Sandler query should generate an error saying that it cannot find that road. But if the user is speaking to their television, the action resolution should result in a list of movies starring actor Adam Sandler being displayed on the television.

In practice, the VREX system grew organically to support a significant amount of complexity. The redesign assessed the current requirements for the system, those expected in the near future, and the complexity in deployment and support of the voice system. The voice recognition system was redesigned to improve the development, deployment, and servicing of the individual components of the system. These changes will be discussed as they affect each portion of the system.

2. Current Voice Relevancy System

Over time the number of domains, the number of languages, and the breadth of possible uttered commands has increased dramatically. So too, has the relative complexity of the system. The initial design supported speaking into a phone app to choose content to watch on a television set. The use case at design time was that a mobile device was paired to a television and all commands were for tuning content on the television. That model has grown to support not only tuning content but most television controls across smart televisions, voice remotes, hands-free devices, and mobile devices. It now covers many languages in multiple countries. In addition to video, the Voice System processes customer experience, home, sports, and other domain requests. Currently the model looks like this:

Anatomy of a Voice Transaction - Before

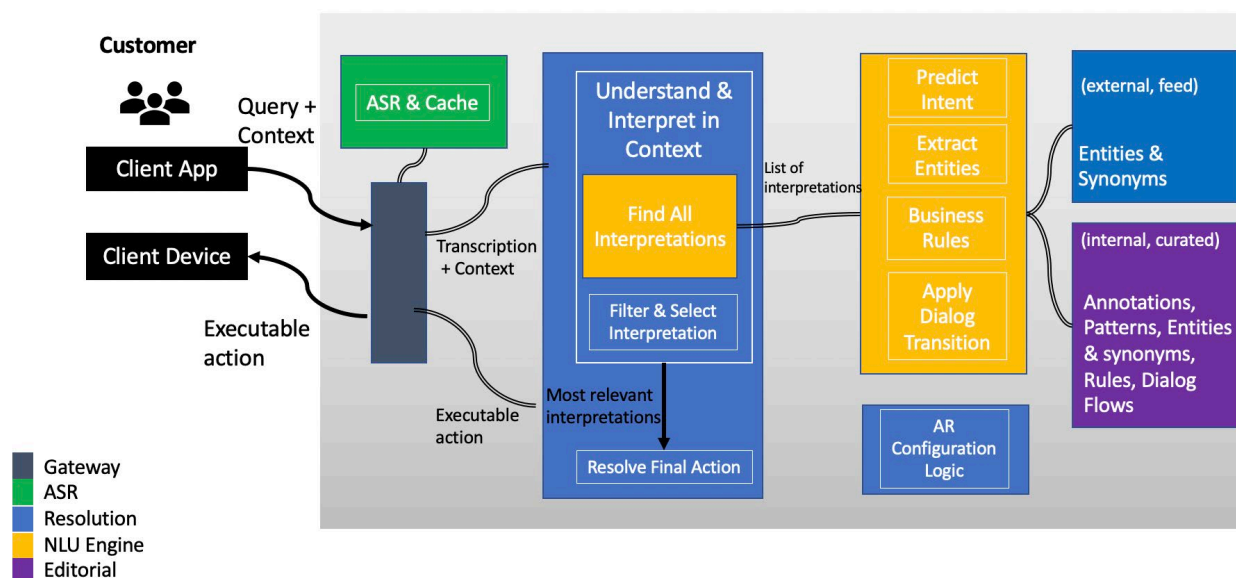


Figure 2 - VREX Voice Transaction (Before)

This drawing is a simplification of the model but includes several parts not included in the earlier simple block diagram. Specifically, it addresses the training of the model including the input of external feeds and the editorial, both of which are a significant piece in managing and training the NLU model. It also includes the configuration necessary to interpret the context of various domains.

Of course, this drawing also excludes several things. There is a significant amount of logic needed for routing, authenticating, and gathering context prior to context processing. While it shows a box for the gateway into the system it does not include the additional external calls to gather context and other needed information. And it does not include any of the testing and observability tools that aid the release and monitoring of the system. But importantly, this drawing adds a level of detail to context processing, NLP, selection, and action resolution. The structure of dealing with a voice command is far from linear. There are loops back to the decision engine with duplication of some common features. This level of interconnectedness, complexity, rigidity, organic system growth, and escalating support costs led to the redesign of the Voice Relevancy System.

3. Redesigning Voice Relevancy System

The goal in redesigning VREX was to improve the development, deployment, and servicing of the components that make up the voice recognition system. This does not mean that the model was simplified, in fact, the drawing below shows there are more components involved. But it allows better separation of work and modularity thus improving flexibility and decoupling and leading to a system that is more reliable and easier to maintain.

3.1. External Interfaces

The gateway for the VREX system had the single purpose of getting queries into the system and returning executable actions, but it grew organically over time and was expanded to cover several disparate cases. The gateway was responsible for directing voice queries to the ASR, gathering metadata from various points in the system (simultaneously to reduce delay), and passing ASR results, context, and metadata to the interpretation and resolution services. It was then responsible for passing the resultant resolution action back to the calling device.

3.1.1. Queries

Anatomy of a Voice Transaction - After

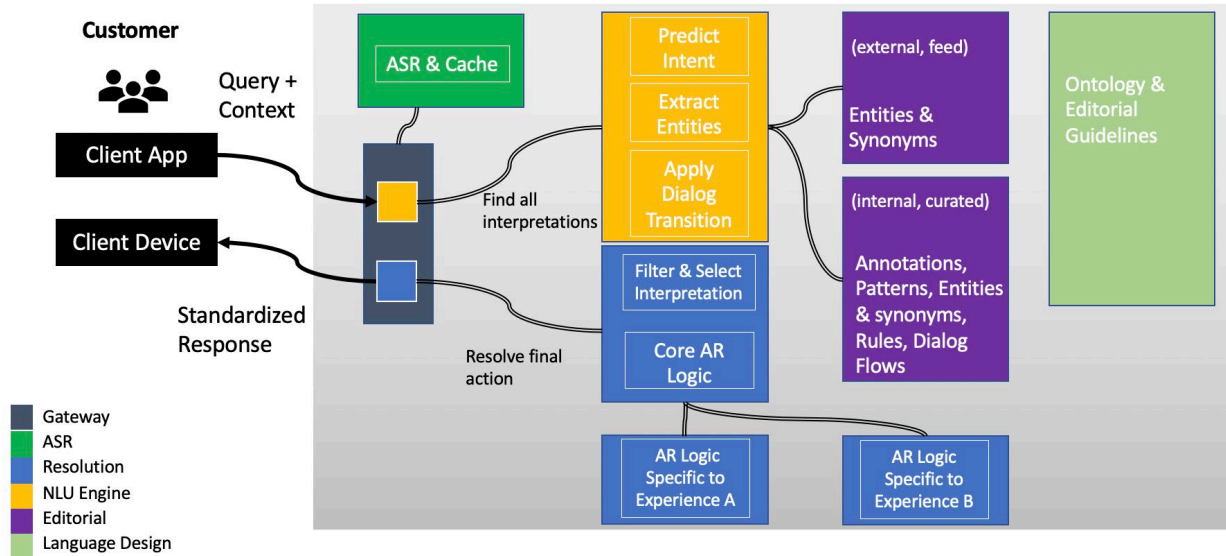


Figure 3 - VREX Voice Transaction (After)

As the client base using VREX expanded, so did the interaction patterns. Rather than just speaking into a mobile device, there was the voice remote push-to-talk, hands-free devices, and text-based chatbots. Some of these clients need ASR, NLU, and AR. But others need only NLU or ASR. The voice system was designed with a single input that gathered all information from multiple sources for every query. It allowed the client to specify the use of ASR, NLU, and AR, but it followed the same pathways regardless. Systems that used the older http requests were on a different, and older, version of the API and the teams had to keep multiple distinct APIs, and the stacks that support them up to date or else allow the older APIs to lag in functionality. It directed transcriptions to a dispatcher that did the domain selection based on the additional context, and answer selection based on the results from the NLU engine. This meant that every query entered the system, gathered all the external information, and traversed the same path through dispatch and interpretation regardless of what the query really required.

This web socket request is still used for voice commands, especially for session-based, multiple message communications. Commands requiring only NLU may prefer to use a lighter weight http call. Most importantly, the gateway can direct the request and so it can be passed where it is needed. If no interpretation is required, then that portion of the system can be skipped entirely.

Being more intentional in what the gateway does, and does not do, allows better use and routing of the services in the system. The client is now responsible for providing whatever additional context is required. The gateway's sole responsibility now is to maintain the contracts for routing information to the right service. It matches the usage of the product with the service and verifies the type of input and output expected. More calls go through the gateway, but it is skinnier and faster. It includes the API to client services, authentication logic, and configurable logic for dispatching messages based on the experience. It

can send audio and context to the speech engine, transcriptions, metadata, and context to the NLU engine, and logical forms to the resolution service.

3.1.2. Responses

A common way of communicating intents back to the client regardless of the experience in question was introduced to decouple the voice system from the experiences they support. The response messages were standardized because prior to the redesign, each client experience could determine what and how they wanted to receive responses from the voice system. This meant that the logical forms created by the NLU had to be translated for each client experience. Many used deeplinks directly to the resultant pages with the result that any resolution often had to be custom designed to support a specific client and changed when the client UI changed. By defining a standard interface, similar, but different from the logical forms created by the NLU system, the entire VREX system has a known language through which it communicates user intents and was separated from client-side changes.

In this case the action resolution provided general intents that include all the information that a client may need to determine the appropriate UI experience to present to the user. It is not prescriptive of the specific page to render, so the same structure can be used by multiple experiences and does not change if there are changes to the UI flow. A result of this is that some of the information is supplied in multiple formats, to allow this flexibility. For example, the rawQuery and the entity could be the same information presented differently for use by different clients. General pseudocode structure of an intent response is below.

```
{
  "type": "xrn:<domain>:<intent>:<receiver>:<action>",
  "intent": {
    "action": "",
    "data": {
      "rawQuery": "",
      "entity": {
        "entityType": "",
        "entityId": "",
        "entityAttributes": []
      }
    },
    "context": {
    }
  }
}
```

Figure 4 - Example Intent Pseudocode

Thin adaptation layers can then be created at the interface where needed. The VREX system is insulated from changes on the client side and the results are more deterministic because the patterns of their resolution and communication are the same.

3.2. NLU and AR

The system was designed for a single client and domain which meant that many of the services were tightly coupled. The speech recognition, NLU, and action resolution are distinct functions and yet, in understanding and interpreting, and doing so within the context of the request, those lines were blurred. This coupling caused cascading changes in the system. For example, when upgrading the ENLP (elastic NLP) which does the pattern matching NLP, the upgrade drove changes into the selection and resolution portion of the system. This was compounded when multiple experiences were supported as the coupling existed for every experience built. Thus, reestablishing the division between the responsibilities of the gateway (as explained above), the NLU, and the filtering and action resolution was needed.

As seen in the ‘Before’ VREX transaction (Figure 2), NLU related functions for finding interpretations was actually part of the resolution services. Everything that was needed to understand, interpret, and resolve actions for the domain was in one place. But this meant there was tight coupling between the interpretation and resolution services, and the NLU engine to find all the interpretations and resolve them. It also meant that everything had to pass through this interpretation even if action resolution was not required

Re-establishing clear boundaries between services within the VREX system meant separating those functions. The NLU logic exists separate from the interpretation and action resolution logic. The gateway contains the per experience information needed to dispatch the queries appropriately. The NLU receives the query text, context, and any metadata it needs from the gateway, and passes back logical forms. The interpretation and resolution services receive logical forms and select answers, determine resolution, and return standardized responses on a per domain basis. The two services, once again, focus on their core responsibilities and communicate only through the defined interfaces in the gateway.

At the same time, the layers of the system were restructured to separate the experience business rules from the core functionality meaning that the individual pieces that need to be built for a specific experience are thinner, allowing for better reuse. Core logic is that which is needed for turning any voice command into an action. Business logic is the configuration and function that is needed to create a meaningful action in each environment.

Before

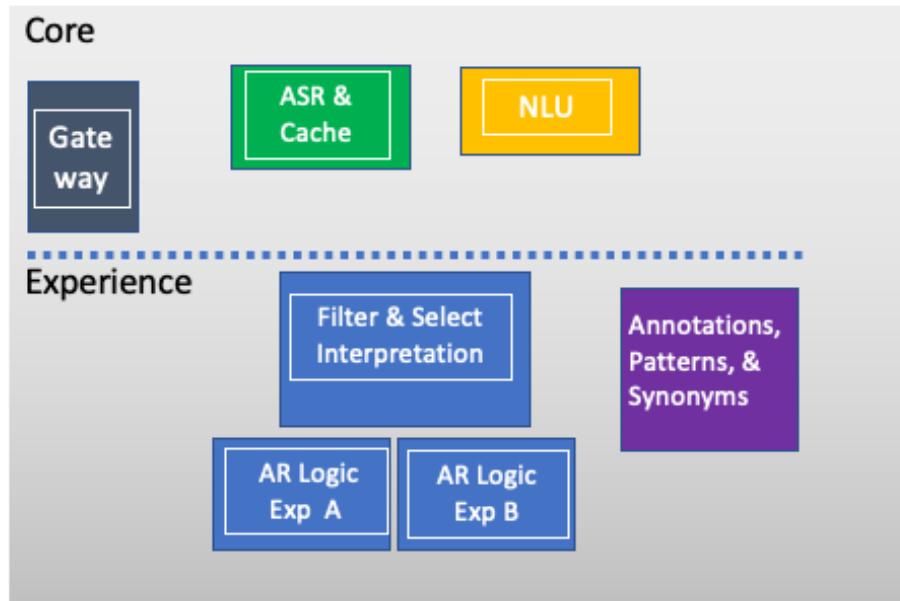


Figure 5 - Core vs Experience Layers (Before)

After




Figure 6 - Core vs Experience Layers (After)

In the above drawings, before the redesign only the ASR and the NLU services were core functionality, and even they had some configuration within the services that required redeployment when changed. Most of the training data input, interpretation, and resolution logic was linked to the specific experience, so each new client or environment involved significant design and deployment resources.

With the redesign, the core filtering interpretation and selection, and the ontology and editorial guidelines, became reusable core functions. This separability of the custom logic from the core logic increases the complexity of the model but greatly reduces the unique work for development and deployment of a specific experience. This then allowed the separation of understanding a query from filtering and selecting the resolution action. This is important because understanding the queries depends on learning processes in the NLU engine and its prediction of the intents based on the appropriate rules. These change only when the model information - the feed, the annotations, or the feedback - change. Filtering and resolution depend on several factors. The first is the fundamental logic for selection and filtering. Understanding what ‘play’ means is near universal. But how to ‘play’ depends very much on the experience in question. The goal here was to, as much as possible, separate that which is universal, the meaning of an utterance, from how to address that user intent in each environment.

Reduced complexity and interdependency resulted. Once most of the specific configuration and logic was pushed as far to the edge of the system as possible, the surface area that needs customization for new experiences was reduced. The entire process of supporting a new function, a new client, or a new experience has a much shorter development time.

3.3. Annotations, Synonyms, Ontology, and Testing

To improve the consistency of the system, a set of ontology and guidelines were specified to simplify the patterns for the editorial work that is used to enrich the dataset. Fallback patterns from optimal behavior to acceptable behavior improve both the flexibility and the resiliency of the system. Finally, each service was shifted to use to a common logging, tracing, and debugging system.

As a system grows it gathers a lot of information. The ontology redesign found several ways of representing similar concepts that were being supported simultaneously. The team took the time to define an ontology and guidelines for the editorial additions that enrich the dataset. This involved combing through the existing rules to define reusable patterns, removing duplication, and designing a process of intro, request, update for creating patterns and entering synonyms. This makes for fewer, more consistent patterns, and makes the job of an editor easier because most additions will fit into a series of well-defined patterns. The ontology changes included using labels that create flexibility for ranking importance of certain structured utterances.

With more commonality in patterns and with the separation of the experience specific logic, and its push to the edge of the system, some behavior can be defined on an experience basis without affecting the core logic.

Additionally, while not expressly mentioned previously, any system is going to have a considerable amount of logging and debugging code. Because of the modular nature of the VREX system a lot of this was done by each individual team in the way that best suited them. But this made it disjointed and harder to debug the voice system end-to-end. The re-architecture was an opportunity to take the best and most useful practices across teams and consolidate them into one structured process including our own testing system and ELK (Elasticsearch and Kibana) cluster for common observability. Note that Comcast collects, stores, and uses all data in accordance with its disclosures to users and applicable laws.

4. Conclusion

VREX was presented exactly 10 years ago at SCTE as a novel way to perform voice resolution for the video space at scale. A large-scale system can be expected to have experienced a lot of organic growth and change over a decade. As systems grow organically to suit changing needs, they often take on unexpected forms.

The Comcast team determined its needs as greater flexibility, faster development time for new features, faster deployment time for new customers, and easier and more consistent support. They then redesigned the system, re-defining the APIs, service boundaries, patterns, and ontology. This fundamentally changed how the services can be called and how they respond. It changed service boundaries to separate core functionality from experience driven code, made configuration handling common, and reduced the scope of the experience layers. Finally, they simplified the patterns and ontology that is used to train the model and make decisions. This redesign, while still in the process of being implemented, has made the development and support of the VREX system faster, and more scalable, by reducing the scope of work needed to support an experience, increasing the reusable code, and isolating the code in the system that needs to be modified closer to the edge of the system.

Abbreviations

API	Application Programming Interface
AR	Action resolution
ASR	Automatic Speech Recognition
BIO	Beginning-inside-outside
ELK	Elasticsearch and Kibana
ENLP	Elastic Natural Language Processor
HTTP	Hypertext Transport Protocol
NLP/NLU	Natural Language Processing/Natural Language Understanding
SCTE	Society of Cable Telecommunications Engineers
UI	User Interface
VREX	Voice Relevance Engine for Xfinity

Bibliography & References

[1] *V-REX – Voice Relevancy Engine For Xfinity*, Stefan Deichmann, Oliver Jojic, Akash Nagle, Scot Zola, Tom Des Jardins, Robert Rubinoff, Amit Bagga, Comcast Labs, SCTE 2012.

Reducing Investigation Time for Researchers, And Enabling Automated Configuration Updates by Digitizing Contextual Information

A Technical Paper prepared for SCTE by

Ajay Gavagal

Data Solutions Architect
Comcast Corporation
1800 Arch Street Philadelphia PA 19102
215-286-3078
Ajay_gavagal@cable.comcast.com

Mehul Patel

Distinguished Architect
Comcast Corporation
183 Inverness Dr West, Englewood, CO
303-658-7826
mehul_patel@cable.comcast.com

Sinan Onder

Vice President
Comcast Corporation
1800 Arch Street Philadelphia PA 19102
267-260-0964
sinan_onder@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Personnel & Motivation	4
3. Defining Contextual Information.....	5
4. Day 1 vs. Day N Scenario	6
5. Proposal/Solution	9
6. Conclusion.....	14
Abbreviations	14

List of Figures

Title	Page Number
Figure 1- Sample flow of telemetry data and personnel involved	3
Figure 2- Exchange of contextual information across personnel.....	5
Figure 3- Alerting application deployed with contextual data within configuration.....	8
Figure 4- Cascading failure due to non-digitization of CI	8
Figure 5- Contextual Data Object	11
Figure 6- Buildup of domain contextual data object.....	12
Figure 7- DCDO enabling a full cycle of automation.....	13

List of Tables

Title	Page Number
Table 1- Sample CI and probable storage systems.....	6
Table 2- Sample list of elements in CDO's	10

1. Introduction

Contextual Information (CI) is an asset to every enterprise for its digitization to be achieved end-to-end. The digitization of contextual information and its changes can build a sustainable growth engine for development of new products and services. It can lower cost of software changes and lead to high productivity. To explore how we can achieve digitization of contextual information, in this paper we start with a scenario, explain the personnel involved, the kind of questions that get asked by these personnel, explain the problem that is in exchange of contextual information, and finally provide a possible solution for it. In this paper we also will explore the mechanics of contextual information and how it can benefit small or large use cases for data analysis and machine learning.

We start with a simple question that requires analysis into telemetry data, investigate the journey of how that question gets answered within an enterprise. Let us assume there is a device “X.” This device is part of the internet protocol (IP) network. This device also has a capability to provide consistent telemetry such as its state. State here refers to the overall condition of the device, for example, is the device online/offline, and if the device is offline then reasons for being offline such as error conditions. Given this telemetry data, business owners can ask specific questions that can help them in business impacting decisions. Such as, evaluating device models from various vendors. A sample question might be to find out if a specific device model breaks down more than others. Given the value that can be achieved from answers to such questions using telemetry data, leadership allocates time and resources to capture data from device “X,” ingest its telemetry into a storage layer and then appropriately make that data available for use cases such as alerting, analysis, and machine learning.

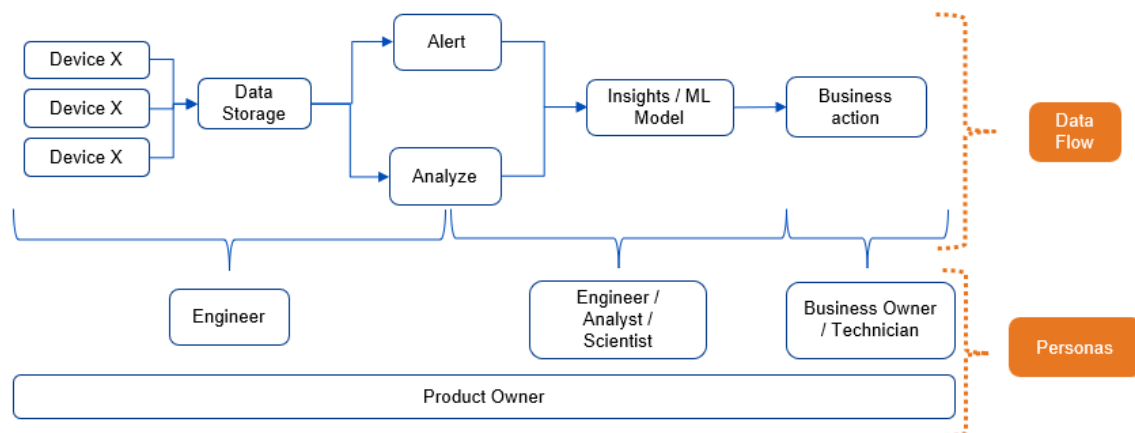


Figure 1- Sample flow of telemetry data and personnel involved

Note: Arrow direction represents data flow.

In the diagram above, which is common flow of data across enterprises, we see five different personnel getting involved in ingesting, storing, alerting, analyzing, and actioning on the insights. In the section below, each of the five personnel are analyzed. Moreover, and some of their motivations behind enabling insights from this data is highlighted.

2. Personnel & Motivation

Business owner(s) are motivated to provide high reliability in the overall service offered to customers. Here device “X” is a critical part of the IP network and should be made reliable and stable. Any disruptions in service due to device outages can negatively impact the customer experience. By choosing the right device models with the highest reliability they look to reduce disruptions. Their eventual goal is to have high customer satisfaction by providing reliable service. Business owners also often have limited budget to achieve the previously mentioned goals. So, they look to strike a fine balance between the highest reliability possible given their allocated budgets.

Product owners are looking to gain value through telemetry data and look for opportunities where the products or features they own, can save cost, or improve productivity. Product owners typically get measured on customer satisfaction, and hence want to see their customers succeed. To ensure customer satisfaction, the speed of deploying features is one of the many criteria’s they focus on, in designing and engineering features. In large enterprises, due to pace of evolution, in many instances this criterion can become the highest priority, while other criteria can take lower priority.

Engineers are specifically interested in maintaining a reliable stream of data to consumption platforms with low latency. Any disruptions in data pipelines are a disruption to data flow where engineers are called upon to fix the issue. Engineers also must optimize data pipelines for low cost of computation while providing maximum data resolution. Given these opposing set of goals, namely cost vs. data resolution, they must make decisions on the resolution of data to fit either compute or storage requirements.

Data consumers on the other hand are keen on getting data that is reliable and of high quality to help them ease their job. Data cleaning is usually a large part of their projects. Given the velocity, variety, and volume of data available within an enterprise, it is safe to assume that data cleaning is a regular routine. However, data consumers are typically/usually demotivated by data cleaning since they view this as time taken away from more exciting tasks, which is to find insights in the telemetry data that provides an accurate representation of the real world.

Finally, network technicians want to see lower repeat issues. They want to ensure lower number of disruptions in the network. They depend on meaningful insights that help reduce avoidable maintenance and repairs. They are usually demotivated with technology and process that is more of a hinderance to their work than assist them, for example: cases where device recalls are made.

Below is a representation of the personnel and exchange of contextual information mentioned above:

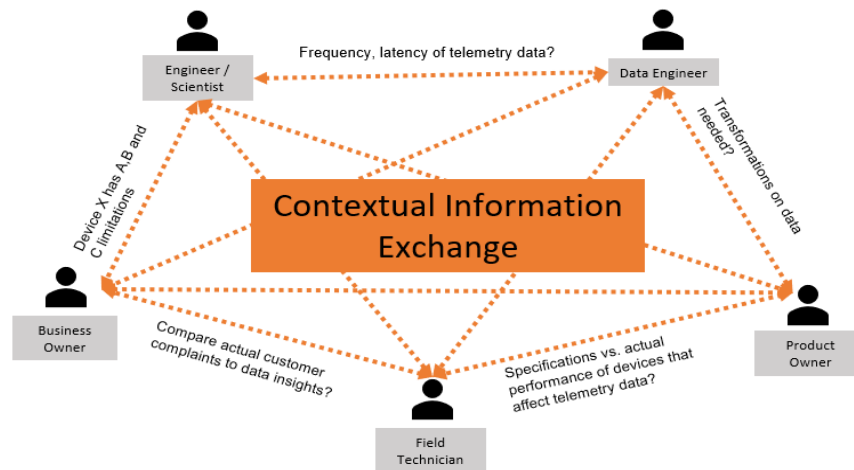


Figure 2- Exchange of contextual information across personnel

Now that we have seen the scenario and personnel involved, let us understand what we mean by contextual information and which of the above personnel might be responsible for such information and possibly where it might be made available in an enterprise.

3. Defining Contextual Information

We start by defining what we mean by contextual information. Contextual information can include but not limited to any metadata that provides context and perspective around telemetry data. For example, application configuration information used to set up the data pipeline, inherent limitations of data, data lineage information such as source and target systems for the entire data pipeline chain, data catalog such as list of data sets available, schemas and their versions, and many more. To better explain this with an analogy, consider contextual information as controls in the control pane of a manufacturing line and the telemetry data as the products being manufactured through the machinery itself.

Continuing with this analogy, in a manufacturing line, the controls (switches and dials) are usually set to certain values to ensure a constant production of the product. To either increase or decrease production that meets demand the demand, these dials in the control pane need to adjust. If the controls are not digitized and automated, then a human intervention is required every time to increase or decrease production. Also, considering the changes in the control values, they are not frequent, but too rare either. If these changes in the control values are not propagated or communicated to the entire manufacturing line, as an example to the packaging department or to the inventory department, the whole manufacturing line fails.

Similarly, consider an application that polls devices in the IP network at every A mins (frequency) to receive B resolution of data. In this case the telemetry data polled by the application is the product, while the polling frequency and the resolution detail polled are the controls in the control pane. If changes to the polling frequency or the resolution of data is not communicated by the application to downstream consumers, the whole data pipeline fails, causing cascading failures.

By nature, Contextual information is usually distributed in multiple systems. These systems may or may not talk to one another. Issues that prevent systems from communicating with one and other can include but are not limited to: mix of legacy and new applications, varying hardware/software platforms, varying cloud environments and varying feature capabilities. Let us consider the first example of telemetry data

from device X discussed in this paper and see how contextual information may be distributed among many systems.

Device vendors publish object identifiers for each device property and define functions, attributes in libraries allowing customers to consume and ingest such information. When telemetry data from the device is ingested, the schema of the data is registered in a schema registry to ensure changes/versions are tracked for field names, descriptions, data type changes. Information such as source, target, transformation logic and resolution of data are stored in application configuration. If data takes few forms such as raw, transformed and aggregated with retention periods increasing with lower resolutions of data then such information might be stored in the configuration of the application as well or in a catalog. Finally, the data catalog could be from well-defined tables or a simple list of storage paths once data is processed and written to a location. The ingestion application is registered in a service catalog that is maintained across the enterprise. Such a service catalog would contain information like, application name, purpose, source, destination, developers, support personnel etc.

Table 1- Sample CI and probable storage systems

Contextual Information	Probable Storage System
Object identifiers, descriptions, limitations, values etc.,	Libraries published by device vendors
Data resolution, polling frequency etc.,	Application configuration
Quality of datasets – raw, enriched, transformed, aggregated etc.,	Wiki pages, word documentation, schemas within DB's, segregation by paths in an object store
Data catalog	Schemas within DB's, segregation by paths in an object store, wiki pages
Application information	Service catalog
Application code and versioning	Version control, code repository
Change information	Tools that support CI/CD such as Concourse

4. Day 1 vs. Day N Scenario

Now that we have outlined the scenario, personnel with their motivations, defined CI, its nature, and the systems/applications where it might be stored, let us have a look at how a data consumer looking to generate alerts on telemetry data from device X discovers CI on day 1 (i.e., when designing the application) and day 2 to day N (when the application is deployed and needs to be maintained).

Initially a “discovery” phase to scrub, understand, document, if necessary, all the above-mentioned contextual information from various systems is required. We will discuss later why it is important to reduce the time involved in this phase. Since the focus initially is on generating alerts atop of telemetry

data from device X, the engineer might take the below attributes into consideration when designing an alerting application:

- Source of device X telemetry data
- Target to provide alerts once processed
- Various conditions within the telemetry data
- Values that signify these conditions
- Any exceptions to the conditions such as occurrences of NULL values or random values
- Applications that provide enrichment information to values within telemetry data (dependency)
- Data types of values
- Frequency of source data & target
- Volume of data at source & target
- Latency of data at source & target
- Modifications/transformations done to values to make them human friendly. For example: consider state information containing '1' for online and '0' for offline or vice versa. To make this more human friendly, one might be converted to "online" and 0 might be converted to "offline" to remove ambiguity for the data consumer
- Source and destination schema of the data

In a perfect world, all the above contextual information and their source systems are digitized and work cohesively (without any linkage breaks) to make the engineers' life easy to consume such information programmatically and design the alerting application in a fully digitized manner. However, it is far from reality since a lot of this information may not be digitized for easy consumption, for example, there might not be easily accessible application programmable interfaces (APIs) that provide information on all the modifications done to data values to make them human readable or a programmable interface that provides the frequency and resolution of the polling application. CI might be available within configuration of applications (as code) or free text documentation or audio/video clips (when recorded as training videos). In such situations, scrubbing through just free form information during discovery phase and then building configuration objects that contain contextual data and using this in alerting applications might be the accepted norm, but is not ideal and does not achieve digitization. This rapid prototyping and is also known as fast go-to-market strategies to justify the acceptance of non-digitization of contextual information. As previously discussed in the section, large enterprises have deadlines, priorities to meet, especially during the development phase since it costs time and money.

Since some of these CI inputs arrive from data discovery for the engineer to use them for building configuration objects manually for the alerting application, we have a breakage in the chain of digitization on day 1 (day 1 here is referring to the application deployed in production) itself. We are interested in showing how using non-digitized CI from day 1 causes a cascading set of problems on day 2 to day N (refer to figure 4). So let us assume that the engineer has configured required CI and successfully deployed the application that generates alerts atop of the telemetry data.

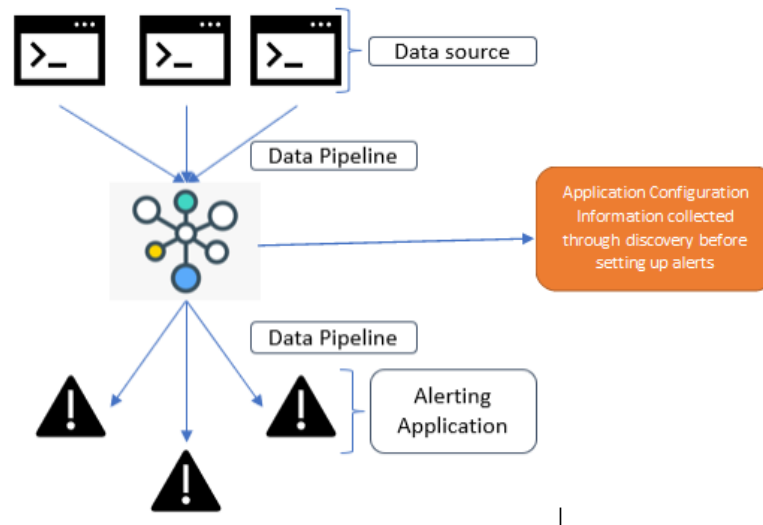


Figure 3- Alerting application deployed with contextual data within configuration

On day two, let us assume there is now a machine learning model set up to consume alerts from the alerting application. Now we have a two-layer dependance on the non-digitized contextual information.

Now let us also assume there is change in multiple attributes of the contextual information such as change in schema (data type changes), source system change, destination system change, data value changes (such as two or more values combined to one or a single value split into multiple), change in frequency of data, change in resolution of data etc. Since CI elements were not fully digitized on day 1, these can be treated as breaking changes for all downstream layers.

Below is a representation of cascading failure can occur on when CI is not fully digitized and transmitted.

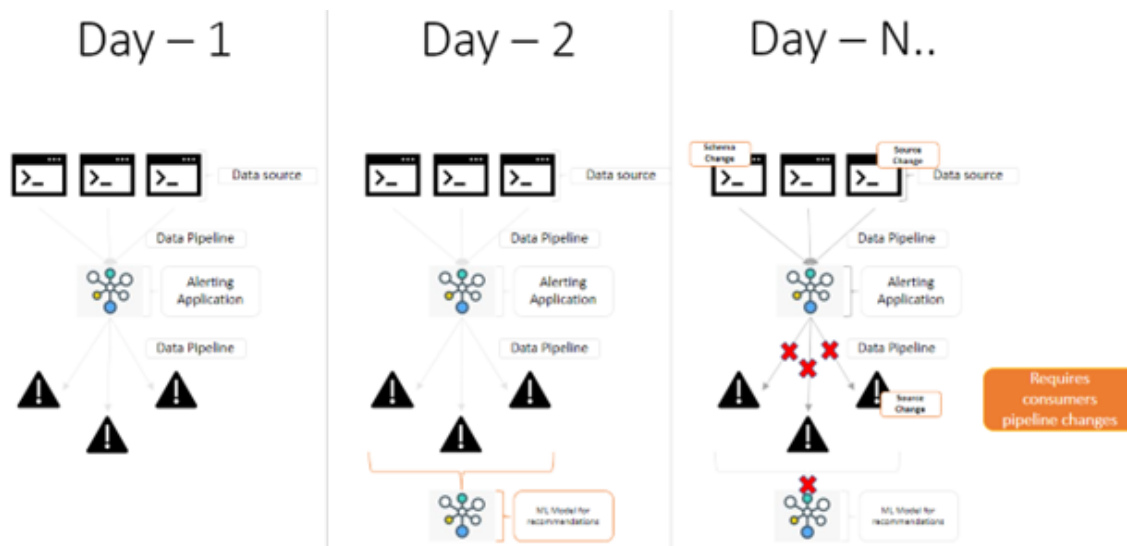


Figure 4- Cascading failure due to non-digitization of CI

If we do not address this problem of non-digitization of CI but kept adding more layers downstream that sends notifications to customers based on recommendations from the machine learning model, this cycle continues until, it collapses due to unmanageable changes in CI.

To give a scale, if there are A number of measurements available from a device going through B number of CI changes and there are C number of dependent data consumers downstream, then we can assume that there would be $A*B*C$ number of changes to manage or intervention due to non-digitization of CI.

This problem only exacerbates if engineering groups are distributed between domains in an enterprise and reporting into various leaders. Groups having many individuals working on a single project with individuals distributed in various geographies. Then there are problems that the enterprise is not in control of such as vendors making devices smarter, software/hardware changes etc., the list goes on. And finally, the growth in volume, variety and veracity of data generation does not help either, where we are having to re-think our traditional well-defined storage formats with well-defined meta data stores to a more loosely bound structure where metadata is distributed all over the place.

If each time CI changes need to be reviewed, handled before re-deploying applications, then it is both very human centric and time consuming. This process is highly error prone and leads to lot of wastage in money and lowers productivity.

5. Proposal/Solution

One of the first things we observe as we review the list of CI elements is that it is diverse, sourced from a lot of systems/applications. Remember CI is meant to provide context and perspective to the telemetry data and its use and hence it includes everything that addresses this requirement.

With the evolution of software, we have an exceptionally good understanding of metadata needed to build applications, microservices and event-driven structures. However, we are lacking a standardized framework of required CI elements, its storage and transmission when cutting across pure software development activities and into more of the realm of data analytics/science activities that prototype and launch machine learning (ML) models that in turn assist software development. The CI elements of interest are varied between these two kinds of activities and hence an encompassing standardized set might be a good one to have to start with. We do not have a standardized list and that is for another day and another research paper. But in the absence of such a list, how should we produce one to address the problem in the short term at least?

One way is to see the kind of questions that get asked to an engineer/analyst/scientist when they present findings:

- What is the source of this data?
- Why did we not use an alternate source to answer the question on hand?
- What does a value mean? For example, in case of state information from a device X, what does online really mean? What does it signify?
- Do we know what is the source of truth between two similar datasets from two different applications?
- What assumptions have been made when generating an alert/analysis/ML model on telemetry data?
- What is the built-in latency into an alert/analysis/ML model? Can we reduce it?
- What thresholds were used to generate this alert?

- Who is the owner of this data?

The personnel asking these questions encountered doubt in the alerting mechanism, or analysis generated or the recommendations from the machine learning model. Based on this, it can be observed that this is just non-digitization of CI manifesting as the problem. The question is not asked to see if answer is known, but it is purely to get more context and perspective on the data. This data could have been transformed from pure raw telemetry to an alert, analysis, or an ML model that is helping drive the insight to help make a business decision.

With that in mind, let us start with a list that solves for these above questions and then build from there. It should be noted that this is not a comprehensive list of CI elements and should not be treated as one. We are merely trying to answer the above questions and providing a means to solve the problem. If these questions are not of high priority within your enterprise, then this list should be reformed to fit your enterprise needs.

Table 2- Sample list of elements in CDO's

Sample CI standardization	Description
Source system	Refers to the device, application etc., providing the telemetry data
Target system	Refers to the storage layer, application etc., where data is persisted after transforming the telemetry data
Data Frequency, Data Resolution etc.,	Explaining how frequently fresh data arrives, how deep can the data go etc.,
Data lineage	The source system/application where the data originated, got transformed, stored etc.,
Data description	Describes the telemetry data and can include field descriptions, value explanations, value limitations etc.,
Data schemas	Structure of the data
Data catalog	If more than one stream in the telemetry data, then list or catalog of those datasets
Exceptions	Exceptions when data might not be transmitted, transmitted with errors etc.,
Transformations applied	Changes to the datasets applied between source to destination

Next, we discuss the approaches we can take to solve for CI list to be transmitted through the entire lifetime of the data within an organization i.e., through various domains:

1. Ensure all contextual information is documented, distributed at regular intervals
 - Time consuming, human intensive and error prone if not for impossible.

2. Ensure all applications transmitting data provides interfaces for not only data, but also contextual information about the data (including metadata)
 - Well defined micro services powering every application within the organization with great programmable interfaces that data consumers can interact with to gain contextual knowledge, we can deem this problem solved.

However, this is not always the case since there will be standalone applications that do not have interfaces but are vital within enterprises. These applications came about as prototypes or as legacy service. Information technology (IT) is no more centralized where one central IT organization sets standards. The distributed nature of modern IT ensures that applications/systems are always built in a diverse way.

Start by bundling everything we determine as necessary contextual information required for the application in question and for downstream applications and build a contextual data object (CDO) and domain contextual data objects (DCDOs). Store and transmit this CDO in a way that it can be easily queried for changes. Also ensure that DCDOs can be enhanced by multiple application owners starting from the data producer, all the way to the last data consumer. Our solution proposal resolves around this third option. This is not necessarily new, since we have various flavors of this solution in usage within enterprises today, but the semantics of how we implement this might need another look.

A typical set of personnel working on making telemetry data smarter through alerts, analysis, or ML models, are distributed in various parts of the organization from functional, hierarchical, and geographical perspectives. Given such a distributed workforce, it would be wise to ensure that we start building the CDOs in bits and pieces throughout the domain and allow each data producer and consumer to decide the kind of CI they would like to add into the CDO. And when the CDO needs to be moved between domains, ensure that domain contextual data objects are used to transmit such information.

We see a way that contextual information can be digitized into an object that gets transformed as it moves along the enterprise systems and applications. First the data producer might produce a CDO as below:

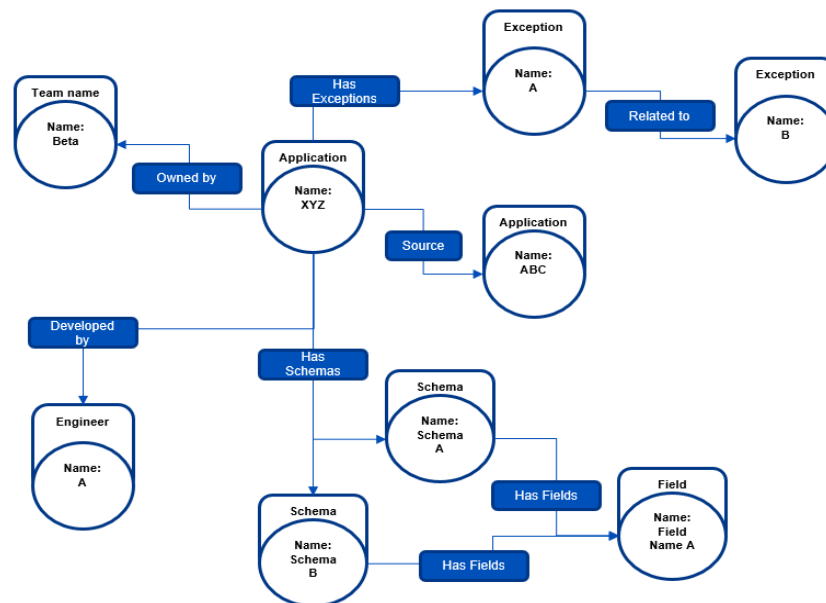


Figure 5- Contextual Data Object

Remember the above CDO is entirely up to the data producer to determine what to populate and what to leave out. In this paper we only define a common set of elements part of the CDO that is available to the producer to add into the CDO as part of the standardization. We want to ensure that this CDO is used internally by the data producer to drive changes to their application. This is an especially crucial factor, since self-use is the best motivator for the data producer to keep the CDO up to date and can help drive automation.

Now as this CDO makes its way through the organization to say, another engineering team that is in a different domain (in this case let us assume the elements in CDO and DCDO are same), that adds alerts atop of the “field name A.” The engineering team enhances the DCDO with alert contextual information.

Observe that each of the DCDO elements connected to one another through a relationship. For example, the source data in application XYZ is FROM the application ABC. This signifies lineage information. The application XYZ was developed by team Beta and the engineer A was responsible for its development. Application XYZ has two exceptions named A and B, while B is related to A. Application XYZ also has two schemas: A and B. Both these schemas have common field A.

Below is a representation of how a DCDO might evolve as it makes its way through the enterprise. All orange nodes and edges are added by a different domain:

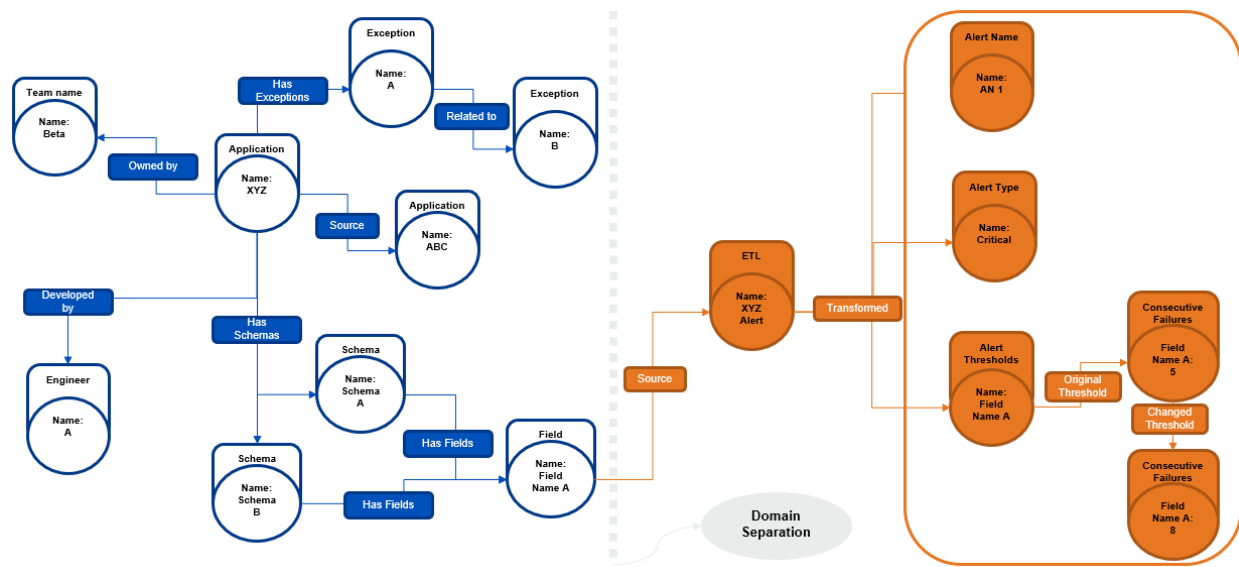


Figure 6- Buildup of domain contextual data object

Input source to an extract, transform and load (ETL) job that performs the function of generating the alert has source XYZ and transforms the data to alert name “AN 1” that has an alert type “critical” with the thresholds set as five consecutive failures originally. The team decided to change the alerting threshold to 8 at some point, which can just be another node signify the date of change.

We can also observe that application that performs the ETL is linked to the source team name and the engineer. This is immensely powerful. This is now helping connect elements that were not intended to be connected in the first place and allows for a lot of questions to be answered. This is the power of having DCDOs.

Expanding on the topic regarding the representation of DCDO and its evolution through the enterprise, an observation can be made that the nature of this data has an entity relationship during its formation. Given this nature of contextual information, two entities (nodes) being linked by a relationship (edge), using a graph storage mechanism for CDOs and DCDOs might be a good idea. As DCDOs evolve we start to build a knowledge graph across the organization giving the consumer access to powerful information that is hidden from many users.

The solution does not stop here. To help data producer teams and the consumer teams to create, interact with CDOs & DCDOs and build on it, we must provide them easy access ways to interact and modify them. Or else, we are just moving the breakage in digitization to a later point in time.

This is only the beginning of the solution. CDOs are built to assist the automation of functions within applications. When this automation controls are needed to be handed over to another domain, we need DCDOs or domain contextual data objects. These DCDOs can be directly used within ML models to ensure that recommendations from the ML models powers the application through automation vs. intervention.

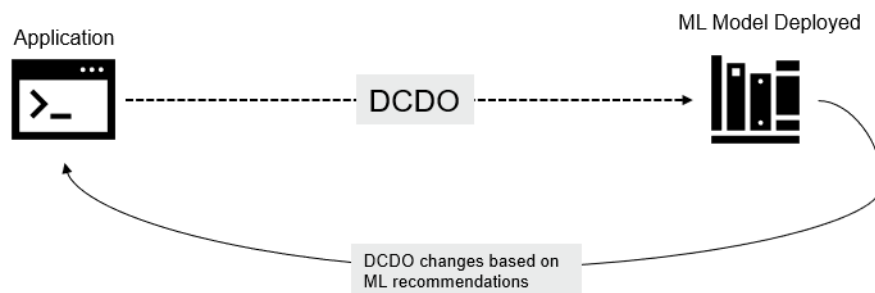


Figure 7- DCDO enabling a full cycle of automation

We should also enforce standardization of CDO/DCDO elements that the enterprise deems necessary and ensure this is adopted from the start for all new applications being developed. This is the hardest part since enforcement of a norm is hard within distributed structures unless it is used by the application.

The way to bring domains to onboard and share their CDOs as DCDOs into a common repository like an open WIKI within the enterprise is to help domains understand the time and cost savings from such an effort. The questions addressed earlier may seem trivial at first, but add in the mix of employee attrition, accidental deletion of information, modifications done for an ad-hoc request, and then having to invest time in discovering all this, the value will speak for itself.

6. Conclusion

Contextual information is vital and should be treated as an asset that needs to handle within an enterprise. Changes in CI is even more vital to enterprises and their propagation through DCDOs is essential for sustainable interaction between domains.

In this paper we first outlined a typical scenario for ingesting, persisting, and acting on telemetry data to derive alerts, insights, and ML recommendations. We also observed personnel and their motivation when interacting with telemetry data and its context. We saw how a cascading failure can occur when CI is not digitized and managed for propagation into downstream systems. Finally, we observed how digitizing CI in the form CDOs and DCDOs not only enables rapid transfer of knowledge between humans but also provides a methodology to handle updates programmatically in fully automated systems.

Digitization of CI also builds resiliency in automation against external factors not under direct control of the organization such as vendor device decommissions and recalls. This is also the case for applications where CI reside needs to evolve at a fast pace. Without effectively tracking changes to CI, models that power ML and automated systems built atop these assumptions can be ineffective. Such ineffectiveness in ML models lead to low return on investment (ROI) and sometimes even negative ROI. Assumptions underlying ML change rapidly with the evolution of network objects, applications, platforms, infrastructure, data pipelines, storage mechanisms and application configuration information. Digitizing CI and storing for easy access, ensures changes can be discovered programmatically by automated systems without intervention. Although we are not entirely solving for end-to-end digitization of all CI in an IP network, we provide means to show how it can be done.

Finally, imagine the amount of time required by engineers, analysts, scientists in discovering contextual information every time a change needs to be made. It is both wasteful and lowers productivity. Digitizing CI as mentioned in this paper by starting to create CDOs for internal use by application owners and DCDOs when transmitting to different domains saves money and time.

Abbreviations

API	application programming interface
CDO	contextual data object
CI	contextual information
CI/CD	change integration and change deployment
DCDO	domain contextual data object
DB	database
ETL	extract, transform and load
IP	internet protocol
IT	information technology
ML	machine learning
ROI	return on investment

Requirements for the IoT Infrastructure in the Customer Premises

A Technical Paper prepared for SCTE by

Rajesh Abbi

Principal Consultant
Duke Tech Solutions, Inc.
111 Fieldbrook Ct. Cary, NC 27519
+1 919 455 4787
rajesh.abbi@duketechsolutions.com

Charles Chapman

Manager of Customer Training and Loyalty
Enersys
3767 Alpha Way
Bellingham, WA. 98226
+1 941 228 5102
chuck.chapman@enersys.com

Sudheer Dharanikota

Managing Director
Duke Tech Solutions, Inc.
111 Fieldbrook Ct. Cary NC 27519
+1 919 961 6175
sudheer@duketechsolutions.com

Kyle Haefner

Lead Security Architect
CableLabs
858 Coal Creek Cir, Louisville, CO 80027
+1 303 661 9100
k.haefner@cablelabs.com

Clarke Stevens

Principal Architect, Emerging Technologies
Shaw Communications, Inc.
1801 N. Broadway, Suite 501
Denver, CO 80202
+1 720 723 2316
clarke.stevens@sjrb.ca

Table of Contents

Title	Page Number
1. Executive summary	3
2. Evolution of cable customer premises	3
3. Emerging IoT applications	4
4. Special needs of IoT applications	6
4.1. Location needs	6
4.2. Security needs	6
4.3. Reliability needs	7
4.4. Timing/Latency needs	7
4.5. Power needs	7
4.5.1. Power sources	7
4.5.2. Backup power considerations	9
4.6. Interoperability needs	9
5. Opportunities for cable operators	9
6. Wiring closet concept	10
7. Wiring closet requirements	10
7.1. Customer premises location requirements	10
7.2. Networking requirements	11
7.3. Security requirements	11
7.4. Powering requirements	11
7.5. Environmental requirements	11
7.6. Installation and support requirements	11
7.7. Reliability requirements	12
7.8. Interoperability requirements	12
8. Conclusions and next steps	12
Abbreviations	13
Bibliography & References	13

List of Figures

Title	Page Number
Figure 1 - Telecom for Wellness Opportunity and Challenges Summary	4

List of Tables

Title	Page Number
Table 2 - Location where services are being offered by current and some future applications	6

1. Executive summary

Customer premises networks have been evolving since the early days of POTS (Plain Old Telephone Service). The introduction of new services like Cable TV, Broadband, and later VoIP, has driven numerous changes in the premises networking requirements. It is therefore no surprise that with the emergence of a vast array of new Internet of Things (IoT) applications and other revenue-generating applications, the premises networking requirements will change yet again. Also, with the advent of these new service offerings such as Telecom for Wellness (T4W) [3], [4], [5] and Smart Cities [12] applications, even the definition of the term “customer premises” itself needs to be realigned.

To better understand these changes in the customer premises networks - as well as plan for them in advance - the SCTE has launched a new initiative under its IoT Working Group called the Wiring Closet Drafting Group (WCDG) [13]. This is a cross-functional team comprising members from the IoT, Smart Cities, Telehealth, and Aging in Place (AIP) working groups. The expectation is to gather and consolidate customer premises networking requirements from all these various classes of IoT applications and make recommendations for future customer premises networks that can be built in a modular fashion using the recommendations from the WCDG.

In this paper, we explore the special needs arising from these recent developments in the customer premises network and outline a framework of requirements that are being compiled by the SCTE WCDG. We highlight key service and business considerations for various application use cases including managed Wi-Fi, home security, home automation, telehealth, aging-in-place, hospital-at-home, smart cities, and their associated security, installation, and support services. We believe this framework can greatly enhance the serviceability and adoption of these new services and can enable significant future revenue opportunities for communications service providers (CSP).

2. Evolution of cable customer premises

Home networks have come a long way over the past few years. They are no longer the domain of technology enthusiasts who cobbled up connectivity between various devices in the home. Today – especially in the wake of the ongoing pandemic – home networks have taken center stage in the life of most people. In addition to the need for ubiquitous Wi-Fi service being the critical driver for these networks – a host of home automation IoT services including sensors, video communications, interactive devices, as well as other next-generation service offerings such as aging-in-place and telehealth are beginning to drive significant complexity in the home network.

From the cable customer perspective, the scope of cable services has also evolved from primarily video services to mainly residential subscribers to now a much broader set of video, voice, and data services to subscribers spanning residential, businesses (offices, retail, hospitality), airports, parks, stadium, hospitals, and even governments (municipal, city infrastructure) etc.

The recent emergence of IoT applications has dramatically expanded the type of services cable networks can deliver. At the same time, the needs of all these disparate services have brought with it significantly increased complexity in the network infrastructure in the customer premises. The future of in-home networking is dependent on the modular deployment of revenue-generating services that not only offer services in a tiered fashion but also are manageable from installation and service assurance points of view. To better understand the needs of the emerging IoT applications and services, and to better prepare for them, the SCTE IoT working group has launched the Wiring Closet Drafting Group (WCDG). This group is a cross-functional team from the IoT, Smart Cities, Telehealth, and Aging in Place (AIP) working groups. The fundamental Premises of WCDG is that by deploying relevant solutions where the

services are deployed a priori, the cable operator can (1) turn up the service faster, (2) customize services to the needs, and (3) can have a competitive advantage in delivering end-to-end services. In this work, we evaluate different needs that arise from such enhancements to the in-home (SFU, MDU), common area (MDUs and planned communities), and smart cities network. We propose a step-by-step upgrade to the Telecom closet, a general-purpose term used for an aggregating unit of the Telecom services offered by the operators. These factors are further explored in this paper.

3. Emerging IoT applications

No new technology in recent history has had the dramatic scope of impact as IoT. This is a technology that touches just about everything we deal with and has the potential to revolutionize life in unforeseen ways. Many ecosystems (such as Telehealth, AIP, smart cities, etc.) are being developed using these IoT devices. Let's take a quick look at some of the more popular IoT applications and services being deployed today.

Home Automation and Security Applications: Some of the most popular IoT applications deployed in homes today are for home automation and security. Home automation enables homeowners to control various elements in their home remotely. These include lights, door locks, thermostats, sprinkler systems etc. While home security has been deployed for many years, most applications in the past involve professionally installed and monitored systems using proprietary technologies. IoT has opened the door to self-installed security systems and devices including sensors and cameras for monitoring the home remotely. These have brought not only significant convenience to homeowners, but they can literally be lifesaving at times.



	 Aging in Place	 Telehealth
Users	Older adults (65+), caregivers	Individuals, providers
Stakeholders	Family members, caregivers, doctors, service personnel etc.	All family members, providers, (payors)
Needs	Communicating, monitoring, service, support, integration	Communicating, monitoring, integrating with provider systems
Challenges	Ease of use, provider network integration, problem solving	Ease of use, device and EMR integration, remote monitoring,
Telecom opportunity	End to end solution, managed services, provider integration	End to end solution, managed services, provider integration

Figure 1 - Telecom for Wellness Opportunity and Challenges Summary

Telecom for Wellness (T4W) applications: The wellness industry is going through a major transformation to modernize the infrastructure, reduce the cost and increase the quality of care. In a series of articles, we have suggested how the Telecom industry can assist the wellness industry[5], [6], [7], [8], [9]. We call this inter-industry collaboration Telecom for Wellness (T4W). Even though the T4W

opportunity is not limited to these two major intersection points, we focus on AIP and Telehealth use cases to illustrate our requirements on the WCDG architecture. (Refer to [1] for six different opportunities that a Telecom operator can address through the T4W.) The SCTE Data Standards Subcommittee, of which two of the authors are members, is actively working on T4W solutions for the AIP and Telehealth areas in working groups three [10] and four [11].

Figure 1 provides a quick summary of the T4W opportunity and challenges from AIP and Telehealth points of view. Many of the needs, challenges, and Telecom opportunities of both markets are similar (refer to the SCTE working group analysis at [10], [11]). Some of the high-level use cases that need to be supported for these two markets include:

- Providing basic communication between the users and the providers/caregivers
- Providing seamless communication between the users and the stakeholders
- Monitoring the users for health, mobility, fall detection, etc.
- Analyzing the data collected from the users and properly notifying the stakeholders
- Assisting the T4W service providers with claims by documenting accountability
- Offering managed services to support installations, product support, and other services to improve adoption and retain customers

The goal of this paper is not to elaborate on these use cases, but to use them to motivate the WCDG modular architectures. For additional information refer to the working group documents.

Retail Business Applications: Retailers have been on the forefront of the IoT revolution trying to exploit the troves of valuable information made available by IoT-enabled devices. From managing the store security, environment, lighting, to monitoring customer traffic and store inventory on a real-time basis – IoT is finding uses in almost every facet of their business today.

Factory/Industrial Automation: Factories and the manufacturing industry has been at the forefront of automation for a long time. Most factories deploy purpose-built machines, tools, robots etc. to automate repetitive tasks. IoT has brought a whole new dimension to this automation. Machines do not just automate their workflow – they can even coordinate seamlessly with other machines and systems on the shop floor, and perhaps even with others across the world.

Inventory/Fleet Management: IoT enables convenient and precise tracking of the location of enabled devices. This has found many applications in the shipping and logistics business for managing product inventory as well as for vehicle and material tracking within a building to all the way across continents.

Smart City Applications: IoT is finding many applications in municipal governance from automated utilities monitoring to smart streetlights, traffic and parking management, energy, and waste management to public transportation applications. Cable operators are uniquely positioned to assist in accelerating many of these applications (refer to Smart City working group at [12]).

Some of these applications intersect in terms of service offerings. For example, smart communities can provide wellness applications along with smart city applications on the managed public Wi-Fi infrastructure.

All the above evolving solutions can provide a huge differentiator for cable operators if they are ahead with their infrastructure to monetize them through flexible and managed offerings at the service delivering locations.

4. Special needs of IoT applications

The vast array of new IoT applications – some of which we outlined above - bring with them many special needs. Let's explore some of the key needs of these IoT applications across various dimensions.

4.1. Location needs

The first dimension we will look at is location. The location where a service is delivered has a major impact on many other requirements across many other dimensions as well. The location determines the scope of the service, and hence the prebuilt modular components that can be deployed in that location are determined. Table 1 gives a high-level view of the typical location where certain types of IoT services are delivered.

Table 1 - Location where services are being offered by current and some future applications

Service	In single-family homes	In multi-dwelling homes	In business premises	In common areas	In public places
Traditional quad-play	X	X	X		
Residential IoT	X	X		X	
Business IoT			X		X
Smart city applications			X	X	X
Telehealth applications	X	X	X		
AIP applications	X	X		X	X

4.2. Security needs

The second dimension we look at is security. Security is another dimension that has broad implications across many areas of an IoT application. IoT devices and applications can expose sensitive data and details about subscribers. Modern networks must be architected to protect and isolate these devices. The network should be capable of identifying the devices that run on it, and where possible, implement security controls based on this device's identity. Devices should be explicitly and securely onboarded on to the network using standards such as Wi-Fi Alliance's Easy Connect protocol, which gives devices unique credentials on the network that can be easily updated on a per-device basis and does not require re-credentialing of all the devices on the network.

In addition to device identity, the network architecture should allow devices to be isolated and segmented based on several factors including device type, function, or risk. For example, medical devices may be put on a network that is separate from other smart home devices. This would allow for access controls scoped to the medical device and prevent generic smart home devices from connecting to the medical device.

Finally, future networks should be able to recognize if the behavior of a device is acting outside of established bounds. This could be accomplished deterministically using a device's Manufacturer Usage Description (MUD) which gives explicit ports and protocols that the device is allowed to communicate with over the network, or probabilistically using heuristics or machine learning.

4.3. Reliability needs

As we move from entertainment services to many more mission critical services in the customer premises, the need for system reliability increases significantly. Telehealth services in the home, industrial IoT and connected vehicle support applications will require high availability.

4.4. Timing/Latency needs

Many real-time applications like industrial IoT and connected vehicle support applications that will require stringent network timing and latency control. Since not all applications have the same level of need, it is best to architect the wiring closet based on the needs of the applications being served.

4.5. Power needs

Once again, the applications requiring high reliability will also require reliable power. Backup power will commonly be needed in many cases. The location of the wiring closet will determine the powering options. Battery backup will commonly be needed in many cases. The type of applications being supported by the wiring closet and the location will determine the battery backup options.

Power has many considerations in this infrastructure. There are power sources, power demands and power distribution which must work in concert.

4.5.1. Power sources

Let's look at some of the typical power sources that can be used for these applications.

4.5.1.1. Utility power

Utility power is the most abundant of the power source which is available at all but a few potential deployment sites. The North American power grid is delivered to customer premises using a 120V/240VAC single split phase deployment. The same grid can deliver this 120/240VAC service to businesses, but 3 phase utility power is also delivered to businesses. These 3 phase deliverables can be modified using a transformer to deliver the common 120/240VAC used elsewhere, so this utility power standard will be referenced throughout this document. The frequency of the utility AC waveform is constant throughout North America using a 60Hz standard.

Most electronic devices considered for this IoT infrastructure deployment support utility power using a universal utility AC power supply or power pack which can directly connect to the utility grid for powering. These packs are reliable, cost effective and energy efficient. The packs come in a wide range of capacity to drive both the direct device needs as well as those that may be hosted by these devices.

Though the utility grid is ubiquitous and capable there can be issues with its delivery. A utility or grid outage is a common term for these delivery issues. In most cases these outages are of a short duration and more of a nuisance for the end user, but the electronics and services supported by this IoT infrastructure system could provide "Lifeline" services, so any interruption of their operation can create a bad scenario. The following steps should be taken to improve resiliency and reduce trouble calls:

- Limiting user intervention to the power source of the IoT infrastructure.
- Avoiding a “switched” outlet and have a dedicated power circuit, if possible, with a clearly marked circuit breaker in the panel.

4.5.1.2. HFC plant power

HFC plant powering is one of the oldest methods of proving CPE power in a cable network. eMTAs and ATAs are also commonly powered by the cable HFC plant. HFC plants typically power the copper network and associated actives using 63VAC or 89VAC power supplies. This power can be routed anywhere the hardline plant is deployed and extended to the customer through the use of power passing taps. The older HFC CPE power passing taps could only support about 300mA, but newer power passing taps can support 2+ Amps of HFC power to the CPE devices.

There are some considerations while using HFC plant powering:

- The load on the HFC plant must be considered in larger, high current deployments. The typical HFC power supply has a capability of delivering 15A @ 90VAC routinely; but if the IoT infrastructure demanded 2 Amps for each site this would quickly exhaust the available power.
- HFC plant powering could be used as a primary or backup source of power, but once again, a high demand would exhaust the HFC powering batteries quickly; additional batteries might have to be added to the HFC powering location to address this additional power demand. The HFC plant is very efficient at propagating RF signals, but slightly less efficient at conveying the HFC powering. Losses occur through the hardline, actives, passives and largely in the customer drop. The end of line voltage range would need to accommodate a voltage range from 45 to 90VAC.

4.5.1.3. Hybrid fiber power

The Hybrid Fiber Power solutions have been deployed in cellular tower deployments as well as similar high voltage low current powering of DSLAMs. PON architectures could provide centralized powering sources and hybrid power cable with converters in the IoT infrastructure cabinet to provide the needed power for the supported electronics. The basic designs leverage several 24 AWG to 12 AWG conductors to allow a high voltage low current transmission, some touting distances of up to 10km. The hybrid fiber powering could be delivered as a primary or backup source of power and communications.

4.5.1.4. Local renewable power

Consumers and businesses are deploying renewable energy resources at their sites and these sources can be leveraged to power our IoT infrastructure hardware. Solar, Wind, Micro-Hydro are the main sources of these renewable energy systems and there are several modes of operation of these systems: Grid Direct, Off grid, and Hybrid. Grid direct and Hybrid are most likely what would be used in concert with our IoT infrastructure; but our connection considerations are very different in those two deployments. Grid Direct systems rely on the utility grid for operation, and the loss of the utility grid causes these systems to shut down for safety reasons. The IoT infrastructure would be connected to a panel as it would in a utility deployment, the major difference is that during renewable production the IoT infrastructure would be powered from the renewable energy source. The Hybrid renewable system includes storage and handles grid loss a bit differently by disconnecting itself and its loads from the Grid to assure safety, but a “Backup Loads Panel” would be powered by the Hybrid system and its storage. In this case all the IoT infrastructure should be tied to the backup loads panel to assure operation through utility outage events.

4.5.2. Backup power considerations

The IoT infrastructure will likely support lifeline services and will require resilient powering. There are several ways of providing resilient backup power. We will look at some of the most commonly available options below. The amount of backup time these systems provide can vary, so the right option should be used depending on the application.

4.5.2.1. Interdevice batteries

One of the longest deployed backup powering scenarios in cable is the use of a backup battery in an eMTA. A similar method could be used to maintain the IoT infrastructure powering. The battery could be placed in the device itself or located externally. The device can be used to charge the battery. The size of the battery should be sufficient to meet the needs of the application.

Battery based backup is very efficient from both a cost and round-trip power utilization point of view. However, it can prove to be an extra burden when the batteries reach their end of life and must be replaced. Most of these battery backup solutions include battery monitoring for run time and battery health.

4.5.2.2. External UPS

External UPS devices are probably the second widest deployment of backup power for cable user equipment deployments. These are external units that are plugged into the utility and then the supported cable hardware is plugged into them. These systems are largely not monitored, which can create a “delayed outage” scenario where the system runs for a time during a utility event, but when the battery is exhausted in these external UPS devices the hosted IoT infrastructure system goes dark. There are both consumer and industrial grades of UPS devices, with the industrial scale devices providing longer runtime better monitoring options and survivability.

4.5.2.3. Whole house backup

Behind the meter storage and other large scale backup solutions are similar hybrid renewable system that can leverage renewable sources. The idea is that a large battery bank is hooked to an inverter/charger which can provide a backup power source for many powered devices in the premises. The system will disconnect from the grid during a power outage and through a backup loads panel can provide service during the utility outage. The IoT infrastructure electronics would need to be connected to the backup loads panel for power.

4.6. Interoperability needs

With the evolution of IoT technology, a broad range of devices, applications, and services have been deployed in the customer premises networks today. The IoT framework of the future must provide some means of integrating all these devices and applications in a consistent manner.

5. Opportunities for cable operators

So why should cable operators care about these emerging new services in their customer premises networks – especially given all the complexity they bring with them?

It is precisely the complexity that has stymied the widespread deployment of these IoT services in the market. Numerous technology vendors are vying to capture the market – but, at best, most have only succeeded in delivering point solutions to one specific application they control.

We believe that cable operators are ideally placed to take full advantage of their central role in the customer premises network to pull all the pieces of the puzzle together and thereby unlock a lot of potential value in their customer's IoT environment.

It is also clear that many proprietary solutions currently being deployed will only find limited success in the marketplace. For any widespread deployment, a standards-based approach will be needed.

6. Wiring closet concept

The traditional infrastructure cable operators placed in the customer premises was simple and compact. All that was needed was a small enclosure or a closet to house it all. This was commonly known as the Wiring Closet or Telecom Closet. As we outlined above, the needs of the customer premises – and even the customers themselves – have evolved significantly. So, what is the new “Wiring Closet” of the future?

It is clear the needs of customers will vary a lot from a single-family home subscribing to basic cable to a municipality supporting various smart-city applications. Obviously, we cannot have a one-size-fits-all solution for all these applications. It is also clear the customer premises infrastructure of the future will not necessarily even be a closet. The idea is to develop a customer premises framework that provides a standard and modular way of scaling from a simple wiring enclosure to a distributed set of resources where necessary.

The following are the three basic characteristics of a Wiring Closet:

- Expandable to current and future IoT applications
- Address the spectrum of needs from in-home to public areas
- Serviceable unit for the cable operators

7. Wiring closet requirements

What are the requirements for such a broad framework to support IoT services in the customer premises network? That is one of the main questions the WCDG group is trying to address at the SCTE. While this is still a work in progress, we highlight a few key requirements in this section that the group feels are necessary for any such framework.

7.1. Customer premises location requirements

As we noted above, various IoT applications need to be supported in different locations in a customer premises. Following locations must be supported by the customer premises framework.

- Single Family Units: Support in indoor and outdoor locations.
- Multi-Dwelling Units: Support in unit indoor and outdoor locations as well as common area indoor and outdoor locations.
- Commercial Units: Support in indoor as well as common area indoor and outdoor locations.
- Public Spaces: Support in a cabinet as well as pole mounted locations.

7.2. Networking requirements

Customer premises framework needs to support both wired and wireless networking infrastructure.

- Wired infrastructure should preferably be supported for premises distribution where wireless distribution is challenging and for fixed devices.
- Wireless infrastructure should be supported for mobile devices. Adequate signal coverage throughout the customer premises should be ensured.

7.3. Security requirements

As noted above, security is critical for many IoT applications. As such, the customer premises framework shall support comprehensive security capabilities including following:

- All communication links shall support authentication, authorization, and encryption capabilities.
- Wireless Access Points shall support authentication and encryption capabilities.
- Gateway devices shall support secure boot and secure upgrade capabilities.
- No default login credentials should be supported. In addition, security infrastructure should be upgraded to the latest available standards.
- Any test/diagnostic access ports shall be disabled by default or require secure login.

7.4. Powering requirements

Many IoT applications have special powering needs. As such, the premises infrastructure should support following powering requirements.

- The premises infrastructure shall support following options as power sources:
 - AC utility mains
 - HFC 60V/90V
 - Power over Coax
 - Power over Ethernet (PoE)

7.5. Environmental requirements

Since the premises infrastructure will have to support a wide range of premises locations, it needs to meet a broad range of environmental needs. Following are some of the key environmental requirements the premises infrastructure will need to support.

- The premises equipment shall support installation in indoor enclosure, outdoor enclosure, outdoor cabinet installation, and outdoor pole mount enclosure.
- The equipment shall meet all necessary environmental and safety requirements per ETL, UL, and NEMA.
- The equipment shall meet all local building and environmental codes.

7.6. Installation and support requirements

Due to the broad range of potential IoT applications that could be installed in the customer premises and the resulting complexity, the premises framework will need to meet numerous installation and support requirements.

- All serviceable equipment must be located in an area offering easy access for service personnel

- Basic infrastructure should be pre-installed in any new building construction
- Auto-configuration processes should be used where possible to minimize user intervention
- Installer should ensure adequate signal coverage throughout the premise (e.g., Wi-Fi signal coverage)
- All installed equipment and infrastructure should be well labeled and documented
- Support responsibilities should be clearly outlined along with contact information for responsible party
- Maintenance responsibility for any backup batteries should be clearly specified along with necessary maintenance schedule
- Installation should support notification of low/failed/missing backup battery condition to service personnel
- Equipment installations should support easy hardware upgrade.
- Devices deploying software should preferably support automatic online software upgrades.

7.7. Reliability requirements

Many IoT applications will potentially support critical healthcare or business support services that will have high reliability expectations. Following are some of the reliability requirements the premises infrastructure needs to meet:

- **Power Reliability:** Backup power should be made available depending on the type of application and location needs.
- **Communication Reliability:** A backup communication link may be necessary in case of failure of primary communication link.
- **Device Reliability:** Devices providing critical services such as healthcare support should have built-in redundancy or a spare available.
- **Service Reliability:** Critical services must be supported by a highly reliable support infrastructure including a 24x7 support team.

7.8. Interoperability requirements

It is not enough for IoT applications to function by themselves. To minimize complexity for the user as well as to make applications easier to use, the customer premises framework should provide a common infrastructure for applications to interface with each other. Following are some of the interoperability requirements the premises infrastructure needs to meet:

- The customer premises framework should support a common infrastructure for applications to interface with each other.
- The framework should support a consistent way of managing all the applications and devices in the customer premises.

8. Conclusions and next steps

In this paper we have clearly outlined the need for developing a comprehensive framework for supporting the vast array of IoT applications and services in customer premises networks from simple single-family homes to large city-wide smart-city networks. Such a framework will enable widespread deployment of a multitude of value-added services in the near future and place the cable operator in a pivotal role in this vast ecosystem.

The first step in getting to such a framework is to develop a comprehensive set of requirements to be supported by this framework. We have highlighted a number of key requirements in this paper, but much more work must be done. The WCDG group is looking for subject matter experts from a broad range of areas to come and contribute to this work.

Abbreviations

AIP	aging in place
ETL	Electrical Testing Laboratories
HFC	hybrid fiber coax
IoT	internet of things
MDU	multi dwelling unit
NEMA	National Electrical Manufacturers Association
PoE	power over ethernet
POTS	plain old telephone service
SCTE	Society of Cable Telecommunications Engineers
SFU	single family unit
T4W	telecom for wellness
UL	Underwriters Laboratories
VoIP	voice over internet protocol
WCDG	wiring closet drafting group

Bibliography & References

- [1] SCTE 266-2021: *IoT Recommended Premises Network Infrastructure Practices for Cable Operators*
- [2] Rajesh Abbi, Changing Landscape with IoT, Dec 2021, available [here](#)
- [3] Duke Tech Solutions market research, *Telehealth market report – A Telecom based opportunity analysis*, available [here](#)
- [4] Sudheer Dharanikota, *Summary of Telecom for Wellness interviews*, Oct 2021, available [here](#)
- [5] Clarke Stevens, Sudheer Dharanikota, *Aging in Place and Telehealth Use Cases from the Cable Operator Perspective*, SCTE Journal, March 2022, available [here](#)
- [6] Sudheer Dharanikota, Ayarah Dharanikota, *Why are cable operators natural fit to support Telehealth – An inter-industry perspective*, 2020 SCTE Expo, available [here](#)
- [7] Sudheer Dharanikota, Ayarah Dharanikota, Dennis Edens, Bruce McLeod, *Aging in Place business case for cable operators*, SCTE Journal, June 2021, available [here](#)
- [8] Sudheer Dharanikota, Ayarah Dharanikota, Dennis Edens, Bruce McLeod, *Telehealth business case for cable operators*, SCTE Journal, September 2021, available [here](#)
- [9] Sudheer Dharanikota, Clarke Stevens, *End to end Telecom for Healthcare architecture – a cable operator perspective*, 2021 SCTE Expo, available [here](#)
- [10] SCTE Data Standards Subcommittee, Working Group 3, Aging in Place, available [here](#)
- [11] SCTE Data Standards Subcommittee, Working Group 4, Telemedicine, available [here](#)
- [12] SCTE Data Standards Subcommittee, Working Group 1, Smart Cities DG, available [here](#)
- [13] SCTE Data Standards Subcommittee, Working Group 1, Wiring Closet DG, available [here](#)

Robust and Resilient Service Assurance System Design with Observability to Improve Enterprise Customer Experience

A Technical Paper prepared for SCTE by

Anil Mohan

Principal Engineer, Product Development Engineering
Comcast Cable
1800 Bishops Gate Blvd, Mt. Laurel, NJ 08054
anil_mohan@cable.comcast.com

Xin Huang

Sr. Principal Engineer, Product Development Engineering
Comcast Cable
1800 Bishops Gate Blvd, Mt. Laurel, NJ 08054
xin_huang@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. High-level Service Assurance System Design and Challenges.....	4
2.1. Out of Sync Device or Port Status between Systems.....	5
2.2. High Response Time and Timeouts for Data Presentations.....	5
2.3. Lack of Mediation Layer to Standardize the Multi-Vendor Log/Event Formats	6
2.4. Least Scalable with Fewer Integration Options.....	6
2.5. Lack of Advanced Correlation between Faults and Performance Metrics to Provide Meaningful Insights	6
2.6. Less Visibility to Different Layers of Network Service.....	7
2.7. Lack of Advanced Troubleshooting Capabilities.....	7
3. Advanced Service Assurance System Design with Observability and Awareness	8
3.1. New Mediation Layer Added	9
3.2. Separate Dedicated Path for Fault and Performance Events.....	9
3.3. Additional Sources for Alarm Validations.....	9
3.3.1. Polling the Device / NMS	9
3.3.2. Adding Heartbeat Events from Device / NMS.....	10
3.3.3. Correlating Events/Alarms against Performance Metrics	10
3.3.4. Correlating against Underlying Network/Transport Layer Alarms / Tickets.....	10
3.4. New Data Delivery Methodology for Faster and Consistent API Responses.....	11
3.5. Hybrid Notification Approach.....	12
3.6. New Message Bus Added	12
3.7. Data Storage Changed from Traditional RDBMS to Big Data system.....	12
3.8. New Raw Log/Event/Message Browser.....	13
4. Improvements for Customers and Operation Teams Experiences.....	13
4.1. Out-of-Sync Device or Port Status between Systems Improvement	13
4.2. Data Presentations Response Time Improvement	13
4.2.1. Customer Level:	14
4.2.2. Site Level:	14
4.2.3. Multiple Sites Level:	15
4.3. Mediation Layer Improvements.....	15
5. Conclusion and Future Work.....	16
Abbreviations	16
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1 – Initial Service Assurance Architecture	4
Figure 2 – Initial Service Assurance Architecture Challenges	5
Figure 3 – Advanced Service Assurance Functional Blocks	8
Figure 4 – Advanced Service Assurance Architecture	9
Figure 5 – Cross-Correlation between Alarms and Performance Metrics	10
Figure 6 – Out of Sync Issue Before and After	13
Figure 7 – API Response Times for Customer Level Before and After	14
Figure 8 – API Response Times for Site Level Before and After.....	14
Figure 9 – API Response Times for Multiple Sites Level Before and After	15
Figure 10 – Null Device Name Issues Before and After	16

List of Tables

Title	Page Number
Table 1 - Hybrid Notification Approach (with Underlying Transport Issue Identified).....	12

1. Introduction

Service assurance (SA), as a subset of the operational support system (OSS), plays an important role in the internet service provider (ISP) ecosystem. However, the rapidly evolving internet service provider (ISP) technologies, enterprise services offerings, and customer expectations bring great challenges to the modern service assurance system design. This paper discusses several general design principles and best practices that are essential to building a robust and resilient service assurance system with observability and awareness that could stay ahead of these fast-paced industry transformations.

First, collecting telemetry from multiple sources helps to avoid single point of failure (SPOF) and improve confidence and accuracy of alerting customers of network/service issues. Second, introducing a unified mediation layer provides flexibility to isolate vendor-specific implementations and prevent bugs from negatively impacting customer experience. Third, making use of cross product correlation and leveraging machine learning (ML) for data analytics, trending and anomaly detections to prevent service interruptions and guarantee accurate customer alerting.

This paper reflects years of SDN-based centralized service assurance system integration design, development, and customer support experience. In this paper, the authors will share ways in which these principles and techniques are applied in our enterprise service product to support business values and keep our customers happy.

2. High-level Service Assurance System Design and Challenges

This section describes the initial service assurance system design and challenges and next section will describe how we resolved these challenges using the advanced service assurance architecture design. Figure 1 below shows the different layers and functional blocks of the initial service assurance architecture design with the key functional roles played by each layer. Figure 2 below shows the different challenges faced at the different layers in the service assurance architecture.

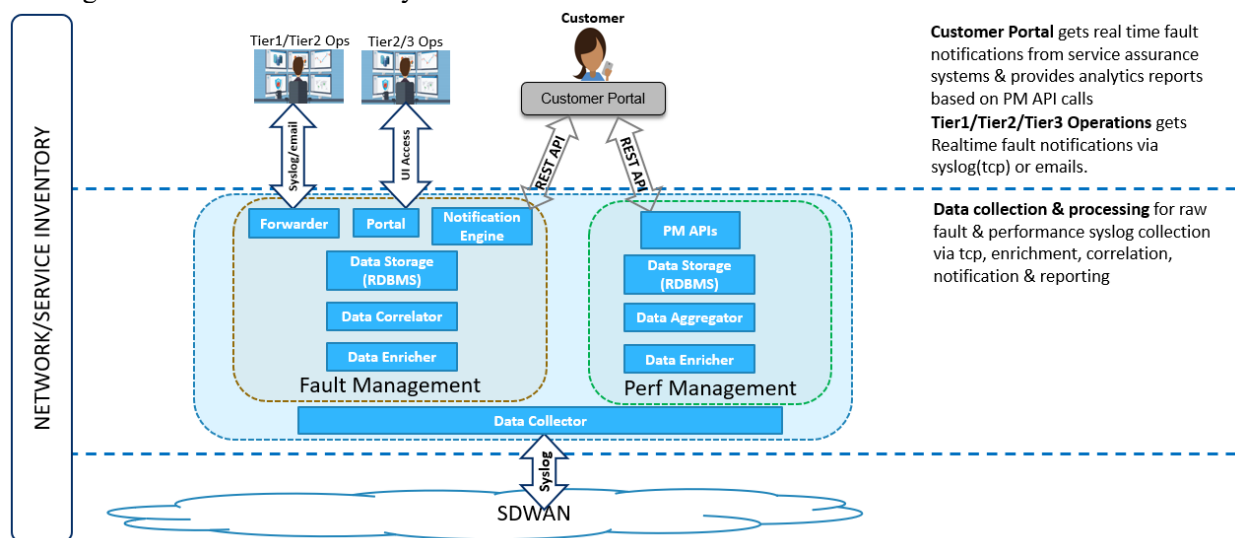


Figure 1 – Initial Service Assurance Architecture

As illustrated in Figure 2, the key challenges faced in the initial service assurance architecture design are:

- Out of sync device or port status between systems;

- ❑ High response time and timeouts for data presentations;
- ❑ Lack of mediation layer to standardize the multi-vendor log/event formats;
- ❑ Least scalable with fewer integration options;
- ❑ Lack of advanced correlation between faults and performance metrics to provide meaningful insights;
- ❑ Lack of visibility to different layers of network service; and
- ❑ Lack of advanced troubleshooting capabilities.

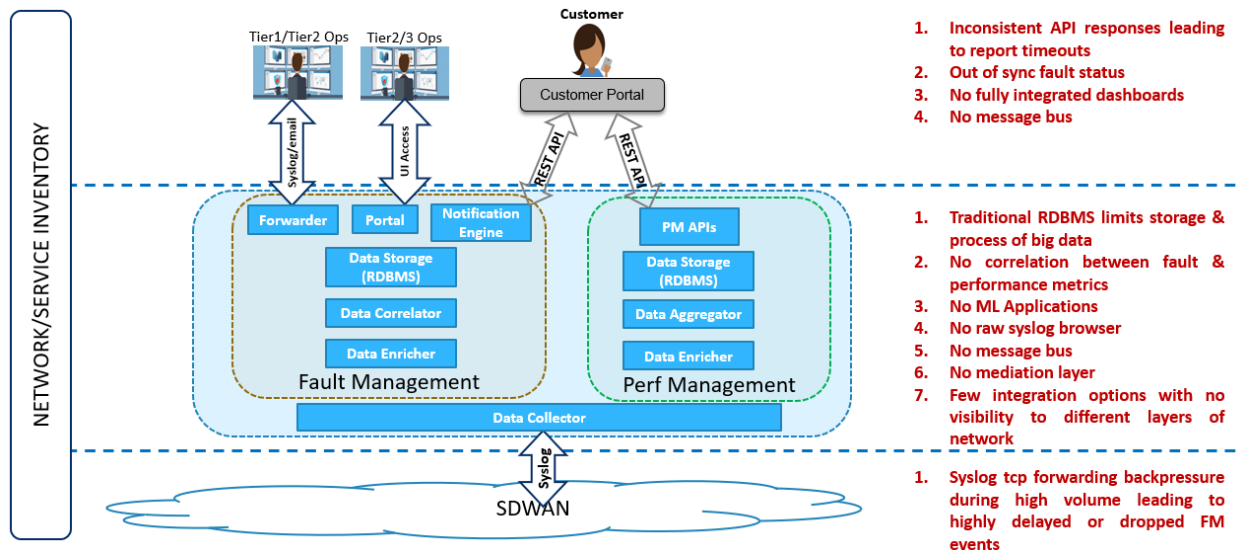


Figure 2 – Initial Service Assurance Architecture Challenges

2.1. Out of Sync Device or Port Status between Systems

With the initial service assurance system, there were some cases of device or port statuses going out of sync between systems. Customer experience may have been impacted by these out-of-sync issues.

Below listed some of the reasons for statuses to go out of sync:

1. When vendor element managers fail to send set or clear event due to bugs.
2. When set and clear events/messages come at the same time.
3. When there is no dedicated path for faults, fault event may be delayed or dropped due to congestion in the analytics data pipeline to service assurance systems.
4. When the vendor does not provide single fault event for a problem and the service assurance system does not have the capability to cross-validate against performance metrics.
5. When producing or consuming systems fail to process the fault events.

2.2. High Response Time and Timeouts for Data Presentations

The initial service assurance architecture design had the potential to cause delays and timeouts in loading data reports for customers.

Below are listed some of the reasons for such high response time or timeouts for the reports:

1. The granular customer data for the top reports were pulled for the selected time and aggregations were run on the fly at the portal reporting platform.

2. Every user click for a similar report triggers corresponding database pulls for the granular customer data. This causes a lot of unnecessary database transactions which is not a normalized approach.
3. The granular customer data was being pulled from non-big data storage systems which caused delays when the customer data were high.
4. The report data were tied with inventory cache where the inventory cache was not up to date all the time, leading to missing or inconsistent data.

2.3. Lack of Mediation Layer to Standardize the Multi-Vendor Log/Event Formats

In the initial service assurance architecture, there was no mediation layer to standardize, filter out, modify, or convert log/event formats to meet the customer needs. This led to higher time to roll out changes because of the need to accommodate log/event format changes.

Below listed some of the reasons for requiring a proper mediation layer in service assurance architecture:

1. Log/event format changes due to vendor upgrades or bug fixes.
2. Filter out log/event due to bugs in existing vendor versions
3. Standardize multiple logs/events into same format for better processing and reporting
4. Define higher priority pipelines for fault events compared to performance metrics
5. Enable tagging and distribution of logs/events

2.4. Least Scalable with Fewer Integration Options

Initial service assurance architecture had scalability limitations at all functional layers, including but not limited to data collection, data storage, data integration with external systems etc. This would become a bottleneck as the network grew and caused delays in processing incoming or outgoing messages leading to bad customer experiences.

Below are listed some of the reasons for requiring a more scalable and highly interfacing architecture:

1. Number of logs/events per customer device is not static but rather dynamic based on customer traffic. The service assurance systems should be scalable enough to accommodate changing volumes of customer device traffic and usage.
2. To achieve a better predictive and reliable service assurance system, there is always the need to integrate multiple datasets from different platforms. This requires multiple integrations / interfacing points.
3. For faster rollout of features, service assurance architecture should be able to consume / collect multiple data types using different protocols, and be able to produce / export data to different interfacing systems.

2.5. Lack of Advanced Correlation between Faults and Performance Metrics to Provide Meaningful Insights

Initial service assurance systems did not have the capability to correlate across fault and performance metrics. Correlating across fault gives meaningful insights into the data along with accuracy and better ML predictions. Alternate options within initial service assurance systems limited the full capabilities one gets from marrying the fault and performance metrics.

Below listed some of the benefits for correlating between faults and performance metrics:

1. For meaningful insights into customer data.

2. Good ML predictions.
3. Better accuracy of fault events.

2.6. Less Visibility to Different Layers of Network Service

The initial service assurance systems were not flexible enough to provide visibility to different layers of network service since there were minimum integration options and no correlation between fault and performance metrics. Visibility to different layers of network service can help in identifying the root cause of any issue, provide better representation of the service to operations or customer, and create a better fit for proactive monitoring of the network service.

Below listed some of the reasons for having better visibility to different layers of network service:

1. Better root cause analysis (RCA).
2. Better visualization of the service.
3. Better proactive monitoring of the service.

2.7. Lack of Advanced Troubleshooting Capabilities

The initial service assurance systems did not have good visualization capabilities to create on-the-fly dashboards on fault, performance metrics or both fault and performance metrics.

Below are listed some of the factors that impaired the troubleshooting capabilities:

1. Lack of visibility to different layers of network services.
2. Lack of correlation between fault and performance metrics.
3. Separate clients for viewing faults and performance metrics, rather than a single pane of glass view.

3. Advanced Service Assurance System Design with Observability and Awareness

This section describes the key enhancements / features of the advanced service assurance architecture that can overcome each of the challenges explained in Section 2 by implementing Architectural design changes or methodology changes.

Figure 3 and Figure 4 describe the high-level service assurance architecture providing observability and awareness.

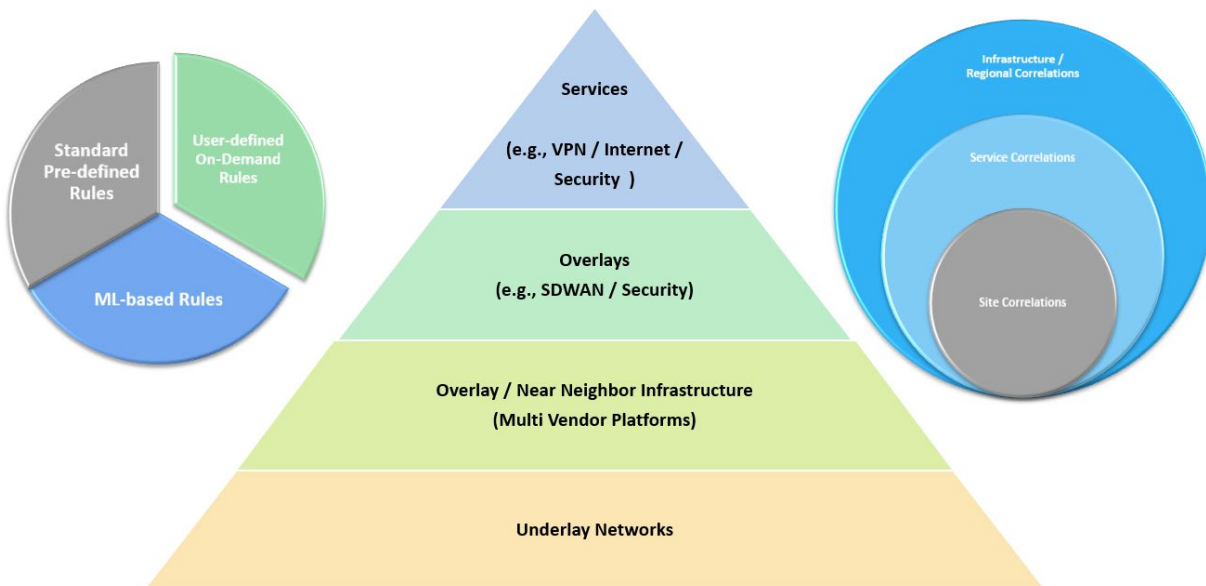


Figure 3 – Advanced Service Assurance Functional Blocks

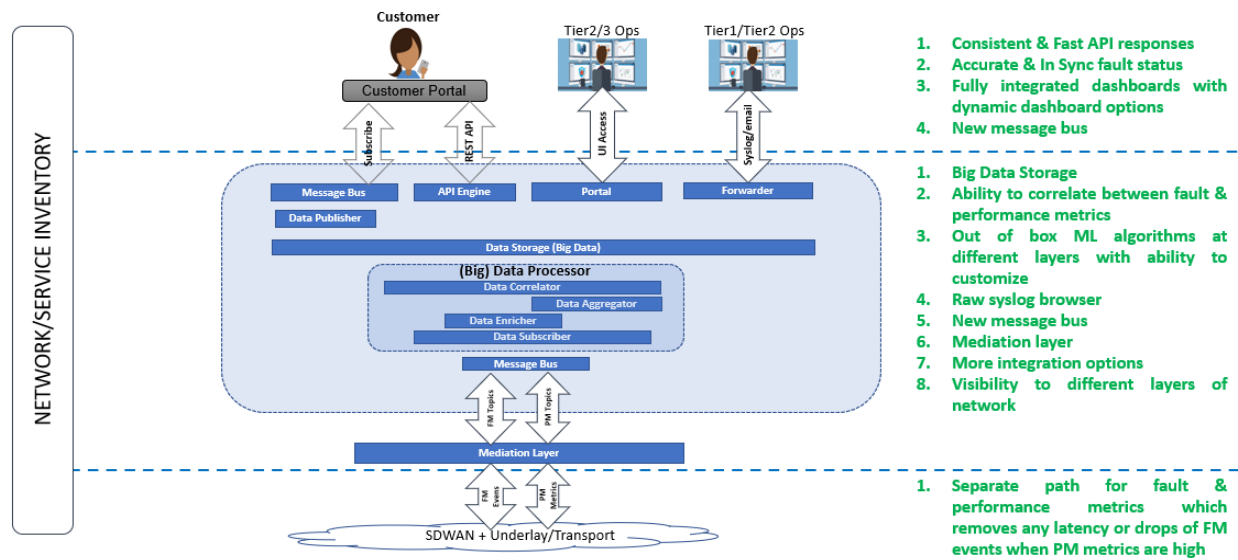


Figure 4 – Advanced Service Assurance Architecture

3.1. New Mediation Layer Added

A new mediation layer added between the network source and the service assurance collection layer as shown in Figure 4. This layer adds the functionalities below to help in creating a more robust service assurance architecture by eliminating some of the challenges listed in section 2. For example:

1. Filtering out messages.
2. Format changes of messages.
3. Supports multiple input and output format for message transfers between systems.
4. Message tagging.
5. Distribution of messages to multiple destinations.

3.2. Separate Dedicated Path for Fault and Performance Events

As part of the advanced service assurance architecture, separate dedicated paths were added for fault and performance events from the source to mediation layer to data collectors in assurance systems. This avoids any latencies on fault events whenever the volume of performance events becomes too high and causes backpressure on those paths.

3.3. Additional Sources for Alarm Validations

As part of the advanced service assurance architecture, more sources were added to fault management to increase the confidence level of the alarm. The below sub-sections explain the 5 major additional sources added to improve the accuracy of device/port/path status. This helped to achieve more consistent and accurate operational statuses.

3.3.1. Polling the Device / NMS

Polling the device/NMS via different supported protocols (e.g., REST API, SNMP, ICMP, etc.) always helps to sync the right operational status within service assurance systems.

3.3.2. *Adding Heartbeat Events from Device / NMS*

Heartbeat events from device / NMS provide a recurring event on the health of the device every 'x' minutes. This helps to keep the service assurance systems in sync on the device health status. This also helps to minimize the dependency on polling since polling device / NMS sometimes can be expensive depending on the device / NMS providers.

3.3.3. *Correlating Events/Alarms against Performance Metrics*

Device / port / path level performance metrics that generated every 'x' mins from device / NMS can be considered as heartbeat messages. The presence of performance metrics can be considered as 'UP' heartbeat while absence of performance metrics from a previously present metric can be considered as 'DOWN' heartbeat. Correlating the events/alarms against these performance metrics will help in improving the confidence level of the corresponding event/alarm. This also helps in minimizing the dependency on polling. Below Figure 5 shows cross-correlation between alarms and performance metrics for finalizing device status.

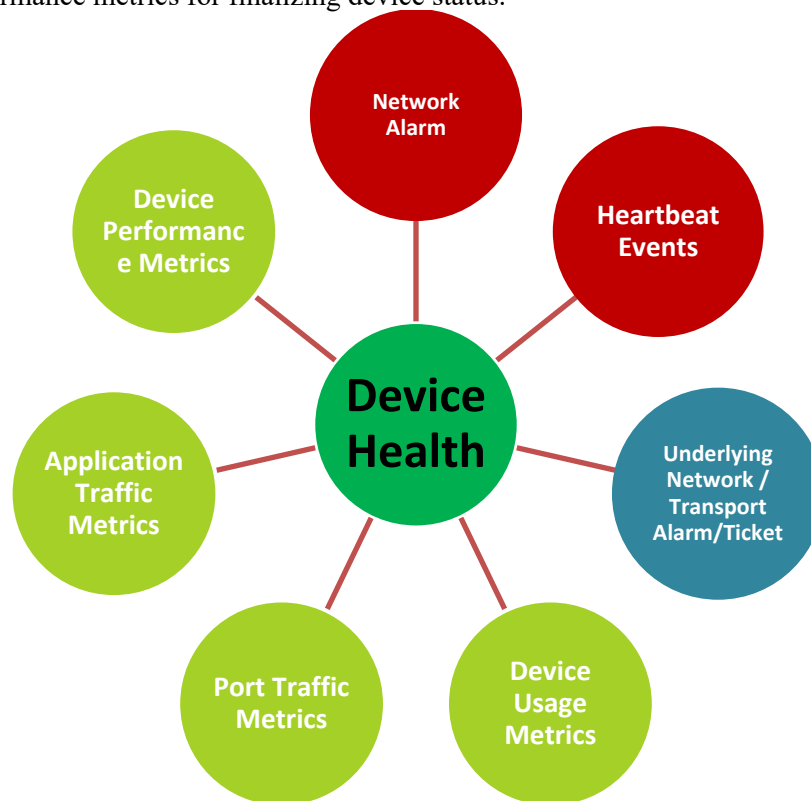


Figure 5 – Cross-Correlation between Alarms and Performance Metrics

3.3.4. *Correlating against Underlying Network/Transport Layer Alarms / Tickets*

Correlating fault events against underlying network / transport layer alarms / tickets always gives better insights into the root cause of the fault/event. This also aids in creating a visual representation of network service and giving the operations team advanced troubleshooting capabilities.

3.4. New Data Delivery Methodology for Faster and Consistent API Responses

In the initial service assurance architecture approach of providing APIs for performance metrics, we have seen multiple challenges leading to delays / timeouts in API responses. In this new advanced service assurance architecture, based on the learnings, we have come up with a new data delivery methodology for performance metrics in reports to customers. The new methodology is much faster and provides consistent API responses.

Below are the key steps for implementing this new data delivery methodology:




1. Identify the different report types to be delivered to customers.
2. Identify the different groupings needed for each of the report, like customer-level, device-level, port-level etc.
3. Identify the pre-defined time buckets with granularity needed.
4. For each type of the above report, pre-stage the data for the report by scheduling these granular data from the database at regular intervals based on the time buckets.
5. Create API request definitions which pick up the latest pre-staged data every time a user clicks a report.

The key benefits of this approach include:

1. No on-the-fly aggregations, leading to faster response times for loading reports.
2. Normalized approach with minimum database calls for reports. No unnecessary database transactions.
3. Removal of the dependency on the inventory cache, thereby resolving any missing or inconsistent data.

3.5. Hybrid Notification Approach

Table 1 - Hybrid Notification Approach (with Underlying Transport Issue Identified)

Category	Status	Realtime / Soaked	Source	Notification Severity	Notification Message
Device / Port / Path Health Notifications	Down	Realtime	Network Set Alarms	Major	Instead of reporting as down, should say something like “there seems to be an issue with the device / port / path. Hold on for further updates”
Underlying Transport Ticket Notifications for the Device / Port / Path Health notifications	Down	Soaked ('X' mins)	PM Metrics / Transport Alarms / Tickets	Critical 	Change severity to Critical and say the device / port / path as down.
Device / Port / Path Health Notifications	Up	Realtime	Network Clear Alarms	Major 	Change severity from Critical to Major, should say something like “issue seems to be resolved. Hold on for further updates”
Underlying Transport Ticket Resolution Notification for the Device / Port / Path Health notifications	Up	Soaked ('X' mins)	Transport Alarms / Tickets	Cleared 	Clear the status of device / port / path

All “Real time” and “Soaked” notifications will go to separate message bus topics. So, consumers can choose to subscribe to specific message topics for each notification type

3.6. New Message Bus Added

Based on the challenges observed in the initial service assurance architecture, message bus is a better fit in the advanced service assurance architecture. Message bus provides a highly reliable, scalable data collector as well as an asynchronous communication platform with decoupling for internal and external data movements. It comes with data persistence and fault tolerance that allows the service assurance system to continue processing data even when different parts of the system fail. Most of the message bus platforms available in market come with a lot of additional features like filtering, enrichment, correlation, and more, providing more intelligence at each layer of the architecture.

3.7. Data Storage Changed from Traditional RDBMS to Big Data system

There are 3 major challenges for traditional RDBMS Storage – data too large, data too complex (multiple data types), and data too fast. As the network grows, the analytics data also grows and as we add more features, there will always be new data types/formats to be processed which makes the analytics data complex with different types/formats. As data grows, data ingestion or retrieval times cannot be compromised as it will negatively affect the customer experience. With these challenges and the added benefits of big data systems, it was an easy decision to move away from the traditional RDBMS to big data.

3.8. New Raw Log/Event/Message Browser

Raw log / event / message browser is a good add-on to any service assurance architecture since it provides the user / operations the ability to skim through historical / real-time alarms / events / performance metrics for troubleshooting purposes / generating on the fly reports / statistics.

4. Improvements for Customers and Operation Teams Experiences

4.1. Out-of-Sync Device or Port Status between Systems Improvement

Figure 6 shows the percentage of occurrences of out of sync device / port alarm status before and after moving to advanced service assurance with the key enhancements or features described in Section 3.




Figure 6 – Out of Sync Issue Before and After


Figure 6 clearly shows the improvements in out of sync issues before (using initial service assurance systems) and after (using the advanced service assurance systems). Before the percentage of out-of-sync statuses varied from 2% - 17% while after there were no such out-of-sync occurrences observed based on 30-day data between April and May 2022

4.2. Data Presentations Response Time Improvement

Below graphs show the API response times at customer / single site / multiple site levels for analytics data consumed by systems external to service assurance systems before and after moving to Advanced Service Assurance with the key enhancements or features described in Section 3.

4.2.1. Customer Level:

Response Times (BEFORE):



Response Times (AFTER):





Figure 7 – API Response Times for Customer Level Before and After

4.2.2. Site Level:

Response Times (BEFORE):



Response Times (AFTER):




Figure 8 – API Response Times for Site Level Before and After

4.2.3. Multiple Sites Level:

Response Times (BEFORE):



Response Times (AFTER):




Figure 9 – API Response Times for Multiple Sites Level Before and After

Figure 7, Figure 8, and Figure 9 clearly show the improvements in API response times at customer / single site / multiple site level for analytics data consumed by systems external to service assurance systems. Below are the key highlights:

- API response times are consistent and lower.
- More aggregation types are supported.
- More flexible filtering with ease-of-use APIs

4.3. Mediation Layer Improvements

The below graph shows the percentage of occurrences of alarms with missing device names for fault events before and after moving to advanced service assurance with the key enhancements or features described in the previous section.




Figure 10 – Null Device Name Issues Before and After

Figure 10 clearly shows the improvements in occurrences of alarms with missing device names for fault events before (using initial service assurance systems) and after (using the advanced service assurance systems). Before the percentage of such missing device names in fault events was very low, varying from 0.01% - 0.09%, while after there were no such fault events with missing device name occurrences based on data collected between Dec 2021 and Feb 2022.

5. Conclusion and Future Work

The initial service assurance architecture design, which faced multiple technical challenges, was changed to include a mediation layer, separate dedicated FM and PM paths, and additional sources of alarm validations, along with a new data delivery methodology for faster and consistent API responses, etc.

This re-architecture has shown a lot of improvements. For example:

- Minimum out of sync issues (reduced from around 15% errors to 0% based on 30-day data between April and May 2022).
- Faster API response times and no time outs for API responses (80% reduction).
- Better integration capabilities.

This new robust and resilient service assurance architecture with observability and awareness has enabled to add more advanced correlations with capability to use ML algorithms to build data models for better prediction, trending, and forecasting.

Abbreviations

API	application programming interface
FM	fault management
ICMP	Internet Control Message Protocol
ISP	internet service provider
ML	machine learning

NMS	network management systems
OSS	operational support system
PM	performance management
RCA	root cause analysis
RDBMS	relational database management system
REST	representational state transfer
SA	service assurance
SCTE	Society of Cable Telecommunications Engineers
SDN	software-defined networks
SDWAN	software defined wide area network
SNMP	Simple Network Management Protocol
SPOF	single point of failure
Syslog	System Logging Protocol
TCP	Transmission Control Protocol
UI	user interface
VPN	virtual private network

Bibliography & References

M. Casado, N. McKeown, and S. Shenker, "From ethane to SDN and beyond", in ACM SIGCOMM Computer Communication Review, vol. 49, issue 5, Oct. 2019, pp 92-95.

L. L. Peterson, C. Cascone, and B.S. Davie, "Software-Defined Networks: A System Approach"; System Approach, LLC.

J. Halvorsen, J. Waite and A. Hahn, "Evaluating the Observability of Network Security Monitoring Strategies With TOMATO," in IEEE Access, vol. 7, pp. 108304-108315, 2019, doi: 10.1109/ACCESS.2019.2933415.

Deep Learning: A Visual Approach, Andrew Glassner; No Starch Press

G. Xu, Y. Cao, Y. Ren, X. Li and Z. Feng, "Network Security Situation Awareness Based on Semantic Ontology and User-Defined Rules for Internet of Things," in IEEE Access, vol. 5, pp. 21046-21056, 2017, doi: 10.1109/ACCESS.2017.2734681.

S. Chickerur, A. Goudar and A. Kinnerkar, "Comparison of Relational Database with Document-Oriented Database (MongoDB) for Big Data Applications," 2015 8th International Conference on Advanced Software Engineering & Its Applications (ASEA), 2015, pp. 41-47, doi: 10.1109/ASEA.2015.19.

N. Kumar, A. Leventer and A. Matatyaou, "Monitoring and Troubleshooting at Scale with Advanced Analytics", SCTE Cable-Tec Expo 2021.

Rural 5G Fixed Wireless Access. Economics Analysis and Methodology.

A Technical Paper prepared for SCTE by

Sanjay Patel

Sr. Director & Distinguished Strategist
CableLabs
858 Coal Creek Circle, Louisville CO 80027
303-661-3488
s.patel@cablelabs.com

Dorin Viorel, Distinguished Technologist - CableLabs

303-661-3357
d.viorel@cablelabs.com

Ruoyu Sun, Principal Technologist - CableLabs

303-661-6789
r.sun@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Service Assumptions.....	3
3. Spectrum and Infrastructure Considerations	4
4. Propagation and Capacity	5
5. Subscriber Household Demand	7
6. Serviceable Households	7
7. Economic Considerations	9
8. Conclusion.....	13
Abbreviations	13

List of Figures

Title	Page Number
Figure 1 - Suburban Macro Tower	4
Figure 2 - 5G In-Home CPE.....	4
Figure 3 – 30m Macro Tower.....	6
Figure 4 – 60m Macro Tower.....	6
Figure 5 - 2.6 GHz Propagation & Capacity (100 Mbps Service, 30m Macro Tower).....	6

List of Tables

Title	Page Number
Table 1 - Calculated 2.6 GHz O2I Propagation (distance from base station) for Target Service Speed and Tower Height.....	5
Table 2 - Average Projected Broadband Usage per Household.....	7
Table 3 - Maximum Supportable Households for a 100 Mbps FWA Service	8

1. Introduction

Fixed Wireless Access (FWA) is an alternative to wired broadband services using radio links between two fixed points. With the introduction of 5G, advancements in spectral efficiency and antenna technologies using massive MIMO (Multiple-Input Multiple-Output) and beamforming have become key enablers. The technology and costs have evolved and FWA services are now available from many service providers ranging from niche wireless Internet service providers (WISPs) to the largest mobile carriers. Today's FWA services compete directly with other wired broadband options from telcos (DSL and FTTH) and cable operators (HFC).

While our companion paper, "5G FWA Technical Performance Analysis for Mid-Band Rural Networks", investigates the technical performance characteristics of Rural FWA, this paper seeks to provide a framework to assess the economic considerations for FWA. The methodology involves:

- Defining Rural FWA service assumptions
- Determining spectrum and infrastructure
- Characterizing propagation and available capacity
- Projecting household level broadband demand
- Solving for supportable subscribers
- Pro Forma Economic considerations

FWA can be a strategic opportunity to deliver broadband in unserved and underserved rural areas. The analysis focuses on FWA using mid-band spectrum delivered from macro towers in rural geographies as defined by household density. Broadband usage is expected to continue to grow for the foreseeable future and a FWA service must account for projected demand. Therefore, we forecast average fixed broadband usage at the household level in 2027 and determine how many households can be supported in a typical macro cell at that point in time.

We explore the revenue drivers, variable and fixed cost components and discuss how a provider may consider cost allocations for spectrum and internally across business units. Finally, we develop an Excel model to help analyze projected cash flows and returns on investment based on variable assumptions.

For questions about this paper and/or to obtain a copy of the Excel model with a sample pro-forma, please contact the author at s.patel@cablelabs.com.

2. Service Assumptions

This paper will primarily focus on a target 100 Mbps download/ 10 Mbps upload (headline rate) 5G NR FWA service. The FWA service is assumed to have a 95% availability at the headline speed. Subscribers would still experience a connection at other times, but at lower than headline rate. The analysis assumes indoor customer premise equipment (CPE) requiring a simple self-installation, outdoor-to-indoor (O2I) downstream and indoor-to-outdoor (I2O) upstream signal propagation. Subscribers are assumed to be households with average broadband consumption on a GB/month basis. We model growth in household level broadband consumption at 25% per year and determine how many subscribers can be supported 5-years-out in 2027.

3. Spectrum and Infrastructure Considerations

FWA can be delivered using licensed or unlicensed spectrum ranging from sub-GHz to mmWave. Lower bands such as 600/700 MHz offer much better propagation but lower capacity. Higher bands offer better capacity but with lower propagation. While carriers can leverage their active spectrum and carrier aggregation for FWA, for simplicity, this analysis assumes dedicated spectrum.

We model the use of 100 MHz of 2.6 GHz spectrum on existing macro towers in suburban (30m towers) and rural (60m towers) areas. The main advantage is the ability to use the same infrastructure that is used for mobile service with an added layer of radios for the FWA service. These sites have existing power and backhaul infrastructure and would require incremental capacity to serve FWA needs. The typical clutter in these environments is 1-4 story residential and commercial buildings and representative foliage ranging from shrubbery to mature tree canopies. The 3GPP model values we use incorporate large-scale assumptions for clutter using probability projections for line of site (LOS) and non-line-of-site (NLOS) between transmitting and receive antennas. The end result is a composite propagation value incorporating LOS and NLOS probabilities for the target region.



Figure 1 - Suburban Macro Tower



Figure 2 - 5G In-Home CPE

We model the use of indoor Customer Premise Equipment (CPE). The CPE is assumed to be placed within 1 meter of the closest outside wall to the transmitting tower and at a height of 2 meters in the subscriber home. An alternative scenario, not considered in this analysis, involves the use of an outdoor mounted antenna at or above the subscriber home roofline. The use of outdoor receive antenna can add up to ~40% incremental propagation for a 100 Mbps service in suburban areas and up to 50% in rural. This typically requires a professional install, adds significant cost to the service and can extend the payback period depending on whether the subscriber bears the cost of installation.

4. Propagation and Capacity

CableLabs has been investigating FWA performance since 2018 and has developed a proprietary MATLAB¹ based simulations engine to characterize the performance of different spectrum bands using multiple assumptions. The simulations engine is backed by two technical modeling tools working in tandem; 1) a system level simulator (SLS) based on a 19-cell site topology, which statistically calculates aggregated system interference and 2) a link level simulator (LLS) modeling 5G New Radio (NR) waveforms and multiple-input multiple-output (MIMO) antenna patterns. The LLS uses the results from the (SLS) to evaluate the signal-to-noise ratio (SINR) and throughput versus link distance. The output from these models provides aggregate performance and propagation data to inform this techno-economic analysis. Technical details of the CableLabs SLS and LLS models can be found in the companion SCTE paper.

Table 1 - Calculated 2.6 GHz O2I Propagation (distance from base station) for Target Service Speed and Tower Height

	50 Mbps	100 Mbps	300 Mbps
30m Tower	1,718	1,418	730
60m Tower	4,625	3,600	1,765

Table 1 shows the calculated propagation at 3 different headline speeds for 2.6 GHz spectrum deployed on a 30 meter macro tower in a suburban environment and a 60 meter macro tower in a rural area. The distances can be considered a service edge where connections can still be made beyond the representative values, but at lower than headline speeds. Capacity and throughput diminish the further the FWA subscriber is located from the base station as depicted in Figure 3 – 30m Macro Tower and Figure 4 – 60m Macro Tower, which show available capacity per user in O2I and outdoor receive antenna scenarios respectively. The CableLabs model assumes 3 sectors per cell

¹ MATLAB is a proprietary multi-paradigm programming language and numeric computing environment developed by MathWorks (www.mathworks.com).

and 8 simultaneously active users per sector. It also assumes that subscriber households are distributed evenly across the cell area.

In addition to propagation, the model calculates available capacity for each sector of the macro cell. Radio frequencies lose signal strength and data carrying capacity in correlation to their distance from the gNB. Figure 3 and Figure 4 show available capacity per user at varying distances from 30m and 60m tall macro towers in the case of O2I CPE and with the use of outdoor receive antenna respectively. In the 30m tower case (Figure 3), we see that a 300 Mbps service can be delivered at a distance of 730m and a 100 Mbps service can reach 1418m. Of note, if 300 Mbps is the desired target FWA service, not all households within the area would be considered serviceable.




Figure 3 – 30m Macro Tower




Figure 4 – 60m Macro Tower

We assume that households are evenly distributed across the respective cell coverage areas based on the average density (households/Km²) in the area of interest. The calculation for mid-cell incorporates the cell edge for a 100 Mbps O2I service, 1419 meters, in the case of 2.6 GHz deployed on a 30m tall macro tower. At roughly 70% of the distance between the gNB and the cell edge ($1419m \times .7 = 1003m$), half of households would be between the base station and mid-cell and half would be between mid-cell and the cell edge. Calculated available capacity at mid-cell is a total of 1,429 Mbps.




Figure 5 - 2.6 GHz Propagation & Capacity (100 Mbps Service, 30m Macro Tower)

5. Subscriber Household Demand

We assume that a mobile FWA household will have similar usage on a GB/month as a fixed broadband household. OpenVault² is a 3rd party firm that collects, tracks and reports average fixed broadband consumption per household. OpenVault's Q4 '21 report shows the average household consumed 504 GB/month downstream and 32 GB/month upstream for a total of 536 GB/month. As we are interested in what consumption will be 5 years in the future, we assume an annual growth rate of 25%³ and calculate that by 2027, the average fixed broadband household will consume 2,045 GB/month of data. This calculates as an average of 6.2 Mbps in 2027 (2,045 GB/ 30.4 days/ 24 hours/ 3600 seconds * 8000 bytes to bits). While the average usage in Mbps is useful, operators design their networks to support peak usage. At CableLabs, we have observed that peak usage has maintained a consistent ratio of 2:1 against average usage for a number of years. This implies peak usage per household of 12.5 Mbps in 2027.

Table 2 - Average Projected Broadband Usage per Household

		2021	2022	2023	2024	2025	2026	2027
Total Usage (GB/month)		536						
Annual Growth	25%		670	838	1,047	1,309	1,636	2,045
Mbps Equivalent		1.6	2.0	2.6	3.2	4.0	5.0	6.2
Peak to Average Ratio	2.0X							
Peak Usage (Mbps)		3.3	4.1	5.1	6.4	8.0	10.0	12.5

6. Serviceable Households

We have shown propagation and available capacity for a given suburban macro cell sector and calculated the projected average household demand five-years out in 2027. A straight-forward exercise of dividing available capacity by household demand would yield the maximum number of subscribers that could be supported in the given sector. The resulting value would exhaust all available capacity. Service providers typically operate with a maximum target fill-rate to allow for some spikes in peak consumption but also to allow a margin of time to either manage existing subscribers and policies or to add incremental capacity. For a FWA network, this could involve activating additional spectrum, improving the efficiency of the existing network with more advanced radios and CPE or densifying the network with additional macro towers or small cells.

² www.openvault.com. OVBI_4Q21_Report.

³ 25% represents a reasonable growth assumption based on recent non-Covid period data and can be adjusted as consumption data is reported over time. The same growth rate for upstream and downstream is used for simplicity.

In our example of a 30m suburban macro tower, we use mid-cell capacity and apply a max fill-rate of 50% as the average available across all households in the footprint (1429 Mbps * 50% = 715 Mbps). We recognize that households nearer the gNB could experience significantly higher throughput, depending on the providers' policies and practices of limiting available speeds to any given subscriber. Households at the service edge would typically experience the headline speed subject to the number of active users in the sector.

715 Mbps would support approximately 57 households per sector or 171 household per cell (715 Mbps ÷ 12.5 Peak Average usage/Household in 2027 = 57 X 3 sectors = 171). Our calculations are based on average household consumption. Service providers may have the ability to segment and target users i.e. lower usage households and to manage how many subscribers are offered the service and supported in a given cell sector.

An alternative view of serviceability for FWA is to understand how many households can be supported in a given area. The 100 Mbps service edge in our 30m suburban tower example above is at a radius of 1,418 meters equating to a coverage area of 6.3 Km². With 171 maximum serviceable households per cell we arrive at a maximum of 27 on a per Km² basis. Table 3 below shows the maximum number of households can be supported (in green) for a 100 Mbps FWA service across a range of household densities and market penetrations. Note that where the number of macro towers in an area of interest does not provide ubiquitous coverage for the given service, an adjustment would need to be made to account for non-addressable households.

Table 3 - Maximum Supportable Households for a 100 Mbps FWA Service

Market Penetration	Household Density/Square Km						
	100	200	300	400	500	750	1,000
2.5%	3	5	8	10	13	19	25
5.0%	5	10	15	20	25	38	50
7.5%	8	15	23	30	38	57	75
10.0%	10	20	30	40	50	75	100
12.5%	13	25	38	50	63	94	125
15.0%	15	30	45	60	75	113	150
17.5%	18	35	53	70	88	132	175
20.0%	20	40	60	80	100	150	200
22.5%	23	45	68	90	113	169	225
25.0%	25	50	75	100	125	188	250

The U.S. Census defines an area with less than ~325 households/Km² as rural⁴. We can see from the results in Table 3 that, in general in suburban areas, available capacity would be exhausted in areas up to 500 households/Km² once 5% of households have signed up for the service. At this point, the provider would either need to manage the demand i.e. halt additional subscriber adds or increase capacity by adding more spectrum, increasing the efficiency of the network or adding more towers in the service area.

⁴ <https://www.ers.usda.gov/topics/rural-economy-population/rural-classifications/what-is-rural.aspx>. We use 2.5 persons/household to arrive at approximately ~325 households/Km².

Demand could be tempered by targeting areas with smaller households size (predominance of single-person households) or segments that have lower overall usage. Over-the-top video is a significant driver of data consumption and providers could target areas or segments with higher traditional linear video subscribership or lower rates of work and/or study from home.

7. Economic Considerations

To understand the business case for FWA, a proforma of discounted cash flows with the following assumptions and factors should be considered. To obtain a sample Excel based pro-forma incorporating the factors below, please contact the author at s.patel@cablelabs.com:

Revenue and Top Line:	
Area of Interest	<ul style="list-style-type: none"> • 3GPP considers different performance models for Urban Macro and Rural Macro scenarios. • Can be defined as a market, defined target service area or macro cell.
Service Offering	<ul style="list-style-type: none"> • ARPUs and incentives may vary depending on whether a subscriber bundles mobile and FWA. • Target speed offerings affect the service edge or how far signals propagate. A 300 Mbps service would offer a smaller serviceable footprint than a 100 Mbps service. • Target availability of headline speeds. There is a direct correlation between propagation and target service availability. • Spectrum to be used for the FWA service including depth.
Household Density	<ul style="list-style-type: none"> • Households/Km² can be translated into addressable subscribers once the service and calculated propagation are defined.
Subscriber Penetration	<ul style="list-style-type: none"> • Drives the calculation for demand. • Consider mobile market share and the type and number of available competing broadband options in the area of interest i.e. FTTH, DSL, HFC, FWA, LEOS. • Some MDUs offer broadband as part of their rent or Home Owners Association (HOA) benefits. • Consider that a certain (small) percentage of households will not subscribe to any broadband service.
Subscriber Ramp	<ul style="list-style-type: none"> • Time it will take in years to reach target penetration rates.
Max Supportable Subscribers	<ul style="list-style-type: none"> • Note in Section 4, we describe how to calculate available capacity in a cell and in Section 6 the maximum number of subscribers that can be supported based on average household demand. This serves as a subscriber ceiling until additional capacity can be created.
CPE Revenue	<ul style="list-style-type: none"> • If the subscriber bears either a partial or full upfront cost or monthly lease cost for CPE.

Variable Costs:	
CPE	<ul style="list-style-type: none"> • Cost per unit for new indoor 5G CPE. • Useful life and/or replacement cycles. • Distribution costs i.e. warehousing, shipping, delivery, returns. • Cost to refurbish and package for churned units. • We model CPE cost as being borne by the operator. Subscribers may bear part or all of the cost and those revenues should be accounted for accordingly.
Customer Acquisition	<ul style="list-style-type: none"> • Retail, channel partners, online acquisition paths and associated costs. • Promotions i.e. X months free service and other tie-ins i.e. bundled OTT video services or other discounts i.e fixed and mobile.
Churn	<ul style="list-style-type: none"> • Rate of monthly disconnects. Given the nature of the service, FWA churn may be higher than typical fixed broadband or mobile subscriber churn rates. • For a model that considers a terminal number of subscribers following a ramp up, a service provider would need to replace disconnecting subscribers on a regular basis to maintain counts and account for CPE and subscriber acquisition costs.
Customer Cost to Serve	<ul style="list-style-type: none"> • Customer support including retail, online channels, call centers, back-office systems, field techs, trucks, consumable tools. • Can be expressed as a percent of revenue for simplicity.
Tower leases	<ul style="list-style-type: none"> • Incremental cost to lease space on existing towers for FWA specific radios, and land for power and backhaul.
Power	<ul style="list-style-type: none"> • Incremental cost for radios, networking equipment.
Backhaul	<ul style="list-style-type: none"> • Incremental cost to add capacity or new links.
Maintenance	<ul style="list-style-type: none"> • For new equipment and incremental infrastructure
Licensing and Support Agreements	<ul style="list-style-type: none"> • As applicable from vendors for hardware and software.

Fixed Costs & Other:	
Spectrum	<ul style="list-style-type: none"> • Acquisition cost or calculated as \$/MHz Pop based on the population for the area of interest. • Population for a given area can be referenced or calculated based on household data. A typical household has ~2.3 residents in the U.S. • Depending on licensing terms, where regulators extend spectrum licenses indefinitely, they could be considered an asset for accounting purposes. In that case, a carrying cost calculated as the provider's cost of capital against the acquisition cost could be applied. For example, in the FCC's recent C-Band auction, the average price paid was \$1.17/MHz Pop. If the area of interest was a market with a population of 1 Mil, and the operator dedicates 100 MHz of spectrum for the service, the spectrum for that market would \$117 Mil. A carrying cost using a cost of capital of 8%/year would equal a monthly spectrum expense of \$780,000 (\$117 Mil * 8% /12). Note that spectrum cost can vary market by market. • Another consideration is whether the operator has indefinite rights to the spectrum. In cases where the spectrum is considered owned, its value could be incorporated into the pro-forma as a future cash flow if the service is being evaluated on a standalone basis i.e. the spectrum could be re-allocated for mobile use or otherwise monetized. We do not make consideration of a future cash flow in our analysis.
FWA Antennas, cabling, power, networking equipment	<ul style="list-style-type: none"> • MIMO characteristics for the antennas. • Hardware and installation. • Useful life or technology replacement considerations. • Distributed units, Central units. • Incremental core networking and management.
Tax Savings	<ul style="list-style-type: none"> • In some countries, spectrum and other capital investments can be amortized and are tax deductible for cash tax purposes. A typical amortization period would be straight line over 15 years. • www.csimarket.com offers data on tax rates for U.S. companies.
Cost of Capital	<ul style="list-style-type: none"> • Also referred to as the Discount Rate used for Net Present Value (NPV) calculations. • See NYU Stern for average cost of capital rates across U.S. industries⁵.

⁵ https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/wacc.html

Financial Metrics	<ul style="list-style-type: none"> • Cash Flow: increase or decrease in cash for a period once revenues, variable and fixed costs are accounted. • Net Present Value: The current dollar value of future cash flows discounted at the cost of capital. • Payback: Number of years for discounted cash flows to become positive. • IRR: Internal Rate of Return.
--------------------------	---

A pro-forma designed with the above assumptions would calculate contribution margin and cash flows from which a FWA service can be assessed on a net present value basis and payback basis. For an Excel-based example, please contact the author directly.

A final key consideration depending on the operator is cost allocation. Many FWA operators are primarily mobile service providers. Where they build excess capacity, it makes economic sense to use that capacity for other revenue generating activities such as FWA as long as they don't degrade the mobile experience. In that case, a consideration can be made as to how much of the spectrum and infrastructure costs should be allocated to the FWA service. Our Excel tool allows for understanding the impacts to ROI with different sensitivities for revenue, cost and allocation assumptions.

8. Conclusion

This paper lays out the economic considerations for a rural FWA based broadband service. It is a companion to our technical paper which is also submitted for Tec-Expo 2022 titled "5G FWA Technical Performance Analysis for Mid-Band Rural Networks". We show how calculations for propagation and capacity based on a number of service assumptions can be applied to understand how much coverage can be created for a given macro cell and market area. We further describe the concept of service edge and show how capacity diminishes with distance from the service antennas. A methodology to calculate peak average demand for a broadband household and forecast usage 5-years out is used to understand how many subscribers can be supported in a given cell and sector. Finally, we present the revenue, variable, fixed and other considerations that drive a financial analysis of the economics for FWA. Should you have any questions or would like to obtain a reference Excel pro-forma model, please reach out to the author directly at s.patel@cablelabs.com,

Abbreviations

3GPP	Third-generation partnership project
5G	Fifth-generation technology standard for broadband cellular networks
ARPU	Average revenue per unit
CPE	Customer premise equipment
DSL	Digital subscriber line
FTTH	Fiber to the home
FWA	Fixed wireless access

HFC	Hybrid fiber coax
GB	Giga-Byte
gNB	Next generation node-B
I2O	Indoor to outdoor
LLS	Link level simulator
LOS	Line of site
MIMO	Multiple in multiple out
Mbps	Megabit per second
NLOS	Non-line of site
NPV	Net present value
NR	New radio
O2I	Outdoor to indoor
SINR	Signal to noise ratio
SLS	System level simulator
OTT	Over the top
WISP	Wireless Internet service provider

Scaling a SCTE-224 Policy Decision System to Accommodate Burst Loads Driven by Marquee Events

A Technical Paper prepared for SCTE by

Madhuvanth Gopalan
Principal Engineer 1
Comcast India Engineering Center
Chennai, India
+91 9677189018
Madhuvanth_Gopalan@comcast.com

Timothy Wilson
Software Development Engineer 5
Comcast Technology Solutions
Chicago, IL
720-502-3789
Timothy_Wilson3@cable.comcast.com

Stuart Kurkowski, PhD, Comcast Technology Solutions

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Linear Rights Background.....	3
3. SCTE 224 Context.....	3
3.1. SCTE 224 Media.....	4
3.2. SCTE 224 MediaPoint.....	4
3.3. SCTE 224 Policy.....	5
3.4. SCTE 224 ViewingPolicy.....	5
3.5. SCTE 224 Audience.....	5
4. SCTE 224 Example Logic.....	5
5. SCTE 224 Indexing.....	6
5.1. Scheduling.....	6
5.2. Entitlement.....	6
5.3. Pre-Conversion and Caching.....	7
6. Concurrency Scaling.....	8
6.1. Rate-Limiting.....	8
6.2. Solution Options.....	8
6.3. High Response Time Requests.....	9
6.4. Serverless Function Warm-up.....	10
7. Final Architecture.....	10
8. Conclusion.....	11
Abbreviations.....	12
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 – SCTE 224 Constructs.....	4
Figure 2 – Entitlement Tracking.....	7
Figure 3 – Requests Per Second.....	8
Figure 4 – Peak Loads.....	10
Figure 5 – Architecture View.....	11

List of Tables

Title	Page Number
Table 1 Response Times.....	9

1. Introduction

Internet linear delivery has introduced the need for out-of-band schedule and entitlements information (SCTE 224) alongside event signaling (SCTE 35). The popularity and adoption of SCTE 224, Event Scheduling and Notification Interface (ESNI) is opening new use cases for which the protocol is a great fit. SCTE 224 has proven itself as an efficient and effective means for machine-to-machine communication of out-of-band linear rights management. This combination of SCTE 224 and SCTE 35 to trigger the in-band signaling allows precision execution of linear rights for content substitution and addressable advertising management. But the data maintained by the system in the SCTE 224 format is also useful for other entitlement use cases. The main use case we look at in this paper is around sports entitlement and using the SCTE 224 objects to determine whether a user has the rights to see a game or not – before even seeing the video – based on their location.

With its increase in popularity, content providers and operators are now delivering internet linear television to millions of subscribers simultaneously. As a result, SCTE 224 decision systems can receive millions of simultaneous decision requests from individual playback devices. Scaling these systems is paramount to effective delivery and subscriber satisfaction. One scaling strategy is the deployment of distributed decision systems as “serverless” functions. However, merely moving the decision logic to serverless functions did not achieve performant scale for loads expected for marquee events like major sports championships. Even though these functions exhibited acceptable elasticity in launch, the decision logic was not performant at runtime, because SCTE 224 objects do not lend themselves to expedient searches for applicable viewing policies. The Comcast Technology Solutions (CTS) team optimized the design by indexing the SCTE 224 objects and took advantage of other cloud features to create an optimal and performant distributed decision engine that scales in a ready state, making it instantly responsive across the video subscriber footprint. Our highly performant infrastructure decision system expands automatically to accommodate peak loads. As a result, this distributed decision system delivers 50 milliseconds responses under a load of nearly 20K requests per second. Furthermore, the system scales back under times of lower viewership, reducing our cloud processing cost.

2. Linear Rights Background

The ESNI protocol is perfectly suited to communicate rules and policies at an audience-based level, thereby providing a substrate to implement linear rights. ESNI is the key mechanism to communicate linear rights around web or over-the-top (OTT) embargoes, regional blackouts, and even dynamic advertising. ESNI also provides policies for specific rules at the precise audience level. ESNI therefore facilitates appropriate rights decisioning at scale from machine-to-machine. The ESNI objects are perfect for determining such things as channels and events a user has the rights to watch in real-time.

ESNI objects are XML (Extensible Markup Language) messages with relevant fields for rights such as ViewingPolicy actions for content switching, blackouts, and playout restrictions. These ESNI messages are managed with a REST (representational state transfer) interface for exchange between the providers and the distributors. The out-of-band ESNI execution components then link these markers with the ESNI instructions. The hierarchy of ESNI Policy to ViewingPolicy and to Audiences allows support of various ways to evaluate linear rights against these audiences.

3. SCTE 224 Context

SCTE 224, Event Scheduling and Notification Interface (ESNI) is an XML based standard that provides a defined protocol for carrying machine-to-machine metadata for video. There are five basic constructs within SCTE 224, as shown in Figure 1. These constructs work together to provide a programmer with a

means to convey video rights for content replacement as well as advertising instructions on the distributor or operator side of the workflow. These five constructs are Media, MediaPoints, Policy, ViewingPolicy, and Audiences. We describe each of these here, and then tie them all together with the various use cases around sports entitlement.

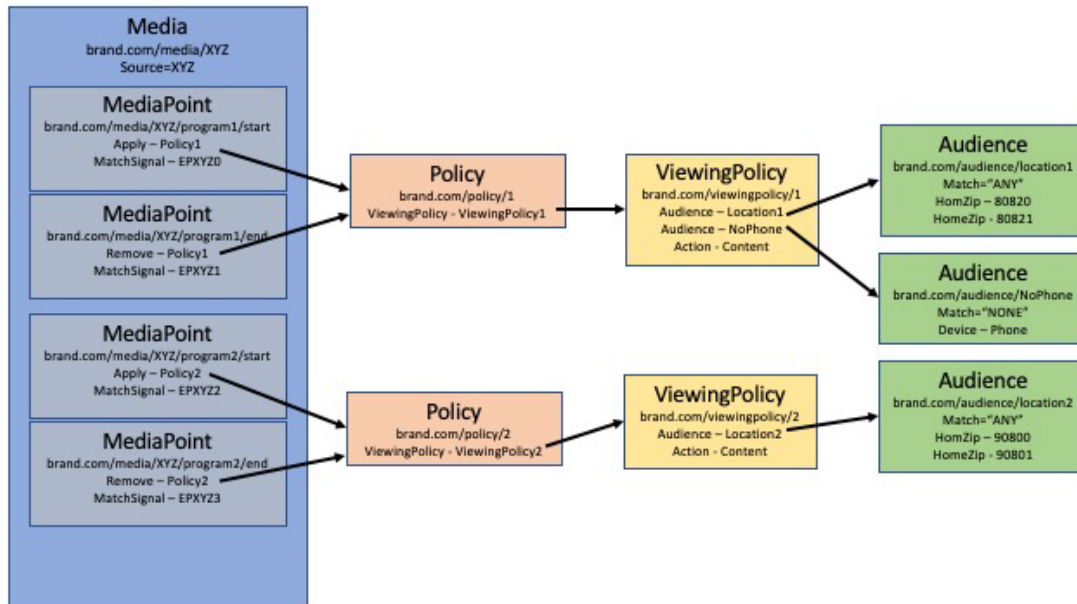


Figure 1 – SCTE 224 Constructs

3.1. SCTE 224 Media

The Media object is a top-level container representing a linear channel whose primary function is carrying all the MediaPoints, so it contains an ordered list of MediaPoint elements as shown in Figure 1. The Media also contains a few key elements like a description and a source for the linear channel it represents.

3.2. SCTE 224 MediaPoint

The MediaPoint object describes a point in the Media when a decision needs to be made or an action needs to be taken. These points can either be time-based (i.e., the presence of a @matchTime attribute in the MediaPoint) or an SCTE 35 in-band signal based for frame accuracy. The signal-based MediaPoints contain a MatchSignal element with XPath matching logic to link the MediaPoint to the presence of the in-band signal. Signals can be reused, because MediaPoints also have an effective/expires window constraining when the MediaPoint can be evaluated.

When a MediaPoint is triggered, based on time or signal, it can either “Apply” or “Remove” one or more Policy objects which affect the state of the linear playback. See the Policy object description below for more details. MediaPoints that “Apply” a Policy do so until another MediaPoint explicitly “Remove” that Policy or they time out based on the duration indicated in the “Apply” statement.

3.3. SCTE 224 Policy

A Policy object is nothing more than a container for defining a set of ViewingPolicy elements that should be acted upon based on this Policy being “Apply” or “Removed” from the Policy stack. The “Apply” or application of a Policy means putting that Policy on that Medias stack via first-in-last-out queue, so multiple Policy objects can be affecting the state at one time. The removal of a Policy then takes it off that stack and out of the state of that Media. SCTE 224 has explicit rules about how to manage the Policy queue in a SCTE 224 execution engine.

3.4. SCTE 224 ViewingPolicy

The ViewingPolicy object is the key SCTE 224 object that associates one or more actions to an audience. These “Actions” can range from telling an audience to go to alternate content, restrict trick mode, or restrict resolution. For advertising specifically, these actions can contain information about the advertising decisioning service (ADS) to use for a particular audience, or various advertisement conflicting rules for a particular audience. The key to a ViewingPolicy is that if the “Audience” criteria is met, then the action must be taken. The SCTE 224 specification maintains a large list of actions allowed within the ViewingPolicy object, many of which are specific to addressable advertising.

3.5. SCTE 224 Audience

The Audience object is a set of characteristics that define a subset of viewers based on criteria such as device-based characteristics (tablet, phone, etc.), general characteristics (local storage, mobile, etc.), or location-based characteristics such as zip codes, postal codes, latitude/longitude, market areas, as well as roles such as distributor or Virtual Integrated Receiver Decoder (vIRD) identifiers that place the audience into groups to dictate specific actions. Audience objects can contain other Audience objects making them a compound Audience. Additionally, logic to associate a client with a viewer or viewers is based on matches of ANY, ALL, or NONE of the characteristics outlined in the Audience, so that characteristics can easily be included or excluded. For example, you can say Match=“ANY” for a list of zip codes to characterize the audience within that area or use Match=“NONE” to characterize an audience outside that area.

4. SCTE 224 Example Logic

So now let’s take those five objects from Figure 1 and run through a scenario of a real-time decision request for the video playout. The trigger is a video signal acquisition system (SAS) seeing an in-band SCTE 35 signal or a user logging into an app. An events triggers the request to a signal decisioning system (SDS) to figure out what it should do. When the SAS calls the SDS it tells the SDS which source it was on, what time it saw the signal, the binary signal, and the client characteristics. Something like “I just saw the signal ‘UhJeafojoihe23edde’ on source XYZ, at 1:00pm and I am encoding for zip code 80820.” Or if it is an app user, they are requesting a determination of entitlement to see if the game they are requesting is or is not on the feed they are requesting it on. The SDS then looks through all the out-of-band SCTE 224 and its Media to find the one for that source (i.e., XYZ). Once it finds the correct Media it looks through its MediaPoints to find the list of MediaPoints that at that time falls within their effectiveness window. Once it has the list of MediaPoints it evaluates each one to see if there is a match with the signal. For the MediaPoint that matches, the decision is either to “Apply” or “Remove” the associated Policy. It then goes from that Policy to the ViewingPolicy where it gets the Audience and sees if the client’s zip code is in or out of that Audience or not based on the Match criteria. If it is in that Audience, the SDS returns the “Action” of the ViewingPolicy to the SAS. In the case of alternate content, it might return to the SAS that it needs to switch over to another source and start encoding the alternate

source. Based on the client information and the construction of the Audience objects, API logic in this use case can give an entitlement answer to a requester.

5. SCTE 224 Indexing

As mentioned in the Introduction, SCTE 224 objects do not lend themselves to expedient searches. To provide a performant service, the SCTE 224 objects need to be transformed into some other form that does support expedient searches. Looking at the types of requests to be processed, there was no single type of transformation sufficient to provide the required performance. Thus, we needed multiple indexing schemes, with each type of request using an indexing scheme that best suits it for the SCTE 224 object.

This is a key concept, because even at the expense of potentially storing the same data in multiple structures, the executional performance outweighed the duplication.

5.1. Scheduling

For linear playout channels, an SCTE 224 Media object typically represents a schedule for that channel, with MediaPoint objects representing individual shows or events. The SCTE 224 specification also allows for metadata about these shows to be contained within the MediaPoint object. Taken together, these two features allow a scheduling interface to determine what is currently playing and to gather metadata on a particular show or event.

The SCTE 224 specification does not make any comments about how MediaPoint objects are stored within a Media object, just how the document order of these MediaPoint objects dictate their precedence. As the MediaPoint objects are stored as a list within the Media object, doing a linear search to determine what is currently playing would take too long on average. For this system under discussion in this paper, the MediaPoints were set up in a true schedule fashion, with earlier events closer to the beginning of the Media object and later events farther down in the document. Additionally, the shows were set up to eliminate any overlap. These restrictions on the MediaPoint objects lead nicely to using a form of self-balancing binary search tree, such as an AVL tree, using time as the search criterion.

Searching for the metadata on a specific show would not lend itself well to using the same binary search tree described above. Therefore, a different form of indexing needs to be applied for this type of request. With the caller supplying the identifier for the show, a hash map object mapping the identifier to a MediaPoint would provide the desired performance.

5.2. Entitlement

Once the information for what is playing and the metadata has been gathered, now it's time to determine if the user is entitled to view the show or event. As with the request for gathering the metadata, the system is provided with the identifier for the show, lending this type of request to rely on an underlying map structure to pinpoint the MediaPoint object. However, unlike the metadata request, these entitlement requests are required to follow the chains of Policy, ViewingPolicy, and Audience objects to make the appropriate decision. Generally, both the Policy and ViewingPolicy objects are succinct, allowing them to be stored as simple data objects. The Audience objects, on the other hand, could contain a list of thousands of zip codes, calling for a more complex structure such as a hash map.

The flow for satisfying an entitlement request would look something like this:

1. Using the show identifier provided in the request, locate the desired MediaPoint object.
2. For each Policy object listed in the Apply elements of the MediaPoint object:

- a. Locate the named Policy object, if it is only supplied as a reference in the MediaPoint object
- b. For any ViewingPolicy associated with this Policy:
 - i. Locate the named ViewingPolicy, if it is only supplied as a reference in the Policy object.
 - ii. For each Audience associated with the ViewingPolicy object, use the zip code provided in the request to look for a match.
 1. If a match is found, use the ViewingPolicy to determine the action for a member of this Audience.
 2. If no match is found, continue searching until all Policy, ViewingPolicy, and Audience objects for the MediaPoint have been examined.
 3. If there is a match in any other Policy, it is handled like step #1 above.
 4. If there are still no matches, then nothing is found, and a default “Not Entitled” response is given.

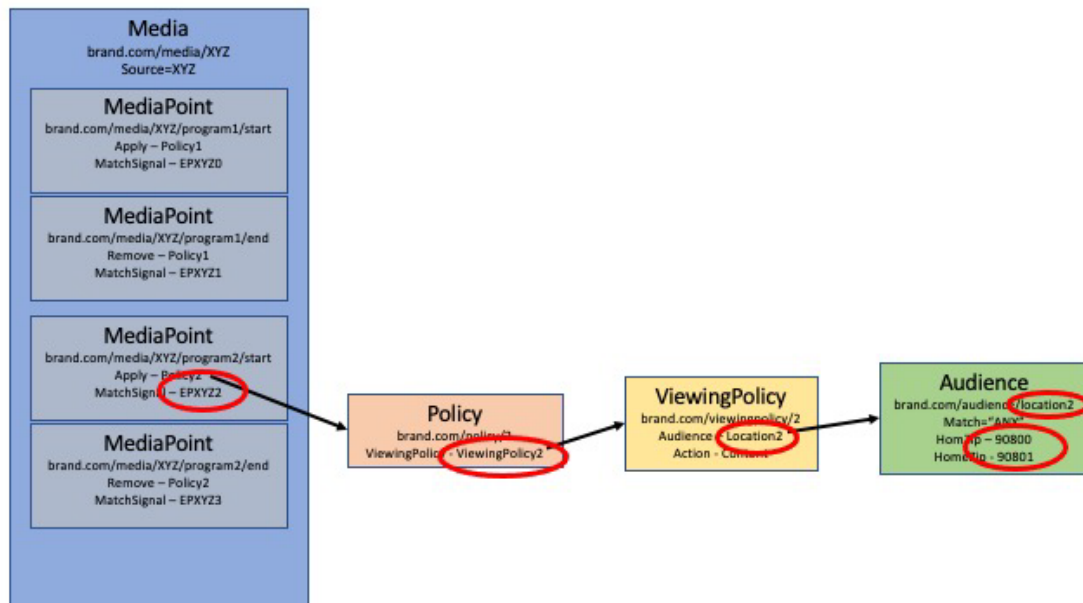


Figure 2 – Entitlement Logical Decision Flow

5.3. Pre-Conversion and Caching

As serverless functions used to supply this functionality are expected to be ephemeral, converting the raw SCTE 224 to structures optimized for SCTE 224, such as AVL or hash maps as described above upon start up excessively impacts the performance of the initial requests. The system was designed to do these conversions externally to the serverless functions providing the searching capabilities. When any update is received by the system, a processor generates the various structures that form the optimized SCTE with the updated data and stores it in a cache that is shared across all instances of the request-handling serverless functions.

6. Concurrency Scaling

Concurrency scaling is what controls the number of concurrent serverless functions that can be set up to handle all of the calls in a scaled event. As the number of connections reaches the limit, throttling of requests can occur during requests that are held until they can be serviced.

6.1. Rate-Limiting

The design choices explained in the indexing section of this paper laid the foundation of this highly performant system and during various load tests that were performed, the system scaled reasonably well until it hit a certain burst load, and requests were being limited.

Most serverless functions are not infinitely scalable and therefore have a configured concurrency limit. There are two types of concurrency limits with serverless functions.

- There is a limit on the number of function containers that can serve requests at a given time. In most cases this value is configurable.
- There is also a burst concurrency limit which is the number of new function containers that can be spun up at a given point in time. This is set at 3,000 in the area these systems were operating in.

In each of the cases above, if either limit is reached during execution, any more requests that need additional serverless functions would not be served. This is reflected during execution as limiting those requests so that available resources can serve the current requests.

The metrics graph below shows, as many as ~16k requests @~11.47AM were throttled, because we hit the burst concurrency limit of 3,000 instances.



Figure 3 – Requests Per Second

6.2. Solution Options

Because the burst levels of requests for these extreme events requests were running up against the limiting, we needed a way to allow the design to work. The goal then was to raise this concurrency limit in some fashion, but there are constraints on the limit, so we looked at two options:

- Increasing the concurrency limit of serverless compute serve. But our functions were being limited because they hit the burst concurrency limit, which is non-configurable.
- Take a closer look at behavior of the functions themselves and identify opportunities for optimization.

It was not enough to just throw more resources at it, because there were cost prohibitions for the majority amount of time when the system is not under load. Since the first option is not configurable, our course of action was to look at the second option around function behavior.

6.3. High Response Time Requests

We decided to take a closer look at:

- What are some of our slowest requests?
- Which part of the code are they spending significant time on?

We already had built some custom instrumentation in our functions, which were basically measuring time taken by different sections of the code and logging them in a structured format to analyze. Querying those logs using analysis tools gave us insights allowing us to answer the two questions above.

#	total_time_ms	db_conn_ms	execution_time_Fin_	execution_time_TreeBySta_
▶ 1	4127.7085	2005.5475	0.00173	33.6695
▶ 2	4100.0024	2002.9073	0.007616	43.5819
▶ 3	4089.9966	2003.8354	0.00227	51.5729
▶ 4	4088.0203	4002.5376	0.002546	64.7034
▶ 5	4085.9028	4003.8772	0.002401	40.0637
▶ 6	4084.3572	4003.59	0.00249	47.2985
▶ 7	4077.8228	4002.887	0.001731	31.4863
▶ 8	4077.6042	4002.8762	0.002939	34.8527
▶ 9	4076.252	3.3288	0.002005	38.6671
▶ 10	4075.7478	4002.984	0.002004	71.1681
▶ 11	4074.7073	4003.7002	0.002398	33.3264
▶ 12	4074.5022	4003.2598	0.0018	30.7119
▶ 13	4071.6487	4002.935	0.001934	36.5997
▶ 14	4071.4495	2002.521	0.002003	37.9416
▶ 15	4071.1333	4002.9124	0.002255	36.8048
▶ 16	4070.7996	2003.0397	0.001996	37.8136
▶ 17	4070.7488	4003.3252	0.00223	42.2515
▶ 18	4069.8464	4002.905	0.002119	38.5778
▶ 19	4069.3198	3.2697	0.005251	36.5757
▶ 20	4068.2205	4003.4956	0.001871	37.597
▶ 21	4067.6438	4003.4858	0.002094	30.18
▶ 22	4067.1667	2002.8424	0.007691	31.651

Table 1 Response Times

It allowed us to see a linear correlation between high processing times and:

- i. Time taken to forward logs to a centralized logging system of choice. Yes, we were using two logging systems; one allowed us to correlate log-events from this system with related log-events from other systems that weren't using central logging tools.
- ii. Time taken to establish connections to the database. A burst of requests meant many different serverless function containers were requesting "new" connections at the same time causing congestion on the server side.

We replaced logging synchronously to the second logging system with logging asynchronously. All the events were logged to the central logging tool, then forwarded to the second using another independent

serverless function.

To address (ii), we implemented a connection pooling mechanism and initialized connections to a database in the initial phase of serverless function, rather than on-demand.

6.4. Serverless Function Warm-up

With all those performance advances described above, there was a significant improvement in response times. However, given that the expected scaling for an event like the Super Bowl will be anything but linear (i.e., tens of thousands of users logging in to watch the broadcast just before the game starts), we knew we might still hit the non-configurable *burst limit* leading to requests being throttled. But an advantage is that it's easy to predict when to expect the burst (before game start).

One aspect of the request-processing life cycle that wasn't in our control, but was consuming significant amounts of time, was the server-less function container initialization time. In other words, the time required to perform a cold start of a new serverless function.

To avoid being limited by cold starts so that bursts could be gracefully handled, we leveraged the provisioned concurrency feature of the server-less function. This instructed the server-less function to spin up X number of serverless function containers, ready to handle the burst of requests as they come.

All of the sound design choices, optimizations for less-efficient parts of the code, and enabling provisioned concurrency allowed us to achieve a P95 response times of <50 milliseconds, when the system was under a peak load of 16-18k requests per second.




Figure 4 – Peak Loads

7. Final Architecture

So based on the information and the success of the above optimization, here is a view of an example overall architecture for an optimized solution.




Figure 5 – Architecture View

8. Conclusion

With its increase in popularity, content providers and operators are now delivering linear television over the internet to millions of subscribers simultaneously. As a result, SCTE 224 decision systems can receive millions of simultaneous decision requests from individual playback devices. Scaling these systems is paramount to effective delivery and subscriber satisfaction. One scaling strategy is the deployment of distributed decision systems as “serverless” functions. However, for us, merely moving the decision logic to serverless functions did not achieve performant scale for loads expected for marquee events like the major sports championships. Even though these functions exhibited acceptable elasticity in launch, the decision logic was not performant at runtime, because SCTE 224 objects do not lend themselves to expedient searches for applicable viewing policies. We optimized the design by indexing the SCTE 224 objects in multiple methods and took advantage of other cloud features such as scaling concurrency and warming up serverless functions before use, to create an optimal and performant distributed decision engine that scales in a ready state, making it instantly responsive across the video subscriber footprint. Our highly performant infrastructure decision system expands automatically to accommodate peak loads. As a result, this distributed decision system delivers 50 milliseconds responses under a load of nearly 20K requests per second. Furthermore, the system scales back under times of lower viewership, reducing our cloud processing costs.

Abbreviations

ADS	ad decisioning service
API	application programming interface
AVL	Adelson-Velskii and Landis binary search tree
ESNI	Event Scheduling and Notification Interface
IRD	integrated receiver decoder
REST	representational state transfer
SAS	signal acquisition system
SCTE	Society of Cable Telecommunications Engineers
SDS	signal decisioning system
vIRD	virtual integrated receiver decoder
XML	Extensible Markup Language

Bibliography & References

SCTE 35, Digital Program Insertion Curing Message for Cable, 2020. https://scte-cms-resource-storage.s3.amazonaws.com/ANSI_SCTE-35-2020-1619708851007.pdf

SCTE 224, ESNI, Event Scheduling and Notification Interface 2021. <https://scte-cms-resource-storage.s3.amazonaws.com/SCTE-224-2021-1620314764331.pdf>

W3C XML Base (Second Edition). W3C Recommendation 28 January 2009.
<http://www.w3.org/TR/xmlbase/>

W3C XML Schema Part 2: Datatypes Second Edition. W3C Recommendation 28 October 2004.
<http://www.w3.org/TR/xmlschema-2/>

W3C XML Path Language (XPath) 2.0 (Second Edition). W3C Recommendation 14 December 2010.
<http://www.w3.org/TR/xpath20/>

AVL Tree: https://en.wikipedia.org/wiki/AVL_tree

Scaling DAA: Automated Network Health Check for vCMTS Platform

A Technical Paper prepared for SCTE by

Marissa Eppes

Data Scientist

Comcast

1800 Arch St, Philadelphia, PA 19103

Marissa_Eppes@comcast.com

Ilana Weinstein

Data Scientist

Comcast

1800 Arch St, Philadelphia, PA 19103

Ilana_Weinstein@comcast.com

Matthew Stehman

Senior Data Scientist

Comcast

1800 Arch St, Philadelphia, PA 19103

Matthew_Stehman@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
1.1. Problem Statement	4
1.2. Solution	5
2. Background	5
2.1. DAA Topology	5
2.1.1. The PPOD	6
2.1.2. The RPD	6
2.1.3. The CM	7
2.2. DAA Telemetry	7
3. Methodology	7
3.1. Network Health Check Overview and Terminology	7
3.1.1. Pre-Check, Post-Check, and Metric Snapshots	7
3.1.2. Service-Affecting vs. Non-Service-Affecting Updates.....	9
3.2. Assessment of Key Performance Indicators	9
3.2.1. Deciding When to Alert	10
3.2.2. Algorithms and Thresholds	11
3.2.3. Custom Algorithm Example – Partial Service	14
3.3. Integration with Software Deployment Automation	15
3.4. Health Check Cloud Environment	16
4. Discussion	17
4.1. Usage Analysis	17
4.2. Analysis of Check Results.....	18
4.3. Lessons Learned.....	19
5. Future Work.....	20
5.1. Continued Optimization of Network Health Check: Mid-Split Updates	20
5.2. Continuous Monitoring	21
5.2.1. Implementation.....	21
5.2.2. Expectations.....	22
6. Conclusion.....	23
Abbreviations	24
Bibliography & References.....	25

List of Figures

Title	Page Number
Figure 1: A simplified view of DAA topology	6
Figure 2: Conceptual relationship among pre-check, post-check, and historical/pre-deployment/post-deployment snapshots	8
Figure 3: Example views of CM status vs. time during SA updates at the RPD-level	9
Figure 4: Some conceptual check outcomes	10
Figure 5: Visuals demonstrating statistical partial service algorithm	15
Figure 6: CI/CD integration with health check API — an example workflow	16
Figure 7: Diagram of cloud environment.....	17
Figure 8: Health check usage	18
Figure 9: Health check results	19
Figure 10: Number of KPIs evaluated in health check over time.....	20
Figure 11: Continuous monitoring workflow.....	22
Figure 12: Proof-of-concept anomaly and model in development	22

List of Tables

Title	Page Number
Table 1: Summary of PPOD-Level KPI Algorithms and Thresholds.....	12
Table 2: Summary of RPD-Level KPI Algorithms and Thresholds	13
Table 3: Summary of CM-Level KPI Algorithms and Thresholds	13

1. Introduction

As Comcast progresses through Gen2 of the Distributed Access Architecture (DAA) initiative, “automating everything” is paramount to the continued growth of the footprint. With the exciting milestones achieved in Gen1, which took the first DAA customers online, came a laundry list of operational improvements needed to realistically achieve the desired scale. The common theme among these needed improvements? Automate. So far, Gen2 has made considerable progress in this regard, automating everything from individual virtual cable modem termination system (vCMTS) cluster standups to software and network changes, to incident detection and mitigation (Krishnamurthy & Medders, 2021).

Since virtualization is the name of the game in the world of DAA, the vast majority of vCMTS maintenance and upkeep is achieved with cloud component software updates, specifically component configuration changes. Achieving near total automation involved transitioning software upkeep over to a DevOps approach, with which changes to individual clusters are managed with cluster-specific continuous integration/continuous deployment (CI/CD) pipelines. With this setup, a single Git commit indicating a configuration change to a specific vCMTS component schedules the entire deployment process, which kicks off a cascade of automated events, including but not limited to silencing alarms, microservice performance checks, the actual software deployment, and a network health check to ensure no customer impact. The goal in using this strategy to orchestrate software deployments is to eliminate as much human interaction as possible, as even the slightest human error introduced at a single step can have monumental impacts to the cluster configuration and network performance down the line and can lead to complex outages and delay feature enhancements and releases.

During software updates, preserving or improving the customer experience is of the highest priority. In Gen1, network health was checked manually; an operator would verify no customer impact by checking telemetry dashboards and manually flagging any signs of service degradation. Not only is this practice not scalable but subjecting this critical process to the risk of human error is undesirable as DAA scales and the customer base grows. Therefore, automating and optimizing network health monitoring surrounding software updates is of utmost importance. This paper highlights a data-driven approach to developing and productionizing an automated network health check for use in vCMTS deployment CI/CD pipelines as part of the Gen2 DAA initiative.

1.1. Problem Statement

As the number of software updates needed to maintain and scale the DAA footprint increases from tens to hundreds to sometimes thousands per day, it is necessary to not only have these automated processes in place, but to ensure that they are optimized to near perfection. With millions of customers already converted to DAA, automated network health monitoring is one of these essential processes, given that software updates can sometimes cause unintended side effects on the network, resulting in a degraded customer experience. Software-related service degradation might manifest as anything from interruption of service to poor traffic throughput, to partial utilization of network capabilities, and more. Even though most software updates pose little risk to the customer experience, catching these occurrences when they do happen so that quick mitigative action can be taken is considered to be a critical capability. Therefore, there is a need for an automated and dependable tool for post-deployment network monitoring, which can validate that software updates do not degrade service for the existing customer base *or* flag the occasional software updates that do.

To deliver true value and support the goal of total automation, this tool needs to be compatible with the vCMTS deployment CI/CD pipelines and to deliver highly accurate results, alerting the ops team only

when there is definite service degradation detected following a deployment, while simultaneously maintaining a level of sensitivity as to not let any undetected service degradation slip through the cracks. Additionally, to adhere to target maintenance window timeframes, this tool must determine and deliver a decision on the state of the network back to the CI/CD pipeline within a matter of minutes following the deployment. Last but not least, if service degradation is detected, the tool must provide a meaningful summary detailing the reasons for service degradation so that an operator can take the appropriate action.

1.2. Solution

To address this need, the data sciences team has leveraged the rich network telemetry available within the DAA system to develop a network health decision engine made available as an application programming interface (API). When called, the API queries cluster-specific telemetry metrics live from a time series database and performs a suite of algorithms to assess the health of the network. The API response flags any incidental impact on already-live customers and alerts an operator within a matter of minutes.

The data sciences team has collaborated closely with the DAA engineering and ops teams to identify relevant network key performance indicators (KPIs)—all in accordance with Data Over Cable Service Interface Specification (DOCSIS[®])—to check as part of this process. The application performs *i.* a pre-check to gather a baseline measure of network KPIs prior to a software update and *ii.* a post-check to draw comparisons among KPIs and ensure that the network remains in a healthy state following a software update. The post-check assesses the state of the network instantaneously by comparing instant post-deployment KPI readings to either instant pre-deployment or historical KPI readings. The API is designed specifically to integrate with the vCMTS software deployment CI/CD pipelines and is robust enough to offer a one-size-fits-all solution to all clusters, regardless of configuration differences, number of customers, differences in downstream topology, etc. The API is also compatible with a variety of software deployment types, regardless of the target component and predicted risk level.

The API is currently integrated and running in a production environment. The application is invoked during each software update for a variety of vCMTS components ranging from server builds to configuration updates to operating system upgrades. As measured in recent analytics, the application is invoked, on average, ~400 times per day and triggers alarms on ~2.5% of invocations. This paper will take a deep dive into the methodology used to develop the tool and tune the rule-based algorithms, present performance metrics, discuss lessons learned, and briefly touch on relevant future work.

2. Background

2.1. DAA Topology

While DAA offers a technologically progressive means of providing service to customers, the distributed architectural setup is rather complex. The access network contains a variety of physical and logical components ranging from headend Kubernetes servers to a series of leaf-spine switches to downstream digital nodes all the way down to the customer premises equipment (CPE). When one or more of these vCMTS components undergoes a software update, the deployment CI/CD pipeline performs a series of checks on the entire cluster footprint. These checks can be grouped into two general categories: *microservice performance* and *network health*. Because this paper focuses on the latter, the intricacies of the DAA architecture and the vCMTS cloud environment will be outside the scope of this paper. However, several past papers cover these topics in detail, namely *Distributed Access Architecture Is Now Widely Distributed - And Delivering On It's Promise* (Howald et al., 2021) and *Node Provisioning and Management in DAA* (Gaydos et al., 2018).

To understand how the automated network health check is performed from a customer impact standpoint, three topological entities must be understood: *i.* the physical point of deployment (PPOD), *ii.* the remote PHY device (RPD), and *iii.* the cable modem (CM). Figure 1 portrays a simplified view of relevant DAA topology.

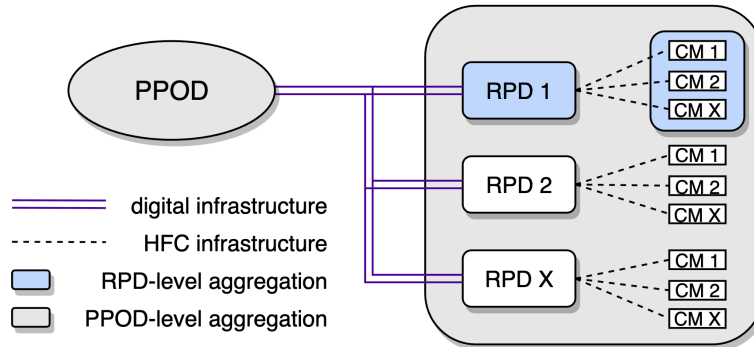


Figure 1: A simplified view of DAA topology

2.1.1. The PPOD

Located at each primary headend, PPODs comprise the actual servers on which DAA software is deployed. PPODs and vCMTS clusters are often referenced interchangeably and, in theory, describe the same technology. However, usage of one term versus the other depends on the context in which the technology is being discussed; a PPOD can be thought of as an abstract deployment unit, whereas “vCMTS” is often used in reference to the physical cluster hardware. Given that this paper discusses DAA from a software standpoint, we will largely use the “PPOD” naming convention going forward.

As Krishnamurthy and Medders discuss, clusters are often spun up with configurational differences, “each with their own slightly different personalities” (2021). Given that any two PPODs might be configured differently, the PPOD is the highest-level architectural component on which it makes sense to perform a series of automated checks and aggregate results. Even when the entire DAA footprint needs to undergo a particular component update, performing PPOD-level checks eliminates any risk of confounding configurational differences into the equation. As such, all DAA deployment CI/CD pipelines are kicked off at a PPOD-level, and subsequent checks are intended to indicate how a particular PPOD “personality” fairs with any given software update. From the network health perspective, this entails checking for any service degradation or incidental impact experienced by customers downstream of the PPOD of interest.

2.1.2. The RPD

Designed to bring digital data transmission as close to the home as possible, RPDs sit on the very edge of the access network and comprise the gateway between the digital system and the hybrid fiber/coax (HFC) network, which eventually reaches the home. RPDs are the most downstream digital component of the access network and can be subject to software updates themselves. Like PPODs, RPDs can have configurational differences and/or come from different vendors, creating variety among them. Therefore, RPD-level data aggregations provide diagnostic value when conducting the network health check. The strategy of grouping together customers downstream of each RPD and analyzing each of these customer subsets separately provides a more thorough network health validation and, if service degradation is noticed, helps the responding operator to determine if the source of the issue lives at a particular RPD.

2.1.3. The CM

The CM, which is often used interchangeably with “CPE”, is the final topological entity involved in the network health check. Individual CM metrics are often the fundamental units used to derive data-based network health algorithms, and all CMs subscribed to the PPOD-of-interest are taken into account during the assessment. However, with millions of CMs already dispersed across the DAA footprint, it would be impractical and disadvantageous to collect, analyze, and report on telemetry data for each and every CM. In analyzing CM metrics, aggregations at both the PPOD- and RPD-levels deliver a more meaningful and statistically robust measure of how the customer experience is fairing across the footprint following a software update. Additionally, if service degradation is detected, this aggregation strategy offers a diagnostic advantage in that it can pinpoint a lowest common ancestor for problematic CMs, allowing for quicker isolation of the issue. For these reasons, almost all algorithms operate at either the PPOD- or RPD-level. The only exception to this aggregation strategy is seen in the analysis of Business Services over DOCSIS (BSOD) customers. This analysis does perform a CM-level evaluation, which will be presented in Section 3.2.2.

2.2. DAA Telemetry

One noteworthy enhancement of the DAA system is the improved real-time telemetry offered across each of the architectural components from the PPOD down to the CM. With the legacy system, telemetry data was only emitted at 5-minute intervals at the very least. With DAA, data is streamed with 15-second resolution. This improvement is key to delivering a speedy assessment of network health, as it eliminates the need to wait for post-deployment telemetry to become available. Additionally, this increased resolution into telemetry data allows the freedom to explore more check algorithms and provide higher quality network diagnostics. In *Solving The Mysteries of the Distributed Access Architecture*, Stehman et al. details the available telemetry across the access network and further expands on relevant topology (2021).

3. Methodology

3.1. Network Health Check Overview and Terminology

3.1.1. Pre-Check, Post-Check, and Metric Snapshots

The automated network health check is meant to detect and alert on any unintended customer impact after a software deployment. This of course entails checking live telemetry metrics directly after the deployment; however, to establish a PPOD-specific baseline on which to compare these post-deployment metrics, data collection and analysis is also needed prior to the deployment. As such, the health check was designed to be a two-part process, consisting of both a pre-check and a post-check. By design, a completed pre-check is required to start a post-check. From the PPOD CI/CD standpoint, this means that the API must be integrated both prior to and after the deployment.

Each of the pre- and post-checks can be further broken down into metric “snapshots”. A metric snapshot can be defined as a collection of related telemetry metrics measured over the same time period and compiled to create a meaningful KPI. Since most of the evaluation algorithms function by comparing a given KPI at two different time points, each time-aligned metric snapshot can be considered a single comparison unit. Pre-check gathers metric snapshots from two different timeframes: *i.* over the course of history for the PPOD and *ii.* in the instant right before the deployment. Respectively, we will refer to these as “historical snapshots” (“historical-snaps” for short) and “pre-deployment snapshots” (“pre-snaps” for short). As historical-snaps often consider a wide timeframe prior to a deployment, typically an aggregation of telemetry data over time serves as the comparative metric value. This could be, for instance, the 10th percentile of a PPOD’s upstream packet rate over the past seven days.

The post-check works by querying and deriving metric snapshots from one timeframe only—in the instant right after the deployment. These snapshots will be referred to as “post-deployment snapshots” (“post-snaps” for short). The post-check also has the capability to load any of the cached historical-snaps or pre-snaps from the pre-check to perform the actual pre-to-post comparisons. Whether or not post-check borrows from the historical-snaps or pre-snaps depends on the KPI being compared. As a general rule of thumb, continuous KPIs, such as traffic flow, tend to be compared post-deployment-to-historical, whereas discrete KPIs, such as number of CMs online, tend to be compared post-deployment-to-pre-deployment. In adherence to DevOps daily maintenance schedules, the target completion time of a post-check, including all queries, data reads, comparison logic, and data writes, is two minutes or less. Figure 2 aims to further clarify the concepts of pre- and post-checks as well as metric snapshots.

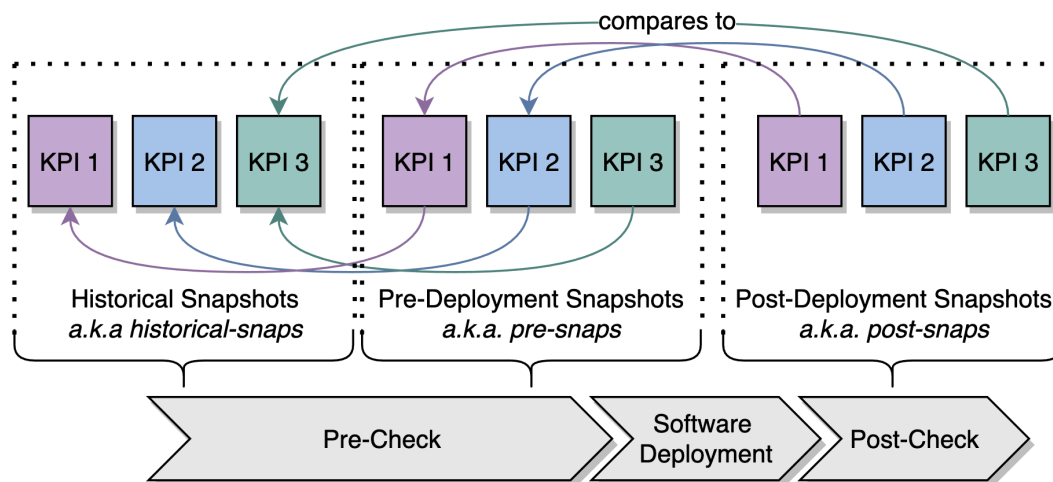


Figure 2: Conceptual relationship among pre-check, post-check, and historical/pre-deployment/post-deployment snapshots

It should be noted that the pre-check currently possesses the capability to compare pre-snaps to historical-snaps and assess network health prior to the deployment. The initial concept of the automated network health check was designed with this capability in mind so that a deployment could be automatically blocked if the network was deemed to be unfit for a software update. As the DAA initiative progressed, this feature was decommissioned, as there are other automated alert systems in place capable of blocking a deployment if the network is not considered healthy enough to undergo an update. Therefore, the main purpose of the pre-check currently is to gather the snapshots needed to perform comparisons in the post-check and to confirm that these snapshots are complete. In other words, the software deployment automation only interacts with network health assessments made in the post-check. Since the data sciences team has primarily focused on developing and optimizing the post-check functionality, this paper will focus on the assessment of KPIs in the post-check.

It should also be noted that the data sciences team’s network health check is not the only post-deployment fail-safe in place within the deployment automation. The DAA engineering team has also incorporated a series of automated checks, more so pertaining to cloud microservice health; however, there can be slight overlap when it comes to the KPIs observed among the various checks. While these microservice checks are deemed to be another essential step within the deployment automation pipeline, they will remain outside the scope of this paper.

3.1.2. Service-Affecting vs. Non-Service-Affecting Updates

All DAA software updates, regardless of the component being updated, can be placed into one of two general categories: service-affecting (SA) or non-service-affecting (NSA). The distinction between the two depends on whether a component update will take customers offline for a brief period. SA updates, which are scheduled during maintenance windows (normally between 01:00 and 04:00 headend local time), typically involve RPD reboots causing all downstream customers to experience a brief expected service interruption. NSA updates, which comprise the vast majority of all DAA software updates, occur on components that often have service-preserving backup units, and will most likely not result in interruptions. In other words, SA updates can be considered high-risk, whereas NSA updates can be considered low-risk. The same automated network health check is performed on both SA and NSA deployments; however, it is important to make this distinction between the two categories when analyzing results, as SA updates tend to trigger more alarms than NSA updates, given that it takes some time for customers to come back online and the network to return to a normal, steady state following SA updates. As SA updates are high-risk and are purposely scheduled on just a handful of PPODs at a time, an operator typically oversees SA updates and interacts with the health check response live. Figure 3 shows examples of typical CM behavior over time during an SA update—specifically the steep drop-off in total number of CMs online during the RPD reboot, followed by a gradual recovery. The expectation for NSA updates is that a service interruption like this should *not* occur.




Figure 3: Example views of CM status vs. time during SA updates at the RPD-level

3.2. Assessment of Key Performance Indicators

The automated network health check, specifically the post-check, works by making rule-based comparisons between two snapshots of the same KPI. Thresholds and comparison algorithms are specific to each KPI and will be detailed in this section. When a particular KPI is evaluated against a threshold and breaks a rule, the KPI is assigned one of two categories: *warning* or *failure*. Whether or not a broken rule indicates a warning or failure depends on the KPI as well as the algorithm being used to assess it, both of which indicate the severity of the anomaly. Warnings can indicate signs of service degradation, but typically to a lesser extent than failures. As will be discussed in the next section, KPIs with warnings are intended to be observed and studied more closely but not quite considered severe enough to automatically alert an operator. The data sciences team has worked closely with DAA subject matter

experts (SMEs) in deciding which KPIs should belong in the failure category versus the warning category.

In addition to warnings and failures, there are a several KPIs which are calculated and presented in the API response for information only, as they are not indicative of service degradation nor are they held to a certain threshold. These KPIs are placed in a category called *info* and are meant to aid in observation and analytics.

3.2.1. Deciding When to Alert

The API response compiles all individual KPI results and provides a list of all KPIs which have failed the check and all KPIs which should deliver warnings. If a single KPI evaluation fails, the result of the entire health check is considered a failure, the PPOD is flagged, and an operator is alerted to take further action. KPIs in the warning category do not automatically fail the entire health check when a rule is broken. Rather, these KPIs are manually observed by DAA operators and SMEs as possible contenders for stricter treatment, further algorithm tuning, or even live operator intervention in the case of SA updates. It is not uncommon for KPIs to start out in the warning category and later be transferred over to the failure category when SMEs confirm the relevancy of the KPI, and an optimal algorithm has been decided. Figure 4 aims to conceptually demonstrate how the overall health check decision is calculated based on KPIs in the failures category *only*, despite KPIs in the warnings category.

Alongside the overall health check decision, a summary of the KPI readings and comparisons is presented to the operator as a diagnostic aid. For PPOD-level evaluations, a summary might include, for instance, CM counts for given states pre- and post- deployment and list any CMs which have changed state following a deployment. If an RPD-level evaluation fails, typically the summary will detail KPI readings and comparisons for the problematic RPD(s).

With this information, an operator can decide on a mitigative course of action in the event that service degradation is detected. For instance, an operator might perform an RPD reboot and retry the post-check to see if detected service degradation has been fixed. If nothing can be done remotely, an operator might decide to send a technician out into the field to intervene with RPD hardware. If a root cause cannot be identified, the operator might ultimately decide it is best to roll back the software update until further troubleshooting can be performed.

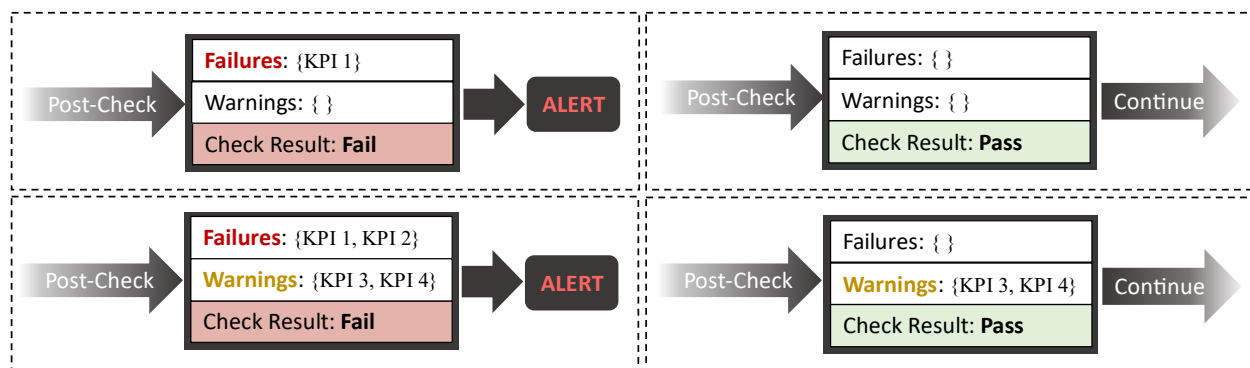


Figure 4: Some conceptual check outcomes

3.2.2. Algorithms and Thresholds

As discussed previously, most algorithms operate at either the PPOD- or RPD-level. This means that more granular CM and traffic metrics will be aggregated to provide a big picture view of the customer experience downstream of the PPOD and each RPD. This might involve, for instance, counting the CMs in a particular state downstream of an RPD, or summing all upstream traffic across all RPDs to calculate a PPOD-level traffic measure.

KPI comparison algorithms can be generally categorized as one of the following: *i.* percent recovery, *ii.* percent increase, *iii.* greater-than-zero, or *iv.* custom. Simply put, algorithms that analyze components in “good states”, such as *online*, *connected*, *synced*, etc., tend to use percent recovery calculations, while algorithms that analyze components in “bad states”, like *partial service*, tend to use percent increase calculations. Some KPIs are analyzed simply by checking if the post-deployment reading is greater-than-zero. This is typically done with continuous traffic KPIs since the algorithm is only limited to an instant sample of data post-deployment, eliminating the possibility of any trend analysis. Custom algorithms, as the name suggests, do not fall into these general categories and are evaluated using customized logic and thresholds. Table 1, Table 2, and Table 3 summarize KPI algorithms, thresholds, and alert categories at the PPOD-, RPD-, and CM-levels respectively. These KPIs are observed in accordance with DOCSIS and are detailed in the latest CableLabs DOCSIS Remote PHY specification and database of DOCSIS Management Information Bases (MIBs). It may be helpful to reference the “Abbreviations” section on page 24 to comprehend the KPIs presented.

As seen in these tables, several KPIs have two evaluation renditions. Typically, this is done when two levels of severity—one that results in a warning and one that results in a failure—are analyzed for the same KPI. Also noteworthy are the several algorithms/thresholds which have scenario-specific exceptions, as highlighted in the “Notes” section in each table. An example of this can be seen when CMs-in-partial-service is analyzed and held to a percent increase threshold ($< X\%$ increase), but the pre-snap count of partial CMs is very low (< 10). In a scenario like this, the addition of just a few more CMs in partial service post-deployment can cause an entire health check to fail, halting the deployment pipeline and alerting an operator. This scenario is likely not indicative of service degradation due to a software update given that it is not atypical to see some random partial service fluctuation. In this scenario, an exception would be programmed which would forgive a few additional CMs in partial service despite technically breaching the percent increase threshold.

Table 1: Summary of PPOD-Level KPI Algorithms and Thresholds

Level	KPI	Algorithm	Threshold	Notes	Category
PPOD	RPDs Online	percent recovery	$> X\%$ RPDs	test/pre-production RPDs omitted from calculation	failure
	CMs Online - Overall	percent recovery	$> X\%$ CMs	uses subset of CMs online in pre-snap	failure
	CMs Online - IP v4/v6	percent recovery	$> X\%$ CMs	IP version breakdown, uses subset of CMs online in pre-snap	warning
	CMs Online - BSOD	percent recovery	$> X\%$ CMs	-	failure
	CMs in Partial Service	percent increase	$< X\%$ increase in partial CMs	custom algorithm for low-CM scenarios	failure
	CPE Types Online	custom	N/A	breakdown by CPE type	info
	CPE Types in Partial Service	custom	N/A	breakdown by CPE type	info
	MD DS/US Traffic	greater-than-zero	packet rate > 0	except when packet rate is historically zero	warning
	Partial Service (Statistical Percentages)	custom	not statistically greater than history ($< 3\sigma$)	calculates historical distributions of percent-CMs-in-partial	warning
	RPDs PTP-Synced	percent recovery	$> X\%$ RPDs	-	failure
	RPD Time Offline	custom	N/A	time each RPD is down during SA event	info
	PCMM Connection 1	greater-than-zero	COPS connected/open > 0	-	failure
	PCMM Connection 2	percent recovery	$> X\%$ COPS connected/open	-	warning
	OFDMA Channels	custom	N/A	breakdown of OFDMA channels	info
	Mid-Split Enabled CMs	percent recovery	$> X\%$ mid-split enabled CMs	includes breakdown of mid-split enablement status pre- and post-deployment	warning
	Mid-Split Utilizing CMs	percent recovery	$> X\%$ mid-split utilizing CMs	includes breakdown of mid-split utilization status pre- and post-deployment	warning
	Mid-Split Enabled RPDs	custom	N/A	breakdown of mid-split enabled RPDs	info

Table 2: Summary of RPD-Level KPI Algorithms and Thresholds

Level	KPI	Algorithm	Threshold	Notes	Category
RPD	CMs Online - Overall 1	percent recovery	> X % CMs per RPD	uses subset of CMs online in pre-snap	warning
	CMs Online - Overall 2	greater-than-zero	> 0 CMs per RPD	except when RPD had zero CMs in pre-snap	failure
	CMs Online - IPv4/v6	percent recovery	> X % CMs per RPD	IP version breakdown, uses subset of CMs online in pre-snap	warning
	CMs Online - BSOD	percent recovery	> X % CMs per RPD	uses subset of CMs online in pre-snap	failure
	CMs in Partial Service	percent increase	< X % increase in partial CMs per RPD	custom algorithm for low-CM scenarios	warning
	CPE Types Online	custom	N/A	breakdown by CPE type	info
	CPE Types in Partial Service	custom	N/A	breakdown by CPE type	info
	MD DS/US Traffic	greater-than-zero	packet rate > 0 per RPD	except when packet rate is historically zero	warning
	DSG Traffic 1	greater-than-zero	packet rate > 0 per tunnel, per channel, per RPD	warns if a single tunnel has zero traffic post-deployment for a single RPD	warning
	DSG Traffic 2	greater-than-zero	packet rate > 0 per tunnel, per channel, per RPD	fails if all tunnels have zero traffic post-deployment for a single RPD	failure
	OFDMA Channels	custom	N/A	breakdown of OFDMA channels	info
	Mid-Split Enabled CMs	percent recovery	> X % mid-split enabled CMs per RPD	includes breakdown of mid-split enablement status pre- and post- deployment	warning
	Mid-Split Utilizing CMs	percent recovery	> X % mid-split utilizing CMs per RPD	includes breakdown of mid-split utilization status pre- and post- deployment	warning

Table 3: Summary of CM-Level KPI Algorithms and Thresholds

Level	KPI	Algorithm	Threshold	Notes	Category
CM	BSOD DS/US Traffic	greater-than-zero	packet rate > 0 per CM	except when CM historical traffic is also zero	warning

One noteworthy feature of our approach is the ease with which custom algorithms and more sophisticated models can be incorporated. Data required to implement a custom algorithm or model is simply gathered and processed, often in-query, to produce the desired data format as an abstract metric snapshot unit. This collection step often takes direction from a configuration file, which might specify the timeframe over which to gather data, for example. The same steps which perform comparisons and evaluate simple algorithms against thresholds can be easily abstracted out to incorporate models instead, all while delivering results and summaries in a format consistent with that of other evaluations. An example of a custom algorithm is discussed in the next section.

3.2.3. Custom Algorithm Example – Partial Service

One particular KPI worth discussing in-depth is partial service. Partial service occurs when a CM is online but unable to operate on one or more downstream (DS) or upstream (US) channels, which may or may not hinder the customer experience. A CM can go into partial service for a variety of reasons, including loss of communication on a channel, inability to acquire a channel, and/or configurational incompatibilities. Despite its negative connotation, partial service is actually a beneficial feature in that it can often allow an impaired CM to have a mostly normal transmit/receive experience on the subset of channels it has available (Volpe, 2011). Nevertheless, partial service is an indicator that a customer is either currently experiencing service degradation or at risk for service degradation in the future; therefore, partial service is an important KPI to monitor during a software update.

As indicated in the algorithm exception example discussed in Section 3.2.2, partial service can be a notoriously difficult metric to analyze pre-to-post deployment. The difficulty lies in the fact that: *i.* the analysis is limited to a brief sample of data instantly after the deployment, as long-term trend analysis is not in compliance with the two-minute check execution window *ii.* partial service can naturally fluctuate due to factors unrelated to DAA software updates, and *iii.* the expected effect of a particular software update on partial service is not always known. As demonstrated in Table 1 and Table 2, there are standard percent increase algorithms in place to analyze partial service. These algorithms vary in efficacy depending on the size of the CM population observed and tend to capture blatant partial service issues but might not adequately flag more subtle post-deployment partial service anomalies. It is not obvious how to define a more sensitive threshold using the percent increase strategy without introducing excessive false positives; therefore, a new-and-improved custom algorithm was developed. This novel partial service algorithm attempts to further capture PPOD-specific partial service anomalies using a statistical approach.

The steps to this approach can be summarized as follows: *i.* perform a historical query in the pre-check to get a sample of partial service snapshots, *ii.* form distributions of percent-CMs-in-partial specific to the PPOD over several time periods, *iii.* assume normality and calculate thresholds for each distribution by considering statistical convention “three standard deviations above the mean” (a.k.a. 3σ) to be the cutoff, *iv.* calculate instant post-deployment percent-CMs-in-partial and compare to the thresholds. Essentially, this algorithm answers the question: “Is partial service outside the normal range following a software update?”. Figure 5 aims to visually depict how these steps are used to evaluate the KPI. An instant pre-snap reference point is also included to indicate the difference pre-to-post deployment in relation to the historical distributions.

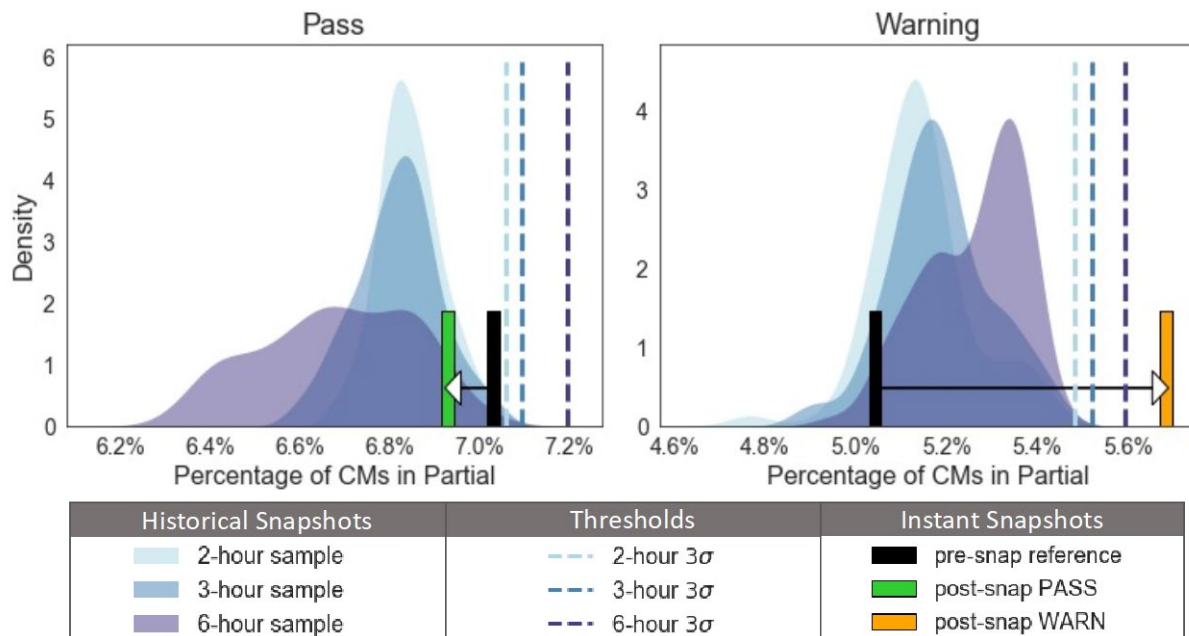


Figure 5: Visuals demonstrating statistical partial service algorithm

3.3. Integration with Software Deployment Automation

As previously mentioned, the network health check integrates with each CI/CD pipeline both prior to and after the software deployment for pre-check and post-check respectively. This setup technically requires four endpoints: *i.* initiation of pre-check, *ii.* polling of pre-check status, *iii.* initiation of post-check, and *iv.* polling of post-check status. The reason for this setup is that each check execution can take up to several minutes, which can exceed the connection time limitations of the cloud resources chosen to implement this application. Check executions are run as asynchronous background processes kicked off by the initiation steps, and the polling steps are intended to deliver a quick indication as to whether the check is complete or still running. Therefore, polling loops are needed within the automation to continue polling a check until the check is complete and delivers results.

When a pre-check is started for a PPOD, a universally unique identifier (UUID) is generated and passed back in the response of the first endpoint call. This UUID is used for reference throughout the remainder of the workflow to ensure each step accesses the correct cached snapshots and process metadata. UUIDs also serve the purpose of representing unique PPOD/software update/timestamp combinations, which is helpful in debugging and analytics.

The application is capable of post-check retries, which take new post-snaps and compare them to the same pre-/historical-snaps when the post-check initiation endpoint is called again. This feature can come in handy particularly during SA updates, when network recovery is gradual and variable from scenario-to-scenario. Typically, a retry is a manual process kicked off after an operator has waited for recovery following an SA update or intervened to correct a detected network issue. Figure 6 demonstrates how the health check features discussed above integrate to form a CI/CD workflow.

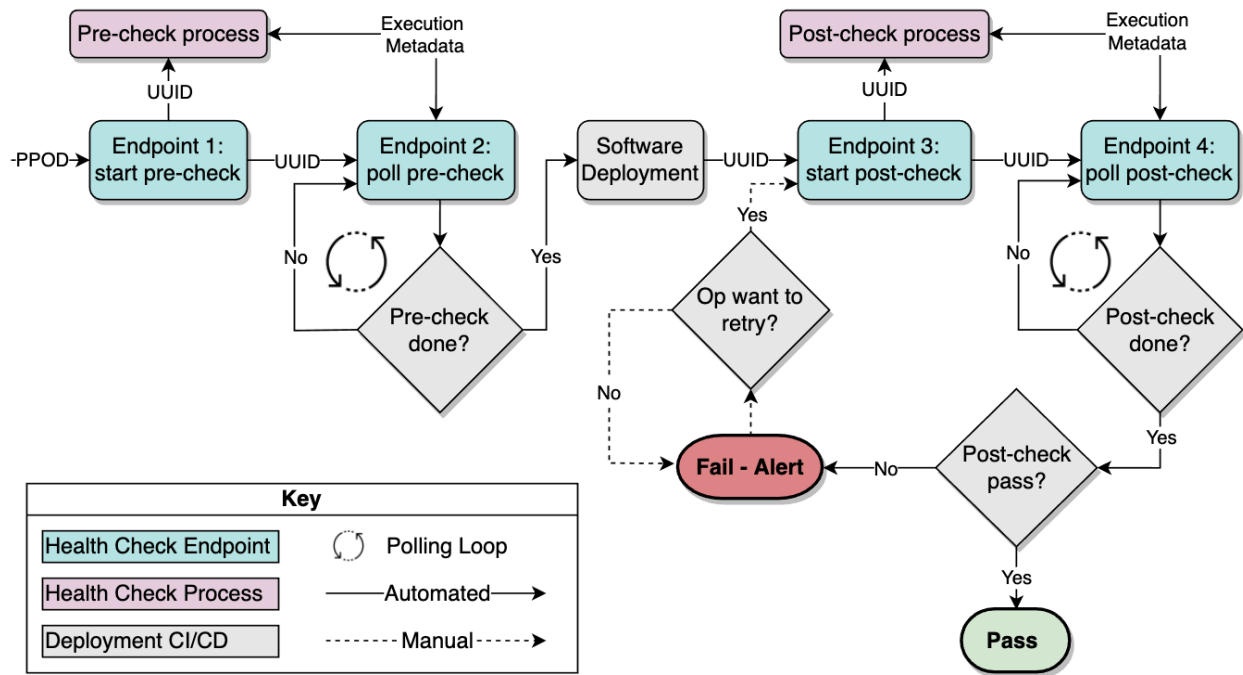


Figure 6: CI/CD integration with health check API — an example workflow

3.4. Health Check Cloud Environment

The health check operates entirely in a cloud environment. The API is hosted using a cloud API gateway service and interacts with a state machine, which orchestrates the health check workflow. A single state machine is used to manage both pre- and post-check, allowing each execution to be defined by the UUID described in the previous section. This is made possible by using callback functionality, which waits for a token to be passed back to the workflow before moving from pre-check to post-check (or from post-check n to post-check $n+1$). Check logic, which includes telemetry queries, comparisons, evaluations, etc., is carried out using serverless cloud workers. Check execution metadata, cached snapshots, and final results are all stored in a cloud database, which can be queried upon calling the API polling endpoints to retrieve the API response. The response should either indicate that a check is still running or deliver the completed check results. Unexpected errors (e.g., errors connecting to or querying from the telemetry database) would also be detailed in the response so that a retry can be performed. Verbose check data intended only for *ad hoc* analysis is also stored in the cloud database. Figure 7 presents a diagram of the health check cloud environment, demonstrating the workflow and interactions among cloud services.

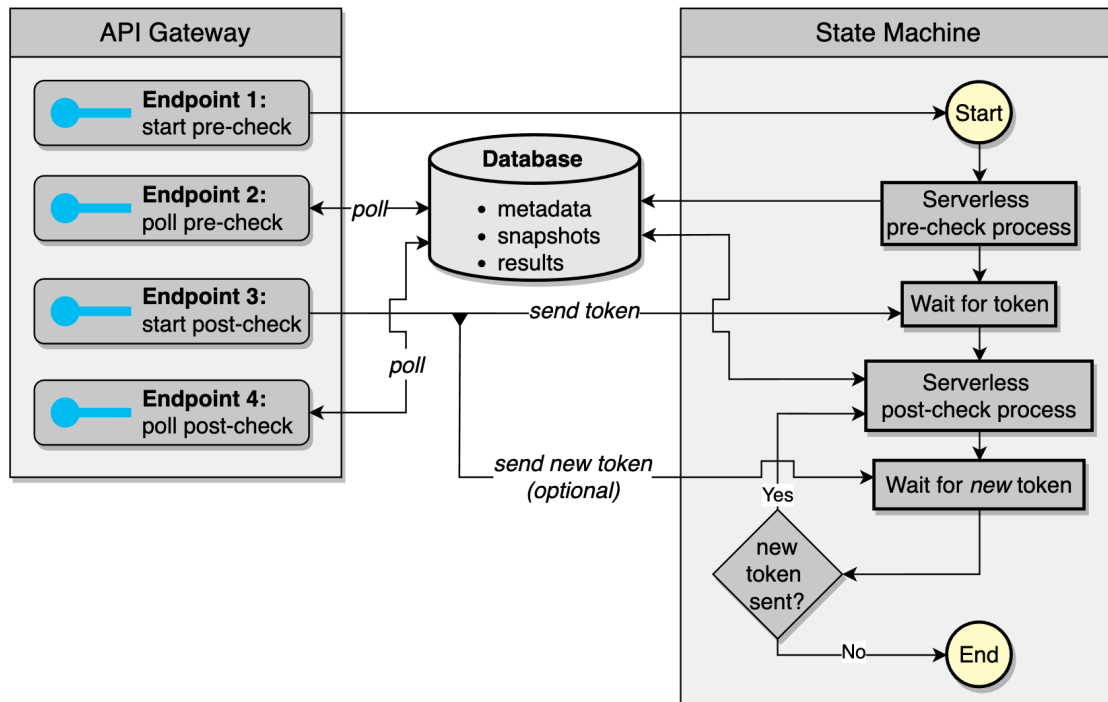


Figure 7: Diagram of cloud environment

4. Discussion

4.1. Usage Analysis

Since the very first production release in 2020, usage of the automated network health check has increased drastically and should continue to increase as Comcast scales DAA deployments. During the initial release stage, API calls were relatively sparse, as the DAA footprint at the time was much smaller, and the automation initiative was only beginning. Nowadays, it is not uncommon to see the API called hundreds, if not thousands, of times per day, depending on the components scheduled for maintenance. The health check is constantly interacting with more and more PPODs, customers, and types of software updates, as DAA continues to scale and engage in automation efforts. As demonstrated in Figure 8, this upward-trending call rate aligns proportionally to the standup of new digital clusters¹, which is a proxy to the growth rate of the DAA footprint.

While this substantial increase in call rate is undoubtedly a testament to the utility of the health check, it is also an indicator that the data sciences team must consistently perform due diligence to ensure that the check infrastructure and selected cloud resources continue to scale to meet the needs of the DAA automation initiative. For instance, cloud environment settings—namely microservice memory, provisioned concurrency settings, and programmed timeouts—are frequently tweaked to maintain reliability and performance of the health check. Additionally, telemetry queries performed in the health check are stress-tested frequently to ensure that the telemetry database can safely handle the large request loads typically seen during a burst of concurrent network health checks.

¹ The health check was released and integrated in 2020; however, health check data was not stored in an optimal format for analytics until later. Therefore, visuals will only show trends for more recent timeframes.




Figure 8: Health check usage

4.2. Analysis of Check Results

To depict failure and warning rates, Figure 9 breaks down the percentages of passes, failures, and warnings for both SA and NSA updates respectively over the past year². As expected, low-risk NSA updates appear to deliver a consistently high pass rate, with a very small percentage of calls resulting in failures. As demonstrated by this data and confirmed by DAA operators and SMEs, the health check has been optimized to a steady state and performs proper due diligence surrounding NSA software updates. The check is capable of alerting the ops team *only* when necessary, but does not over-alert with false indicators of service degradation. Achieving optimal sensitivity is especially important for NSA updates, as they comprise ~91% of all DAA software updates.

As expected, SA updates typically show a higher failure rate, given that customer recovery following an RPD reboot can be a gradual process and the automation can prematurely run the post-check before recovery is complete. Additionally, high-risk SA updates tend to need more network intervention than do low-risk NSA updates, even after a recovery period is observed. While these frequent failures might seem burdensome, they are typically no hindrance to the automation process, given that an operator oversees SA updates as part of maintenance protocol. In these failure scenarios, the main value of the network health check is in the post-check retry capability, not necessarily the alert functionality. The retry capability allows the operator to intervene and repeatedly check all 36 network KPIs with the click of a button, until customer recovery is complete and service is fully restored. As previously mentioned, the operator may also choose to act on warnings during SA updates.

The warning rate for NSA updates consistently hovers around 20%, while SA updates show a more variable warning rate. In both NSA and SA updates, a slight uptick can be seen in recent months, which is attributed to the launch of Comcast's mid-split trials and relevant KPIs recently introduced to the check for testing. These warnings are not necessarily causes for concern, but rather demonstrate the troubleshooting and finetuning process when new KPIs are added. The results of newly added KPIs are often studied on a case-by-case basis, validated against telemetry dashboards, and discussed with DAA

² Prior to August 2021, *warnings* were inconsistently defined and consisted of general comments to the end user in addition to observed KPIs. Therefore, results prior to this date will be omitted from analysis for the sake of consistency.

SMEs in terms of optimal sensitivity. Using this iterative process, we might adjust thresholds, add programmed exceptions, or even break out a single KPI into two KPIs, as seen done in Table 1 and Table 2. In Section 5.1, we will take a deeper dive into the mid-split initiative and discuss next steps for mid-split-related checks.




Figure 9: Health check results

We attempted a correlation analysis to better understand individual KPI result trends, but the findings are somewhat skewed, given that not all KPI results are independent of one another. Additionally, the large variety of component update types further confounds the analysis, but we will continue to seek out ways in which we can glean insight from the expansive data we have collected. We can, however, state that the two most common *independent* modes of check failure are PPOD-level/RPDs-online, which currently requires 100% RPD recovery following a software update, and PPOD-level/CMs-in-partial-service. These two failures have trickle-down effects on other KPIs; for instance, a missing RPD will also manifest as a DSG traffic failure and likely a CM-online failure, but the converses are not necessarily true. As warnings are constantly being tweaked and do not always have optimal thresholds, we will not discuss them in this context.

4.3. Lessons Learned

During the initial development stage of the health check, it was believed that only a few basic metrics—namely variations of RPD recovery, CM recovery, and CMs in partial service—would need to be observed. The initial design, which consisted of a handful of functions, was not well future-proofed, and quickly became unmanageable as DAA progressed and the need to observe new KPIs was realized. The team ended up refactoring the source code using an object-oriented approach to modularize the steps taken—namely metric snapshot querying, comparing snapshots, and evaluating against thresholds—so that new KPIs could easily be added without compromising the existing code. In retrospect, this strategy proved essential in optimizing the check to its current state, which evaluates more-than-triple the number of KPIs at inception. Figure 10 demonstrates the extent to which KPIs have been added over time.



Figure 10: Number of KPIs evaluated in health check over time

Another key lesson learned was “make everything configurable”. Not only are algorithm thresholds configurable, but so are historical-snapshot timeframes and applied percentiles, the ability to assign KPI evaluations to the *warning*, *failure*, or *info* categories, the ability to turn on/turn off programmed algorithm exceptions, etc. This strategy greatly aided in the maintenance of unit testing and allowed finetuning of algorithms with a simple update to a configuration file.

Last but not least, optimizing storage of check data has greatly aided in our analytics, not only for informational purposes (like this paper), but also for debugging specific scenarios and further optimization of the health check based on retrospective analyses. As previously mentioned, we store all API responses, complete with health check results and summaries, in a database, and each is tied to a specific UUID for easy traceability. A verbose version of the check results, which contains more detail than the DevOps team requires but is helpful for our internal analytics and debugging, is also stored per each UUID. Additionally, placing check results in a relational database has vastly improved the efficiency of these analytics and has also enabled network health check results to be featured in other data science applications, such as Stehman et al.’s “Sherlock” analytics tool (2021).

5. Future Work

5.1. Continued Optimization of Network Health Check: Mid-Split Updates

As alluded to previously, the data sciences team is in constant communication with DAA SMEs to discuss the addition of new KPIs and ways in which current KPI algorithms can be improved. This is expected to be an ongoing process as DAA progresses and Comcast launches other initiatives that overlap with DAA, particularly in support of 10G and full duplex (FDX) technology.

On the road to 10G and FDX, Comcast is actively working on deploying mid-splits to enable higher US and DS speeds within its digital footprint. Simply put, this is accomplished by incorporating an orthogonal frequency division multiple access (OFDMA) channel into the broadband spectrum, effectively doubling the available US spectrum (Olfert, n.d.). Enabling mid-split for customers across the DAA footprint requires a handful of configuration changes to each cluster’s custom resource document (CRD). As such, the mid-split enablement process is orchestrated with the standard PPOD software update CI/CD automation, and the network health check is run after each attempted enablement.

To monitor the effectiveness and stability of each attempted mid-split enablement during the trial period, we started by adding several *info* and *warning* KPIs at both the PPOD- and RPD-levels, as it was not immediately clear what pass/fail criteria should be or if they were even needed. These KPIs looked at things like: CMs which successfully became enabled, CMs which became enabled but went into partial

service on the OFDMA channel, CMs which started seeing OFDMA traffic flow, OFDMA channels connected per RPD, etc. As referenced, in Section 4.2, we have seen a slight uptick in overall warnings rates due to this update, as we experiment and continue to tune these KPIs.

Observations like these have prompted further analysis and troubleshooting of the mid-split enablement process and have helped the data sciences team to brainstorm a few new purposeful rule-based algorithms. For instance, one KPI we intend to monitor more strictly in the future is CMs which go into partial service on the OFDMA channel due to mid-split enablement, as this is indicative that field technicians might need to intervene with an RPD's hardware to return network performance to a stable state. As the mid-split initiative expands, we will continually aim to identify meaningful KPIs and finetune algorithms, just as we have been doing throughout the lifetime of the network health check. The hope is that we can add KPIs and algorithms that will help operators to determine intervention strategies when mid-split specific service degradation occurs. Additionally, we will continue to perform analytics on mid-split results—*failure*, *warnings*, and *info*—to look for ways in which network monitoring and diagnostic reports can be improved.

5.2. Continuous Monitoring

The network health check adds tremendous value to the DAA initiative in that it provides thorough network monitoring in the moments *right after* a software update and supports the goal of total automation by eliminating the need for manual observation during maintenance windows. This tool has been optimized to catch many signs of service degradation so that an operator can take quick mitigative action. However, some signs of software-related service degradation do not manifest until several hours, or even days after the deployment. Similarly, some impairments are not easily detected with a quick telemetry sample and require long-term trend analysis to detect. Given that the network health check is only intended to take a quick on-demand snapshot of the network health, it is not a good tool for continuous monitoring. Therefore, there is a need to develop a new tool, which will expand on the network health check approach and add the capability for long-term network health monitoring of each PPOD following a software update. This tool will differ from other production network monitoring tools in that it will try to pinpoint signs of service degradation specifically caused by or correlated with software updates, as opposed to all signs of service degradation in general. This distinction is key in delivering an optimized diagnostic tool designed specifically for DAA operations.

5.2.1. Implementation

With continuous monitoring, post-deployment anomalies will be detected using rule-based and machine learning (ML) algorithms, as well as time series analysis techniques, pushing the network monitoring initiative to the next level with artificial intelligence (AI). The proposed workflow for continuous monitoring is illustrated in Figure 11, which depicts an anomaly detection algorithm consistently observing the network for a set period after deployment. In a similar manner to the live health check, the DAA DevOps team will receive alerts when service degradation is detected or concerning anomalies are noticed. However, the workflow will differ slightly in that the continuous monitoring tool will run as a cron job and periodically push alerts to a publish/subscription (pub/sub) service available to the DAA ops team.

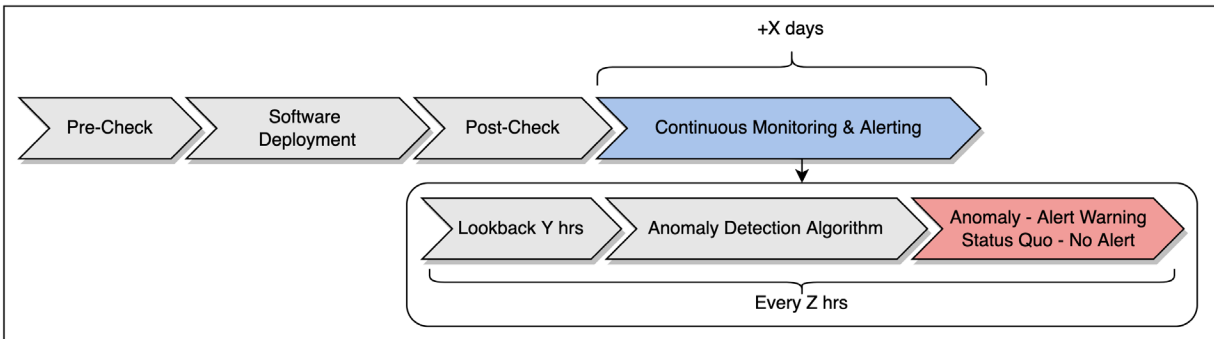


Figure 11: Continuous monitoring workflow

Continuous monitoring is considered an extension of the network health check but will be an independent product due to the magnitude and nature of monitoring metrics, algorithms, and alerts. Metrics in development for continuous monitoring include CM online status, partial service, MD traffic, customer contact metrics, and more. Figure 12 displays a recently noticed anomaly overlayed by a deep learning model in development to detect it. This particular anomaly is characterized by a select subset of CPE device types experiencing random brief service interruptions as a result of an erroneous configuration setting pushed through the deployment automation. This is an example of a scenario that went undetected with the instant telemetry analyses provided by the live network health check. An anomaly like this requires pattern analysis over time, making it an exemplary candidate for continuous monitoring.




Figure 12: Proof-of-concept anomaly and model in development

5.2.2. Expectations

The data sciences team is currently testing the waters with some proof-of-concept models intended to capture anomalies following deployments in the long run. As these models improve and software-related anomalies are better understood, the hope is that continuous monitoring can play an even bigger role in the automation initiative—possibly even rolling back software automatically if it can be determined with near certainty that an erroneous deployment resulted in service degradation. With this initiative, we also hope to perform analytics at an even higher-level aggregation by examining groups of PPODs over the same timeframe following a common software update. Although we had mentioned that this would not be

a good approach for the live network health check, we hope that advanced time series techniques and possibly unsupervised methods, in conjunction with the expansive amount of data available, might be able to help identify problematic PPOD configurations, or even more specifically, erratic configuration/software interactions. In summary, the expectation for continuous monitoring is twofold: *i.* continue to bolster DAA automation by detecting and alerting on known anomalies and *ii.* use advanced analytics to increase our understanding of anomalous network patterns resulting from rapid scaling and constant enhancements occurring as part of the DAA initiative.

6. Conclusion

In response to the DAA engineering team's call to automate, the data sciences team has developed a live network health check meant to replace eyes-on-glass network health monitoring surrounding software updates. This was made possible thanks to the near real-time network telemetry streaming across the DAA footprint, allowing for quick and nimble analysis of network health KPIs. With the guidance of DAA SMEs and consistent feedback from the DevOps team, the data sciences team has been able to finetune and optimize the health check algorithms to achieve the dependable, steady decision engine currently in production today. While there have been a handful of key operational improvements that have supported the expansion of DAA, automated network health monitoring has been particularly impactful given the strict need to preserve the customer experience while performing updates and maintenance. With improvements like this health check, the footprint has been able to expand considerably, as made evident by the substantial increase in vCMTS clusters launched since the health check was first integrated. Despite the rapid growth of the DAA footprint, operational manpower needed to sustain the DAA initiative has mostly remained steady or even reduced in some scenarios, demonstrating the utility of the automated network health check.

While the health check is a valuable tool meant for use during brief maintenance windows, it is limited in that it cannot perform continuous monitoring over a more expansive timeframe following a software update. This is the next key need that the data sciences team will aim to tackle in support of DAA expansion and automation. With the continuous monitoring initiative, we intend to go beyond the scope of quick rule-based evaluations and enter the domain of anomaly detection and ML to deliver an even more thorough, diagnostic view of network health following software updates.

Abbreviations

API	application programming interface
AI	artificial intelligence
BSOD	business services over DOCSIS
CI/CD	continuous integration/continuous deployment
CM	cable modem
COPS	common open policy service
CPE	customer premises equipment
CRD	custom resource document
DAA	Distributed Access Architecture
DOCSIS	Data Over Cable Service Interface Specification
DS	downstream
DSG	DOCSIS set-top gateway
FDX	full duplex
HFC	hybrid fiber/coax
IP	internet protocol
KPI	key performance indicator
MD	MAC domain
MIB	management information base
ML	machine learning
NSA	non-service-affecting
OFDMA	orthogonal frequency division multiple access
PCMM	PacketCable MultiMedia
PPOD	physical point of deployment
PTP	precision time protocol
RPD	remote PHY device
SA	service-affecting
SME	subject matter expert
US	upstream
UUID	universally unique identifier
vCMTS	virtual cable modem termination system

Bibliography & References

Humanoids Optional: Deploying vCMTS at Scale with Automation, Bhanu Krishnamurthy and Gregory Medders, SCTE Expo 2021

Distributed Access Architecture Is Now Widely Distributed - And Delivering On It's Promise, Dr. Robert Howald et al., SCTE Expo 2021

Node Provisioning and Management in DAA, Robert Gaydos et al., SCTE Expo 2018

Solving The Mysteries of the Distributed Access Architecture, Matthew Stehman et al., SCTE Expo 2021

Data-Over-Cable Service Interface Specifications, DCA - MHA v2, Remote PHY Specification. CM-SP-R-PHY-I14-200323

<http://mibs.cablelabs.com/MIBs/DOCSIS/>. Accessed June 2, 2022.

Brady Volpe. "DOCSIS 3.0 Partial Service". The Volpe Firm, December 7th, 2011, <https://volpefirm.com/docsis-3-0-partial-service/>. Accessed June 15, 2022.

Matthew Olfert. "Getting Started with OFDMA". Broadband Library, n.d., <https://broadbandlibrary.com/getting-started-with-ofdma/>. Accessed June 15, 2022.

SCTE Smart Amplifier Project

Extend Proactive Network Maintenance to the Outside Plant

A Technical Paper prepared for SCTE by

Doug Jones
Principal Architect
CableLabs
858 Coal Creek Circle, Louisville, CO. 80027
+1 303.661.9100
d.jones@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Network Management Architecture.....	4
3. SCTE 1.8 GHz Amplifier Standard.....	5
4. Smart Amplifier Model.....	6
4.1. Introduction.....	6
4.1.1. Embedded DOCSIS (eDOCSIS) Model.....	6
4.1.2. Other Transponders.....	7
4.2. Embedded Amplifier (eAMP).....	8
5. Semantics Around Modeling.....	8
5.1. Introduction.....	8
5.2. Information Model.....	8
5.3. Data Model.....	8
5.4. FCAPS Modelling.....	9
5.4.1. Fault Management.....	9
5.4.2. Configuration Management.....	9
5.4.3. Accounting Management.....	9
5.4.4. Performance Management.....	9
5.4.5. Security Management.....	10
5.5. RESTCONF / YANG.....	10
6. The Smart Amplifier Information Model.....	10
6.1. Background.....	10
6.2. Introduction.....	10
6.2.1. Status and Configuration Modules.....	11
6.2.2. Initial Setup.....	11
6.3. System Group Information.....	11
6.3.1. System Status Information.....	11
6.3.2. System Configuration Information.....	12
6.4. RF Group Information.....	13
6.4.1. RF Status Information.....	13
6.4.2. RF Configuration Information.....	14
6.5. Networking Group Information.....	15
6.6. Additional Information Model Areas.....	16
6.6.1. Reset.....	16
6.6.2. Events.....	16
6.6.3. File Management.....	16
7. Conclusion.....	16
Abbreviations.....	17
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 – DOCSIS PNM Architecture.....	4
Figure 2 – Outside Plant PNM Architecture.....	5
Figure 3 – Logical Model for Smart Amplifier.....	6
Figure 4 – eDOCSIS Model for Smart Amplifier.....	7

Figure 5 – Smart Amplifier Top-Level Information Model	11
Figure 6 – System Status UML Diagram	12
Figure 7 – System Configuration UML Diagram	13
Figure 8 – RF Status UML Diagram.....	14
Figure 9 – RF Configuration UML Diagram	15
Figure 10 – Network Status UML Diagram	15

List of Tables

Title	Page Number
Table 1 – SCTE Standards Supporting DOCSIS 4.0 Technology	6
Table 2 – CableLabs Embedded DOCSIS Specifications	7

1. Introduction

A “smart amplifier” defines the information model needed to perform both remote configuration and the gathering of status information from RF amplifiers. At the time of this writing, the working group is actively completing its work based on the SCTE 1.8 GHz amplifier standard, [SCTE 279].

This proposed scope of the project is to define a standard information model using the Universal Modeling Language (UML) and then create the YANG data model for communications with broadband amplifiers used in hybrid fiber-coax (HFC) networks. The modeling can apply to all HFC networks including DOCSIS 4.0 applications and is intended to include all required monitoring and control communications with an amplifier, whether accessed over the HFC control plane or locally via direct wired or wireless connection.

The cable industry benefits from having a standardized information model that is used for amplifier configuration and status. A goal of this project is to create an information model that is applicable to both stand-alone distribution amplifiers and to launch amplifiers inside fiber nodes, leveraging existing CableLabs operations system work from the Distributed Access Architecture (DAA) project.

The smart amplifier features and capabilities can be leveraged to enable measurement and reporting of network conditions such that the system can be made more reliable. 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.

2. Network Management Architecture

A smart amplifier extends the concept of proactive network maintenance (PNM) to the outside plant (OSP) elements including both the fiber node and amplifiers. PNM introduced data gathering and algorithms for DOCSIS components of the network that attached to the coaxial cable, these being the cable modem termination system (CMTS) and cable modem (CM).

Figure 1 show the DOCSIS CMTS and CMs sending telemetry data to a data store where algorithms can be performed on the information.

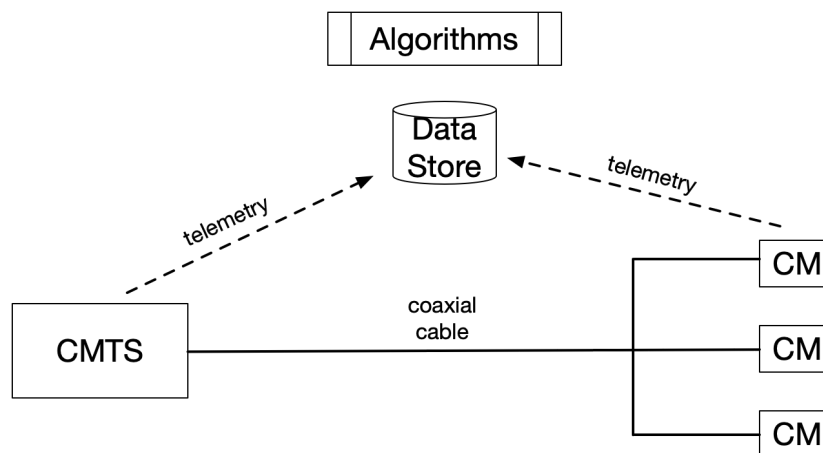


Figure 1 – DOCSIS PNM Architecture

Note that while the DOCSIS equipment directly connects to the coaxial cable network, the traditional elements making up the OSP, including the fiber node and amplifiers, are not sending information to that data store.

Figure 2 shows how the OSP, including both fiber node and amplifiers, can be included in the PNM solution.




Figure 2 – Outside Plant PNM Architecture

The CableLabs DAA project has instrumented the fiber node so that data is available for the part of the CMTS in the node, as well as the RF launch amplifier in the node. The SCTE smart amplifier project is instrumenting the coaxial cable network amplifier, including both the downstream and upstream paths.

With both DAA and smart amplifiers, the active elements of the OSP are instrumented and can provide data to the operator. With these new sources of data, new methods for managing both the OSP and the DOCSIS network can be developed. It is expected that new algorithms will use the combined data from both DOCSIS equipment and OSP equipment for proactively monitoring the cable broadband network.

A goal of the smart amplifier standard is to provide capabilities in the amplifier so monitoring can be performed remotely, for example in a network operations center (NOC), without having to dispatch a technician to the amplifier and open the lid to gather readings. The development of these use cases can be used to make future adjustments to the [SCTE 270] amplifier information model.

3. SCTE 1.8 GHz Amplifier Standard

[SCTE 279] was recently published and defines standard mechanical, environmental, and electrical characteristics for RF amplifiers that support DOCSIS 4.0 frequency division duplex (FDD) capabilities with downstream operation to frequencies up to 1794 MHz and upstream operation to frequencies up to 684 MHz.

The SCTE smart amplifier standard includes a management model for the [SCTE 279] amplifier to allow both remote configuration and status gathering from the amplifier. The same standard tools used for the fiber node information modem in the DAA project are used to develop the amplifier information model and align with the fiber node work.

[SCTE 279] rounds out a group of SCTE standards tied to DOCSIS 4.0 technology that include:

Table 1 – SCTE Standards Supporting DOCSIS 4.0 Technology

Number	Title
[SCTE 264]	Broadband Radio Frequency Hardline Taps for Cable Systems
[SCTE 265]	Broadband Radio Frequency Hardline Passives for Cable Systems
[SCTE 273]	Generic Access Platform (GAP)
[SCTE 279]	1.8 GHz Broadband Radio Frequency Hardline Amplifiers for Cable Systems

In addition to the standards listed, numerous other SCTE standards have been updated for both 1.8 GHz and 3.0 GHz operation as part of the effort for the DOCSIS 4.0 specifications.

4. Smart Amplifier Model

4.1. Introduction

[SCTE 38-10] was the first attempt to create standard instrumentation for an amplifier. The smart amplifier work both builds and extends this early work.

Figure 3 shows the smart amplifier modeled to include both a transponder and an embedded amplifier (eAMP). The eAMP includes amplifier functions defined in [SCTE 279] for both the downstream and upstream signal paths and is the focus of smart amplifier. The transponder is not defined in [SCTE 279] and the choice of transponder is left as a business choice by a cable operator. The transponder could be a DOCSIS modem or some other type of modem.




Figure 3 – Logical Model for Smart Amplifier

The choices of transponder and amplifier capabilities should be made carefully. The transponder and eAMP are separate logical entities and have separate information models. The smart amplifier work focuses on the eAMP logical entity and does not define the transponder.

To support the management model in the smart amplifier standard, it is recommended the transponder implement a high-speed IP connection that supports the types of networking used by the cable operator. Note that the DAA fiber node has direct digital fiber connection and does not use a cable modem as a transponder.

4.1.1. Embedded DOCSIS (eDOCSIS) Model

If the transponder is a DOCSIS modem, then the embedded DOCSIS (eDOCSIS) model is available to operators and suppliers. The eDOCSIS model has been updated in support of the smart amplifier and the relevant CableLabs specifications are listed in Table 2.

Table 2 – CableLabs Embedded DOCSIS Specifications

Designation	Title
CANN	CableLabs Assigned Names and Numbers specification
CANN-DHCP-Reg	CableLabs DHCP Options Registry specification
eDOCSIS	Embedded DOCSIS specification

The eDOCSIS model has successfully been used for embedding modems for digital voice, set-top boxes, and other types of equipment. Cable operators are familiar with this model and have it in use.

Figure 4 shows the transponder being an embedded cable modem which is defined in [eDOCSIS].




Figure 4 – eDOCSIS Model for Smart Amplifier

A short discussion on the DOCSIS version of the eCM transponder. As noted previously, the eCM and eAMP could have different capabilities. The amplifier defined in [SCTE 279] operates to 1794 MHz. A DOCSIS 3.0 eCM is required to operate to only 870 MHz (though many DOCSIS 3.0 modems operate to 1002 MHz) and is not required to have the PNM tools to report on OFDM and OFDMA carriers, which were introduced in the DOCSIS 3.1 specification. Even if the eCM is a DOCSIS 3.1 cable modem, it would be specified to operate to 1218 MHz and not to 1794 MHz.

[SCTE 279] discusses amplifier properties intended for a DOCSIS 4.0 network that include operation to 1794 MHz. [SCTE 279] does not discuss vector signal analyzer (VSA) functions such as OFDM symbol decoding or reporting RxMER on OFDM subcarriers, or other PNM functions that are associated with DOCSIS modems.

Hence the choice of DOCSIS transponder includes discussion on both the frequency range of that transponder as compared to the amplifier, and what if any DOCSIS PNM capabilities might be desirable in that transponder.

Additionally, while a DOCSIS CM supports various PNM capabilities, for these to come into play on each physical port of the amplifier, that CM would need to attach to those physical ports. This could entail a switch that could allow the CM to sample the RF at each port. Figure 3 does not identify a particular port the transponder is attached to and how the transponder may report is not specifically identified.

4.1.2. Other Transponders

The transponder does not have to be a DOCSIS CM; other transponders are possible. Note that the transponder should support a high-speed Internet Protocol (IP) connection.

One example is the SCTE Hybrid Management Subcommittee (HMS) working group specification of a transponder in [SCTE 25-1] and [SCTE 25-2]. The HMS transponder is in use for applications such as OSP power supplies and amplifiers.

Other types of transponders can be provided by suppliers; there are several deployed.

4.2. Embedded Amplifier (eAMP)

Looking again at Figure 3, the smart amplifier work defines an information model for the eAMP logical entity; smart amplifier does not define the transponder. There are no DOCSIS functions in the eAMP. The eAMP represents the [SCTE 279] amplifier functions including parameters such as RF test points, attenuators, and equalizers. From the standpoint of the eAMP, the transponder is simply a method to provide communications to the eAMP.

5. Semantics Around Modeling

5.1. Introduction

An information model is a way of representing and structuring information available from the device being modeled. An information model (as compared to a data model) is an abstraction and only provides a high-level view of relationships and things of interest (i.e., information). It aids in understanding the scope of functions of the device being modeled.

Data models usually specify items in more detail and include protocol specific constructs. The level of abstraction does not depend on the language being used (e.g., XML, IDL, YANG). Such languages allow modeling both at high and low (i.e., detailed) levels. The smart amplifier data model, using YANG, is planned to be available later this calendar year.

5.2. Information Model

The information model is designed using the unified modeling language (UML). A description of how to use UML can be found in [UML Guidelines].

Many network elements have an information model including CMTS, CM, Wi-Fi access points, servers, routers, and switches. This is how the network is managed, through interacting with the information models of these network elements. The DOCSIS CMTS and CM devices have had information models since the DOCSIS 1.0 specifications which have evolved as the DOCSIS specifications have evolved.

5.3. Data Model

A data model is a specific implementation of an information model using a specific set of protocols. For example, the early DOCSIS specifications had data models for the CMTS and CM devices using structure of managed information (SMI) and simple network management protocol (SNMP) which were early network management protocols, designed in the 1990s and were included in DOCSIS 1.0 technology and are still available with the DOCSIS 4.0 specifications.

SMI provides the rules for structuring the information, these structures are commonly referred to as a management information base (MIB), and these MIBs were managed using SNMP. While MIBs and SNMP were used in the past, the migration is toward YANG models and RESTCONF for the future.

YANG and RESTCONF are modern, web scale protocols which were included with DOCSIS 3.0 technology for more saleable network management. YANG is a method of representing the data model and RESTCONF is a protocol based on HTTP (hyper-text transfer protocol) that has proven to be very scalable for use in large networks and includes its own security model.

5.4. FCAPS Modelling

The FCAPS model is a widely used framework that organizes management functions into the categories of fault, configuration, accounting, performance, and security management. Telecommunications operators, including cable operators, commonly use the FCAPS model to manage large networks of devices. This specification uses these management categories to organize the requirements for the configuration and management of the smart amplifier.

- Fault management seeks to identify, isolate, correct, and record system faults.
- Configuration management modifies system configuration variables and collects configuration information.
- Accounting management collects usage statistics for subscribers, sets usage quotas, and bills users according to their use of the system.
- Performance management focuses on the collection of performance metrics, analysis of these metrics and the setting of thresholds and rate limits.
- Security management encompasses identification and authorization of users and equipment, provides audit logs and alerting functions, as well as providing vulnerability assessment.

Each of these management categories is discussed in more detail in the following sections.

5.4.1. *Fault Management*

Fault management is a proactive and on-demand network management function that allows non-standard/abnormal operation on the network to be detected, diagnosed, and corrected. A typical use case involves network elements detecting service-impacting abnormalities; when detected, an autonomous event (often referred to as an alarm notification) is sent to the NOC to alert the operator of a possible fault condition in the network affecting a customer's service. Once the operator receives the event notification, further troubleshooting and diagnostics can be performed by the operator to correct the fault condition and restore the service to proper operation.

5.4.2. *Configuration Management*

Configuration management enables system configuration building and instantiating, installation and system turn up, network and device provisioning, auto-discovery, backup and restore, software download, status, and control (e.g., checking or changing the service state of an interface).

Configuration management is primarily concerned with network control via modifying operating parameters on network elements such as the smart amplifier. Configuration parameters could include both physical resources (for example, a gain stage) and logical objects (for example, a description of the amplifier).

5.4.3. *Accounting Management*

Accounting management allows operators to measure the use of network services by subscribers for the purposes of cost estimation and subscriber billing. No Accounting management has been defined for the smart amplifier at this time.

5.4.4. *Performance Management*

Performance management is a proactive function to gather and analyze data for the purpose of monitoring and correcting the behavior and effectiveness of the network, network equipment, or other equipment and

to aid in planning, provisioning, maintenance, and the measurement of quality. A Performance management use case might include the NOC performing periodic (15 min, for example) collections of output level measurements from each port on the amplifier to perform monitoring and identification of any potential performance issues. With the historical data that has been collected, trending analysis can be performed to identify issues that may be related to certain times of day or other corollary events.

Additional performance management functions include monitoring the amplifier power supply or managing an ingress detection mechanism enabled by the amplifier.

5.4.5. Security Management

Security management provides for network and operator security, as well as providing an umbrella of security for the telecommunications management network functions. Security management functions include authentication, access control, data confidentiality, data integrity, event detection, and reporting. For example, [SCTE 279] requires the amplifier to have a USB port for configuration, and while outside the scope of this work, that USB port needs to be secured from tampering. A sensor is required to provide notification when the amplifier enclosure is opened and closed.

5.5. RESTCONF / YANG

Interfaces to the smart amplifier are defined using YANG-based data models. To manage that data model, RESTCONF is expected to be chosen as the protocol for creating, reading, updating, and deleting instances of YANG objects in the smart amplifier. The YANG models are not yet defined and will be available in the smart amplifier standard.

6. The Smart Amplifier Information Model

6.1. Background

The information model diagrams shown were still under development when this paper was written. While they are substantially complete, changes may occur as the SCTE smart amplifier standard is completed.

For the sake of brevity, not all the UML diagrams will be presented. The intent is to introduce the concepts of an information model and show the applicability to an RF amplifier. The complete UML with text descriptions will be available in the SCTE smart amplifier standard when that becomes available.

Aspects of the information model have been borrowed from the CableLabs Distributed Access Architecture (DAA) project which is defining “Smart Fiber Nodes” for the next generation HFC network. The intent is to have active elements of the OSP, including both fiber nodes and amplifiers, be managed using common systems.

6.2. Introduction

As shown in Figure 5, the smart amplifier information model is organized around three areas:

- SystemGrp, which is system group information including model number and software version
- RfGRP, which is RF Group information and settings for the RF functions on the eAMP
- NetworkingGrp, which provides networking information known to the eAMP




Figure 5 – Smart Amplifier Top-Level Information Model

Each of the three areas will be described in more detail later in this document.

6.2.1. Status and Configuration Modules

The UML diagrams are separated into status models, which provide read-only information, and configuration models which provide read-write information and support configuration of the eAMP.

Status is about reading data information from the eAMP. Status diagrams can show both capabilities of the amplifier and the actual status of components of the eAMP, for example, RF test points on a physical port, the diplex filter in use, as well as any other information described in the UML diagrams and [SCTE 279]. Another example is the amplifier power supply provides status information, to check the input voltage to the power supply, as well as get information on the power rails used inside the amplifier.

Capabilities depend on the specific amplifier and will vary from supplier to supplier. These are capabilities the amplifier can provide.

Configuration allows the amplifier to be set-up or simply adjusted from a remote operations system. This includes changing the output settings of the amplifier or configuring an upstream ingress attenuator, and any other settings specified in [SCTE 279].

Looking back at Figure 5, both the system group and RF group support both status and configuration information. The network group supports just status information as the information is configured during the eAMP DHCP process and is not directly configurable.

6.2.2. Initial Setup

[SCTE 279] requires the amplifier to use a USB port for configuration via an application provided by the supplier. Initial setup is expected to be performed by a technician at the site where the amplifier is installed. Parameters set during this initial configuration will be reflected in the data model of the eAMP that then can be remotely accessed, either to verify the configuration by reading status objects or by adjusting the configuration of the eAMP by setting configuration objects.

6.3. System Group Information

6.3.1. System Status Information

System status information is read-only, and these items are not configurable by the operator.

Figure 6 shows the system group UML diagram and under SystemGrp there are two branches. To the left are system capabilities and beneath are system status information.




Figure 6 – System Status UML Diagram

System capabilities are what is designed into the amplifier and are not changeable by the operator. This information includes the type of amplifier, which is set at manufacture and cannot be changed. [SCTE 279] lists four types of amplifiers including:

- Multiport amplifier which could be a trunk or bridger amplifier.
- Line extender amplifier which typically has one or two outputs.
- Booster amplifier which typically has one output with lower gain and power consumption.
- Compact amplifier can be mounted in a cabinet and can have one or multiple outputs.

System status information includes supplier-provided information about the amplifier including model number and serial number, as well as software and hardware versions running in that amplifier. Information on the power supply is also available, including monitoring the input voltage as well as the output voltages of the various power rails available inside the amplifier.

Unique to smart amplifier is the capability of having sensors in the amplifier that the operator can get status data from. There is a defined lid sensor, to indicate if the housing is open or closed. And there can be additional sensors that could include a temperature or moisture sensor.

6.3.2. System Configuration Information

System configuration information is read-write and can be configured either at the amplifier location or remotely, for example, from a NOC. Figure 7 shows the system configuration UML diagram.

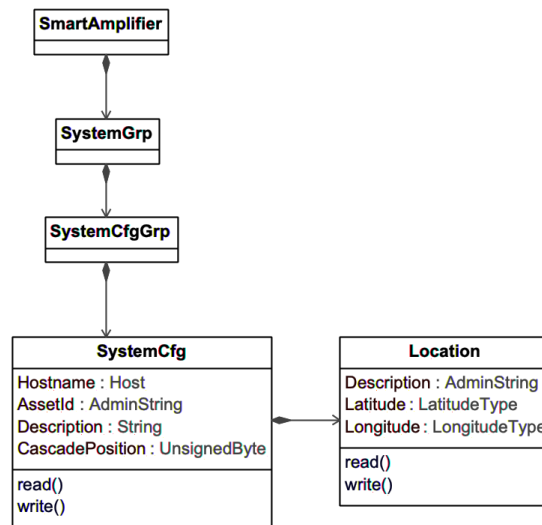


Figure 7 – System Configuration UML Diagram

Items that are writeable include identifying information that can be unique to the operator naming conventions such as hostname, asset identifier and a description of the amplifier. Additionally, the location of the amplifier and its position in the cascade can be written to that amplifier.

6.4. RF Group Information

The RF group contains the details for managing the RF capabilities of the eAMP including RF signal processing functions such as attenuators, equalizers, tilt control, and duplex filters, as well as diagnostic circuitry.

Within the RF UML diagrams, each physical port on the amplifier has logical ports including:

- Downstream logical port
- Upstream logical port
- Bi-directional logical port

This classification exists in both the RF status and configuration UML diagrams.

6.4.1. RF Status Information

Figure 8 shows the RF status UML diagram and under **RfGrp** there are two branches. The RF status information provides a view of how the amplifier is operating. To the right side are RF capabilities and beneath is RF status information.

The RF capabilities include what the amplifier is capable of in terms of downstream, upstream, and duplex filters. These are physical capabilities of the amplifier, including

- the operating frequency ranges
- available duplex filters
- upstream and downstream attenuation and equalization capabilities. Since the [SCTE 279] amplifier is electronically controlled, there are no plug-ins.
- downstream automatic gain control (AGC) capabilities
- upstream ingress attenuator capabilities, i.e., a “wink” switch

- downstream AGC settings
- upstream level control settings

The RF capabilities information is read-only because it is built into the amplifier by the supplier.

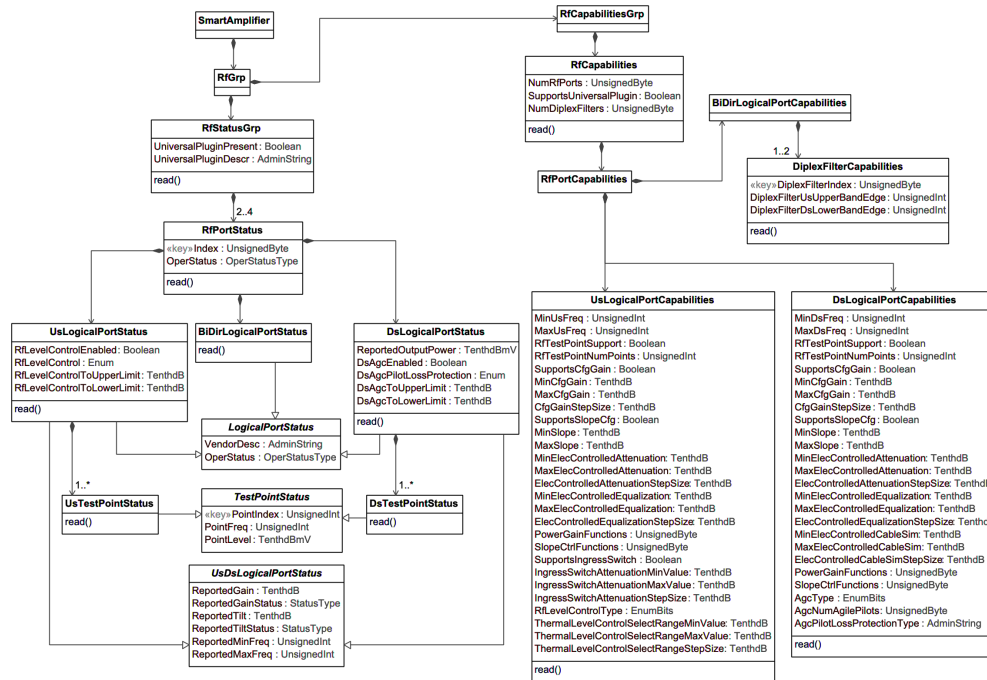


Figure 8 – RF Status UML Diagram

Under the RF status group, available information includes the amount of headroom left in the downstream automatic gain control (AGC) and upstream level control, to allow the operator to learn how the station is performing. Additionally, the upstream and downstream RF test points can be remotely read to give the operator a picture of the output on each port without having to send a tech to that location.

6.4.2. RF Configuration Information

Figure 9 shows the RF configuration information model. These items are described in [SCTE 279] and are read-write to allow remote configuration of the amplifier. For each physical port, common downstream and upstream parameters include adjusting attenuation, and equalization.

Additionally level control is available including downstream AGC including the location of the pilots, and upstream level control. There is also a setting to mute either (or both) the downstream or upstream signals associated with a physical port.

If there is a switchable duplex filter, this is where that choice of duplex filter would be made.

And finally, on the upstream an attenuator can be switched in or out of any upstream path to support locating ingress.

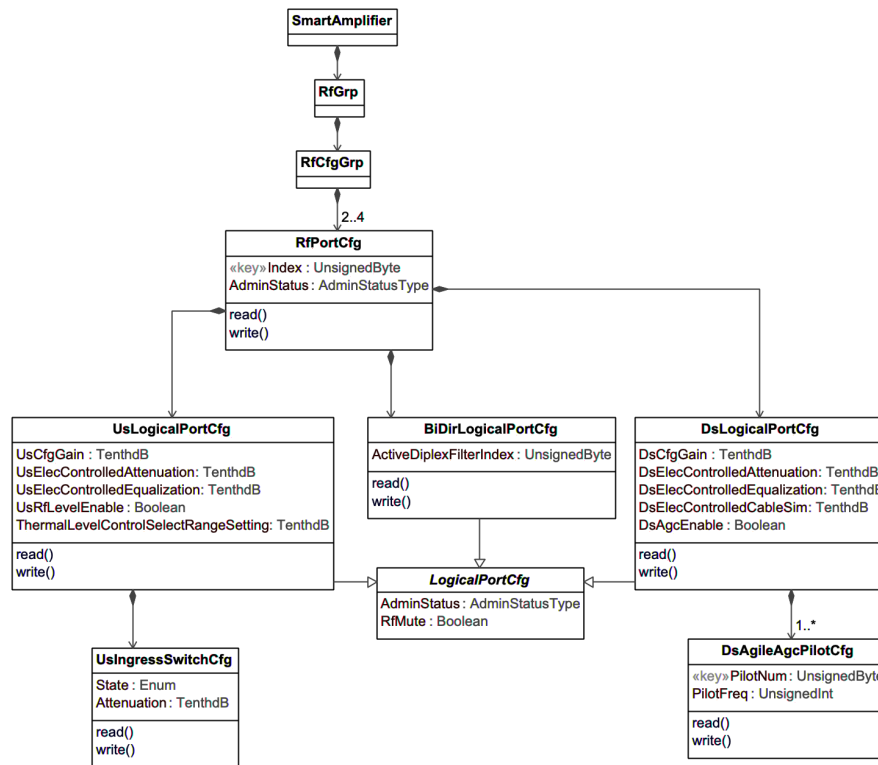


Figure 9 – RF Configuration UML Diagram

These configuration objects allow the eAMP to be adjusted to perform its functions.

6.5. Networking Group Information

Figure 10 shows the networking information model for the eAMP. These items are read-only and are intended to be set during the DHCP exchanges the eAMP has with operator-managed configuration servers.

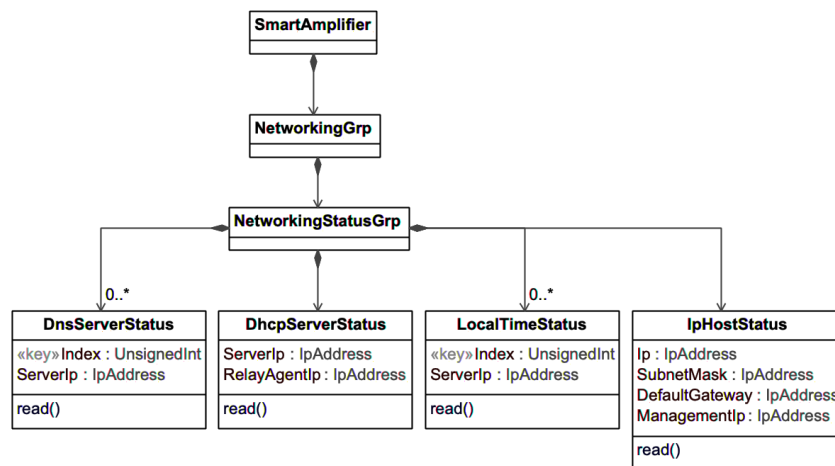


Figure 10 – Network Status UML Diagram

The network status UML diagram shows both internet information for the eAMP and servers the eAMP interacts with, including:

- the IP address of the eAMP
- the DHCP server that provided the IP address to the eAMP
- the IP address of an amplifier management server
- a DNS server that the eAMP uses to resolve IP addresses.
- a server where time is distributed

This information can be useful in debugging network connectivity issues and verifying the eAMP is using the correct servers to send and receive information.

6.6. Additional Information Model Areas

6.6.1. *Reset*

The information model includes capabilities to remotely reboot the amplifier.

6.6.2. *Events*

The information model includes capabilities for a local event log, and for the amplifier to send events to the cloud including a syslog server or using other protocols. Additionally, standard events have been defined for the amplifier.

6.6.3. *File Management*

The amplifier has the capabilities to create two types of files that could be remotely accessed by the operator.

1. A debug file in case the amplifier unexpectedly reboots or has a fault. This file would contain information to allow debugging what caused the issue.
2. A configuration file that provides supplier-specific details on the actual internals of the amplifier and how they are configured.

The information model provides a means to learn which files are on the amplifier, and fetch, rename or delete them.

7. Conclusion

The smart amplifier project provides a method to remotely gather status and configure the amplifier defined in [SCTE 279]. The model is designed to both provide standard management functions and provide room for suppliers to innovate on new products and services.

The smart amplifier work aligns with the CableLabs DAA project so that both fiber nodes and amplifiers have a common management framework. Combined with DOCSIS PNM capabilities, new methods can be developed to increase the reliability of the HFC network.

Abbreviations

AGC	automatic gain control
CM	cable modem
CMTS	cable modem termination system
DAA	distributed access architecture
DHCP	dynamic host configuration protocol
DNS	domain name service
DOCSIS	data over cable service interface specifications
eAMP	embedded amplifier
eCM	embedded cable modem
eDOCSIS	embedded DOCSIS
FCAPS	fault, configuration, accounting, performance, and security
FDD	frequency division duplex
GAP	generic access platform
HFC	hybrid fiber-coax
HMS	hybrid management subcommittee
HTTP	hyper-text transfer protocol
IP	Internet protocol
MHz	megahertz
MIB	management information base
NOC	network operations center
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OSP	outside plant
PNM	proactive network maintenance
RESTCONF	representational state transfer configuration protocol
RF	radio frequency
RxMER	received modulation error ratio
SCTE	Society of Cable Telecommunications Engineers
SMI	structure of managed information
SNMP	simple network management protocol
UML	unified modeling language
VSA	vector signal analyzer
YANG	yet another next generation

Bibliography & References

[CANN]	CableLabs' Assigned Names and Numbers, CL-SP-CANN-I20-200715, July 15, 2020, Cable Television Laboratories, Inc.
[CANN-DHCP]	CableLabs DHCP Options Registry, CL-SP-CANN-DHCP-Reg-I15-180509, May 9, 2018, Cable Television Laboratories, Inc.
[eDOCSIS]	eDOCSIS Specification, CM-SP-eDOCSIS-I30-190213, February 13, 2019, Cable Television Laboratories, Inc.
[RESTCONF]	IETF RFC 8040, RESTCONF Protocol, A. Bierman, et. al., 2017
[SCTE 25-1]	ANSI/SCTE 25-1 2017, Hybrid Fiber Coax Outside Plant Status Monitoring Physical (PHY) Layer Specification v1.0
[SCTE 25-2]	ANSI/SCTE 25-1 2017, Hybrid Fiber Coax Outside Plant Status Monitoring Media Access Control (MAC) Layer Specification v1.0
[SCTE 38-10]	ANSI/SCTE 38-10 2017, Outside Plant Status Monitoring SCTE-HMS-RF-AMPLIFIER-MIB Management Information Base (MIB) Definition
[SCTE 264]	ANSI/SCTE 264 2020, Broadband Radio Frequency Hardline Taps for Cable Systems
[SCTE 265]	ANSI/SCTE 264 2020, Broadband Radio Frequency Hardline Passives for Cable Systems
[SCTE 273-1]	SCTE 273-1 2021, Generic Access Platform Enclosure Specification
[SCTE 273-2]	SCTE 273-1 2021, Generic Access Platform Modules Specification
[SCTE 273-3]	SCTE 273-1 2021, Generic Access Platform Systems Integrator Best Practices
[SCTE 279]	ANSI/SCTE 279 2022, 1.8 GHz Broadband Radio Frequency Hardline Amplifiers for Cable Systems
[UML Guidelines]	UML Modeling Guidelines, CM-GL-OSS-UML-V01-180627, June 28, 2018, Cable Television Laboratories, Inc.
[YANG]	IETF RFC 7950, The YANG 1.1 Data Modeling Language, M. Bjorklund, 2016

Security and Privacy IoT Vulnerabilities:

The Danger of Too Many Entry Points

A Technical Paper prepared for SCTE by

Mangesh Bhamre, Senior Manager of Product (Cybersecurity), Plume Design, Inc.,
325 Lytton Ave, Palo Alto,
CA 94301
mangesh@plume.com
1-250-863-0297

Table of Contents

Title	Page Number
1. Introduction	4
2. Consumers' Security Concerns	4
3. Understanding IoT Security Challenges	5
3.1. IoT Ecosystem Challenges	6
3.1.1. Identity and Authentication	6
3.1.2. Compute Power	6
3.1.3. IoT Device Heterogeneity	6
3.1.4. User Awareness	6
3.2. Work-from-home Trends	7
3.3. Attacker Motivations and Challenges	7
4. The Evolving Landscape	8
4.1. Threat Landscape	8
4.2. CSP Security Services	11
5. Vulnerabilities Life Cycle and Risks	11
5.1. Risk exposure phases	11
5.1.1. High Risk	11
5.1.2. Elevated Risk	12
5.1.3. Medium Risk	12
5.2. Vulnerabilities and Open doors	12
5.2.1. Open ports	12
5.2.2. Vulnerability Numbers on the Rise	14
5.3. Vulnerability Prioritization	17
6. Vulnerability And Attack Taxonomy	18
6.1. IoT architecture	18
6.1.1. Application	18
6.1.2. Middle layer	18
6.1.3. Network	18
6.1.4. Device	18
6.2. The Taxonomy of Vulnerabilities in IoT	19
6.2.1. Weak authentication mechanisms	19
6.2.2. Insecure network services	19
6.2.3. Lacking privacy and data protection	20
6.2.4. Weak device hardening	20
6.3. IoT Attack Flow	21
6.3.1. Use-case of Ransomware	21
6.3.2. Taxonomy of Ransomware	23
7. Solution	24
7.1. Vulnerability Detection and Protection	24
7.1.1. Detection	25
7.1.2. Prevent and Protection	25
7.1.3. Reporting	27
8. Conclusion	28
9. Abbreviations	28
10. References	29

List of Figures

Title	Page Number
Figure 1 – Consumer’s Security Concerns	5
Figure 2 – Limited Consumer Security Awareness	6
Figure 3 – Work-From-Home Security Concerns	7
Figure 4 – Attack Motivations	8
Figure 5 – Growth in IoT Devices at Home (Source: Plume)	9
Figure 6 - Cyber-attack Trends	10
Figure 7 – Cyber-attack Types	10
Figure 8 – Attack Lifecycle and Exposure Zones	12
Figure 9 - Open UPnP Ports Globally (Shodan)	14
Figure 10 – Vulnerabilities Reported Each Year	15
Figure 11 – Common Attack Vectors	16
Figure 12 – Most Exploited Vulnerabilities	16
Figure 13 – IoT Reference Model	19
Figure 14 – IoT Reference Model Taxonomy	21
Figure 15 – Vulnerabilities exploited in Log4J	22
Figure 16 – Log4J Attack Timeline	23
Figure 17 – Ransomware Attack Path	23
Figure 18 – Vulnerability Detection and Protection Solution	25

List of Tables

Title	Page Number
Table 1 - Top 10 Open IoT Ports	13
Table 2 – Distribution of All Vulnerabilities by CVSS Scores	17
Table 3 – Vulnerabilities Exploited in Ransomware Attack Path	22

1. Introduction

While the Internet of Things (IoT) solves some important business concerns for consumers, it also poses significant risks because IoT devices are attractive targets for attack. IoT devices have a history of being vulnerable, they can't be intrinsically protected like less constrained devices, and because they are configured by non-professional/layman users they are ripe for exploitation. Many IoT devices, including everyday objects like kitchen appliances, thermostats, baby monitors, and light control systems, have minimal security built in as compared to full-featured smart devices and are mostly unprotected.

Because they are inexpensive and of limited purpose, IoT devices may have unpatched software flaws. They often have resource-constrained environments with limited processing, memory, and power that make them challenging to secure. Users are mostly non-technical and often lack the knowledge it takes to manage the IoT devices on their networks.

The decline in the overall security profile of homes and offices makes IoT devices a low-hanging fruit for cyberattacks. Attackers can easily get a foothold on the device, exploiting a vulnerability like a weak password or other software flaws. Once a cybercriminal gets access to one device, they can use lateral movement techniques to find other vulnerable devices in the home and conduct severe attacks like ransomware, crypto-mining, password-stuffing, and remote code execution.

There is a critical need for an effective solution that can address the consumer's security concerns and provide state-of-art, enterprise-grade security to homes and business owners. The solution should be able to proactively detect and protect against the security vulnerability which is the primary attack vector in IoT. Communications Service Providers (CSPs) are ideally positioned to play a critical role in mitigating cyberattacks on IoT devices by providing an end-to-end, integrated solution encompassing discovery, detection, monitoring, and resolution.

2. Consumers' Security Concerns

These attack trends have made users wary of the consequences of security breaches, reducing the adoption of IoT. For any transformation, it is essential to build consumer's trust and ensure security is in-built from the design stage. According to a recent survey, consumers have raised concerns (Figure 1) and have called out cyber protection as a requirement.



Figure 1 – Consumer's Security Concerns

IoT clearly presents a security challenge for both homes and small businesses. Based on data gathered from the over 41 million homes and small businesses currently connected to Plume's network, around 67% of homes have high/critical vulnerabilities to attacks on IoT devices. According to analysts, small businesses (SMBs) are the primary targets of cyberattacks and face the most damage in terms of money and customer trust. Around 66% of SMBs in US have been impacted by at least one incident between 2018-2020.

Software vulnerabilities are the primary attack vector for IoT devices and provide an easy foothold for cybercriminals. A recent BotenaGo IoT attack [21] was identified by AT&T Alien Labs in November 2021 as a new malware that exposed millions of IoT devices. The BotenaGo backdoor vulnerability exploits IoT through the open networking port or related modules. There have been countless such incidents in the past and the trend suggests an ongoing increase in these attacks.

3. Understanding IoT Security Challenges

The majority of IoT devices are not built with security-first design principles [23]. As a result, these devices have inherent software vulnerabilities. Many IoT devices cannot be patched with security fixes and, as a result, almost all devices will be at risk. Hackers are now actively targeting IoT devices such as routers and webcams because their inherent lack of security makes them vulnerable and easy to compromise.

These IoT security challenges are partly due to the technical nature of the IoT ecosystem as well as unique security requirements. The technical ecosystem must deal with scalability, distribution, heterogeneity, low energy, and the omnipresent nature of IoT devices. Authentication, confidentiality, integrity, and end-to-end security, on the other hand, are inherent security requirements. Fulfilling all security requirements is difficult given the constraints and limitations in computational and power resources within the devices.

3.1. IoT Ecosystem Challenges

3.1.1. Identity and Authentication

IoT devices need a unique identity on the network to provide mutual authentication, however, there is no consistent mechanism for this. An academic survey found that there are more than 80 different authentication mechanisms proposed or implemented [24]. There is no authentication standard at this point and when many entities (i.e., devices, humans, software, etc.) are involved, authentication becomes difficult. Authentication can also become more complex due to the scale and size of the IoT fabric.

3.1.2. Compute Power

Because IoT devices have limited computing and power capabilities, designing, and implementing encryption or authentication methods is difficult. For maximum IoT security, these cryptographic algorithms must be able to work on small devices and be compatible with the device's compute capabilities. Lightweight and pluggable solutions should be created and deployed to match the limited compute power of IoT devices.

3.1.3. IoT Device Heterogeneity

IoT devices are heterogeneous in their capabilities, communication protocols, technical interfaces, etc. This poses serious challenges when trying to provide an end-to-end security solution that requires devices to share information and collaborate.

3.1.4. User Awareness

Consumers lack awareness about the connected devices installed in their homes and businesses. Often, the devices are unpatched, have weak or default credentials, have vulnerable open ports and services, or are exposed to the internet. Most users are non-technical and lack the expertise to understand the security implication, patch vulnerabilities, or fix the security issues (Figure 2).




Figure 2 – Limited Consumer Security Awareness

3.2. Work-from-home Trends

The pandemic has transformed business functions forever and working from home (WFH) has been growing in popularity. However, WFH comes with a slew of security risks, according to the CISO magazine (Figure 3): The rise of the remote workforce has multiplied the attack surface by adding more endpoints that can be vulnerable to security breaches and device access over insecure network connections. This makes WFH users vulnerable and an easy target for attackers, making the task of security providers even more challenging.

However, rising to this challenge is critical as many users are now doing business-critical work—using sensitive customer data—on machines connected to networks with unknown security weaknesses and populated by many unvetted devices. A joint advisory issued by the US Cybersecurity and Infrastructure Security Agency and the UK's National Cyber Security Centre says the rise of WFH during the COVID-19 pandemic has seen an increase in bad actors targeting individuals and organizations [11].

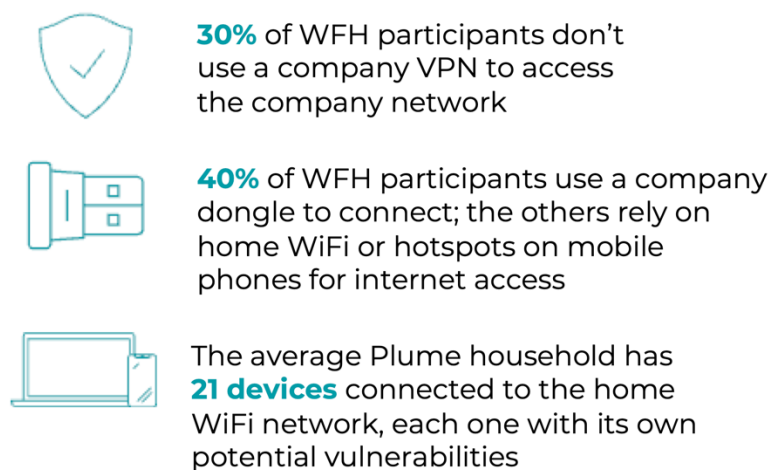


Figure 3 – Work-From-Home Security Concerns

3.3. Attacker Motivations and Challenges

Organizations fail to give enough importance to IoT security as attacker motivations and security risks are not well understood. Security is often completely ignored when, instead, efforts should be focused on critical areas. It is important to understand that the motivation for attacks depends on the attacker and the domain. Attacker types can range from novice, petty thieves, and hactivities to professional crime facilitators. A new class of malicious actors that are knowledgeable and well-funded is emerging.

The primary motives for attackers in the context of homes and small businesses include:

- Financial gain – Attacks geared towards stealing personal and financial information, followed by monetization.
- State-sponsored hacking -
- Recognition and popularity
- Revenge

The challenge lies in intelligently narrowing down the attacker type and motivation for any given home or business. A careful audit and profiling of the assets is essential to identify what would be attractive to attackers. This profiling needs to include the inventory of confidential and sensitive data and the industry that the user belongs to.



Figure 4 – Attack Motivations

4. The Evolving Landscape

The IoT landscape—including IoT adaptation, threats, and solutions—is evolving at a rapid pace. While attacks are becoming more evasive and easier to conduct, the solutions should also evolve using AI-driven, data-centric, cloud-based solutions that allow protections to be adaptive. CSPs need to keep up with the demands of the ever-changing consumer threat landscape.

4.1. Threat Landscape

Everyone with access to the internet must have cybersecurity protection. Cyber-attacks reached a peak in 2021 as data breaches grew by over 17% from 2020, according to the Identity Theft Resource Center. In fact, cybersecurity is now considered a growing human rights issue, with the UN Security Council holding its second-ever cybersecurity meeting in 2020.

Businesses and government organizations aren't the only ones who should be worried about cyber threats. Consumers are increasingly vulnerable to these attacks as they fill their homes with connected devices. Our research has found that the number of devices per Plume-powered US household increased by 38% during the pandemic, with an average of 18 devices per household. And we're not just talking about laptops and smartphones. There was a 223% increase in virtual reality devices, a 132% increase in fitness bikes and trainers, and a 110% increase in smart light bulbs.



Figure 5 – Growth in IoT Devices at Home (Source: Plume)

It was hypothesized that criminals would likely take advantage of the fear, confusion, and increased use of the internet during a pandemic. Sadly, it turns out this was true. Figure 6 below shows the increase in various types of threats before and after COVID-19, with several attack types doubling in frequency. In fact, across the period of this study, 87% of the homes connected to Plume's network experienced some type of cybersecurity attack.




Figure 6 - Cyber-attack Trends

As cyber criminals create new ways to wreak havoc on smart homes, CSPs must stay one step ahead of them. A multi-layered approach to security, with services that act across every potential threat area—from anomaly detection and device quarantining to cyber-intrusion protection—is the answer. Analysis of data taken from the Plume Cloud showed that 87% of households were attacked. 85% of those attacks were DNS-based while preventing dangerous outbound and inbound IP events accounted for 37% and 8% respectively.

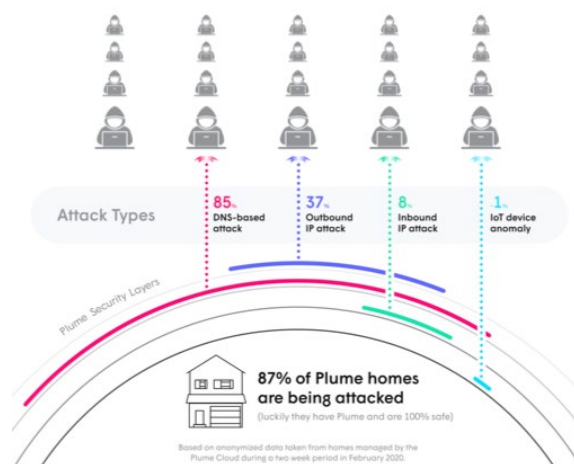


Figure 7 – Cyber-attack Types

4.2. CSP Security Services

Plume predicts that by 2024, connected devices in Plume homes will grow from the current 21 (or so) to exceed 38 devices across all categories—computers, mobile phones, tablets, set-top boxes, voice assistants, smart TVs, printers, surveillance cameras, game consoles, and more. Each of these devices opens a unique door to potential attacks. These attacks could be from websites and the servers they connect to. They could be the result of exploiting weak or reused passwords and unpatched software, and they could take the form of targeted phishing, spam, fraud attacks, and more.

CSPs are in the best position to deliver security to homes and small businesses. In addition, they can differentiate themselves by helping customers create a more secure environment—with more control over, and visibility into, their personal IT security. CSPs that adopt this role stand a good chance of achieving significant revenue gains; it's a win-win for CSPs and consumers.

CSPs can choose to build a complex and robust architecture that delivers a unified, multi-layer security service network, CPE, and endpoint that takes care of all the intricacies for the end users. This is time consuming and slows the delivery and adoption of services. Alternately the CSP can adopt a cloud-based, security-as-service model eliminating the dependencies and limitations of hardware-based solutions. The ongoing silicon shortage [22] continues to affect consumers' ability to get new hardware—be it a dishwasher, laptop, or network server. CSPs who rely too heavily on their hardware may have their hands tied by this issue. But those who offer cloud-based solutions can continue to adapt and expand.

Now that consumers are welcoming cloud-based services into their homes, CSPs should leverage this shift and set up the delivery platform to offer new value-added services quickly, using the same hardware already deployed on customer premises.

5. Vulnerabilities Life Cycle and Risks

Vulnerabilities are the primary source of attacks. A vulnerability lifecycle is divided into the following phases: Discovery, Disclosure, Patch, and Patch installed. Each of these life cycle phases has a corresponding risk exposure phase with unique characteristics and criticality. Risk exposure phases are discussed below (Figure 8).

It is challenging for a non-technical user to track vulnerabilities across the entire life cycle; however, they are constantly at risk right from the discovery until the patch installation phase.

5.1. Risk exposure phases

5.1.1. High Risk

High-risk phases live between the Discovery and Disclosure phases of the vulnerability lifecycle. In this phase, a subset of motivated attackers could be aware of this vulnerability and would already be engaged in developing an exploit. There is no public knowledge, resolution, or patch available in this phase. In this phase, behavior anomaly detection techniques are most relevant and effective to detect attacks and protect against compromises.

5.1.2. Elevated Risk

This phase exists when a vulnerability is discovered and disclosed publicly. Vendors are still working on a fix/patch and release plans while users wait. Exploits could already be publicly available for ready use by a wider community of attackers. It is very likely that the attacks are already active, making it the riskiest zone. In addition to the protection techniques of the high-risk phase, proactive vulnerability scanning, detection, and protection are the keys to safeguarding against attacks in this phase. A virtual patch to remediate the vulnerability is effective in restricting the exploits before the official patch is available.

5.1.3. Medium Risk

This is the period between patch availability and patch installation. Exposure during this period is under the direct control of users and vendors. The successful protection strategy includes publishing the vulnerability and patch availability, providing a patch installation mechanism, and ensuring that the patch is applied to the devices. Success in this phase relies on a collaborative approach between users and vendors.

A good security solution works on all the above phases and provides an end-to-end solution by breaking the cycle with complete vulnerability remediation.

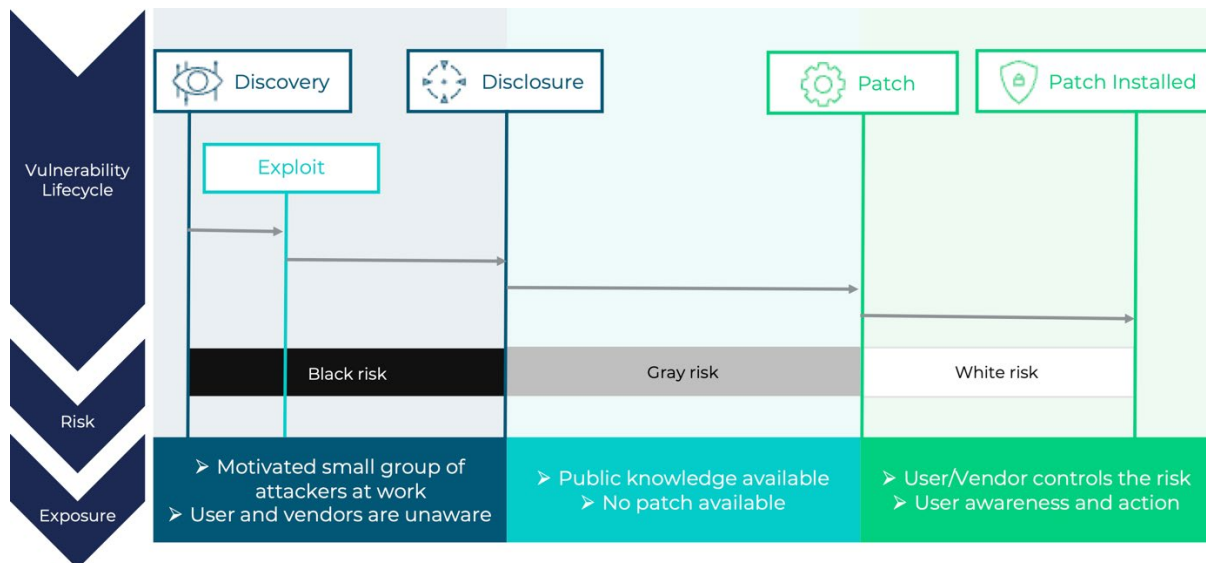


Figure 8 – Attack Lifecycle and Exposure Zones

5.2. Vulnerabilities and Open doors

5.2.1. Open ports

According to the latest stats from Shodan there are too many IoT devices operating with open ports. These open ports are like ticking time bombs, ready to be scanned by automated botnet crawlers, uncovering the known vulnerabilities, and leading to compromises and attacks.

Open and exposed ports are problematic as the services listening on the ports are often vulnerable to exploits. The open ports increase the attack surface for an attacker to exploit. As per a previous study by F5 [3] in total, 2,171,934 IoT ports were found to be exposed. Focusing on the most important 119 IoT ports, the top 10 exposed ports and their services are shown in Table 1 below. The listed 10 ports account for 84.7% of exposed ports in the Irish IP address space.

Table 1 - Top 10 Open IoT Ports

TCP Port	Service	Ports Open	% of Overall Exposed
443	HTTPS	772,258	35.6%
80	HTTP	670,789	30.9%
22	SSH	184,848	8.5%
3389	RDP	40,893	1.9%
8443	HTTPS-Alt	391,000	1.8%
8080	HTTP_Alt	30,502	1.4%
21	FTP	30,059	1.4%
8081	HTTP_Alt	27,187	1.3%
25	SMTP	23,901	1.1%
8000	Applications	21,028	1.0%

UPnP port is another critical port that is dangerous and a popular attack vector. UPnP allows zero-configuration connection implying no authentication is required to establish connections. Ports are forwarded automatically to establish a connection for a UPnP request, making it easy for attackers to establish an internet connection with the devices behind the firewall and exploit vulnerabilities. While the intended purpose of UPnP is convenience, it poses a serious threat to device security. UPnP is yet another technology that trades convenience for security. The enormity of the problem is apparent from the fact that there are 6M+ open UPnP ports worldwide. These are easily discoverable over Shodan (Figure 9). As per the July 2022 stats, most of these open ports are reported from devices in USA (12%) and China (11%) [4].

Total Results

6,201,899

Top Countries

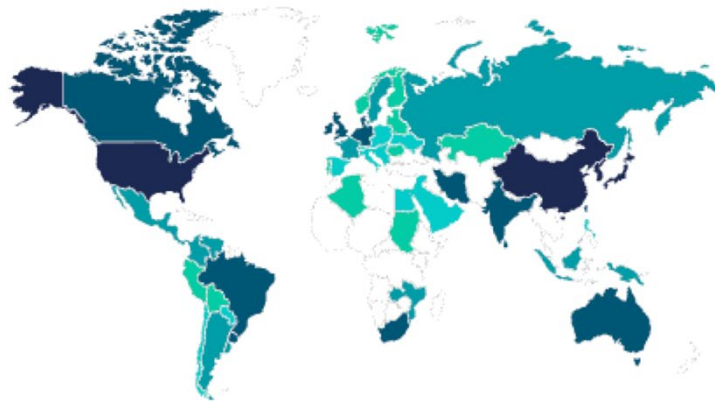


Figure 9 - Open UPnP Ports Globally (Shodan)

From the analyst reports, the top IoT device types that are exposed are routers, media devices, game console, NAS, printers, and smart TVs. The brands include some of the top names.

A recent attack (in the long list of attacks) was discovered in February 2022. Eternal Silence [1] (a UPnP based attack campaign) exposed 1.7 million devices to attacks via UPnPProxy abuse. UPnPProxy was reported back in 2018 by Akamai [2] researchers. The attackers targeted the routers vulnerable to UPnPProxy and exploited the unpatched vulnerabilities—EternalBlue (CVE-2017-0144) and EternalRed (CVE-2017-7494)—on unpatched Windows and Linux systems.

5.2.2. Vulnerability Numbers on the Rise

The number of vulnerabilities is increasing exponentially every year. In 2021, 20,169 new vulnerabilities were reported. Halfway through 2022, we already have 13,948 new vulnerabilities reported and counting.




Figure 10 – Vulnerabilities Reported Each Year

[12] Software vulnerabilities are one of the primary attack vectors in cyber incidents. Software vulnerabilities—together with phishing—make up 70% of the attacks [6] while software vulnerabilities alone contribute to 31%. Between high risk phase and the medium risk phase (Figure 8), attackers scan the network for known vulnerabilities and exploit them even before vendors can release and apply the patches. Time to patch these vulnerabilities is getting shorter and the vulnerability exploits are getting much faster, almost practically coinciding with the patch. For example, Palo Alto Networks released a Threat Prevention signature for the F5 BIG-IP Authentication Bypass Vulnerability (CVE-2022-1388), and within just 10 hours, the signature triggered 2,552 times due to vulnerability scanning and active exploitation attempts [6].




Figure 11 – Common Attack Vectors

The CISA [5] maintains a list of known exploited vulnerabilities. As per the list, 789 vulnerabilities have been recently exploited in the last year. New vulnerabilities are added to this list frequently based on the evidence of active exploitation.




Figure 12 – Most Exploited Vulnerabilities

Log4j remains the most highly exploited zero-day vulnerability in 2021-2022.

5.3. Vulnerability Prioritization

The following table shows vulnerability severity scores and their percentage distribution over the past year. It is important to note that more than 30% of the vulnerabilities reported were high- or critical-severity issues (CVSS >7). This implies that 30% have high attack and damage potential. It does not take advanced technical skills to exploit these vulnerabilities.

Table 2 – Distribution of All Vulnerabilities by CVSS Scores

CVSS Score	Number of Vulnerabilities	Percentage
0-1	1,007	0.60
1-2	1,196	0.70
2-3	8,327	4.60
3-4	9,460	5.20
4-5	42,973	23.70
5-6	34,116	18.80
6-7	27,136	15.00
7-8	36,000	19.90
8-9	895	0.50
9-10	19,978	11.00
Total	181,088	

What is CVSS?

The Common Vulnerability Scoring System (CVSS) is an open framework for communicating the characteristics and severity of software vulnerabilities. The NIST National Vulnerability Database [8] (NVD) defines the CVSS score as a mechanism for organizations to properly assess and prioritize their vulnerability management processes.

CVSS comprises three metrics: Base, Temporal and Environmental. Base metrics provide a static score ranging from 0 to 10. Temporal scores define metrics that change over time due to events external to the vulnerability. Environmental scores define metrics customized to reflect the impact of the vulnerability on your organization. The Base score can then be modified by scoring the Temporal and Environmental metrics. The NVD does not currently provide “temporal scores” or “environmental scores”. NVD does provide a CVSS calculator to allow an organization to compute the Temporal and Environmental score data.

Consumer organization and security solutions should augment the CVSS system with a risk-based vulnerability prioritization system for isolating the immediate focus areas. Exploitability is the primary factor that decides the vulnerability priority and includes an analysis of essential factors such as threat landscape, attack taxonomy, industry vertical and geolocation, etc. In addition, the

security solutions also need to leverage the data feeds from the sources like CISA [5] and prioritize remediating vulnerabilities that are actively exploited.

6. Vulnerability And Attack Taxonomy

6.1. IoT architecture

A generic IoT architecture is a hierarchical model with four layers: Application, Middle, Network, and Device [9].

6.1.1. Application

This layer implements different applications for different IoT scenarios and business verticals. As this is the front end for IoT solutions, the security issues mainly arise from authentication, illegal access, data theft, and permissions. Attackers can exploit software vulnerabilities to attack systems and disrupt functionality.

6.1.2. Middle layer

The middleware layer or the service support layer sits between the network layer and backend cloud systems. This layer obtains the data from the network layer and connects the system to the cloud and data repositories. This layer is also responsible for data processing and storage. Data repository security and cloud security are the main concerns in the middleware layer, as these can affect the quality of service in the application layer.

6.1.3. Network

This layer is responsible for the connectivity of the IoT infrastructure. It also collects data from the device layer and transmits it to the upper layer. The transmission medium can be wired or wireless, and the main technologies are ZigBee, Wi-Fi, Bluetooth, 3G, and so on. Attacks on the network layer are diverse, typically affecting the coordination of work and information-sharing among devices.

6.1.4. Device

The main challenges for this layer are the attacks on sensors and identification technology, which interfere with the collection of data from devices. An attack can send incorrect device states and other crucial statistics interrupting the entire system.

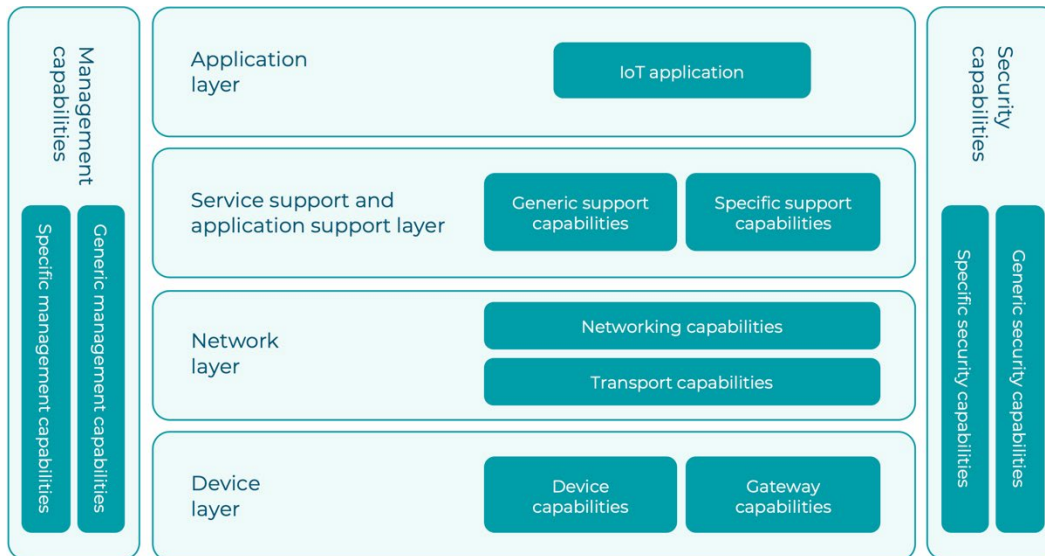


Figure 13 – IoT Reference Model

6.2. The Taxonomy of Vulnerabilities in IoT

To understand the IoT vulnerabilities and their impact, this section discusses vulnerability taxonomy. This taxonomy provides a good reference for CSPs and network defenders.

6.2.1. Weak authentication mechanisms

The interfaces in the IoT ecosystem like mobile, cloud, firmware, and web should be secured with strong authentication mechanisms. Weak, guessable, and hardcoded passwords on these interfaces give attackers unauthorized access to the IoT ecosystem. These vulnerabilities can be exploited in numerous ways and are used as one of the entry points into the network. Some common attacks that leverage authentication vulnerabilities are DDos attacks, Dictionary attacks, Sybil attacks, etc.

The Mirai attack [13], which almost brought down the internet, was carried out by compromising various IoT devices that were configured with default weak credentials (say: admin/admin).

6.2.2. Insecure network services

Unwanted and vulnerable network services exposed and listening on devices are easily compromised to gain access to the device itself, to inject malicious code, modify firmware, bypass security, move laterally to scan other devices, and infect more devices. A wide range of attacks can be launched via these open ports.

Such vulnerabilities are fatal, especially when exposed over the internet. The attackers can gain access to the devices and remote control the network.

An SSH port (22) with an old and vulnerable service version is one of the main attack vectors used in Ransomware attacks.

6.2.3. Lacking privacy and data protection

Users' personally identifiable information (PII) and other personal data are stored everywhere in the ecosystem including devices, the middle layer, and the network layer. If the data is stored, accessed, or processed without proper access control policies and encryption, it can lead to data breaches/losses with a significant damaging impact on personal lives and businesses.

Data breach incidents have been rated topmost when it comes to revenue, brand value, and customer trust loss. In addition, stolen personal information is available and sold on the dark web [17] for minimal cost.

Most IoT devices use wireless communication media, like Zigbee, LoRa, 802.11. a, SigFox, and 802.15.4. These protocols are less reliable and as a result, the devices are more susceptible to data leakage attacks.

Organizations must comply with regulatory guidelines such as CCPA and GDPR when handling personal and business sensitive data. Country-specific guidelines should be followed to safeguard the users when handling of their personal information. IoT devices require a central security policy [16] to correctly handle personal data including detecting, operating, collecting, and storing data. PKI should be used, where possible, to provide a robust encryption methodology rather than relying on hard-coded secrets to authenticate. The data-in-motion should be securely transmitted over the network ensuring that confidentiality, integrity, and availability are guaranteed.

6.2.4. Weak device hardening

The device boot process is vulnerable to attacks if the secure boot is not implemented. Malicious actors can compromise the firmware, boot loader, and boot process sequence by replacing legit executables with a malicious component. With a Secure Boot process in place, the reboot process would identify the malicious executable file and prevent it from running.

Different types of rootkits load at different phases of the boot process: Device manufacturers should follow the best practices and implement secure boot on the devices. Various types of rootkits have been discovered and discussed previously [15]

The recent UFEI rootkit, CosmicStrand [14], discovered in July 2022 is a classic example that rootkits are not rare.

The discussion above explains how IoT vulnerabilities open up a plethora of attack vectors for adversaries and the possible attacks that originate from each vulnerability. Alternatively, the figure below (Fig: 14) provides a view of the layer-based examination of the potential attacks concerning each layer of the IoT architecture.

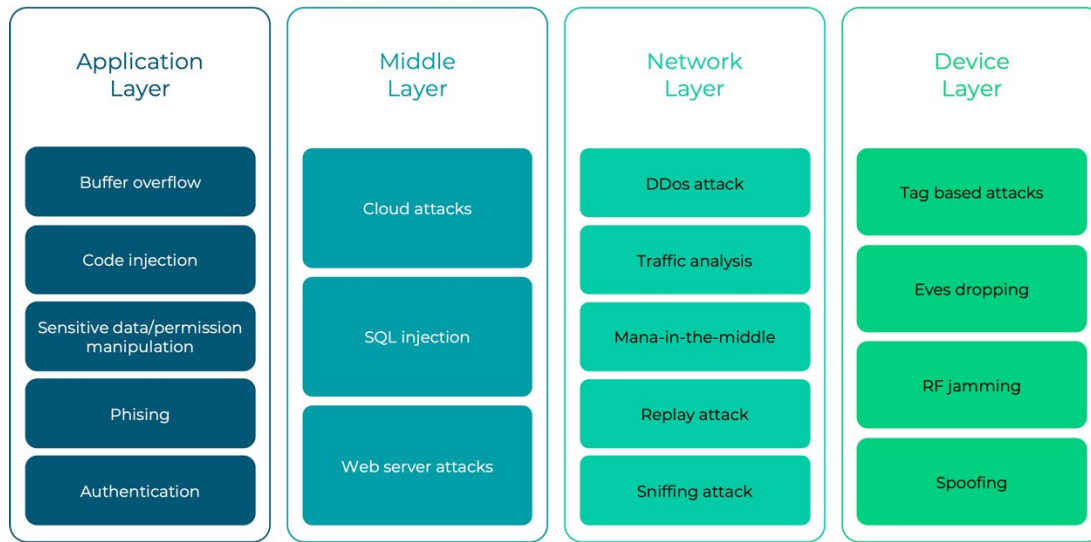


Figure 14 – IoT Reference Model Taxonomy

6.3. IoT Attack Flow

6.3.1. Use-case of Ransomware

According to leading analysts' reports, the use of ransomware is on the rise, with an average of 144% increase in demands. Attackers use multi-extortion techniques including encryption and name-shaming.

Ransomware is emerging as a productive business model enabling even a novice attacker with little to no technical knowledge to rapidly launch an attack. There are ransomware kits and services available to cybercriminals that remove technical hurdles and lower the bar for participation.

Software vulnerabilities are the primary attack vector for ransomware. In 2021, attackers exploited multiple high-profile vulnerabilities to gain a foothold in homes and small businesses. The duration between vulnerability disclosure and exploit availability has reduced considerably. If vulnerabilities are available, attackers will exploit them and launch attacks. This has made it challenging for organizations and security solutions. Security solutions need to detect and remediate vulnerabilities rapidly.

Log4J is an example of how vulnerabilities are weaponized and exploited at a rapid pace. The following vulnerabilities were exploited in a ransomware attack path [7].

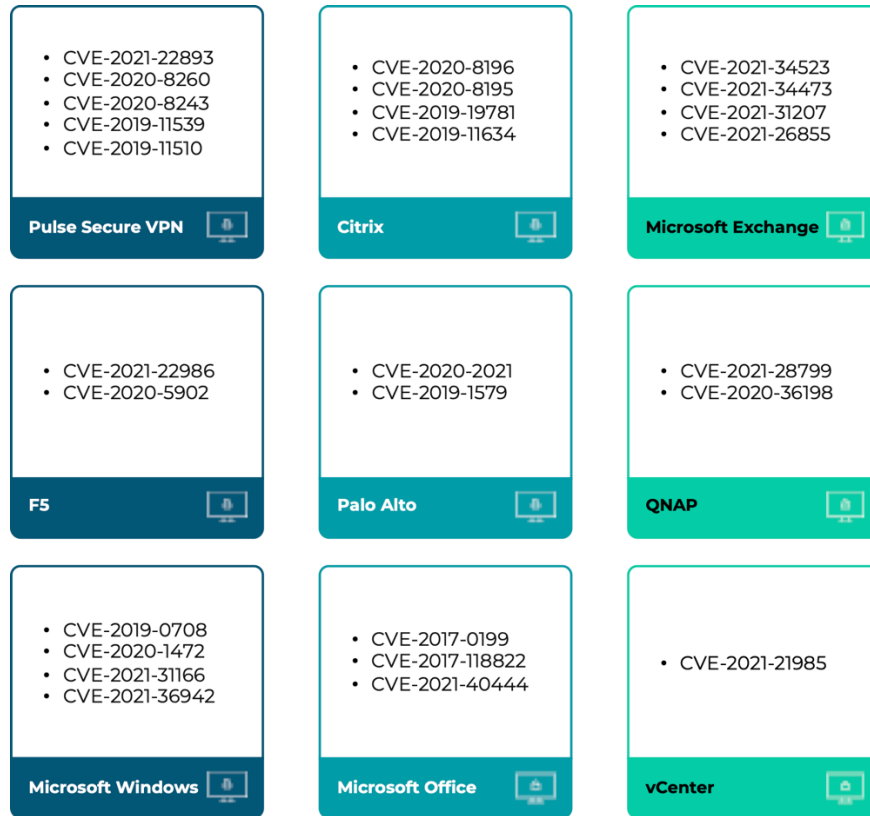


Figure 15 – Vulnerabilities exploited in Log4J

Table 3 – Vulnerabilities Exploited in Ransomware Attack Path

Vulnerability identified	CVE-2021-44228 , CVE-2021-45046 , CVE-2021-44832 , CVE-2017-5645 , CVE-2021-45105 , CVE-2019-17571
Vulnerability category	Remote code execution, denial of service

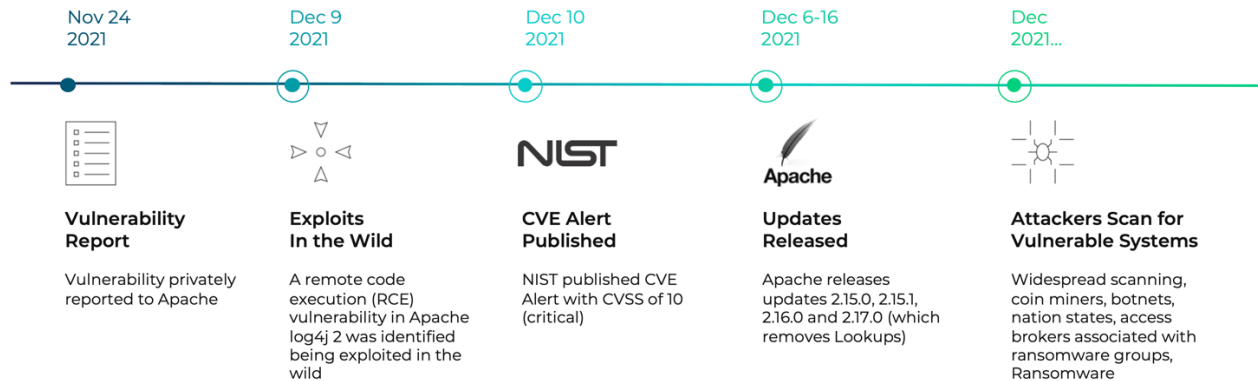


Figure 16 – Log4J Attack Timeline

6.3.2. Taxonomy of Ransomware

A ransomware attack typically follows a path with three stages. In stage one, it tries to gain access to the home or business. This is done using internet-facing devices, vulnerable devices, and devices with default credentials. In the second stage, it moves laterally, infecting other devices and scanning for business-critical, sensitive, or personal data. In the last stage, the attacker launches the attack by stealing, locking, or destroying data.

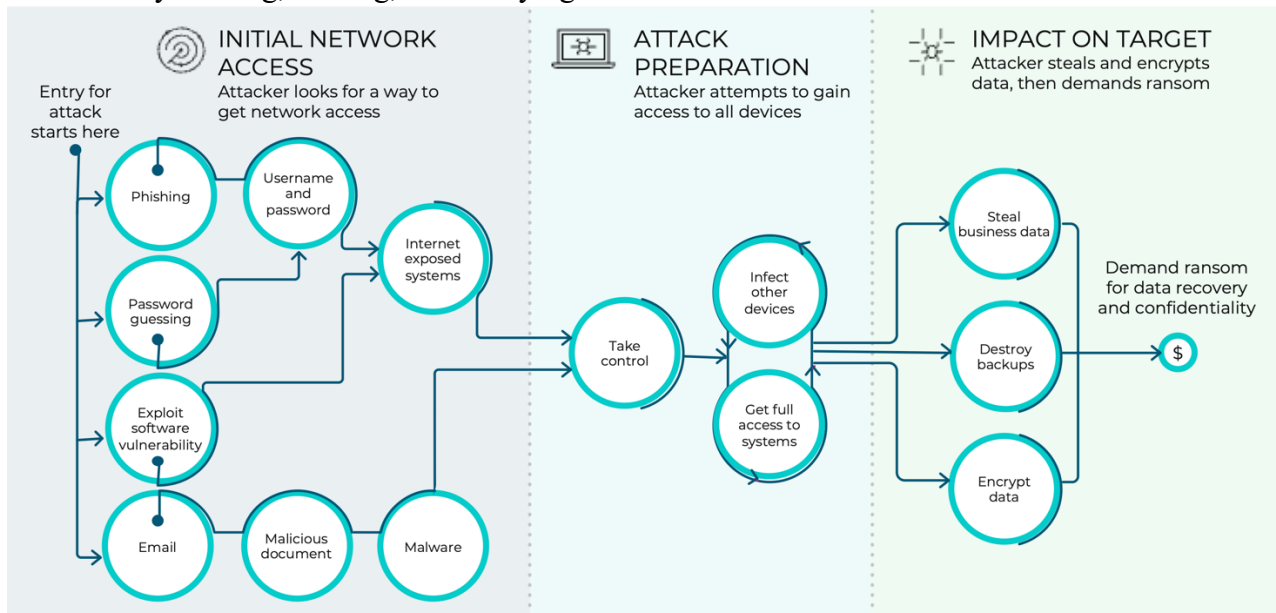


Figure 17 – Ransomware Attack Path

7. Solution

In the ITU IoT (Figure 13) reference models, different layers have different security requirements:

- Application layer – This layer needs the authorization, authentication, application data confidentiality, and integrity protection, privacy protection, security audit, and anti-virus.
- Network layer – This layer needs authorization, authentication, user data and signaling data confidentiality, and signaling integrity protection.
- Device layer – This layer need authentication, authorization, device integrity validation, access control, data confidentiality, and integrity protection.

Specific security capabilities are closely coupled with application-specific requirements, for example, mobile payment and security requirements.

When designing the security for IoT devices, Plume considers it important to take custom security procedures into account in addition to conventional security procedures. It would be best if the solution assures device security, network security, and the overall security of the IoT architecture and system.

Top analysts rate software vulnerabilities as the most popular attack vector among the cybercriminal community. The following section describes how CSPs can solve many of the problems associated with insecure IoT by providing proactive and continuous vulnerability protection.

7.1. Vulnerability Detection and Protection

Vulnerability insight is part of a multi-layered security strategy and improves the overall network defenses and zero-trust. It is a preventive strategy, and its job is to proactively inform users of known vulnerabilities in devices and help them fix the issues by taking proactive, preventative action.

The solution should protect routers and other devices in the network that are subject to attacks from open ports (capable of infiltrating network defenses) putting user data, finances, and privacy at risk.

A vulnerability management solution has three key components: Detection, protection, and reporting. The section below discusses them in detail (Figure 18). Asset discovery and inventory is a pre-requisite for this solution. Organizations can leverage their existing methods — the details of doing that are out of the scope of this paper.



Figure 18 – Vulnerability Detection and Protection Solution

7.1.1. Detection

Visibility and an accurate picture of the network is the key to helping CSPs provide an effective solution. The detection phase is critical to identifying, exposing, and prioritizing the possible threats and weaknesses in the network.

Discovery scan

Detection requires continuous scanning and monitoring of the devices inside the home and small business network. The goal of this scanning is to generate insights about the security posture of the home or business. This is achieved by detecting the device characteristics like firmware, OS, open ports, the services running on those ports, and any port forwarding (UPnP) configuration. All these characteristics should be analyzed and mapped to the known vulnerabilities on the devices. Then vulnerable devices can be automatically blocked from accessing other devices on the home and business network.

Vulnerability scan

Detect if the open ports and the services running on them are vulnerable to brute force or password stuffing attacks. To do this, CPS should be able to find out if the OS running on the devices has any known vulnerabilities. Use opensource or homegrown network scanning tools, like nmap [18] to identify the ports open on the devices on the network and the service versions running on them. Leverage the vulnerabilities databases like NIST [19] and MITRE [20] to identify any known vulnerabilities reported in the discovered services.

In addition, scan the services to detect if the service login credentials are weak or default. Any services running with weak passwords are vulnerable to brute force attacks.

7.1.2. Prevent and Protection

Protection approaches can vary based on the technical knowledge and capability of the user. They can range from a complex patch update to step-by-step guided remediation. The goals resolve the

issue with minimal user engagement. The following are the possible remediation methods that a good vulnerability protection solution should provide:

Vendor advisory or patch

Provide the link to the vendor advisory or patch where available. The user can follow the advisory and resolve the issue themselves. Example: Firmware upgrade steps, a password change, or temporarily removing a device from the network.

Credential brute force protection

Alert the user if the credentials on the devices are default, weak, and easily hackable. Alert users if the same username/passwords are used for multiple devices/services.

Best practices and guidelines

Provides device-type-specific and a well-curated list of best practices. Provide password hygiene guidelines like a reminder to enable 2FA on the devices, a reminder to change passwords periodically, and a reminder not to repeat passwords for multiple devices.

Virtual patch - continuous protection

The complete vulnerability detection and protection life cycle is very long. It includes the following vulnerability exploitation factors->patch availability-> patch application. The system is vulnerable to attacks during this phase. The virtual patch is a concept that provides vulnerability-specific rules to prevent malicious traffic from targeting the device.

- Pre-patch protection - Detect and prevent traffic exploiting the vulnerability
 - IDS/IPS rules for the vulnerabilities found on the device.
 - Alert when the rule triggers to update the user that they have been protected against the vulnerability.
- Post patch protection - be on the lookout for the vendor patch availability and update the user when the patch is available.

Enhanced security level

Create and deploy a security policy for enhanced and stricter controls for vulnerable devices and homes. Possible policies include:

- Full DPI-based malware protection.
- Block malicious IP/domains.
- Prioritize DOOR alerts for the impacted devices.

Limit device internet exposure

Alert the user if the vulnerable device is exposed on the internet. Proactively limit network access paths to the device and consider disabling UPnP on the vulnerable device.

Internal honey pot

Install a passive detection system in the network. This creates a virtual IP with the most common ports open. Any connection attempt to these services will throw an alert to flag a potential threat. Users can blacklist the connecting IP to proactively protect the network.

Chatbot

Provide a specialized and trained chatbot to guide the user through the remediation steps.

- Provide the user with the unique problem ID w.r.t. The vulnerability
 - E.g.: Say for [CVE-2020-25687](#), Id is “Vuln-2020-25687”
- Problem specific guidelines
 - The user can use the chatbot with the given ID and get step to resolve the problem.

Call support

Solution providers should provide a support center where the user can get help resolving issues. Customers can choose to call vendor support directly or the solution provider. Alternately, CSPs can partner with the vendor support and report issues on behalf of the end user.

7.1.3. Reporting

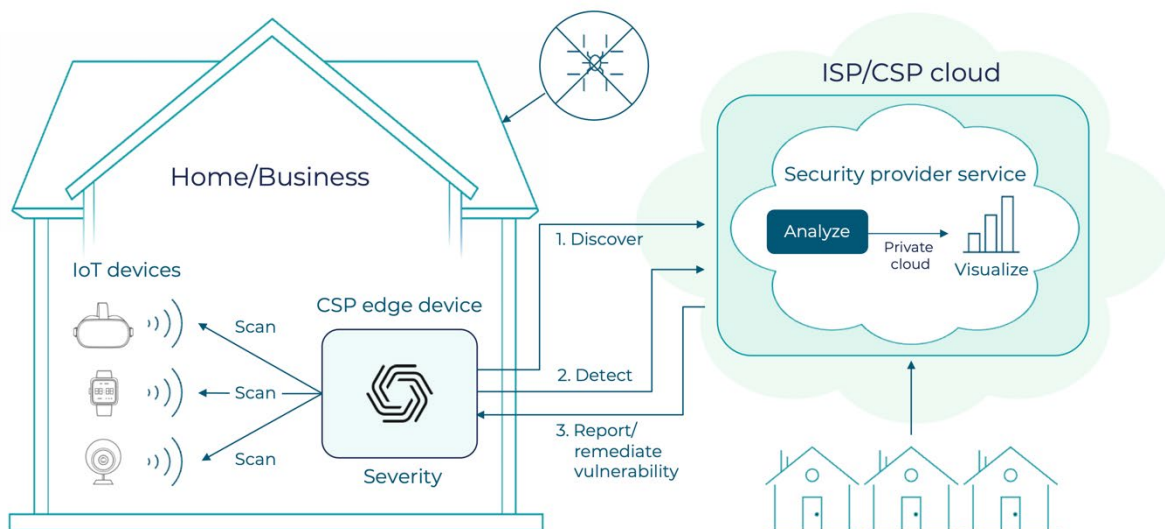


Figure 19 – Reporting as part of the vulnerability protection solution

Reporting is a critical part of the vulnerability protection solution and should focus on providing an actionable, prioritized subset of vulnerabilities to the user or administrator. An overall threat score would give a consolidated view of the security posture of the network. The security score would comprise of several factors to measure the security state i.e., number and severity of vulnerabilities, exploitability, time of exposure to a vulnerability, and more.

8. Conclusion

IoT for homes and small businesses has its merits, but also comes with significant security challenges due to rapid technology transformation and a lack of knowledge. IoT is like an open door with lucrative low-hanging fruit for attackers. Openly available attack kits are further lowering the bar for attackers. Software vulnerabilities, weak passwords, and open ports are the primary attack vectors. It is possible, however, for CSPs to provide value-added security services and safeguard consumers against cyber threats. This is a win-win scenario for both CSPs and users. We need an intelligent vulnerability detection and remediation solution that can continuously scan and flag vulnerabilities and well scan for weak passwords and provide remediation guidance to users. Considering the evasive nature of attacks and the short time spans within which they take place, we need to invest in an AI-based behavioral solution that will provide proactive alerts when a device exhibits malicious behavior and take action to protect the network.

Let's use IoT safely and improve our productivity. Let's close the door on attackers before they close your business.

9. Abbreviations

AI	artificial intelligence
CSP	communications service providers
CVE	common vulnerabilities and exposures
CVSS	common vulnerability scoring system
DNS	domain name system
IoT	Internet of Things
IT	information technology
ITU	International Telecommunication Union
MITRE	Massachusetts Institute of Technology Research and Engineering
NIST	National Institute of Standards and Technology
NVD	National Vulnerability Database
TV	television
WFH	work from home

10. References

- [1] <https://www.akamai.com/blog/security/upnproxy-eternal-silence>
- [2] <https://www.akamai.com/site/en/documents/research-paper/upnproxy-blackhat-proxies-via-nat-injections-white-paper.pdf>
- [3] <https://www.f5.com/labs/articles/threat-intelligence/iot-vulnerability-assessment-of-the-irish-ip-address-space>
- [4] <https://www.shodan.io/dashboard>
- [5] <https://www.cisa.gov/known-exploited-vulnerabilities-catalog>
- [6] <https://start.paloaltonetworks.com/2022-unit42-incident-response-report>
- [7] cisa.gov <https://www.cisa.gov/stopransomware>: Stop Ransomware; Cybersecurity and Infrastructure Security Agency
- [8] NIST <https://nvd.nist.gov/vuln-metrics/cvss>: National Vulnerability Database: National Institute of Standards and Technology
- [9] ITU <https://www.itu.int/rec/T-REC-F.748.5-201511-I/en>
- [10] <https://ieeexplore.ieee.org/document/9189773> , 2020: IoT Vulnerability Assessment for Sustainable Computing: Threats, Current Solutions, and Open Challenges
- [11] CISOMAG <https://cisomag.com/ciso-mag-survey/>, 2020: 1 in 3 Employees Don't Use VPN to Connect to Company Network While Working from Home: CISO MAG Survey; CISOMAG
- [12] CVE Details <https://www.cvedetails.com/>: Current CVSS Score Distribution For All Vulnerabilities; CVE Details
- [13] <https://www.csoononline.com/article/3258748/the-mirai-botnet-explained-how-teen-scammers-and-cctv-cameras-almost-brought-down-the-internet.html>
- [14] <https://securelist.com/cosmicstrand-uefi-firmware-rootkit/106973/>
- [15] <https://www.blackhat.com/docs/asia-17/materials/asia-17-Matrosov-The-UEFI-Firmware-Rootkits-Myths-And-Reality.pdf>
- [16] https://www.researchgate.net/publication/325562955_Smart_IoT_Devices_in_the_Home_Security_and_Privacy_Implications: 2018, ResearchGate
- [17] https://www.researchgate.net/publication/331867659_Dark_Web_and_Its_Impact_in_Online_Anonymity_and_Privacy_A_Critical_Analysis_and_Review, 2019: ResearchGate
- [18] <https://nmap.org>: NMAP

[19] NIST <https://nvd.nist.gov> : National Vulnerability Database: National Institute of Standards and Technology

[20] MITRE <https://cve.mitre.org>: Massachusetts Institute of Technology Research and Engineering

[21] IOTWorld <https://www.iotworldtoday.com/2021/11/16/botenago-malware-targets-millions-of-iot-devices/>

[22] IEEE spectrum <https://spectrum.ieee.org/chip-shortage>

[23] CEPRO <https://www.cepro.com/news/57-percent-iot-devices-vulnerable-attack/>

[24] DarkReading <https://www.darkreading.com/iot/babel-of-iot-authentication-poses-security-challenges>

Service Stability

A Data Analytics Approach

A Technical Paper prepared for SCTE by

Jeffrey Lee

Manager, OSS Data Acquisition and Reporting
Shaw Communications
2400 32 Ave. N.E., Calgary, Alberta T2E 6T4
jeffrey.lee@sjrb.ca

Achintha Maddumabandara

Principal Data Analyst
Shaw Communications
2400 32 Ave. N.E., Calgary, Alberta T2E 6T4
achintha.maddumabandara@sjrb.ca

Stuart Mann

Network Architect
Shaw Communications
2400 32 Ave. N.E., Calgary, Alberta T2E 6T4
stuart.mann@sjrb.ca

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Statistical Process Control	4
2.1. Control Charts	4
3. CM RF vs Service Metrics.....	5
4. CM RF Control Charts.....	7
5. Internet Service Control Charts.....	9
6. Operationalization	14
7. Future Work.....	14
8. Conclusion.....	14
Abbreviations	16
Bibliography & References.....	16

List of Figures

Title	Page Number
Figure 1 - Future of Connectivity.....	3
Figure 2 - Anatomy of a Control Chart	5
Figure 3 – CM RF Signal Quality vs Internet Speed.....	6
Figure 4 - CM DS OFDM Channel SNR vs Error Percentage	7
Figure 5 - CM DS SC-QAM SNR vs Error percentage	7
Figure 6 - Control Chart of CM Flap.....	8
Figure 7 - Control Chart of CM DS OFDM Channel SNR.....	8
Figure 8 - Control Chart of CM DS OFDM Channel Receive Level.....	8
Figure 9 - Control Chart of CM DS OFDM Channel Error%	9
Figure 10 – Control Chart of CM US Speed Test of Higher Stability Service.....	10
Figure 11 - Control Chart of CM US Speed Test of Lower Stability Service	10
Figure 12 - Control Chart of CM DS Speed Test of Higher Stability Service.....	11
Figure 13 - Control Chart of CM DS Speed Test of Lower Stability Service	11
Figure 14 - Control Chart of CM DS Latency of Higher Stability Service	12
Figure 15 - Control Chart of CM DS Latency of Lower Stability Service	12
Figure 16 - Control Chart of CM DS Jitter of Higher Stability Service	13
Figure 17 - Control Chart of CM DS Jitter of Lower Stability Service	13
Figure 18 - Speed to Service Stability Readiness	14

1. Introduction

Technologies such as wearables, smart homes, smart cities, and smart industries are changing customer expectations for internet services. These new technologies prioritize consistency and responsiveness over sheer volume. Customers expect their smart devices to be responsive and always on. Customers want their wearables and smart home devices such as thermostats, windows, doors, smoke & fire detectors, and security systems to respond to them in near real time. Similarly, smart city and smart industries also have business requirements for low latency and internet service consistency. Customers and businesses are less tolerant of issues regarding network connectivity and latency variability as the demand for “connectedness” is ever increasing amongst internet users and machines.

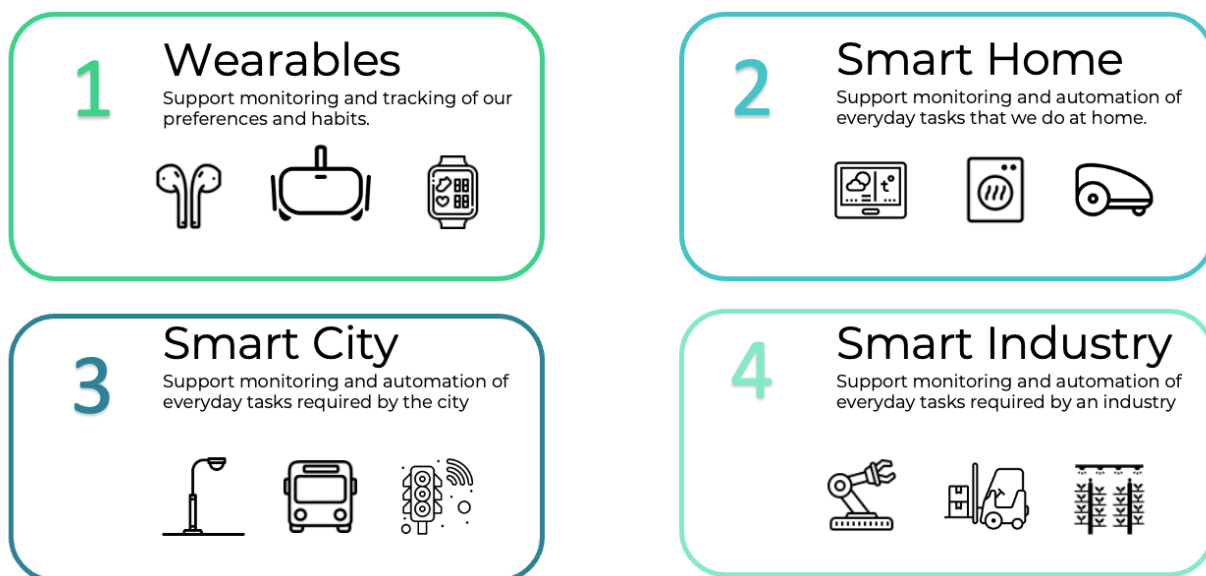


Figure 1 - Future of Connectivity

Operators are developing new methods to better understand their customers’ service experiences. Operators will need to measure variability of service in addition to threshold performance.

Traditional radio frequency (RF) diagnostic tools measure current and historical values based on manufacturer thresholds and are effective at fault detection. This means assessing if a CM has exceeded a performance threshold that may cause performance degradation. We need a different method for determining the stability of service or degree of variation in service. The lower the variability means the more reliable the service. High variability would indicate that additional measures are required to stabilize the service.

Statistical Process Control (SPC) is a cornerstone of heavily adopted process improvement methodologies such as Six Sigma. SPC can provide operators with the tools to measure and monitor service stability, which can result in improved service quality and reliability to meet the needs of advancing technologies.

2. Statistical Process Control

SPC is a statistics methodology for measuring and monitoring the variability of a desired process and can be implemented as a tool. SPC can provide us with insights about variability based on changes of averages, upper control limits, and lower control limits calculated using conventional statistics methodology. Control charts in SPC are an effective way to visually communicate variability in a concise and easily understandable graph.

In the 1920s, American physicist, engineer, and statistician Walter A. Shewhart developed the control chart used in statistical process control. The application of his detection methodology at Bell allowed engineers to observe and measure the stability of their system and thereby improve quality and reliability of Bell's transmission systems.

Today Shewhart's control charts are still in use in process improvement methodologies such as Six Sigma. In Six Sigma control charts are used to determine quality characteristics.

The primary rule or Shewhart's rule identifies assignable cause whenever a single point falls outside the three-sigma limits. Shewhart—and later Lloyd S. Nelson and Western Electric—developed additional rules that can be used to further increase detection sensitivity to variance at the cost of a higher false positive rate.

Using SPC we can improve process quality and reliability by measuring and monitoring the effectiveness of applied controls to our processes.

2.1. Control Charts

Control charts are graphs that can help us visualize the state of control for a given process. There are various types of control charts used to depict different aspects of SPC.

The center line denoted as \bar{x} is the average of individual values of the samples collected and is calculated using:

$$\bar{x} = \frac{\sum_{i=1}^n x_i}{n - 1}$$

x_i = Each of the values of the data

n = The number of data points

The upper control limit (UCL) or 3σ and the lower control limit (LCL) or -3σ are calculated using the standard deviation of a sample formula:

$$\sigma \text{ (sample rather than population)} = \sqrt{\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n - 1}}$$

n = The number of data points in sample

x_i = Each of the values of the data

\bar{x} = The mean of x_i

Common-cause variation (or noise) is the natural or expected variation in a process. Generally, noise is found between the upper and lower control limits. There are exceptions where distinct patterns within the upper and lower control limits have been identified by Shewhart, Nelson, and Western Electric as special cause variation; however, those rules will not be covered in this paper.

Special-cause variation (or signal) is the unexpected variation that results from unusual occurrences. Generally, signal is found above the UCL or below the LCL; however, there are exceptions where this would not be the case.




Figure 2 - Anatomy of a Control Chart

3. CM RF vs Service Metrics

CM RF metrics describe RF signal quality and performance. These metrics can be collected from CMs via simple network management protocol (SNMP). Regular collection of RF metric data has a negligible effect on the service performance of a cable modem. Hence, it is possible to collect this data at a higher frequency than internet service metrics which results in a higher resolution analysis of RF signal quality. It is important to understand that RF signal quality is strongly correlated with internet service quality but cannot be used directly quantify the internet service experience.

Consider the following:

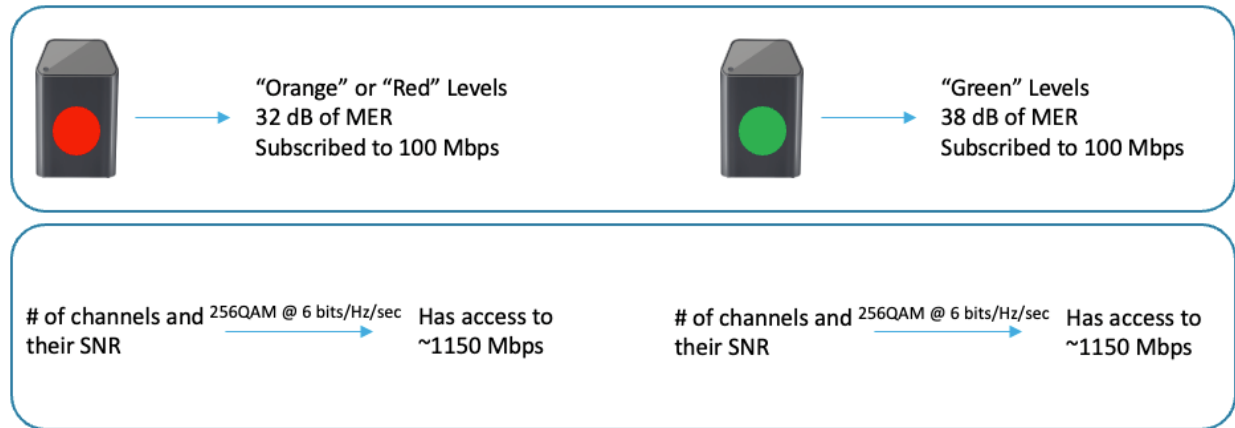


Figure 3 – CM RF Signal Quality vs Internet Speed

When CM RF metrics degrade, they do not correlate to an equivalent degradation of internet service metrics. For example, the CM RF signal to noise ratio (SNR) degrading a few dB does not equate to a consistent drop in upload or download speeds. There are thresholds, such that once exceeded, there is an observable impact to internet service metrics as seen below in Figure 4 and Figure 5.

Internet service metrics such as upload speed, download speed, latency, and jitter which are a true representation of the internet service experience is more challenging to collect. The collection of this data can be customer impacting during the collection process and thus, would be scheduled less often and during non-peak hours which may not be representative of internet service under stress.

The scatter plots below depict a correlation between SNR and uncorrectable codewords (denoted as Error %). We observe an increase in the presence of uncorrectable codewords as the SNR degrades below ~32dB for the downstream orthogonal frequency-division multiplexing (OFDM) channel in Figure 4, and below ~35dB for downstream single carrier quadrature amplitude modulation (SC-QAM) in Figure 5.

The figures below depict variability in the correlation strength of data points, which suggests that individual CMs may be impacted differently by degraded SNRs, leading to varying amounts of Error%.

SPC would allow an operator to measure each CM against itself to determine variance of a given metric. CMs that have a high degree of variability in RF or internet service metrics would require further diagnostics.




Figure 4 - CM DS OFDM Channel SNR vs Error Percentage

Figure 5 depicts more variability in correlation strength between SNR and Error %, hence the addition of a regression line (in red) to better represent this relationship.




Figure 5 - CM DS SC-QAM SNR vs Error percentage

4. CM RF Control Charts

CM RF metrics such as modem flaps, SNR, receive power level, and % of errors are captured in the control charts below. CM RF metrics are sampled at regular intervals, except for codewords, which are calculated as a percentage from a cumulative counter residing on the modem. An increase or decrease in variation in one metric does not reflect consistently in other metrics.

CM flap or CM de-registration occurs when a CM loses connection to the cable modem termination system (CMTS). This metric can also be considered an internet service metric as it is an important indicator of the customer service experience. An increase in the average count or range between the UCL and LCL of this metric would likely impact the customer experience and should be prioritized over other RF metrics.




Figure 6 - Control Chart of CM Flap

SNR, transmit, and receive power levels can fluctuate based on many external variables. The tolerance of these metrics can be vendor, model, and device specific. Variability in these metrics below or above the manufacturer specification can result in intermittent service but would be a lower priority than CM flaps and error %.




Figure 7 - Control Chart of CM DS OFDM Channel SNR




Figure 8 - Control Chart of CM DS OFDM Channel Receive Level

Uncorrectable error % is an important RF metric that can impact internet service experience. Even a small increase in this metric can be impacting to customers (between 0 and 1% of total codewords).




Figure 9 - Control Chart of CM DS OFDM Channel Error%

CM RF control charts depict variation of RF signal quality over time. The upper and lower control limits help us to quickly visualize the range of variation and if metric values can be classified as special cause. Many operators may find it difficult to determine if the overall effect of a device firmware change was positive or negative for a given device. Using this methodology, an operator can quickly visualize and quantify if a network change has improved or degraded service stability for devices in its network.

We would expect that a good change will narrow the range of variation in RF metrics meaning the UCL and LCL will be closer together and align the center line closer to the mean of the manufacturer specifications. Conversely a bad change would increase variability and move the center line closer to the threshold limits.

In either case, the control charts will continue to be effective as the UCL and LCL are calculated using the third standard deviation of the metrics. Hence, as the variation decreases then the sensitivity will increase and vice versa.

5. Internet Service Control Charts

The following control charts depict internet service stability using data from upstream speed test, downstream speed test, jitter, and latency. CM RF metrics are normally sampled at regular intervals on CMs in the field. However, at the time these control charts were created, CM RF metrics were not being collected from CM devices at regular intervals and as such we are unable to correlate insights between the CM RF and CM internet service metrics directly.

Upstream (US) and downstream (DS) speed tests have a direct relationship to the customer service experience. US and DS speed tests are affected by many external variables and tend to vary over time. A control chart is an ideal method for evaluating US and DS speed test performance over time. The control chart centerline should meet or exceed the advertised speed. The UCL and LCL should not exceed the degree of variation that the internet service provider (ISP) has specified and would allow ISP to quantify and proactively address service experience issues before a customer may notice.

Figure 10 and Figure 11 depict a control chart for a 15 Mbps US service package. The centerline for Figure 10 is at 16 Mbps while the centerline for Figure 11 is at 14 Mbps. Since Figure 11 has achieved an US speed tier (16 Mbps) higher than the advertised US speed (15 Mbps), it is possible that this CM may not have registered as a priority for service maintenance compared to other services which may have achieved a lower speed tier. However, the UCL and LCL clearly depict a different story. The service

shown in Figure 11 is clearly unstable and shows a large degree of variation which would be impacting the customer's service experience.




Figure 10 – Control Chart of CM US Speed Test of Higher Stability Service

Upper control limit was cropped from this chart as upstream speed is capped at 16mbps whereas the upper control limit was 25mbps




Figure 11 - Control Chart of CM US Speed Test of Lower Stability Service

The control charts shown below for DS speed tests tell a similar story to the US speed tests seen above. Both Figure 12 and Figure 13 show services that have met and exceed the advertised DS speed tier (150 Mbps) however, the amount variation seen in Figure 13 and the frequency of tests crossing the LCL is concerning.



Figure 12 - Control Chart of CM DS Speed Test of Higher Stability Service




Figure 13 - Control Chart of CM DS Speed Test of Lower Stability Service

Latency and jitter are increasingly important metrics for driving positive customer experiences. The control charts below clearly depict the level of variation as well as the centerline target. Figure 14 and Figure 16 are highly stable services with a satisfactory UCL and LCL whereas Figure 15 and Figure 17 are unstable and poorly performing services.




Figure 14 - Control Chart of CM DS Latency of Higher Stability Service




Figure 15 - Control Chart of CM DS Latency of Lower Stability Service




Figure 16 - Control Chart of CM DS Jitter of Higher Stability Service




Figure 17 - Control Chart of CM DS Jitter of Lower Stability Service

Control charts provide an intuitive way of measuring variance of internet service metrics. Using control charts, we can immediately observe the mean and range of the internet service metrics delivered to a customer. Improving internet service metrics average and tightening CLs will lead to improved reliability and overall customer experience by minimizing bandwidth margins. More importantly, operators that leverage SPC will be able to deliver services ready for the next generation of “always on” users and compete to win and retain customers.

6. Operationalization

The effort to operationalize these techniques at Shaw are in progress. The control charts shown in the previous sections are rendered from historical samples of Shaw production CM data and visualized using Tableau. The post processing engine is written in python and the logic can be incorporated into current RF and service metric performance indicators. Control charts can be rendered on the current data to evaluate the service stability of Shaw’s internet services using current data.

An effort is in place to develop the ability to initiate and collect CM US and DS speed test data on a regular scheduled interval so that we may directly measure customer experience using a more statistically relevant data set. Shaw also believes that increasing the data acquisition interval of CM RF data will improve the visibility of stability fluctuations in RF metrics further helping to improve CM reliability and performance.

These insights will drive new processes to improve service stability.

7. Future Work

The results from the operationalization of these techniques will be covered in a future paper. Prioritization and resolution of service issues will differ when we start to leverage SPC as seen in section 5.

8. Conclusion

Advancements in technologies like wearables, smart home, smart city, and smart industry are increasing the extent to which customers rely on connectivity for their livelihoods and well-being. It is expected that these customers will prioritize service quality and reliability over upload and download speeds, and as such, operators will be differentiated by their abilities to provide quality connectivity experiences. Detecting and controlling special-cause variances in networks and network devices are important steps in improving service stability and reliability.



Figure 18 - Speed to Service Stability Readiness

Control charts are well suited to observing and monitoring service stability, and control charts capturing internet service metrics are particularly interesting. While RF control charts can be useful in identifying problem areas for CMs, Internet service capabilities such as upload speed, download speed, latency, jitter, and bandwidth determine the competitiveness of the product offered. As seen in Section 5 of this paper, each of these internet service metrics can be captured within a control chart to be measured and validated to determine quality of service.

Operators that choose to leverage SPC will gain access to immediate eye-opening insights on how to truly measure and validate the quality of the services provided to their customers, thereby improving the overall customer experience, and preparing their networks for technologies of the future.

Abbreviations

CL	control limit
CM	cable modem
CPE	customer premises equipment
DS	downstream
LCL	lower control limit
OFDM	orthogonal frequency-division multiplexing
RF	radio frequency
SC-QAM	single carrier quadrature amplitude modulation
SNMP	simple network management protocol
SNR	signal to noise
SPC	statistical process control
UCL	upper control limit
US	upstream

Bibliography & References

Economic Control of Quality of Manufactured Product, Walter A. Shewhart

Shaw's Technology Forward Future

A Technical Paper prepared for SCTE by

Clarke Stevens
Principal Architect
Shaw Communications
2400 32 Ave NE, Calgary, AB T2E 6T4
clarke.stevens@sjrb.ca

Goutam Agarwal
Principal Enterprise Architect
Shaw Communications
2400 32 Ave NE, Calgary, AB T2E 6T4
goutam.agarwal@sjrb.ca

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Industry Perspective.....	4
3. Framework Approach at Shaw.....	7
3.1. Shaw's Emerging Technology Strategy Introduction.....	7
3.2. Network Modernization.....	8
3.2.1. Data Throughput.....	8
3.2.2. Data Latency.....	8
3.2.3. Data Density.....	8
3.2.4. Preparing for the Future.....	9
3.3. Data Ecosystem.....	9
3.4. Decision Intelligence.....	10
3.5. Autonomous Action.....	10
3.6. Total Experience.....	10
3.7. Trust.....	11
4. How to Operationalize?.....	12
5. Conclusion.....	15
Abbreviations.....	16
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1 - Current Business Environment.....	3
Figure 2 - Emerging Technology Research Process.....	4
Figure 3 - Industry Perspective - Emerging Technology Themes.....	5
Figure 4 - Shaw's Emerging Technology Focus Areas.....	7
Figure 5 - Emerging Technology Deep Dive.....	11
Figure 6 - Emerging Technology Operationalization Approach.....	12
Figure 7 – Hypothetical Shared Responsibility Model.....	13

1. Introduction

Over time, the rate of advancement in new technology tends to accelerate. As technology companies, we are continually faced with the challenge of not only keeping up, but also keeping ahead of the technology wave so we can use it to drive success. To do this, we need to support a continuous flow of new ideas into our companies and vet these ideas to determine relevance and assign priority. The ideas that rise to the top need to be quickly integrated without exceeding budgetary constraints, overwhelming employees, and negatively impacting customers.

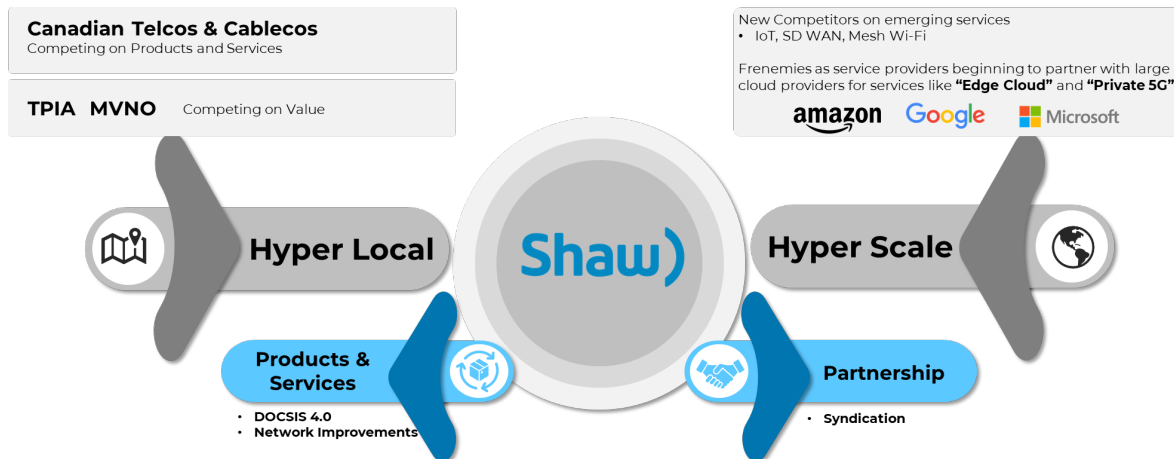


Figure 1 - Current Business Environment

Service providers face competition from two primary sources. The first source is direct competition from other service providers in the same primary business and geography. This hyperlocal competition is evident through price competition and feature competition. The price competition tends towards selling certain popular products and services and little or no profit in the hope of enticing customers to switch providers. Such offers are clearly not sustainable in the long run, so the bilateral churn leads to commoditized pricing.

The second major competition is from *Hyperscalers*. These are large multinational companies that can afford to compete on price because they are pursuing markets that are not primary sources of revenue. In addition, they often negotiate preferential pricing from vendors due to their scale. This competitive landscape severely constrains the competitive options for service provider.

The competitive landscape may require us to radically change our approach towards technology as a whole and adopt emerging technologies intelligently. If we can accomplish this, there are numerous benefits. We can anticipate rather than react to market changes. We can increase internal efficiency and improve customer products to benefit the customer experience. And as an additional bonus, we can provide employees with more interesting and valuable work, which consistently establishes a reputation for an interesting work environment and attracts new talent. Then the cycle repeats.

In this paper, we will discuss how we are trying to leverage and improve the intelligent adoption of new technology and this virtuous cycle at Shaw Communications.

2. Industry Perspective

Cable operators work in a competitive environment and must respond to global and industry-specific challenges such as supply chain issues due to COVID 19, geopolitical events, the rise and encroachment of *Hyperscalers* in the telecom domain, evolving customer demands and expectations, and cybersecurity and privacy issues. As these challenges increase, so too does the pressure to commercialize our core products and services. To meet these challenges, cable operators need to find sustainable advantages.

Our Emerging Technology Strategy is formed in combination of self-reflection (our strengths and areas of opportunities) and researching where the industry is heading (through peers, technology pioneers, and academics). As part of our research, we found multiple points of view that helped us identify our areas of focus.

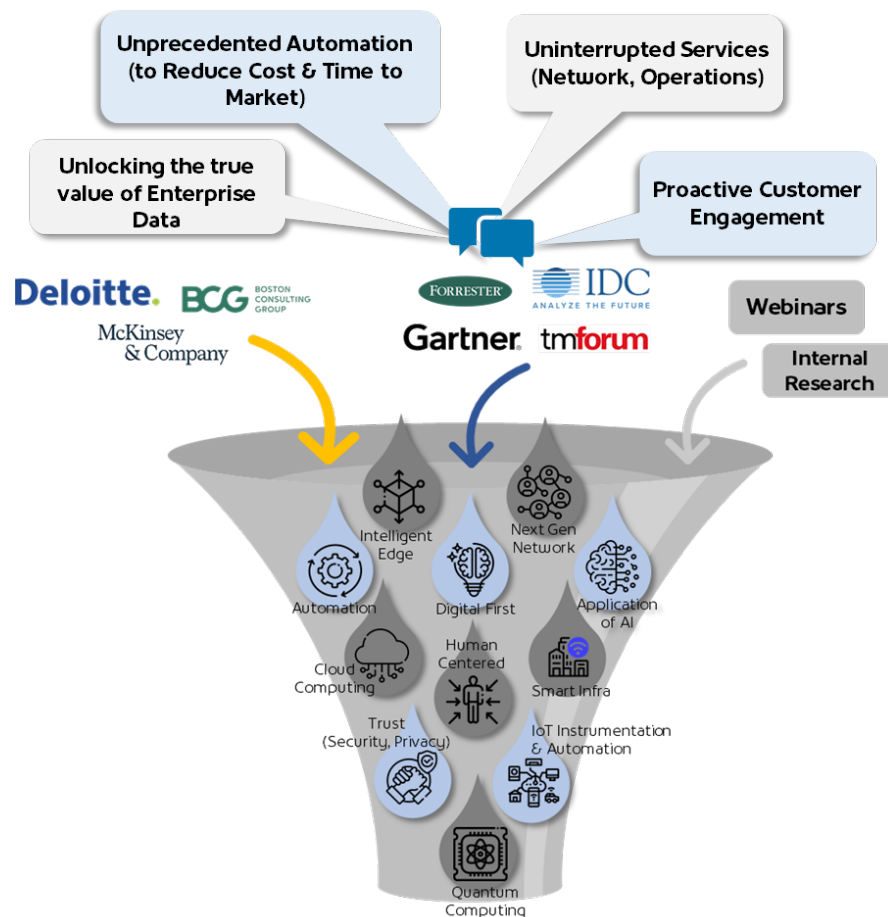


Figure 2 - Emerging Technology Research Process

Gartner's Technology trends 2021 [1] and 2022 [2] helped us narrow our research down to the following nine technology themes that operate together to build and reinforce each other.

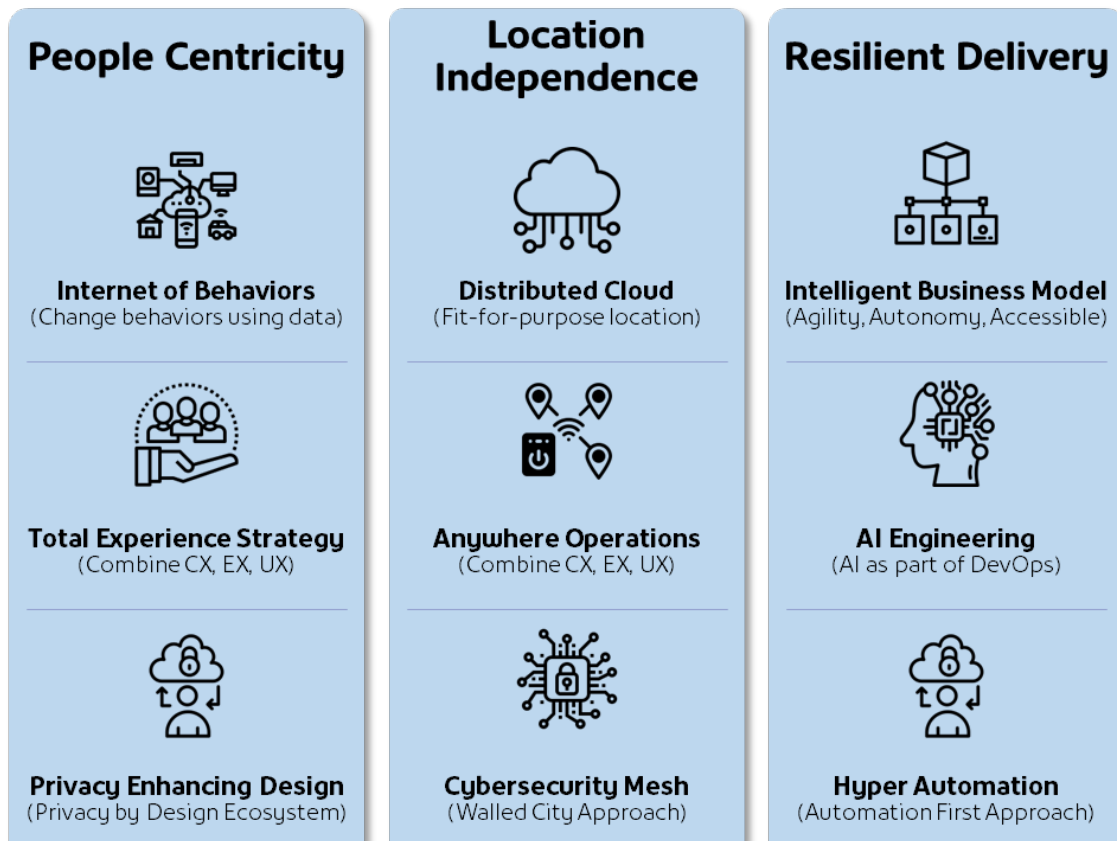


Figure 3 - Industry Perspective - Emerging Technology Themes

- **People Centricity:**
 - **Internet of Behaviors** – Internet of Behaviors is taking Internet of Things (IoT) to the next level. It refers to the ability to use the data collected to determine patterns and influence customer decisions.
 - **Total Experience Strategy** – Traditionally, experience strategies have been very siloed. The concept of *Total Experience* combines these disciplines (customer experience, employee experience, agent experience, and user experience) and links them to create a better overall experience for all parties.
 - **Privacy Enhancing Design** – Privacy Enhancing Design refers to creating an ecosystem of trust that enables organizations to safely store, share, process, and protect customer and employee data that is growing at an exponential rate.
- **Location Independence:**
 - **Distributed Cloud** – Distributed Cloud refers to providing an efficient multi-cloud environment or an “invisible cloud” environment. Essentially, users should be responsible for the maintenance, operation, and evolution of their services and not worry about physical infrastructure. The right cloud (public, private, or hybrid) backbone will

power users' needs around latency, performance, resiliency, and data sovereignty perspective.

- **Anywhere Operations** – Anywhere Operations refers to an operating model designed to support end users anywhere and everywhere and manage the deployment of business services across a robust distributed infrastructure. The model for Anywhere Operations is “digital first, remote first”.
- **Cybersecurity Mesh** – The Cybersecurity Mesh is a distributed architectural approach to scalable, flexible, and reliable cybersecurity control. It enables any person or device to securely access products and services, no matter the location, while providing an optimum level of security.
- **Resilient Delivery:**
 - **Intelligent Business Model** – Intelligent Business Model refers to agility and resiliency in the business process that can withstand disruptions. Business models and process should provide the following traits: agility, autonomy, and accessibility.
 - **AI Engineering** – AI Engineering refers to combining the best of both the worlds—human and software engineering—using a human-centered design powered by the performance and scale of software engineering. A robust AI engineering strategy will facilitate performance, scalability, interpretability, and reliability to ensure it can be used in day-to-day life and reap the full value of the investments.
 - **Hyper Automation** – Hyper-automation or an automation first approach is a practice in which organizations automate as many processes as possible using modern tools. For example, using automation to hyperscale agility and time to market.

Starting from these themes, we researched how large enterprises are responding to trends and quickly concluded that the competitive landscape requires us to radically change our approach towards technology as a whole and adopt emerging technologies intelligently into our day-to-day business operations.

3. Framework Approach at Shaw

The following sections describe a framework for innovation that leverages existing core infrastructure and builds upon that foundation.

3.1. Shaw's Emerging Technology Strategy Introduction

Shaw's Emerging Technology strategy and existing areas of focus suggest that we should leverage the core strength of our network as a foundation, while building technology services to provide a competitive advantage to differentiate us from our competition. Shaw's approach to emerging technology is based on the following building blocks:




Figure 4 - Shaw's Emerging Technology Focus Areas

- ✓ **Network Modernization** – Construction of networks that optimize speed, latency, and density.
- ✓ **Data Ecosystem** – Collection of more data from more sources more often.
- ✓ **Decision Intelligence** – Installation of tools to interpret the collected information both in real-time and offline using big data principles.
- ✓ **Autonomous Action** – Enabling the ability to act on the interpreted data for ultimate optimization.

- ✓ **Total Experience** – Creation of radically simplified and differentiated products using these tools.
- ✓ **Trust** – Achievement of the above with security and privacy as absolute requirements designed in from the start to build trust.

3.2. Network Modernization

Over the past several years, Shaw has pursued an aggressive plan to modernize our networks. Our networks are not only the basis of our business, but over time—as communication networks have increased in speed, latency, availability, and reliability—they have become essential infrastructure in the world economy. They are the basis of business for many of our customers. The COVID-19 pandemic proved this beyond any doubt as people became more reliant on networks to communicate and get business done when in-person activities were restricted.

The pressures of the pandemic have begun to ease, but many of the network activities that were initially undertaken to address pandemic related concerns continue to operate on the network now by choice. These activities were generally expected to increase operational efficiencies and reliability on the network until the pandemic drove the need for fast improvements.

The way forward involves improving all aspects of the network. This includes data throughput (how long it takes to send a certain amount of data), latency (how long it takes a specific piece of data to travel from point A to B), and density (how many streams of data can be moving on the same network at the same time).

Let us look at each of these features individually.

3.2.1. Data Throughput

Traditionally, this is the measurement that is commonly referred to as speed. If you run a speed test on your network connection, you will get a download and upload number in megabits-per-second (Mbps). This figure is an average performed over several seconds and does not provide a lot of detail about how long it took an individual message to traverse the distance. If we think of the network path as a freeway, throughput is the capacity of the freeway in general; it does not consider fast lanes or on- and off-ramps.

3.2.2. Data Latency

Data latency measures the time it takes an individual message to travel from point A to B. In many cases, this may not be critical. When loading a web page, for example, you will likely be more interested in loading the entire page rather than any particular word on the page. If, however, you are relying on the communication of a particular command to your self-driving car or the remote surgery to remove an appendix, timing is critical. Advanced networks can reduce latency and even potentially make some guarantees to meet time-critical requirements.

3.2.3. Data Density


A dense data network can run more simultaneous network paths in the same amount of space. This feature is useful in locations like stadiums, where thousands of people may be trying to access the network at the same time. Networks that allow more simultaneous networking streams will enable more users to be on the network at the same time in the same amount of space.

3.2.4. Preparing for the Future

The networking features identified above are some of the most common features in modern networks such as 5G wireless and advanced wireline networks.

Fortunately, Shaw's emphasis on modernizing our networking infrastructure sufficiently prepared us in advance of the pandemic. Maintaining and expanding our network during a pandemic was challenging, but our preparation kept the networks viable when they were under extreme pressure.

We were not alone in this, however. The networks of market competitors also performed well.



Source: Speedtest Intelligence Q1 2022 [3]

Preparing our network is essential to our success, but it is not sufficient. Networking will always be under commoditization pressure, and as service providers, it is important that we move up the stack.

3.3. Data Ecosystem

While basic networking is always improving, most emerging technologies of particular interest to cable operators sit above the basic Open Systems Interconnection (OSI) stack. The enabling networks are assumed to be available and able to meet requirements of complex applications. An increasingly common thread of these emerging technologies is a reliance on data—a lot of data.

“Big data,” as it is commonly identified, is data that is so abundant that it is difficult to comprehend without extensive processing by computers. Big data processing can identify subtle correlations and anomalies in a data set that would be nearly impossible to notice otherwise. These identifying features of large data sets can be used to extract important insights. Computers, however, are not always aware of the

relevance of data features to human priorities. Data scientists are needed to decide how to process the data and what information to look for.

Often, important insights are not clear when the data is originally collected. Therefore, it is important to store the data for long periods of time. While the cost of data storage can be expensive, it has dropped dramatically in recent years and helped to make the use of big data more practical.

Simultaneously, the advent of the IoT has led to the instrumentation of everything imaginable. Part of this instrumentation includes the collection and networked transmission of the data for immediate processing and instantaneous interpretation, as well as data storage for later analysis.

3.4. Decision Intelligence

Automatic analysis of massive amounts of data can not only bring timely insight, but in many cases, can also be interpreted with enough reliability that it can be used to make immediate decisions. Artificial intelligence (AI) can be used to make both analytical and intuitive decisions. Machine learning (ML), as a form of AI, uses massive amounts of data to draw correlations between observations and actions. This can be leveraged by training an ML algorithm with data that is clearly tied to specific outcomes. Additionally, as more data is collected, it can be incorporated into the training and used to improve accurate analysis over time.

In fact, some correlations are particularly predictive and can allow AI algorithms to notice potential problems before they become actual problems. For example, telltale vibrations coming from a piece of machinery can indicate that a part needs to be replaced soon. The part can then be replaced at a convenient time without bringing the machine down during production or risk causing further damage.

3.5. Autonomous Action

Autonomous action is the ability to act on interpretation. It refers to self-managing and self-healing systems that use the insights of and provide feedback to the decision intelligence systems, so that it can learn from its surroundings. This helps form a closed loop environment that focuses on end-to-end services and not just segments. The concept of autonomous actions offers a unique opportunity to provide highly resilient and repeatable services at scale. Allowing our people to be more successful and our customers to be more satisfied with their services and ensuring we get the most out of our resources and technology investments.

Organizations cannot scale manually at the rate that is required to meet the business needs and customer expectations. Here, relying on governed and monitored automation can help us to not only keep up with the pace of change, but also deliver cost effective, reliable, and productive solutions (e.g., automatically handle simple problems, notification and actions). Human intervention can be reserved for more complex problems.

3.6. Total Experience

In today's world, organizations do not independently decide how customer experience (CX), or employee experience (EX) works. It is best determined by the end-users. Total Experience that unifies the end-user experience could not only help control costs and provide sustainable growth, but also build long-term relationships with our customers and employees and improve brand trust.

- ✓ **Customer Experience** – Customer experience aims to provide an experience where we know our customers and are delivering a simplified, personalized, relevant experience to them through the channel of their choice.
- ✓ **Employee Experience** – Employers can use talent and culture as a competitive differentiator. We strongly believe in investing in people first, and that our talent is scarcer than other resources. Studies have shown a positive correlation between talent experience and organizational performance. That is, happy employees respond to customer needs more quickly and adapt to business changes faster.

3.7. Trust

Companies cannot expect to retain customers in the long-term without trust. Building and maintaining trust is of paramount importance to an organization as customer behaviors change and trusted products increasingly include claims like “locally-sourced” and “sustainable”. For service providers, taking a pragmatic and proactive approach towards trust (i.e., privacy, security, and data protection) is not an option, but a necessity.

The concept of building trust with our customers is not a new idea, but with increases in digital working environments and amounts of data being collected and processed—and the emphasis on and complexity of building trust—it has increased manyfold. As the “new normal” and actions of organizations come under intense scrutiny, companies’ actions have greater consequences and influence customer opinions. These actions will determine if existing customers choose to continue their relationships with a brand and whether new customers will choose to start them. With so much at stake, it is important to rethink our approach, skills, processes, and technologies that support a Trust Ecosystem.

Shaw undertook a current state assessment to understand where we stand on these areas of focus and to identify opportunities and mechanisms to operationalize our Emerging Technology Strategy.

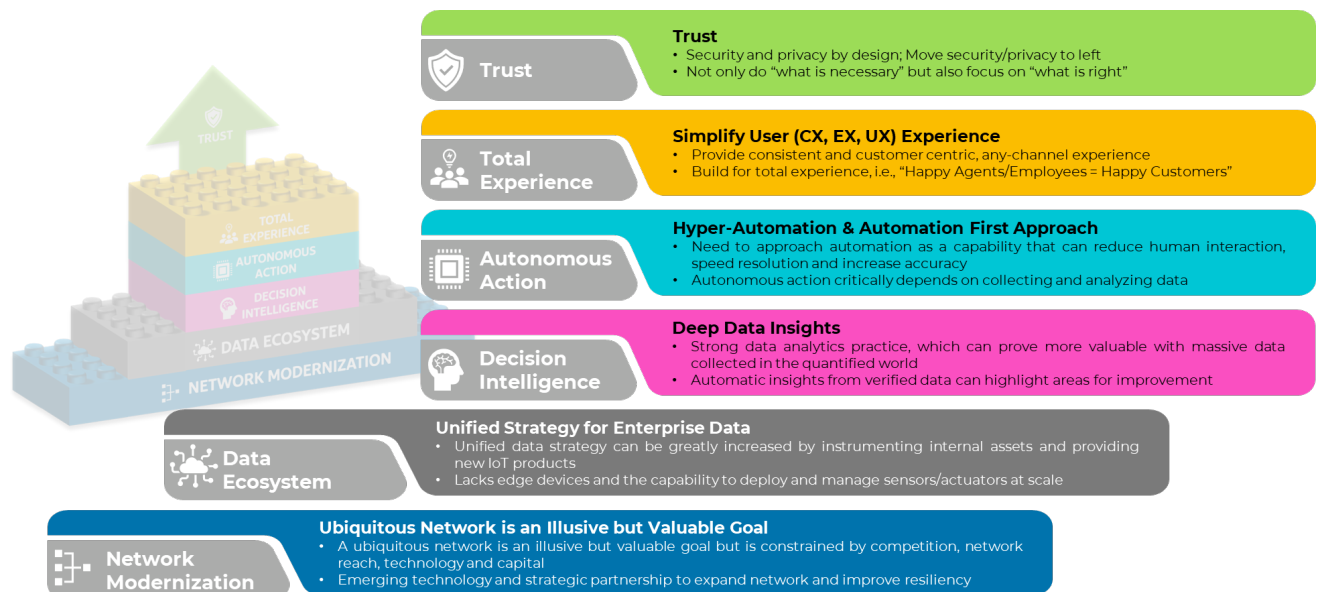


Figure 5 - Emerging Technology Deep Dive

4. How to Operationalize?

In this paper, we have laid out a logical structure for integrating emerging technologies into a cable operator's business model and suggested the benefits of doing so. However, to operationalize this structure and successfully integrate new technologies into the business, it is important to consider current operations. While it may be possible and at times necessary to swap out technologies completely and quickly (for example, after a security breach), a well-considered transition that minimizes current operational disruption would produce better, more sustainable results.

It is important to have a long-term strategy. When a new technology is introduced, there are serious practical considerations. Why is the change necessary? How much will it cost? What are the expected cost savings or revenue increases? How will employees and customers react? A long-term strategy should directly address these and similar questions. Resistance to change will be reduced when the reasons for the change are clear.

Communication is also critical. Adoption of a new technology cannot be accomplished without the assistance of those who will be implementing it. They need to understand the motivation, but also the information to plan thoroughly to minimize downtime for employees and customers. A proof-of-concept (PoC) trial may be advisable to verify the details of how the technology will be installed and connect with existing systems. Furthermore, the deployment plan must be carefully created and executed. For example, it may make sense for the technology to be introduced using an already planned operational upgrade rather than creating new projects.

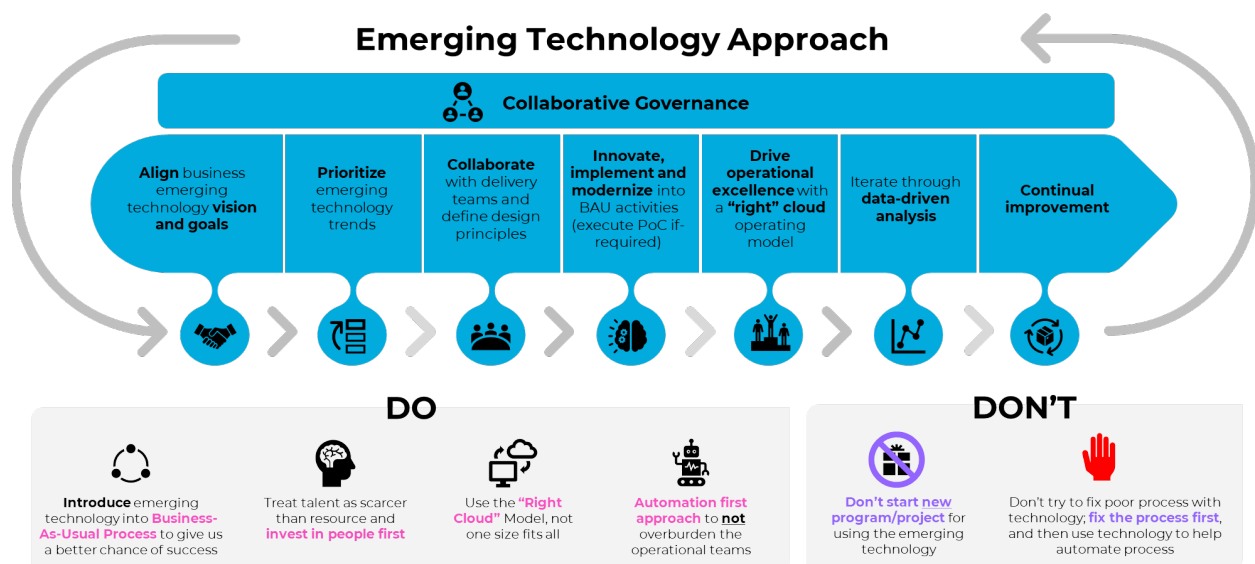















































Figure 6 - Emerging Technology Operationalization Approach

The introduction of the technology through the deployment plan should be clearly communicated to all impacted groups—and potentially even to groups who would not be impacted by the upgrade but would support if technical issues arise. When people are aware that a change is happening, regardless of how minimal it is expected to be, they can anticipate that change and plan appropriately.

It is also important to establish roles and responsibilities. Those impacted by the transition should know exactly who to ask when they have questions. Who is responsible for making decisions? Who has the information that will inform those decisions? Who needs to complete which tasks and on what timeline? Are there executives or other people who need to be informed of the progress and alerted when problems arise?

Operating Model Stage	Shared Responsibility Model			
	Responsible	Accountable	Consulted	Informed
 Collaborative Governance				
 Business Alignment			 	
 Prioritization			 	
 Planning & Design	 			
 Innovate, Implement & Modernize	 			
 Operational Excellence	  			
 Data Driven Analysis	 			
 Continual Improvement	   			

Internal Teams

-  Business Team
-  IT Team
-  Operations Team

External Teams


-  Technology Partners

Figure 7 – Hypothetical Shared Responsibility Model

For those who will be using the new technology, a thorough training plan is likely necessary. Training should be planned sufficiently in advance of the actual change so employees are as comfortable as possible during the transition.

With instrumentation and data collection intrinsic in modern business, all operational considerations can be handled more efficiently, and sometimes even automatically.

Referring to our framework, the network lies at the foundation. Information must be instantly and reliably communicated. This includes not only information between people, but also information between machines, sensors, actuators, and data systems. Data must be collected from every conceivable source, the verity of that data must be dependable and timely, and it must be reliably delivered to the appropriate destinations over the network. Where possible, the data should be used for accurate analysis and automated action. Finally, all these interactions need to be protected with strong security that can quickly react to breaches.

The fundamental operations of the business do not change. Businesses still create products and sell them to customers. For cable operators, we have both end customers who use our products to enrich their lives and business customers who use our products to support their own customers. We can better serve both types of customers when we operationalize new technologies into our business.

New products often rely on new technology and new capabilities, but many new technologies will now include instrumentation that will allow products to be monitored for health and usage. In some cases, problems can be diagnosed and fixed before the customer even notices them. Important data about how

the product is being used and how it can be improved can keep product teams informed on what improvements to consider and additional markets that will embrace the same product. The product development cycle does not fundamentally change; it becomes more efficient. In this new cycle, the product team can better meet the needs of customers, customers are more satisfied with the products, and customer care teams can more accurately diagnose problems or avoid them altogether since they can be diagnosed and repaired without human interaction.

Operationalizing emerging technology is a success-based cycle. With a careful integration plan, initial disruption can be minimized. As the benefits of the new technology are observed, enthusiasm for the technology increases for all involved parties, which encourages the adoption of more relevant technologies. As employees see that technologies are being adopted to solve problems rather than for technology's sake, they are more willing to trust management. As the value of emerging technologies is operationalized, employees at all levels have greater job satisfaction, which in turn attracts new employees in a very tight job market.

5. Conclusion

While not all emerging technologies are based on networking, data, and IoT, a very large number of technologies are improved by these features. The cable industry has been extremely successful over many decades because we observed early on that our networks, originally designed for broader distribution of live television, could be used for much more than that. Indeed, we find ourselves at the nexus of the information age with advanced, industry-leading, reliable networks that are becoming essential infrastructure to modern society.

As with any valuable resource, many companies want a share. The value of networks will always be subject to competitive forces that will push them toward commoditization. In order to overcome the gravity of commoditization and continue to operate as profitable businesses, cable operators need to continually look for new revenue opportunities.

At Shaw, we have observed that there will be an increasing need to measure, analyze and automatically react to information, regardless of the source. Fortunately, we are well-positioned at the center of that revolution. There is a great opportunity now for cable operators to adopt fundamental technologies and create new lines of business in order to participate in new waves of opportunity in the future.

As cable operators, it is essential that we build the flexible infrastructure required for the information, intelligence, and automation that will ride on our networks. We need integration, deployment, maintenance, and automation systems that will scale to accelerating consumer demand. This will provide new revenue opportunities in the future and set us up for the next information revolution—one that even we have yet to imagine.

Abbreviations

AI	artificial intelligence
BAU	business as usual
BCG	Boston Consulting Group
COVID	Coronavirus disease
CX	customer experience
DOCSIS	Data Over Cable Service Interface Specifications
EX	employee experience
IDC	International Data Corporation
IoT	Internet of Things
IT	information technology
ML	machine learning
OSI	Open Systems Interconnection
PoC	proof of concept
UX	user experience

Bibliography & References

- [1] <https://www.gartner.com/smarterwithgartner/gartner-top-strategic-technology-trends-for-2021>
- [2] <https://www.gartner.com/en/information-technology/insights/top-technology-trends>
- [3] <https://www.speedtest.net/global-index/canada#fixed>

Smart Amplifiers – Are They Worth It?

A Technical Paper prepared for SCTE by

Mike Darling
Principal Engineer
Shaw Communications
2728 Hopewell Place NE, Calgary AB T1Y 7J7
mike.darling@sjrb.ca

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Background	4
2.1. Hybrid Fibre Coax Networks	4
2.2. Amplifier Functionality	4
2.3. Amplifier Block Diagram.....	6
3. Smart Amplifiers	7
3.1. What Constitutes a Smart Amplifier?	8
3.2. Smart Setup	8
3.3. Communication	8
3.4. Monitoring.....	8
3.5. Control.....	9
3.6. Automation	9
4. Operational Considerations.....	10
4.1. Amplifier Operations.....	10
4.1.1. Initial Setup	10
4.1.2. Maintenance.....	10
4.1.3. Troubleshooting	11
4.2. Shaw Statistics	12
4.3. Potential Value of Smart Amplifiers	17
4.3.1. Node Size Reduction	17
4.3.2. Proactive Network Maintenance	17
4.3.3. Full Band Capture	18
4.3.4. Gateway Architecture.....	18
5. Smart Amplifier Systems	19
5.1. Operational Shift	19
5.2. Maintenance.....	19
5.3. Additional Powering Requirement.....	19
6. Strategic Value	19
6.1. Smart Amplifiers in the Path to DOCSIS 4.0.....	20
7. Financial Analysis.....	20
7.1. Costs Avoided	20
7.2. Smart Amplifier Costs	20
7.3. Net Present Value and Sensitivity Analysis	21
7.3.1. Costs Avoided	21
7.3.2. Maintenance Costs	21
7.3.3. Incremental Amplifier Cost.....	22
7.4. Context.....	22
8. Conclusion.....	23
Abbreviations	23
Bibliography & References.....	24

List of Figures

Title	Page Number
Figure 1 – Hybrid Fibre Coax Topology	4
Figure 2 – Cable Attenuation with Frequency – CommScope QR540 [2]	5
Figure 3 – Insertion Loss with Frequency – ATX GigaXtend XST-24-20 [3]	5

Figure 4 – Amplifier Output Levels.....	6
Figure 5 – Amplifier Cascade.....	6
Figure 6 – Simplified Amplifier Block Diagram.....	7
Figure 7 – Downstream Troubleshooting	11
Figure 8 – Upstream Troubleshooting	12
Figure 9 – Keyword Occurrence in Ticket Creation	13
Figure 10 – Keyword Occurrence in Ticket Completion	14
Figure 11 – In-Home vs Outside Plant Tickets	14
Figure 12 – Outside Plant Referral Hierarchy	15
Figure 13 – Outside Plant Referral Tickets	15
Figure 14 – Outside Plant Incident Hierarchy	16
Figure 15 – Outside Plant Incident Tickets	17
Figure 16 – In-home Networks.....	18
Figure 17 – Costs Avoided Sensitivity	21
Figure 18 – Maintenance Cost Sensitivity.....	22
Figure 19 – Incremental Amplifier Cost Sensitivity	22

List of Tables

Title	Page Number
Table 1 – DS to US Ratio of Different Diplex Splits	9
Table 2 – Cost Avoidance Assumptions	20

1. Introduction

As the newest generation of the Data over Cable Service Interface Specification (DOCSIS) ecosystem takes shape, multiple system operators (MSOs) are contemplating upgrading their hybrid fibre coax (HFC) networks. Whether an operator chooses the full duplex (FDX) or frequency division duplex (FDD) variant of DOCSIS 4.0, all amplifiers in the network will need to be replaced by newer versions. Historically, amplifiers were simple devices set up during install and only revisited during routine maintenance or troubleshooting. However, there is an opportunity to enhance the functionality of these devices with the aim of decreasing the total cost of ownership (TCO) of the HFC network and to reduce the incidence and duration of customer impacting outages. This paper will present a fresh perspective on smart amplifier functionality and evaluate the potential benefits of deployment.

2. Background

2.1. Hybrid Fibre Coax Networks

In HFC networks, signals are combined in a hub site and transmitted over fibre optic cable to an optical node in the field. An analog optical node performs an optical to radio frequency (RF) transition and sends signals onto coax cables. Coax cable design follows a tree-and-branch topology, which was developed prior to the introduction of fibre optics when the function of the network was to distribute analog television channels [1]. Signals propagate through a cascade of amplifiers separated by coax cable and passives such that the gain of the amplifiers cancels out the loss of the cable and passives. This high-level topology is shown in Figure 1.




Figure 1 – Hybrid Fibre Coax Topology

HFC networks are designed to deliver signals to customer premises equipment (CPE) roughly at equal power across frequency to optimize signal quality. A complicating factor is that coax cable and passive devices have higher loss at higher frequencies. To correct for this effect, amplifier outputs are tilted so that signals at higher frequencies have higher output power than signals at lower frequencies.

2.2. Amplifier Functionality

Coaxial cable attenuation increases with the square root of frequency, as shown for CommScope QR540 cable in Figure 2.




Figure 2 – Cable Attenuation with Frequency – CommScope QR540 [2]

Similarly, passives such as taps, splitters, and couplers have higher insertion loss at higher frequencies, as shown in Figure 3 for an ATX GigaXtend XST-24-20 model tap.




Figure 3 – Insertion Loss with Frequency – ATX GigaXtend XST-24-20 [3]

To optimize signal quality while compensating for frequency specific losses, amplifiers accept an input signal, condition it to be as flat across frequency as possible, amplify the signal, and then tilt the output. This allows a signal to be transmitted through a cascade of amplifiers and ultimately to homes with optimal quality. HFC networks are generally designed so that each amplifier has the same output levels, signified by the power level in the lowest and highest downstream channels. For instance, RF output levels can be 35/49 dBmV at 54/1000 MHz, as shown in Figure 4.




Figure 4 – Amplifier Output Levels

Amplifier signals are designed to stay within an optimal range of power level, such that they are not too close to the noise floor on the low end, or in danger of distortion at the high end. When designed and implemented correctly, signals can pass through a long cascade of amplifiers with little degradation in performance. Figure 5 shows the power level of a signal over frequency as it passes through a cascade of amplifiers and a tap to the customer premises.




Figure 5 – Amplifier Cascade

2.3. Amplifier Block Diagram

Signals pass through many discrete steps inside an amplifier. A simplified amplifier block diagram is shown in Figure 6.




Figure 6 – Simplified Amplifier Block Diagram

Starting from the left, a downstream signal enters the amplifier and is directed along the top path by a diplex filter, signified as a box with an H (high frequency) and L (low frequency). An equalizer removes any frequency tilt remaining in the signal after traversing the cable and passives in the previous network segment, and an attenuator or pad ensures the signal level is in the optimal range for the preamplifier gain stage. Interstage attenuation controls the output level, while slope control sets the output tilt. The signal then goes through an additional diplex filter to be reunited with the upstream signal and exits the amplifier.

In the upstream, the signal flow is simpler, as losses at low frequency are smaller, allowing for a single gain stage. Starting from the right, an upstream signal enters the amplifier and is directed to the upstream gain stage, after which the signal is equalized, attenuated, and reunited with the downstream signal. Attenuation and equalization are generally accomplished using plug-in components such as attenuators and equalizers, which can be varied to achieve the desired levels.

An additional function not shown in Figure 6 is temperature compensation, commonly included in amplifiers as the attenuation of components such as cable increase with higher temperatures. Without temperature compensation, amplifiers that are set up on a hot day may amplify signals beyond specified levels on a cold day, and amplifiers set up on a cold day may not amplify signals enough as the temperature rises. Signal level changes due to temperature swings are more significant in the downstream because the attenuation is greater. Downstream temperature compensation systems use a feedback control loop, which attempts to keep a specific portion of the signal at a specified power level. If the level drops the system increases the gain, and if the level increases the system lowers the gain. In the upstream, a thermal attenuator can provide some compensation without the need for a control loop.

Amplifiers have historically been relatively simple devices that are set up during HFC network construction and only visited during routine maintenance or troubleshooting. However, smart enhancement can provide opportunities for greater network reliability and improve the customer experience. Now that amplifier functionality has been examined, smart enhancement will be discussed in following section.

3. Smart Amplifiers

Many amplifiers come with smart features that aim to improve the functionality of these devices. But what exactly is a smart amplifier, and what benefits do they bring to operators?

3.1. What Constitutes a Smart Amplifier?

The term smart amplifier is often used in two main contexts. The first is smart functionality that is local to the device, such as electronically selectable attenuation and equalization settings. It is expected that all 1.8 GHz amplifiers will have these features. The second is the ability to remotely communicate with the amplifier. This paper briefly discusses the first context but focuses on the second, as communications capability is expected to be optional, and MSOs will have to choose whether to make this investment.

3.2. Smart Setup

As mentioned earlier, amplifiers have generally used discrete plug-in components to vary settings such as attenuation and equalization. This requires HFC technicians to carry a stock of these components for setup and maintenance. In addition to added expense, this creates potential logistical issues. It is possible for a technician to *not* have the correct value component and be forced to use the closest value in their possession, resulting in a non-optimal setup. Plug-in components are a possible source of failure, as are their connections to the amplifier board. Even small failure rates of these components can cause an operational headache, and intermittent failures can be especially difficult to troubleshoot. For example, an attenuator that fails at high temperature may be tracked to an amplifier, only to have the problem cease when the amplifier is opened and the internal heat dissipated. Additionally, the mechanical joint between the component and the amplifier board can create unexpected problems when an incomplete connection is made. Wiggling or reseating the component may fix the problem, but a technician may end up discarding the component and installing a new one, fearing the problem will return.

An amplifier can be designed so that some or all configurable features are electronically selectable, either through a wired or wireless connection to a hand-held device. There are benefits to this type of system beyond avoiding issues with plug-in components. A greater range of potential settings and a larger number of incremental steps can be designed, allowing for a more precise setup. Additionally, authorization to make changes to amplifier setup can be more easily controlled, either by access to the required hardware or through restricted access via software.

3.3. Communication

While amplifiers with electronically configurable components can be referred to as smart amplifiers, a fundamental capability for this category of amplifier is the ability to communicate. Communication can be achieved either through a proprietary signaling protocol that uses upstream and downstream frequencies separate from end-user signals, or in-band through DOCSIS. There are pros and cons to both methods. Proprietary systems can use inexpensive transponders, especially if the quantity of data sent back and forth is small, but they require exclusive use of some amount of spectrum. DOCSIS communication can be more costly but has the benefit of sharing spectrum with other services, meaning that bandwidth is only used when amplifiers are sending or receiving data. In both scenarios, communication enables monitoring and control of smart amplifiers, however it adds both additional cost and complexity.

3.4. Monitoring

There are many different data points that can be monitored in a smart amplifier. These include:

- Configuration settings
- Power status
- Lid status (open/shut)

- Input/Output RF levels
- Diplex filter frequencies
- Temperature compensation

Input and output RF level monitoring requires additional hardware to evaluate signal strength. This can be done at specific frequencies or across the band with spectrum analysis functionality.

3.5. Control

It is possible to implement different levels of control in a smart amplifier. The simplest is remote setup via electronically configurable components. Combined with the ability to monitor input and output levels, a standard network setup can be achieved by a technician in an operations center, rather than sending a field technician to perform the work locally. More complex systems that control diplex frequency based on an evaluation of traffic requirements can also be designed.

DOCSIS 4.0 specifications allow for several diplex frequencies between 204 and 684 MHz [4]. A diplex frequency at 204 MHz would accommodate more downstream and less upstream traffic, while a diplex frequency at 684 MHz would accommodate more upstream and less downstream traffic. The ability to remotely change the diplex frequency would enable each node to be optimized for its traffic demand. For example, a node with many business customers might have a more symmetric traffic pattern when compared to a residential node. In the latter case, a diplex frequency of 300 MHz might be optimal, while in the former case, 684 MHz would likely best match the traffic pattern. Upstream and downstream spectrum allocation and associated ratios for each diplex frequency is found in Table 1.

Table 1 – DS to US Ratio of Different Diplex Splits

US Diplex Frequency (MHz)	US Spectrum (MHz)	DS Spectrum (MHz)	DS/US Spectrum Ratio
204	199	1536	7.7
300	295	1422	4.8
396	391	1302	3.3
492	487	1188	2.4
684	679	960	1.4

This table assumes worst-case downstream lower band edges and uses spectrum, not capacity, in calculating downstream to upstream ratio. Since upstream modulation rates tend to be lower than downstream modulation rates, the downstream to upstream capacity ratio would be slightly higher.

3.6. Automation

Electronically configurable attenuation and equalization, along with monitored signal levels, would allow for automated setup in smart amplifiers. This setup could be achieved locally at the amplifier through an application or push-button functionality. This would require logic to be implemented locally to the amplifier or in a hand-held device. Alternatively, if the device has communication capabilities, this logic could exist in a back-office system where the device is automatically configured to a predefined state when communication is first established after install.

4. Operational Considerations

While additional levels of intelligence in amplifiers can be beneficial, it is important to recognize that this intelligence will also change the operational model for setup, maintenance, and troubleshooting of amplifiers, with the goal of lowering lifecycle cost.

4.1. Amplifier Operations

In the absence of smart amplifier communication, all activities concerning those amplifiers require a technician to be physically present at the device. Adding smart features changes the way that HFC networks are set up and maintained.

4.1.1. Initial Setup

During the setup of new HFC networks, each amplifier is installed and configured by field technicians in the order of its position in the cascade. This is because the configuration of all amplifiers upstream will impact signal levels downstream. Consequently, amplifier setup happens sequentially, and a single technician may start at the first amplifier and physically move down the cascade, installing and setting up each device. However, if the amplifier is capable of remote configuration and signal level monitoring, amplifier setup can be completed by staff in an operations center. This would allow field technicians to install the amplifiers in the most efficient manner and configuration to happen remotely, without the requirement of being present at each device.

4.1.2. Maintenance

Routine maintenance of amplifiers is accomplished through sweep programs, where a group of technicians visit each amplifier in a node and confirm that signal levels are appropriate. Signal levels may change over time owing to different causes such as equipment aging, temperature swings, and deliberate changes that solve localized problems but create others. For example, a technician may increase signal level at a drop by raising the output signal level at the preceding amplifier, solving the immediate problem but potentially causing issues further downstream in the cascade.

Sweep programs are time-consuming as technicians must drive to each amplifier location, access the amplifier, open it up, connect it to a field meter and confirm whether setup is appropriate, and make any changes necessary. In the case of aerial plant, a bucket truck is required to access devices, whereas in underground plant, vaults and pedestals need to be accessed. Access can also be restricted in some extenuating circumstances, such as where amplifiers in buildings require special access, or when underground infrastructure is flooded. In the worst cases, access can require days or weeks of advance notice and fees to be paid. In addition, HFC spectrum has expanded and overlapped with mobile spectrum, creating another source of interference—through both mobile signals to the HFC network and HFC signals to the mobile network—and opening amplifier housings during a sweep program can exacerbate this issue. The extent to which MSOs utilize a sweep program can vary significantly, with some enacting rules about how often an amplifier is visited over time by a sweep program, while others may avoid them completely.

If a smart amplifier can monitor RF levels and be configured remotely, then a sweep program would be unnecessary or performed from the operations center, removing the requirements for field travel and amplifier access, ultimately creating operational efficiencies. However, the extent to which this eliminates effort depends on how a sweep is implemented at the operator level.

4.1.3. Troubleshooting

Amplifiers are potential points of failure as well as locations to subdivide the network during troubleshooting. Troubleshooting downstream issues can be achieved by working back from where the issue is reported. For example, if a modem is not receiving sufficient signal strength, a technician may visit the tap. If levels are poor at the tap, the technician will move upstream to the amplifier. If the amplifier has poor output levels, the technician can continue upstream through the amplifier cascade, ultimately to the node or hub site, to find the source of the problem. In the downstream, modem statistics can also be used to localize an issue. If there is only one modem suffering from low signal level, the problem is likely in-home or at the drop. If many modems are experiencing the issue, the common point of failure can be determined by an operations center, and technicians dispatched directly to the problem.




Figure 7 – Downstream Troubleshooting

Troubleshooting noise in the upstream cannot be accomplished by correlating modem statistics, as all upstream signals terminate at the same location in the hub site or at the node in distributed access architecture (DAA). In this case, technicians use a brute-force technique, working out from the node and looking at each output leg for noise. The noise is followed until it no longer presents, at which point it is determined that the noise is entering the network between that location and the last place it was observed.




Figure 8 – Upstream Troubleshooting

Chasing intermittent noise can be difficult, time-consuming, and frustrating. It is not uncommon for a technician to be dispatched to troubleshoot upstream noise, only to have it disappear. The noise may return shortly after the technician has moved on to other activities, or it may not return for days, weeks, or at all.

The ability of smart amplifiers to aid in finding noise remotely is very useful. A remote system makes use of a spectrum analyzer in the hub site as well as the ability to temporarily increase the attenuation on each leg of a multiple-port amplifier one-by-one. If the noise presenting at the spectrum analyzer decreases when a specific amplifier leg is attenuated, it can be deduced that the noise is coming from that leg. This uses the same brute-force method used by field technicians, but eliminates the time required to drive to each location. It also allows a technician in the operations center to quickly address the issue when it is present in the case of intermittency.

To support the ability to remotely chase upstream noise, operators could also choose to only place modems or transponders in multiple-output amplifiers where noise would be entering the device from a single leg. The extent to which this strategy would work depends on how often the HFC network makes use of passive splitters instead of multiple-output amplifiers to split the network. Even in the case of splitter use, putting a modem or transponder in the first single-output amplifier downstream of a splitter leg would allow full functionality of a noise chasing system, but this could create logistical challenges as it can be difficult to ensure that the transponder or modem is put into the correct amplifiers.

4.2. Shaw Statistics

To estimate the amount of effort that could be eliminated with the implementation of smart amplifiers, it is important to have good statistics on technicians' activities. Ticketing systems for field technicians are used to assign work but are often not designed for statistical analysis. For example, drop-down boxes with a subset of potential problem categories and solutions often do not allow the flexibility to properly describe their activities, while free-form entries will vary in description from person to person and are

difficult to analyze. Despite these issues, this data remains the best source of intelligence on how technicians use their time and where smart amplifiers can be most impactful.

Field technicians are often grouped by the function they provide, or by the area of the network they focus on. Some possible delineations include in-home, service, maintenance, construction, and plant. This section will only focus on troubleshooting activities and will exclude installs and disconnects.

Service calls are generated when a customer calls in with an issue that cannot be resolved over the phone. Customer service representatives create a ticket, enter a brief description of the issue from the customer's point of view, and include details from tools they have access to. The top 10 keywords are shown in Figure 9 as a percentage of use in ticket creation.




Figure 9 – Keyword Occurrence in Ticket Creation

Keyword analysis can give a high-level view of the issues that customers are experiencing but is inherently problematic as these words can be used in many contexts. As an example, the keyword “RF” can exist in the description “Poor RF Signal Levels” and “RF Signal Levels OK”. Despite this, it can be observed that a common description from a customer perspective is the behaviour of the hardware or service. Descriptions such as “Modem won’t turn on” or “Internet down” are common. Tickets are then assigned to field technicians who will investigate the issue at the customer premises. After investigating and resolving the issue, the ticket is closed, and an additional description is included in completion notes. The top 10 keywords used in ticket completion notes are shown in Figure 10.




Figure 10 – Keyword Occurrence in Ticket Completion

The ticket completion notes provide insight into what the solution to the issue was. From the use of the keywords “swap”, “install”, and “replaced”, it can be inferred that changing out CPE is the most common solution to customer issues. The keywords “splitter” and “connector” are likely indicators of issues within the in-home wiring. The use of “tap” and “drop” point to issues outside of the home. Service or in-home technicians generally do not troubleshoot beyond the tap, so if an issue is traced to a cause beyond the tap, a ticket is opened for a maintenance or outside plant (OSP) technician to continue troubleshooting.

Outside plant tickets can also be created by the operations center when alarms occur or when customer issues are correlated to a single source. The number of in-home tickets is much larger than the number of outside plant tickets, as shown in Figure 11.




Figure 11 – In-Home vs Outside Plant Tickets

This figure represents the absolute number of tickets, and not time spent or cost incurred. In general, outside plant tickets take longer to resolve than in-home tickets and are more costly.

Outside plant referrals can be organized in a hierarchy using categorization data to estimate the number of tickets where work could be aided by smart amplifiers. Figure 12 shows the categories in the hierarchy pertinent to the discussion.

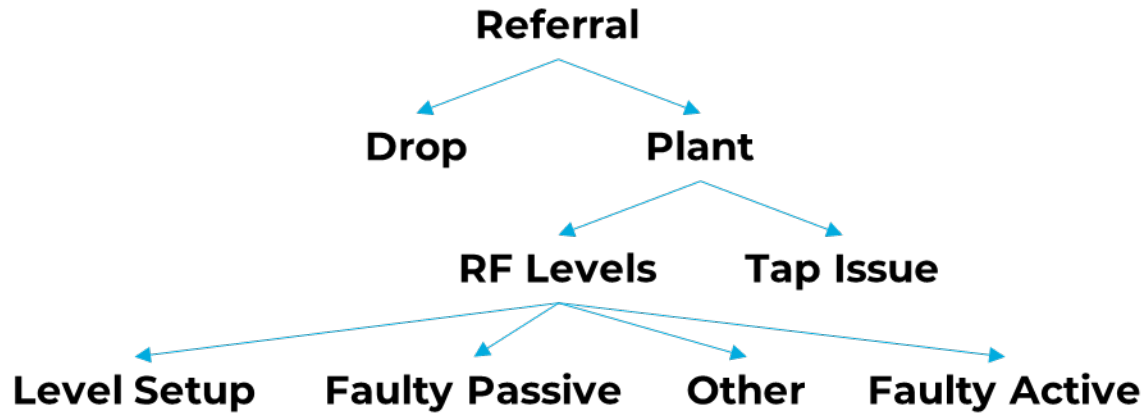


Figure 12 – Outside Plant Referral Hierarchy

One-third of plant referrals are for drop replacements while two-thirds involve issues at the tap or further into the network. Figure 13 is a Sankey diagram that visualizes the many-to-many relationship between categorizations within outside plant tickets.




Figure 13 – Outside Plant Referral Tickets

The “Other” category above includes issues such as fibre-to-the-home (FTTH) tickets, logistical requests (e.g., for installing a larger pedestal), requests from the planning department, and no fault found scenarios.

While there are many different causes of plant issues, the majority will require checking RF levels at multiple locations. Smart amplifiers could be of help in this case, both by allowing remote RF level observation to find problems and for RF levels to be changed remotely. The degree to which a smart amplifier system would aid in solving the problem depends on where the specific issue is found. In the case of a damaged tap, the solution ultimately requires the tap to be changed out by a technician, but a smart amplifier system could localize the problem to a specific plant segment, saving troubleshooting time. Systems that correlate modem levels can also be used to determine the location of the problem, potentially to a greater degree owing to the larger number of modems.

A plant incident ticket hierarchy is shown in Figure 14. Only categories pertinent to upstream noise are shown.




Figure 14 – Outside Plant Incident Hierarchy

Approximately one-third of incident tickets are for downstream outages while two-thirds are for upstream noise, which can be caused by a variety of different issues. Much of the ingress from homes is caused by telco noise, which is created when customers are connected to the HFC network and using their in-home coax wiring to pass signals using HFC upstream frequencies. The solution is to disconnect their drops at the tap, reconnecting them only if they become HFC customers in the future. Noise also enters the HFC network through weaknesses such as loose connectors and damaged cables. Again, a Sankey diagram can be used to see the many-to-many relationship between incident ticket categorizations in Figure 15.




Figure 15 – Outside Plant Incident Tickets

As noted earlier, the clearest benefit of a smart amplifier system is in chasing upstream noise, as this cannot be achieved through correlating modem levels or any other systems due to the noise funneling effect. Noise tickets are estimated to be approximately 3% of all tickets involving HFC in-home and outside plant troubleshooting. While this is a small number, it must be stated that noise chasing is among the most challenging and time-consuming activities that field technicians deal with.

4.3. Potential Value of Smart Amplifiers

Smart amplifiers can provide benefits by reducing required components such as attenuators and equalizers, saving time for field technicians, and reducing outage times for customers. The scale of the potential opportunity has changed over time owing to multiple factors that will be discussed below.

4.3.1. Node Size Reduction

The process of increasing capacity per home passed by building fibre deeper into the HFC network has reduced node sizes, both in terms of homes passed and the number of amplifiers. This has the effect of reducing the failure domain, increasing the uptime for all customers, and reducing the physical scope of troubleshooting. A node that has 1,000 homes passed and 100 amplifiers inherently has more potential sources of failure and is more difficult to troubleshoot than a node with 250 homes passed and fifteen amplifiers. Node sizes have decreased substantially in the past decades and will continue to do so going forward in response to increased traffic demands.

4.3.2. Proactive Network Maintenance

Proactive network maintenance (PNM) programs use performance statistics from modems to detect and fix HFC network problems before they cause customer impacting events [5]. Upstream equalizer settings can be correlated to find and fix impedance mismatches caused by weaknesses like cracked cables and loose connectors. A properly executed PNM program can ensure that the HFC network is in good shape, reducing issues in the network and direct impacts to customers.

4.3.3. Full Band Capture

Modems with full band capture (FBC) provide remote access to spectrum analyzer functionality. This allows for remote troubleshooting of several potential issues. Downstream troubleshooting, as described in 4.1.3, can be accomplished by using modems instead of taking readings at amplifiers.

4.3.4. Gateway Architecture

Legacy in-home wiring can consist of several splitters connecting a number of video, data, and phone CPE. In place of this arrangement, many operators are moving to a gateway architecture, where a single DOCSIS gateway connects to the HFC network and all services are accessed from this device. These two scenarios are shown in Figure 16.




Figure 16 – In-home Networks

Moving to a gateway architecture has two main benefits. The first is that the signal level required to support a single CPE per home is lesser when compared with legacy architectures that involve splitting and will therefore suffer additional loss. This puts less pressure on the HFC network to provide sufficient signal level to support multiple boxes and reduces the need for drop amplifiers, which can amplify upstream noise in addition to desired signals. Secondly, a smaller number of endpoints reduces the potential for noise to enter the HFC network.

Smart amplifiers would have been more useful in years past when HFC networks were larger and more difficult to troubleshoot. Their value may have increased, however, due to the shift to work-from-home

and hybrid-work situations driven by the COVID-19 pandemic, which has made customers less tolerant of unplanned outages.

5. Smart Amplifier Systems

Smart amplifiers that are remotely monitored and controlled require software to facilitate operation. There is potentially a large amount of data to be polled from amplifiers, and the number of amplifiers in an HFC network is large, requiring any system to be designed for usability. As of today, there are no industry standards, which means that systems are proprietary and only cover certain amplifier models. HFC networks tend to have amplifiers from multiple manufacturers for either historic reasons or in order to maintain multiple suppliers from a strategic sourcing perspective. The Society of Cable Telecommunication Engineers (SCTE) Smart Amplifier project, which launched last year, aims to create standards that the industry can use to develop smart amplifiers for DOCSIS 4.0.

5.1. Operational Shift

While smart amplifiers can reduce the workload for field technicians, it will create more work for technicians in operations centers. Amplifier setup and maintenance, and HFC technical work in general, has a large craft component and field technicians will often learn best practices over time with senior staff. This body of knowledge needs to be transferred to operations centers, as the ability to change the configuration of amplifiers in the field remotely should not be accessible to the uninitiated.

5.2. Maintenance

Making amplifiers smart adds new maintenance requirements and potential failure modes. Modems or transponders can go offline, requiring a technician to visit the amplifier to reset them. Power supply monitoring modems often suffer from this issue, but when they go offline the network continues to function. Because there is no direct customer impact to these types of modems going offline, it can be a low priority send a technician to bring them back online. This leaves operational teams with an incomplete picture of the health of the network. There is a significantly greater number of amplifiers than power supplies in the network, and as such the problem has the potential to be a magnitude larger, underscoring the need for smart amplifier components to be extremely reliable.

5.3. Additional Powering Requirement

The addition of a transponder or DOCSIS modem to each amplifier will create incremental powering requirements that will both increase operational costs due to the increased power usage and drive the need for additional power supplies. While this requirement is small in the overall context of powering the HFC network, it adds to other demands expected over the next few years. Additional power demands are expected to come from the transition to DAA, the use of small cells powered and backhauled using the HFC network and moving to 1.8 GHz. In addition to the added power requirement, a DOCSIS modem or transponder will add more heat that needs to be dissipated in an amplifier housing—something that may be challenging with 1.8 GHz amplifiers.

6. Strategic Value

HFC maintenance programs are often difficult to justify on financial arguments alone, especially as metrics such as “truck rolls avoided” are theoretical and subject to many different factors. Smart amplifiers also have strategic value that is not necessarily reflected in a financial analysis.

6.1. Smart Amplifiers in the Path to DOCSIS 4.0

Increasingly, HFC networks are competing against FTTH networks. In Canada, telcos have built out FTTH networks to most of their homes passed, and in the United States, telcos are not far behind. Current HFC networks can compete with gigabit passive optical network (GPON) downstream speeds, and high-split HFC networks can enable upstream speeds competitive with GPON. To compete with Ten Gigabit Symmetric Passive Optical Network (XGS-PON), HFC networks will have to be upgraded to DOCSIS 4.0, which enables 10 Gbps downstream and 6 Gbps upstream [6].

As customers look beyond speed when selecting a service provider, features such as latency and reliability become more important. While DOCSIS 4.0 will come with improvements in these areas, smart amplifiers have the potential to further improve HFC networks by reducing the incidence and duration of outages, making the overall network more reliable.

7. Financial Analysis

A financial analysis can be undertaken using data from ticketing systems and assuming high-level costs for smart amplifiers and field technicians' activities. This analysis should be considered directional as there are many uncertainties using both ticketing data and high-level cost assumptions.

7.1. Costs Avoided

The activity that is most affected by the introduction of smart amplifiers is upstream noise chasing. Smart amplifiers would allow noise to be located to a specific section of plant, at which point a field technician could be dispatched much closer to the source of ingress than would be the case otherwise. It can be assumed that 75% of upstream troubleshooting effort is eliminated by smart amplifiers, and a sensitivity analysis can be performed to gauge how that assumption impacts the net present value (NPV). While there is the potential to eliminate downstream troubleshooting efforts and prevent outside plant issues from causing needless in-home service calls, most of this functionality is available through PNM or similar modem performance analysis tools. Because of this, the assumption is made that smart amplifiers would not eliminate any effort for in-home or downstream issues. The estimated percentages for each category of activity are shown in Table 2.

Table 2 – Cost Avoidance Assumptions

Category	Percent Avoided
In-home	0%
OSP – DS Issues	0%
OSP – US Noise	75%

Assuming \$300 for outside plant activities and \$100 for in-home activities, the costs potentially avoided on a per-home passed, per-year basis is calculated as \$0.75.

7.2. Smart Amplifier Costs

If an incremental cost of \$50 per smart amplifier is assumed and there are 20 homes passed per amplifier, this leads to a figure of \$2.50 per home passed. Yearly costs of a smart amplifier system in terms of maintenance and licensing are more difficult to estimate as they will vary with operator and solution. A cost of \$0.50 per home passed per year with no capital start-up costs is assumed. An operator could alternatively develop a system in-house, in which case the yearly operational costs would be lower but the start-up costs would be potentially higher.

7.3. Net Present Value and Sensitivity Analysis

With the assumed costs and savings spread over a ten-year amplifier lifespan, the NPV of a smart amplifier system comes in at approximately \$0. This means that for a positive NPV, either the costs must be lower than assumed, or the benefits be greater. A simple sensitivity analysis shows how the NPV is impacted by a change in the input variables.

7.3.1. Costs Avoided

Assuming other inputs remain the same, Figure 17 shows how the NPV is impacted by the costs avoided per home passed.




Figure 17 – Costs Avoided Sensitivity

A savings of \$0.75 per home passed per year is a conservative estimate, which does not consider cost savings due to foregoing a sweep program, or due to the ability to remotely troubleshoot amplifiers that are difficult to access locally.

7.3.2. Maintenance Costs

Yearly maintenance costs, which include software licensing fees and costs related to system upkeep, would have to be higher than \$0.50 per home passed per year before the system NPV would be negative, as shown in Figure 18.




Figure 18 – Maintenance Cost Sensitivity

Maintenance costs are where companies offering smart amplifier back-office systems recoup their investment in development. Any systems that are over-engineered may risk becoming uneconomic and therefore unappealing to MSOs. Operators with in-house development capacity may be able to design and build their own systems once smart amplifier standards become available.

7.3.3. Incremental Amplifier Cost

The incremental cost of a smart amplifier would have to increase to over \$50 per amplifier to turn the NPV negative.




Figure 19 – Incremental Amplifier Cost Sensitivity

An effort has been made to use conservative costs and savings in this high-level financial analysis. The directional conclusion is that the value of a smart amplifier system under the assumptions used above is neutral, and a smart amplifier system would have to be carefully designed and properly executed to ensure a return on investment.

7.4. Context

It has been reported by Jeff Baumgartner at Light Reading that the costs to upgrade to DOCSIS 4.0 may reach \$250-\$400 per home passed [7]. The NPV of smart amplifiers should be thought of in this context,

and whether the ultimate NPV is slightly positive or slightly negative, it represents a very small fraction of expected HFC network expenditure.

8. Conclusion

Smart amplifiers can reduce HFC network TCO and increase reliability by lowering the number of network outages and reducing the time required to troubleshoot issues. While PNM systems can provide this functionality in the downstream using modem-level data, locating the source of upstream noise is more challenging. This represents the greatest value for smart amplifiers. Analyzing data from operational ticketing systems can help in estimating the scope of potential savings. Using this analysis, a simple financial model shows that costs saved by reducing upstream troubleshooting effort pays for the investment in smart amplifiers. This must be taken in the context of DOCSIS 4.0 deployments, which will allow MSOs to compete against FTTH competitors. Any smart amplifier development should be designed carefully to avoid delaying equipment availability and made optional to allow MSOs to choose whether to make the investment.

Abbreviations

CPE	customer premises equipment
DAA	distributed access architecture
dB	decibel
dBmV	decibels relative to a millivolt
DOCSIS	Data Over Cable Service Interface Specification
FBC	full band capture
FDD	frequency division duplex
FDX	full duplex
FTTH	fibre-to-the-home
GHz	gigahertz
GPON	gigabit passive optical network
HFC	hybrid fibre coax
HP	home passed
MHz	megahertz
MSO	multiple system operator
NPV	net present value
OSP	outside plant
PNM	proactive network maintenance
RF	radio frequency
SCTE	Society of Cable Telecommunication Engineers
TCO	total cost of ownership
XGS-PON	ten gigabit symmetric PON

Bibliography & References

- [1] *Modern Cable Television Communications-Video, Voice and Data Communications*, 2nd Ed. – Walter Ciciora, James Farmer, David Large and Michael Adams; Morgan Kaufmann, 2004
- [2] QR540 JCA Specifications (<https://www.commscope.com/globalassets/digizuite/114515-p360-5506302-external.pdf>), CommScope 2002
- [3] GigaXtend XS Series (https://www2.atxnetworks.com/l/73612/2019-09-11/18g134/73612/1632518328KCdJzE6s/ANW1403_GigaXtend_2GHz_XS_Taps.pdf), ATX 2022
- [4] CableLabs – Data-over-Cable-Service Interface Specifications DOCSIS 4.0 Physical Layer Specification
- [5] <https://www.cablelabs.com/technologies/proactive-network-maintenance>
- [6] Maximizing Returns on the Path to D4.0 – Mike Darling, Shaw Communications, SCTE Fall Technical Forum, October 2021
- [7] DOCSIS 4.0 network upgrades could reach \$300 per homes passed - analyst – Jeff Baumgartner, Lightreading.com ([https://www.lightreading.com/cable-tech/docsis-40-network-upgrades-could-reach-\\$300-per-home-passed---analyst-/d/d-id/777724](https://www.lightreading.com/cable-tech/docsis-40-network-upgrades-could-reach-$300-per-home-passed---analyst-/d/d-id/777724))

Solving Automation and Orchestration for Large-Scale Hybrid Cloud Environments

A Technical Paper prepared for SCTE by

David Grizzanti

Distinguished Engineer
Comcast Cable
1800 Arch St. Philadelphia, PA
215-356-2354
david_grizzanti@comcast.com

Matthew Morrissey

Principal Engineer
Comcast Cable
1800 Arch St. Philadelphia, PA
215-435-7648
matthew_morrissey@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Problem Statement.....	3
2.1. Operating in a hybrid cloud environment	3
2.2. Simplifying application developer choice	4
2.3. Benefits of time series metrics over logs	4
3. Building the platform	5
3.1. Standardized toolset over cloud native	5
3.2. Choosing the open source path	6
3.3. Centralized versus distributed.....	6
4. Building modular components.....	7
4.1. Terraform Modules	7
4.2. Container Images	8
4.3. Image Builds.....	8
4.4. Client Configuration.....	8
4.5. Client Onboarding	8
5. Deployment & Continuous Delivery	8
5.1. Pipeline Generation.....	9
5.2. Concourse Interface.....	9
5.3. Delivery Automation	10
6. Conclusion.....	10
Abbreviations	11
Bibliography & References.....	12

List of Figures

Title	Page Number
Figure 1 - Observability Toolchain	3
Figure 2 - Time Series Data Graph.....	5
Figure 3 – Centralized Query Layer.....	7
Figure 4 – Concourse User Interface.....	9
Figure 5 - Deployment Orchestration.....	10

1. Introduction

The idea of using automation to reduce human error and speed up infrastructure deployments is not new, however, choosing the “right” tools can be challenging. Tool selection has, in some ways, become an engineering challenge on its own. Many large cable operators are not all in on a single cloud provider, automation framework, or container orchestration engine.

The authors will review a set of cloud agnostic infrastructure-as-code (IaC) modules, continuous deployment pipelines, and a GitOps/Configuration Management interface that facilitates changes. This framework can deliver an internal observability platform that hosts millions of time-series metrics, including visualization and a unified search interface.

2. Problem Statement

As cable operators expand workloads into public cloud providers, teams will be tasked with providing a solution to time series metrics collection that is cloud agnostic. Many cable operators already use some form of on-premises infrastructure and the expansion to public cloud provides an opportunity to re-think how monitoring infrastructure is built. One primary goal is to give clients a storage and query interface to their metric data, without having to manage the system on their own. This will in turn enable client teams to focus on the things that bring the most value to their customers. This platform must be robust enough to handle outage scenarios such that the clients can access their time series data to perform troubleshooting or triage for their application.

2.1. Operating in a hybrid cloud environment

A best practice for time series platforms is to collect and store samples as close to the source application as possible. Given that this team was building a platform for other teams within the company, they would need to be able to deploy wherever the clients are deployed. This meant that the platform would need to deploy into both public and private clouds. Deploying into a hybrid cloud environment can get complicated quickly if you want to maintain a consistent way to deploy across all clouds. Each cloud has its own APIs and preferred methods for interacting with the infrastructure that you’re building. In addition, service offerings such as load balancing, secret management, and auto-scaling are different across each cloud.

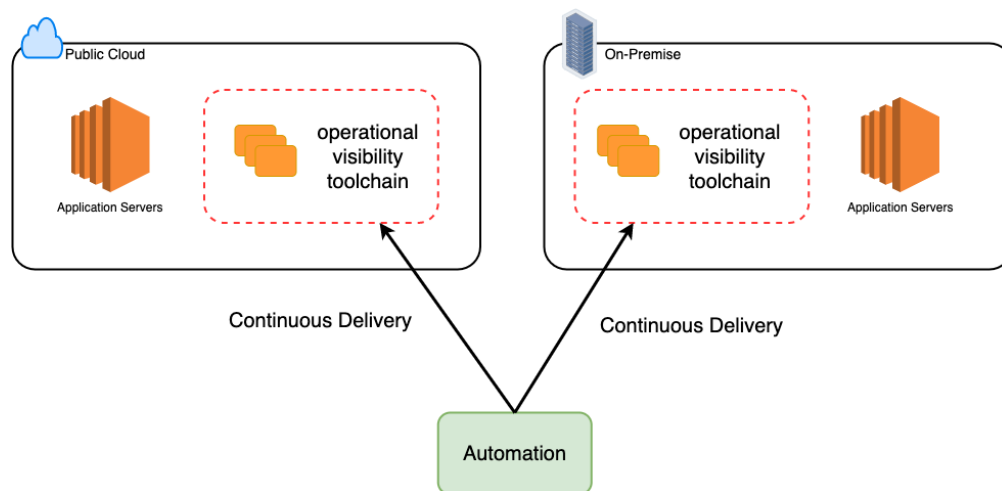


Figure 1 - Observability Toolchain

Choosing a set of abstractions can help with these challenges. The team chose to use Terraform and treat the Terraform configuration just as we would any other piece of code. This practice is commonly referred to as Infrastructure as Code. This enabled them to predictably and repeatably build infrastructure into any cloud that had a Terraform provider. This significantly decreased the complexity of deploying to multiple clouds. Once this pattern was established, they were effectively able to extract the functionality into a set of Terraform modules and fill in provider specific details.

2.2. Simplifying application developer choice

Large companies tend to have a huge breadth of choice in terms of tool sets and platforms application developers can use to deploy and monitor their applications. This can frequently lead to analysis paralysis from a developer's perspective since there can be so many options, some seemingly competing with one another. It can also be difficult as a developer to discover what tools or platforms are available to you since documentation for these tools or platforms tends to be siloed within a given team or organization. From the company's perspective this spread of similar tools and platforms doing the same or similar things, and their potential lack of discoverability, is challenging in terms of budget dollars. There could be lots of engineering hours spent developing a tool set or platform that might already exist. Teams or organizations might purchase a vendor solution for a problem that has already been solved in other parts of the company.

A goal is to avoid these pitfalls, making use of their own platform so easy that it would not even be presented as a choice. Users would be compelled to use the system because there would be little to no friction involved in getting it set up or interacting with it. This meant that not only did they need to spend a significant amount of time up front thinking about the user experience, but they also had to make sure that every decision made would not negatively impact the user experience if it could be avoided. For the team, that meant hiding as much of the complexity as they could from the user. To achieve that goal, they built a web portal for the clients to interact with that would guide them through the process of providing the required information to deploy the infrastructure into their chosen environment. Rather than require users to setup a meeting or fill out a ticket, they were given the ability to go to a web page and onboard in an entirely self-service manner.

2.3. Benefits of time series metrics over logs

Time series data is a collection of numeric measurements, made over time. There are many uses for time series data, plotting points on a graph being one of the most popular. Time series data exists in many domains outside of software systems, stock prices, heartbeats per minute, etc. For a web application, time series data may be requests per second, or for a database this might be number of active connections or active queries. [1]

Historically, applications relied heavily on application logs for observability. Application logs being text-based messages that contain informational events, warnings, and errors. Whether this was for application alerting, creating dashboards, or debugging the system, logs were the standard. The team was tasked with building a time series collection platform to reduce log ingestion and shift teams towards a "metrics first" approach. Most application teams would write verbose logs and extract metrics from those logs in downstream log analytics tools. As a result, there are often high log volumes and large storage costs associated with many applications.

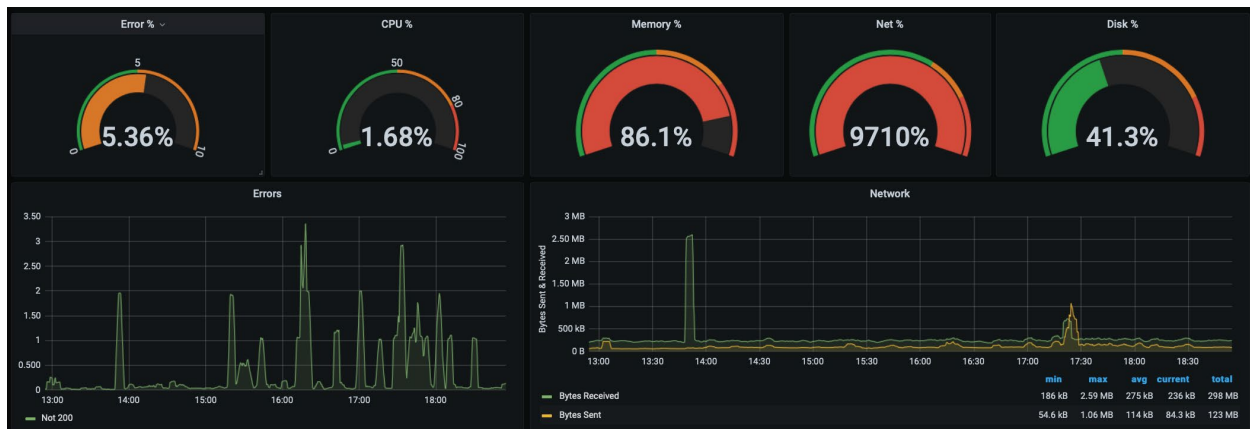


Figure 2 - Time Series Data Graph

“By and large, the biggest advantage of metrics-based monitoring over logs is that unlike log generation and storage, metrics transfer and storage have a constant overhead. Unlike logs, the cost of metrics doesn’t increase in lockstep with user traffic or any other system activity that could result in a sharp uptick in data.

With metrics, an increase in traffic to an application will not incur a significant increase in disk utilization, processing complexity, speed of visualization, and operational costs the way logs do. Metrics storage increases with more permutations of label values (e.g., when more hosts or containers are spun up, or when new services get added or when existing services get instrumented more), but client-side aggregation can ensure that metric traffic doesn’t increase proportionally with user traffic.” [2]

3. Building the platform

At the outset of the project, the team made the choice to build a metrics collection platform, instead of buying an off-the-shelf product. They also decided to avoid so called “cloud native” toolsets that existed within public cloud platforms. Each public cloud offers a suite of observability tools that integrate into their infrastructure services.

A primary goal of the project was to build an easily maintainable and scalable platform to deploy metrics infrastructure for application teams. The target for this infrastructure ranged from on-premises datacenters to public cloud providers, with the initial target being on public clouds. Many vendors offer platforms that provide metrics storage “as a Service”, however, the team’s research found the cost and compatibility with the various internal applications and tools incompatible. Many public cloud vendors offer their own hosted tools that provide these features; however, they wanted a single solution for all teams regardless of where they chose to host their application.

3.1. Standardized toolset over cloud native

The team chose Prometheus [1] as the time series database, primarily for its wide adoption as an open-source standard for monitoring and alerting. Using Prometheus as the base allowed them to focus their efforts on building out the platform around a common tool versus worrying about how to automate and support each cloud vendors specific product. In addition, this allowed them to offer a single solution regardless of an application teams intended deployment location.

In addition to choosing a standard tool to deploy across clouds, the team also decided to keep their deployment simple by using Virtual Machines (VM) instead of cloud specific orchestration platforms. Since each cloud platform offers its own unique orchestration service and on-premises datacenters are unique, choosing the lowest common denominator simplified the architecture. There were a few drawbacks to this, namely losing key features like auto-scaling, but they were able to work around these with proper alerting and health checks.

Moreover, this enabled the team to build upon a single common image that could be deployed to whichever cloud, be it public or private, with minimal effort to add cloud specific instruction sets for the VM. This gave the benefit of only needing to fetch a single image which has been patched or updated and build on top of it. These changes could then be easily propagated out to the client stacks without needing to remember or think about cloud VM specific patching.

3.2. Choosing the open source path

As mentioned earlier, Prometheus was selected as the base software tool for metrics collection and storage. The next focus was to choose how they wanted to manage and deploy Prometheus plus any customizations and client specific configuration that would be needed. They decided to continue with open-source, modular tools and made the following selections:

- **Terraform:** For Infrastructure as Code automation
- **Packer and Ansible:** Building Base Cloud AMIs/Images
- **Git:** Version control of Terraform modules and client configuration
- **Concourse:** Continuous Integration and Deployment

Using Terraform allowed them to take advantage of existing providers for many public clouds and, on-premises private clouds like OpenStack.

3.3. Centralized versus distributed

Another challenge they faced was whether the collection and storage of application metrics should be centralized versus distributed. More specifically, whether the collection and storage of an application's metrics should stay within the cloud and region where they originate. If they chose to distribute, then querying could be challenging, but if they chose to centralize, they'd have to deal with a centralized cost and chargeback model.

The team chose to keep collection and storage local to the application's virtual networking environment. This had a few advantages, namely:

- Prometheus uses a pull model for metrics collection, so this simplified network and access control as scrapers were local to an applications network.
- Cost was managed by deploying into a client's local network, so no chargeback or show back was necessary
- A network outage would not cause any loss in metrics collection

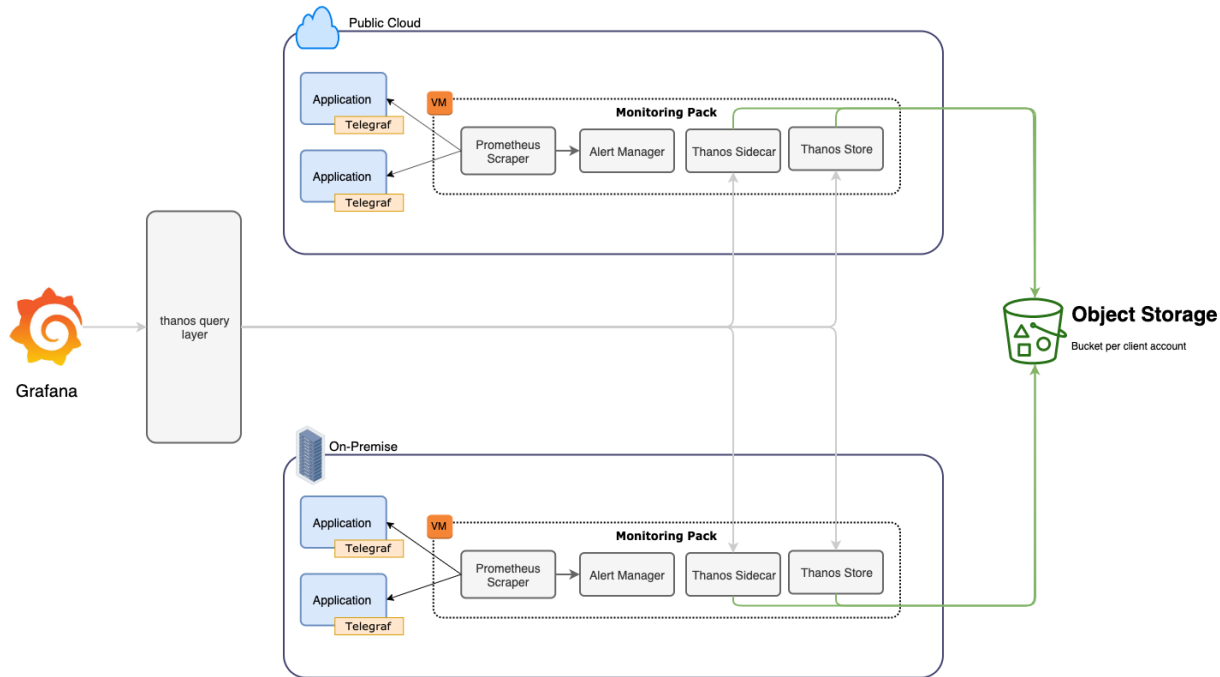


Figure 3 – Centralized Query Layer

A distributed model posed a challenge however for querying all metrics for a given client at once. This challenge was solved by deploying Thanos, an open-source Prometheus compatible tool with global querying capabilities as shown in Figure 3. The Thanos Query component, specifically, enabled deployment of a centralized set of instances that could access all regions, across all clouds and offer a single interface for client queries.

It was also decided to use the selected tool chain to deploy the monitoring as a client of infrastructure. This provided valuable insights into pain points the clients might experience, find places where there were gaps in the automation, and get comfortable with using the tools were being recommended that other teams adopt. This experience created a robust set of documentation, which in turn helped to reduce the support burden on the team.

4. Building modular components

Segmenting each concern into its own module allowed the platform to be deployable to any current or future cloud platform. The platform was broken down into Terraform modules, application container images, VM image builds, and client configuration.

4.1. Terraform Modules

The Terraform modules are primarily responsible for deploying a set of application container images, via docker, onto a virtual machine in the specified cloud environment. The logic for each cloud provider is encapsulated into each module, which includes details on instance types, disk sizes, cloud initialization instructions, etc. The set of container images can be customized per cloud and per client as needed.

4.2. Container Images

The application container images, which includes Prometheus, are the main software tools in the system. They provide the functionality for the metrics collection and storage. It's worth noting however, that the platform is agnostic to the container images deployed. Adding new container images or swapping all of them out for a different set of software tools would be trivial.

4.3. Image Builds

Since each cloud platform requires its own image type, cloud specific logic was extracted from the automation. A set of images were then built to be deployed across all clouds at once, with the same set of requirements driven by Ansible playbooks.

4.4. Client Configuration

Each client of the platform has a set of custom configurations that need to be applied at deployment time. These range from what hosts Prometheus needs to scrape, alerting rules, metric retention time, etc. This configuration is kept in version control (git) and updates to these files drive deployments to the various clouds and regions that are supported.

4.5. Client Onboarding

Clients are provided with a single URL at which they can onboard and manage their stacks. This was done to simplify how a client can start using the stack. The website has links to all the relevant documentation for any pre-requisite steps they might need to perform, as well as a simple form which takes in all the required information for the automation to build the stack for them. With a single click the client has their git directory and file structure created, has the minimal set of configurations for the containers to run, has the pipeline to deploy the stack created, and the stack is deployed. The client is presented with links to all the relevant sites and can immediately configure their stack for their use case. The whole process is entirely automated which allows the platform developers to focus on providing more value to the client and allows the client to more quickly get their stack running and configured since they don't need any touch points with the platform team.

5. Deployment & Continuous Delivery

The final step is integrating all the various components and modules into continuous delivery pipelines, using Concourse. This is where the core intelligence and logic of the platform comes together. At the outset of the project, they built each pipeline manually for each client by stringing together:

- The set of necessary Terraform modules for the target cloud platforms
- The logic for what regions were necessary
- The set of container images that matched the client configuration

This process became very tedious over time, especially when the number of clients grew beyond a handful and the number of regions for each client increased. To address to this, a templating application was created using Golang to automate the creation of each client pipeline.

5.1. Pipeline Generation

Concourse pipelines are written using YAML, which can be very verbose given its declarative nature. To simplify the management of the pipeline YAML files, which could range from 5,000 to 20,000 lines, it was decided to write an application that generate these files, given a defined set of inputs. The inputs included:

- Target cloud platform(s)
- Cloud region(s)
- Container image(s)
- Cloud platform account details

Given this set of inputs, a templated YAML file was created that was fed into the Golang program. When a client onboards or adds additional stacks to their deployment the program will re-generate the pipeline entirely rather than trying to insert just the pieces that change. This is possible because the pipeline generation is deterministic.

5.2. Concourse Interface

The generated YAML is used to create a pipeline within the Concourse platform as shown in Figure 4. This interface allows for manual triggering a region deployment if necessary, viewing logs, and tracking progress of deployments.

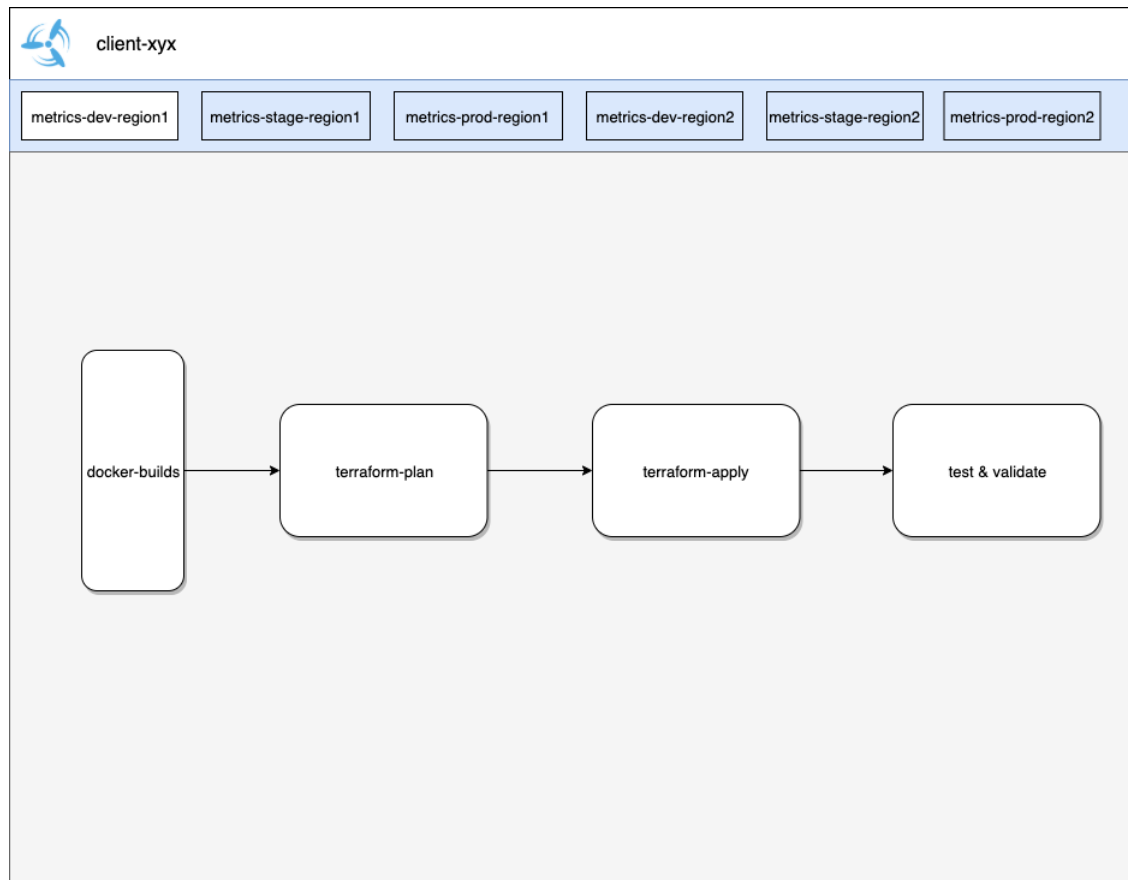


Figure 4 – Concourse User Interface

5.3. Delivery Automation

Once the Concourse pipeline is generated and applied in the Concourse platform, updates made to client configuration in git will trigger a deployment. The trigger is executed via webhooks through the version control system. A change made to a file in git will send a HTTP request to Concourse and kick off a deployment. The set of jobs run for each cloud do not need any manual intervention from the team, unless an error occurs. Any errors will trigger an alert to the team and error logs can be viewed within the Concourse user interface (UI).

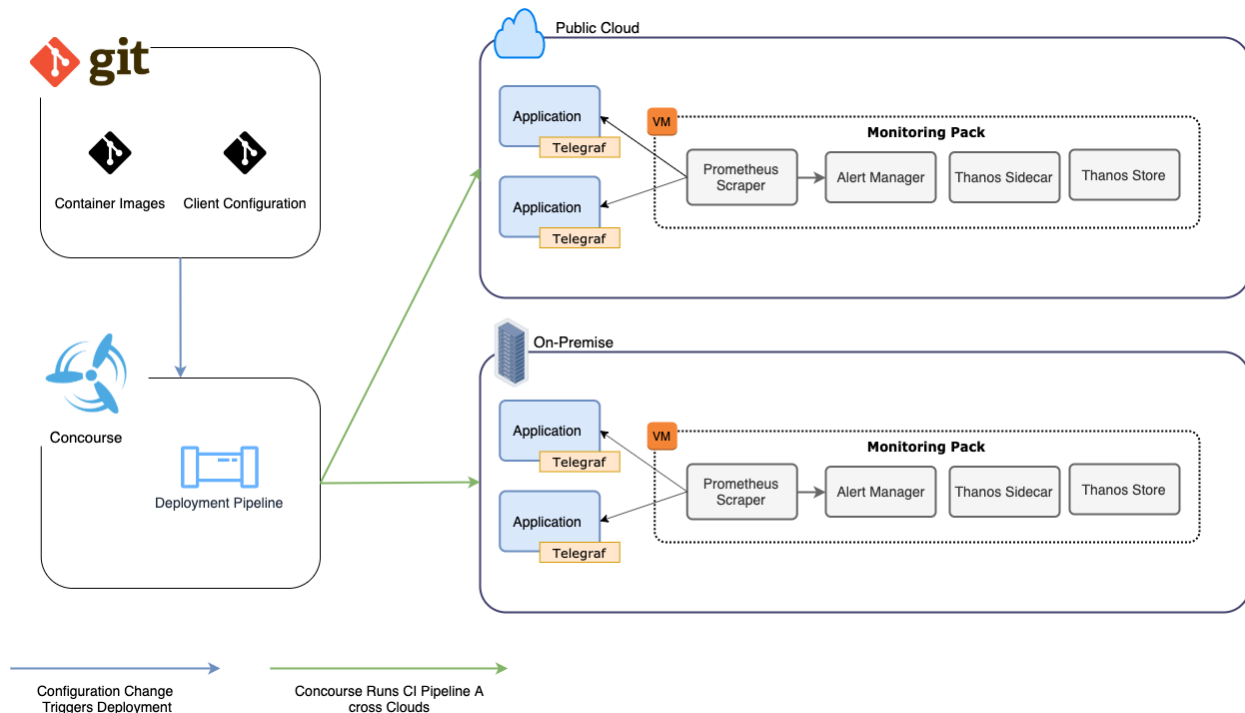


Figure 5 - Deployment Orchestration

6. Conclusion

As cable operators grow their on-premises cloud offerings and expanded into public clouds, there becomes a need for simplified abstractions to deploying infrastructure across environments. This team was tasked with providing this abstraction, specifically for observability tooling, available to all teams. They approached the problem by breaking down the areas into modularized components, capable of abstracting away the complexities of the target cloud environment, application type, and regionality. Focusing on clients' needs helped frame how to abstract the application architecture and functionality. This made the platform easy to use, quick to change and deploy seamlessly. All of this was possible because at each decision point, the software landscape was evaluated. A decision was made to go with proven, open-source tools that would meet the case. This, in turn, allowed the team to rapidly develop and scale the platform to serve any clients who need it, wherever their applications are running.

Abbreviations

API	Application Programming Interface
CI/CD	Continuous Integration and Continuous Delivery
HTTP	Hypertext Transfer Protocol
IaC	Infrastructure as Code
SCTE	Society of Cable Telecommunications Engineers
UI	User Interface
VM	Virtual Machine
YAML	Text based markup language

Bibliography & References

[1] *Prometheus Overview*, Prometheus Authors 2014-202
<https://prometheus.io/docs/introduction/overview/>

[2] *Distributed Systems Observability*, Cindy Sridharan; O'Reilly Media, Inc.

Strategies and Techniques for Ensuring Network Reliability for Enterprise Customers

Dhirendra Singh Kashiwale

Principal Engineer

Comcast

1800 Bishops Gate Blvd, Mount Laurel, NJ

+1 856 571 9134

Dhirendra_kashiwale@cable.comcast.com

Yogesh Ade

Manager 1

Comcast

1800 Bishops Gate Blvd, Mount Laurel, NJ

+1 215 863 1907

Yogesh_Ade@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. General Terms and Concepts about Reliability	4
2.1. Common Terms.....	4
2.1.1. Failure:	4
2.1.2. SLA (Service Level Agreement):.....	5
2.1.3. SLO (Service Level Objective):.....	5
2.1.4. Failure Rate:.....	5
2.1.5. MTBF (Mean Time Between Failures):	5
2.1.6. MTTF (Mean Time To Failure):.....	5
2.1.7. Risk:	5
2.1.8. Maintainability:	5
2.1.9. Observability:	5
2.1.10. MTTR (Mean Time To Repair / Restore):	5
2.1.11. Availability:	6
2.1.12. Supportability:	6
2.1.13. Minimal Path Set:	6
2.1.14. Minimal Cut Set:	6
2.1.15. SDP (Sum of Disjoint Product):.....	7
2.1.16. MVI (Multiple Variable Inversion):	7
3. Reliability As A Mindset.....	7
4. Useful Life Of Components	8
4.1.1. Early Failures:	9
4.1.2. Chance Failures:	9
4.1.3. Wear-out Failures:.....	9
5. Failure Distribution	10
6. Reliability Evaluation of the Network.....	11
6.1. Reliability Evaluation of a large Network.....	12
7. Reliability Cost Optimization	16
8. Conclusion:.....	18
Abbreviations	19
Bibliography & References.....	19

List of Figures

Title	Page Number
Figure 1: Minimal Pathset and Minimal CutSet.....	6
Figure 2: Boolean Lemmas	7
Figure 3: Reliability Cost Savings	8
Figure 4: Example of Bathtub curve created by different life stages	9
Figure 5: Example distribution of categories of events causing outages.....	10
Figure 6: Series System.....	11
Figure 7: Series Parallel System.....	11
Figure 8: Non-Series Parallel System.....	12
Figure 9: Small Sample NSP Network for Reliability evaluation.....	12
Figure 10: Example Point 2 Point link between customer locations	13

Figure 11:Simplified network topology	14
Figure 12 :Reliability cost categories [7]	16
Figure 13: Cost curves for the service/product [7]	17
Figure 14: Optimal preventive maintenance time	18

List of Tables

Title	Page Number
Table 1: k-Shortest Paths.....	14
Table 2: Cut Sets for the reduced network graph	15

1. Introduction

Contemporary networks are playing a critical role in sustaining business continuity in every vertical of the industry. The customers are becoming more sensitive to outage or degradation in service performance. In simple language, customers rely on networks for their day to day business. That brings an opportunity and responsibility on operators to assure that customers can rely on them. This assurance shall be reinforced with a tangible estimate and measure of the reliability metric.

Reliability heavily depends on the knowledge of Statistics, Physics, and Engineering. However, in order to systematically implement Reliability Engineering, the service operators have to evolve a mindset of viewing every aspect of the organization as contributor to the improvement of reliability.

This paper will discuss strategies and techniques to enhance network reliability for enterprise customers in the present state of network design and performance. This paper highlights one approach to reliability and many additional factors are involved in ensuring overall system reliability. These models can then be utilized to evaluate the impact of hardware, software, and machine learning components on reliability of network to the end customer. This paper will delve into the factors that drive optimized reliability goals such as cost, complexity, maturity, redundancy, and operational efficiency, and will illustrate the reliability of networks from conceptual, architectural, monitoring, and cost optimization perspectives.

2. General Terms and Concepts about Reliability

While reliability has been understood and interpreted from varying perspectives, the most widely accepted definition of reliability is stated by Electronics Industries Association (EIA) as follows:

The Reliability of an item (a component, a complex system, a computer program or a human being) is defined as the probability of performing its purpose adequately for the period of time intended under the operating and environmental conditions encountered.[1]

Study and analysis of reliability presents an opportunity to integrate Reliability, Availability, Maintainability and Safety (RAMS) at every stage in the product lifecycle to achieve excellent quality, optimal product reliability, and customer delight.

Reliability engineering is a mindset homogeneously internalized within an entire organization. It requires a collaborative team effort where contributors of diverse perspectives, skillsets, backgrounds, and functional departments synergize to produce superior reliability for the product.

It is accepted industry-wide that reliability of the product or services shall be examined and analyzed at the earliest stage of development. Every missed opportunity translates into a cost increase of a multiple of 10. If a team misses the opportunity to identify a reliability issue during the design stage, then it will cost 10 times more to remediate in the development stage. In essence, reliability shall be applied at every stage in the lifecycle of product.

2.1. Common Terms

2.1.1. Failure:

Failure is an event when the element/service is no longer available to perform as per SLO/SLA.

2.1.2. SLA (Service Level Agreement):

SLA is an agreement written between service provider and customer. This agreement clearly determines the measurable metrics of service quality, penalties, and remedies, along with the roles and responsibilities of both parties in maintaining the mentioned quality of service.

2.1.3. SLO (Service Level Objective):

SLO is a commitment where the service provider declares its intention of maintaining certain level of measurable metrics.

2.1.4. Failure Rate:

The frequency at which failures occur.

2.1.5. MTBF (Mean Time Between Failures):

Average time between the occurrence of failures. This can be calculated as following:

$$MTBF = N_s(t_i) [t_{i+1} - t_i] / (N_s(t_i) - N_s(t_{i+1})); t_i < t \leq t_{i+1}$$

$N_s(t_i)$: Number of Survivors at time t_i

2.1.6. MTTF (Mean Time To Failure):

This metric is similar to MTBF, measuring the average amount of time a non-repairable element operates before it fails.

2.1.7. Risk:

The estimate of likely loss due to failure influenced by the reliability of one or more components of the system

2.1.8. Maintainability:

The probability that an element/service can be retained in, or restored to, a specified operable condition within a specified interval of time when maintenance is performed in accordance with the prescribed procedure. Maintainability is the characteristic of design, installation, and operation.

2.1.9. Observability:

It is a capability of measuring/estimating the internal state of the system by measuring/monitoring the external outputs of the system.

2.1.10. MTTR (Mean Time To Repair / Restore):

This metric is applicable only to repairable element/service. It measures the average time it takes to repair/restore a failure. This metric is an indicator of operational efficiencies and maintainability of the element/service.

2.1.11. Availability:

A measure of time that a system is operating versus the time that the system is targeted to operate

2.1.12. Supportability:

The capability of provider to maintain inbuilt reliability and to perform scheduled and unscheduled maintenance according to the Network maintainability with minimum cost.

2.1.13. Minimal Path Set:

Minimal Path Set is a set whose elements are paths. The System is available if all components of any element (path) are available. Refer to Figure1.

2.1.14. Minimal Cut Set:

Minimal Cut Set is a set of set of nodes. The System is unavailable if all nodes within an element of the Cut Set are unavailable. Refer to Figure 1.

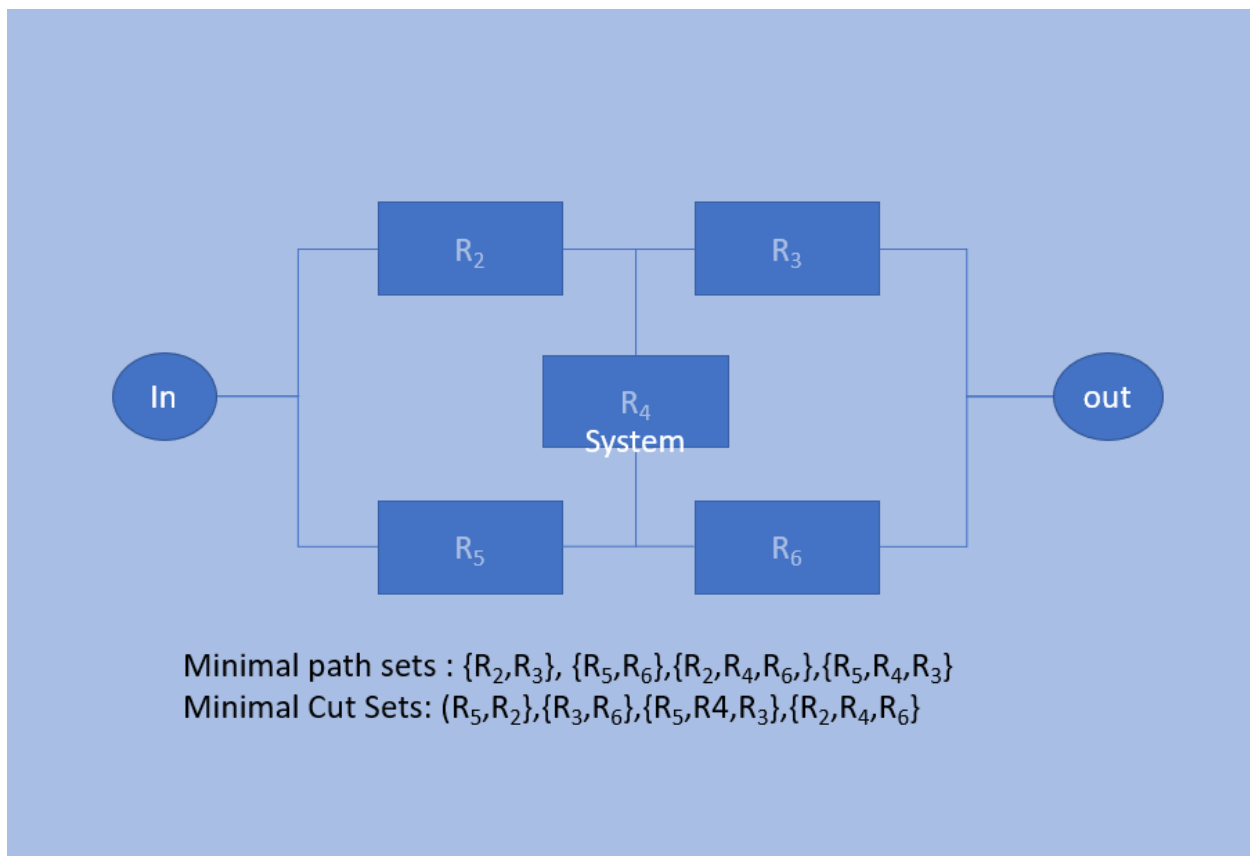


Figure 1: Minimal Pathset and Minimal CutSet

2.1.15. SDP (Sum of Disjoint Product):

The Joint Probability formula

$$\Pr(E1 \cup E2 \cup E3 \dots \cup En) = \Pr(E1) + \Pr(E2) + \Pr(E3) + \dots + \Pr(En)$$

is easy to calculate and valid only when Events E1, E2, E3,...En are mutually exclusive.

Referring to Figure 1, let P1 = R₂R₃, P2 = R₅R₆, P3 = R₂R₄R₆, P4 = R₅R₄R₆ be the Path Sets between Input and Output. Where R₂, R₃, R₄, R₅, R₆ are the components of the network (System).

In order to calculate the probability of success of the network a Boolean expression can be written such that all the terms of that expression are disjoint. This method of evaluating reliability is called Sum of Disjoint product. This disjoint form has one to one mapping with the probability expression.[7]

2.1.16. MVI (Multiple Variable Inversion):

MVI is a technique based on Boolean algebra used to generate a compact expression of SDP terms. In this technique a group of variables are inverted simultaneously. This results in generating a compact Boolean expression at very efficient processing time. The mentioned Lemmas in the illustrated table are used by MVI techniques to extract compact and disjoint form of Boolean expression.[8]

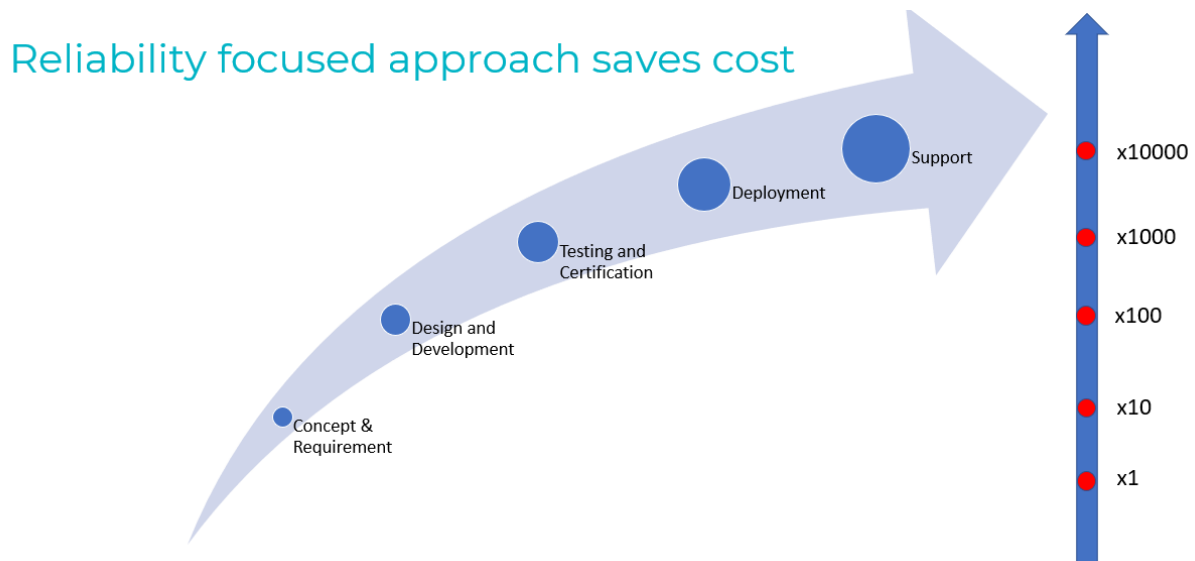
S. No	Lemma	Explanation
1.	$\overline{U}V.U \equiv \emptyset$	
2.	\overline{UV}	$\overline{UV} = \overline{U} + U.\overline{V}$
3.	$\overline{UV}.\overline{U} \equiv \overline{U}$	$(\overline{U} + U.\overline{V}) \overline{U} = \overline{U}$
4.	$\overline{UV}.U \equiv U\overline{V}$	$(\overline{U} + U.\overline{V}) U = U.\overline{V}$
5.	$\overline{UV}.\overline{UW} = \overline{U} + U.\overline{V}.\overline{W}$	$(\overline{U} + U.\overline{V})(\overline{U} + U.\overline{W}) = \overline{U} + \overline{U}.U.\overline{W} + U.\overline{V}.\overline{U} + U.\overline{V}.\overline{W}$
6.	$UV + UW = U.(\overline{U} + (V + W))$	$U.\overline{U} + U.(V + W) = U.(\overline{U} + (V + W))$

Figure 2: Boolean Lemmas

3. Reliability As A Mindset

It is common knowledge that Reliability is synergy of Statistics, Physics and Engineering. This concept of Reliability is absolutely true. However, besides the pure objective part of Reliability Engineering, there is also a subjective aspect. That subjectivity is associated with the mindset that any reliability-aware organization should have intentionally evolved. This mindset involves viewing every part of the organization as contributor to reliability.

This vision of reliability not only enhances the customer experience but also provides enormous cost savings. The following chart will give an idea:



4. Useful Life Of Components

There is a usual failure pattern in the lifetime of components when they are placed in service. These patterns are resultant of weaknesses in the components resulting in early-stage failures, normal random failures due to natural phenomenon of physics, and failures due to aging of the component. The following graph shows the three types of failures with their distributions and also the combined failure:




Figure 4: Example of Bathtub curve created by different life stages

If we take a large sample of components and operate them under constant conditions, a statistical pattern emerges with three regions. Each region providing its own interesting failure/ hazard behavior.

4.1.1. Early Failures:

This pattern of failures is also called burn-in, or debugging period. These failures are related to weakness in hardware, software, or design. Issues that arise during this period can be stabilized. These failures fit in a Weibull distribution. The **failure rate** in this region **decreases very rapidly**.

4.1.2. Chance Failures:

The useful life of the component starts after the burn-in period. This is the period where failure rates are at the minimum level. During this period the failure rate is constant which tells us that chance failures cannot be prevented by any replacement policy. Also, due to **constant failure rate**, these failures fit into exponential distribution.

4.1.3. Wear-out Failures:

The wear-out begins when the element has lived its life in terms of age, stress, or cycles of operation. **Failure rate starts increasing very rapidly** with the start of wear-out time. The simple indicator of wear-out time is a period during which approximately one-half of the total population will fail. These failures fit in Lognormal distribution.

The golden rule of reliability is to replace the component as it fails during the useful life and proactively replace them before the end of their useful life. The actual algorithms and optimization techniques of proactively replacing the component are not in the scope of this discussion but plenty of literature is available in this regard.

5. Failure Distribution

For an operating network it is important to understand the distribution of categories of failures. An example of a distribution is illustrated as follows:




Figure 5: Example distribution of categories of events causing outages

This distribution allows us to identify the opportunities to enhance the reliability of the network. This will provide a very good data point to teams to start with their Failure Mode and Effect Analysis, and Fault Tree analysis.

Few examples from above illustration produces these observations:

- Failure cause code 1531 – Power. This indicates that the failures are attributed to the failure of electrical power managed by the service provider. This opens the opportunity to prevent power failures either by providing redundancy, back up, or in some cases, simple routine preventive maintenance of a backup power generator.
- Failure cause code 150 – Scheduled Maintenance. This failure cause code is attributed to failures that occur during scheduled maintenance. The opportunities for improvement are multiple including the potential for enhancement of method of procedures, network impact analysis, customer impact analysis, and data integrity of the network inventory management systems to name a few.
- Provisioning Incomplete 119 – This failure cause code can be attributed to the process of provisioning all service attributes in the system. It opens a number of systemic and

communication opportunities. Many times, these scenarios bring into light the issues like flexibility of product or need for automation.

6. Reliability Evaluation of the Network

Reliability of the system is the sum of reliabilities of its components. The components themselves are dependent on the reliabilities of their elements. There are numerous techniques and algorithms to determine overall reliabilities of the system.

A system can be a small system with very few components in it, or the components may be connected in a simple manner such as a series. Any component failure will cause system failure.

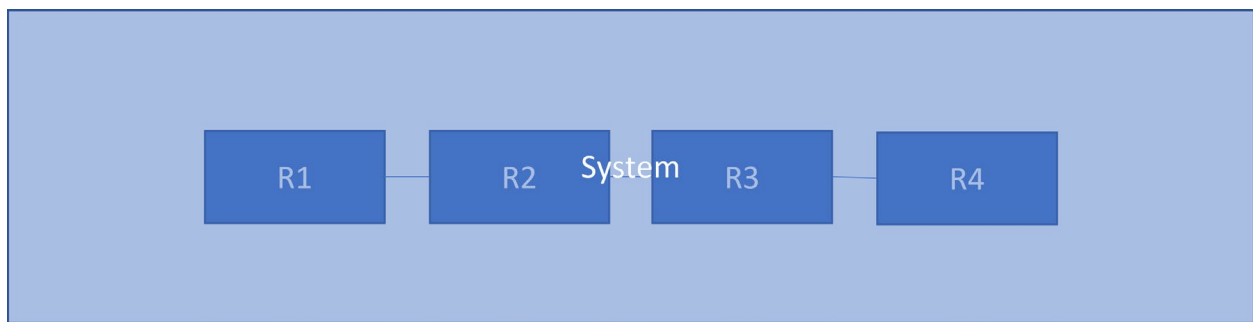


Figure 6:Series System

Another configuration could be that a few components are connected in parallel for the sake of redundancy. This system is termed as Series Parallel System.

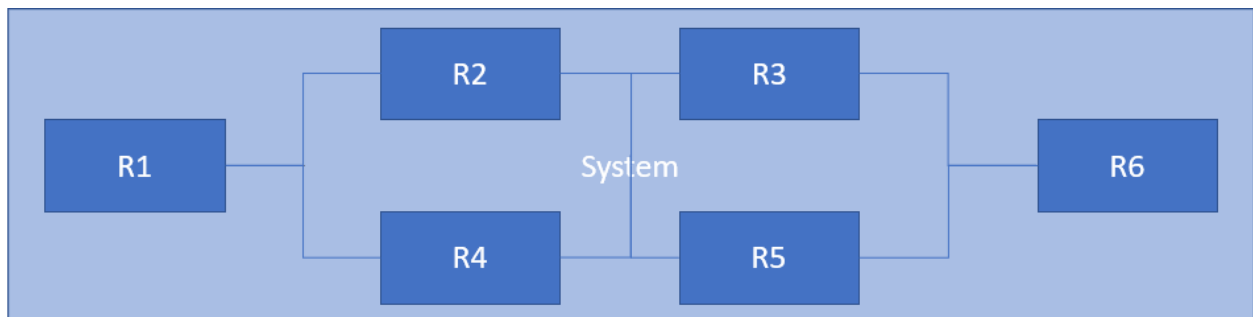


Figure 7:Series Parallel System

The next configuration is a Non-Series Parallel system (NSP). This is the configuration that presents us a reliability evaluation problem in most of our scenarios. The example is illustrated in the following picture.

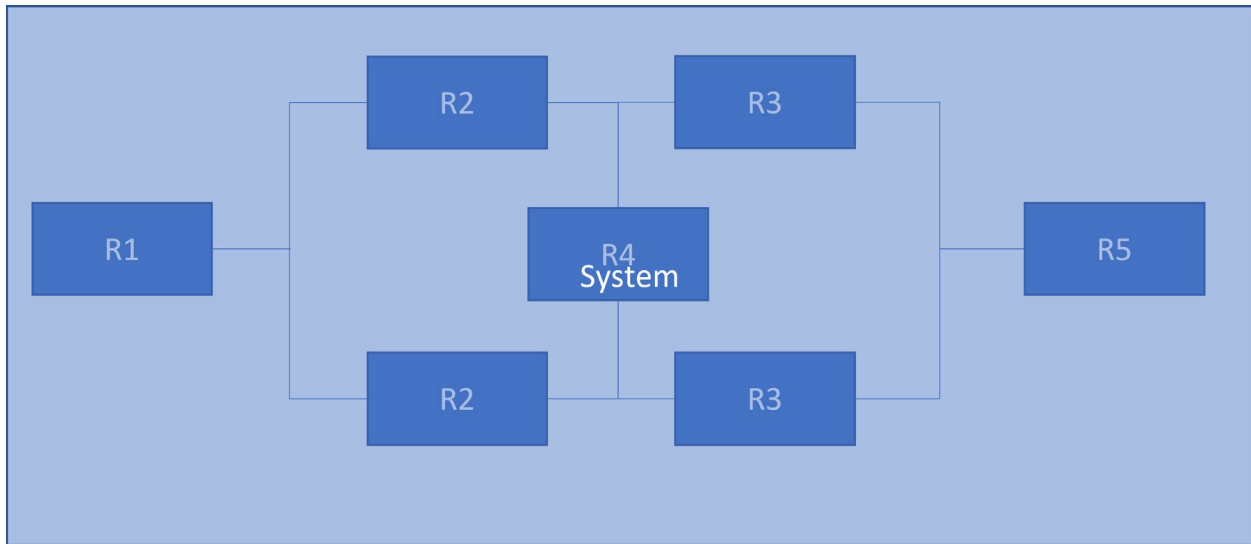


Figure 8: Non-Series Parallel System

6.1. Reliability Evaluation of a large Network




Figure 9: Small Sample NSP Network for Reliability evaluation

The above picture shows a network with n1, n2 as input and output nodes respectively. Nodes n3, n4 are intermediate nodes. A, B, C, D, and E are edges.

In order to evaluate the reliability of the above network following two assumptions are made in order to simplify the mathematics:

1. Edges failure (success) are statistically independent
2. Nodes are perfectly reliable

Let p_a, p_b, p_c, p_d, p_e be success Probabilities of Edges A, B, C, D, E respectively.

Also, q_a, q_b, q_c, q_d, q_e be failure Probabilities of Edges A, B, C, D, E respectively.

Using available reliability evaluation methods, the reliability of the system is calculated as follows:

$$R(z) = p_a p_b + p_a p_c q_b p_d$$

There are several algorithms and techniques for evaluation of reliability of large networks. They had their own advantages and disadvantages in terms of efficiency, scalability, and accuracy.

With the recent contributions and advancements, Graph Theory has been playing a very important role in the field of reliability evaluations of large networks.

In this paper we will evaluate the reliability of a very large network using Sum of Disjoint Product (SDP) and Multi Variable Inversion (MVI) techniques. This is a three-step process.

1. Create reduced network topology
2. Extract minimal path set or minimal cut set from the topology
3. Evaluate reliability from the path set or cut set extracted from step 2 using SDP and MVI

Step1.

For the current networks spanning large distances, operators are using the model of access network delivering to a full or nearly fully meshed core network as illustrated below:




Figure 10: Example Point 2 Point link between customer locations

From this illustrated network it is evident that there are potentially hundreds of millions of paths sets or cut sets. This makes any reliability evaluation algorithm reach its processing limits and accuracy suffers due to roundup errors.

In this paper we address this problem by reducing the size of this network into a focused network. We are aware of the situation where network boundaries are limited by the latency introduced by the links. If the total latency introduced is beyond the service level objective, then that path is a failed path (infinite horizon). Based on this philosophy we have extracted path sets belonging to k-shortest paths. Following table list the 8 shortest paths between node 1 and node 2011:

Table 1: k-Shortest Paths

Path	Latency(ms)
[1, 25, 62, 2011]	12.20812742
[1, 53, 62, 2011]	12.20812742
[1, 25, 37, 62, 2011]	12.22414687
[1, 53, 65, 62, 2011]	12.22414687
[1, 25, 65, 62, 2011]	12.22414687
[1, 53, 37, 62, 2011]	12.22414687
[1, 25, 34, 2011]	12.40271075
[1, 53, 34, 2011]	12.40271075

These paths are then merged into a new reduced subgraph as illustrated below.




Figure 11: Simplified network topology

Step2.

At this stage we must decide whether we should utilize a cut set or path set approach. It has been suggested by Aggarwal, Chopra, & Wajwa, (1982) that for a network of n nodes and l links, the number of cut sets between any pair of nodes would be of the order of 2^{n-2} , whereas the number of path sets is of the order of 2^{l-n+2} . From this recommendation the estimated cut set = 64 and path set = 256.

Hence, we have decided to use cut set approach.

The cut sets are enumerated as per algorithm described in Ahmad, (1990).

The following cut sets were returned by using the algorithm:

Table 2: Cut Sets for the reduced network graph

Cut Sets		
2 8	1 7 8 9 10 14	1 5 7 8 9 12 14
6 7	2 3 7 11 12 13	2 3 4 7 9 11 12
2 11 14	3 4 6 9 12 14	2 3 5 7 10 11 13
1 3 4 5 8	3 5 6 10 13 14	3 4 5 6 9 10 14
1 7 9 10 11	3 6 8 11 12 13	3 4 6 8 9 11 12
3 6 12 13 14	1 2 4 6 10 13 14	3 5 6 8 10 11 13
1 2 6 9 10 14	1 2 5 6 9 12 14	1 2 4 5 6 12 13 14
1 3 4 5 11 14	1 3 4 10 11 12 14	1 3 9 10 11 12 13 14
1 3 4 8 10 12	1 3 5 9 11 13 14	1 4 5 7 8 12 13 14
1 3 5 8 9 13	1 3 8 9 10 12 13	2 3 4 5 7 9 10 11
1 4 7 10 11 13	1 4 5 7 11 12 13	3 4 5 6 8 9 10 11
1 5 7 9 11 12	1 4 7 8 10 13 14	

We can visually verify that the links in each cut set is isolating network with node 1 (Source) with node 2001(Sink).

Step3.

Finally, we can now evaluate reliability of the network by using CAREL algorithm (Soh, S. & Rai, S., 1991 [7]). The input that mentioned algorithms needed are minimal cut sets and failure probability of the link. We assumed that all links fail with 0.1 probability.

After applying CAREL algorithm to the minimal cut sets derived in Step 2, the Reliability / Unreliability of the Mentioned P2P links is calculated as following:

System Unreliability = 0.02082693568

System Reliability = 0.979173064

With total disjoint paths = 76

7. Reliability Cost Optimization

Every organization aspires to achieve maximum profit while keeping customer satisfaction at an optimum level to sustain the profit margin. Reliability engineering is the scientific tool that management can use to keep balance between customer satisfaction and cost of the product or services.

The following picture is illustrated to enumerate important categories that contribute to the cost incurred by implementing Reliability enhancement methods.

Reliability Costs				
Internal Failures	Prevention Cost	Administrative Cost	Testing and Detection Cost	External Failure Cost
<ul style="list-style-type: none"> • Root Cause Analysis • Testing Cost • Fix , Upgrade cost 	<ul style="list-style-type: none"> • Vendor evaluation Cost • Device Certification Cost • Network Design certification Plan • Customer Impact Analysis • Data collection, storage and Analysis 	<ul style="list-style-type: none"> • Reviewing Contracts • Preparing Paring Budgets • Forecasting • Management 	<ul style="list-style-type: none"> • Cost of Monitoring and Detection(Operation Centers) • Cost of Testing Infrastructure 	<ul style="list-style-type: none"> • Cost of Customer care • Cost of customer credits • Cost of replacement and credits in leu of damage control • Cost of repair

Figure 12 :Reliability cost categories [7]

It is not always profitable to increase reliability to achieve perfection. There is always an optimum point where a balance should be made. That decision of balance shall be dictated by the facts emerging from reliability analysis as shown in this diagram.




Figure 13: Cost curves for the service/product [7]

There are several cost reliabilities functions in the literature. A few have been listed here:

- Misra et al Function
- Tillman et al Function
- Aggarwal et al Function
- Fratta et al's Function
- Majumdar et al's Function
- Llyod and Lipow's Function

A hypothetical study is illustrated here on a model trained on real failures from a very large network. Here the goal is to showcase an optimum time at which the preventive maintenance shall be done on the devices. In this example the cost of preventive maintenance is 5 and cost of corrective maintenance is 200.

Cost model assuming as good as new replacement ($q=0$):
The minimum cost per unit time is 0.0092
The optimal replacement time is 1095.22




Figure 14: Optimal preventive maintenance time

8. Conclusion:

As we discussed in this paper reliability of the network needs to be measured, monitored, tested, and investigated continuously. The reliability should be investigated for each product right from its inception. We have talked about different life stages of components. Reliability evaluation technique is demonstrated with an optimized approach for large networks.

Referring to Figure 14, the reliability studies can optimize the opex and capex by making optimally calculated decisions.

There is an enormous opportunity for improvement, development, and enhancement for telecom networks' reliability and availability.

Abbreviations

MVI	Multiple Variable Inversion
NSP	Non-Series-Parallel systems
RAMS	Reliability Availability Maintainability and Safety
SCTE	Society of Cable Telecommunications Engineers
SDP	Sum of Disjoint Products
SLA	Service Level Agreement
SLO	Service Level Objective
SVI	Single Variable Inversion

Bibliography & References

Reliability Engineering: Theory and Practice 8th ed. 2017 Edition: Alessandro Birolini [1]

Quality And Reliability Engineering International. VOL. X. 349-354 (1992): Economical Design of Reliable Systems – Some Practical Situations: K.K. Aggarwal [2]

Shen, Y. (1995). A New Simple Algorithm for Enumerating all minimal Paths and Cuts of a Graph. Microelectronics Reliability, Vol. 35(6), 973–976. [3]

Mishra, R., Saifi, M. A. & Chaturvedi, S. K., 2016. Enumeration of Minimal Cutsets for Directed Networks with Comparative Reliability Study for Paths and Cuts. Quality and Reliability Engineering International, Vol.32(2), pp.555–65. DOI: 10.1002/qre.1772 [4]

Ahmad, S. H. (1988). Simple Enumeration of Minimal Cutsets of Acyclic Directed Graph. IEEE Transaction on Reliability, R-37, 484–487. [5]

Ahmad, S. H. (1990). Enumeration of Minimal Cutsets of an Undirected Graph. Microelectronics Reliability, Vol.30(1), 23–26 [6]

Soh, S. & Rai, S., 1991. CAREL: Computer Aided Reliability Evaluator for Distributed Computing Networks. IEEE Transaction on parallel and Distributed Systems, Vol.2(2), pp.199–213. [6]

Reliability Engineering: K.K. Aggarwal [7]

Network Reliability Measures and Evaluation: Sanjay K. Chaturvedi [8]

Irredundance Graphs: C.M. Mynhardt; A. Roux; arXiv:1812.03382v3 [math.CO] 7 Apr 2021 [9]

Teach Me to Fish: The Role of Virtual Training Environments in Workplace Learning

A Technical Paper prepared for SCTE by

Abbie O'Dell

Sr Dir, Learning Services: Field Operations
Charter Communications
6399 S Fiddler's Green Cir
Greenwood Village, CO 80111
720-482-4205
Abbie.odell@charter.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Overview	3
3. Literature Review	4
3.1. Adult Learning Theories and Practices	4
3.1.1. Self-Directed Learning Theory	5
3.1.2. Learning Transfer.....	5
3.2. Software Training for Adult Learners: Early Inquiries	6
3.3. Technology for Adult Learning: Simulations and Learner-Centered Design	6
3.4. Software Video Tutorials in the K-12 Environment	7
3.5. Conclusions.....	8
4. Research Methods	8
4.1. Participants.....	8
4.2. Permission.....	10
4.3. Data Collection	11
4.4. Data Analysis	11
5. Results	12
5.1. Categories	14
5.1.1. Safe Environment	14
5.1.2. Learning from Mistakes.....	15
5.1.3. Business Impact.....	15
5.1.4. Confidence	15
6. Discussion	15
6.1. Authenticity of Learning Environment and Learning Transfer.....	16
6.2. Learning from Mistakes and Safe Learning through Self-Direction	16
6.3. Study Limitations	16
6.4. Future Research	17
7. Conclusion.....	17
Abbreviations	17
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - Example Reporting Structure	10
Figure 2 - Question Responses – Trainee Confidence Related to Presence of VTEs	13
Figure 3 - Question Responses – Trainee Proficiency Related to Presence of VTE	14

List of Tables

Title	Page Number
Table 1 – Population Data – Total Headcount by Job Title	9
Table 2 – Sample Data – Total Headcount for Selected Regions	9
Table 3 – Participant Data – Job Title Distribution.....	10
Table 4 – Response Rate by Job Title.....	12
Table 5 – Demographic Data	12

1. Introduction

This qualitative study considers the relationship between an immersive Virtual Training Environment (VTE) and the post-training confidence of learners, through examining the perceptions of trainers. Study participants were selected from trainers in a large telecommunications organization who provide instruction on software and systems, including those both with and without a VTE. Six participants responded to an online survey containing closed- and open-ended questions that gathered their perceptions of post-training confidence for learners relative to use of a VTE during their training course, and responses were analyzed to identify key topical patterns. The findings indicate a perception that the presence of a VTE provides significant positive impact to the learner experience, and that the absence of one can be equally detrimental. Concepts from the literature, including exploratory and participatory learning, self-directed learning, learning transfer and others were found throughout the participant responses. The findings indicate a clear preference for VTEs and suggest an opportunity for future inquiry to establish the validity of this instructional method.

2. Overview

The use of computer software and applications has become intertwined with the work performed by most employees in both large and small organizations in the United States. A 2001 study by the US Department of Commerce estimated that 65 million adults use a computer in order to perform their job, and by 2003 that number had risen to 77 million, accounting for 55.5% of total employed persons (US Department of Labor, 2005). The ensuing years have only continued this upward trend. The transition to remote work as a result of COVID-19 further shifted the workforce to an online computer-supported work model, with many workers performing all job duties remotely via computer. Despite this clear reliance on software and applications in the workplace, many organizations still struggle to provide effective training to employees on software and applications that are necessary to perform their work.

Companies invest significant capital in proprietary software systems but still suffer implementation failures, some of which can be related to employees' inability to effectively use the systems (Marler et al., 2006). Participants often lack confidence and proficiency in the use of the software tools at the conclusion training programs, which contributes to errors and impacts both employee performance measures and the business overall. Why are some companies willing to fund the creation of new software and applications, yet hesitant to fund the corresponding training programs needed to ensure their success?

While we may consider this a modern problem, studies in the early 1990s clearly indicated a need to conduct further research to correlate failures of workplace technology to communication and training methods, and that workplace computer and software training should incorporate an understanding of learner needs in their design (Martocchio & Webster, 1992; Turnage, 1990). More recent works explore relationships between instructional methods such as video tutorials to post-training performance, learning transfer, and learner self-reported satisfaction (Van der Meij et al., 2018; Roumell, 2018; Lavendels et al., 2014).

Research has been conducted on the methods of instructional design to use in the creation of simulations or training videos (Van der Meij, 2013a, 2013b, 2014; Van der Meij et al., 2018), but minimal research has been performed relative to the concept of a virtual training environment (VTE), here defined as a full-function, separate instance of the software used to create an immersive learning experience. Virtual training tools do exist for complex procedures and systems in the medical field, and a 2017 study of a tool used to teach radiotherapy indicated that despite the availability of a virtual environment for training, many of the advanced features of the training system were unused and some organizational skepticism existed (Bridge et al., 2017). Further, current studies of adult learners in a workplace setting related to

use of VTEs for software and application training are extremely limited. The work of Van der Meij (2013a, 2013b, 2014) and Van der Meij et al. (2018) provides a fascinating and relevant starting point, but focuses on the experience of K-12 learners rather than workplace learners. Further inquiry into the experience of adult learners in the workplace could help expand the collective understanding of this important and timely topic.

Based on my own observations and anecdotal conversations with both instructors and learners during my eighteen years as an adult learning professional in a large technology organization, learner confidence and proficiency post-training may have a relationship to the methods used in the training program. Specifically, learners who are afforded an opportunity to use a VTE report greater confidence in its use after the completion of the program. Learners who have been presented only video simulations have struggled to effectively use new software and their self-reported confidence post-training is low. The continued increase of complex and proprietary software applications as an integral part of the modern workplace indicates that further research specifically targeting the efficacy of VTEs for software and application learning is warranted and necessary.

Exploring the relationship between the presence of VTEs and efficacy of training programs can help workplace learning professionals understand the need for the initial investment in a VTE through a consideration of post-training performance. Further, instructional designers can better customize the design of software training programs to include structured VTE activities that improve learner experience and provide greater learning transfer. Businesses may be hesitant to invest capital in VTEs without a clear relationship to a return on that investment, and closer examination of the impact of VTEs on post-training trainee confidence, proficiency and workplace effectiveness can help with that justification.

3. Literature Review

Data from the US Department of Labor (2005) indicate that workplace use of software and applications will only continue to trend upward in the future, and companies seeking to remain vital and retain a strong workforce should cast a critical eye to the training practices used for these systems. In reviewing current literature specifically relevant to the topic of VTEs and software training methods, we will consider two main areas. First, we will examine literature related to the core principles of adult learning theory and practice, including self-directed learning theory and learning transfer. Second, we will explore current literature specifically investigating practices and methods used for software and application training in both K-12 and workplace settings.

3.1. Adult Learning Theories and Practices

What specifically differentiates the concept of a Virtual Training Environment (VTE) from other methods of software and application instruction such as videos or e-learns? How might existing research help to differentiate a need for this instructional methodology? An answer may be found in early concepts of andragogy and adult learning theories first explored by Knowles, who presented a series of five foundational characteristics of the adult learner which have since been incorporated into the collective understanding of adult learning theory (Merriam & Bierema, 2014). Among these principles are that adults have “an independent self-concept. . . and can direct [their] own learning” and that they are “problem-centered and interested in an immediate application of knowledge” (Merriam, 2001, p. 5). These two foundational principles directly support the use of software VTEs in the workplace learning environment, since one of the differentiating factors of a VTE compared to other instructional methods is that the learner is afforded an opportunity to experience an immersive environment in which unstructured self-directed learning can occur, and practice can directly mimic very real and practical skills that can be immediately put into use back on the job.

3.1.1. Self-Directed Learning Theory

Self-Directed Learning (SDL) as a learning theory is a fundamental concept and has been extensively researched and practiced for over fifty years (Merriam & Bierema, 2014). Tough's 'Adult Learning Projects' during the 1970s provided an in-depth study of how adult learners work through the process of selecting the topic(s), resources, and materials needed to support a specific learning endeavor (Abdullah et al., 2008). A more recent piece by Hendriks et al. (2018) focuses specifically on the workplace learning experiences of customer-facing employees, as related to SDL and integration of technology. The study found a correlation between participant comfort with technology and a perspective that technology is integral to work and career development.

Caruso (2018) further explored the correlation between the effectiveness of learning outcomes and the use of technology-supported SDL methods during learning events. She found that the use of Web 2.0 technologies (defined in the scope of the paper as socially-driven resources or applications such as media sharing, discussion boards, search tools and the like) are effective in supporting both structured and informal learning events in the workplace if they are strategically and deliberately used and the guidelines around their use are clear to the learner. Learner affinity and comfort with technology also directly impact the effectiveness of technology-supported SDL and in order to remain nimble in a quickly changing marketplace, employees need to be afforded the technological tools to enable them to use SDL principles and informal learning to meet the needs of the organization (Fleming et al., 2014).

3.1.2. Learning Transfer

While the concepts of SDL instructional methods provide a compelling case for the use of VTEs from a participant experience perspective, an even stronger case can be made related to the concept of learning transfer, which is most simply explained as the ability to put the new skills learned in the classroom environment into practice (Roumell, 2019). Foley and Kaiser (2013) defined the different levels of learning transfer, including near and far, positive and negative, and high- and low-road transfer. The context of the VTE aligns closely with the concepts of near transfer, meaning the newly-experienced situation is similar to the original learning; and low-road transfer, where the technique or skill is practiced extensively in the learning environment so that its replication is nearly automatic in the new experience (Foley & Kaiser, 2013). Foley and Kaiser also noted that instructional practices can become barriers to learning transfer, specifically noting situations in which opportunities to practice transferable skills are lacking in the learning environment. To remediate this, the authors recommend several techniques, the most relevant of which to our current context is the concept of scaffolding, which affords learners structured tools to enable them to construct their learning. This aligns with the concept of a VTE, in that the learner is empowered to construct their own learning experience while still receiving support and guidance (scaffolding) of the learning process by the instructor.

Hardré (2013) explored the concept of learning transfer specifically in the context of technology training and proposed that the effectiveness of learning transfer is related to the concept of authenticity, defined as the realism of the training experience compared to the actual environment. Hardré further noted that effectiveness of software training videos or e-learning courses is limited when there is low authenticity to actual tasks or work environments. Relevant to the scope of this work is Hardré's concept of the authenticity of the learning environment, considering that often the learning environment holds less distractions, variables and errors than the "real world" and as a result, the learner may not be able to effectively transfer their skills outside of the classroom. The concepts of authenticity of representation and authenticity of interactivity are key to the potential value of a VTE, since they specifically indicate the need for the technology learning experiences to be as true-to-life as possible, affording the learner an experience that effectively mimics their ultimate experience.

3.2. Software Training for Adult Learners: Early Inquiries

Beyond the core adult learning theories and practices that support the concept of VTEs, research and inquiry into the nuances of software training for adult learners provides important insights. A need for research into the effectiveness of software training was recognized in the early 1990s, and researchers sought to explore how training methodologies and course design could be used to maximize learning outcomes (Martocchio & Webster, 1992; Turnage, 1990). Martocchio and Webster (1992) proposed that “cognitive playfulness” during the learning process is related to greater effectiveness, indicating that learners “exercise and develop skills through exploratory behaviors, resulting in enhanced task performance” (p. 557). They conducted a study with 68 individuals employed at a large public university who were enrolled in a training course on a word processing program, and their findings indicated a strong correlation between learners high in playfulness to positive outcomes both in test scores and post-training performance. They recommended further research on the topic, noting that outcomes may vary between students and those in a workplace learning setting. The early need for a customized and enhanced approach to workplace training to suit the computer-based work environment was also recognized by Turnage (1990) who somewhat presciently stated:

“Training will change as computer based training becomes more prevalent with new applications including embedded training, computer literacy, interactive video disc, and electronic lectures. Intelligent computer-assisted instruction (CAI), authoring systems, hand-held computers, speech processing, and new telecommunication technologies will also shape the future direction of automated instruction in the workplace.” (p. 176)

3.3. Technology for Adult Learning: Simulations and Learner-Centered Design

While these early works provide us important context, recent studies allow for more relevant and timely inquiry into the topic of software training methods for adult learners. Hardin et al. (2013) explored the effectiveness of computer simulated software training systems (CSSTS) and investigated the connection between software self-efficacy and post-training effectiveness. In this context, a CSSTS is defined as “a specific type of e-learning self-study system that has become immensely popular for facilitating software instruction” (Hardin et al., 2013, p. 4). Important to note here is the difference between the CSSTS and the VTE discussed within the scope of this inquiry - specifically, the CSSTS is a *simulation* of the software, which provides a more structured experience for the learner; whereas VTE is a fully immersive instance of the *actual* software that allows the learner greater self-direction and control of the learning activities. The research reviews user perceptions of a CSSTS that uses a model of “Teach Me, Show Me, Let Me Try” to guide learners through the steps of the learning process, moving from a verbal instruction, demonstration, and unaided practice. The findings of the study indicate that learners with a high self-reported software self-efficacy (SSE) score (in other words, those who indicated on a survey that they were confident in their ability to learn a new system) were less likely to utilize the “Teach Me” and “Show Me” portions of the CSSTS, but instead were more likely to proceed directly to the “Let Me Try” feature.

Lavendels et al. (2015) explored the use of an online learning methodology in the insurance industry for employees learning to use a complex software solution that contains sensitive customer data. The research proposes a remote training process in which the trainer and the trainee are in different locations, and the trainer utilizes a screen-share application to review the trainees’ work and provide feedback. Due to the intensive nature of the program discussed, Lavendels et al. specifically called out the need to have synchronous trainer oversight and coaching for the trainees as they learn, since their activities are performed within the production system rather than a VTE. The concerns raised about training in a

production environment align with the scope of this inquiry related to a VTE, although the remote instruction aspect will not be addressed here.

Bridge et al. (2017) studied user perceptions of the Virtual Environment for Radiotherapy Training (VERT), an immersive 3-D software solution designed to teach medical personnel skills and procedures needed for administration of radiation therapy to patients. Users of the VERT tool were surveyed, and while overall feedback on individuals' perceptions of the value of the tool was positive, the researchers found that 32% of those surveyed reported that their organizations had a perception that VERT was not actually useful. The researchers were not able to identify a clear reason for these organizational impressions, and greater inquiry into this type of impression of virtual training systems and software could be beneficial to helping identify why companies may be hesitant to invest in this type of learning solution.

The concept of participatory design (PD) in technology learning experiences was explored by Inguva et al. (2018), citing the value of including various self-directed and experiential learning methodologies to create a more learner-centered environment in a university engineering laboratory. Specifically, the researchers studied the use of a practical, hands-on learning model called the "Knowledge Laboratory" in which undergraduate students were provided an opportunity to work in small groups in a self-directed fashion through exploratory and experiential learning. The researchers surveyed students who participated in a more traditional instructor-centered course design, and those who participated in the PD model. Those in the instructor-centered model reported that they did not feel challenged by the delivery method, had less interest in the content, and struggled to find relevance in the content and topics. Conversely, students who were provided the opportunity to approach their learning in a more participatory, self-directed method with more realistic practice scenarios reported they were better able to understand the relevance of the work, and that their overall understanding of the processes being taught were greater.

Both Inguva et al. (2018) and Bridge et al. (2017) noted that an overarching benefit of their respective virtual environments was the ability for learners to gain confidence and skills through realistic simulations in the learning environment, rather than in a higher-stakes environment after the learning. This closely aligns with the concept of the VTE, where learners are provided an opportunity to practice in the safe learning environment of the classroom without fear of making mistakes that could impact customers or the business.

3.4. Software Video Tutorials in the K-12 Environment

Extensive and detailed research exists in the K-12 space on design and implementation of software training, although the research is centered around the creation of software videos rather than a VTE. In a 2013 work by van der Meij and van der Meij, the authors proposed a series of eight guidelines for the design of instructional videos for K-12 software training and provide a definition for the term "video tutorial" to mean a "set of videos that together form an instructional package" (p. 207). A notable difference between this paper's definition of a VTE and the definition from van der Meij and van der Meij is that the video tutorial is a guided, structured learning experience; whereas the VTE is an open-ended unstructured environment in which the learner is free to explore and self-direct the process of learning. Regardless, many of the design concepts and guidance discussed in the work are equally applicable to the concept of a VTE. Specifically, Guideline 2 from the work calls out the need to ensure that the video tutorial "gives the user the same image that he or she is likely to be facing when trying to execute the task" and Guideline 3 specifies the need to incorporate an element of user control of the video tutorial (van der Meij & van der Meij, 2013, p. 210). Additionally, the researchers note that their design principles align both with the concept of constructivism and with multimedia design principles from Mayer (2003). Both Guideline 2 and Guideline 3, and the constructivist and engagement theories are

aligned with the use of a VTE since it is an immersive instance of the software. A later work by van der Meij et al. (2017) elaborated on the guidance from the 2013 study and noted that video tutorials not adhering to these design principles failed to successfully prepare the learner to use the software, even when the steps were clearly and accurately portrayed in the video.

3.5. Conclusions

Existing research and inquiries provide critical insights into learner needs in the software and application training space. While adult learning theories and practices have a long history, recent works continue to support the needs identified by Knowles in regards to adult learner self-direction, exploratory or experiential methods, and the importance of learning transfer (Foley & Kaiser, 2013; Hardré, 2003; Roumell, 2019). Core andragogical theories such as self-directed learning theory and learning transfer provide overall conceptual support for the value of a VTE, relative to the need for authenticity of the training experience and appropriate learner involvement in constructing meaning. Beyond the theoretical considerations, current research describing training methods or practices provides additional perspective, but there is still a gap in the collective body of work. Specifically, most research has been done on the modality of e-learns, videos or other more guided/structured teaching methods, rather than the open-ended exploratory environment a VTE provides. In addition, a significant amount of the available research focuses on participants in a K-12 environment, rather than the workplace environment of adult learners.

The purpose of this qualitative study is to describe the impact of the presence or absence of a Virtual Training Environment (VTE) on the post-training confidence and proficiency of trainees, by exploring opinions of trainers at a large national telecommunications organization. To do so, we pose the following questions:

- What is the trainer perception of the post-training confidence level for telecommunications employees who are provided access to a Virtual Training Environment?
- What is the trainer perception of post-training confidence level for new telecommunications employees who are not provided access to a Virtual Training Environment?

4. Research Methods

As detailed in the literature review, there is a lack of current research that directly addresses the topic of Virtual Training Environments (VTEs) in workplace learning. Seeking to gain a better understanding of the impact of these tools through the perspective of trainers required an open-ended approach that did not begin with the end in mind, as traditional quantitative research typically does. Instead, a qualitative method allowed for a more curious and exploratory approach, which provided an opportunity to learn from the experiences and attitudes of the study participants. Qualitative study enables insight into the human experience and perspective, which is a foundational element of what this research sought to identify (Creswell, 2013). It allowed a level of detail that quantitative study would not provide by capturing the very personal perspective of the participants, to help broaden the understanding of the central phenomenon we are studying. Finally, since qualitative research begins with an open mindset, the findings of this research may help inform future study on this topic by identifying themes and patterns which warrant further inquiry (Creswell & Guetterman, 2019).

4.1. Participants

The participants in the study are Training professionals in a large telecommunications organization, who support the delivery of coursework to new and experienced employees within the Field Operations

business unit. In their roles, they provide training on software and systems, including some that currently have a virtual training environment (VTE) and others that do not. As a result, their perspectives and attitudes on trainee confidence post-training relative to VTEs provides valuable insight into this research.

These individuals are a diverse group of males and females in various locations across the 41 states in which the company operates, and range in age from mid 20s to early 60s. Since the study is being performed within a workplace setting, and age, gender, and race are considered protected data in this scope, participants were not asked to provide this specific demographic information as part of the study.

Table 1 – Population Data – Total Headcount by Job Title

Job Title	Total
Field Tech & Safety Trainer	73
Training Manager	11
Senior Field Tech & Safety Trainer	79
Training Supervisor	32
Technical Service Trainer	12
Total	207

The total population of participants within the organization was approximately 200 at the time of the study (see Table 1). Because the research was intended to explore the trainers' ideas on VTEs, the concept sampling method was selected. This method enables purposeful selection of participants, areas or sites in order to uncover information about the research topic (Creswell, 2003; Creswell & Guetterman, 2019). Three of the total eleven regions were selected to participate, targeting those that have trainers who were known to be actively teaching coursework on software and applications that have VTEs, including the Northwest Region, Northeast Region, and Great Lakes Region. The Training Managers and Human Resources leaders for each of the selected regions were notified via email that they could select up to five training professionals from their team to participate in the survey. The leadership were not provided any specific criteria for selection and were free to select any of their employees. One Manager indicated that their selection was based on those employees who frequently use training environments (J. Knapp, personal communication, April 6, 2021). Northeast and Great Lakes each provided five, and Northwest provided three, for a total of 13 participants. This represented 6.2% of the total training staff population, and represented 33% of Great Lakes training staff population, 23.8% of the Northeast training staff population, and 18.7% of the Northwest training staff population (see Table 2).

Table 2 – Sample Data – Total Headcount for Selected Regions

Region	Total	Number Selected	% of Regional Training Population
Field Ops Great Lakes	15	5	33.0%
Field Ops Northeast	21	5	23.8%
Field Ops Northwest	16	3	18.7%

In addition, the distribution of job titles was five Field Technical & Safety Trainers, (7% of the total population of this job title), five Senior Field Technical & Safety Trainers (6% of the total population) and three Technical Service Trainers (25% of the total population), as shown in Table 3.

Table 3 – Participant Data – Job Title Distribution

Job Title	Total	Number Selected	% of Total Population
Field Technical & Safety Trainer	73	5	7%
Senior Field Technical & Safety Trainer	79	5	6%
Technical Service Trainer	12	3	25%
Total	13	164	8%

4.2. Permission

The Senior Directors of Human Resources leadership within each Region were the gatekeepers who authorized access to the study participants detailed in the Participants section. They directly manage the training departments, and the recruiting and training process as a whole within their respective regions (see Figure 1). Their permission was necessary to gather data from the participants (indicated in green in Figure 1), since they are the senior leaders for that organization (indicated in orange in Figure 1). In addition, they are uniquely suited to provide this permission since their positions are responsible for effective training for new and existing employees. Permission from the leaders within the regions whose employees were surveyed was obtained via email using the format shown in Appendix A. Training Managers for each Region were included on the email requesting access to the study participants, but the final gatekeepers were the Human Resources Directors (see Figure 1). A side benefit of using these individuals for permission for the current study is that interest may be generated, which may open the door to future studies on the central phenomenon being explored.




Figure 1 - Example Reporting Structure

Note. Example of typical reporting structure including gatekeeper (orange) and study participants (green)

4.3. Data Collection

Respondents were sent a cover letter via email, which provided background on the importance and purpose of the study, as well as including language that fulfilled the need for informed consent and confidentiality considerations based on examples from Creswell & Guetterman (2019). Respondents were given 10 days to complete the survey, after which time the survey was closed and no new responses were accepted.

Data was collected via an online survey using SurveyMonkey™, which included a series of open- and closed-ended questions. An online survey method was selected to allow greater time for coding and analysis of responses in two significant ways. First, by removing the need to transcribe a recorded interview or notes, more time was available for review and analysis of the data. Second, an online survey allowed participants to respond asynchronously during a window of time, thereby avoiding scheduling challenges and lengthy phone or online interviews. Beyond reasons of efficiency, this method also improved the credibility of the research through greater number and diversity of responses, compared to fewer but lengthier one-on-one interviews often used in qualitative research. There is a notable challenge with qualitative methodology and gathering greater numbers of responses, since greater volumes of information require more time for analysis to identify and interpret themes (Creswell, 2003; Creswell & Guetterman, 2019).

The survey began with an open-ended question asking the participants to provide their own interpretation and definition of a training environment or VTE, which was designed to help identify whether a common understanding of this concept existed amongst the participants. As noted previously, the survey did not include demographic questions such as sex or age, since it was distributed within a workplace environment in which these data points are considered sensitive or protected. Instead, the survey contained a question on current job title and a question on tenure in a training role.

In addition to the modified demographic questions, the survey also contained questions that began with a closed-ended question, followed by an open-ended question to enable more exploration of the answer, as demonstrated in Creswell and Guetterman (2019). For example, participants were asked to read a statement (e.g. “Having a training environment for practice during class helps participants be more confident after they’re back on the job.”), then used a Likert scale to indicate their relative level of agreement or disagreement with the statement. The subsequent question asked the participants to explain their response to the prior question in more detail. This was designed to help with narrowing the focus to key categories that were anticipated to emerge during the data analysis.

4.4. Data Analysis

After the survey concluded, the data was exported into a spreadsheet from the survey system. Each respondent was assigned a number (e.g. “Respondent 1”), which was used throughout the rest of the study to notate and track that individual’s response. The closed-ended questions were coded using Likert scale responses (e.g. 5 - Strongly agree, 4 – Agree, 3 - Neither agree nor disagree, 2 – Disagree, 1 - Strongly disagree), see Appendix C for full question and answer details. The open-ended questions were coded through reviewing the actual response text to determine common words/phrases and develop categories of ideas (Creswell, 2003; Creswell & Guetterman, 2019). After determining the broad categories apparent in the open-ended responses, major themes were developed and linked back to the research questions.

5. Results

The average time to complete the survey was 23 minutes, 50 seconds; the longest response time was 57 minutes for Respondent 1 and the shortest response time was 12 minutes for Respondent 3. Of the 13 participants invited to participate in the survey, six responded, for overall response rate of 46%. By job title, the highest response rate of 67% was among the Technical Service Trainers, and the lowest response rate of 20% among the Senior Trainers, with only one of the five responding (see Table 4). Technical Service Trainers are the job title most likely to train frequently on software and applications, and many of the tools they provide training on have VTEs so the higher response rate in this job title was not unexpected.

Table 4 – Response Rate by Job Title

Job Title	Number of Recipients	Number of Respondents	% Response
Field Technical & Safety Trainer	5	3	60%
Senior Field Technical & Safety Trainer	5	1	20%
Technical Service Trainer	3	2	67%
Total	13	6	46%

Of the six respondents, only one indicated that they had been in a training role greater than eight years, and the other five respondents indicated they had between two and five years in a training role. Participants were also asked to indicate how frequently they teach any software or applications in their classes, using a Likert scale (5 – Very Frequently, 4 – Frequently, 3 – Sometimes, 2 – Infrequently, 1 – Never). The mean score across all participants was 4, with only one participant (Respondent 2) indicating anything other than “Frequently” or “Very Frequently.” The respondents were asked to indicate whether the software or applications they currently provide training on have a VTE available or not. Four indicated that they teach some programs that do include a VTE and some that do not (indicated by “Both” in the Use VTEs column of Table 5) while the other two indicated that none of the software/applications they teach currently have a VTE (indicated by “No” in the Use VTEs column of Table 5). Participants were asked to rate the frequency of post-class support provided to trainees using a Likert scale (3 – Frequently, 2 – Sometimes, 1 – Never) and the mean score was 2.3, with all participants responding with either “Frequently” or “Sometimes”.

Table 5 – Demographic Data

Name	Job Title	Years in Role	Frequency of Software Training	Use VTEs	Post-Class Support
Respondent 1	Technical Service Trainer	>8	Very frequently	Both	Frequently
Respondent 2	Field Technical & Safety Trainer	2-5	Infrequently	No	Sometimes
Respondent 3	Field Technical & Safety Trainer	2-5	Frequently	No	Sometimes
Respondent 4	Sr Field Technical & Safety Trainer	2-5	Frequently	Both	Sometimes
Respondent 5	Field Technical & Safety Trainer	2-5	Very frequently	Both	Sometimes
Respondent 6	Technical Service Trainer	2-5	Frequently	Both	Frequently

Participants were asked to read two statements and indicate the extent to which they agreed or disagreed, using a Likert scale (5 – Strongly agree, 4 – Agree, 3 – Neither agree nor disagree, 2 – Disagree, 1 – Strongly disagree). For the statement “Having a training environment for practice during class helps participants be more confident after they’re back on the job” the mean response was 4.8, with only one participant responding “Agree” and all others responding “Strongly Agree” (see Figure 2).




Figure 2 - Question Responses – Trainee Confidence Related to Presence of VTEs

The same results were found for the statement “Having a training environment for practice during class helps participants be more proficient using the software after they’re back on the job,” with a mean of 4.8 and the same participant (Respondent 1) responding “Agree” and all others responding “Strongly Agree” (see Figure 3).




Figure 3 - Question Responses – Trainee Proficiency Related to Presence of VTE

5.1. Categories

While the closed-ended questions provided a clear foundation of the survey participants' perspectives on the efficacy of VTEs, significant information supporting the research questions was also found in the open-ended responses from the survey. In reviewing the text of these responses, four main categories were identified. These categories closely relate to the research questions and will be discussed here.

5.1.1. *Safe Environment*

The clearest and most common theme across the responses to the open-ended questions was the concept of safety within the learning experience. The first question in the survey asked the participants to provide their own definition of a VTE. Respondent 1, who is the most tenured trainer of those surveyed, defined a VTE as "a safe, controlled environment." Respondent 5 also included the phrase "safe environment" in their definition. Respondent 3 noted that a VTE provides the participant an opportunity to learn without the "consequences of a live environment," and Respondent 4 stated that having a VTE means that the participants do not have to be "scared" or "worry about blowing things up." Respondent 6 noted that a VTE provides an opportunity to "comfortably learn."

The theme of safety in the learning space emerged again in the responses to the open-ended question related to post-training confidence of the learner. Respondent 2 characterized a VTE as a "safe controlled environment" in their response, echoing Respondent 1. Respondent 3 indicated that having an opportunity to use a VTE helps alleviate the "fear of repercussions or breaking something." Respondent 4 noted that having a VTE provides a setting where learners have "little fear of ruining things or screwing up."

5.1.2. Learning from Mistakes

The value of making mistakes during the learning process was another clear pattern that emerged in the participant responses. Respondent 5 included this in their definition, stating that a VTE is an environment where learners can “play with and learn to control said program and make mistakes.” Respondent 2 noted that the positive impact to learner confidence comes from “making mistakes and learn[ing] from them rather than [sic] affecting a real account.” Regarding the concept of building proficiency through use of the VTE, Respondent 2 stated that the “inability to make mistakes in a controlled environment will hinder the learning experience because people tend to learn more from mistakes than [sic] successes.” Respondent 5 also indicated the value not only of being afforded an opportunity to make mistakes during the learning, but also that a VTE can allow the opportunity to learn how to correct mistakes, as follows:

“They also need to know how to fix what’s wrong. Especially if the mistake was their fault. Having a place to show them those errors and how to fix them will allow us to teach how to not only how to do that, but make them good at it before they reach the live application.”

5.1.3. Business Impact

Another category identified in the analysis of the open-ended responses relates to the topic of impact to the business. Respondent 6 noted in their definition that a VTE allows learners to practice navigation in the software “without impact to the business unit,” and Respondent 1 specifically called out that use of a VTE avoids “negative customer impact” by better preparing the learners for their job. Three of the respondents included statements related to minimizing the effect on the live environment/production environment in their response to participant confidence back on the job.

5.1.4. Confidence

The category of confidence with use of the tool was a major element found in many of the open-ended responses. The respondents noted that the ability for trainees to practice in a realistic simulation of the live environment and have repetitive activities was key to building their confidence and skills during training (Respondent 4 characterized it as “muscle memory”), with the word “comfort” or “comfortable” being used six times across multiple respondents and “confident” or “confidence” used four times. Respondent 5 noted that “operating one of our programs requires confidence. Having a training platform to build that confidence and make mistakes will allow the employee to become familiar and comfortable in that program.” Respondent 6 summed it up succinctly:

“For programs that we currently have a training environment for, I’ve noticed that my trainees are more confident going into their job duties. They are more comfortable with navigation and more willing to attempt job tasks that they may not be as familiar with. For programs that we do not have a training environment to use, the trainees are typically less comfortable heading into their job duties. I usually do not see them interacting with the tool as comfortably, nor using it unless they are specifically directed to and someone is there to work with them as they execute the task.”

6. Discussion

The overwhelming pattern that emerged from the responses of the survey participants to both the closed- and open-ended questions was a consensus on the value and benefit of the use of VTEs in software and application training. Even those survey participants who indicated that they teach software less frequently, or do not teach software that currently offers a VTE for training were clear on this point. The research questions related to the post-training confidence level of trainees as observed by the trainers, and

the responses to the closed-ended questions showed a clear preference for the presence of a VTE, with a mean score of 4.8 on both proficiency and confidence. Beyond just a close match to the research questions, responses to the open-ended questions contained the same concepts and ideas as those found in the literature, including the concepts of scaffolding, authenticity of the learning environment, and self-direction or exploration.

6.1. Authenticity of Learning Environment and Learning Transfer

Hardré (2003) noted that the authenticity of the learning environment vis-à-vis the actual tasks or work impacts the effectiveness of software training videos or e-learning, and several respondents alluded to this in their responses. Respondent 4 stated that by using a VTE “it’ll look familiar to them in the real world when they get out there . . . they won’t be completely lost when they see a screen that looks nothing like the old material pictures that were shown in training.” Respondent 6, who is one of the more frequent users of VTEs also noted that the VTE must be as close to the real world as possible, and that “if the training environment is too different from the live system, it can create more confusion than good.” This aligns closely with Hardré’s findings that authenticity of representation and authenticity of interactivity are critical to the ultimate success of the learning experience.

Foley and Kaiser (2013) highlighted the importance that extensive practice of skills in the learning environment ensures that replication is nearly automatic in the new experience, and the survey participants noted this in their responses as well. Respondent 4 noted “the more they practice with something, the better they will be,” and Respondent 6 stated “they are more willing to utilize the skills and translate them into live functions” if they have access to a VTE. The concept of scaffolding as described by Foley and Kaiser (2013) is also present in the responses to how the trainers teach the software today, with several respondents providing detailed descriptions of processes by which they build a series of practice exercises and scenarios whereby the learner is presented with increasingly complex tasks while receiving support and coaching from the instructor.

6.2. Learning from Mistakes and Safe Learning through Self-Direction

In much of the literature, the concepts of self-directed learning and the ability for learners to gain confidence and skills through an open-ended, exploratory environment in which they can safely make mistakes was key to success (Bridget et al., 2017, Hardin et al., 2013, Inguva et al., 2018, Merriam & Bierema, 2014, Martocchio & Webster, 1992). This concept was found throughout the open-ended responses and was one of the main categories identified in the data analysis. The consensus of the survey respondents was that without an opportunity for learners to explore, try things, and fail in a safe environment, they are not as effective or confident with the program after class. A key component of the safety of the learning environment as noted by the respondents was that the learners were able to practice skills without the risk or fear of negatively impacting customers or the business in a production environment of the program.

6.3. Study Limitations

There are at least two potential limitations related to the scope of this research. A first limitation relates to the number of respondents to the survey, which represents only a small sample of the overall population of trainers at the specific organization. While the purposeful selection of this group allowed us to gain a greater understanding of the phenomenon at the center of our inquiry, seeking insights from a larger group could increase the diversity of perspectives and create a deeper understanding of the impact of VTEs on workplace learning. A second potential limitation concerns the selection of exclusively trainers for this survey. By using the concept of triangulation and gaining perspectives from different

sources (e.g. supervisors and/or the trainees themselves) the validity of the data analysis could be improved.

Despite these limitations, the results suggest a clear connection between the presence of a VTE and trainee performance. The implication to organizations that are engaged in workplace software and application training is that the presence of a VTE can be the differentiating factor between confident, correct use of the software post-training, or potential business- or customer-impacting mistakes and errors.

6.4. Future Research

Although the findings of this particular study support the value of VTEs in workplace learning, the most important contribution may be that it raises awareness of this instructional method and creates an opportunity for further inquiry. If, as this study suggests, the presence of a VTE significantly improves both the learning experience itself and the learner performance post-training, organizations seeking to improve accuracy in software usage would be wise to continue the line of inquiry begun in this research. A recommendation would be to consider a quantitative study to better understand the actual performance measures of trainees who are afforded access to VTEs during training, and compare these to the performance of those who do not have a VTE. Alternatively, mixed-method research combining both the personal feedback of individuals and focus groups combined with quantitative performance measures may provide more robust insight into the topic. As noted earlier, organizations will require a clearly-articulated business case to justify the potential investment needed to design and deploy VTEs for their software solutions; and further research by learning professionals will be critical to help shape the future of software instruction in the workplace.

7. Conclusion

The findings of this research show a pattern of trainer perception that the presence of a VTE during software and application training provides significant positive impact to the learners' experience, and that the absence of one can be equally detrimental. Concepts from the literature, including exploratory and participatory learning, self-directed learning, learning transfer and others were found throughout the participant responses, indicating that respondents have a strong sense of what works (or doesn't work) in the learning environment based on their own experience as professional educators of adults. Although future study will be needed to further explore this topic, the present study has enhanced the understanding of the relationship between virtual training environments for software training and trainee confidence and provided clear support for the value of this instructional method.

Abbreviations

VTE	Virtual Training Environment
SDL	Self-Directed Learning
K-12	Kindergarten through twelfth grade

Bibliography & References

- Abdullah, M., Koren, S., Muniapan, B., Parasuraman, B., & Rathakrishnan, B. (2008). Adult participation in self-directed learning programs. *International Education Studies*, 1(3), 66–72.
<https://doi.org/10.5539/ies.v1n3p66>
- Bridge, P., Giles, E., Williams, A., Boejen, A., Appleyard, R. & Kirby, M. (2017). International audit of virtual environment for radiotherapy training usage. *Journal of Radiotherapy in Practice*, 16(4), 375–382. <https://doi.org/10.1017/S146039691700022X>
- Caruso, S. (2018). Toward understanding the role of Web 2.0 technology in self-directed learning and job performance. *Contemporary Issues in Education Research*, 11(3), 89–98.
<https://doi.org/10.19030/cier.v11i3.10180>
- Creswell, J. W. (2003). *Research Design: Qualitative, quantitative, and mixed methods approaches* (2nd ed.). Sage Publications, Inc.
- Creswell, J. W. (2013, February 19). *What is mixed methods research?* [Video]. YouTube. <https://www.youtube.com/watch?v=1OaNiTIpyX8>
- Creswell, J. W. & Guetterman, T. C. (2019). *Educational research: Planning, conducting, and evaluating quantitative and qualitative research*. Pearson Education, Inc.
- Fleming, D., Artis, A., & Hawes, J. (2014). Technology perceptions in employees' use of self-directed learning. *Journal Of Services Marketing*, 28(1), 50–59. <https://doi.org/10.1108/JSM-03-2012-0062>
- Hardré, P. L. (2013). Considering components, types, and degrees of authenticity in designing technology to support transfer. *New Directions for Adult and Continuing Education*, 2013(137), 39-47.
<https://doi.org/10.1002/ace.20043>
- Hendriks, S., Sung, S., & Poell, R. (2018). Learning paths of customer-facing professionals in the digital age. *Journal of Workplace Learning*, 30(5), 377–392. <https://doi.org/10.1108/JWL-01-2018-0023>
- Kearsley, G., & Seidel, R. J. (1985). Automation in training and education. *Human Factors*, 27(1), 61-74.
- Lavendels, J., Krischuk, M., Sitikovs, V., & Bulins, Z. (2014). Virtual observation environment for training and monitoring of insurance software end-users. *Applied Computer Systems*, 15(1), 32–35. <https://doi.org/10.2478/acss-2014-0005>
- Marler, J. H., Liang, X., & Dulebohn, J. H. (2006). Training and effective employee information technology use. *Journal of Management*, 32(5), 721-743.
<https://doi.org/10.1177/0149206306292388>
- Martocchio, J. J. & Webster, J. (1992). Effects of feedback and cognitive playfulness on performance in microcomputer software training. *Personnel Psychology*, 45(3), 553–578.
<https://doi.org/10.1111/j.1744-6570.1992.tb00860.x>
- Mayer, R. (2003). The promise of multimedia learning: Using the same instructional design methods across different media. *Learning and Instruction*, 13(2), 125–139. [https://doi.org/10.1016/S0959-4752\(02\)00016-6](https://doi.org/10.1016/S0959-4752(02)00016-6)

- Merriam, S., & Bierema, L. (2014). *Adult learning: Linking theory and practice*. John Wiley & Sons, Incorporated.
- Miller, S. (1973). Ends, means, and galumphing: Some leitmotifs of play. *American Anthropologist*, 75(1), 87–98. <https://doi.org/10.1525/aa.1973.75.1.02a00050>
- Roumell, E. A. (2019). Priming adult learners for learning transfer: Beyond content and delivery. *Adult Learning*, 30(1), 15-22. <https://doi.org/10.1177/1045159518791281>
- Turnage, J. (1990). The challenge of new workplace technology for psychology. *The American Psychologist*, 45(2), 171–178. <https://doi.org/10.1037/0003-066X.45.2.171>
- United States Department of Commerce & National Telecommunications and Information Administration. (2002). *A nation online: How Americans are expanding their use of the Internet*. <https://www.ntia.doc.gov/report/2002/nation-online-internet-use-america>
- United States Department of Labor. (2005). *Computer and internet use at work in 2003*. <https://www.bls.gov/news.release/ciuaw.toc.htm>
- Van der Meij, H. (2013a). Do pedagogical agents enhance software training? *Human-Computer Interaction*, 28(6), 518–547. <https://doi.org/10.1080/07370024.2013.789348>
- Van der Meij, H. (2013b). Motivating agents in software tutorials. *Computers in Human Behavior*, 29(3), 845–857. <https://doi.org/10.1016/j.chb.2012.10.018>
- Van der Meij, H. & van der Meij, J. (2013). Eight guidelines for the design of instructional videos for software training. *Technical Communication*, 60(3), 205-228.
- Van der Meij, H. (2014). Developing and testing a video tutorial for software training. *Technical Communication*, 61(2), 110–122.
- Van der Meij, H., van der Meij, J., Voerman, T., & Duipmans, E. (2018). Supporting motivation, task performance and retention in video tutorials for software training. *Educational Technology Research and Development*, 66(3), 597–614. <https://doi.org/10.1007/s11423-017-9560-z>



Creating Infinite
Possibilities.

Test Environments and Methods for Validation of DOCSIS 4 Devices

Jon-En Wang

Sr. Director, Physical & Environmental Product Evaluations
Comcast


Jon-en_wang@comcast.com

Introduction

The Comcast Physical and Environmental team is responsible for hardware evaluation for deployment into Comcast.

Our main functions are

- Research
- Product evaluation and change management
- Root cause analysis
- Standards development (CMD)
- Quality verification assurance




Our technology environments




**ELECTROMAGNETIC
COMPATIBILITY**



MECHANICAL



**MATERIAL
ANALYSIS**



**END-TO-END
ENVIRONMENTAL**

DOCSIS 4.0 implementations

Extended spectrum DOCSIS

- Downstream spectrum extended to 1.8 GHz
- Attenuation increases with frequency

Full duplex DOCSIS

- Upstream spectrum extended to 684 MHz
- Echo cancellation mitigates interference between transmission and reception


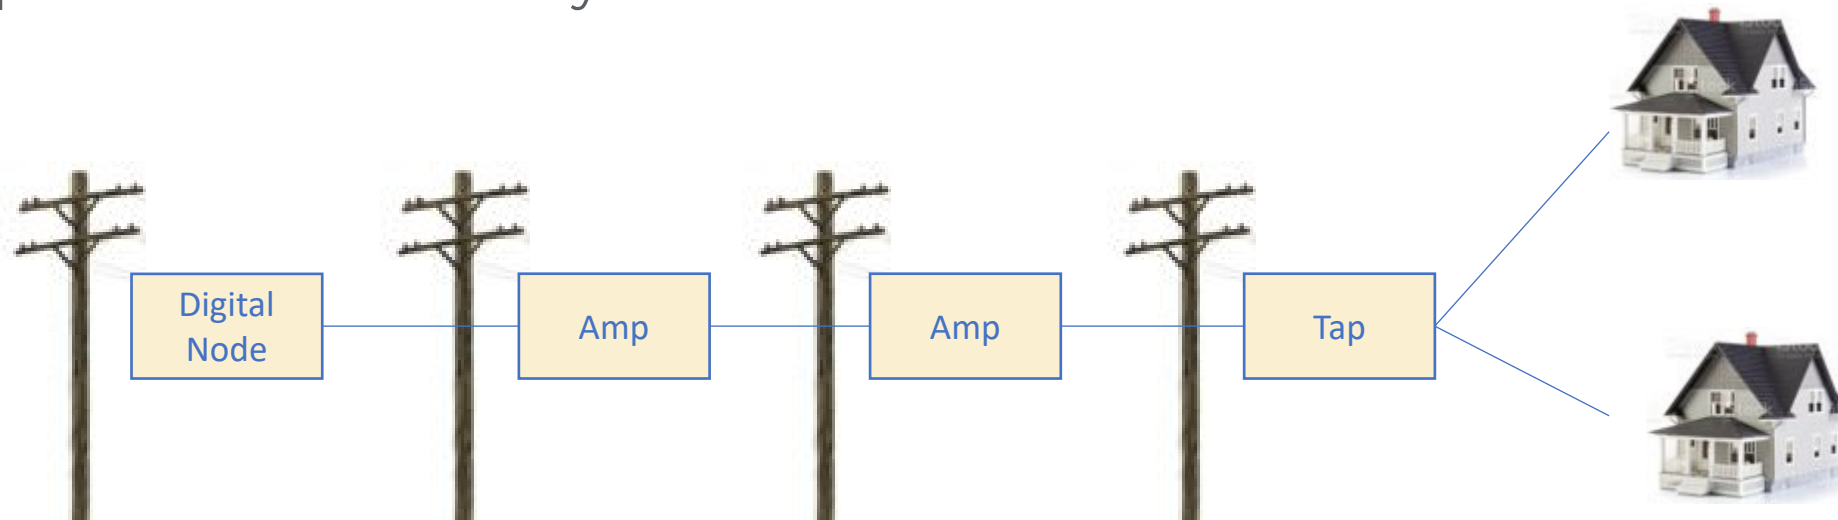


Figure 25 - Configurable FDD Upstream Allocated Spectrum Bandwidths

Source: CableLabs Data-Over-Cable Service Interface Specifications DOCSIS® 4.0

Our Environment

- Variety of networks due to geography, power, and density
- Cascade of plant components (nodes, amplifiers, taps, cables, connectors) contribute to RF performance
- Impairments from craftsmanship or other causes
- Temperature and humidity



Dense MDU - high density




Dense MDU

Frequency Domain - Single Tap

Frequency domain measurement shows return loss degrades as taps are added

Time element needs to be measured differently




Dense MDU

Time Domain – Individual Taps

Time domain measurement shows 2 reflections from each tap, corresponding to the interface between the connector and the input and output ports of the tap

Each tap was the same body style, resulting in the same travel time for the reflected signal from each interface




Dense MDU

Time Domain –Cluster of 4 Taps

Time domain measurement shows additional reflections at longer times as taps are added.

The added reflection times correlate to the extra path lengths for the reflected signal to travel.





Dense MDU

Frequency Domain - Multiple Tap Clusters

Frequency domain measurement shows return loss degrades as taps are added

Time element needs to be measured differently





Dense MDU

Time Domain – Multiple Tap Clusters

Second cluster of taps added with ~80 feet of hardline cable in between tap clusters.

Calculated round trip delay based on RF signal propagation over 80 feet of cable is 187 nanoseconds.

Return loss of the second cluster of taps is attenuated by coaxial cable as well as tap's through port loss.



Summary

Most contribution to reflection amplitude is from the closest tap cluster

Time domain analysis shows that reflections from multiple taps in a cluster, and from multiple clusters are important and must be recreated

Complete plant models can be recreated in environmental chambers for testing to outside plant conditions





Creating Infinite Possibilities.

Thank You!

Jon-En Wang

Sr. Director, Physical & Environmental Product Evaluations
Comcast
Jon-en_wang@comcast.com

Testing Wi-Fi Upgrades for Latency and Throughput

Evaluating OFDMA for Latency Improvements

A Technical Paper prepared for SCTE by

Matt Brooks

Distinguished Software Engineer
CommScope
3871 Lakefield Dr. Suwanee, GA 30024
Matt.Brooks@commscope.com

David Williams

Senior Staff Software Engineer
CommScope
3871 Lakefield Dr. Suwanee, GA 30024
David.Williams@commscope.com

Table of Contents

Title	Page Number
1. Abstract	4
2. Introduction.....	4
3. Latency Overview.....	5
3.1. Latency Values and Test Techniques.....	6
4. OFDMA Overview	7
4.1. OFDMA: Resource Units.....	8
4.2. OFDMA: Process	9
5. Test House Setup.....	10
6. Test Tools.....	10
7. Test Methodology.....	13
7.1. Test Constants	13
7.2. AP Settings for Comparison.....	14
7.3. Client Conditions for Comparison	14
8. Test Results	15
8.1. Non-Rate Limited Test Scenarios	17
8.2. Rate-Limited Test Scenarios with Channel Utilization Set to 90%	22
8.3. Rate-Limited Test Scenarios with Channel Utilization Set to 50%	25
8.4. Rate-Limited Test Scenarios with Channel Utilization Set to 10%	28
8.5. Rate-Limited Test Scenarios with LAN Hosted Gaming	32
8.6. Wi-Fi 6E Testing with Rate-Limited Test Scenarios.....	35
8.7. Wi-Fi 6E Rate-Limited Test Scenarios with LAN Hosted Gaming	37
9. Additional Tests to Consider	38
10. Upgrading to Wi-Fi 6E and enabling OFDMA for Wi-Fi 6	38
11. Conclusion.....	39
Abbreviations	42
Bibliography & References.....	43

List of Figures

Figure 1 – Example OFDMA Tone Allocation with HE SU vs. HE OFDMA.....	9
Figure 2 – Example iperf3 JSON to python3 Dictionary for a Single Flow on One Client	12
Figure 3 – Example Ping Data Received During a Test	12
Figure 4 – Counter Strike Global Offensive Screen Overlay with RTT Latency	13
Figure 5 – UDP Upstream No Rate Limits	17
Figure 6 – Per Ping Sample OFDMA Disabled vs. Enabled.....	18
Figure 7 – Per Ping Sample Wi-Fi 5 vs. OFDMA Disabled/Enabled with half Wi-Fi 5 Clients	19
Figure 8 – TCP Downstream No Rate Limits.....	20
Figure 9 – Per Ping Sample Wi-Fi 5 vs. OFDMA Disabled/Enabled with half Wi-Fi 5 Clients	21
Figure 10 – TCP Upstream, 90% Limit Others, CUT Limit 50 Mbps	22
Figure 11 – Per Ping Sample OFDMA Disabled vs. Enabled.....	23
Figure 12 – TCP Downstream, 90% Limit Others, CUT Limit 500 Mbps – AP 2.....	24
Figure 13 – TCP Downstream, 50% Limit Others, CUT Limit 500 Mbps.....	25
Figure 14 – UDP Downstream, 50% Limit Others, CUT Limit 50 Mbps	26
Figure 15 – Per Ping Sample OFDMA Disabled vs. Enabled.....	27

Figure 16 – TCP Upstream, 10% Limit Others, CUT Limit 500 Mbps	28
Figure 17 – Per Ping Sample OFDMA Disabled vs. Enabled with half Wi-Fi 5 Clients	29
Figure 18 – TCP Downstream, 10% Limit Others, CUT Limit 5 Mbps – AP 1	30
Figure 19 – TCP Downstream, 10% Limit Others, CUT Limit 5 Mbps – AP 2	31
Figure 20 – TCP Upstream, 90% Limit Others, CUT Limit 1 Mbps + Game Play	32
Figure 21 – CS:GO Reported RTT Latency Over 1 Minute of Game Play	33
Figure 22 – Per Ping Sample OFDMA Disabled vs. Enabled During CS:GO Game Play	34
Figure 23 – TCP Upstream, 90% Limit Others, CUT Limit 50 Mbps – AP 3 with Wi-Fi 6E	35
Figure 24 – Per Ping Sample OFDMA Disabled vs. Enabled – AP 3 with Wi-Fi 6E	36
Figure 25 – CS:GO Reported RTT Latency Over 1 Minute of Game Play – Wi-Fi 6E	37
Figure 26 – Latency Change Observed with OFDMA and Downstream Traffic	40
Figure 27 – Latency Change Observed with OFDMA and Upstream Traffic	40

List of Tables

Title	Page Number
Table 1 – RTT Latency Added from WAN	4
Table 2 - Common Application Latency Expectations for RTT Latency	7

1. Abstract

Subscriber satisfaction is increasingly about delivering services with the right latency and throughput characteristics. So, how does Orthogonal Frequency-Division Multiple Access (OFDMA) technology deliver this in the home? This paper analyzes empirical testing within a home using typical Wi-Fi clients, focusing on latency and throughput with real-world traffic patterns.

OFDMA in Wi-Fi 6 and Wi-Fi 6E allows for simultaneous transmissions on all spatial streams for each device and subsequently results in lower latency and less contention of airtime. When using both Wi-Fi 6 clients and legacy clients with OFDMA enabled, the latency and jitter improvements are still seen. Wi-Fi 6E access points and clients fully realize OFDMA latency reduction with 6 GHz greenfield spectrum.

The results of testing empirically focus on latency, throughput, and application performance for each Wi-Fi standard and OFDMA setting and showcase use cases of traffic patterns in a house with different client utilizations. By showing latency for a given client while changing only the access point mode and OFDMA settings, a latency reduction is realized in homogeneous client environments as well as mixed client environments. This paper will explain a method for evaluating OFDMA without expensive test equipment or ideal lab setups while providing decision points for when to invest in Wi-Fi 6 or Wi-Fi 6E access points and when to expect an improvement from enabling OFDMA.

2. Introduction

Subscribers are increasingly dissatisfied with their quality of experience (QoE), using programs and applications that have sufficient bandwidth in the home yet are plagued by a problem they do not fully understand: too much latency. Multiple System Operators (MSOs) traditionally sell speed or bandwidth tiers but not latency tiers. Much of the latency conversation is traditionally regarding the wide area network (WAN) access layer to the internet, whether this is fiber, Data Over Cable Service Interface Specifications (DOCSIS), digital subscriber line (DSL), or Satellite. However, the Wi-Fi connection from a gateway/access point (AP) is a common medium by which users access the internet regardless of the WAN link being used.

Using a Wi-Fi 6 or Wi-Fi 6E AP, with OFDMA support, along with devices that also support Wi-Fi 6 or Wi-Fi 6E, allow for a lower latency access to the shared Wi-Fi medium for a more responsive experience during congested or multi-client situations. Improving one layer of the larger network can have a dramatic increase in the QoE of the end-user. All latency values reported in this paper were tested without a contribution from WAN latency [3] referred to in Table 1, which should be considered in addition to values reported in the test results of this paper.

Table 1 – RTT Latency Added from WAN

Last Mile Connection	Latency Contribution
Fiber	10-20 ms
DOCSIS	15-40 ms
DSL	30-65 ms
Satellite	45-500 ms

OFDMA can be used between an AP and client that both support OFDMA. Improvements in latency are possible in mixed legacy client populations by allowing OFDMA capable clients to have some improved latency while less capable clients are still using legacy Wi-Fi standards and continue to tie up access to the whole channel for part of the time. OFDMA is not a feature that is meant to increase throughput and, in many scenarios, can decrease speeds in the current generation of chipsets and software. Decisions on

grouping OFDMA capable clients and other scheduler decisions to prioritize speed or latency over one another, as well as serving legacy clients, can change the total throughput seen with OFDMA enabled, negatively impact latency, or cause variation in test results from run to run. In some cases, the data will demonstrate latency is improved, while in other cases it is about the same or worse.

A real house with OFDMA capable clients using a common Wi-Fi client chipset was used to run a few APs through tests to show the relative difference for the same channel but with different OFDMA settings and AP modes used. While changing the wireless mode between Wi-Fi 5, Wi-Fi 6, and Wi-Fi 6E, several variables were tested. This included OFDMA enabled and disabled, different traffic patterns to a client under test, different channel utilizations from other clients, different packet sizes, and different protocols in the downstream and upstream directions while measuring RTT latency on the client under test. Total throughput was also observed, and although Wi-Fi 6 and 6E increased the MCS data rates to support 1024 QAM, as well as allowed better usage of the spectrum with smaller subcarrier width and more efficient use of the spectrum for data subcarriers, throughput was not a focus of the paper.

Results in this paper will show that as you increase the channel utilization or the throughput needs of the client under test, both conditions are more conducive to showing a decrease in latency if OFDMA is enabled. The reduction in average round trip time (RTT) latency is sometimes seen but more often a dramatic decrease in the maximum RTT latency measured is observed with OFDMA enabled. This is also accompanied with a lower median deviation of RTT or variance of the latency known as jitter. This means that even in cases where the average RTT is about the same, the maximum and median deviation RTT are often meaningfully lower providing a smoother QoE for the end-user with applications less likely to encounter latency spikes that cause issues with the QoE.

3. Latency Overview

Much effort is put into emphasizing network throughput achieved or bandwidth available, but latency is also one of the most important and often overlooked metrics. A consistent and good QoE reduces churn for MSOs, and latency is a main contributor to QoE. A user may be able to achieve the same throughput, that is data over time, on different networks, but one network could have less latency and be more responsive and pleasant to use.

When discussing latency, it's important to define terminology. This paper will discuss round trip time (RTT) vs. one-way delay (OWD), as well as the difference between idle pings and pings under load, and lastly a consideration for RTT of each transmission control protocol (TCP) stream or session vs. RTT of a ping to and from the source and destination.

Latency means different things to different people even in the networking industry. End-users outside the networking industry are that much further removed from knowing what is good or bad and may blame the wrong metric for a poor user experience. Latency can be defined simply as the amount of time for information or a packet to arrive from a source to a destination and is comprised of components such as: propagation delay (time required for the signal to travel over the medium), transmission delay (time required to push the bits into the link), processing delay (time required to process packet headers), and queuing delay (amount of time a packet is waiting in queue to be processed and gain access to the physical medium) [1]. This can be defined over a subset of network segments or a full network delay of a packet going to the destination over multiple hops locally or through the internet. This delay is also different under a working load vs. while idle. Latency observed in an idle network or idle Wi-Fi channel does not correlate with the user experience on the same network under a different load. On the wireless medium, latency is an especially fluid number changing constantly based on traffic passed, the number of clients accessing the shared medium, un-coordinated channel collisions, distance, and retransmissions.

Latency can be defined by the time it takes for a packet to be sent from a source and received at a destination; this is referred to as OWD. OWD measurements are usually accomplished with user datagram protocol (UDP) and require clock synchronization between source and destination to measure the elapsed time without having feedback to the source's clock. A timestamp is added in the data of the outgoing packet which is measured against the destination's clock when it is received. For this method of OWD measurements to be accurate, the same synchronized time must be used on source and destination devices. One such tool that uses this approach is an application called NUTTCP when used with UDP.

Latency can also be defined by timing how long it takes for a packet to go from source to destination and back to source with a response. This method is useful as well because most packets in networks are going to a destination to solicit a response back to where it came from. This may not be latency in its purest sense, but it is a total latency that most users would experience with a real application. This full elapsed time is measured and referred to as round trip time (RTT). Internet control message protocol (ICMP) pings and TCP RTT are common ways of measuring RTT. Pings over long multi-hop connections are not necessarily a true reflection of the latency in a loaded network, because network equipment could choose to prioritize or de-prioritize handling ICMP pings or cached information in routers could speed up subsequent pings. However, on a local network with a single hop, pings are still an essential tool to measure RTT delay. Iperf3, a networking throughput testing tool, can also report RTT per TCP flow while data is being sent if sourcing from a Linux Ethernet port.

3.1. Latency Values and Test Techniques

Considering only the instantaneous or just the average latency measured is not enough to characterize QoE of the end-user in a real wireless network. In addition to the average values of RTT or OWD, consideration of the minimum, maximum, and jitter, or a measure of the variability of the latency reading during a measurement period, is necessary. Sometimes, the average or minimum latencies won't change much in test scenarios, but the maximum latency observed can be far higher than the average and responsible for problems in user experience [1]. Many applications can handle a higher consistent latency, but if a packet suddenly has a much higher delay than average, this can cause extremely noticeable problems depending on the application being used. If an application can buffer without the user noticing, such as video playback, the end-user may not notice high jitter or higher latencies as easily.

Consider the difference in effect of high latency packets while watching a movie that buffers a significant amount, a voice call or video call that buffers very little, a virtual reality headset tracking a person's movements, or online game play that experiences a sudden skip in graphics displayed or a delayed reception of an action of a player. General expectations for good and bad latency values per application type are found in Table 2 below [4,5,6]. This includes full network round trip time latency including WAN, while the values in this paper are the latencies inside the house with Wi-Fi and do not include a WAN component.

Table 2 - Common Application Latency Expectations for RTT Latency

Application	Excellent	Good	Fair	Bad	Recommended
Gaming (FPS)	<20 ms	20-50 ms	50-150 ms	150+ ms	x
Gaming (MMO)	<20 ms	20-100 ms	100-250 ms	250+ ms	x
Gaming (RTS)	<20 ms	20-100 ms	100-200 ms	200+ ms	x
Cloud Gaming *	<20 ms	20-30 ms	30-40 ms	40+ ms	x
Voice Call	x	x	x	x	<100 ms
Video Call	x	x	x	x	<100 ms
Video Stream	x	x	x	x	<100 ms

* Latency requirements for cloud gaming are more stringent due to control input

A loaded network, especially in wireless networking, will have increased delays when traffic is being sent and while serving multiple wireless clients on the shared medium. Enhanced distributed channel access (EDCA) provides the mechanism for how backoffs are observed between different APs and clients trying to access the same channel. The more clients attempting to use the channel the greater the chance for one of them to need to backoff and wait even longer before winning an opportunity to transmit on the channel.

A network with traffic during tests or pings is referred to as experiencing load or being loaded. For example, a Wi-Fi 5 network with just 4 clients and very low utilization on the channel of 10% can still achieve 2 to 5 ms RTT pings with a very low 0.5 Mbps iperf3 data flow in the downstream direction. However, the same very low 0.5 Mbps TCP data flow to a Wi-Fi 5 client while the channel is 90% loaded, causes the same flow to incur an average of 19 to 24 ms RTT pings, up to a max of 187 to 198 ms possibly. A loaded network's delay is a more accurate view of the worst-case latency experienced by an end-user and is a preferred situation to characterize latency improvements. Loaded networks have more delay in queues/buffers, from backoff timers for fair access of other clients, and from retransmissions. Without load, it is more difficult to make discernable differences in outcome or QoE with different technology or settings when later used in a real environment.

In Wi-Fi channels, for this paper, the load refers to channel utilization which does not refer to Mbps. Channel utilization could be very high, even for a very low Mbps being transferred, if a slower, more robust MCS rate is being used for a client that is far away. Referring to load as airtime utilization abstracts away problems in test setups that specify loads as Mbps which are not the same anywhere else. Scheduler algorithms in AP chipsets are proprietary and have thresholds for considering usage of OFDMA to begin with, and idle or low usage clients are often not allocated any resource units (RUs). In this paper, many channel utilizations and throughput levels to a client under test were characterized.

The above methods for determining latency are simplifying what can be complicated. The absolute RTT or absolute OWD may change with a different number of TCP flows or different TCP window sizes, better time synchronized clocks, and with special treatment of ICMP pings by each operating system. In addition, if separate Ethernet links were used to source to each client, further differences may be observed. However, the relative differences between AP Wi-Fi modes and OFDMA settings, while keeping test methods and conditions the same accompanied with automation of test execution, allows for compelling comparisons.

4. OFDMA Overview

The primary feature being tested empirically in this paper is OFDMA. This technology first available for 802.11 networks in the Wi-Fi 6 standard allows for simultaneous transmissions to and from clients in the same channels on different subcarriers. This contrasts with previous generations of Wi-Fi standards,

which are for the most part round robin transmissions. Multi-user multiple input multiple output (MU-MIMO) in Wi-Fi 5 was able to achieve simultaneous transmissions, but it was not able to predictably realize gain in many real environments and often causes too much self-degradation and overhead to be beneficial especially with more than 3 clients. It did however improve latency in many cases even with increased retransmissions or lower throughput as client counts increased. To isolate contributions of OFDMA to latency, MU-MIMO was disabled for very high throughput (VHT) and high-efficiency (HE) modes in downlink and uplink directions.

OFDMA is often presented in settings of the web user interface of an AP with control separated to enable downlink OFDMA, that is AP to client, as well as uplink OFDMA, that is client to AP. For this paper both downlink and uplink OFDMA will be enabled or disabled together. The AP scheduler, proprietary in each chipset's firmware, is responsible for determining what technology should be used or which groups to form for simultaneous transmissions. For each transmit opportunity (TXOP), the AP will choose if the transmission will be single user (SU) legacy orthogonal frequency-division multiplexing (OFDM) traffic, synchronized multi-user OFDMA uplink traffic or synchronized multi-user OFDMA downlink traffic. The OFDMA modes are referred to as high-efficiency multiple user (HE-MU). This is because of the synchronized aspect of OFDMA transmissions, which drives efficiency, instead of legacy carrier sense multiple access with collision avoidance (CSMA/CA). Each OFDMA capable device is assigned a subset of subcarriers, or tones, by the AP for that device to simultaneously transmit on at the same time as others do on different subcarriers in the same channel.

With OFDM in Wi-Fi standards prior to Wi-Fi 6, the subcarrier spacing is 312.5 KHz, however, with Wi-Fi 6 and OFDMA the subcarrier spacing is 78.125 KHz. The subcarriers are 4 times closer together in OFDMA, and the symbol time is increased to be 4 times longer from 3.2 microseconds to 12.8 microseconds. The reduction in size of each subcarrier also results in efficiency gains in the channel itself because of less spectrum being used for pilot subcarriers and null guard carriers. The trigger frame from the AP to the clients specifies which RUs are allocated for simultaneous uplink transmission. In Wi-Fi 6 on 2.4 GHz and 5 GHz bands, the data exchanges can occur with OFDMA, however, many control and management frames must still use legacy OFDM to hold off and notify legacy clients about the channel being used [2]. The exception is for HE specific control frames such as buffer status report (BSR), clear-to-send (CTS), and block acknowledgements (ACK) which can occur simultaneously on RUs that are assigned.

4.1. OFDMA: Resource Units

Resource units (RUs) are groups of OFDM subcarriers, referred to as tones, and are predefined in these allocations: 26, 52, 106, 242, 484, 996, or 2x996 tones. The location of the RU within the channel is further defined by an RU index contained in a trigger frame that lets the client know what part of the spectrum is set aside for it in the single TXOP. Most current generation APs use a 242 tone RU as the smallest RU allocation in the 5 GHz and 6 GHz bands when 80 MHz channels are in use. A total of four 242 tone RU assignments can exist in a single group if the channel is 80 MHz. The 2.4 GHz band was not considered in this paper, but naturally uses smaller than 242 tone RUs to achieve multi-client OFDMA transmissions in 20 MHz and 40 MHz channels.

The AP can create additional groups to simultaneously transmit or receive from different sets of clients using the same RUs but in different groups, and therefore at different times. Multiple groups using OFDMA within each group are still advantageous to a certain point, but it does reduce the benefits of OFDMA and a study on multiple groups and its effect on latency reduction is left for further study. There would be a certain point when clients have unequal bandwidth needs in a specific TXOP that warrant scheduling a single group with lower RU designations than 242 tones. However, currently 242 tone RUs

are preferred most often in simultaneous 4 client scenarios as the AP doesn't have to allocate extra null guard carriers or pilots in the spectrum as it does with smaller 106 tone RU assignments. When 106 tone RUs are assigned to be used in an 80 MHz channel, there are 8 to 13 total clients supported in the group if the AP allocates 26 tone RUs to 5 of the clients. Some APs are not allocating smaller RUs with larger RUs in the same group and would leave gaps in the spectrum unused. However, using RU assignments of less than 242 tones is more advantageous for latency reduction with many clients at the same time. In Figure 1 below, the difference between single-user and multi-user OFDMA is shown with respect to time; one or more clients will be allocated tones across an 80 MHz channel. This example shows 996, 484, and 242 tones being used over time.

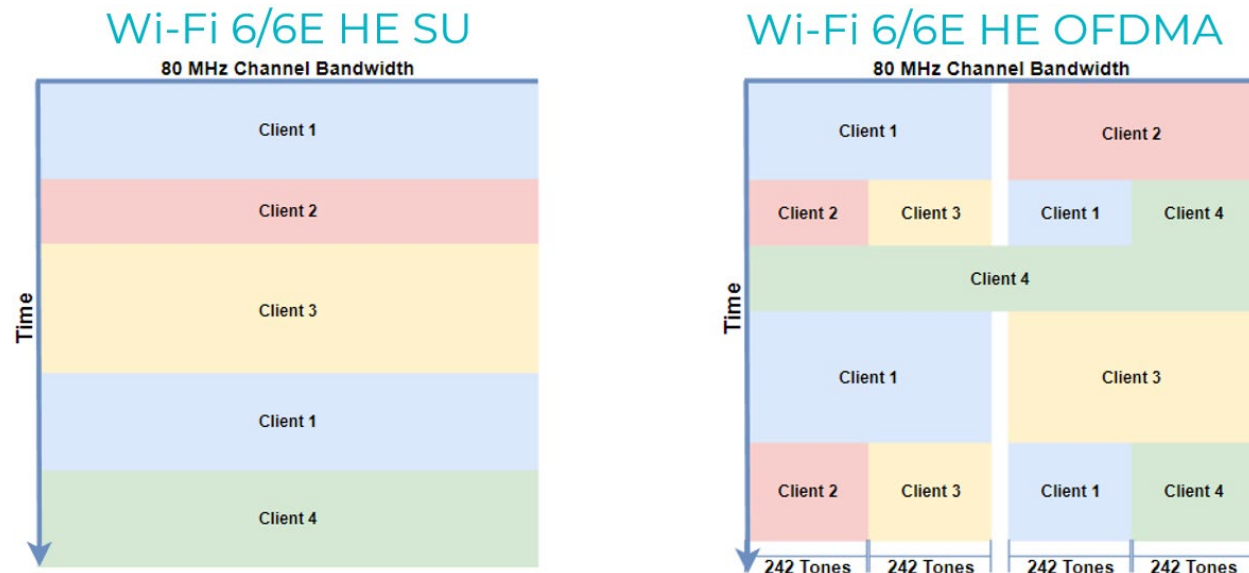


Figure 1 – Example OFDMA Tone Allocation with HE SU vs. HE OFDMA

The AP's scheduler is constantly managing if it prefers two clients at 484 tone RUs each in two separate groups to alternate between, or if it prefers to have four clients in a single group of 242 tone RUs each. This was observed to constantly change and is unknown exactly what proprietary reasons, other than the amount of traffic needing to be sent, that the AP is using to make decisions to use 2 groups of 484 tone RUs each or a single group of 242 tone RUs each. Receive power differences per client is one such reason for an AP to prefer to group certain clients together. However, when seeing spontaneous and rapid changes in RU assignments while clients are physically static and traffic is at a set limit, the rapid RU assignment changes and group changes are something that should be improved upon and will as AP schedulers are matured.

4.2. OFDMA: Process

For downlink OFDMA, non-Wi-Fi 6 clients are aware of the time they must be silent because after the AP has won a TXOP, it will send a multi-user request-to-send (MU-RTS) frame to clear the channel for the length of time of the full OFDMA exchange. The physical header and mac layer will contain the RU assignments, but they can also be specified in other trigger frames, or buffer status report frames. Other Wi-Fi 6 clients receive this frame and respond with their own CTS on their designated RUs at the same time. Next, the AP will send simultaneous data or multi-user downlink physical layer convergence procedure (PLCP) protocol data units (MU DL-PPDU) on each client's RUs. Clients will auto block ACK or wait for a block ACK request (BAR) from the AP and send upstream a block ACK at that time [2].

For uplink OFDMA, the AP still must win a TXOP, and only then can it schedule simultaneous uplink transmissions with uplink OFDMA on the clients that support it. A buffer status report poll (BSRP) is sent from the AP to the clients and solicits a BSR from each client. The BSRP contains the RU designations to be used for each client to respond with its own BSR which contains information about how much and what quality of service (QoS) of data needs to be sent upstream. This BSR information can also be unsolicited and indicated to the AP using a QoS control field in a data frame. A MU-RTS trigger frame may be sent from the AP and assigns RUs to each station as well as serves to notify the legacy clients of the upcoming transmission. The AP then waits for a CTS to return on each clients' assigned RUs. The MU-RTS is optional and can be skipped; the AP can go straight to a basic trigger frame which contains information about which RUs each client can use, the power each should try to use, as well as spatial streams and MCS rates to send their data upstream simultaneously. Next, the uplink PLCP protocol data unit (UL-PPDU) from each station is received for the same amount of time and each client will pad data if there is empty time. Finally, a multi-STA block ACK is broadcast so all clients can discover which frames need to be resent and the AP may optionally choose to send individual block ACKs to each client [2].

The goal of any network is to maintain high enough throughput and low, consistent latency, and low packet loss to provide a smooth, responsive, and predictable user experience. OFDMA technology in Wi-Fi 6/6E provides a mechanism to realize decreased maximum and average latencies in the Wi-Fi layer, as well as decreasing the variance or range of latencies experienced in more heavily utilized Wi-Fi channels.

5. Test House Setup

The Wi-Fi test house used for this testing is over 4500 square feet and consists of 3 stories including the finished basement and typical build materials such as wood floors, carpeted floors, sheetrock walls, and furniture throughout; it also has clean channels available with 0% utilized airtime because of the distance to neighbors. Microsoft Windows test clients were used to represent a more typical use case and good control of clients. The AP location was in a front corner of the house.

In contrast to normal lab testing techniques that seek to setup ideal conditions including precise location of clients, angles that present equal power to and from each client, and intricate test instrumentation, this paper was meant to determine if improvements can be realized in normal and non-curated test conditions with easy-to-use test applications. The locations of clients for this test were also chosen to represent a normal use case of clients in simultaneous use in the following locations: a client on the 2nd floor above the AP, a client in the basement below the AP, and a client in the same room as the AP.

A fourth client, the client under test (CUT), was located 26 ft away on the same floor and was the farthest client in this test on purpose as end-users are not usually right next to the AP. This client was used to ping and monitor more closely to investigate when the client was able to see better latency or not. The client was also used to play the game on for the test case involving playing an online game. This client under test would represent an end user's experience while other clients are using the same Wi-Fi network at various loads. The AP was not rotated to find ideal angles since comparison tests were planned, and therefore relative differences only were being considered.

6. Test Tools

Many test tools and scripts were evaluated including subscription-based tools we routinely use in our lab environments that are not freely available. However, there is value in being able to show improvements with tools that are freely available, such as iperf3. Another free tool NUTTCP was considered, but we had problems with ensuring time synchronized clocks with a LAN NTP server in the Wi-Fi test house.

This negated the advantage to using NUTTCP with UDP to measure OWD as this tool can do with well synchronized clocks. Also, because NUTTCP did not show the RTT of each TCP stream and instead only found RTT before the test started, it would have necessitated more complexity by using a second tool to get TCP RTT results. IxChariot was also considered but it proved harder to allow the type of automation required to test and then send back certain values into the next test automatically.

When using a source computer running Ubuntu Linux combined with iperf3, we were able to get RTT per TCP flow in the downstream direction, AP to Client, and this proved valuable enough to combine with ping data to select iperf3 as a tool. In the upstream direction, since the clients in use for this paper were Microsoft Windows to show applicability to real world use cases, we were missing the ability to get per TCP stream RTT. However, we still had the ability to run data in the upstream direction and use pings to evaluate the channel.

Another appealing reason to select iperf3 was the easy-to-use javascript object notation (JSON) output for aiding use in automation. The large number of test iterations planned called for creating some automation to organize and execute tests and remove any possibility for user error in setup of test parameters. Python3 was used to automatically setup certain iperf3 settings and rates, get data, and calculate different rates based on the data from a prior baseline run. Python3 also allowed for easy parsing of JSON output and presenting data in an easy to consume format, even the exact format needed to organize data in Excel. Using a Linux Ubuntu 20.04 operating system source Ethernet on LAN allowed for more precise timing control of pings used at a rate of 10 per second and was started just after traffic began and terminated just before traffic ended. The following network config parameters were changed on the linux machine to allow a larger TCP window size: *net.core.optmem_max = 524287; net.ipv4.udp_rmem_min = 8192; net.ipv4.udp_wmem_min = 8192; net.core.rmem_max = 16777216; net.core.wmem_max = 16777216; net.core.rmem_default = 2097152; net.core.wmem_default = 2097152; net.ipv4.tcp_rmem = 4096 87380 16777216; net.ipv4.tcp_wmem = 4096 87380 16777216; net.ipv4.tcp_mem = 1638400 1638400 1638400*. Using python3 allowed precise control of iperf3 sessions to all clients while simultaneously saving the TCP RTT data and using controlled pings to gather RTT information and throughput information for each controlled test cases.

The version of iperf used was 3.9. An example of the iperf3 commands that the custom written python3 scripts would launch with 'Popen' simultaneously and receive JSON output from is seen below. The rate-limited bitrate is per flow, and in this example, it was 4 flows per client at the bitrate listed in -b. The first 4 seconds of each run were always discarded with the -O omit flag. -J command was used to specify to receive results in JSON format. -R --get-server-output was used for an upstream direction test to specify receiving data at the Ethernet source from the clients.

- */usr/bin/iperf3 -c 192.168.0.238 -i4 -P4 -w2M -M1460 -b12.5M -t60 -O4 -J -T cut*
- */usr/bin/iperf3 -c 192.168.0.2 -i4 -P4 -w2M -M1460 -b62.13M -t60 -O4 -J -T client2*
- */usr/bin/iperf3 -c 192.168.0.3 -i4 -P4 -w2M -M1460 -b62.205M -t60 -O4 -J -T client3*
- */usr/bin/iperf3 -c 192.168.0.4 -i4 -P4 -w2M -M1460 -b62.3625M -t60 -O4 -J -T client4*

The JSON output received back per client was converted to a simple python3 dictionary object for parsing. The data for each stream to a client was then processed to take the average RTT of the streams, the maximum of the max RTT, and the minimum of the min RTT to represent that client's data in the test. An example of the data contained, per stream is seen in Figure 2.

```
"sender": {
  "socket": 9,
  "start": 0,
  "end": 59.999983,
  "seconds": 59.999983,
  "bytes": 467795968,
  "bits_per_second": 62372813.405630462,
  "retransmits": 128,
  "max_snd_cwnd": 770880,
  "max_rtt": 32568,
  "min_rtt": 16234,
  "mean_rtt": 21106,
  "sender": true
```

Figure 2 – Example iperf3 JSON to python3 Dictionary for a Single Flow on One Client

An example ping command, launched one second after the simultaneous iperf3 flows started is: *ping 192.168.0.238 -i.1 -c570 -w57 -D*. The pings were sent at a rate of 10 times per second and were set to end at 57 seconds with -w flag regardless of being able to complete the 570 pings specified by -c flag. This was necessary to make sure any delays or loss did not result in pings being measured after iperf3 flows had ended. The total time of traffic being sent from iperf3 to clients was 64 seconds, with the first 4 seconds being discarded.

Below in Figure 3 is an example at the end of the ping data received during a test; this allowed for parsing and graphing per ping sample.

```
[1656698977.360137] 64 bytes from 192.168.0.238: icmp_seq=560 ttl=128 time=45.0 ms
[1656698977.440025] 64 bytes from 192.168.0.238: icmp_seq=561 ttl=128 time=24.8 ms
[1656698977.542482] 64 bytes from 192.168.0.238: icmp_seq=562 ttl=128 time=26.3 ms
[1656698977.659009] 64 bytes from 192.168.0.238: icmp_seq=563 ttl=128 time=42.3 ms
[1656698977.750418] 64 bytes from 192.168.0.238: icmp_seq=564 ttl=128 time=33.2 ms
[1656698977.862114] 64 bytes from 192.168.0.238: icmp_seq=565 ttl=128 time=44.6 ms
[1656698977.951455] 64 bytes from 192.168.0.238: icmp_seq=566 ttl=128 time=33.1 ms
[1656698978.036393] 64 bytes from 192.168.0.238: icmp_seq=567 ttl=128 time=37.8 ms

--- 192.168.0.238 ping statistics ---
567 packets transmitted, 567 received, 0% packet loss, time 56941ms
rtt min/avg/max/mdev = 6.434/28.538/248.917/20.525 ms, pipe 3
```

Figure 3 – Example Ping Data Received During a Test

A particular online game was chosen as well, Counter Strike Global Offensive (CS:GO), because it was hosted easily on the Ethernet LAN connected Linux server that also sourced iperf3 traffic, and it was able to be played without interaction with the WAN internet at all. This game also has a developer mode that overlays latency and variance on the screen during play. This game is also not graphic intensive and was used only for the networking measurement aspect to show any relevant differences experienced in a real application under different channel load conditions and AP settings.




Figure 4 – Counter Strike Global Offensive Screen Overlay with RTT Latency

Figure 4 above shows an example of the overlay used to get the game's measurement on the RTT of data which changes during a given 60 second run. Its value was recorded every 3 seconds for analysis during traffic for the given game play test scenarios. This information along with iperf3 RTT data and ping data was used to characterize latency in this real application while the channel was set to various utilization levels and the AP under test was used in different modes.

7. Test Methodology

7.1. Test Constants

This paper sought to identify which scenarios can show lower latency, lower max latencies, or a more consistent delay with much lower variation in latency. This setup was specifically not in a lab or in a heavily curated test with equal power clients and equal traffic to each client. 2.4 GHz was not considered for this paper. For 5 GHz testing, channel 100 was used as it was a completely clear channel in the test house which allowed for repeatability. For Wi-Fi 6E tests channel 37 was chosen. A channel bandwidth of 80 MHz was used because this paper was about evaluating AP modes, OFDMA, and latency, not throughput. Increasing the bandwidth of the channel to 160 MHz would inflate the total throughput required on each test case to max out the channel and leads to less deterministic or repeatable results. However, some chipsets and firmware today can aggregate even more clients on 242 tone RUs in a single group with 160 MHz. This will be left to one of many additional tests for a later date.

Another test constant decided for this paper was to completely disable MU-MIMO including VHT MU-MIMO as well as HE MU-MIMO. MU-MIMO contributions to latency or throughput was not the focus of this paper and would have only served to confuse the contributions of OFDMA in various test scenarios including when baselining the ability of a channel with three clients, vs. adding in one more client under test in the same channel. The additional client under test added to a scenario would change the interference experienced and grouping of the other clients negating the baseline if MU-MIMO was enabled beforehand. The various non-repeatable conditions and baselining issues in addition to the overhead of null data packet announcement frames and beamforming report poll frames made disabling MU-MIMO for this paper an easy decision.

Default QoS was used by not setting QoS at all, to not intermix this layer of prioritization in with the evaluation of the OFDMA feature in these test conditions and iterations already defined. Other constants included the fixed physical positioning of the AP and clients between representative loads and setting changes. A certain AP was used to ensure coverage of Wi-Fi 6E data, however, the maturity of firmware on the Wi-Fi 6E AP is still lagging APs with just 5 GHz Wi-Fi 6. Two other APs for a good sampling of today's AP abilities were used to test in 5 GHz.

7.2. AP Settings for Comparison

The AP under test was controlled and tested against the following modes, settings, and client modes:

- Wi-Fi 5
- Wi-Fi 6 with OFDMA DL & UL disabled
- Wi-Fi 6 with OFDMA DL & UL enabled
- Wi-Fi 6 with OFDMA DL & UL disabled and two clients set to Wi-Fi 5 (802.11AC) only
- Wi-Fi 6 with OFDMA DL & UL enabled and two clients set to Wi-Fi 5 (802.11AC) only
- Wi-Fi 6E with OFDMA DL & UL disabled (if available)
- Wi-Fi 6E with OFDMA DL & UL enabled (if available)

At the beginning of a test with the AP set to a particular mode and OFDMA setting, a set of baseline tests were executed. The baseline tests determined what the max average rate was for each of the three control clients running traffic together. This was baselined separately for each AP mode and OFDMA setting and directly preceded each set of test iterations listed in the client conditions for comparison section.

The baselining served a couple purposes. Each client was in a different location which represented a more typical use case and therefore each client's capability to achieve a certain Mbps throughput was different. For example, the same throughput achieved in Wi-Fi 6 was more than what was achievable in Wi-Fi 5; it would not be a fair comparison of channel utilization to set the Wi-Fi 6 speeds at the same rate limits used in the Wi-Fi 5 testing, because the channel would be at a different airtime utilization in the comparison. A better approach was to determine what speeds the clients could achieve during each set of AP settings, then discount the rates of those clients to achieve an approximate percentage of channel utilization; this would be a different bitrate for each client and each AP mode, packet size, OFDMA settings, traffic direction, and client mode settings. Another reason for doing this is certain settings in use would cause different schedulers to be employed and different technology or decisions to be used. For example, in UL OFDMA the AP is instructing the clients which MCS to use amongst other parameters, however, with OFDMA disabled, the clients are choosing their own rates.

For these and other reasons it was decided to not rate limit other clients based on Mbps blindly for all test cases, but instead to rate limit based on percentages of max achieved average rates when the other three clients were used simultaneously without rate limits. After determining the throughput on average achieved by each of the 3 clients with no rate limits, this allowed a known discounted rate to approximate channel utilization to a given percentage for subsequent tests. This method should also allow extrapolation of findings to similar test scenarios with clients that are far away and not able to do high rates, but they could of course use up certain percentages of the channel's airtime just the same.

7.3. Client Conditions for Comparison

During the test iterations for all the AP settings defined above, different parameters for the other clients and the CUT were then tested over the following test parameter iterations:

- Downlink and uplink traffic direction
- TCP and UDP
- Client under test rate limits of 500 Mbps, 50 Mbps, 5 Mbps, and 1 Mbps while also playing a low bandwidth LAN Ethernet hosted game.
- Other three clients rate limited to 10%, 50%, and 90% of their capability determined in prior baseline tests with just the three clients running traffic without rate limits
- 536 byte and 1460 byte size payloads (Used for both TCP and UDP tests)

The different rate limits chosen were to represent high, medium, low, and very low rates for the CUT. The rate of 500 Mbps was specifically chosen to be below the full channel capacity for most test conditions, but also higher than what could be achieved in a 242 tone RU assignment for a 2x2:2 client at the client locations chosen. This was to expose some areas where OFDMA can reduce throughput as schedulers struggle to figure out how to divvy up tones and groups for multiclient scenarios with high data rates for some clients but not all clients. The scheduler must decide how to handle a highly utilized client in the face of lower utilized clients. There is not a right or wrong answer here, and we believe it is a source of some unknown expectations and ongoing improvement.

The highly utilized client could be transmitted to using the entire channel and then rotated between the other OFDMA transmissions of the 3 other clients using their RU designations for short periods of time. However, the highly utilized client could also be put in a smaller than needed RU designation for the sake of preserving low latency access to others and put in the same group as the lower RU designated clients; this would cause throughput degradation in the name of reducing total latency for all clients. Finally, it could decide to only briefly use a smaller RU combined in the same group with other lower utilized clients for part of the time, and then switch back to full usage of the channel causing higher latency on the other clients.

The channel utilization percentages chosen included 10% for the other clients to represent low utilization either in the house or to an extent, overlapping neighboring networks on the same primary or secondary channels of a similar percentage. Utilizations of 50% and 90% would represent busy and very busy channels which will uncover at what points of channel utilization OFDMA starts to help with latency reduction. The same testing re-executed while creating controlled amounts of out-of-network traffic on the same primary and secondary channels are left to future testing. This would be to characterize the difference in results with uncorrelated utilization and EDCA backoffs being observed between different APs in other houses or even multi-AP mesh systems sharing channels within a house.

The different packet sizes represent a range of payloads, small and large, to determine if benefits are only seen based on the size of packets sent in the channel. The same payload was chosen for UDP payload and TCP max segment size (MSS) tests for parity's sake even though the UDP data could have been smaller than the lowest TCP MSS and 12 bytes larger than the largest TCP MSS. The packet size for a test scenario is used to all 4 clients for the test case, not just the CUT. iperf3 TCP flows were set with 4 flows except for the CUT at 1 Mbps and 5 Mbps where multi-flow was not beneficial; in this case just 1 flow was used. iperf3 TCP window size was set to 2 MB in all cases. UDP was also tested to isolate much of the other direction's traffic except for MAC layer 2 acknowledgments (control frames) to see if behavior is much different without the constant interruptions to send and receive TCP acknowledgments as data frames in the opposite direction. Lab test scenarios often start with UDP for more exact control, but it is not indicative of real-life application flows this paper sought to primarily evaluate.

8. Test Results

Different AP chipsets were used to do this testing and to investigate what amount of channel utilization or amount of client under test utilization can show improved latency with the OFDMA feature in current generation firmware. The test environment was not overly curated to create ideal lab test conditions in a chamber or test cases that included every client receiving the same low to medium throughput with UDP and equal power as many lab test cases and vendors would and should execute. This paper evaluated a non-lab test scenario in an actual house of four clients of controlled and different throughput with TCP and UDP while at varying AP receive power, AP transmit power, and distance. This simulated a real situation but with repeatable, controlled utilization of clients to facilitate finding relative differences between OFDMA settings and AP modes in a real environment.

The relative comparisons and charts were made on the same chipset for direct comparison and consideration of OFDMA's contributions to latency. There were extreme differences in chipset abilities, algorithms, and scheduler decisions, but comparisons were made on the same AP. 260 to 364 tests created approximately 5,000 to 6,000 data points per AP and protocol. This was executed on three APs with both protocols creating total data points of well over 30,000 between the modes and different APs tested. Each test also had additional raw data saved that represented the per sample data making up the averages, min, and max; some of the per sample data was also analyzed and graphed. The data was examined to find generalizations to illustrate improvements that could be expected in certain scenarios.

There is much data created from this type of exhaustive testing, as well as what would be created from the myriad of proposed follow-up testing scenarios. Many of the scenarios tested were not expected to show a latency RTT improvement. Executing each scenario, however, was the only way to find at what amount of channel usage and traffic level the client under test would start to realize a latency improvement in today's chipsets and firmware. Improvements and more testing will be continuously sought, while drawing more conclusions to the data already collected.

The most compelling RTT latency reduction improvements included the scenarios with high channel utilization, including maxed out, non-rate limited iperf3 flows for each client as well as many of the 90% channel utilization test cases. Also, the higher usage of the client under test provided the highest improvement in latency RTT for that client. The other generalized improvement was when using a larger MSS of 1460 bytes compared to smaller MSS of 536 bytes. Many test results did not show favorable results and in many scenarios, this was expected. For example, in a low utilized channel the airtime is not causing contention and clients already get access to the channel without the need of OFDMA. Test results selected to be shown in this paper are just a very small subset chosen to represent a mix of UDP and TCP, large and small packets, downstream and upstream, different rates to the CUT, and different channel utilizations.

For each chart consider the following way to look at the chart. The categories on the X axis are used to group test results for a given AP mode and client mode scenario including OFDMA enabled or not. Each test result for that given test mode is then listed front to back, starting with latency values as indicated and labeled per parameter as a category to the right, and ends with total throughput for reference in the back of the chart. The Y axis is being use for both Mbps for the last category per group and serving as a millisecond reading for all the other categories with latency.

8.1. Non-Rate Limited Test Scenarios

The tests with no iperf3 rate limits to all four clients were sourced from or destined to a single Ethernet 1 Gbps port on the LAN of the AP. Non-rate limited scenarios are perhaps where many test efforts start and is worthwhile to characterize, but it is not a likely scenario to occur in a house. However, since this test case is craved by many, it was also evaluated. In general, when comparing unidirectional UDP vs. TCP results, the ping RTT is lower for the same test case using UDP due to most of the data being in one direction.

When testing in the upstream with TCP, as previously discussed, the only indication we have of RTT, because of using Microsoft Windows clients with iperf3, is the ping RTT statistics. When testing the same non-rate limited scenario in the upstream direction, we do see improvements with OFDMA enabled on TCP and UDP.

Shown in Figure 5 is UDP in the upstream direction at UDP Payload of 1460 bytes, the average RTT of a ping is reduced by more than half for OFDMA enabled scenarios in the maxed-out channel test. For example, with all Wi-Fi 6 clients used and OFDMA disabled, the average RTT of a ping was 14.41 ms to the client under test (CUT). With OFDMA enabled on the AP for the same four Wi-Fi 6 client test, the average ping to the CUT was reduced to 6.76 ms. Even better results are seen with UDP payloads of 536 bytes in this same test scenario, however, the throughput is better with the UDP payload of 1460 bytes.




Figure 5 – UDP Upstream No Rate Limits

In this completely saturated airtime scenario, the per ping sample to the CUT was graphed in Figure 6; it is quite revealing in the OFDMA disabled to OFDMA enabled comparison for upstream traffic in a completely used airtime scenario. The AP scheduling upstream transmissions allows for very quick and continuous access to the channel. In Figure 6 below, a much tighter and lower latency can be seen with the blue line with OFDMA enabled and is a visual representation of the median deviation of RTT latency changing from 15 ms to 3 ms which was revealed in the previous Figure 5 as well as the max RTT recorded reducing 3 fold.




Figure 6 – Per Ping Sample OFDMA Disabled vs. Enabled

In Figure 7 below, the same scenario of UDP and maxed out upstream traffic is compared between AP set to Wi-Fi 5 (gray line) vs. OFDMA enabled (blue line) and disabled (orange line) while half the clients are set to Wi-Fi 5 only mode. This test mode was tried in all scenarios to evaluate if benefits to Wi-Fi 6 clients can be realized in multi-client scenarios that include previous generation Wi-Fi 5 clients. The latency reduction benefit is still seen to the Wi-Fi 6 client under test with half the clients being unable to use OFDMA but still sending traffic upstream.




Figure 7 – Per Ping Sample Wi-Fi 5 vs. OFDMA Disabled/Enabled with half Wi-Fi 5 Clients

Next is an example of downstream TCP traffic while maxing out the utilization of the channel with no rate limits being used. There was a particular improvement seen in the mixed client population test scenario where half of the network's clients, are set to Wi-Fi 5 mode only. The other two clients, including the CUT, are left to support Wi-Fi 6 and OFDMA. In the following scenario TCP was used in the downlink direction at 1460 byte MSS with no rate limit set on iperf3 while measuring the client under test's RTT latency. This scenario again shows a benefit can still be seen on a Wi-Fi 6 client under test while half the clients in the test are Wi-Fi 5 clients and don't have an ability to use OFDMA. The AP can schedule the OFDMA clients together, in this case 484 RU each in a group while Wi-Fi 5 clients are served in alternating fashion with the OFDMA group of size 2 as transmit opportunities are won.

Shown in Figure 8, the maximum RTT latency (orange bars) recorded during a 60 second test showed a latency reduction of about 50% with OFDMA enabled and was seen in both max ping RTT as well as max data RTT (blue bars) to the Wi-Fi 6 client under test. The average RTT (gray bars) only showed slight improvements in OFDMA enabled case, however the max RTT recorded was meaningfully lower. This reduction in the max RTT recorded is very likely to improve the QoE for the end user in the most demanding maxed out scenarios of fully used channel airtime from mostly downstream traffic. The mean deviation of the latency (green bars) or jitter as reported from the approximately 567 pings in the 1 minute test time showed a reduction to just 8.42 ms with OFDMA enabled while OFDMA disable was at 22.26 ms and legacy Wi-Fi 5 was at 50.98 ms.




Figure 8 – TCP Downstream No Rate Limits

The same scenario of the AP at Wi-Fi 6 but half the clients being set to Wi-Fi 5 is graphed below in Figure 9 on a per ping basis to the CUT and shows a good improvement in reduction of frequency of RTT latency spikes with OFDMA enabled (blue line) vs. OFDMA disabled (orange line). The AP set to Wi-Fi 5 mode (gray line) is shown as well to show how variable Wi-Fi 5 networks can be when loaded.




Figure 9 – Per Ping Sample Wi-Fi 5 vs. OFDMA Disabled/Enabled with half Wi-Fi 5 Clients

In summary, for the non-rate limited scenarios with heavy in-network channel utilization, the OFDMA enabled tests showed improvements in average RTT and max RTT as well as mean deviation of RTT latency. For upstream traffic utilization UDP traffic tests were the easiest to see drastic improvements while for the TCP traffic tests the most latency reduction was seen in downstream direction test cases. In most non-rate limited test scenarios there was an improvement for Wi-Fi 6 client latency even when half the clients receiving background traffic were Wi-Fi 5 clients.

8.2. Rate-Limited Test Scenarios with Channel Utilization Set to 90%

In this section's test scenarios, the bitrate limits for each of the other 3 clients were set to a limit of 90% of what each had achieved without rate-limits in baseline testing, per AP mode. This caused approximately 90% channel utilization as compared to the baseline unlimited test case. Two of this scenario's results were selected to discuss below.

In the following scenario with results shown in Figure 10, the client under test was set to send 50 Mbps upstream with the channel utilization at approximately 90% with upstream traffic from the other clients. This scenario represents a very highly utilized channel from the other clients while a moderate throughput demand in the upstream is created from the CUT. The OFDMA enabled test, shows a reduction in the max RTT ping (orange bars) from 313 ms to 121 ms to the CUT with an average RTT ping (gray bars) that reduced from 50 ms to 37 ms. The median deviation or jitter (green bars), also decreased from 44 ms to 21 ms for the same OFDMA enabled vs. disabled comparison to the CUT. Additionally, the data in Figure 10 records a decrease in the max RTT ping seen when OFDMA is enabled with half the clients set to Wi-Fi 5 mode; this shows some improvement can still be realized in mixed client populations of Wi-Fi 5 and Wi-Fi 6 clients.




Figure 10 – TCP Upstream, 90% Limit Others, CUT Limit 50 Mbps

In Figure 11 below, the per ping sample between OFDMA disabled (orange line) and enabled (blue line) is graphed for the same scenario depicted in Figure 10 and shows visually the drastic improvement in the max RTT latency observed with much lower spikes in latency and a much tighter jitter or median deviation of the latency. The median deviation was shown previously to have reduced from 44 to 21 ms from OFDMA disabled to enabled and is seen below with a smaller range of values during the test, meaning the blue line's spikes in latency were much lower with OFDMA enabled.




Figure 11 – Per Ping Sample OFDMA Disabled vs. Enabled

The next scenario, for an approximately 90% utilized channel, was with downstream TCP traffic at a smaller 536 byte MSS and was more compelling on a different AP shown below in Figure 12. Other clients were limited to approximately 90% and the rate-limit of the CUT was set to 500 Mbps. This test scenario is like the maxed-out scenario in that the total allowed rates are well over the ability of the channel to support but serves as a test to see if the CUT can get more throughput in the channel with less latency on that data when OFDMA is enabled. Indeed, in addition to improving the throughput achieved, latency was also reduced for the CUT. The CUT throughput (dark blue bars) seen in this scenario for each AP mode included Wi-Fi 5 achieving 69 Mbps, Wi-Fi 6 with OFDMA disabled achieving 153 Mbps, and Wi-Fi 6 with OFDMA enabled achieving 173 Mbps. The scheduler preferred to give more equal access to the channel preserving lower latencies for all clients while still allowing more traffic through on average to the CUT in the OFDMA enabled test. This contrasts with assigning a higher number of tones for this client and others in two groups or just 996 tones to the CUT to try to allow the higher TCP limit of 500 Mbps, however it did not do this. In Figure 12 below, both the iperf3 data max RTT (light blue solid bar) and iperf3 average RTT (orange bars) have reduced by about 30 - 40%.




Figure 12 – TCP Downstream, 90% Limit Others, CUT Limit 500 Mbps – AP 2

In summary, the approximately 90% utilized channel scenarios showed most improvements in TCP upstream and UDP upstream scenarios. With UDP downstream scenarios the results were usually worse at all CUT bitrates tested and may allude to the fact that the scheduling overhead is just too much to see a benefit over OFDMA disabled test cases. In UDP downstream tests, low latency is already able to be achieved to begin with since the AP is already in control of scheduling traffic one at a time to each client without much interruption in the other direction since only layer 2 acknowledgements need to be received. With downstream TCP scenarios the lower utilization tests of the CUT were not able to realize much of an improvement, however, with slightly higher bitrates to the CUT such as 50 Mbps or 500 Mbps, some moderate improvements were measured.

8.3. Rate-Limited Test Scenarios with Channel Utilization Set to 50%

In this section's test scenarios, the bitrate limits for each of the other 3 clients were set to a limit of 50% of what each had achieved without rate-limits in baseline testing, per AP mode. This caused approximately 50% channel utilization as compared to the baseline unlimited test case. Two of this scenario's results were selected to discuss. As channel utilization is reduced from 90% to 50% the RTT latency benefits are seen to less of an extent but are still noticed if the CUT has a high bitrate.

In the following scenario, the CUT is set to receive at a 500 Mbps limit while the other clients are set to a limit of 50% to represent approximately 50% channel utilization before the CUT's traffic. The 500 Mbps rate limit set on the CUT is a higher demand or throughput limit than available unused channel airtime and higher than what can be achieved in only a 242 tone RU alone. This tests if the scheduler will give higher RU allocation, such as 484 tones, to this client under test to achieve its higher throughput needs. Or, if the scheduler would cause throughput degradation to the CUT by assigning a smaller than required RU allocation such as 242 tone to maintain a single group of all 242 tone RUs to each of the four Wi-Fi 6 clients, in the OFDMA enabled test case.

Since the client throughput (dark blue bars) can achieve a rate over time that is higher than the rate possible in just a 242 tone RU, this scheduler sacrificed some latency reduction potential that could have been obtained for all clients, at the expense of creating multiple groups with higher number of tones for each client. In this way, a higher throughput on a client demanding more than the others was accommodated. Therefore, the latency reductions shown in Figure 13 are more muted, and a great example of scenarios that are not expected to show a big RTT latency reduction because of the test conditions and scheduler deciding to allow as much of the CUT's higher throughput traffic as possible.




Figure 13 – TCP Downstream, 50% Limit Others, CUT Limit 500 Mbps

In another scenario of approximately 50% channel utilization allowed on other clients, downstream UDP traffic is used with a moderate limit of 50 Mbps to the CUT. The average RTT latency (gray bars) on the CUT is about 1 to 2 ms higher with OFDMA enabled. However, the max RTT latency (orange bars) did fall from 104 ms to 39 ms with OFDMA enabled, and from 86 ms to 54 ms in the mixed Wi-Fi 5 client tests with OFDMA enabled. This is an example of a scenario that may not show great average latency reduction on the Wi-Fi 6 CUT but does show a lower max RTT latency and would result in a better QoE for the end user. These results are shown below in Figure 14.




Figure 14 – UDP Downstream, 50% Limit Others, CUT Limit 50 Mbps

To illustrate the moderate benefit of a lower max RTT for this scenario, a per ping sample graph in Figure 15 is shown below. The spikes in latency are still seen but not as high, and the slight increase in average RTT ping is also able to be seen.




Figure 15 – Per Ping Sample OFDMA Disabled vs. Enabled

In summary, the approximately 50% utilized channel scenarios showed most improvements with OFDMA enabled when the CUT was set to a higher throughput limit such as 500 Mbps. This was certainly the case with UDP in the upstream direction, but not as convincingly seen in the downstream direction. UDP upstream isn't really a use case seen at this high bitrate in real-life and therefore focus was to show UDP and TCP downstream scenarios described above. UDP downstream showed a small improvement in RTT latency with the CUT receiving 50 Mbps. For TCP traffic tests, the only cases improving with approximately 50% channel utilization were in the downstream direction and at 500 Mbps to the CUT. The TCP upstream tests, for the 50% channel utilization scenarios, for the most part showed slightly worse latency.

8.4. Rate-Limited Test Scenarios with Channel Utilization Set to 10%

In this section's test scenarios, the bitrate limits for each of the other 3 clients were set to a limit of 10% of what each had achieved without rate-limits in baseline testing, per AP mode. This caused approximately 10% channel utilization as compared to the baseline unlimited test case. Two of this scenario's results were selected to discuss including one to show the difference between two different AP's in the same scenario.

In the least utilized channel test scenario with other clients set to 10% rate limits, there were less situations found with meaningful improvements in latency for the CUT. One of the only scenarios that showed a compelling difference for a low utilized channel and highly utilized client was in the mixed client mode test with and without OFDMA enabled. In the scenario shown in Figure 16, TCP is used with a 1460 byte MSS and a 10% utilization limit for the other clients while the client under test is set to a rate limited 500 Mbps in the upstream direction. This represents a very high throughput demand in the upstream while the channel utilization is quite low in the same direction on the other clients. The mixed client mode cases of half Wi-Fi 5 clients and half Wi-Fi 6 clients showed a compelling difference in the latency experienced on the CUT during this high throughput of a lightly utilized channel in the upstream. The CUT was able to achieve the 500 Mbps upstream in both cases, yet the RTT max latency (orange bars) as well as the RTT average latency (gray bars) both were reduced approximately 3-fold.




Figure 16 – TCP Upstream, 10% Limit Others, CUT Limit 500 Mbps

The same scenario is graphed with per ping RTT latency data in Figure 17 below and really illustrates the 5 times lower RTT median deviation in the ping data to the CUT with OFDMA enabled (blue line) in the mixed half Wi-Fi 5 client test. The data set to flow upstream from the CUT was able to achieve 500 Mbps in both scenarios, but the OFDMA enabled case did it with a much lower average RTT latency as well as considerably less jitter.




Figure 17 – Per Ping Sample OFDMA Disabled vs. Enabled with half Wi-Fi 5 Clients

As mentioned before, not all scenarios showed reduced latency and many also show increased latency, especially with lower utilized channels or lower rates to the CUT. Most notably, differences between different APs are very easily seen when compared in the same test conditions and scenarios. The same scenario with two different APs will illustrate this extreme difference that can be seen.

The scenario compared here was TCP 1460 byte MSS with 10% traffic in the downstream direction with the CUT receiving a rate-limited 5 Mbps. So, this represents lightly used channel conditions and a low throughput on the CUT. This is actually a very common scenario and may represent the most often situation experienced in the home.

The first AP's results, in Figure 18, show that with OFDMA enabled or disabled the same average RTT latency (solid light blue bars) for the iperf3 data flows is seen at 4 ms. The max RTT latency with OFDMA enabled increased from 5 ms to 7 ms and reduced max RTT latency from 7 ms to 5 ms in the mixed Wi-Fi 5 client mode case. All these latency values, both average and max RTT are acceptably low for the local LAN segment for most applications but are shown here to illustrate the improvements are not usually seen in low to medium use cases with 4 clients. Also, this simple result is given as a contrast to the result with the same scenario on a different AP.




Figure 18 – TCP Downstream, 10% Limit Others, CUT Limit 5 Mbps – AP 1

In Figure 19, the same scenario is tested using a different AP and revealed some issue with increased latency for the OFDMA enabled test cases with all Wi-Fi 6 clients and again in the mixed half Wi-Fi 5 client mode test as well. This is likely from scheduler problems with deciding if the clients should be in a group or not and the overhead associated with creating and tearing down groups when low bitrates are used. The scheduler in this AP should be improved to create groups early, even with low bitrates, to receive upstream TCP acknowledgements coming back as data frames in OFDMA transmissions.




Figure 19 – TCP Downstream, 10% Limit Others, CUT Limit 5 Mbps – AP 2

An area of further testing to be completed is to investigate at what bitrate threshold AP schedulers are considering adding clients to OFDMA groups. A client at distance could be achieving a low bitrate but are using a lower MCS rate and airtime usage would still be high; these low bitrate and high airtime clients should not be denied scheduling into OFDMA groups. A threshold based on predicted airtime usage would be a better approach for such situations when clients are not close to the AP and low MCS rates with high airtime usage are being used. The data shows that highly utilized channels can benefit greatly by OFDMA and highly utilized refers to airtime usage, not bitrates. Therefore, it is likely beneficial that chipsets use thresholds that are airtime based, not throughput based if the scheduler requires such thresholds.

In summary, the approximately 10% utilized channel scenarios showed some small improvements with OFDMA enabled when the CUT was set to 5 Mbps and TCP was used with half Wi-Fi 5 clients in either direction. It was also an easy result to showcase differences in implementation of OFDMA in different chipsets. In other TCP tests, with the bitrate to the CUT set to 50 Mbps or 500 Mbps, latency generally was slightly worse with OFDMA enabled with either direction of traffic. However, an exception was the mixed client scenario with upstream traffic of 500 Mbps to CUT and 10% utilization with other clients, a large reduction in latency was seen with OFDMA enabled and was shown in Figure 16. UDP traffic scenarios with CUT limits of 1 Mbps, 50 Mbps, and 500 Mbps showed slight improvements in the

downstream direction tests while the upstream direction only showed an improvement with 500 Mbps from the CUT.

8.5. Rate-Limited Test Scenarios with LAN Hosted Gaming

Online gaming requires quick and responsive connections to have a good QoE. As explained in more detail earlier, a LAN Ethernet hosted CS:GO game was played within the LAN on the CUT while monitoring the RTT Latency reported on the game screen in real-time. Simultaneous to that game play is iperf3 data of 1 Mbps flowing to or from the same CUT and pings just as in prior scenarios. This provides information about the responsiveness from the CUT in the traditional RTT iperf3 data and ping data, in addition to the CS:GO RTT latency reported on the screen overlay.

The 10%, 50%, and 90% usage levels of the other clients as well as iterating over the two packets sizes and AP modes showed similar results to testing without the game play. The scenario that showed an obvious improvement in game play was the heavily utilized 90% TCP upstream traffic scenario, the results are shown in Figure 20. A real-life example of this scenario might include one or more clients in the house backing up data or uploading videos to the cloud while another client is trying to play an online game. Notice the max RTT ping (orange bars) reductions from 212 ms to 69 ms with OFDMA enabled and from 196 ms to 74 ms with half the clients being Wi-Fi 5 clients. The RTT median deviation ping (green bars) also shows a reduction of approximately 50% which alludes to more consistent and lower pings achieved in this scenario to the CUT with OFDMA enabled regardless of half the clients being Wi-Fi 5 clients. The average RTT ping (gray bars) was cut in half when OFDMA was enabled with all Wi-Fi 6 clients.




Figure 20 – TCP Upstream, 90% Limit Others, CUT Limit 1 Mbps + Game Play

Consistent with the above measured max and average RTT ping times improving with OFDMA enabled in the heavily utilized channel in the upstream with TCP, the CS:GO game play itself on the CUT also showed improved latency in both OFDMA enabled modes. The average ping recorded during the same sample period of the game play was very close to the RTT Latency reported on the screen overlay during the CS:GO game play itself. This sampling of the CS:GO RTT latency over the 60 seconds is graphed below in Figure 21 and shows the two OFDMA enabled cases performing much better with all Wi-Fi 6 clients (blue line), as well as the mixed half Wi-Fi 5 clients (green line). Both OFDMA enabled tests show an obvious consistent latency or low jitter to the CUT during game play while the channel has very low available airtime with heavy upstream traffic. This was also evident by the reduction in median deviation of RTT pings shown in Figure 20 above.




Figure 21 – CS:GO Reported RTT Latency Over 1 Minute of Game Play

In Figure 22 below, the same game play scenario is graphed per ping RTT latency for OFDMA enabled (blue line) vs. OFDMA disabled (orange line). The per ping data shows the same increasing latency for OFDMA disabled at the same points on the graph as were reported by the overlay on screen of CS:GO and depicted in Figure 21 above. In Figure 21 at 36 seconds into the OFDMA disabled test scenario (orange line) latency begins to rise; this correlates with the ping data in Figure 22 for the OFDMA disabled test (orange line) at around the 360th sample which also starts to have more spikes in latency. The ping data samples over time in Figure 22 below show that pings are a good representation of what is also occurring to the actual game flow's latency even if the absolute values are different.




Figure 22 – Per Ping Sample OFDMA Disabled vs. Enabled During CS:GO Game Play

In summary, during online game play via Wi-Fi, OFDMA enabled really improved the variance of latency, or jitter, in highly utilized channels in the upstream for both TCP and UDP. This improvement was measured solely on the LAN as game play was limited to within the LAN. The improvement on the CUT playing the game was realized even with half the clients in the network being Wi-Fi 5 clients. The consistency of RTT latency is much improved with OFDMA enabled in this scenario that represents several other devices uploading data and using approximately 90% of the channel airtime at the same time as a CUT is playing an online game.

8.6. Wi-Fi 6E Testing with Rate-Limited Test Scenarios

Wi-Fi 6E testing revealed some throughput degradation perhaps from range degradation at 6 GHz or from immaturity of firmware on AP or client for 6 GHz. The specific AP used for this testing showed some unexpected degradation in the upstream with OFDMA enabled even in 5 GHz. However, relative differences are still able to be seen by rate-limiting to certain percentages and comparing relative latency during the same approximate channel utilization across AP modes and scenarios.

Like other findings with APs in 5 GHz band, the highly utilized channel of 90% showed the most improvement to the CUT when enabling OFDMA in the 6 GHz band. Figure 23 shows 6 GHz testing with 90% utilization on control clients during TCP 1460 byte MSS testing in the upstream direction with 50 Mbps from the CUT. All stats improved with pings for OFDMA enabled cases, including the median deviation RTT (green bars), average RTT (gray bars), and max RTT (orange bars).




Figure 23 – TCP Upstream, 90% Limit Others, CUT Limit 50 Mbps – AP 3 with Wi-Fi 6E

The same scenario's per ping RTT latency is graphed in Figure 24 below to show Wi-Fi 6E OFDMA disabled (orange line) vs. OFDMA enabled (blue line). The reduction in median deviation of RTT pings as well as the reduction of the max RTT ping recorded is evident with a more consistent and lower average latency for the OFDMA enabled test.




Figure 24 – Per Ping Sample OFDMA Disabled vs. Enabled – AP 3 with Wi-Fi 6E

8.7. Wi-Fi 6E Rate-Limited Test Scenarios with LAN Hosted Gaming

While playing CS:GO on the CUT during the Wi-Fi 6E test scenario, the RTT latency improved the most during upstream 90% channel utilization scenarios while the downstream version of the same test showed about 50% worse RTT latency. Shown below in Figure 25 is the 1460 byte MSS scenario and shows results from a 90% utilized channel in 6 GHz before and after enabling OFDMA. The 536 byte MSS results were similar. This chart is the CS:GO reported RTT latency during a 60 second test while approximately 90% utilization is occurring in the downstream or upstream as labeled by a different colored line. Improvements in latency on a CUT playing a game are not always seen based on if the channel is heavily utilized with downstream traffic or if it is the same percent utilization with upstream traffic. It is also likely something this AP would improve upon with later firmware over time.

In Figure 25, the blue line shows game play latency improved while heavy traffic was flowing in the upstream direction using 90% of the channel. This is to be compared to the orange line showing the same test with OFDMA disabled.

However, the 90% utilization set in the downstream with OFDMA disabled produced a lower average RTT latency shown as the green line in contrast to the OFDMA enabled case shown as the red line, where RTT latency was consistently higher. These results further underscore the point that improvements in latency are not necessarily the same based on the direction of the traffic that is loading the channel. During this game play the downstream traffic load produced lower RTT latency with OFDMA disabled, while the upstream traffic load test produced lower RTT latency with OFDMA enabled.




Figure 25 – CS:GO Reported RTT Latency Over 1 Minute of Game Play – Wi-Fi 6E

9. Additional Tests to Consider

There are many opportunities to extend this work and help chipset vendors to improve schedulers for real world scenarios and expectations. A list of additional tests to explore include:

- Create approximate defined percentages of out-of-network traffic on the same primary and secondary channels
- Create approximate defined percentages of out-of-network traffic on only the second half of an 80 MHz channel in the non-primary channels
- Traffic load set on the channel in the opposite direction of the direction of traffic to/from a client under test
- The same test conditions in this paper but with 5, 8, 9, and 13 clients at 80 MHz and 160 MHz to test grouping and usage of mixed RU allocations
- The same test conditions in this paper but with Linux operating system on the clients to get RTT of each TCP data flow in iperf3 in the upstream as well
- 20 clients with 12 clients doing constant 0.1 Mbps traffic to stay utilized while testing other 8 clients with various traffic and channel utilizations
- Characterize how a scheduler allocates RUs as you gradually increase bitrates or airtime percent in use
- Locating clients at near and far locations for 2 groups of power ranges
- Locating clients at near, mid, and far locations for 3 groups of power ranges
- Creating varied but defined traffic load percentages for the other clients to represent low usage and medium usage on many clients in addition to a client under test usage
- Send high data rate to a client while alternating low data rates to another and then vice versa to see if groups are too sticky or if RU allocations are not changed fast enough
- Enable VHT MU-MIMO for evaluation against OFDMA benefits with clients in various locations where VHT MU-MIMO can operate ideally
- Enable HE MU-MIMO for evaluation of scenarios that should use HE MU-MIMO instead of or at the same time as OFDMA – run each separately to determine if schedulers are picking the right mode when both are enabled
- Set up a reliable NTP for better time synchronization between source and clients to get reliable one-way-delay UDP measurements with NUTTCP
- Evaluate the same test case repeatedly to evaluate consistency of results with schedulers for a given test method and given scheduler implementation – is it doing the same thing each time?
- Statically/manually set OFDMA groups and RU number of tones per client to evaluate best case 100% OFDMA performance and compare to results of automatic scheduler operation in throughput and latency
- Use a separate physical Ethernet port for each of the four client's iperf3 traffic
- Use a different Wi-Fi 6/6E client chipset for all the above
- Use one of the many other test tools to evaluate all the above

10. Upgrading to Wi-Fi 6E and enabling OFDMA for Wi-Fi 6

Considering the testing described in this paper, when is the best time to enable OFDMA or deploy Wi-Fi 6E? The answer depends on the needs of the end-user. If the end-user is in a crowded RF environment with a lot of out-of-band Wi-Fi traffic and neighboring APs, upgrading to Wi-Fi 6E now is appropriate. OFDMA alone may not be enough to help overcome persistently crowded channels especially if it is out-

of-network traffic on the same channels in use. There are more than enough Wi-Fi 6E devices and laptops available now to take advantage of the extra channels and clean airtime if the end-user needs to solve a problem with airtime availability and can change to a Wi-Fi 6E client or Wi-Fi 6E client adaptor. The clients using Wi-Fi 6E have endured some growing pains with sparse driver update availability and security differences in how clients connect, but these are temporary issues and have already improved greatly over the last year. OFDMA enabled in 6 GHz will continue to improve as chipsets and firmware mature and much of the development and progress on-going in 5 GHz will directly carry over to 6 GHz without much delay.

In environments not overly crowded from out-of-band overlapping 5 GHz traffic, enabling OFDMA on a Wi-Fi 6 AP is a great way to allow what appears to the end user as a more responsive channel while serving the same clients and bandwidth needs in the home. This appears to be the case from the testing completed in this paper, especially if the channel is crowded from in-network traffic. Even mixed client populations of Wi-Fi 5 and Wi-Fi 6 can achieve latency improvements with OFDMA enabled if the Wi-Fi 6 clients are served quicker and simultaneously. Most of the traffic on the internet is TCP, and with TCP traffic the number of situations where latency worsened with OFDMA enabled were about the same as the situations that improved. The situations that didn't improve weren't necessarily situations that needed to improve to have a better QoE to begin with. AP chipset scheduler improvements over time should alleviate many of the issues encountered, because the channel is not contentious in many of the situations in which latency worsened instead of improving. It is yet to be determined if the same benefits observed and described in this paper are also seen with the additional test conditions listed in the Additional Tests to Consider section.

11. Conclusion

This paper sought to define a wide range of test cases to determine which scenarios with today's AP chipsets can achieve reduced latency when OFDMA is enabled. The results showed highly utilized channels stand to benefit the most from enabling OFDMA. Furthermore, higher utilization of the client in use improves the chance of realizing a latency reduction. The upstream tests with OFDMA enabled tended to show more impressive reductions in latency because the clients were not fighting each other for TXOPs in the upstream. The downstream tests did not show the same extent of improvements as were seen in the upstream. In the downstream with OFDMA disabled, the AP is already scheduling transmissions individually as the queues demanded and was less contentious as the AP was in charge in the downstream. VHT and HE MU-MIMO may be more advantageous for latency reductions in the downstream for the environments and schedulers that can use it effectively. The lower bitrates sent to or from the CUT also didn't improve latency as much as the higher bitrates tested and sometimes incurred even higher RTT latency, indicating some chipset thresholds to cause OFDMA scheduling of the CUT were not met.

Test scenarios with 10% utilization from other clients didn't have major problems with latency to begin with and didn't usually improve the latency to the CUT when OFDMA was enabled. When the airtime is highly congested, or a lot of frames are being sent to or received from many clients, the ability to win a TXOP becomes harder, and it delays the information being sent. In this setup, greatest improvements in latency were realized in high airtime usage scenarios. The greatest improvements in reduction are naturally seen when comparing to situations that have the highest latency to begin with.

In conclusion, the charts below in Figure 26 and Figure 27 summarize when an end-user may expect to see improved latency with OFDMA enabled in a house with four clients. This summary was created while generalizing the AP chipset used, ignoring differences in packet size outcomes, and generalizing results from pure Wi-Fi 6 client tests and mixed half Wi-Fi 5 client tests to provide a simpler summary

with just a couple variables. The jitter or mean deviation of RTT latency and max RTT latency were heavily considered in addition to the average RTT latency to indicate if latency was seen improving or worsening. This is because the changes were often more dramatic in max RTT and jitter of RTT samples even if the average RTT didn't change meaningfully.

The color of the bars represents the CUT throughput while the four clusters of bars in each chart represent each channel utilization scenario, indicated at the bottom of each chart. The taller the bars the larger the difference in RTT latency, one way or another. If the bar, starting from the middle, is going upward, a general improvement in latency was seen. If the bar is going downward from the middle, a general degradation in latency was seen. The middle represents about the same latency observed and would be a missing bar in this set of charts, to indicate it generally was not better or worse.




Figure 26 – Latency Change Observed with OFDMA and Downstream Traffic




Figure 27 – Latency Change Observed with OFDMA and Upstream Traffic

The test iterations were large in number and defined many scenarios and settings, but was rigid in location of clients, number of clients, type of clients, and measurement methods. The rigid test constraints served the purpose of allowing for repeatable testing as well as allowing for relative testing so results of each mode could be compared to each other. There may be some results that are only realized in a setup like this paper assumed, and other results and improvements that were not seen from perhaps moving a client five feet in another direction, using a different client, or using 20 clients.

This paper sought to identify scenarios that improve with OFDMA enabled with the setup described herein, but it is not the only setup there is. In fact, there is not another house with the same layout, furniture, clients, AP or client positioning, or any number of many other factors which could lead to different results. There is included in this paper an extensive list of follow-up testing to consider, in order to investigate the many other ways and scenarios OFDMA can help reduce latency. In this paper, a given static setup and careful adherence to keep many things constant while iterating certain chosen parameters only, allowed for relative comparisons to be made and conclusions drawn.

Abbreviations

6E	802.11ax Wi-Fi 6 in 6 GHz band
AC	802.11ac Wi-Fi 5
ACK	acknowledgement
AP	Wi-Fi access point
AX	802.11ax Wi-Fi 6
BSR	buffer status report
BSRP	buffer status report poll
CS:GO	Counter Strike: Global Offensive by Valve/Hidden Path Entertainment
CTS	clear-to-send
CSMA/CA	carrier sense multiple access with collision avoidance
CUT	client under test
DL	downlink (AP to client) direction
DL-PPDU	downlink PLCP protocol data unit
DOCSIS	Data Over Cable Service Interface Specifications
DS	downstream (AP to client) direction
DSL	digital subscriber line
FPS	first person shooter
HE	high-efficiency 802.11ax
JSON	javascript object notation
LAN	local area network
Mbps	millions of bits per second
MCS	modulation coding scheme
mdev	median deviation
MHz	millions of hertz
MMO	massive multiplayer online
MSO	multiple systems operator
MSS	max segment size tcp data payload
MU	multi-user
MU-MIMO	multi-user multiple input multiple output
MU-RTS	multi-user request-to-send
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OWD	one-way delay
PLCP	physical layer convergence procedure
PPDU	PLCP protocol data unit
QoE	quality of experience
QoS	quality of service
RTS	real-time strategy
TCP	transmission control protocol
TXOP	transmit opportunity
UDP	user datagram protocol
UL	uplink (client to AP) direction
UL-PPDU	uplink PLCP protocol data unit
VHT	very high throughput 802.11ac
WAN	wide area network

Bibliography & References

- [1] https://www.bitag.org/documents/BITAG_latency_explained.pdf
- [2] ISBN: 978-1-119-80787-2: Wi-Fi 6 & 6E For Dummies®, Extreme Networks Special Edition; David Coleman
- [3] <https://hpbn.co/primer-on-latency-and-bandwidth/>
- [4] <https://www.cox.com/residential/internet/guides/gaming-performance/ping-testing.html>
- [5] <https://parsec.app/blog/description-of-parsec-technology-b2738dcc3842>
- [6] <https://docs.microsoft.com/en-us/skypeforbusiness/optimizing-your-network/media-quality-and-network-connectivity-performance>

The Cable Home is the Wellness and Telemedicine home – lets now deliver these new solutions

A Technical Paper prepared for SCTE by

Sudheer Dharanikota
Managing Director
Duke Tech Solutions, Inc.
111 Fieldbrook Ct. Cary NC 27519
+1 919 961 6175
sudheer@duketechsolutions.com

Charles Cheevers
CTO, Home Networks
CommScope Inc.
charles.cheevers@commscope.com

Table of Contents

Title	Page Number
1. Executive summary	3
2. Introduction.....	3
4.1. Hardware and software components	6
4.1.1. Device facing interface.....	7
4.1.2. Network connectivity interface	8
4.1.3. In-home controller interface	10
4.1.4. Internal functionalities	11
4.2. Inter-industry collaboration interfaces	11
5. Conclusions and recommendations	12
6. Bibliography & References.....	13

List of Figures

Title	Page Number
Figure 1 Telecom for Healthcare opportunity and challenges summary	4
Figure 2 DTS's Telecom for Wellness Environment Framework (DTEF) based components.....	5
Figure 3 Sensor Network Gateway architecture	6
Figure 4 Unifying the wellness resources with the in-home and access networks.....	7
Figure 5 Making a home an extension of Telehealth services through standardized interfaces.....	12

1. Executive summary

Telecom for Wellness (T4W) is a multi-trillion-dollar opportunity for cable operators. This opportunity begins with the connectivity and device base that is already in place and expands simply to move into new high-value revenue opportunities in the wellness, telemedicine, and aging-in-place markets. In this paper, we analyze different in-home networking architectures that can be added to the existing Broadband and Wireless networks in the home to support T4W services such as Aging in Place (AIP) and Telehealth. To minimize CapEx investment for this new service evolution we will show how the existing infrastructure can support sensor-based networks, audio networks, video networks, and IoT for medical device networks in a simple incremental architecture from today's quad-play services. The paper will provide an analysis of the home networking components needed for the AIP and Telehealth applications to overlay into existing backend services, processes, and even technician in-home support. This paper will leverage the work done within SCTE Working Groups to derive an outline of the standardization of cable industry premises devices with wellness capabilities used for T4W services. It will highlight the interfaces for wellness solutions to be added to Broadband Gateways, STBs to provide access to resources like BLE and Wi-Fi, such as simple secure onboarding, messages and notifications on the TV / Video and Audio networks. It will also look at the cloud-to-cloud interfaces that can be standardized to allow partnerships between Service Providers and Cloud wellness partners to create inter-industry opportunities.

2. Introduction

The Wellness industry is going through a major transformation to modernize the infrastructure, reduce the cost and increase the quality of care. In a series of articles, we have suggested how the Telecom industry can assist the Wellness industry (Refer to [1], [2], [3], [4], [5], [6], [7]). We call this inter-industry collaboration Telecom for Wellness (T4W). Even though the T4W opportunity is not limited to these two major intersection points ([13]), we focus on Aging in Place (AIP) and Telehealth use cases to illustrate our thoughts on the end-to-end T4W architecture in [8]. (Refer to [9] for six different opportunities that a Telecom operator can address through the T4W architecture covered in this paper.) The SCTE Data Standards Subcommittee, of which the authors are members, is actively working on T4W solutions for the AIP and Telehealth areas in working groups three [10] and four [11].

[1], [2] provides a quick summary of the T4W opportunity and challenges from AIP and Telehealth points of view. Many of the needs, challenges, and Telecom opportunities of both markets are similar (refer to the SCTE working group analysis at [10], [11]). Some of the high-level use cases that need to be supported for these two markets include:

- A. Providing basic communication between the users and the providers/caregivers
- B. Providing seamless communication between the users and the stakeholders
- C. Monitoring the users for health, mobility, fall detection, etc.
- D. Analyzing the data collected from the users and properly notifying the stakeholders
- E. Assisting the T4W service providers with claims by documenting accountability
- F. Offering managed services to support installations, product support, and other services to improve adoption and retain customers



	 Aging in Place	 Telehealth
Users	Older adults (65+), caregivers	Individuals, providers
Stakeholders	Family members, care givers, doctors, service personnel etc.	All family members, providers, (payers)
Needs	Communicating, monitoring, service, support, integration	Communicating, monitoring, integrating with provider systems
Challenges	Ease of use, provider network integration, problem solving	Ease of use, device and EMR integration, remote monitoring,
Telecom opportunity	End to end solution, managed services, provider integration	End to end solution, managed services, provider integration

Figure 1 Telecom for Healthcare opportunity and challenges summary

Many of these use cases are elaborated on in [9]. In the use cases paper, we presented AIP, independent living, and hospital at-home use cases (the extreme ends of the T4W) with the consideration on what are the opportunities for the cable operators.

In the next sections, we summarize the T4W architectural needs, provide a framework, discuss individual components, and start discussing the **sensor network gateway** concept.

3. End-to-end high-level T4W architecture

Figure 2 provides a high-level end-to-end architecture proposed by Duke Tech Solutions (DTS) in their market analysis [13] based on different T4W market opportunities. The framework is elaborated in [8].

To understand the end-to-end T4W architecture, first, we need to understand the users, the service providers, and the other stakeholders as shown in Figure 1. We adopt the architectural framework provided in [8], DTS's Telecom for Wellness Environment Framework (DTEF), to evaluate the in-home components proposed in this paper. We will use AIP and Telehealth use cases (as provided in section 2) to do this.

1. In-home healthcare/wellness aware gateway: From use cases A, B, and C, it is clear that there needs to be a gateway in the T4W home. This gateway, as shown in Figure 2, acts as an integration point for monitoring the sensor devices (e.g., motion sensors, remote patient

monitoring equipment) and integrating with the interactive services endpoints (such as unified communication services). This **Sensor Network Gateway (SNG)** can be a standalone device or integrated with other vendor equipment such as the set-top box or residential gateway. In this paper, we treat it as a logically separate device.

2. **T4W aware network infrastructure:** Again from the use cases A, B, and C it is clear that the T4W requires connections between the users, the providers, and the other stakeholders. This requires not only reusing the existing telecom infrastructure but will also need to meet reliability, security, and privacy requirements specified by the T4W architecture. The communications infrastructure will have to meet the needs of the sensor network traffic, unified communications traffic, and notifications to the different stakeholders. To differentiate (or to keep the focus on) the T4W needs, we call this the **T4W Sensor Network Infrastructure**.

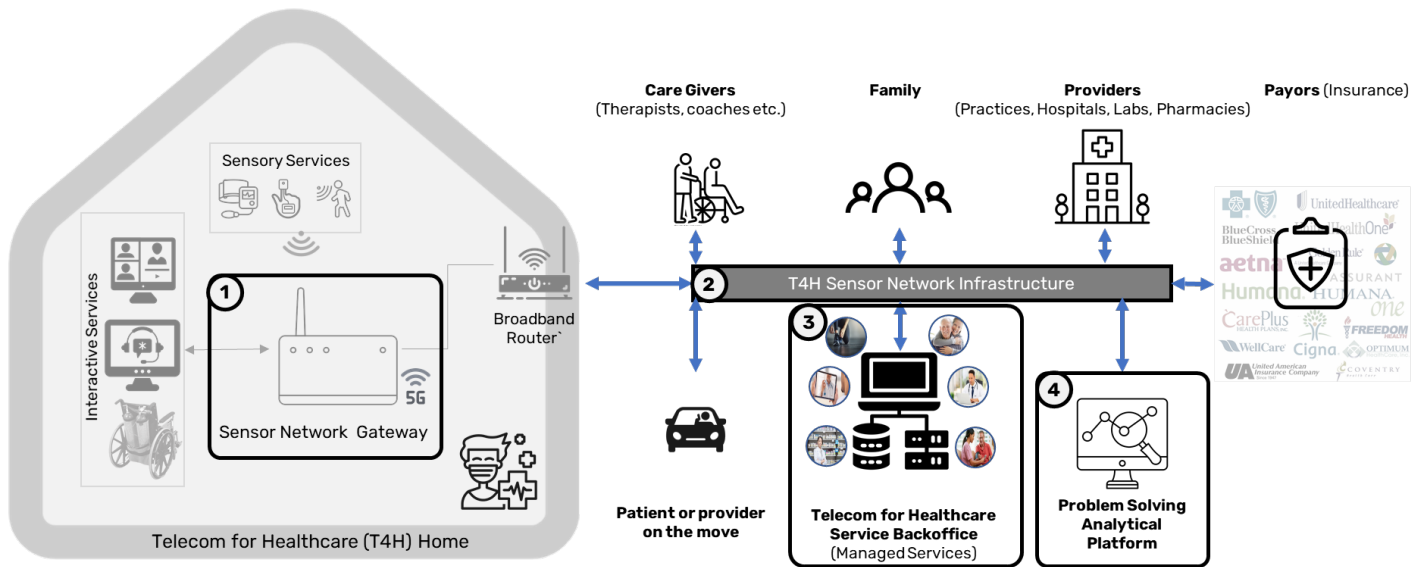


Figure 2 DTS's Telecom for Wellness Environment Framework (DTEF) based components

3. **T4W aware service back office:** The cable operators have all the required infrastructure for managing end-to-end services. As mentioned in use case F, it is essential to turn the fragmented, gadget-oriented point solutions into a well-oiled managed service. This can only be accomplished by Telecom operators who have access to such infrastructure and have been managing communications infrastructure for 90+% of the households in the US. We call such infrastructure as **T4W Service Backoffice**.
4. **T4W aware problem solving analytical platform:** Finally, as mentioned in use cases D and E, this infrastructure attempts to solve the problems stakeholders are facing. These problems and related algorithms may be unique to the healthcare/wellness industry, but the infrastructure is similar to the infrastructure the telecom operators use today. We call this repurposed analytical platform the **T4W Problem Solving Analytical Platform**.

In the following sections, we highlight the sensor network gateway's hardware, software, and cloud interface needs to enable inter-industry opportunities.

4. Sensor network gateway components

Figure 3 shows different tasks that need to be done in a T4W capable home, as briefly discussed below:

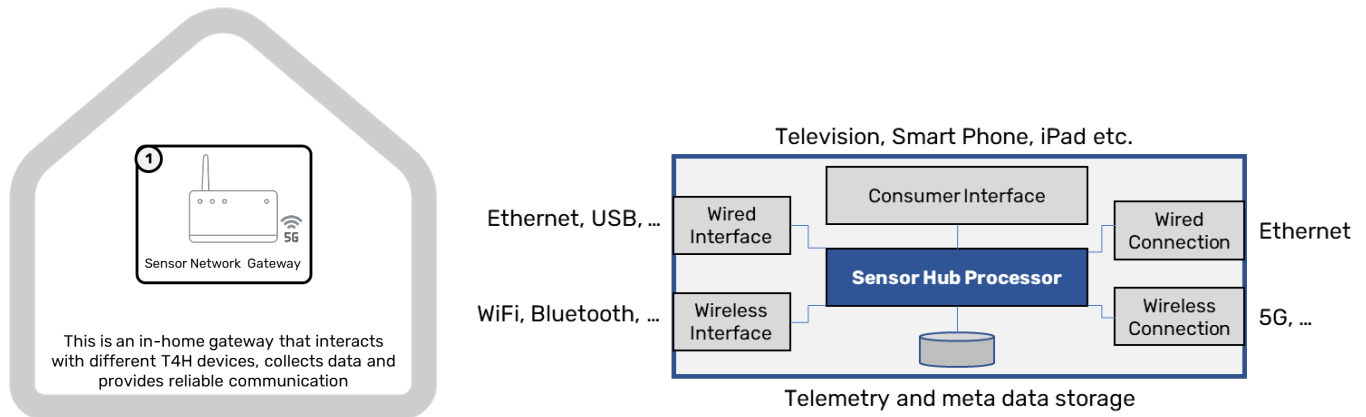


Figure 3 Sensor Network Gateway architecture

- Support for different data streams: The AIP and Telehealth infrastructure needs to support typical data streams generated in a T4W home. These include sensor and actuator data streams, streams to record events, and real-time streams such as video and audio communication between T4W stakeholders.
- Communication with existing in-home broadband devices: These might include consumer consoles (such as TVs or smart speakers), and smart home devices (such as smart locks, lights, or video doorbells). To increase the adoption of T4W solutions and for ease of use, the T4W in-home components need to be on the same logical network.
 - The T4W physical networking can be dependent upon the use case for any particular device. Most components are likely best connected with an in-home broadband network, but certain devices (such as a locator device) may need to be connected even if the user is beyond the limits of the in-home network. The important thing is that the networked devices can communicate with each other on a secure logical network. Critical components may require a secondary backup network connection in case the primary network fails. Cloud-based services also protect against exclusive dependency on the in-home network.
 - The in-home solutions shall integrate remote patient monitoring devices, sensor devices (such as fall detection, motion sensors, etc.), and other IoT devices that are used for wellness needs.
 - Additionally (to increase the utility and ease of use of the system), the T4W solutions shall be integrated with the frequently used consumer consoles (such as Television for the elderly), smartphones, and other handheld devices. Again, a cloud-based solution simplifies an experience that can be duplicated on whatever console is convenient.
- Provide installation and support services: The operator shall also streamline the installation and support services to improve the ease of use of the integrated solution.

4.1. Hardware and software components

As shown in Figure 4, the cable operators already have been deploying many standardized protocols, interfaces, and devices for different in-home solutions. Extending them to T4W solutions will be a natural extension for them. But they have to offer many subtle features to be adopted by the wellness industry. These extensions are elaborated on in the following section. The required tasks of the solution are broken into the device-facing interface, network connectivity interface, in-home controller interface, internal

capabilities, and inter-industry collaboration interfaces (represented by the Cloud2Cloud interface in the figure).

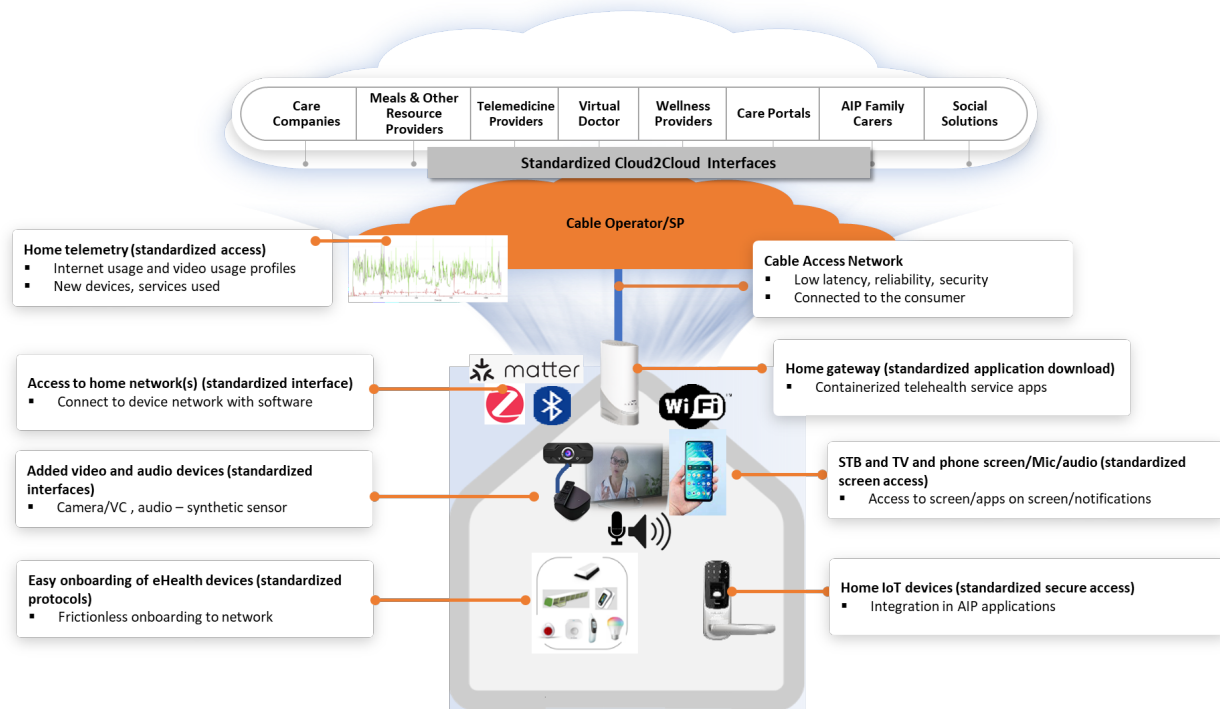
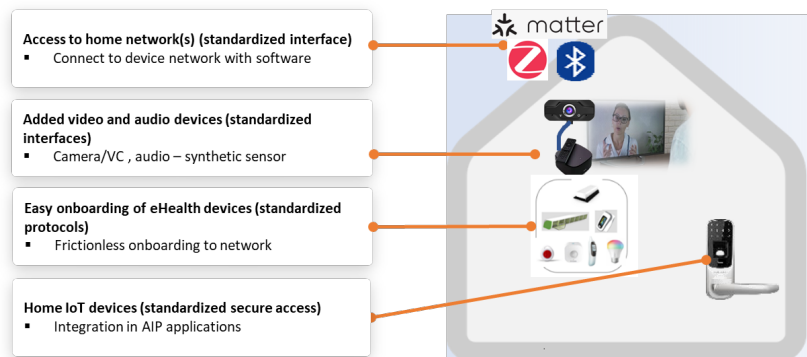


Figure 4 Unifying the wellness resources with the in-home and access networks

4.1.1. Device facing interface

Based on the above high-level needs, we propose that a Sensor Network Gateway (SNG) functionality be developed for supporting T4W solutions. The block diagram of such a gateway is presented in Figure 3 with a bit more details and how it interacts with other devices visually in Figure 4. This gateway will have wired (Ethernet, USB, etc.) and wireless (Wi-Fi, Bluetooth, BLE, etc.) interfaces to integrate the T4W devices and other IoT devices (such as turning on a light, placing a phone call to a family member). Note that if the SNG functionality is integrated with the residential gateway these requirements apply to the residential gateway.



Functionalities

At a broad level, there will be two types of southbound facing devices to the SNG. These are the two-way communication devices such as the devices used for Telehealth and the one-way communication devices such as sensor devices used typically for monitoring. The following discussion applies to both such categories.

The SNG should have the device-facing capabilities to:

- **Register** a new device and **onboard** with proper credentials before letting it use the network
- Learn the **capabilities** of the device as part of the device registration
- **Authenticate** a device against the registered credentials
- **Detect** the wired and wireless devices when they either come alive or connected
- **Securely** communicate with the device for different information exchanges for status gathering, data collection, etc.
- Identify the **status** and different levels of availability of the device during the communication.
- **Collect** on-demand and continuous **data** from the device

Types of interfaces

The SNG will act as a gateway to the T4W applications. This requires the SNG to support some of the following interfaces. Note that only some of the interfaces are required depending on the positioning of the SNG. It is not in the scope of this document to provide the number of such interfaces.

The SNG should support the device-facing (Note that it is not in the scope of this paper to recommend one protocol or the other):

- **Wired** interfaces such as Ethernet or USB connection with potential hub functionality
- **Wireless** interfaces such as Wi-Fi, Bluetooth, BLE, Matter, etc. with potential hub functionality

Serviceability

The serviceability of the SNG and its subtended devices are essential for making the T4W services more sustainable. In this section, we highlight the serviceability needs of the subtended devices.

The SNG should support different protocol serviceability constructs such as:

- **Monitoring the reachability/connectivity** status of different wired and wireless devices
- **Debugging the connectivity** and other **serviceability** issues for the subtending devices

Other requirements

Other requirements from the SNG towards the device interface include:

- **Charging** the connected devices where possible for high-priority devices, and

4.1.2. Network connectivity interface

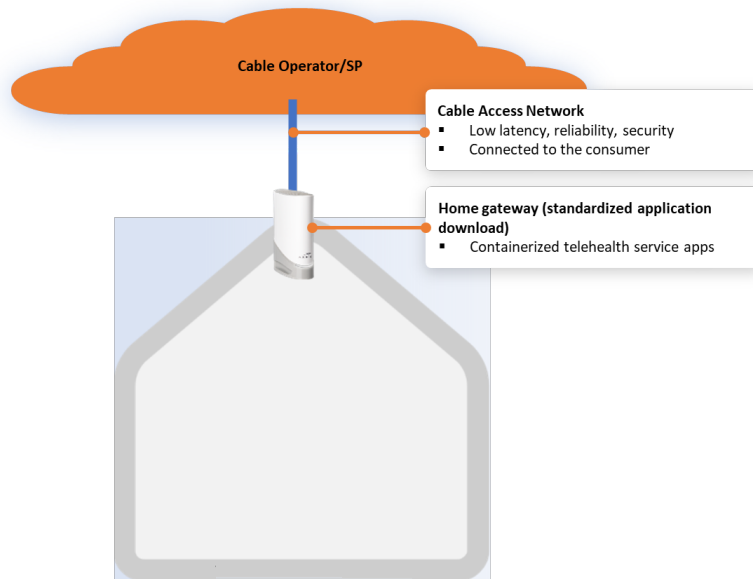
The gateway shall provide high-quality redundant Internet connectivity such as with an Ethernet interface to wired broadband and a 5G or other wireless connection as a backup.

Functionalities

The SNG will have northbound facing internet connectivity for providing access to the network or for in-home devices. The following discussion applies to both such categories of connectivity.

The SNG should have the network-facing capabilities to:

- Provide **high-quality differentiated communication** based on the priority of the application
- Offer differentiated services to both for **one way and two-way communications**
- Provide **high availability connectivity** based on the importance of the communication (for example, backup the cable connectivity with the 5G connectivity for some of the sensor devices as opposed to the Telehealth video communication solutions)
- Provide per **session-based secure communications** at the device level (Note that there may be multiple stakeholders for every session)



Types of interfaces

The SNG should support Quality of Service (QoS) supported redundant interfaces for high availability connectivity, such as:

- A **wired primary interface** for high QoS enabled interface (for example, Ethernet, DOCSIS, PON based)
- A **wireless secondary interface** (for example, 5G interface) for high-priority traffic communication during primary interface failure
- A **failover mechanism** that provides the service continuity protection for high-priority traffic during primary interface failures

Serviceability

The SNG should support different techniques to monitor and test service continuity, such as

- Remote access loopbacks (e.g., using 802.1x)
- Service level network connectivity polling

Interaction with other northbound interfaces

In addition, the SNG should interface with different in-home components, such as

- The **wiring closet**, as discussed in [14]
- The **broadband router**, if SNG is not integrated into the router, and

- An external or internal device that provides a failover redundant communication

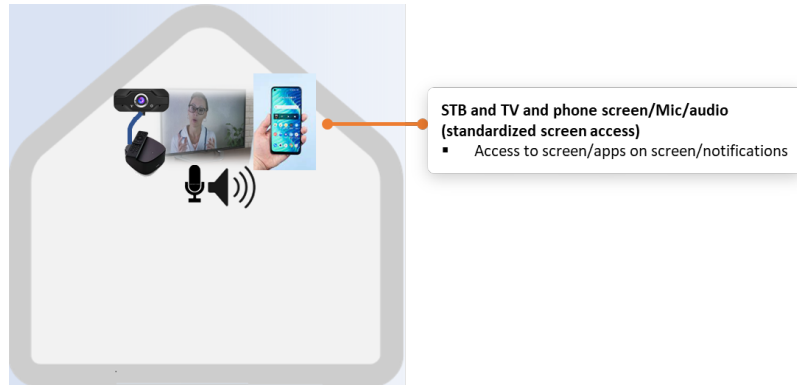
Other requirements

The SNG and its downstream applications may also require an accurate time of day and date. This can only be achieved if the SNG has access to the

- Network timing information through the timing protocols as a pass-through from the router

4.1.3. In-home controller interface

The SNG will also need to integrate the consumer access interfaces (TVs, smartphones, iPads, etc.) to increase the adoption of the T4W applications.



Functionalities

The SNG functionality and many other T4W solutions require some of the following functionalities for ease of use. The SNG should

- Provide **video console** (such as TVs, iPads, etc.) **integration** to offer as a console to the management of the T4W services.
 - This includes all the unified communication aspects of the interactions such as the notifications, two-way conversations, and certain process flows such as interrupting the activity that is being performed on the video console.
- Interact with an **easy-to-use remote controller** to manage the on-screen process flows
- Offer, where possible, **voice-enabled interactions** with the SNG
- Be able to **split the feeds among the stakeholders, on-demand** basis, based on the stakeholder privileges

Types of interfaces

The SNG shall offer different interfaces to communicate with the in-home controller device. It should

- Provide a **wired interface** to connect with the video console (such as USB, Ethernet, etc.)
- Provide a **wireless interface** to connect with the video console (such as discovering and connecting to the Smart TV over Wi-Fi)

Serviceability and debugging

The SNG should extend the service continuity evaluation to the in-home controller interface for ease of service assurance.

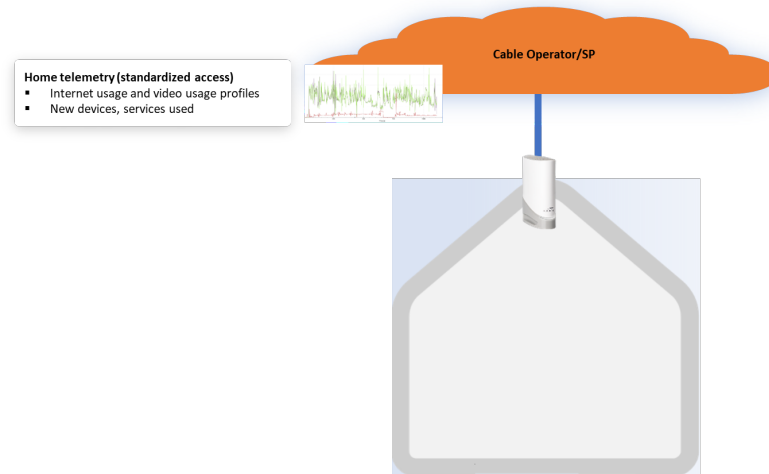
4.1.4. Internal functionalities

The SNG will need limited internal storage for the temporary storage of sensor data and to perform local analytics on time-critical events. The gateway shall offer an easy installation process and support self-install where possible.

Functionalities

The SNG should

- **Register** the gateway device with the T4W platform for different service capabilities that it can enable and is allowed to perform
- Provide support for the **unique addressability of the device** (IP address, etc.)
- Be capable of providing different **QoS** marking, prioritization, and fulfillment capabilities
- Be capable of **self-installing** during initial bootstrapping and later during the upgrading
- Provide a **management interface** to configure and manage different features of the device
- Provide **extensive stats collection** capabilities as highlighted in [12]
- Support some **basic analytics** to perform local analysis and notification capabilities
- **Store and provide access** to the information collected for at least **one day** in case of loss of internet connectivity



Types of interfaces

The SNG should

- Support **service** (a physical or a logical) **interface** to access the management interface
- Support a local **interface** for the field technician if the device is not reachable through the internet
- Allow **access to different debugging tools and statistics** that are collected on the gateway

Other requirements

The SNG should support

- AC and DC power with potential for battery backup based on the scope of the product

4.2. Inter-industry collaboration interfaces

The SNG is the customer-facing part of the end-to-end Telecom for Wellness solutions as shown in Figure 2. A cable operator or someone working along with the operator can create a platform with these components. Once the platform is created, it can be used to enable different stakeholders and solve their problems, and also can be used as a conduit for different services from other value-added wellness-related service providers.

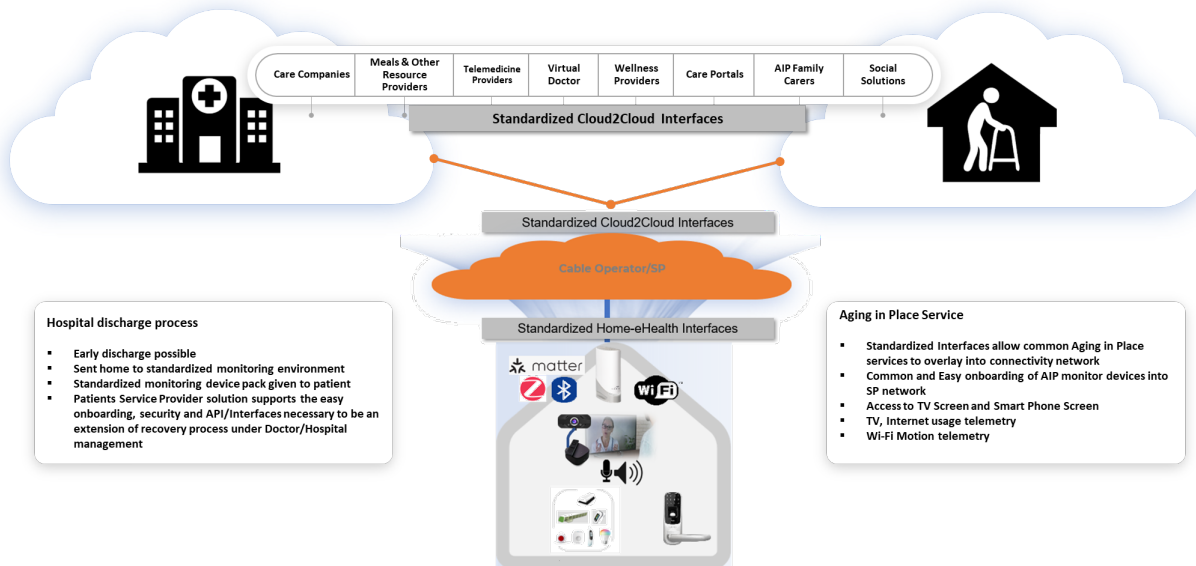


Figure 5 Making a home an extension of Telehealth services through standardized interfaces

The T4W provider can offer this platform for internal use and to the marketplace of the service providers through a set of services. Figure 5 provides two such use cases that can be enabled using an inter-industry collaboration. More information on these use cases and the opportunity for the cable operators are detailed in [9]. The SNG will in this case work as the subtended device in the house for the set of newer services. In this case, the SNG and the cloud platform shall support the following functionality through different service enablement constructs:

- Faster service enablement
 - **Onboarding** newer service providers
 - Establishing the service capabilities with the **same look and feel**
 - Creating constructs for **single billing** capabilities
 - **Extending** the proposed **serviceability interface** to the external service providers
- Knowledge sharing
 - Enabling **access to different monitoring** points based on the registration
 - Creating different status monitoring and **data collection points per service** (or provider)
 - Setting threshold crossing alarms and providing **access to those notifications**
- Giving access to the T4W platform's problem-solving infrastructure based on the service agreements
- Enable cloud-to-cloud interface between service providers through standardized interfaces

5. Conclusions and recommendations

As presented in the previous papers, T4W is a multi-trillion dollar opportunity for cable operators. In this paper, we extended the details of the DTEF framework's in-home component, namely the SNG, requirements to enable T4W. These requirements highlighted can be implemented in a standalone device or the functionalities can be integrated with one of the intelligent in-home devices (such as a residential gateway or a Set Top Box). The cable industry and partners are uniquely positioned to standardize and

implement such solutions with their ubiquitous presence, in-home deployments, service management workforce, and end-to-end infrastructure. To bring this opportunity to life, we recommend developing standardized SNG solutions that

- Enables different one-way and two-way communicating devices,
- Offers standardized, redundant north-bound interface,
- Provides an in-home customer (and their service) management interface,
- Offers constructs to different data collection, problem-solving, and management of the device
- Provides standardized cloud-to-cloud interfaces to enable inter-industry collaborations

We would like the cable operator ecosystem to actively engage with us at the SCTE working groups, implement these solutions and enable the Telecom for Wellness opportunities.

6. Bibliography & References

- [1] Sudheer Dharanikota, Ayarah Dharanikota, *Why are cable operators natural fit to support Telehealth? An inter-industry perspective*, SCTE Expo, 2020
- [2] Ian Wheelock, Charles Cheevers, Sudheer Dharanikota, *The Business Case for Aging in Place with Cable Operators*, SCTE Expo, 2020
- [3] Sudheer Dharanikota, Ayarah Dharanikota, Dennis Edens, Bruce McLeod, *Aging in Place business case for cable operators*, SCTE Journal, June 2021
- [4] Sudheer Dharanikota, Ayarah Dharanikota, Dennis Edens, Bruce McLeod, *Aging in Place Market Landscape from a Cable Operators Perspective*, SCTE Journal, June 2021
- [5] Sudheer Dharanikota, Ayarah Dharanikota, Dennis Edens, Bruce McLeod, *Telehealth business case for cable operators*, SCTE Journal, September 2021
- [6] Sudheer Dharanikota, Ayarah Dharanikota, Dennis Edens, Bruce McLeod, *Telehealth Market Landscape from a Cable Operators Perspective*, SCTE Journal, September 2021
- [7] Sudheer Dharanikota, *Summary of Telecom for Wellness interviews*, Oct 2021, available at <https://duketechsolutions.com/telecom-for-healthcare-SCTE/>
- [8] Sudheer Dharanikota, Clarke Stevens, *End to End Telecom for Healthcare Architecture – A Cable Industry Perspective*, SCTE Expo, 2021
- [9] Clarke Stevens, Sudheer Dharanikota, *Aging in Place and Telehealth Use Cases from the Cable Operator Perspective*, SCTE Journal, March 2022
- [10] Data Standards Subcommittee, Working Group 3, Aging in Place
- [11] Data Standards Subcommittee, Working Group 4, Telemedicine
- [12] Sudheer Dharanikota, Jason Page, *Metadata/Telemetry support from Cable Operators to address Telecom for Healthcare opportunity*, SCTE Expo, 2021
- [13] Duke Tech Solutions market Research, *Telehealth market report – A Telecom-based opportunity analysis*, available [here](#)
- [14] Rajesh Abbi, Charles Chapman, Sudheer Dharanikota, Kyle Haefner, Clarke Stevens, *Requirements for the IoT Infrastructure in the Customer Premises*, SCTE Expo, 2022

The Coming Convergence of Broadband, Energy, and Transportation

A Technical Paper prepared for SCTE by

Ralph W. Brown
Founder
Brown Wolf Consulting, LLC
1355 S Foothills Hwy
Boulder, CO 80305
+1-303-517-6711
ralph@brownwolfconsulting.com

Scott Caruso, VP Strategic Ventures, CableLabs[®]

Hunter Albright, Co-Founder, Curve10

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Background	4
3. Key Drivers of Convergence	5
4. The Point of Convergence Across Contexts	6
4.1. Convergence at the Home	7
4.1. Convergence in Urban Areas	7
4.1. Convergence in Rural Areas	8
5. Conclusion.....	9
Abbreviations	10

List of Figures

Title	Page Number
Figure 1 - Point of convergence of broadband and energy at the home	5
Figure 2 - Point of convergence of broadband and energy in urban areas	7
Figure 3 - Point of convergence of broadband and energy in rural areas	9

1. Introduction

The coming convergence of broadband, energy, and transportation has the potential to transform all three industries in unexpected and powerful ways – with significant impact on CAPEX, OPEX, and value creation opportunities. However, the potential of this transformation cannot be fully realized without successful collaboration across these three industries. The purpose of this technical paper is to raise awareness of this coming convergence and to stimulate conversations across these industries to take advantage of what, to this point, may be nonobvious synergies.

In the United States, the Infrastructure Investment and Jobs Act (IIJA) represents a once in a generation investment in infrastructure for these three sectors and a unique opportunity to speed this convergence by creating an opportunity to completely reimagine the country's underlying infrastructure. Historically, each industry has developed its infrastructure independently. However, as broadband has been a key driver in the transformation of many industries, as they digitize and move to the cloud, rethinking shared infrastructure can change the capabilities that the three industries can bring to market. It is the increased connectivity, reliability, bandwidth, and security of broadband that is enabling innovation at scale across virtually all sectors, not only in the U.S., but globally. Most importantly, the Energy and Transportation sectors are already poised for massive disruption and transformation and, as we will show, broadband is and will be a key enabler for this.

This technical paper will identify the key drivers of this cross-industry convergence and explore how it will come about by focusing on the three most important contexts for this convergence:

- Home: How will convergence impact services and behaviors in the home?
- Urban: How will cities and suburbs change and develop?
- Rural: How will rural areas adapt and thrive?

This technical paper serves as a call for cross-industry collaboration on the opportunities of this cross-industry convergence. As the Broadband Industry's innovation engine, CableLabs has decades of experience in bringing industry-transforming innovation to market through collaboration within and across industries and looks forward to uncovering and collaborating on nonobvious synergies across the industries.

One obvious area where cross-industry collaboration could stimulate innovation are the common expectations of communications capabilities that are limited to the “lowest common denominator”, as opposed to envisioning the possibilities and needs of the future. The cable industry is in the midst of its 10G [10 Gbps] industry initiative¹, addressing not only speed, but reliability, security, and latency. What applications and services could be developed for the transportation and power industries if they were not constrained by the limits of today's broadband technology?

This technical paper identifies several possible synergies and opportunities. We are certain that you will envision many more. We welcome your feedback or interest by contacting the authors here: [Ralph W. Brown <ralph@brownwolfconsulting.com>](mailto:ralph@brownwolfconsulting.com), [Scott Caruso <s.caruso@cablelabs.com>](mailto:s.caruso@cablelabs.com), and [Hunter Albright <hunter@curve10.com>](mailto:hunter@curve10.com).

¹ See: CableLabs 10G Platform: <https://www.cablelabs.com/10g>

2. Background

Perhaps the most important insight into this coming convergence is from an understanding of how broadband access network technology (both fixed and wireless) intersects with the power grid and the shared opportunities that exist. Today, energy companies have a lack of visibility into the performance in the last mile of their power distribution grid. On the contrary, while less well known, broadband access network providers are consumers of energy from the grid, particularly in the last mile, which provides great insight into the last mile of the grid. To deliver broadband services to consumers and businesses, broadband access networks have active network elements that are powered by the grid. To provide broadband services in periods where there may be a power outage, broadband network operators integrate resilient, battery-backed power systems to keep these network elements operational (see figure below). These points of connection between the broadband network and the power grid provide unique insight into the performance of the power grid where power companies lack visibility. As an example of what is possible from these points of connection, Gridmetrics, Inc.² leverages them to provide visibility into the status and performance of the power grid at a geographical and temporal resolution not previously available.

The figure below provides an example of the convergence of broadband and energy utilized by Gridmetrics. It shows the power distribution grid along with a Hybrid Fiber-Coax (HFC) network and mobile small cell. It also shows their connection to the home with an EV charging station and potentially distributed energy sources. The HFC node and the mobile small cell are among the active elements in the fixed and mobile networks. The HFC network can act as a backhaul network for the mobile small cell and provide broadband services to the home. A backup power system monitors the power grid and provides backup power to both the HFC node and mobile small cell in case of power outages. These backup power systems are connected to the broadband network to report on their operational status. The backup power system can not only monitor the availability of power from the grid but can also monitor the quality of the power signal and more. Because this is part of the broadband access network, it can harness this broadband connectivity to aggregate information where it can be analyzed. This aggregated power measurement data provides high quality, fine grain insight into the health of the power grid. With the growth in distributed energy resources (DERS) this insight will be critical for the management of the power grid. It should be noted that the EV charging station at the home or the distributed energy resources (e.g., battery, solar, wind) shown in this figure can serve a similar function for monitoring the grid if they are connected to the broadband network and have some form of backup power.

² See *Gridmetrics White Paper* by CableLabs[®] and Gridmetrics[™] and <https://gridmetrics.io/>




Figure 1 - Point of convergence of broadband and energy at the home

In subsequent sections we will show how this point of convergence between broadband and energy is extended to include transportation and how it fits into the Urban/Suburban and Rural contexts. First, we will cover the key drivers of this coming convergence of broadband, energy, and transportation.

3. Key Drivers of Convergence

There are multiple drivers of the convergence of broadband, energy, and transportation, including:

1. Two-Way Flow - Grid Modernization, the digital transformation to enabling the bi-directional flow of electrons along the distribution power grid
2. Distributed Energy - The shift from centralized to distributed energy resources (e.g., solar, wind, battery, etc.)
3. Charging Infrastructure - The electrification of vehicles and the deployment of the supporting EV charging infrastructure, both at home and through commercial EV charging networks
4. Energy Transactions - The evolving EV charging requirements for connectivity to support “reservations and payments”
5. Transactive Energy - The evolving Transactive Energy Ecosystem that is tightly coupled to a modernized grid and EV needs
6. Electrification - The shift from natural gas to electric for traditional uses of natural gas (e.g., heating, hot water, cooking, etc.)
7. Autonomous Vehicles - The evolution of autonomous vehicles with the resulting transformation of public transit and shipping
8. Cyber Threats - The need to address cyber threats to critical infrastructure
9. Broadband Demand - Increasing demand for ubiquitous broadband connectivity with the pressing need to address the digital divide and rural broadband
10. Investment - Substantial investment resulting from the Infrastructure Investment and Jobs Act (IIJA)

Any number of these drivers could be sufficient to stimulate convergence, the combination of them all makes it essential that we envision this convergence now and take the appropriate steps to realize it in the coming years. Any investment in infrastructure should anticipate the future vision for these industries.

4. The Point of Convergence Across Contexts

It is important to understand where the physical point of convergence occurs across these three industry segments within the identified contexts. The insight here is that the power grid and the broadband network physically come together at those points where a broadband network connection is made that draws power from the grid, e.g.:

- At the home with the broadband modem/gateway, home EV charging stations, and/or distributed energy resources, such as solar and/or battery
- At public, commercial or fleet managed EV charging stations
- At the access points for commercial and MDU buildings
- Along roads and highways where Vehicle to Infrastructure (V2I) small cells are placed

It is these points of connection that enable the convergence of the broadband, energy, and transportation industries. To further illuminate the convergence, consider the deployment of an EV charger, either for use at home, commercially or publicly. Along with supplying power to the charging station, implicit is a communications link. Communications is essentially required to monitor and manage the EV charging (scheduling, notifications, troubleshooting, payments, etc.) The communications component is understood to be part of the delivered power charging solution, and with forethought and collaboration, the opportunity exists to create a more secure, robust, and efficient charging solution as well as a comprehensive data exchange capability. There is a natural, symbiotic interdependence among the industries. Consider:

- Real time monitoring of the quality of the power available and delivered to a vehicle
- Manage the sources of power available to the charging stations
- Scheduling charge times, particularly to optimize costs of power
- Reserving charging stations (public or commercial)
- Payments for on-demand charging
- Securing the charging stations, including video observations
- Providing data uploads to vehicles (including manufacturer specific updates, on-board entertainment, current navigation data, etc.)
- Providing data downloads from vehicles (telemetry data to manufacturer, insurance companies, fleet operators, etc)

Underlying each of these considerations is the communications network. Considering the protocols, security, resilience, and reliability of the communications networks as a key component of the complete solution benefits the entire EV charging ecosystem. In essence, addressing the complete solution with a holistic view can prevent the creation of inherent silos that fail to provide the optimal power delivery, reduce the costs, improve the software interfaces for the EV manufacturers and their consumers, ensure the security of the charging stations and its users and secure the power transactions. Minus an intentional desire to collaborate, the likely outcome is fractured solutions designed by the EV charger suppliers, car manufacturers, state/local specific solutions. The loser is the consumer who must manage the various interfaces with inconsistent outcomes.

4.1. Convergence at the Home

A primary point of convergence and one that is taking on increased importance with the changes in where and how people work, is the home. Figure 1 shows the diversity of assets that convergence at the home and we can expect these to grow.

The convergence of vehicle charging and distributed energy sources at home enable:

- Grid monitoring at the home
- Grid management for home-based energy sources (solar, battery, etc.)
- Automated Demand Response (ADR) for the home
- Secure broadband services for both consumer, grid monitoring/management, and autonomous EV connectivity (stationary)
- Consumer & commercial visibility into energy consumption & contribution
- Two-way power transactions between grid and EVs

It should be noted that there are also potential construction synergies (dig once for resiliency) in green field and new housing developments or in areas where investment is being made to upgrade or harden the grid and broadband infrastructures.

4.1. Convergence in Urban Areas

Moving outside the home, there are significant opportunities for convergence in urban areas across private, public, and fleet assets. The figure below shows the points of convergence across Urban areas.




Figure 2 - Point of convergence of broadband and energy in urban areas

In this context, V2I communication enables:

- Addressing traffic safety, congestion issues, and traffic signal management
- Parking and fleet management

- Transformation of public transit
- Transformation of package delivery
- Broadband connectivity to autonomous vehicles (in motion)
- Grid monitoring at the transportation V2I small cell location

Urban EV charging infrastructure enables:

- Grid monitoring at the charging station
- Security monitoring at the charging station (safety concern)
- Secure broadband services for autonomous EV connectivity (stationary)
- Two-way power transactions between grid and EVs

Vehicle charging and distributed energy sources at the commercial building or MDU enables

- Grid monitoring at the enterprise
- Grid management for enterprise-based energy sources (solar, battery, EV, etc.)
- Automated Demand Response (ADR) for the enterprise
- Secure broadband services for enterprise
- Secure broadband services for grid monitoring/management and autonomous EV connectivity (stationary)
- Commercial visibility into energy consumption & contribution

Similarly, there are construction synergies (dig once for resiliency) during construction or maintenance of transportation infrastructure.

4.1. Convergence in Rural Areas

Lastly, there will be convergence in rural locations that will be driving by engagement with devices while in transit and shifting population with a reimagining of the future of work with recent technology and behavior shifts. The figure below shows the points of convergence across Rural areas.

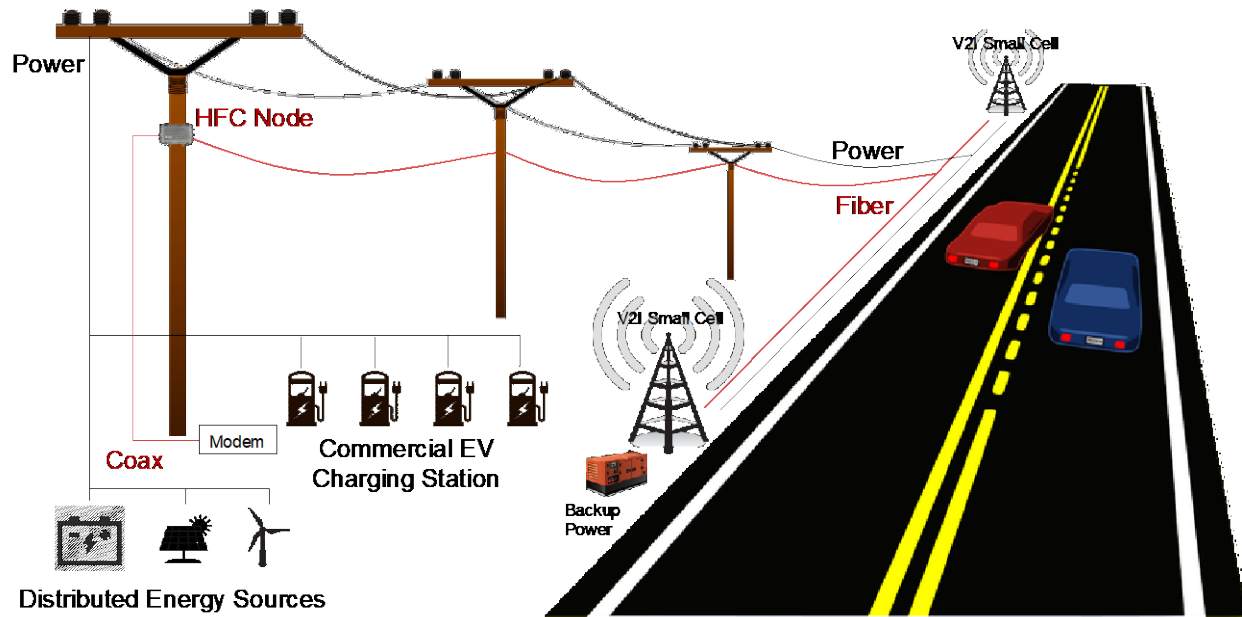


Figure 3 - Point of convergence of broadband and energy in rural areas

Like the Urban and Suburban contexts, V2I Vehicle to infrastructure communication enables:

- Addressing traffic safety and congestion issues
- Transformation of long-distance public transit
- Transformation of long-haul shipping
- Broadband connectivity to autonomous vehicles (in motion)
- Capability to identify and reserve charging stations (in motion)
- Grid monitoring at the transportation V2I small cell location
- Middle mile broadband connectivity for rural communities

As in the previous context, there are construction synergies (dig once for resiliency) during construction or maintenance of transportation infrastructure.

5. Conclusion

As noted in the Introduction, we are at a unique point in time to maximize the impact of investment into smart infrastructure across three major industries. This impact can only be fully realized through cross-industry collaboration. While we do not minimize the challenges of this kind of cross-industry collaboration, we do know that collaboration cannot occur unless the lines of communication are opened, both literally and figuratively.

CableLabs has a long history of creating standards and specifications to leverage communications infrastructure across broad markets. As the electrification of the transportation markets evolves, CableLabs is an ideal point of interface to ensure the communications networks are prepared to deliver the underlying protocols, transport, security, reliability, and resilience of this critical infrastructure.

Abbreviations

IIJA	Infrastructure Investment and Jobs Act
HFC	Hybrid Fiber-Coax
DERS	Distributed Energy Resources
V2I	Vehicle to Infrastructure
ADR	Automated Demand Response

The Evolution of the Edge – Why Edge Compute and Networking Should Tightly Integrate into a Cable Edge Cloud

A Technical Paper prepared for SCTE by

Idris Jafarov
Delivery Team Leader
DriveNets
New Jersey
302-559-5952
ijafarov@drivenets.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. What is the Edge?	3
2.1. Evolution of the Edge	4
2.2. Key Players	5
2.3. Use Cases.....	6
2.3.1. Remote Work	6
2.3.2. Smart City Applications and Autonomous Vehicles.....	6
2.3.3. OTT Content Delivery and Gaming	6
2.3.4. Industry 4.0 Use Cases.....	6
2.3.5. Content Delivery Networks	8
3. Infrastructure Ecosystem for the Edge.....	9
3.1. Data Centers	9
3.2. Cable Headends and Wireless Towers.....	10
3.3. Networking and Edge InterConnect	10
3.4. Alliances and Partnerships.....	10
4. Network Cloud.....	12
4.1. Evolution of Network Functions.....	12
4.2. Edge Cloud Enablers	12
4.2.1. Edge Hardware	13
4.3. Multi-Services Network Cloud Architecture for the Edge.....	15
5. Conclusion.....	18
Abbreviations	19
Bibliography & References.....	20

List of Figures

Title	Page Number
Figure 1 - AI of Things.....	5
Figure 2 - Precision Agriculture.....	7
Figure 3 - Security and Tracking via Drone	7
Figure 4 - Digital Twin Concept as Applied to Industrial Automation.....	8
Figure 5 – Content Delivery Network	9
Figure 6 - Components of the Edge Cloud	11
Figure 7 - Edge Continuum	11
Figure 8 - Evolving from Bare Metal to VNFs to CNFs	12
Figure 9 - DDC/DDBR Architecture	13
Figure 10 - Network Cloud	14
Figure 11 - Packet Forwarder	14
Figure 12 - Fabric Forwarder	14
Figure 13 - Edge Locations for Network Cloud	15
Figure 14 - Software Architecture of the Network Cloud	16
Figure 15 - Abstraction Layer of the Network Cloud.....	17
Figure 16 - An example of multiple applications running over the Network Cloud	18

1. Introduction

For the past decade, and within the halls of this very event (Cable-Tec Expo), technologists, researchers and experts across the compute and communications industries have been discussing and analyzing the “Edge”. Even though this term is not complex in and of itself, it, nevertheless, is fairly fluid. As much as the edge is about compute, it’s about communications, and even more about fusion of compute and communications for distribution of intelligence. Hence, the edge must be solved collectively by the communications providers (traditionally the cablecos and telcos) and the compute providers (the cloud), with a laser focus on consumption, treating the underlying tech merely as tools that can be swapped in and out.

Edge adoption will ultimately be driven by business value. Many organizations want their familiar cloud services brought to the edge where the data they want to process is created, on their choice of infrastructure — and with the flexibility of cloud consumption models. They want the cloud experience for the ever evolving workloads that are best suited to be at the edge. This “cloud edge” infrastructure might be in on-premises data centers, more of it will be in new edge data centers, embedded in edge devices, or even built right into the telecom infrastructure. Regardless of where it is, however, users will be able to consume services on cloud edge infrastructure the same way they consume services on traditional cloud infrastructure. For that reason, tremendous infrastructure investments are needed to support the growing device and infrastructure edge demand. Based on estimates, between 2019 and 2028, cumulative capital expenditures of up to \$800 billion USD will be spent on new and replacement IT server equipment and edge computing facilities [1]. These expenditures will be relatively evenly split between equipment for the device and infrastructure edges.

Many revenue-generating applications such as gaming, healthcare, IoT, AR/VR require low-latency networks with compute resources close to end users, leading to improved user experience and service quality. This brings multiple challenges regarding where to locate edge data centers, how to scale capacity and services within space- and power-limited environments, and how to provision and operate an increasing number of data centers and resources at scale. The answer is a cloud-like, virtualized, shared infrastructure combining compute and networking resources and supporting multiple container-based applications.

In this paper, we will analyze the evolution of the edge, emerging use cases that will accelerate and drive innovation, and key players in the ecosystem. We will touch upon the infrastructure ecosystem for the edge cloud and deep dive into the Network Cloud that will enable operators to utilize network functions on a shared pool of resources.

2. What is the Edge?

There is a biased view that edge and cloud are competing solutions. Actually, they’re part of the same continuum of putting colocation, compute, networking and storage in the most effective place, done in the most efficient way. Edge computing is comprised of combinations of systems that span a wide range of locations and conditions and support a diverse set of use cases. Certain use case might demand high-powered Graphics Processing Units (GPU) for AI (Artificial Intelligence), another one might demand low power consumption to lengthen battery life.

There's not just one answer to the question of the definition of "edge cloud". Each definition points out a unique and important concept in the world's computing infrastructure. For example, edge data centers are small data centers that are located close to the edge of a network. They provide the same hardware found

in traditional data centers, but are contained in a smaller footprint, closer to end users and devices. Distributed edge of 5G where a decentralized cell network made of edge data centers can help provide low latency for use cases with high device density, is another form of edge cloud. Cable operators, also, aim to leverage the edge and in this paper we will deep dive into multiple approaches in which this can be accomplished.

2.1. Evolution of the Edge

Historically, all the data and applications that enterprises needed was stored and processed locally in their on-premises data centers. Due to these resources being available locally, the latency and bandwidth was typically not a concern. These issues became apparent when users connected to the corporate network via a Remote Access VPN. Then, as the cloud computing paradigm took off, some of these applications moved to the cloud compute environment while the rest remained on-premises. This hybrid cloud model certainly has its benefits, but as digital transformation is driving enterprise applications and processes to the cloud, and particularly to public clouds the value of having data centers on site diminishes. Therefore, cloud infrastructure needs to become more distributed and ubiquitous to meet future enterprise needs as demand for cloud-based compute and storage grows. The first generation of distributed edge clouds is being deployed in metro data centers to provide enterprises with a greater choice of compute locations, both within and between countries. Such clouds are serving the needs of enterprises that want to process certain types of data at the edge for regulatory compliance reasons and/or that do not want to incur the cost of backhauling large amounts of data to centralized clouds. The proximity of the cloud infrastructure to the residential customers is, also, of high importance as the "work from home" shift is here to stay and there is a plethora of emerging uses cases like online gaming/streaming, AR/VR that are picking up the pace.

However, beyond this first wave of 'edge clouds', enterprises will require a more distributed, edge-native compute fabric that is available in every conceivable location, not just in a few hundred distributed cloud data centers worldwide, in order to support emerging software applications and software architecture. A new generation of data- and event-driven edge-native services will depend on the existence of a seamless fabric of edge-native clouds that can process their data very close to where it is needed. These edge-native services include but not limited to AI/ML, computer vision, autonomous mobility and applications that support collaboration across multiple contexts and devices, such as multi-player, AR gaming. Such edge-native services are being developed using a very different architectural approach to that used for today's web applications that run in the public cloud, and they have in common the need to be deployed to specific edge cloud locations to execute with the right level of security, latency and compliance.

The new class of edge-native services will serve an increasingly hyperconnected web of devices that is producing and consuming more and more data. AI/ML applications will turn this wealth of data into intelligence far more quickly than humans could and will replace manual steps that introduce friction into processes and experiences. At the end of this decade, industry visionaries expect an 'AI of Things' [2] (Figure-1) to have emerged, whereby everything in our lives will be able to exchange data with everything else, thereby enabling 'magical' and unprecedented levels of automation. The net effect of more data, and the application of AI/ML to it, will be the better optimization of business processes and user experience.



Figure 1 - AI of Things

Edge clouds enable enterprises to run semi-autonomous, latency-sensitive operations locally, to integrate applications into site-based processes more easily and to use server resources more efficiently at individual branches or sites. Cablecos and Telcos have driven early infrastructure edge demand as they virtualize and cloudify their networks. Initially core and transport networks are being transformed with disaggregated distributed networking solutions to enable massive scale and reduction of cost per bit. On top of that, ISPs are uniquely positioned with geographically distributed network infrastructure, which is well suited for infrastructure edge cloud implementations.

2.2. Key Players

There are three primary infrastructure service providers that play a significant role in the buildout of the edge cloud: data center operators, last mile operators, and cloud and edge computing service providers.

Virtualized networking services (NFV, VNF) have also had an early start in edge computing. However, they did not deliver on their promise due to multiple concerns which we will touch upon later. A vast majority of the computing service providers have allocated significant capital into the development of advanced software management platforms. This enables edge computing services to be delivered on-demand and remotely just like with cloud computing. Computing service providers are the most common tenants in early generation edge deployments, which is not an accidental event: they are the ones working directly with end user customers and ultimately hosting and managing the workloads and applications that are being placed at the edge.

Also, the Content Delivery Networks (CDNs) that already have substantial data center footprint predate the modern edge cloud. When CDNs were conceived in the 1990's, it was to enable video on demand and save peering and transport costs and thereby host storage closer to users. This was one of the first applications of edge computing. Now, there is a significant demand for plethora of other applications at the edge.

The CSPs/ISPs can themselves become providers of cloud infrastructure because they will be placing edge-native clouds across geographic locations in order to effect end-to-end, cloud-based transformations of their access and transport networks. Operators can then expose such edge-native clouds to support third-party application pipelines in an infrastructure-as-a-service model (IaaS). Operators are actively partnering with the vendors of disaggregated cloud-native networking solutions to provide a more

extensive edge Network-as-a-Service (NaaS) solution leveraging open hardware that can support use cases that need a high degree of mobility across geographies and can enable them to spin up network functions within a pool of shared resources [3]. We will dive deep into this partnership in a later section.

2.3. Use Cases

There is plenty of demand today for enhanced connectivity and edge computing in multitude of geographical locations. Here are some of the use cases that edge computing will enable and optimize. It is, however, paramount to understand that use cases are different than business cases, as edge cloud ROIs are still a hotly contested item that we are in many cases looking for.

2.3.1. Remote Work

The Work-from-Home (WFH) trend is here to stay. The latest global health concerns have only underlined the importance of minimizing the need for onsite expertise in data centers of any type, from core to edge, and this has hastened the development of tools for remote monitoring, provisioning, repair and management, which could greatly reduce the cost of edge computing. Our solution allows to spin network functions remotely without having anyone on site thereby easing the requirements for on-premises presence.

Also, enterprises have to assure the security of the teleworkers' workstations and secure connectivity to the corporate digital assets. Due to the fact that majority of the resources employees need to access and work on are hosted in the cloud, and depending on the location of the workers, the latency might cause inadequate application response. The necessary workloads have to move closer to the end users for optimal performance.

2.3.2. Smart City Applications and Autonomous Vehicles

The widespread adoption of autonomous vehicles is inevitable and will require a proliferation of edge computing services along roadways to ensure the cars can navigate properly. This will become particularly important in rural areas, where infrastructure will need to be purposefully built out to support autonomous cars in areas that may have no connectivity at all. Similarly, small rural towns and villages will need enhanced edge computing capabilities to implement smart city applications such as turning street lights on or off or monitoring traffic cameras [4].

2.3.3. OTT Content Delivery and Gaming

One of the consumer-driven use cases for edge computing that became a crucial benefit for many during the pandemic is content delivery and online gaming/streaming. Edge computing can help OTT service providers deliver better user experiences for customers and can improve online gaming experiences by reducing latency.

2.3.4. Industry 4.0 Use Cases

As a natural extension of cloud computing, the edge cloud construct is increasingly viewed as a key enabler for the "Fourth Industrial Revolution" in which the widespread deployment of the Internet of Things (IoT), the global sharing economy and the increase of zero marginal cost manufacturing deliver unprecedented communication-driven opportunities with massive economies of scale.

Despite its slow adoption rate, precision agriculture is a great example of this. Farmers using precision agriculture tools and applications will generate thousands if not millions of data points from things like

measuring soil moisture and nutrient content to plant growth and stress, to temperatures and precipitation—to even operating autonomous combines [4] (Figure-2).

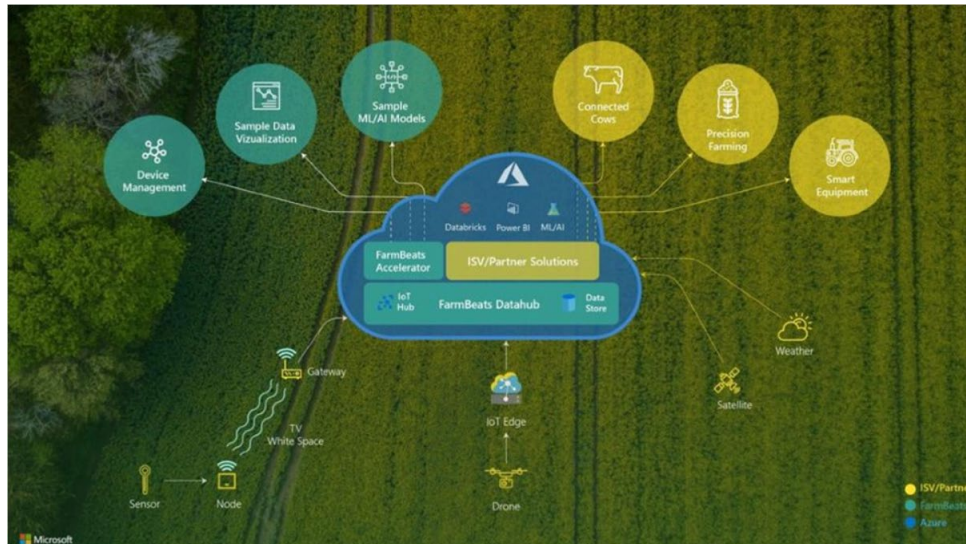


Figure 2 - Precision Agriculture

Other promising use cases are security via drone or camera across large agriculture sites, tracking herds and detecting imminent births or injuries (animal not moving for a long period), optimized feeding for dairy herds and worker safety. All of that data needs to be managed locally at the farm, for example, but it doesn't necessarily need to travel off the farm. Such applications will require connectivity, of course, but because the needs are localized, that connectivity can be served by an array of local network technologies (Figure-3).




Figure 3 - Security and Tracking via Drone

One of the most fascinating use cases is the notion of a digital twin: a near-real-time digital image of a physical object or process that helps optimize business performance. The digital twin can allow companies to have a complete digital footprint of their products from design and development through the end of the product life cycle.

A man wearing a VR headset is shown interacting with a futuristic, glowing blue mechanical assembly. The assembly appears to be a complex, multi-part device, possibly a motor or a sensor, with various components and wires visible. The background is a high-tech laboratory or workshop, with various equipment and tools visible. The overall scene is illuminated with blue and purple light, creating a futuristic and immersive atmosphere.

2.3.5. Content Delivery Networks

Typically, major CDN provider manages and operates a large number of content servers geographically distributed across regions, countries or even continents. Besides the origin server, copies of the content are cached on different servers. Cached content offers clients a faster loading experience since their requests can be answered by geographically closer servers.

© 2022, SCTE® CableLabs® and NCTA. All rights reserved.




Figure 5 – Content Delivery Network

3. Infrastructure Ecosystem for the Edge

At the edge, the physical infrastructure is disparate, vendor-agnostic and comes in various shapes and sizes. The layers are built on top of each other and form a harmonious relationship. From the underlying wholesale data center to the cloud infrastructure that it houses; to the multiple sources of connectivity that connects end users and moves data from the core to the edge; to the real estate that is able to support all these complex requirements. At the edge, each critical infrastructure component is crucial in and of itself, but they work together as part of a single integrated ecosystem.

3.1. Data Centers

Roughly everything important to data center design comes down to three elements: carbon, real estate footprint and scale. The overwhelming tendency in the industry has been to pack increasing amounts of power into ever smaller spaces without overwhelming the cooling systems, which themselves have become progressively more efficient. The end result of all these optimizations has been to increase the density of the compute, networking and storage equipment a data center is capable of handling.

The “rack” in “rack density” refers to a standard 7-foot equipment rack, the kind found in virtually every data center around the world, while “density” refers to how much equipment can be packed into that rack. Rack density, then, becomes a way of expressing the power density of a data center and this concept is being extended to new form factors.

For instance, a typical enterprise data center might average 6-12 kW per rack, whereas hyperscale data centers can handle densities upwards of 50 kW per rack. A data center with 50 kW rack density can consume four times the power of a typical enterprise data center on a rack-by-rack basis.

At the edge, rack density becomes important because space is scarce and expensive. The higher the density, the more we can do at the edge. Several approaches exist that can eliminate the limitations posed by the constraints at the edge by hosting multiple cloud-native network functions on a shared infrastructure which we will touch upon later.

3.2. Cable Headends and Wireless Towers

Real estate is gaining more and more importance in locations adjacent to cable headends, wireless towers, and fiber aggregation points, thereby attracting landowners who will likely become players in the ecosystem. Historically, landlords were involved in building wholesale data centers on their land with cloud infrastructure providers. Now, landowners are providing strategic real estate for edge cloud, however, at a multitude of locations due to the vast amount of viable use cases. Landowners provide the real estate, micro data center operators provide the colocation facilities, and then cloud providers, networking software vendors provide the necessary equipment and applications that operate in those facilities.

Edge cloud locations will also pave a way for new forms of interconnection and peering. Traditional hosting data centers turned out to be meeting points for networks, and akin to that the micro data centers at cable headends and wireless towers that will drive edge computing are generally located at the intersection of connectivity paths. These network regions will become very attractive for local interconnection and edge exchange, enabling more optimal routes for data.

3.3. Networking and Edge InterConnect

The rise of edge cloud is going to require interconnection to move from its traditional centralized Internet Exchange (IX) model, typically in primary locations within major metros like Frankfurt, London, Hong Kong and etc., to an Edge Exchange (EX) model [1]. End users and devices out at the edge are far away from primary IX points and the distance it takes for traffic to travel to these locations degrades performance and also increases transport costs significantly. To solve this problem, interconnection of networks will need to happen in edge data centers near the last mile network in much closer proximity to the end user. CableCos are using multiple real estate locations (and are perfectly positioned for/to take advantage of) such as headends for aggregation and backbone, and hubs that are hosting vCMTS' and other access equipment which can be used to deploy a next-generation cloud-native solutions to enable a plethora of edge use cases.

We will see edge exchanges emerge to allow peering and data sharing at the edge without necessarily involving the core. Edge cloud is not managed in isolated and independent fashion. It will allow more traffic to remain local, but it will also become an interdependent extension of the traditional IX.

Functions of the applications at the edge will be divided between the edge (real-time processing) and the core (primary functions, analysis, archiving) and only the required data will move back and forth through interconnection services. Having an exchange point at the core that is also tied to the edge will enhance performance and drive cost efficiencies.

3.4. Alliances and Partnerships

The diagram below shows the range of network access points and possible topologies, with a range of expected latency and network access:

COMPONENTS OF THE EDGE CLOUD

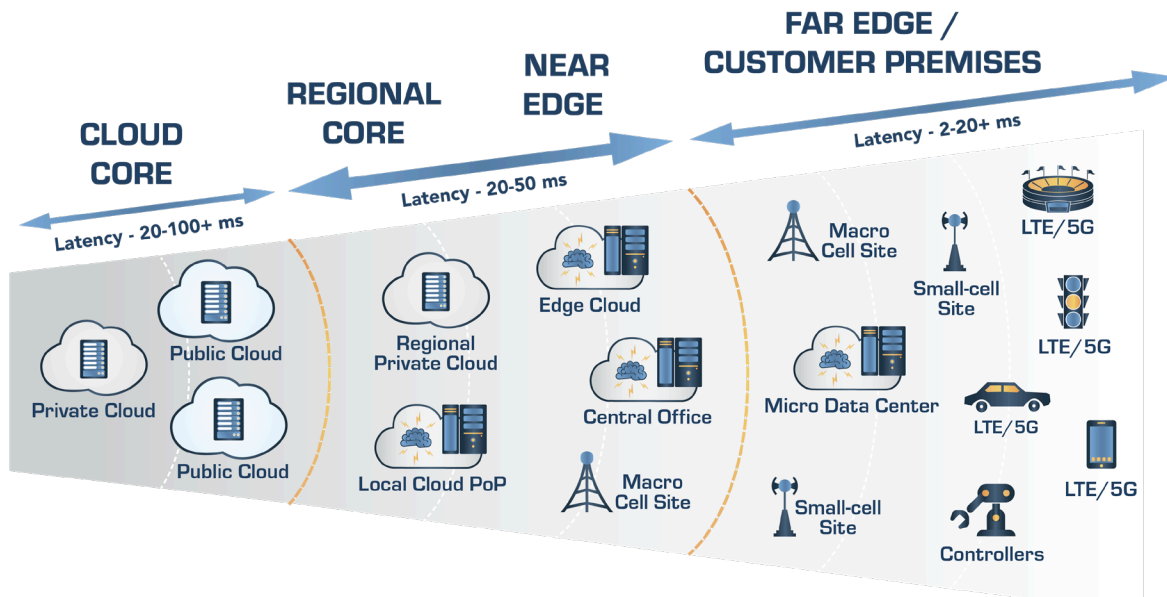


Figure 6 - Components of the Edge Cloud

The reality is that latency choices will depend on the applications. In some cases, it makes sense that compute functions are going to happen at the edge, whether that's a local node or even a device. In one example, retail analytics will use edge compute nodes at the location to deliver specific data or applications to the customers. Yet another example is a connect vehicle, which will process a lot of data on board while sharing telemetry and other information with the cloud. It's all about the data and where it's needed. Much of the data that's gathered at the edge has a short lifespan and really isn't economical to bring back into the core cloud.

At the end of the day, it is partnerships that will bring all these disparate pieces together. Wireless tower operators have considered owning data centers, but the most efficient way for them to get to market is to partner. Micro data center vendors are partnering with the service providers that host and manage cloud infrastructure and bringing them on as tenants. Meanwhile, micro data center operators are working with wireless tower and cable companies to access the real estate at headend/hub and tower sites.

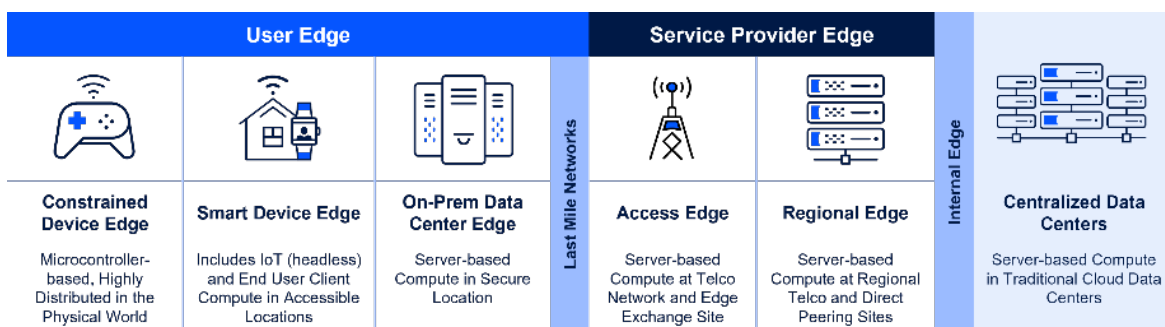


Figure 7 - Edge Continuum

An exceptional example of this kind of partnership is Open Grid Alliance (OGA), where industry leaders and innovative disruptive startups join forces to rearchitect the Internet. This alliance will work to evolve the Internet to be a global, shared platform that distributes compute, data, and intelligence to when and where it's needed, on demand [5].

4. Network Cloud

4.1. Evolution of Network Functions

A key consideration for the industry is how to shift Network Functions (NFs) from physical instances to containerized microservices and disaggregated business models. This poses new requirements for the orchestration and control of highly disaggregated and distributed capabilities.

Network Functions were originally deployed on dedicated physical appliances. Then with the widespread adoption of compute and storage virtualization Network Function Virtualization (NFV) initiative started to pick up steam and Virtual Network Functions (VNF) ran on x86 servers (on CPU) which was sub-optimal. It could scale to multiple servers, but the performance was not up to par. As microservices architecture became increasingly popular, Cloud native Network Functions (CNF) were able to run on the public cloud infrastructure and this enabled further growth. However, these NFs still ran on COTS x86 servers.

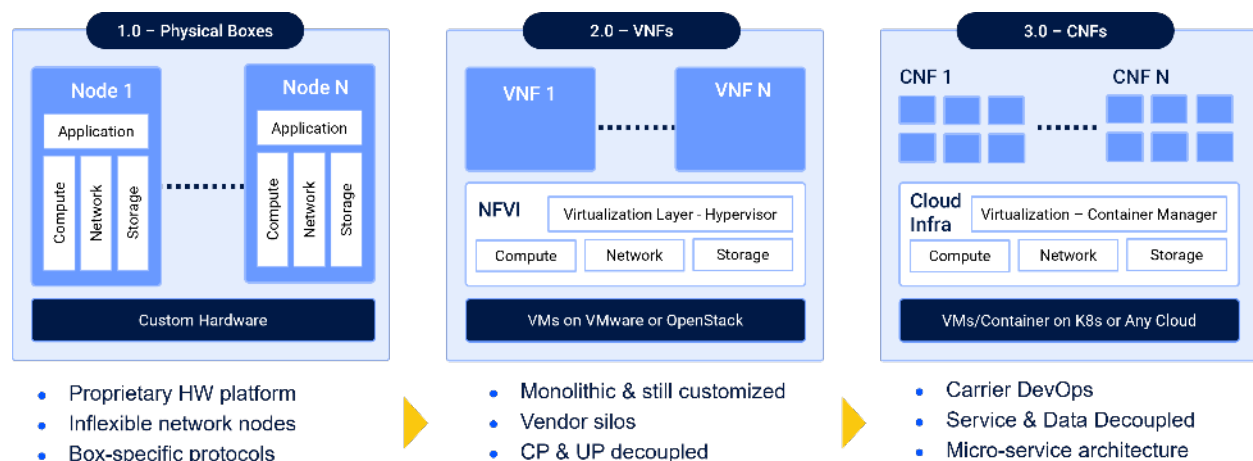


Figure 8 - Evolving from Bare Metal to VNFs to CNFs

Today it is commonly proposed to offload the NF to the Network Processing Unit (NPU) on a dedicated Input/Output (I/O) device. This can also be done through a “smart NIC” provider to offload the NF to the Data Processing Unit (DPU) and accelerate the NF as a result.

4.2. Edge Cloud Enablers

Hyperscalers have spearheaded the development and implementation of the open hardware and software data center solutions. Organizations like Open Compute Project (OCP) [6] and Telecom Infra Project (TIP) [7] have introduced Disaggregated Distributed Chassis (DDC) [8] and Disaggregated Distributed Backbone Router (DDBR) [9] architecture. It is based on whitebox architecture where each component of the standalone monolithic chassis router is distributed into individual components:

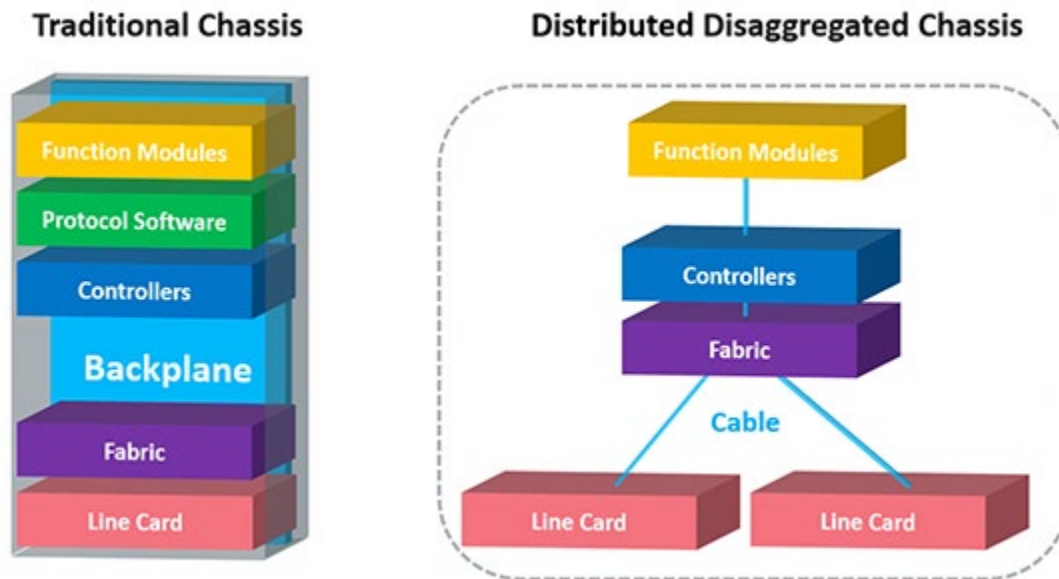


Figure 9 - DDC/DDBR Architecture

This allows the operators to scale out this re-architected chassis by simply adding whiteboxes that perform fabric and packet forwarding functions.

Evidently, this is an architectural improvement related to cloud-native networking, and potentially another indication of a future in which carrier networking is highly modular and disaggregated. The true value of such a disruptive model is not necessarily limited to up-front cost-savings. Rather, the value lies in the optionality that comes from building an infinitely scalable and malleable network cloud platform that is pay-as-you-go.

4.2.1. Edge Hardware

Hardware deployed at the edge has historically been purpose-built for specific workloads, frequently CDNs (Content Delivery Networks) or IoT. As edge computing grows in popularity and new use cases emerge, general purpose infrastructure is also being deployed to run cloud-like workloads.

Edge systems must take the path of the hyperscale cloud buildout – low-cost commodity hardware combined with high-powered software automation and distributed orchestration.

The above described DDC/DDBR architecture is perfectly suited to be deployed at the edge. A variety of Original Design Manufacturers (ODM) [10] have introduced open hardware based on merchant silicon certified by OCP and TIP which is already carrying production traffic in the backbone and aggregation networks of major cable and telecom operators.

Due to the disaggregated nature of the architecture where control and forwarding planes run on different physical platforms, there is only a need for a data plane to be installed at the edge whereas control functions run in the core cloud.




Figure 10 - Network Cloud

Packet and fabric forwarders (PF and FF) which comprise data plane, have very high port density (40x100GE, 36x400GE):



Figure 11 - Packet Forwarder



Figure 12 - Fabric Forwarder

By strategically placing these disaggregated clusters at the edge facilities (cable headends/hubs, wireless towers) operators can utilize this platform to run multiple CNFs in a cloud-like manner thereby eliminating the need to install separate physical appliances for each network function and utilize any port of the system for any service. The latter point is of a particular importance for the edge environment, because the cost of deployments is higher and rack density, power, cooling, scale, flexibility and even physical access to the site is very limited.




Figure 13 - Edge Locations for Network Cloud

4.3. Multi-Services Network Cloud Architecture for the Edge

Initially, it was believed that low-latency needs would drive edge compute, however, we believe the real advantage of network cloud edge is the flexibility, scale and automation provided in deploying new services wherever they are needed, using standardized and affordable software with cloud-based networking and automation software.

Many network edge projects have failed because of efforts to deploy proprietary platforms or complicated architectures and high costs. Following the cloud will be key: leveraging existing cloud-native technologies such as Kubernetes, APIs, and public cloud services and extending them out to the edge.

From a networking perspective, a distributed disaggregated carrier-class networking operating system predicated on cloud-native software principles and standard white-box functionality is required at the edge. This NOS should be predicated on the following cloud-native software design guidelines:

- **Natively Distributed NOS:** marks a departure from the monolithic network hardware designed around a proprietary model of equipment and separates logical routing functionality from the physical infrastructure.
- **Extensive Use of Containers:** designed from the onset to leverage Docker containers to ease development, deployment, and upgrade.

- **Optimized Resource Utilization:** anchored on small, highly cohesive, and loosely-coupled microservices.

There are solutions in the market that address majority of the inefficiencies at the edge and offer operators a way to significantly increase their resource utilization while gaining service & architecture flexibility, optimal scaling and software-paced innovation and time to market. Network Cloud for the edge architecture addresses many challenges by combining networking and compute resources over a shared, cloud-like infrastructure. It allows operators to put greater functionality at the network edge, even with space and power limitations.

The foundation for this significant value is the way the Network Cloud is built – in a disaggregated, cloud native architecture. This means the hardware resources of a cluster of multiple white boxes are abstracted by the NOS to a level in which it is consumed as a virtualized resource pool, based on the DDC/DDBR architecture mentioned above. Each networking function, which runs a containerized Service Instance (SI), can be allocated with its required hardware resource (Physical interfaces, NPU, CPU, TCAM, QoS etc.) out of the underlying shared hardware infrastructure.

The following diagram illustrates the software architecture that allows this:

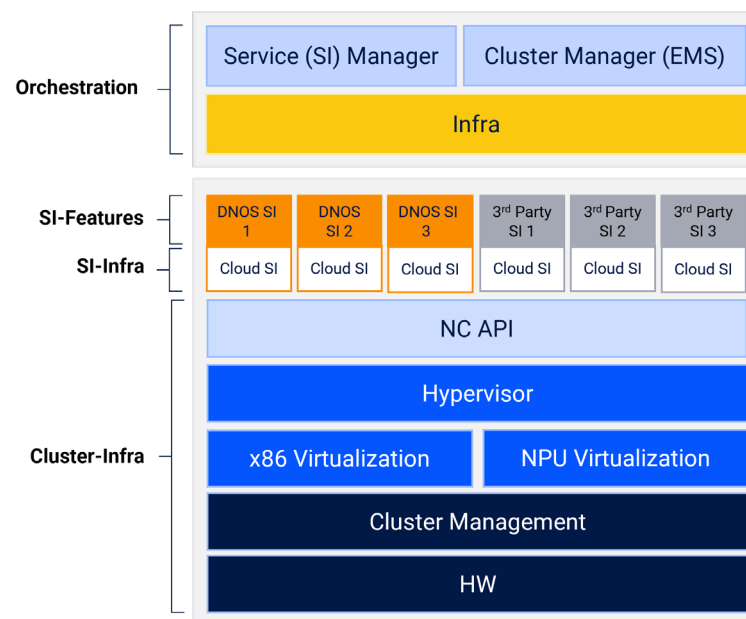


Figure 14 - Software Architecture of the Network Cloud

While different hardware resources are abstracted and different cloud SIs run over it, the system represents each network function as a standalone node, keeping the network manageable. This means that a single cluster can integrate multiple networking functions that are physically collocated, yet logically separated.

A multiservice architecture creates a separation of the data plane from the service plane (or control plane). Multiple services can coexist over the same virtualized physical infrastructure and as long as resources can be made available for a service to perform, the network cloud can launch this service into production through a centralized orchestration system.

This shared use of resources enables operators to place such scalable cluster at the edge locations while lowering footprint, power consumption, space requirements while reducing the need for engineers get on site for installation and troubleshooting activity. Additional network services can be enabled on this infrastructure on-demand in an automated fashion via an orchestrator that has a full view of the topology and hardware resources of all the components of the cluster.

In order to implement a multiservice architecture a couple of virtual entities were introduced:

- **Service Instance (SI):** a network function (e.g., PE router, VPN Gateway), which runs independently of any other networking function co-located on the same cluster. While the service instance is logical, it is assigned with dedicated resources out of the shared pool of hardware resources in the network cloud cluster.
- **Inter Service Link (ISL):** A logical connection between two SIs, which is represented as a (physical) link connecting the two networking nodes. This link allows control plane interconnection between the instances as well as data path features, including QoS and access lists. The ISL is totally logical, hence no cabling work is needed.

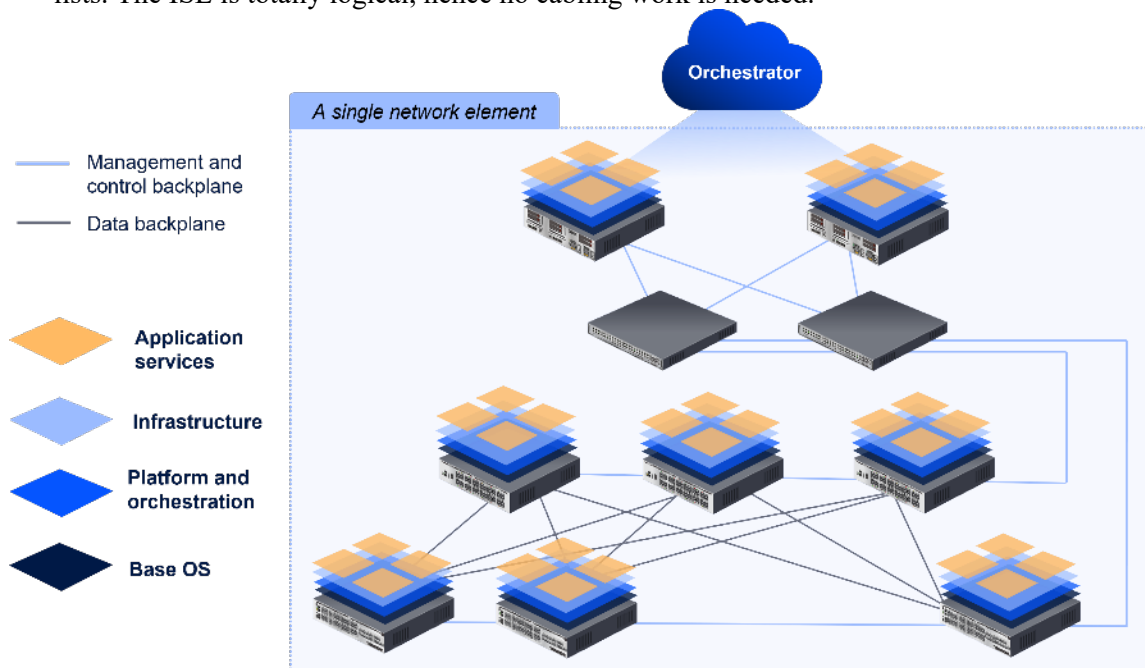


Figure 15 - Abstraction Layer of the Network Cloud

The multiservice solution allows for replacing the model of multiple network appliances at the edge, serving different functions and networks (e.g., Network Security, Mobile Backhaul and Broadband), with a single unified infrastructure that can support all the networking services on a software-based network and cloud-native technologies (i.e., microservices, containers).

Let's provide an example based on a TCAM allocation scheme. Three different services are allocated to the same physical ports (Yellow could be a VPN service, Orange – a Mobile Backhaul and blue – a DDoS Mitigation Application). Their TCAM requirements are different, so there might be a case in which a packet arrives to a port, under the yellow service and there are no TCAM resources available. In such an occasion, it will use a TCAM resource from a different whitebox, which has required resources at its disposal. This will be done over the cluster's fabric, as illustrated below.

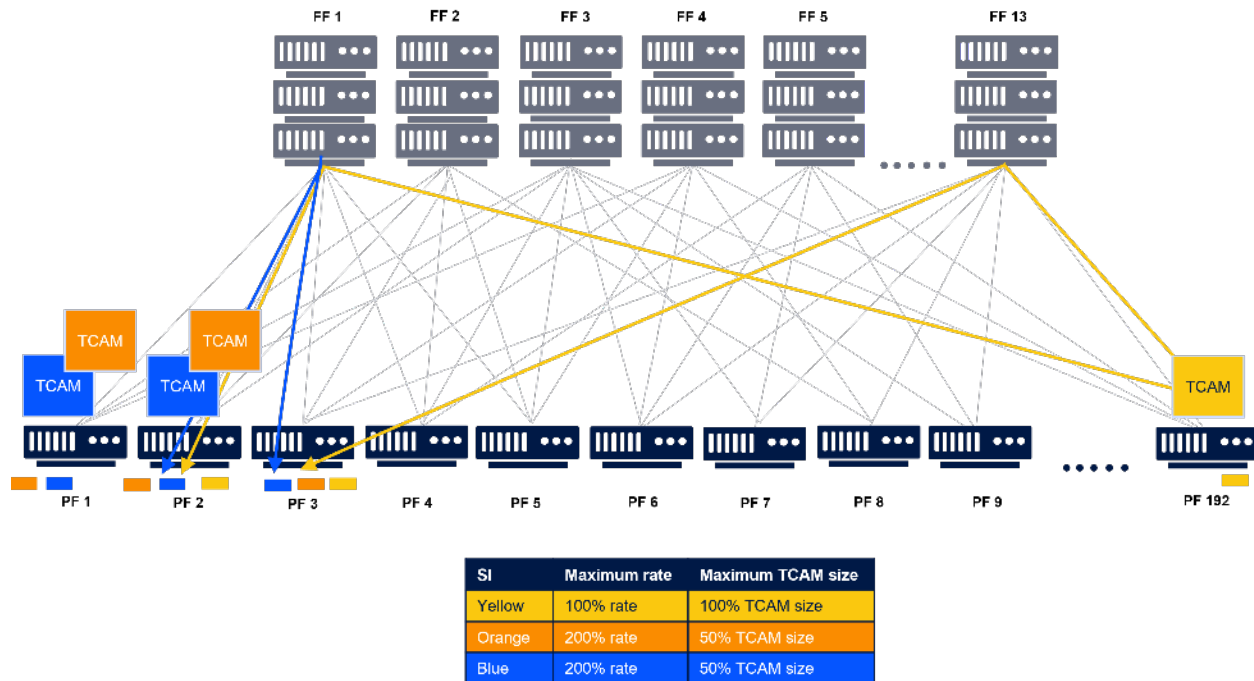


Figure 16 - An example of multiple applications running over the Network Cloud

The same edge cloud platform will also run the network's increasingly autonomous operational and management functions. This is providing operators with the opportunity to create a distributed cloud fabric that can unify multiple network domains and which they can then make available to third-party application pipelines [11].

We believe that operators who have already deployed this architecture in other places of their production environment will have an upper hand as they can utilize the same hardware and leverage new innovative features of the software like the multiservice functionality demonstrated above.

Rest assured – the new network edge is taking shape. But the most successful networking and compute edge deployments will have the characteristics of solving the challenges by simultaneously delivering new applications capabilities and radically lowering the cost of deploying these technologies.

5. Conclusion

The global expansion of internet infrastructure is happening at a rapid pace, and the pandemic was at the center of a junction of events that are accelerating the trends already in place. The next decade will see an explosion of activity as a result. There will be more opportunities for creation of a new infrastructure locations at the edge and in strategic and difficult-to-access places. Moving to the edge is fundamentally about getting infrastructure closer to end users, whether that be in new and emerging markets or wherever critical masses of end users cluster. The underlying infrastructure will be provided by a hyperscale platforms, internet service providers and innovative startup companies.

Operators are in a strong position because their networks are, by nature, distributed and ubiquitous across their geographic footprints. Network cloudification necessarily entails building a distributed, edge-native

cloud fabric across that footprint too. Cablecos and telcos are engaged in a series of network transformations that will eventually introduce cloud into different network domains (backbone, aggregation, mobile, fixed access and transport). Leading operators have a unified vision for such a network cloud, and anticipate that over the next few years, they will put in place a common, distributed and hybrid cloud-based platform to support all software-only network functions and their management and operational systems. The network cloud must extend to the network edge in order to support access network functions, but it will also need to be instantiated in many other locations, and operators will require a holistic way of orchestrating network function workloads across thousands of cloud locations to provide a ‘network’.

Abbreviations

AI/ML	Artificial Intelligence/Machine Learning
API	Application Programming Interface
AR	Augmented Reality
CDN	Content Delivery Network
CNF	Cloud native Network Function
CPU	Central Processing Unit
CSP	Communications Service Provider
DDBR	Disaggregated Distributed Backbone Router
DDC	Distributed Disaggregated Chassis
EX	Edge Exchange
FF	Fabric Forwarder
GPU	Graphics Processing Unit
I/O	Input/Output
IoT	Internet of Things
ISL	Inter Service Link
ISP	Internet Service Provider
IX	Internet Exchange
NaaS	Network as a Service
NF	Network Function
NFV	Network Function Virtualization
NIC	Network Interface Card
NOS	Network Operating System
NPU	Network Processing Unit
OCP	Open Compute Project
ODM	Original Design Manufacturer
OGA	Open Grid Alliance
OTT	Over The Top
PF	Packet Forwarder
QoS	Quality of Service
SI	Service Instance
TCAM	Ternary Content-Addressable Memory
TIP	Telecom Infra Project
vCMTS	Virtual Cable Modem Termination System
VNF	Virtual Network Function
VPN	Virtual Private Network

VR	Virtual Reality
WFH	Work from Home

Bibliography & References

- [1] *The State of the Edge report 2021*; The Linux Foundation
- [2] <https://www.forbes.com/sites/bernardmarr/2019/12/20/what-is-the-artificial-intelligence-of-things-when-ai-meets-iot/?sh=592e3cf8b1fd>; Bernard Marr, Forbes.
- [3] https://about.att.com/story/2020/open_disaggregated_core_router.html
- [4] *The State of the Edge report 2022*; The Linux Foundation
- [5] <https://www.opengridalliance.org/>
- [6] <https://www.opencompute.org/>
- [7] <https://telecominfraproject.com/>
- [8] *Hardware Specifications and Use Case Description for J2-DDC Routing System*, Tuan Duong; Open Compute Project
- [9] *Distributed Disaggregated Backbone Router (DDBR) Technical Requirements Document*, Eva Rossi, Jose Angel Perez, Kenji Kumaki, Ryuji Matsunaga, Yuji Sonoki, Ahmed Hatem, Diego Marí Moretón; Telecom Infra Project
- [10] <https://drivenets.com/news-and-events/press-release/drivenets-partners-with-industry-leaders-broadcom-ufispace-edgecore-and-delta/>
- [11] <https://drivenets.com/products/multiservices/>

The Full Duplex DOCSIS Amplifier – Why, How, and When

A Technical Paper prepared for SCTE by

Richard S. Prodan, Ph.D.
Engineering Fellow
Comcast Cable
183 Inverness Dr. Englewood, CO
+1 (720) 512-3742
rich_prodan@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Origins of FDX DOCSIS.....	3
3. Beyond Node + 0: The FDX Amplifier.....	6
4. FDX Amplifier with Echo Cancellation Cascade Analysis.....	9
5. SNR Calculations for the Node + N FDX Cable System.....	11
6. FDX Moves into the Field.....	16
7. Conclusion.....	16
Abbreviations	18
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - DOCSIS 3.1 OFDM Transmission and Reception	3
Figure 2 - Node + 0 passive cable system.....	4
Figure 4 - FDX node operation with echo path interference	5
Figure 5 - FDX node with echo cancellation signal processing	5
Figure 6 - Upstream Front End Dynamic Range and Residual Downstream Echo/Noise Floor	6
Figure 7 - Node + N amplifier cable system.....	6
Figure 8 - Node + N amplifier cable system upgrade to support FDX.....	7
Figure 9 - FDX Amplifier Echo Interference.....	7
Figure 3 - Node + x cable system with FDX amplifiers.....	8
Figure 10 - FDX amplifier with echo cancellation of the downstream echo.....	9
Figure 11 - FDX Amplifier Transmitted Downstream Echo Interference	9
Figure 12 - Cascaded Noise of Amplifier Echo Cancellation Residuals	10
Figure 13 - Bit-Loading Contour Map – Spec Bit-Loading SNR Thresholds	10
Figure 14 - Bit-Loading Contour Map – PMA Bit-Loading SNR Thresholds.....	11
Figure 15 - Single-Family Unit (SFU) Model for N + 6.....	12
Figure 16 – Return Loss of the Most Common Tap OEM Examples – First Tap from Active	12
Figure 17 - Node Echo Cancellation and SNR with Node + 6 (5 dBmV/6.4 MHz)	13
Figure 18 - Node Echo Cancellation and SNR with Node + 6 (13 dBmV/6.4 MHz)	13
Figure 19 - Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade	14
Figure 20- Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade	14
Figure 21 - Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade	15
Figure 22 - Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade	15
Figure 23 - Net Upstream Throughput by Bandwidth Allocation and Cascade Depth vs Receive Level ...	16

1. Introduction

The advent of Full Duplex DOCSIS[®] (FDX) technology is here as realized in a custom ASIC successfully implemented in an operational FDX R-PHY node reference design. This ASIC is also currently being ported to a node design for trial deployments in Node + 0 networks. A version of this ASIC is also being adapted for a prototype FDX amplifier utilizing the same echo cancellation technology employed in the node.

The original constraint in the DOCSIS 4.0 FDX specification anticipated deployments exclusively in a passive Node + 0 network without amplifiers. The initial evaluation in a demonstration of the implemented prototype FDX R-PHY node incorporating the first FDX RPD ASIC exceeded performance expectations [2]. The resultant performance of the echo cancellation technology suggests use in an FDX amplifier as a “repeater” of bidirectional FDX signals.

This paper considers this implementation for an FDX amplifier. Why this greatly expands the use of FDX technology in conventional Node + X amplifier networks is introduced. How the FDX amplifier is implemented and the resulting performance greatly increasing upstream capacity is analyzed. The performance of cascading FDX amplifiers is shown. Finally, the path to upgrading existing Node + X cable networks when the next generation FDX amplifier ASIC currently under development becomes available is discussed.

2. Origins of FDX DOCSIS

CableLabs launched a project to examine the potential of full duplex simultaneous upstream and downstream transmission and reception within the same cable system spectrum. The DOCSIS 3.1 specification provided the Physical Layer (PHY) transmission and reception based on OFDM/OFDMA as shown in Figure 1.

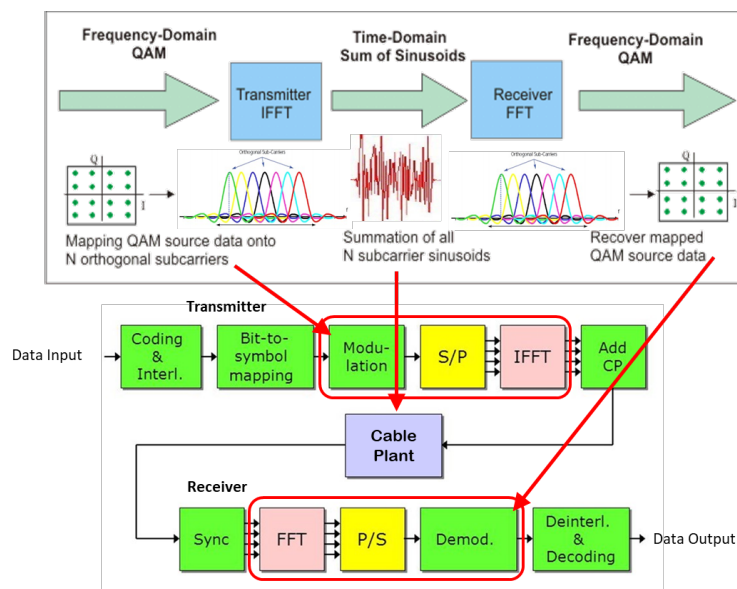


Figure 1 - DOCSIS 3.1 OFDM Transmission and Reception

The difference was the sharing of overlapping spectrum for simultaneous bidirectional transmission and reception. The addition of this spectrum sharing methodology was initially developed as an appendix to the existing 3.1 spec.

Eventually the addition was separated into a new DOCSIS 4.0 spec. The FDX node transmits downstream and receives upstream at the same time in the same bandwidth. FDX transmissions from node to modem and modem to node are overlapped both in time and frequency where new interference cancellation technology in the node allows this simultaneous bidirectional communication.

However, the simultaneous bidirectional communication between the node and the cable modems was limited to a passive coax plant without amplifiers due to the lack of duplex filters to separate upstream from downstream transmission. This is known as a Node + 0 plant where the node is connected to a passive cascade of taps and cable without additional amplifiers as shown in Figure 2.

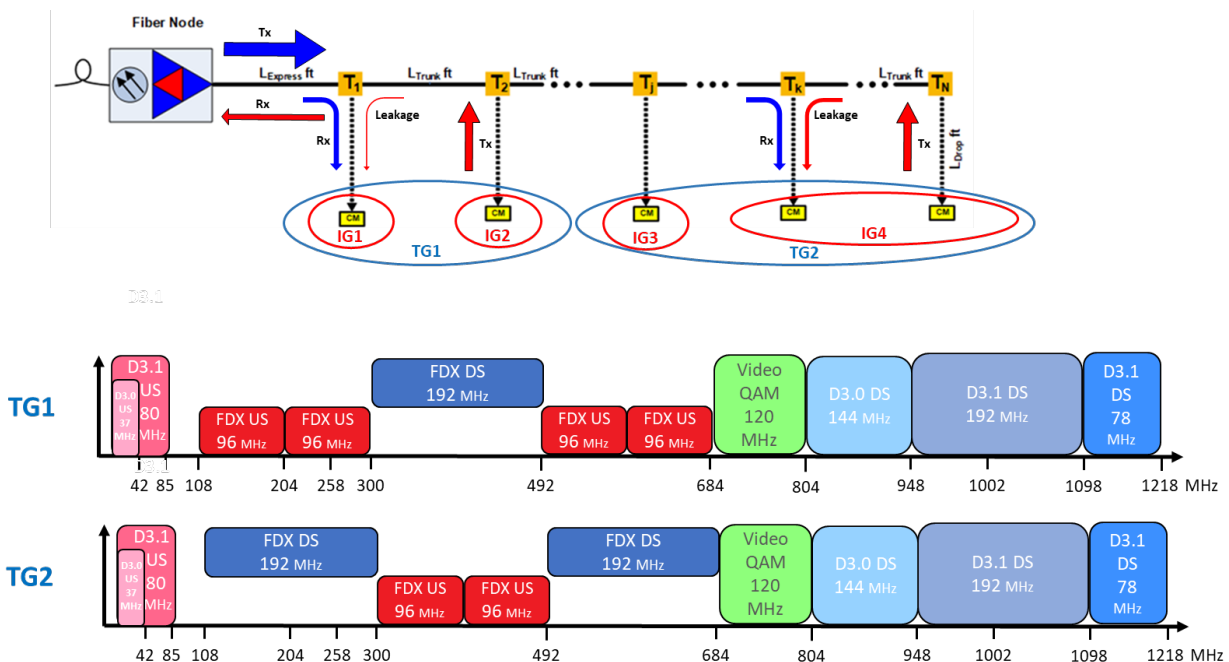


Figure 2 - Node + 0 passive cable system

Simultaneous true full duplex transmission and reception within the same frequency band only occurs at the node. Each cable modem utilizes Frequency Division Duplex (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 each 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 a different tap or groups of taps with sufficient isolation from tap port-to-output leakage between them known as an Interference Group (IG) as depicted in Figure 2. Thus, the downstream receiving cable modem has no reference for the upstream transmitting cable modem signal making cancellation within the same frequency band impossible as discussed in [1].

The high-level node downstream signal is transmitted into the cable system of Figure 2 where that signal propagates down the cable transmission lines and encounters multiple taps with impedance mismatched to the coax transmission line by the tap return loss. A reduced amplitude reflection from each tap results as shown in Figure 3 using the model simulation approach derived in [1].

The node upstream received, downstream transmitted and cable system downstream reflected interference is shown on the right. All these signals combined from each TG occupy the same frequencies in the cable spectrum simultaneously as shown below.

Note in the upper right chart that the echo is well below the node transmission but is significantly higher (around 20 dB) than the low-level upstream received signal from the modems resulting in a negative upstream signal to downstream echo SNR.

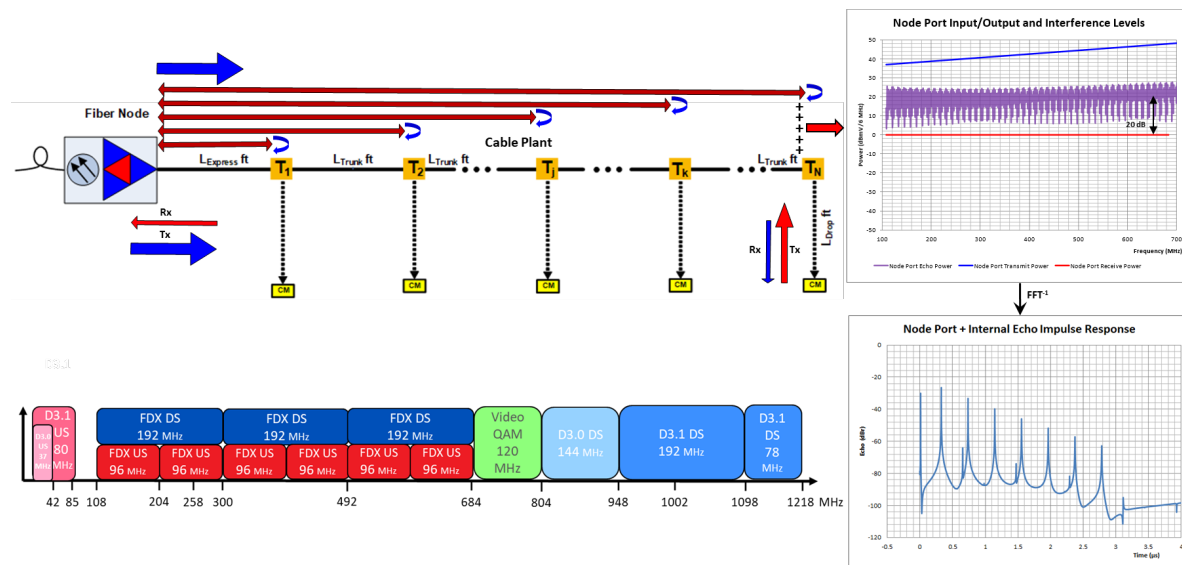


Figure 3 - FDX node operation with echo path interference

The node separates the received modem upstream signals from the transmitted downstream node signal using a directional coupler plus Echo Cancellation (EC) digital signal processing. The function of the FDX node is shown in Figure 4.

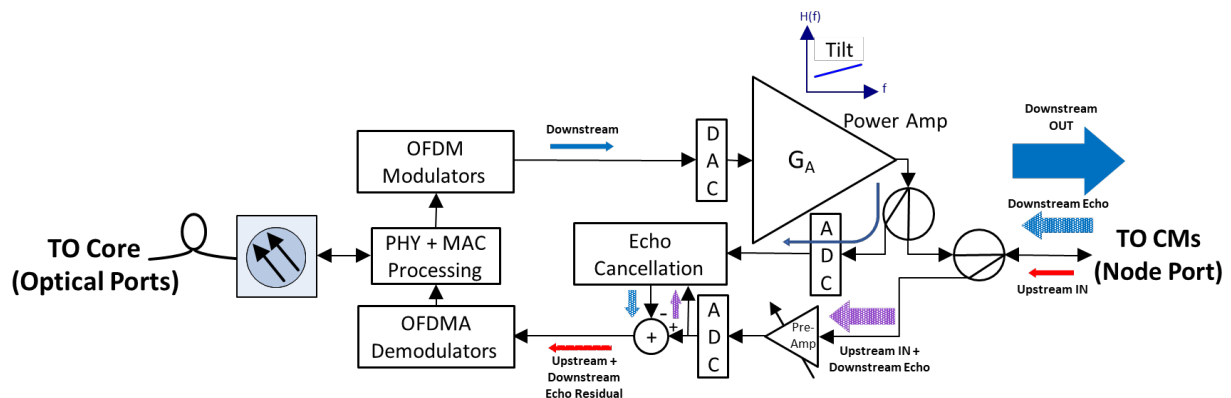


Figure 4 - FDX node with echo cancellation signal processing

A sample of the downstream signal is used as a reference for the EC of the reflected downstream signal. The upstream signal is filtered by an adaptive filter using the downstream reference to remove the echo in the same band. The filtering has a finite depth of EC which is a residual that is now lower than the upstream received signal as shown in Figure 5 resulting in a positive upstream signal to downstream echo residual SNR.

The first prototype FDX Node + 0 system with echo amplitudes producing a negative SNR with respect to the upstream received signal realized an EC depth resulting in a 35 dB SNR in a high downstream launch power Node + 0 cable system design. This demonstrated the power of the EC implemented in the first Broadcom FDX RPD IC as sufficient to consider embedding the same EC technology into an FDX amplifier IC.




Figure 5 - Upstream Front End Dynamic Range and Residual Downstream Echo/Noise Floor

3. Beyond Node + 0: The FDX Amplifier

A traditional cable system with amplifiers is depicted in Figure 6.




Figure 6 - Node + N amplifier cable system

Upgrading the cable system to FDX by replacing the traditional node with an FDX node and replacing the traditional amplifiers with FDX amplifiers is depicted in Figure 7. Note that the taps are unchanged and the amplifier spacing is maintained.




Figure 7 - Node + N amplifier cable system upgrade to support FDX

A similar echo environment found in an FDX node exists for an FDX amplifier. The echo interference paths for an FDX amplifier in a cascade is pictured in Figure 8.

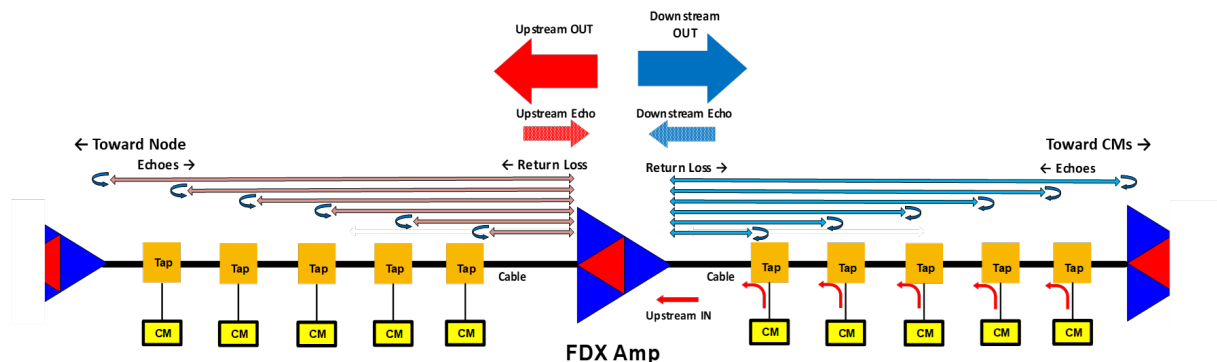


Figure 8 - FDX Amplifier Echo Interference

In the case of a bidirectional amplifier, the downstream transmission level is like that of the node. This results in very similar echo levels in the FDX band. Hence echo cancellation performance required is similar.

The addition of FDX amps requires a different TG grouping of cascaded taps and amps as shown in Figure 9. When a cable modem just upstream of an amp transmits, some of the upstream transmission leaks from the tap port to the output port due to limited tap port-to-output isolation, as depicted in the top amp in the figure. This upstream leakage enters the amp interfering with the intended downstream signal. The SNR is determined by the input downstream-to-upstream leakage ratio. This SNR will propagate downstream of this amp limiting bit-loading (spectral efficiency) thereby lowering downstream capacity. The tap or several taps before the amp and all cascaded taps and amps following the first amp become a single IG. This has been called “IG extension” due to the addition of FDX amps.

To optimize downstream capacity, the entire leg is assigned as a single TG which prevents downstream interference from simultaneous upstream transmission and downstream reception in the same frequency band at the cable modem. Different legs of Figure 9 are combined at the node. Each leg is assigned to a separate TG. The interference from the upstream transmission in one leg into the downstream transmission of the other leg(s) determines the downstream SNR.

Downstream capacity is increased due to the significantly higher isolation between legs of the node splitter/combiner. Traffic engineering of upstream bandwidth grants among the increased number of cable modems in a single TG needs to be considered in long cascades of amplifiers. Also, as explained in the following sections, the degradation of upstream SNR due to the number of amps in a cascade (i.e., Node + x) determines the upstream capacity at the node.

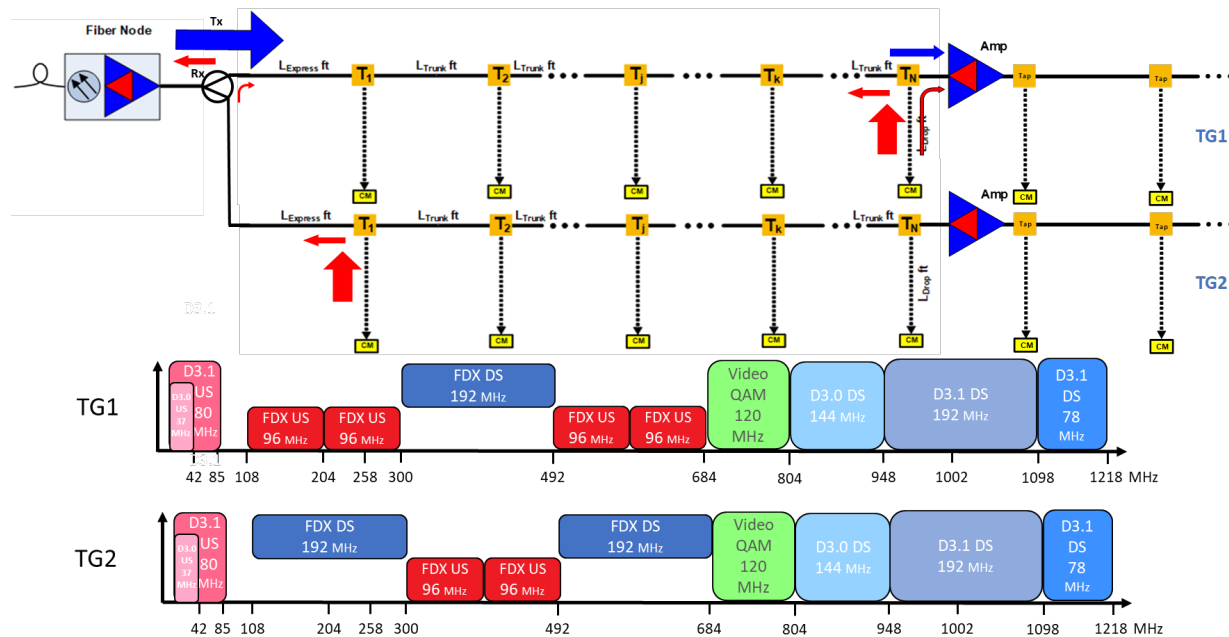


Figure 9 - Node + x cable system with FDX amplifiers

The upstream transmission similarly produces echoes in the received downstream. However, these upstream transmissions negate the reception of the downstream signal in the same band since the signals are in the same transmission group and cannot receive downstream signals in the same band as upstream signals when present from a transmitting modem in that leg of the plant.

The functional diagram for the FDX amplifier is shown in Figure 10. Note the similarity in the EC for the node of Figure 3 on the downstream port of the FDX amplifier. The same EC processing is used here resulting in a similar echo residual after cancelling the echo. This residual is present in the upstream signal. Gain and tilt are applied, and the upstream signal is then launched toward the next amp nearer to the node. Note that this launched upstream signal also produces an echo into the downstream. However, since the upstream transmission is granted in the single TG of the node leg containing the amp (due to IG extension discussed previously), the concurrent downstream signal is not received in the same frequency band.

The same situation applies for a transmitting cable modem producing echo interference in the TG of the node leg containing the amp cascade. This is depicted in Figure 10 where cable modem on the tap before

the north port of the amp transmits upstream introducing interfering leakage into the downstream due to limited tap port-to-output isolation.

The cascade of amplifiers each add additional downstream EC residual to the upstream signal. The nature of this EC residual noise buildup in an amplifier cascade and the resultant degradation to the upstream signal arriving at the node is derived in the following section.




Figure 10 - FDX amplifier with echo cancellation of the downstream echo

4. FDX Amplifier with Echo Cancellation Cascade Analysis

A model for the added noise due to downstream echo interference and the residual noise after echo cancellation is shown in Figure 11.

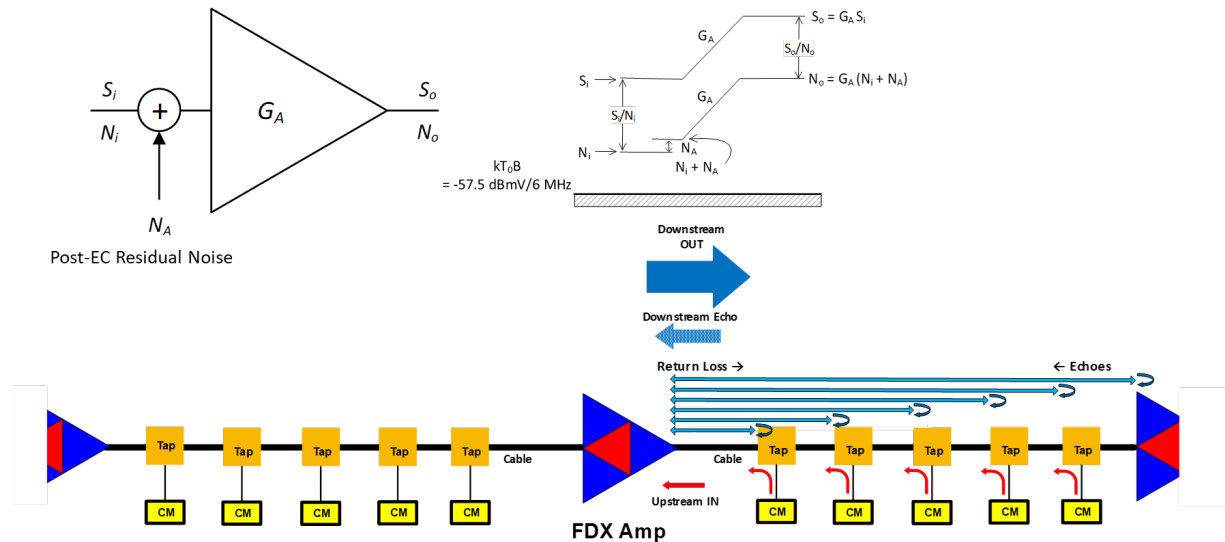


Figure 11 - FDX Amplifier Transmitted Downstream Echo Interference

The upstream input signal S_i to the downstream amplifier is accompanied by input noise N_i with an input SNR of S_i / N_i . Using downstream EC as shown in Figure 10, the EC reduces the echo to a level N_A in

Figure 11 which is substantially above the thermal noise floor. The added downstream echo residual noise N_A is added to the upstream input noise N_i .

The amplifier gain G_A including tilt compensation is applied to the input signal S_i and the combined noise $N_i + N_A$. Thus, the resulting output noise is $G_A (N_i + N_A)$ and the output SNR₀ is given by $S_i / (N_i + N_A)$.

This residual EC noise adds for each amp in the cascade as shown in Figure 12.

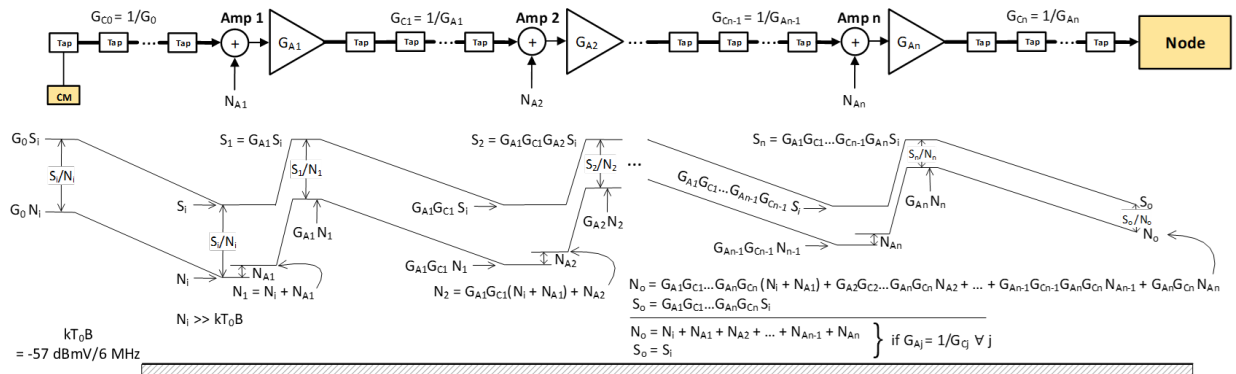


Figure 12 - Cascaded Noise of Amplifier Echo Cancellation Residuals

The node of Figure 4 receives the amplifier cascade upstream signals and accumulated EC residual noise resulting in the signal and noise levels shown at the node port input with an amp cascade CNR of S_o/N_o . This is plotted for a range of values in the bit-loading contour map of Figure 13 below.




Figure 13 - Bit-Loading Contour Map – Spec Bit-Loading SNR Thresholds

Node Upstream Background CNR vs. Downstream SIR (Signal/Echo Ratio)

The node produces its own downstream echo yielding a node signal/interference ratio (SIR) given by the upstream signal/downstream echo. The node reduces its own downstream echo and the node EC residual adds to the amp cascade EC residual. These two noise sources are uncorrelated and add on a power basis. The resulting bit-loading for the DOCSIS 3.1 spec threshold of the OFDMA receiver is shown in Figure 13 for all combinations of amp cascade CNR vs Node SIR.

Comcast utilized a Profile Management Application (PMA) using lower thresholds from lab and field measurements. Downstream PMA is typically 3 dB lower than spec. Current production OFDM PMA implementation leverages 3 dB of available headroom vs spec in setting downstream OFDM MER thresholds (4k-QAM @ 38 dB vs 41 dB; 2k-QAM @ 34 dB vs 37 dB, etc.).

Upstream is expected to be similar. The result for the bit-loading thresholds using upstream PMA is shown in Figure 14. Using PMA results in higher bit-loading values for the lower PMA thresholds.




Figure 14 - Bit-Loading Contour Map – PMA Bit-Loading SNR Thresholds
Node Upstream Background CNR vs. Downstream SIR (Signal/Echo Ratio)


5. SNR Calculations for the Node + N FDX Cable System

A representative system for most medium density designs is the single-family unit (SFU). An example of such a system design is shown in Figure 15.



A simulation was performed of the amp cascade upstream signal plus the EC residual accumulation (amp cascade CNR) as depicted in Figure 12 combined with the node EC residual SIR. Echo levels at the downstream port from the tap cascade were modeled using our return loss measurements of the most

[illegible]

Most Common Tap OEM Examples – First Tap off of Active



12



3.1 Spec



Average Bit-Loading Post-EC Echo + D3.1 US Rx + CNR 108 - 684 MHz 8.2 bits/subc

3.1 PMA

Average Bit-Loading Post-EC Echo + D3.1 US Rx + CNR 108 - 684 MHz 10 bits/subc
--

Node and Amp Upstream Input Level: 5 dBmV/6.4 MHz
Cable System (Tap Cascade) Return Loss: 24 dB

Figure 17 - Node Echo Cancellation and SNR with Node + 6 (5 dBmV/6.4 MHz)



3.1 Spec

Average Bit-Loading Post-EC Echo + D3.1 US Rx + CNR 108 - 684 MHz 10 bits/subc
--


3.1 PMA

Average Bit-Loading Post-EC Echo + D3.1 US Rx + CNR 108 - 684 MHz 11 bits/subc
--


Node and Amp Upstream Input Level: 13 dBmV/6.4 MHz
Cable System (Tap Cascade) Return Loss: 24 dB

Figure 18 - Node Echo Cancellation and SNR with Node + 6 (13 dBmV/6.4 MHz)

A series of 5, 8, 11, and 13 dBmV/6.4 MHz is shown in the following figures:



D3.1 Spec OFDMA SNR Thresholds




D3.1 PMA OFDMA SNR Thresholds


Tier 2 SFU with 24 dB tap cascade return loss and 13 dBmV/6.4 MHz upstream input level (node and amps)

*Figure 19 - Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade
(13 dBmV/6.4 MHz Upstream Input Level)*

Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade
(11 dBmV/6.4 MHz Upstream Input Level)




D3.1 Spec OFDMA SNR Thresholds




D3.1 PMA OFDMA SNR Thresholds

Tier 2 SFU with 24 dB tap cascade return loss and 11 dBmV/6.4 MHz upstream input level (node and amps)

*Figure 20- Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade
(11 dBmV/6.4 MHz Upstream Input Level)*




D3.1 Spec OFDMA SNR Thresholds




D3.1 PMA OFDMA SNR Thresholds

Tier 2 SFU with 24 dB tap cascade return loss and 8 dBmV/6.4 MHz upstream input level (node and amps)

*Figure 21 - Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade
(8 dBmV/6.4 MHz Upstream Input Level)*



D3.1 Spec OFDMA SNR Thresholds



D3.1 PMA OFDMA SNR Thresholds

Tier 2 SFU with 24 dB tap cascade return loss and 5 dBmV/6.4 MHz upstream input level (node and amps)

*Figure 22 - Node Echo Cancellation and SNR with N+1 to 6 Amp Cascade
(5 dBmV/6.4 MHz Upstream Input Level)*

The upstream throughput calculations using PMA bit-loading for an upstream receive level of 5, 8, 11, and 13 dBmV/6.4 MHz with an overhead assumption of 25% (75% of PHY Throughput) are shown in Figure 23.

Note that for a six-amp cascade, a minimum 11 dBmV/6.4 MHz upstream receive level across the entire 108 to 684 MHz FDX band (6 upstream channels) is needed for 5 Gbps throughput. Similarly, a minimum 11 dBmV/6.4 MHz upstream receive level across the 108 to 300 MHz FDX band (2 upstream channels) is needed for 2 Gbps throughput. Alternatively, a minimum 5 dBmV/6.4 MHz upstream receive level across the 108 to 396 MHz FDX band (3 upstream channels) is needed for 2 Gbps throughput.

Amp Cascade	96 MHz OFDMA Channels	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5	Channel 6
N+1	N+1 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204
N+3	N+3 OFDMA PHY Rate Mbps	960	960	960	960	960	960
	Total Cumulative NET Mbps	1172	1892	2612	3332	4052	4772
N+6	N+6 OFDMA PHY Rate Mbps	960	960	960	960	960	960
	Total Cumulative NET Mbps	1172	1892	2612	3332	4052	4772

(a) Upstream Rx @ +5 dBmV/6.4 MHz

Amp Cascade	96 MHz OFDMA Channels	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5	Channel 6
N+1	N+1 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204
N+3	N+3 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204
N+6	N+6 OFDMA PHY Rate Mbps	960	960	960	960	960	960
	Total Cumulative NET Mbps	1172	1892	2612	3332	4052	4772

(b) Upstream Rx @ +8 dBmV/6.4 MHz

Amp Cascade	96 MHz OFDMA Channels	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5	Channel 6
N+1	N+1 OFDMA PHY Rate Mbps	1152	1152	1152	1152	1152	1152
	Total Cumulative NET Mbps	1316	2180	3044	3908	4772	5636
N+3	N+3 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204
N+6	N+6 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204

(c) Upstream Rx @ +11 dBmV/6.4 MHz

Amp Cascade	96 MHz OFDMA Channels	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5	Channel 6
N+1	N+1 OFDMA PHY Rate Mbps	1152	1152	1152	1152	1152	1152
	Total Cumulative NET Mbps	1316	2180	3044	3908	4772	5636
N+3	N+3 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204
N+6	N+6 OFDMA PHY Rate Mbps	1056	1056	1056	1056	1056	1056
	Total Cumulative NET Mbps	1244	2036	2828	3620	4412	5204

(d) Upstream Rx @ +13 dBmV/6.4 MHz

Upstream Receive @ (a)+5 dBmV/6.4 MHz, (b)+8 dBmV/6.4 MHz, (c)+11 dBmV/6.4 MHz, (d)+13 dBmV/6.4 MHz

Figure 23 - Net Upstream Throughput by Bandwidth Allocation and Cascade Depth vs Receive Level

6. FDX Moves into the Field

A prototype FDX node reference design in a Node + 0 multiple tap cable system was demonstrated in our labs last year. Echo cancellation technology in the FDX node was evaluated with results exceeding expectations [2]. Following the FDX node demonstration, multiple cable modems were evaluated in the Node + 0 FDX proof of concept cable system with a cable modem reference design in each of two transmission groups (complementary upstream transmit, downstream receive Resource Blocks). This included:

2021: Full End-to-End Proof of Concept of 10G FDX

- DOCSIS 4.0 vCMTS - Software Upgrade
- FDX Node Reference Design
- FDX CM Reference Design

Plans for this year include FDX amplifier development and several trials in the field, including:

2022: Move to the Field and Introduce FDX Amps (N+x)

- Ongoing: Addition of FDX features per DOCSIS 4.0 spec into the vCMTS code
- Now: First full FDX RPD Node
- Q2: First Comcast FDX MTA
- Q3: Prototype FDX Amplifier
- Multi-Gig symmetric over N+x
- Q4: Qualification of FDX Node Hardware and Software
- Q1 2023 and beyond: Trials
- Tech Trial Multi-Gig Symmetric N+0
- Tech Trial FDX over N+x

7. Conclusion

The successfully demonstrated echo cancellation (EC) technology developed for the FDX node enables the extension to FDX amplifiers. An EC noise model for FDX amplifiers has been developed to calculate the achievable bit-loading and resultant capacity for varying numbers of cascaded amplifiers.

Sufficient fidelity in reasonable amp cascades can ensure multi-gigabit symmetric speeds up to Node + 6. A powerful PMA engine enables true, adaptive bandwidth/capacity optimization. Straightforward RF guidelines for allocations of FDX upstream OFDMA spectrum can deliver Speed Tier objectives for 2 to 5 Gbit/second symmetric peak throughput.

A Full End-to-End Proof of Concept of 10G FDX cable system was demonstrated in 2021. A prototype FDX amplifier (Node + x) will be introduced in 2022. An FDX Node + 0 trial and Node + x amp trial is planned in 2023.

Abbreviations

ASIC	Application Specific Integrated Circuit
BER	Bit Error Ratio
CM	Cable Modem
CMTS	Cable Modem Termination System
CNR	Carrier-to-Noise Ratio
dBmV	Decibel millivolts
EC	Echo Cancellation
FDX	Full Duplex DOCSIS
Gbit	gigabit
Gbps	Gigabit per second
IC	Integrated Circuit
IG	Interference Group
MER	Modulation Error Ratio
MHz	Megahertz
MTA	Multimedia Terminal Adapter
N + x	Node + x amplifiers
OFDM	Orthogonal Frequency Division Multiplex
OFDMA	Orthogonal Frequency Division Multiple Access
PHY	Physical Layer
PMA	Profile Management Application
RPD	Remote PHY Device
R-PHY	Remote PHY
SFU	Single Family Unit
SIR	Signal-to-Interference Ratio
SNR	Signal-to-Noise Ratio
TG	Transmission Group
vCMTS	Virtual CMTS

Bibliography & References

[1] R.S. Prodan, *Full Duplex DOCSIS PHY Layer Design and Analysis for the Fiber Deep Architecture*, SCTE 2017 Cable-Tec Expo

[2] R.S. Prodan, *10G Full Duplex DOCSIS Implementation Exceeds Expectations*, SCTE 2021 Cable-Tec Expo

The Impact of Wi-Fi 7 on Cable Networks

A Technical Paper prepared for SCTE by

Steve Harris

Vice President, Market Development
SCTE
140 Philips Road, Exton, PA 19341-1318
610-594-7324
sharris@scte.org

Paul Rodrigues

Director, Field Education
SCTE
140 Philips Road, Exton, PA 19341-1318
601-594-7306
prodrigues@scte.org

Table of Contents

Title	Page Number
1. Introduction.....	4
2. The 10G pillars.....	4
3. Why Wi-Fi 7?.....	5
4. Release 1 Features.....	6
4.1. Multi-Link Operation.....	6
4.2. Low Complexity AP Coordination.....	9
4.3. 320 MHz Channels.....	11
4.4. 4K QAM.....	11
4.5. OFDMA Enhancements.....	11
4.5.1. Multiple RUs Per STA.....	11
4.5.2. Preamble Puncturing.....	13
4.5.3. PPDU Frame Format.....	14
5. Release 2 Features.....	15
5.1. MIMO Enhancements.....	15
5.1.1. Channel Sounding Optimization.....	15
5.2. HARQ.....	16
5.3. Low-Latency Operation.....	16
5.4. Advanced AP Coordination.....	17
5.4.1. Multi-AP Joint Transmission and Reception.....	18
5.4.2. Coordinated Beamforming.....	18
6. Security.....	18
7. Benefits to Customers.....	19
7.1. MDUs.....	19
7.2. Business Services.....	20
8. Wi-Fi Performance Management.....	21
9. Conclusion.....	23
Abbreviations.....	24
Bibliography & References.....	26

List of Figures

Title	Page Number
Figure 1 – 802.11be Release Timeline.....	4
Figure 2 – 10G.....	5
Figure 3 – Multi-link Device.....	6
Figure 4 – MLO Architecture.....	7
Figure 5 – Asynchronous Multi-Link Channel Access.....	7
Figure 6 – Synchronous Multi-Link Channel Access.....	8
Figure 7 – Multi-Link Frames.....	9
Figure 8 – 802.11ax Spatial Reuse Drawback.....	10
Figure 9 – 802.11be Coordinated Spatial Reuse.....	10
Figure 10 – 80 MHz Tone Map.....	12
Figure 11 – Wi-Fi 6 Preamble Puncturing.....	13
Figure 12 – Preamble Puncturing Frame.....	14

Figure 13 – Wi-Fi 7 PPDU.....	14
Figure 14 – Multi-Link Joint Mode.....	17
Figure 15 – MDU.....	19
Figure 16 – SCTE GAP.....	20
Figure 17 – Heatmap.....	21
Figure 18 – Signal to Noise Ratio.....	22
Figure 19 – RSSI.....	22
Figure 20 – Wavelength of 2.4, 5, and 6 GHz.....	23

List of Tables

Title	Page Number
Table 1 – OFDMA Tone Map.....	12
Table 2 – Multi-Rus for Large-Size RUs.....	13

1. Introduction

Today, Wi-Fi is the predominant way telecommunication customers connect their untethered devices to the Internet, as well as access content like video, virtual reality, and gaming. In the near future, operators will shift compute closer to the end user, expanding the requirements for better Wi-Fi edge compute connectivity. The original Institute of Electronics Engineers (IEEE) 802.11 standard was released in 1997. Since then, the IEEE 802.11 working group has developed and released many amendments to keep up with the growing bandwidth demands of our customers and the industry. During this time, telecommunication operators, CableLabs[®], the SCTE[®], and vendors developed and deployed DOCSIS and other access network technologies to keep up with the same bandwidth demands.

Today operators are moving towards multi-gigabit offerings and are deploying Wi-Fi solutions using Wi-Fi 6 and Wi-Fi 6E customer premises equipment (CPE) to support these offerings. With the telecommunication industry on the road to 10G and beyond, the wireless industry is working on the next generation Wi-Fi standard, IEEE 802.11be, extremely high throughput (EHT). The IEEE 802.11be task group, or just 802.11be, is leading the development of the EHT amendment to the IEEE 802.11 standard. 802.11be is the amendment on which the Wi-Fi Alliance[®] Wi-Fi 7 certification will be based.

The 802.11be amendment is currently being developed and is scheduled to be released in May 2024. To help meet the tight development timeline, 802.11be has two sets of release features known as Release 1 and Release 2. The Release 1 features are outlined in Draft 1.0 and Draft 2.0 of the amendment. The Release 2 features will be defined in Draft 3.0 and Draft 4.0. Figure 1 depicts the development timeline, which the 802.11be task group continually evaluates. The Wi-Fi Alliance Wi-Fi 7 certification requirements will be released closer to the end of the 802.11be development cycle, close to the final amendment.

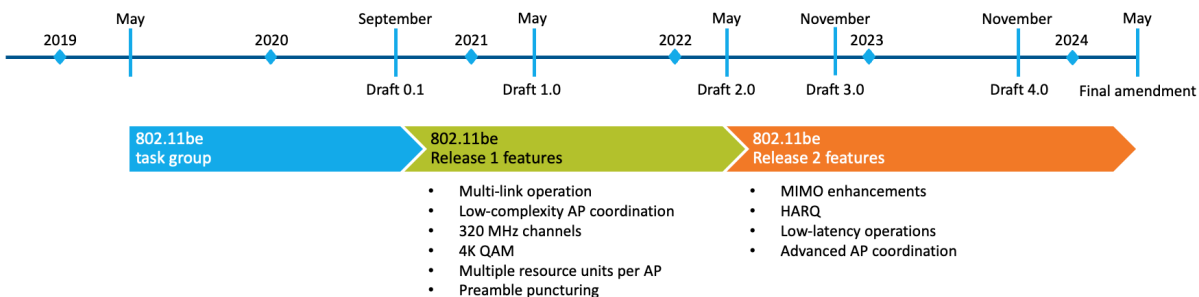


Figure 1 – 802.11be Release Timeline

This paper will explore the 802.11be releases and their features, along with key media access control (MAC) and key physical (PHY) layer techniques. Furthermore, the paper will explore how Wi-Fi 7 will impact the customer quality of experience (QoE) and enable telecommunication operators to support their 10G platform initiatives all the way to the customer premises.

2. The 10G pillars

As connectivity, in particular, Wi-Fi, becomes more crucial to how we live, work, learn, and play, we need a network that will support the hyper-connectivity of our future customers. In 2019, the industry announced the next leap forward in connectivity with the 10G platform, which also encompasses Wi-Fi.

The NCTA–The Internet & Television Association, CableLabs, GIGAEurope, and SCTE have been driving this initiative and the platform forward in the industry.

The 10G platform initiative is a combination of technologies that will deliver multi-gigabit symmetrical Internet services to our consumers in the United States and around the globe. The Wi-Fi 7 certification will dovetail neatly with our intelligent 10G platform initiative in the telecommunication industry. 10G will ensure our customers gain access to the most advanced infrastructure in the marketplace and over Wi-Fi. The pillars of the 10G network span across throughput/speed, low latency, reliability, and scalable security. These pillars will allow our customers to take full advantage of future services over a Wi-Fi network. These customers may be residential, multiple dwelling units (MDU), or business services.

The throughput/speed pillar focuses on consumer bandwidth, enabling Wi-Fi EHT as well as multi-gigabit symmetrical experiences. The low latency pillar delivers on customer QoE, enabling future augmented reality (AR), virtual reality (VR), artificial intelligence (AI), and new digital Wi-Fi experiences like holodecks and lightfield displays. A poor Wi-Fi connection that adds a delay of a few milliseconds (ms) will produce negative user experiences. Lower latency will reduce jitter and delay in a Wi-Fi network, optimizing its performance. While the reliability pillar addresses Wi-Fi network issues proactivity before they inhibit our networks and customers. Wi-Fi network reliability will be required to grow as our consumer devices increase per household or per business. Finally, the security pillar strengthens our more complex Wi-Fi networks' confidentiality, integrity, and availability (CIA) triad, enabling safer symmetrical communications for all.



Figure 2 – 10G

3. Why Wi-Fi 7?

The simple answer to “Why Wi-Fi 7?” is speed. While that is a good answer, especially when considering new features such as wide 320 MHz channels and dense 4K quadrature amplification modulation (QAM). These features enable a potential data rate of 30 to 40 Gbps. However, that does not tell the entire story here.

Many of the features for Wi-Fi 7 are designed to make better use of the available RF bandwidth. With the release of the 6 GHz band, up to 1.2 GHz of new spectrum became available in the USA and other parts of the world. Some areas, such as Europe, only have 500 MHz available initially, but that is still an incredible amount of new bandwidth. However, when designing the 802.11be amendment, the IEEE made some decisions on optimizing the use of all the spectrum. For example, a Release 1 feature called multi-link operation (MLO) allows a device to operate, transmit and receive in more than one band at a time. The AP may access the available links and determine which link will provide the best connection for sending data packets. The AP may even spread the data over multiple links to increase throughput and reduce latency.

Wi-Fi 7 also enables devices to better coordinate their use of the available spectrum. By coordinating resources and adjusting transmit power to reduce interference, nearby devices can share the available bandwidth and spectrum. This can provide significant benefits in high-density deployments, such as MDU and college campuses.

So, to answer the question “Why Wi-Fi 7?”. Wi-Fi 7 is designed to enable Wi-Fi networks to work together to maximize the available Wi-Fi airtime resources to provide faster and more reliable connections.

4. Release 1 Features

As mentioned earlier, the Release 1 features are defined in Draft 1.0 and 2.0 of the amendment. Some features improve upon Wi-Fi 6 features; these include 320 MHz channels, 4K QAM, multiple resource units (RU) per station (STA), and orthogonal frequency division multiple access (OFDMA) enhancements. The other features, multi-link operation (MLO) and low-complexity AP coordination, are new features. Let us explore the new features first.

4.1. Multi-Link Operation

Since Wi-Fi 4 (IEEE 802.11n), wireless APs and STAs have had two radios; one operating in the 2.4 GHz band and another operating in the 5 GHz band. With the release of Wi-Fi 6E and the addition of the 6 GHz band, STAs may have three radios, or potentially more with multiple 5 GHz radios. Regardless of how many radios an STA has, it will only use one radio at a time to carry traffic. MLO looks to change this.

The concept is simple: use any available radio to transmit or receive data. A challenge for multi-link is that the upper layers must treat the combination of the radio interfaces as a single interface with a single MAC address. To accomplish this, a new device concept, the multi-link device (MLD), is set in the amendment.

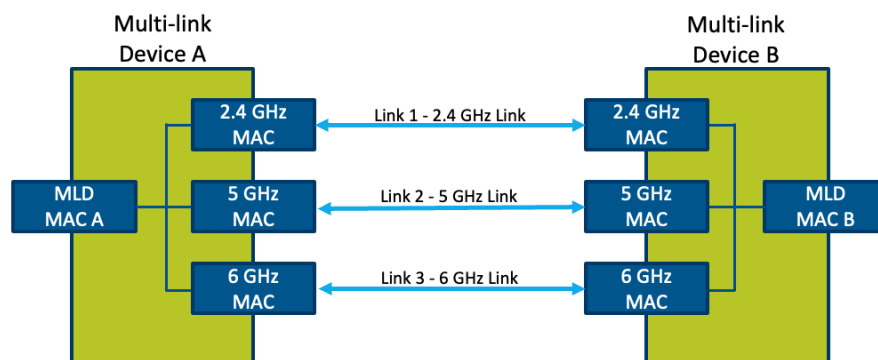


Figure 3 – Multi-link Device

The MLD splits the MAC layer into two parts the upper MAC (U-MAC) and the lower MAC (L-MAC). The L-MAC is the MAC that is tied to the PHY interface, and each wireless radio will have a unique L-MAC address. The U-MAC aggregates the L-MAC to a single MAC and performs link agnostic operations. The figure below shows the L-MAC and U-MAC layout and how they correspond to the upper layers.

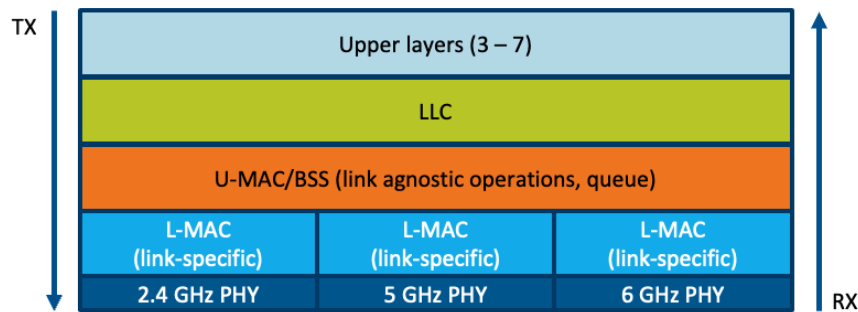


Figure 4 – MLO Architecture

The MLD must determine how to send the traffic. The MLD may be configured in two modes of operation, restricted mode and dynamic link switch mode. The restricted mode uses each link to send different types of frames. One link acts as a data plane and is only used for the data frames and the ACKs. The other link acts as a control plane and is used for the management and the other control frames. In the dynamic link switch mode, each link can send data or control frames, enabling Wi-Fi load balancing.

MLD's channel access can be either asynchronous or synchronous. Asynchronous is the preferred method since it provides more throughput, and when this method is used, it is called the simultaneous transmission reception (STR) mode. When using asynchronous channel access, the MLDs can transmit and receive at the same time using different bands. As noted in the image below, the 2.4 GHz link starts transmitting downstream, and then the 6 GHz link begins receiving data upstream. A little later, the 5 GHz band begins sending data downstream, while the 2.4 GHz sends different data downstream, and the 6 GHz link receives data upstream.

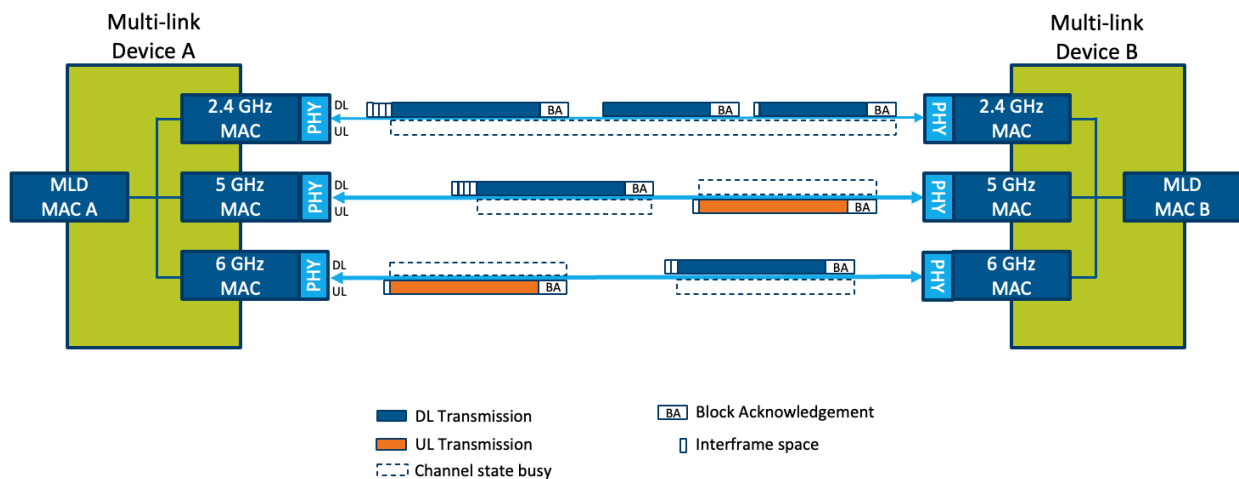


Figure 5 – Asynchronous Multi-Link Channel Access

Synchronous multi-link channel access sends the same data on both links. Synchronous channel access adds redundancy at the cost of reduced throughput. Synchronous channel access is better suited for environments with higher interference. If one link has a high amount of interference, the receiving MLD will examine each packet of the data stream and keep only the good ones. When the MLD uses synchronous channel access, it is in non-simultaneous transmission reception (NSTR) mode. In the image

Links do not have to stay in either asynchronous or synchronous mode. They can switch back and forth based on the link conditions. For example, an MLD has two links operating in asynchronous mode. During the transmission, link #1 gets a high number of errors. The MLD responds by switching to a synchronous transmission to ensure all the data is received. Once the conditions improve and the link quality is restored, the MLD can change back to the asynchronous transmission mode.

Multi-link channel access can potentially increase the time it takes for an STA to scan and discover an AP and its capabilities since it would have to send information for each link. To reduce the need for the STA to scan each interface, the reduced neighbor report (RNR) is used in the management control frames. Each link will provide information about the other links in the RNR, removing the need for the STA to scan the other interfaces.

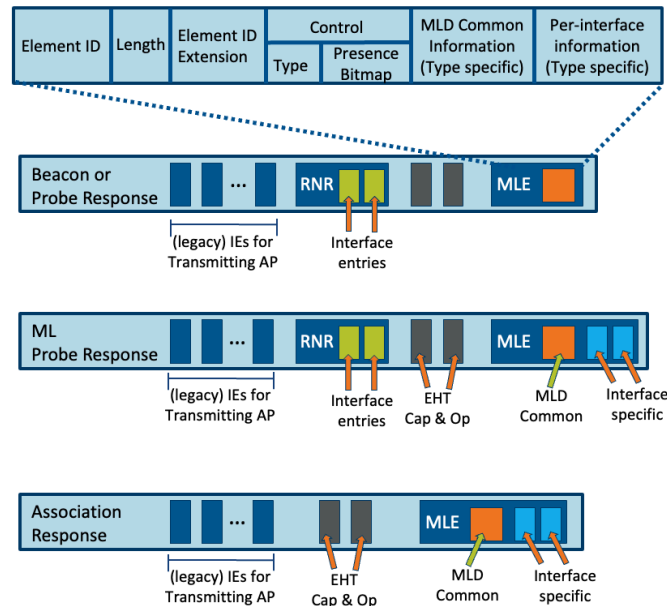


Figure 7 – Multi-Link Frames

Another important aspect of the multi-link channel access operation is the power-saving ability. Having multiple RF radios sending and receiving frames is not an efficient use of power, especially on battery-powered handheld devices. 802.11be will use the traffic indication map (TIM) and the target wake time (TWT) features to address this.

TIM uses beacons to inform the STAs that the AP has information for them. TIM uses an STA ID that is stored in a bitmap. In that bitmap, there is a bit that indicates if there is data for that STA. A binary one indicates there is, and the STA must wake up. A binary zero means there is no data for the STA, and it can stay in snooze mode. For TIM to work with MLDs, a link indication field is added to the bitmap. The link indication informs the STA which link has the data waiting for it.

TWT is based on a TWT schedule that is negotiated between the AP and the client. The TWT schedule includes the wake-up time, the wake interval, and the wake duration for the clients. With multilink, the MLD will negotiate the TWT schedule for each link with the AP. If all links follow the same schedule, then the MLD only needs to negotiate one TWT schedule.

4.2. Low Complexity AP Coordination

Environments such as MDUs where multiple APs are using the same channel and transmitting unique service set identifiers (SSIDs) create the potential for a high amount of Wi-Fi interference. Each AP is a basic service set (BSS), and when they overlap, this creates overlapping BSS (OBSS) interference, which impacts the quality of the wireless signal.

AP coordination can significantly improve Wi-Fi performance in these environments. Due to the complexity of AP coordination, the 802.11be task group split the features into two parts. Release one establishes the technologies used for low-complexity AP coordination, and release two sets the standards for advanced AP coordination. The AP coordination proposed in the 802.11be amendment identifies requirements for primary AP and secondary APs. These APs can be connected via cabling, but they do

not need to be connected. However, the secondary APs need to be able to communicate with the primary AP. The secondary APs do not need to hear each other.

The first method for AP coordination in the amendment builds on the spatial reuse (SR) feature introduced in 802.11ax. 802.11ax uses features such as BSS coloring and power management to handle OBSS interference. It controls the interference by managing the power of the secondary device. When the interference is detected, the primary AP will grant the secondary AP an opportunity to transmit simultaneously but using a lower power setting.

The drawback with this solution impacts a device trying to communicate with the secondary AP and can hear the primary AP. As shown in the figure below, STA 2, which is connected to AP 2, may interpret the signal from AP 1 as interference. If the signal is strong enough, it will not be able to communicate with AP 2.




Figure 8 – 802.11ax Spatial Reuse Drawback

802.11be looks to improve the 802.11ax SR with a coordinated spatial reuse (CSR) system. With CSR, when an AP that is part of a coordinated group receives a transmit opportunity (TXOP), it will coordinate with the other APs to share the TXOP. To reduce the impact of OBSS interference, the coordinated APs will adjust their power levels. Using the scenario mentioned above, AP 1 and AP 2 will adjust their transmit power relative to STA 2. AP 1 will reduce its transmit power to reduce the interference for STA 2, so it can communicate with AP 1.




Figure 9 – 802.11be Coordinated Spatial Reuse

The second method proposed for low complexity AP coordination is coordinated orthogonal frequency division multiple access (Co-OFDMA). OFDMA was introduced as part of the 802.11ax amendment, and it enabled APs to schedule time and frequency slots called resource units (RU) for devices to transmit and receive data. Co-OFDMA enables coordinated APs to share those resource units. APs will work together

to schedule RUs for devices to use. APs could schedule the same RU slots as long as the signals do not interfere with each other.

CSR and Co-OFDMA are being developed for Release 1 of the 802.11be amendment. We will look at the other AP coordination features proposed for Release 2 later in this paper.

4.3. 320 MHz Channels

Starting with Wi-Fi 4 (IEEE 802.11n), channel bonding was used to increase the potential data rate. Wi-Fi 4 added 40 MHz channels, using two 20 MHz channels. Wi-Fi 5 (802.11ac) added the capabilities for 80 and 160 MHz channels. In contrast, Wi-Fi 6 (802.11ax) did not expand the channel size options. However, it added the 6 GHz band, which made using one or two 160 MHz channels possible.

Wi-Fi 7 looks to expand the channel bandwidth to a massive 320 MHz channel (sixteen 20 MHz channels). The 320 MHz channels will only be possible in the 6 GHz band. The 320 MHz channels can also be non-contiguous, with two 160 MHz channels. If 320 MHz is unavailable, another option is a 240 MHz channel, which is made using one 80 MHz channel and another 160 MHz channel. Having a super wide 320 MHz channel sounds great. However, it will not always be possible, especially in dense AP environments. In these environments, channel reuse is important, and despite the 1.2 GHz of bandwidth in the 6 GHz band, there are only three 320 MHz channels.

4.4. 4K QAM

Another enhancement to previous Wi-Fi generation PHYs is increasing the modulation from 1024 QAM to 4096 QAM (4K QAM). 4096 QAM can potentially increase the data rate by 20%, compared to 1024 QAM. The challenge with 4096 QAM is that the SNR must be extremely high. Specifically, 4096 QAM requires an SNR of 40 dB, which will be difficult for many wireless networks to achieve. It can happen using beamforming in non-crowded environments, but it will be challenging in the best environments.

4.5. OFDMA Enhancements

One of the most significant additions to Wi-Fi 6 was OFDMA. Wi-Fi 7 looks to improve OFDMA to use the bandwidth more efficiently. This next section will look at these enhancements, including multiple RUs per STA, preamble puncturing, and a new physical protocol data unit (PPDU).

4.5.1. Multiple RUs Per STA

OFDMA schedules RUs for devices to use for uplink or downlink communication. The AP will allocate an STA the RU size based on the STA's bandwidth requirements. A RU is made up of a group of subcarriers called tones, and the size of the RU varies based on the number of tones. For a 20 MHz channel, the smallest RU has 26 tones, and the largest has 242 tones. The 26-tone RU is called an RU26 tone map, allowing 9 RUs in a 20 MHz channel. The table below shows the tone maps.

Table 1 – OFDMA Tone Map

RU Tone	Number of Tones			
	20 MHz Channel	40 MHz Channel	80 MHz Channel	160 MHz Channel
RU26	9	18	37	74
RU52	4	8	16	32
RU106	2	4	8	16
RU242	1	2	4	8
RU484	N/A	1	2	4
RU996	N/A	N/A	1	2
RU2 x 996	N/A	N/A	N/A	1

The image below shows the possible tone maps for an 80 MHz channel. The yellow lines between the tones are direct conversion, guard, and null tones.




Figure 10 – 80 MHz Tone Map

The 802.11ax standard allowed multiple devices to be allocated a single RUs for each TxOp. An example cited in the *Current Status and Directions of IEEE 802.11be, the Future Wi-Fi 7* paper, has two STAs sharing an 80 MHz channel. The AP grants a 242 tone RU to the 1st STA, and the 2nd STA can be granted a maximum of a 484 tone RU. That leaves 25% of the bandwidth unused. If the 2nd STA has more data to transmit, it must wait for its next TXOP.

802.11be looks to improve the spectrum efficiency by enabling APs to assign multiple RUs to a single STA. For the example cited above, the 1st STA would be assigned a 242-tone RU, and then the 2nd STA could be assigned a 484-tone RU and a 242-tone RU. This effectively gives the 2nd STA 726 tones to send data. With the move to larger channels, assigning multiple RUs to a single STA can significantly improve the ability of APs to use the available bandwidth efficiently. The task group is investigating how many RUs will be assigned to a single STA.

802.11be separates the RUs into two groupings based on their size:

- Small-size RUs: RUs less than 242 tones, ex. RU26, RU52, RU106
- Large-size RUs: RUs equal to or greater than 242 tones, ex. RU242, RU484, RU996

The 802.11be draft amendment lists all the possible multi-RU (MRU) possibilities in an MRU index. The only small-size MRUs are RU78 (RU52+RU26) and RU 132 (RU106+RU26). The 802.11be draft standard identifies the large-size MRUs based on the channel size. The following are mandatory RUs specified in the standard.

Table 2 – Multi-Rus for Large-Size RUs

Bandwidth	RU	Mandatory in non-OFDMA for:
80 MHz	484+242	AP, STA
160 MHz	996+484	AP, STA
	996+(484+242)	AP, STA
240 MHz	3×996, 2×996+484, 2×996 (any 2)	AP, STA
320 MHz	4×996, 3×996+484, 3×996 (any 3)	AP, STA

4.5.2. Preamble Puncturing

Wi-Fi 5 introduced dynamic channel bandwidth to go along with the large 80 MHz and 160 MHz channels. When an STA connects to a client, it connects on the primary 20 MHz channel. If the secondary 20 MHz channel is available, the connection expands to be a 40 MHz channel. Next, it will attempt to connect to the secondary 40 MHz. It will first connect to the primary 20 MHz of the secondary channels and then connect to the secondary 20 MHz of the secondary channels. When all four 20 MHz channels are connected, the STA has an 80 MHz channel to use. However, if the initial secondary channel is busy during this process, only the initial 20 MHz primary channel is used. The AP cannot use any of the secondary channel bandwidth, even if it is available.

Wi-Fi 6 introduced preamble puncturing, which enables an STA to avoid using a 20 MHz portion of a large channel (40 MHz, 80 MHz, and 160 MHz) that is busy. Using the 80 MHz channel example again. The STA connects to the AP on the primary 20 MHz channel. The AP will then check if the secondary 20 MHz channel is busy. If the channel is busy, the AP will skip the channel and attempt to connect to the secondary 40 MHz channel. Wi-Fi 6 identified eight bandwidth modes. Modes 0 – 3 are for the standard channel bandwidths with no puncturing, 20 MHz, 40 MHz, 80 MHz, and 160 MHz. Modes 4 – 7 are the puncturing modes, which are pictured below.

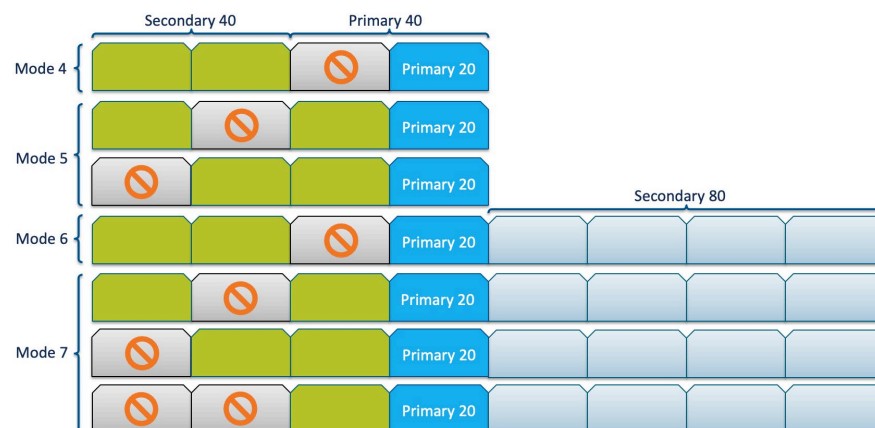


Figure 11 – Wi-Fi 6 Preamble Puncturing

802.11be will build on existing puncturing techniques by expanding the modes to cover the 240 MHz and 320 MHz channels. Also, preamble puncturing is being applied to primary channels. This technology is called preamble puncturing because the preamble is removed from the frame for the channel that is being punctured.



Figure 12 – Preamble Puncturing Frame

4.5.3. PPDU Frame Format

When designing the PPDU, the 802.11be task group aimed to have backward compatibility with previous Wi-Fi standards and provide a framework for future standards. The backward compatibility is maintained by the first four parts of the frame, and the rest of the frame is designed to support future Wi-Fi 7 and future versions.




Figure 13 – Wi-Fi 7 PPDU

The 802.11be PPDU begins with using the same legacy preamble training fields as previous Wi-Fi versions. These fields are the legacy short training field (L-STF) and the legacy long training field (L-LTF), and they are used for frame detection and synchronization. The next field is called the legacy signal field (L-SIG), and it is a carry-over from 802.11a to indicate the MCS and frame length. However, since Wi-Fi 4, the values in the field were not related to the MCS. The data in there is fake data as placeholders to maintain compatibility. Wi-Fi 4 and 5 send the MCS and frame length data in the following field. Wi-Fi 6 transmits it after the following field, the repeat legacy signal field (RL-SIG). Wi-Fi 7 will also use the RL-SIG field.

The PPDU frame variation for 802.11be starts with the universal signal (U-SIG) field. The U-SIG field is created with an eye toward compatibility with future Wi-Fi standards. The U-SIG has two parts: the version independent and the version dependent fields. The version independent field includes a PHY version identifier, an uplink (UL)/downlink (DL) flag, the BSS color field, the TXOP duration, and the bandwidth. The 802.11be task group is still defining the information sent in the version-dependent fields of the U-SIG. Some of the information may include the guard interval duration, EHT-STF, EHT-LTF size, space-time block coding, and information about specific 802.11be features.

The EHT-SIG field will provide information not included in the U-SIG but is needed to implement 802.11 features. The EHT-SIG has two parts: the common field and the user-specific field. The common field has the RU allocation information, the MCS, the guard interval duration, and other signal information. The user-specific field will be present in multi-user (MU) frames and will have information directly related to an STA. The last two fields of the EHT preamble are the EHT-STF and EHT-LTF. These training fields are similar to the training fields at the beginning of the PPDU. They provide synchronization for multiple in multiple out (MIMO) configurations.

5. Release 2 Features

As per the original timeline, the Release 2 features will be added in Draft 3.0 and 4.0. Some of these features build on existing 802.11 standards and the Release 1 features. These features are still being investigated at the time of this writing, and the 802.11be task group has not given final approval. They are reviewing multiple proposals for implementing these features. The features that are part of Release 2 include MIMO enhancements, hybrid automatic repeat requests (HARQ), low latency operation, and advanced AP coordination.

5.1. MIMO Enhancements

Since MIMO was introduced with 802.11n (Wi-Fi 4), MIMO has been enhanced with each release. Wi-Fi 7 continues that trend by doubling the number of spatial streams from 8 to 16. This can potentially double the throughput an AP can handle. However, there are some restrictions on the 16 spatial streams. The first restriction is that the maximum number of STAs using spatial multiplex per AP is eight. The second restriction is that each STA can have a maximum of 4 spatial streams. This will not restrict client capabilities significantly since most mobile devices support a maximum of two spatial streams for power efficiency reasons. As the bandwidth needs of the devices increase, manufacturers can add additional spatial streams up to 4. The other reason for limiting the number of spatial streams is to limit complexity and overhead.

5.1.1. Channel Sounding Optimization

For MU-MIMO to operate effectively, the AP and clients must have accurate information about the channel quality. This is done through channel sounding, and previous 802.11 standards implemented two different versions, implicit and explicit sounding. 802.11ac and 802.11ax use explicit sounding, which replaced the implicit sounding introduced in 802.11n.

The explicit sounding used in 802.11ax works by the AP sending out a null data packet announcement (NDPA). This informs the STA about the null data packet (NDP) that the AP is sending. The STAs use the NDP to assess the channel's status and provide channel state information (CSI) for the channel to the AP. The CSI is sent via a beamforming report (BFR). The BFR provides the SNR for each spatial stream, a comparison of the SNR for the subcarrier compared to the spatial stream, the Givens rotation angles (ϕ and ψ), and other relevant channel quality information. The CSI for a 160 MHz channel with one spatial stream will provide BFR information for every 16th subcarrier, for a total of 128 subcarriers. As the number of spatial streams and channel size increases, the amount of CSI data sent from the clients to the AP will increase. The overhead will increase dramatically, and the information will not be relevant as the channel state constantly changes.

At the time of writing this paper, the 802.11be task group is investigating multiple methods to reduce the overhead. Some of these solutions use implicit sounding. Even though implicit sounding was first introduced in 802.11n, it was never really implemented in APs. The reason was that the APs could not perform the AP-self calibration needed. Self-calibration is needed since the UL and DL signals will vary slightly for each antenna. So, the APs must compensate for these differences. Implicit sounding works by having the STA send NDP sounding information in the UL. The AP will measure the channel using the sounding information. It takes less time to send the NDP than it does to send the BFR, especially when there are many spatial streams and STAs. Reducing the time needed to receive the channel information is more current.

To improve the self-calibration, newer local AP self-calibration techniques can be used where the STAs are not involved in the calibration process. The AP will select a reference antenna and send out a pilot signal from every antenna. It estimates each antenna's baseband to RF gain variations and makes the necessary adjustments. An example of how implicit sounding is used in one of the proposals begins with the AP sending out a trigger frame requesting the STAs to send the UL NDPs. The AP analyzes the NDPs from all the STAs. It then can send beamformed data to each STA. This method is estimated to save approximately 60% of the airtime.

5.2. HARQ

Hybrid automatic repeat requests (HARQ) aims to reduce the amount of traffic on the wireless network by reducing retransmissions. Current Wi-Fi implementations require an entire packet to be retransmitted if the receiving device did not receive a complete packet. This will continue until the receiver gets a complete packet and can lead to a high number of retransmissions and impact the performance of a network. HARQ looks to improve this by combining the failed- transmissions, which will improve the SNR, and the receiver should be able to rebuild the complete packet with the information that it has. In the end, this will reduce the number of retries on the network. Another benefit is that since the SNR is being improved, a higher MCS can be set for the connection.

The task group is evaluating three methods to implement HARQ: chase combining (CC), punctured CC, and incremental redundancy (IR). The lowest complexity solution is CC. With CC, every retry has the same information as the failed packet. Since the information is the same, it makes rebuilding the packet less complex. However, there is the potential for more traffic than other HARQ methods. The second method being examined is punctured CC, which reduces the amount of overhead compared to CC. With punctured CC, the transmitter only repeats the part of the failed transmission. On the receiver side, this version will require more computation than CC. The most complex and bandwidth-efficient method is IR. IR requires the transmitter to use different codewords to represent the same information. The receiver gets more information to rebuild the original packet by using a different set of codewords.

5.3. Low-Latency Operation

Network latency is an essential part of all networking solutions. Residential applications such as virtual reality gaming and commercial applications such as industrial IoT are real-time applications (RTA), and they need networks to have very low latency. The IEEE has done a lot of work looking at time-sensitive networking (TSN) applications. They have established an 802 TSN task group to study these solutions. Most of their work applied to 802.3 Ethernet networks, but now the 802.11 working groups are looking at ways to implement TSN solutions with Wi-Fi.

The solutions being investigated focus on improving the worst-case latency instead of the average latency. Many of the 802.11be features outlined in this paper are a big part of the solutions. Multi-link operation, AP coordination, multi-RU per station, and other techniques are designed to use the spectrum more efficiently, resulting in data packets waiting less to be transmitted.

Two operations scenarios are identified by the 802.11be task group, managed and unmanaged. Examples of unmanaged scenarios are home networks and Wi-Fi hotspots. Managed scenarios are more the corporate and enterprise networks that use controllers to manage the network. The big difference in these scenarios is that the controllers and the APs can better deal with interference and optimize the environment. Unmanaged networks can do things to control latency and jitter, but they have little control of the impact of outside interference.

One way to reduce the latency for RTAs is to use QoS. Wi-Fi uses enhanced distributed channel access (EDCA) to provide QoS in Wi-Fi networks. Currently, EDCA uses four categories of traffic; background, best effort, video, and voice. Voice has the highest priority for getting channel access. The 802.11be task group is investigating adding additional categories. While voice traffic will continue to have higher priority, some applications need to have higher priority than video. A typical example used is gaming. Gaming does not require as much bandwidth as high-resolution video, but the latency must be much less. Adding a category for gaming that is a higher priority than video can reduce the worst-case latency for gaming applications.

The challenge is getting access to the medium. In Ethernet networks, collisions can clear the medium so that the higher priority traffic can get access. Wi-Fi does not have collision detection. The device using the medium cannot do any sensing on the channel while transmitting. The solution for this is different on managed and unmanaged networks. With managed networks, the controllers can coordinate the channel access. They can give access to the channel to the device that has the higher priority traffic. This can be done via scheduling using hybrid coordinated function-controlled channel access. This function is not used in WLANs today. Another option is using a trigger frame. The controller can allocate medium access to high-priority applications using short trigger frames.

On unmanaged networks getting access to the medium can be more difficult. 802.11be enhancements like multi-link operation will help with this since there are multiple mediums to carry traffic. One solution is a busy tone. A busy tone can let a device that is transmitting lower priority traffic know there is higher priority traffic waiting to be sent. When the busy tone is detected, the device stops transmitting, so the higher priority traffic can be sent. Multi-link devices can use another link to send the busy tone. Another way multi-link is being used to reduce latency is by implementing asynchronous communication, also known as joint mode. The AP will use multiple links to send different data in joint mode.




Figure 14 – Multi-Link Joint Mode

5.4. Advanced AP Coordination

Release 2 looks to add more advanced ways of doing AP coordination. The two methods being investigated for Release 2 are coordinated beamforming (CBF) and multi-AP joint transmission and reception.

5.4.1. Multi-AP Joint Transmission and Reception

Multi-AP joint transmission and reception creates a way for multiple APs to serve a single STA. This is a similar concept to multi-link operation, but in the reverse scenario, where there is one STA being served by multiple APs. However, the challenge for this system is the synchronization of the APs. This challenge is extremely difficult in the DL as it requires a high-speed, low latency backhaul to create synchronization among the APs. The backhaul is needed because it must provide the data to all the APs to transmit. The APs need to be in synchronization so they know which parts of the data were transmitted and received. For the UL, multi-AP systems distributed successive interference coordination (SIC) can be used to improve the data reception at the APs. With SIC, each AP receives the data from their STAs. The data is then sent to the other APs for interference subtraction. The APs remove the interfering signal from the received signal to get their data. The other method being investigated for the uplink is joint frame reception. With this method, the APs process the data from all the STAs. The high-bandwidth backhaul is needed to synchronize the APs for handling the STA traffic. The APs share signal time stamps to stay in synchronization.

5.4.2. Coordinated Beamforming

CBF is also called null steering. It is called null steering because the AP sending the traffic will attempt to cancel or null the interference from neighboring STAs. The nulling of the interference happens while the AP is creating the beamforming signal to the STA. To know what interference there is, the AP must collect CSI information for all nearby STA, including the ones it is not serving.

For UL transmissions, each AP must collect interference information for neighboring STAs. During the UL reception, the AP configures its receiver so it can receive data from its STAs while ignoring interference from other STAs.

6. Security

The Wi-Fi Alliance sets security standards for Wi-Fi networks as part of its certification process. The standards identify which Wi-Fi protected access (WPA) version and modes that a device must support. The Wi-Fi Alliance has not announced any new security requirements around Wi-Fi 7. At the time, it is expected that the Wi-Fi 7 device certification will have similar security requirements as Wi-Fi 6 and Wi-Fi 6E. Those certifications brought significant updates to Wi-Fi network security. Wi-Fi 6 added the following security updates:

- WPA3-Personal
 - Simultaneous authentication of equals (SAE) for authentication and association
 - Management frame protection (MFP) used to combat deauthentication attacks
- WPA3-Enterprise
 - 802.1X/EAP
 - Management frame protection (MFP)
 - 192-bit security key optional
- Enhanced Open
 - Provides encryption on open networks

To obtain Wi-Fi 6 certification, the Wi-Fi Alliance requires that devices support these new and legacy security standards such as WPA and WPA2. To enhance the security of devices operating in the 6 GHz band, the Wi-Fi 6E certification requires that devices only support the new security standards. Wi-Fi 6E devices do not support legacy standards.

Some of the questions around Wi-Fi 7 security are:

- Will the same security requirements from Wi-Fi 6 and 6E be carried over as they are for Wi-Fi 7?
- Will there be a separate Wi-Fi 7E certification, similar to 6E?
- Will multi-link devices support legacy security standards on non-6GHz channels?

The Wi-Fi Alliance has not established a time frame for Wi-Fi 7 certification and security requirements. It is expected to be ready in early 2024 to align with the release of the 802.11be standard.

7. Benefits to Customers

As noted in previous sections of the paper, Wi-Fi 7 increases the potential maximum data rate. Wi-Fi 7 takes advantage of the extended available spectrum and increases its efficiency for spectrum using wider channels, preamble puncturing, and 4K QAM. Besides the throughput/speed benefits, Wi-Fi 7 offers lower latency and better reliability in high interference environments using Co-OFDMA, RUs, and tones. The benefits in reliability extend beyond residential and into MDU and business environments. Wi-Fi 7 builds on the security benefits of Wi-Fi 6/6E, making it a viable choice for untethered connectivity. The increased throughput/speed, lower latency, reliability, and solid security in WPA3 makes it easier for operators to achieve their 10G platform pillars.

7.1. MDUs

In MDU environments, where there are many APs operating close to one another, it is important to identify areas of interference, like co-channel interference (CCI) or adjacent channel interference (ACI). One thing that is not changing with Wi-Fi 7 is, as with previous Wi-Fi versions, performing a Wi-Fi site survey at the MDU is important to understand the operating environment. Most Wi-Fi site survey tools offer a Wi-Fi heatmap to represent the coverage and RF signal strength visually. These heatmaps are often overlaid with MDU floor plans, increasing their effectiveness as a deployment and troubleshooting tool. This overlay also provides operators a visual map of the location of trouble zones, the location of APs and other relevant survey data.



Figure 15 – MDU

Besides the high areas of interference, the MDU is a dense environment. Wi-Fi 7 has access to an increased amount of RF spectrum and tools to use this RF spectrum better. Preamble puncturing is one of those features mentioned, along with Co-OFDMA. These features greatly benefit those using Wi-Fi to

communicate in MDU environments. Wi-Fi 7, like Wi-Fi 6E, takes advantage of the 6 GHz spectrum and may leverage MLO to increase throughput/speed and reliability in MDUs.

As mentioned, adjusting AP transmit power can provide big benefits in high-density deployments, reducing OBSS interference. Environments such as MDUs where multiple APs are using the same RF channel and transmitting unique SSIDs create the potential for a high amount of Wi-Fi interference. Furthermore, MDUs will benefit from the low complexity AP and advanced AP coordination features of Wi-Fi 7.

7.2. Business Services

Business services are transforming their enterprise networks with a host of digital elements that require higher data rates than residential. Wi-Fi 7 has a roadmap of throughput/speed options that extend the certification out to 40 Gbps. Today, service offerings like private/public/hybrid cloud, software-defined networking (SDN), and virtualization will benefit from the added Wi-Fi 7 throughput/speed options. Technologies like the 3rd Generation Partnership Project's (3GPP) 5G, citizens broadband radio service (CBRS), along with Wi-Fi 7, will be used to transform the access network edge for enterprise services using SCTE's generic access platform (GAP) enclosure (ANSI/SCTE 273-1 2021) and module (ANSI/SCTE 273-2 2021) specification standards. Further transformation of the edge will occur with the reality of 3GPP's 6G and Wi-Fi 7 Release 2.

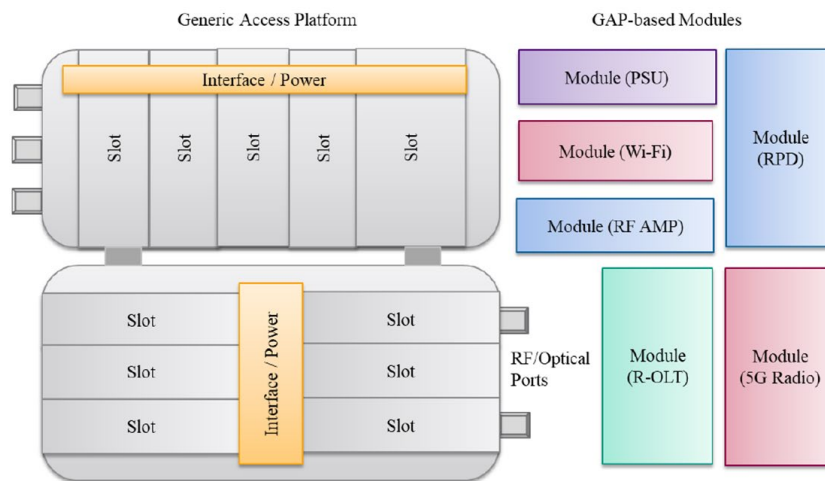


Figure 16 – SCTE GAP

For improving reliability in the enterprise, Wi-Fi 7 offers a toolkit here as well. Business enterprise time-sensitive applications require latency that is deterministic, offering high reliability and quality of service (QoS). Deterministic latency offers predictable jitter and network delays for enterprise networks. Operator enterprise Wi-Fi 7 networks must be designed to cope with the compromise between throughput/speed and deterministic latency. Wi-Fi 7's latency features and Co-OFDMA will benefit the explosion of IoT and industrial (IIoT) smart devices making their way into the enterprise.

The enterprise features mentioned will improve reliability, a key performance indicator (KPI) on the health of a Wi-Fi network.

8. Wi-Fi Performance Management

Subscribers want a consistent wireless/Wi-Fi experience, making performance management an important topic for Wi-Fi 7. A few of the top issues for operators with in-home or business service Wi-Fi are throughput/speed, range, interference, congestion, and compatibility.

In any Wi-Fi network, a high number of users or STAs may overload an AP, reducing the available RF energy in a BSS. In addition, as the STA distance from an AP's radios increases, the receive signal strength indicator (RSSI) value decreases towards the noise floor of -90 dBm. This is because extending the range from the STA to the AP reduces the power due to free space path loss (FSPL). Each 3 dB of signal loss in the RSSI equates to a reduction in milliwatt (mW) power by one-half (1/2). A reduction in RSSI also reduces the data throughput/speed of the AP, like how cellular devices operate with a base station. Leveraging the benefits of Wi-Fi 7 MU-MIMO, preamble puncturing, MLO, and AP coordination will increase performance in these scenarios. An additional consideration that might benefit performance here would be load balancing STAs, additional APs, and decreasing STA to radio distance.

In high multi-path environments, interference and signal fading are always top of mind and an operator's priority. Multi-path environments, or indirect line of sight (LoS) communication, occur from reflected, scattered, and other RF extrinsic factors in a Wi-Fi environment. Wi-Fi 7 offers 16x16 MU-MIMO and spatial diversity to reduce interference and signal fading in high multi-path environments. Spatial diversity allows an AP to select from multiple input signals for the best reception. Furthermore, to identify interference, always perform a Wi-Fi site survey to understand the operating environment and talk with the subscriber about the Wi-Fi coverage in the BSS. A site survey is a great step to providing a good Wi-Fi service in residential, MDU, and business service enterprise environments. A survey identifies items that may reduce RF coverage, range, or performance in the Wi-Fi network. The survey tool offers data to discuss options with subscribers.

Not all Wi-Fi site survey tools are the same. Most survey tools offer a heat map characterization, as well as a Wi-Fi extender and mesh characterization. The heat map produces images utilizing a color map to show the AP's RSSI and RF throughput levels. With most Wi-Fi networks deploying extenders to create mesh networks for range and coverage improvements, we must characterize their range and throughput/speed. Heatmaps will show ways to improve range, coverage, and throughput/speed for larger dwellings. Heatmaps also look for CCI and ACI, as well as an elevated noise floor.




Figure 17 – Heatmap

WI-FI networks operate between -30 dBm and -90 dBm, where -90 dBm is the noise floor. The higher the RSSI, the better the performance of the Wi-Fi network. For example, a -80 dBm is considered poor, while a -60 dBm may be considered acceptable. A dB is a decibel or relative power measurement used to compare two dBm values like the SNR metric. SNR refers to the Wi-Fi signal level of the received signal versus the ambient noise on the RF channel, meaning measurable RF energy without a signal, referred to as noise. For example, Figure 18 below shows a plot of a Wi-Fi signal and the noise floor. The noise floor is a little higher than normal at -79 dBm, and the signal is around -41 dBm, which makes the SNR for this connection 38 dB. Depending on the Wi-Fi design, operators may want an S/N of 35 dB or greater.




Figure 18 – Signal to Noise Ratio

The dBm is a decibel relative to a milliwatt (mW), where 0 dBm equates to 1 mW of RF power. The mW is an absolute power measurement often used in combination with dBm in Wi-Fi networks. RSSI and S/N will still apply to Wi-Fi 7 networks and will vary based on the radios used, 2.4 GHz, 5 GHz, and 6 GHz.

<div> <div>Wi-Fi</div> <div>6 / 28 Radios</div> <div>17 Scan</div> <div>01:54 Length</div> <div>SCTE-5G 9C:C9:EB:99:96:E1</div> <div>CH 149 (149-161) / 80 MHz 0.8 μs / 5 GHz / 5745 MHz</div> <div>802.11 ax d e h i w / 6. Gen</div> <div>MCS Rates 36 / 1201 Mbps</div> <div>102.4 ms US AES / AES / SAE</div> </div>									
Network Name	BSSID	RSSI	SNR	CH	802.11	Width	SS		
CableLabs	9C:C9:EB:99:96:E2	-37 dBm	58 dB	149	ax	80 MHz	2		
CableLabs	9C:C9:EB:99:96:C2	-29 dBm	72 dB	6	ax	20 MHz	2		
Kyrio	9C:C9:EB:99:96:E3	-37 dBm	58 dB	149	ax	80 MHz	2		
SCTE	9C:C9:EB:99:96:C1	-29 dBm	72 dB	6	ax	20 MHz	2		
SCTE-5G	9C:C9:EB:99:96:E1	-37 dBm	58 dB	149	ax	80 MHz	2		
SCTE-Guest	9C:C9:EB:99:96:E4	-37 dBm	58 dB	149	ax	80 MHz	2		

Figure 19 – RSSI

As with Wi-Fi 5/6, it is operating in the 5 GHz as opposed to 2.4 GHz reduces the range of a Wi-Fi network. Then we have newer STAs that will take advantage of 6 GHz from Wi-Fi 6E and Wi-Fi 7,

further reducing the range of the RF signal and the RSSI value. This is because the wavelength of 6 GHz is much shorter than 2.4 GHz. Mesh networking will be an important component of the deployment of these Wi-Fi 7 networks.




Figure 20 – Wavelength of 2.4, 5, and 6 GHz

Other tests may be conducted to determine airtime fairness implementation in the AP. These tests look into the throughput/speed of the AP and how it allocates bandwidth to each of the STAs. These types of tests help operators understand scenarios where a good QoE or QoS may not be delivered. A different kind of test that is useful is determining the band steering capabilities of an AP. Band steering will become more important as Wi-Fi 6E/Wi-Fi 7 utilize multiple radios, especially when multi-link capacity is available. Having the AP and STA choose the best bands to operate in will be crucial for providing the best possible QoS.

9. Conclusion

Residential, MDU, and business service wireless communications will support the mission of the 10G platform, extending the benefits of throughput/speed, low latency, reliability, and security. In addition, Wi-Fi 7 will offer a more carrier-grade version of wireless while driving a more positive customer QoE. In all the access network technologies an operator deploys, it is the premises' Wi-Fi experience that will determine the image of an operator. The Release 1 features like PHY enhancements, OFDMA, MLO, and preamble puncturing will be critical to the next generation of connectivity. Release 2 features will solidify the experience with MIMO enhancements, channel sounding, HARQ, low latency, and AP coordination.

Managing the impact of Wi-Fi 7 will require benchmarking devices with partners like Kyrio and truly understanding the metrics that lead to a healthy Wi-Fi ecosystem. With any new Wi-Fi certification, workforce education and credentialing will be the tools to manage our talent teams, further pushing positive interactions and experiences with Wi-Fi 7. Any such educational program (e.g., SCTE BWS and SCTE CWNP) must measure the return on investment (ROI) to quantify the reduction of call volumes,

truck rolls, etc. related to wireless technology in our industry. Stay tuned as advancements in Wi-Fi 7 take place, and look to the trusted applied science leader in telecommunication to keep you ahead of the curve on this exciting technology.

Abbreviations

3GPP	3rd Generation Partnership Project
ACI	adjacent channel interference
AI	artificial intelligence
AP	access point
AR	augmented reality
bps	bits per second
BFR	beamforming report
BSS	basic service set
CBF	coordinated beamforming
CBRS	citizens broadband radio service
CC	chase combining
CCI	co-channel interference
CIA	confidentiality, integrity, and availability
CPE	customer premises equipment
CSI	channel state information
CSR	coordinated spatial reuse
dB	decibel
dBm	decibel relative to one milliwatt
DL	downlink
EDCA	enhanced distributed channel access
EHT	extremely high throughput
FSPL	free space path loss
GAP	generic access platform
HARQ	hybrid automatic repeat requests
IEEE	Institute of Electronics Engineers
IIoT	Industrial Internet of things
IoT	Internet of things
IR	incremental redundancy
KPI	key performance indicator
L-LTF	legacy long training field
L-SIG	legacy signal field
L-STF	legacy short training field
L-MAC	lower MAC
LoS	line of sight
MAC	media access control
MCS	modulation code scheme
MDU	multiple dwelling unit
MFP	management frame protection
MIMO	multiple in multiple out
MLD	multi-link device

MLE	multi-link element
MLO	multi-link operation
MRU	multi-resource unit
ms	milliseconds
MU	multi-user
mW	milliwatts
NDP	null data packet
NDPA	null data packet announcement
NSTR	non-simultaneous transmission reception
OBSS	overlapping basic service set
OFDMA	orthogonal frequency division multiple access
PHY	physical
PPDU	physical protocol data unit
QAM	quadrature amplitude modulation
QoE	quality of experience
QoS	quality of service
RL-SIG	repeat legacy signal field
RNR	reduced neighbor report
ROI	return on investment
RSSI	receive signal strength indicator
RTA	real-time applications
RU	resource unit
SAE	simultaneous authentication of equals
SCTE	Society of Cable Telecommunications Engineers
SDN	software-defined networking
SIC	successive interference coordination
SNR	signal to noise ratio
SR	spatial reuse
SSID	service set identifier
STA	station
STR	simultaneous transmission reception
TIM	traffic indication map
TSN	time-sensitive networking
TWT	target wake time
TXOP	transmit opportunity
U-MAC	upper MAC
UL	uplink
U-SIG	universal signal
VR	virtual reality
WLAN	wireless local area network
WPA	Wi-Fi protected access

Bibliography & References

Current Status and Directions of IEEE 802.11be, the Future Wi-Fi 7, E. Khorov, I. Levitsky and I. F. Akyildiz, *IEEE Access*, vol. 8 (2020)

IEEE 802.11be: Wi-Fi 7 Strikes Back, Garcia-Rodriguez, Adrian & Lopez-Perez, David & Galati Giordano, Lorenzo & Geraci, Giovanni,.. *IEEE Communications Magazine*. 59 (2021)

802.11be and other IEEE 802.11 Working Group related documents: <http://www.ieee802.org/11/>

IEEE 802.11be Multi-Link Operation: When the Best Could Be to Use Only a Single Interface, Á. López-Raventós and B. Bellalta, 2021 19th Mediterranean Communication and Computer Networking Conference (MedComNet) (2021)

IEEE 802.11be Wi-Fi 7: New Challenges and Opportunities, C. Deng et al., *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4 (2020)

The Operational Impacts of Supporting a Disaggregated, Distributed, Cloud-based Network Architecture

A Technical Paper prepared for SCTE by

Aliraza Bhimani

Principal Network Engineer/Team Lead
Comcast Cable
New Jersey
609 929-3346
aliraza_bhimani@comcast.com

Idris Jafarov

Delivery Team Leader
DriveNets
New Jersey
302-559-5952
ijafarov@drivenets.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Disaggregation as a success criteria for ISPs/CSPs	5
2.1. DDC and DDBR	5
2.2. Difference between DDC/DDBR and traditional routers	6
2.2.1. Modular Chassis	6
2.2.2. Disaggregated Routing System	8
2.3. Processes and skills needed.....	10
2.3.1. Skills	10
2.3.2. Processes	14
3. Operational Considerations and Challenges	15
3.1. Challenges with traditional router architectures	16
3.1.1. Supply chain concerns	16
3.1.2. Traditional nature of the components	16
3.1.3. Chassis limitations	18
3.1.4. Openness and ease of upgrades.....	19
3.2. Comparison of DDC/DDBR and Spine-Leaf architectures	20
4. Orchestration, Automation and Analytics.....	23
4.1. Automated Operations	23
4.2. Health Monitoring and Assurance	24
4.3. Telemetry	24
5. Real-life deployment of the DDC/DDBR model	25
6. The future of the disaggregated solutions.....	29
7. Conclusion.....	30
Abbreviations	30
Bibliography & References.....	33

List of Figures

Title	Page Number
Figure 1 - DDC/DDBR placements in the network.....	6
Figure 2 - Front and Back View of a Modular Chassis [5]	7
Figure 3 - High-Level View of a DDC Routing System	8
Figure 4 - 3-stage non-blocking Clos topology	9
Figure 5 - Traditional versus Microservices Architecture	12
Figure 6 - High-level view of Linux Containers	13
Figure 7 - Overview of YANG and NETCONF	14
Figure 8 - Traditional Routers	16
Figure 9 - Non-linear growth of ports [4]	18
Figure 10 - Folded 3-stage Clos network (also known as Spine-Leaf).....	20
Figure 11 - Large National Backbone Network	26
Figure 12 - Spine-Leaf cluster within a Backbone Site	27
Figure 13 - Virtual Chassis within a Backbone Site	28

List of Tables

Title	Page Number
Table 1 - DDC/DDBR cluster sizes	21

1. Introduction

Managing and maintaining highly scalable networks has historically been a challenging task. A plethora of ISPs/CSPs have been trying to simplify processes and procedures; yet this task gets more complicated as they are faced with the growing cost pressure of supporting today's IP network traffic demands (driven by video, gaming, and remote working) and future 5G/6G cellular traffic volumes. Based on several reports, the internet usage has increased by 1,355% over the last 22 years [1]! With the coming deployment of the Full Duplex DOCSIS 4.0 system in the access networks [2] and the proliferation of the 400/800Gb/s Ethernet technology [3] in the aggregation/core layers, traffic utilization could continue to accelerate in the upcoming decade.

As Mannan Venkatesan, a Distinguished Engineer from Comcast Cable, notes in his NANOG N81 2021 presentation:

“There is a significant amount of port growth we need to support. We all know the internet traffic has exploded within the last decade. The ports you would need to support the traffic are also proportional to the traffic volume, so we must make sure we stay on top of the port capacity we support on core and aggregation routers” [4].

All these trends and market forces are impelling cable operators to rethink and rearchitect their existing IP networks and operations to maximize performance and efficiency.

Network operators must also consider the total cost of ownership of various hardware and software options, not just that of an individual component or software option. Network components don't operate in a vacuum, so choosing feature-light hardware may necessitate greater investment in hardware and vice versa. There is always a tradeoff between feature-rich software and bare bones hardware. It's critical for network operators to understand what's needed for end-to-end delivery, rather than just individual component costs and specs, which may not paint the whole picture.

Due to the above-described forces, a new innovative system is desired to solve these problems. The primary requirements for such a system are that the cost per bit and time to market of new services be reduced, and yet be able to permit the simplicity of the network operations while minimizing the blast radius or failure zone of the network. Losing one router has the potential to impact millions of customers and puts the network in a hazardous condition. Another failure can isolate the whole network, either the one managed by the operator or a peer's network. The main goal is to reduce the blast radius and impact of one router or component and increase the availability of the network.

Disaggregated Distributed Chassis and Disaggregated Distributed Backbone Router (DDC and DDBR) are powerful solutions based on open-source specifications that solve current operational challenges when building and scaling IP backbone/aggregation networks [5, 6]. These two options addresses and mitigates many of the issues described above. These models arm operators with the flexibility of selecting the best breed of IP products in the market.

In this paper we dive into how cable operators can leverage open transport building blocks across different segments of their transport networks (access, aggregation, backbone) and implement concepts of real DDC/DDBR architecture, while utilizing orchestration, automation, and analytics. We will also touch upon how they can overcome operational challenges and pitfall considerations while proactively positioning themselves for future disaggregated network solutions.

2. Disaggregation as a success criteria for ISPs/CSPs

2.1. DDC and DDR

Disaggregated Distributed Chassis and Disaggregated Distributed Backbone Router are open-source specifications for carrier grade routing systems put forth by Open Compute Project (OCP) and Telecom Infra Project (TIP), respectively, which are global collaborative communities focused on innovation and development of open, disaggregated, and standards-based technology solutions [7, 8]. One of the main motives for the inception of these organizations was the huge influx of new registered users and data which resulted in an exponential growth of services and platforms developed and deployed by hyperscalers. These developments incurred unforeseen control costs and energy consumption which can only be optimized by the benefits of open source and open collaboration to hardware. This collaboration model is now being applied to advance the telecommunications industry as ISPs/CSPs experience similar pressures of unprecedented traffic growth.

Telcos and Cablecos are rigorously researching, testing, and deploying disaggregated networking solutions across their footprint. To cope with the tremendous demand, they have been envisaging an evolution path to their core/aggregation networks to introduce innovation, efficiency and mostly openness. Ideas have been entertained to initiate a shift in the IP backbone architecture from the traditional single chassis-based routing systems to more disaggregated distributed ones.

ISPs/CSPs not only desire to meet current needs when deploying core/aggregation transport networks but also staying ahead of the evolving trends in terms of resiliency, capacity scaling & End-To-End (E2E) network automation.

Before going into the nitty gritty details of the challenges that operators face, let's review the key points of the DDC/DDBR solution [5, 6]:

- **Disaggregation is driving competition:** creating opportunities for new players in the market driving costs down.
- **Pay as you grow:** model that allows operators to purchase capacity incrementally as it is needed.
- **Innovation:** open software and hardware to improve flexibility and innovation while reducing time to market.
- **Operational Efficiency:** Taking advantage of centralized control and monitoring tools.
- **Reliability:** Always targeting higher availability & multi-level redundancy while minimizing blast radius impact to decrease customer impact and outages.

This system can be placed not only in the IP/MPLS backbone but also in the access layer and act as an Internet Gateway Router (IGW), as depicted in Figure 1. The same hardware can be utilized for these routing functions regardless of Network Operating System features or implementation.

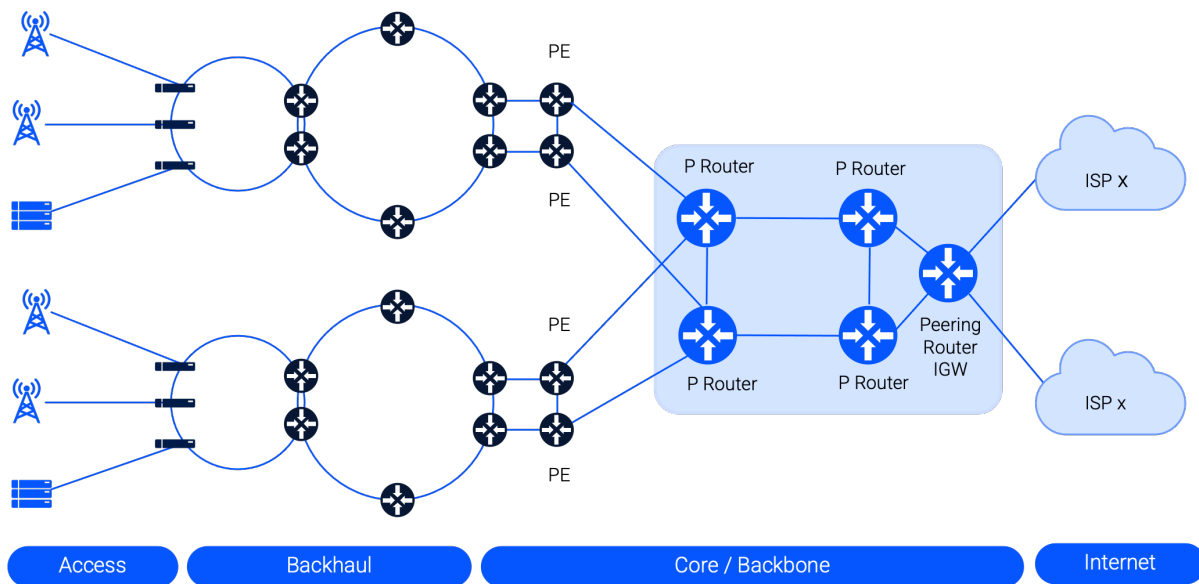


Figure 1 - DDC/DDBR placements in the network

2.2. Difference between DDC/DDBR and traditional routers

2.2.1. Modular Chassis

To understand why DDC/DDBR is an evolutionary architecture, a reasonable convenient starting point would be to analyze the traditional routing platform architecture first.

Figure 2 shows a front and back view of a finished product view of a typical modular chassis routing system from one of the vendors in the ecosystem.




Figure 2 - Front and Back View of a Modular Chassis [5]

There are multiple components that comprise this system which reside inside the chassis and are not visible: line cards, fabric modules, Route Processor modules, PSUs, and FAN FRUs to name a few. Modular chassis routers use pluggable line cards to scale up capacity and service features and form the foundation of ISPs/CSPs' backbone and aggregation networks today. If a service provider needs more ports, more linecards can be bought and inserted into the system. The growth ceiling for this type of system is dependent on how many available slots there are on a modular chassis. One more option to scale up is to purchase a totally new chassis with higher port density and replace the old one due to the scarcity of slots for line cards. An example would be replacing a ten-slot chassis with a twenty-slot chassis.

Edson Erwin invented the highly scalable Clos architecture/Clos Network in 1938 and which then was formalized by Charles Clos in 1952. Later in 1953, Charles Clos published a paper titled "A Study of Non-blocking Switching Networks" in the Bell System Technical Journal [21] where he describes a method of designing arrays of cross points for use in telephone switching systems. In recent years, the variation of a Clos network had been widely deployed by hyperscalers and is now being adopted by the largest telcos and cablecos.

Another possible course of action is the scale-out (horizontal disaggregation) model. This model decomposes the chassis system into a spine and leaf Clos architecture which has been pioneered and successfully adopted in data center designs by hyperscalers. ISPs/CSPs are now making a significant effort to take these concepts and apply them to the routed Wide Area Network (WAN).

As can be seen, these traditional routers mainly scale up hardware, with some scale out options existing such as multi-chassis racks or back-to-back multi-router options. It is viable to architect incremental growth, but then there is a need to account for the same incremental up-front cost for cooling, power, space along with rack, stack, and cabling installation costs. The power per rack may be limited at the facilities and new next-generation single chassis systems have very high-power demands, more cooling is needed and there may be space limitations at the headend or datacenter. This traditional option is not portable.

At one time, this routing system (RS) architecture was an effective way to build backbone/aggregation networks. It was a one-stop shop to have all your port, fabric, and software needs in one fixed chassis. This made troubleshooting easier as well since the networking vendor would be responsible for any software or hardware issues encompassing the whole chassis. The limitations of this system lie in its mechanical design to solve thermal and spatial challenges and live Online Insertion and Removal (OIR) of modules within the system which sometimes can cause traffic impact and may require multiple maintenance windows to ensure business continuity. On top of that, the mentioned hardware design and the software controlling the system are solely proprietary to the specific vendor thus preventing potential ecosystem players to participate and compete in either HW or SW market.

However, with the evolution of networks and traffic patterns, ISPs/CSPs had to adapt and adopt disruptive and innovative technology solutions. DDC/DDBR RS architecture that resembles Clos Spine-Leaf (S/L) topology in terms of interconnections between elements but operates as a single routing system is the optimal way for operators to accelerate business agility and drive revenue growth. This solution also allows per port or speed step function which we will discuss later.

2.2.2. Disaggregated Routing System

So, you may wonder, what is a Distributed Disaggregated Routing System? All the elements that were shown in the Figure 2 traditional refrigerator single chassis router; namely the line cards, fabric modules and route processors are all disaggregated onto separate physical “pizza” whitebox and commercial off-the-shelf (COTS) x86 server components that run virtual machines and software containers. This variation of Spine/Leaf Clos design allows various scalability options. Clos has been already deployed previously in datacenters, backbone/core networks for many decades as we will learn more about later. Now, we have the ability of taking it one step further to the regional area networks or closer to the edge.

In Figure 3 we can see a conceptual diagram that demonstrates such a system from a birds-eye View [5]:




Figure 3 - High-Level View of a DDC Routing System

In this graphic, the disaggregated distributed RS is comprised of a certain number of fixed-RU Packet Forwarders (PF) which represent the leaf of the cluster and act as the line cards, and several Fabric Forwarders (FF) which represent the spine of the cluster and act as the backplane. The “brains” of the system are running on redundant on-premises COTS x86 servers or as containers in a cloud-native fashion which execute all the sophisticated route computation algorithms and performs management of the whole cluster. The external ethernet switches provide distributed communications channels between all 3 described components. This routing system can be implemented across multiple physical racks with proximity at a headend or data center, thus making it portable and having the flexibility to overcome any potential space or rack issues.

Packet forwarders perform strictly forwarding functions and do not store protocol state information which allow them to support line-rate forwarding across all ports without any limitation. Fabric forwarders are responsible for interconnecting all PFs in a full mesh and transporting cells between them. The interconnections between the PFs and FFs are not IP based but rather leverage proprietary cell-based technology from multiple merchant silicon manufacturers to achieve optimal redistribution of traffic across fabric links.

Figure 4 demonstrates an example of how PFs can be interconnected to FFs in 3-stage non-blocking Clos topology. The S/L based architecture has helped the web-scale companies to efficiently grow their infrastructure to a massive scale that can deal with many technical challenges they have faced before. There is an ability to scale vertically by adding more FFs or spines to increase the fabric throughput and redundancy, while also allowing to scale horizontally by adding more PFs or linecards to accommodate port demands and future growth.

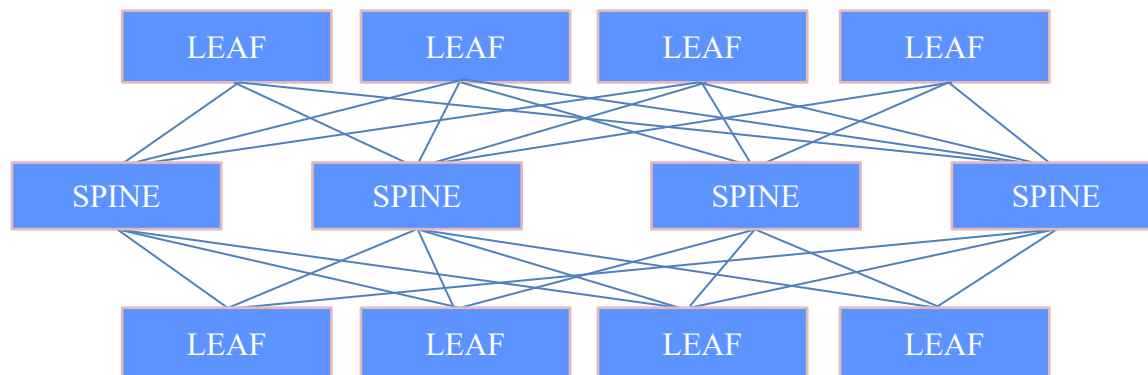


Figure 4 - 3-stage non-blocking Clos topology

However, the Spine & Leaf architecture by itself has multiple drawbacks when trying to apply it to the carrier grade networks. We will provide a detailed comparison of these 2 architectural approaches in an upcoming section.

The combination of the Spine & Leaf based architecture and disaggregation can lead to phenomenal advantages such as disjointed innovation paths between data plane and control plane. As control plane is completely decoupled from the data plane, the innovation of these 2 components can happen independently. For example, with the evolution of the ecosystem ISPs/CSPs can easily change Routing Network Operating System vendors according to their current needs while using the same hardware or even combining hardware from different Original Design Manufacturers (ODM). In recent years, merchant silicon manufacturers have developed products that are on par with custom silicon developed by Original Equipment Manufacturers (OEM). Disaggregation enables operators to purchase simplified hardware based on merchant silicon resulting in significant cost savings as compared to traditional modular chassis routers. Now, ISP/CSP's do not have to be bound to a specific router vendor's hardware/software roadmap and can freely mix and match different whitebox server hardware with custom vendor software applications/containers.

Multiple modern network operating systems (NOSs) are based on Linux, however, most of them use a custom-built command-line interface which means they don't have a look and feel of Linux. However, since operating system components are running as services in containers, engineers should learn how to deploy container management software to create, deploy, and scale and them.

This combination also promotes more open-source configurations and concepts where different API's can "hook" into this custom routing platform. Some NOS's even make the whole disaggregated cluster made up of different components seem like a "virtual" chassis, providing the best of both worlds [\[9\]](#).

These options allow potential for seamless integration into the existing network while maintaining small failure domains or a blast radius. While N+1 redundancy gives you an extra component, node, or link to failover to, N+M design provides a whole node or cluster redundancy giving a whole other level failover capability. This allows hitless maintenances and updates. Now, it is possible to direct traffic away from a segment of spines and leaves and upgrade the software or configurations while having traffic run across the rest of the S/L cluster. This reduces the failover impact of traffic and does not put such a hazardous condition strain on the system, like it used to with a single chassis system. This S/L design makes an In-Service Software Upgrade (ISSU) more reliable than in single chassis routers. This allows any time maintenance, even daytime since it will be Non-Service Affecting. Now, a single failure of a component doesn't impact traffic in multiple directions and reduces the impact to services.

2.3. Processes and skills needed

2.3.1. Skills

It is no secret that to keep up with innovation, engineers and technologists must set out for a lifelong learning and development journey. Having the right skills has always been important in IT, especially if you want to stand out in the industry. This is certainly going to become even more critical as the pace of change and innovation continues to accelerate. Technical experts will come to find that the technology they have come to know on such a deep level is constantly evolving. Carrier-grade networks are highly dependent on the expertise of network professionals

who build, operate, and maintain them. The skillset of these engineers had been evolving for a while and now with the proliferation of new innovative disaggregated network solutions there is a new set of skills and capabilities which they must master. They now need more DevOps, programmability, and scripting skills to leverage automation tools because the newer open software is more API/GNMI/Model driven. There is a trend and requirement by management to move away from traditional CLI network device management, where you have the chance of repetitive human error. Based on the “Annual outage analysis 2021” report by Uptime Institute an aggregated year-on-year average of 63% of failures are due to human error [\[32\]](#). Network engineers now need to blend their skill sets and learn and utilize server sysadmin skills to make them future-prone, ready, and stay marketable in the job market.

2.3.1.1. *Microservices*

Traditional architecture has been a classic software design pattern since the origins of the industry in which the user interface and the source code are combined into a single program. Though in the past this was convenient to have the networking vendor be responsible for any software issues, there are many drawbacks to this approach such as complexity of the code, scale limitations, reliability among many others. A software code bug may affect both the SW and HW, but when decoupling the two and disaggregating them, there is now a clear demarcation point for the hardware components versus the software. The industry is moving towards the microservices architectural style for developing applications. Figure 5 provides an illustrative example of these two approaches.

Microservices allow a relatively large application to be divided into smaller parts, having their own autonomy and realm of responsibility. With microservices in containers, it's simpler to take advantage of hardware and easily orchestrate services, including networking. Network engineers and architects should learn how to operate and maintain the container orchestration platforms like Kubernetes and Docker because many NOSs designed for disaggregated networks are based on the cloud-native, distributed architecture. Cloud Native Computing Foundation (CNCF), which is the vendor-neutral hub of cloud native computing, states that microservices and containers form the foundation for the cloud-native application development [\[10\]](#). Therefore, it is paramount for network professionals to develop their Linux skills and System/Server Admin skills so they can easily operate on both sides of the domain.

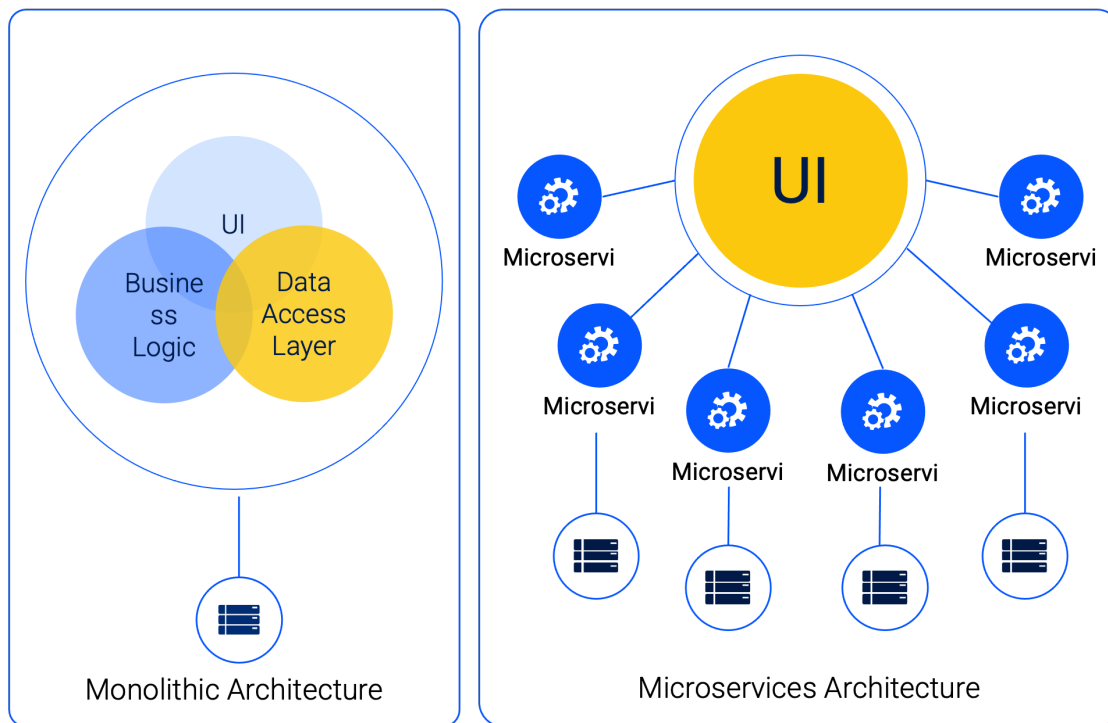


Figure 5 - Traditional versus Microservices Architecture

2.3.1.2. Linux

Refer to Figure 6 for the high-level view of Linux containers. Open organizations like OCP are actively developing and deploying new Linux distributions that are specifically targeted and tailored for network equipment [\[11\]](#). Many vendors in the ecosystem have already started offering NOSs based on these distributions. Network engineers and architects are expected to understand and implement the fundamentals like navigation in the filesystem from the shell, which is the server CLI interface, manipulation of files and directories, running programs and working with background services, also known as daemons. On top of that, engineers must learn how to manipulate network interfaces, how to view and manage routing on a Linux system because interface and routing configuration go together. Also, due to the openness of the hardware, a huge number of logs and event traces are available which can be viewed “under the hood” via the Linux shell. To be able to troubleshoot and analyze the sequence of events or even take the packet captures of control and data plane traffic, an engineer must be familiar with text editors which offer the tools to sift through large amounts of data and even observe events in real-time. Additionally, many of the tools that we will discuss in the future section have their origins in Linux and require to be run from the Linux system.

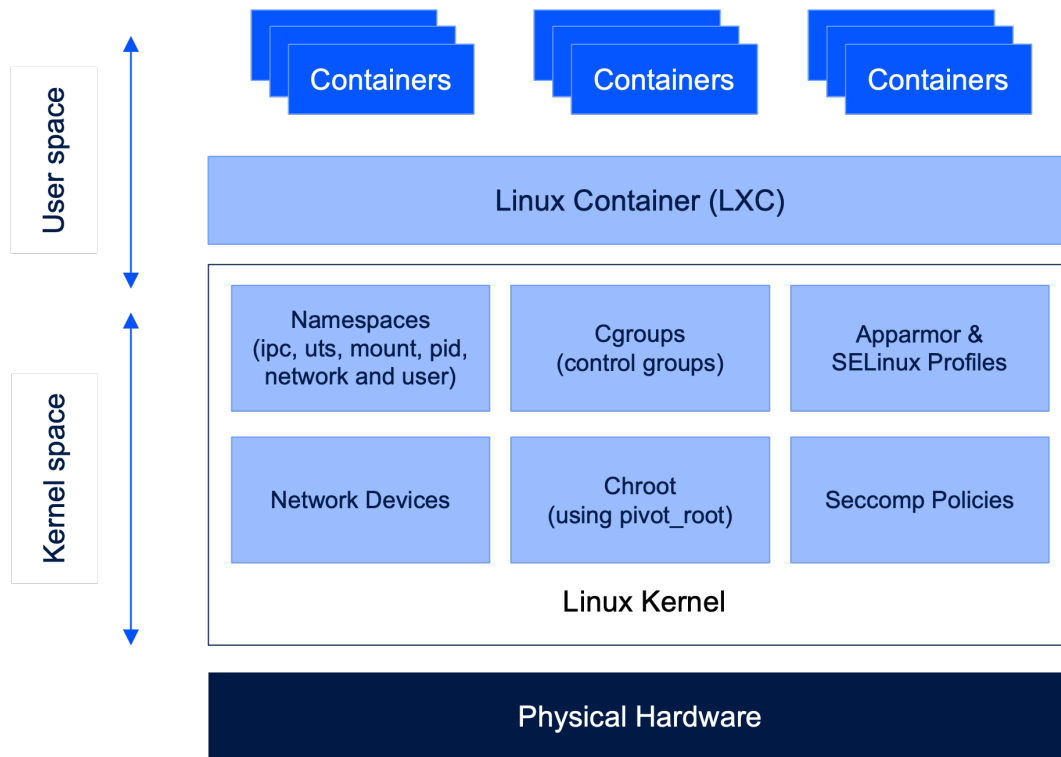


Figure 6 - High-level view of Linux Containers

2.3.1.3. Data Formats and Models

Protocols like BGP, OSPF, ISIS, and TCP/IP were created out of need for devices on the network to have a single language to communicate across a globally distributed system – the Internet! The data formats were conceived for a similar reason – for computer systems to freely understand each other.

With the significant increase of the quantity of network devices and the proliferation of the Internet of Things (IoT) it became apparent that it is beyond the bounds of possibility for humans to manage them [\[12\]](#). The solution here is to automate configuration management and maintenance as much as possible and streamline the deployment of new services. Engineers should get familiar with network automation tools and methodologies as many modern NOSs expose programmatic interfaces that offer an API. Traditionally, each network device is closed (locked from installing third-party software, as an example) and only has a command-line interface (CLI). Although the CLI is still a well-known and even preferred method of access to the router by network professionals, it does not offer the flexibility required to truly manage and operate this brave new Internet.

Imagine operating a disaggregated cluster comprised of fifty or so components or nodes. Each node may need its own Loopback IP for management and its own configuration file. In the past with traditional routers, the network engineer had to create the fifty separate config files and load them onto each component in the cluster. This also opens the system up to human error and the configuration files can be prone to typos and errors. Pushing configurations and updates to

multiple nodes in a S/L system might result in unexpected errors, takes a lot of time to prepare configurations or writing scripts to automate the process. One way to easily automate it is to set up a central system and network platform or orchestrator. This system can now treat all the nodes like a virtual cluster, where the Orchestrator pushes the software and configurations to all the components of the cluster. This buys the best of both worlds where there are benefits of a traditional system with the flexibility of S/L. Thusly, automation becomes key and “table-stakes” to implement any type of disaggregation.

One of the most optimal ways to push and retrieve the configuration from the devices is via the HTTP-based APIs like RESTful and non-RESTful APIs [13]. Another one is NETCONF [14] which is a network management protocol conceived specifically for configuration management and retrieving operational state data. In order to leverage these tools, professionals need to understand the data formats like XML and JSON [15, 16], YAML [17] and data modeling language YANG [18]. Refer to the Figure 7 for the overview of these tools.

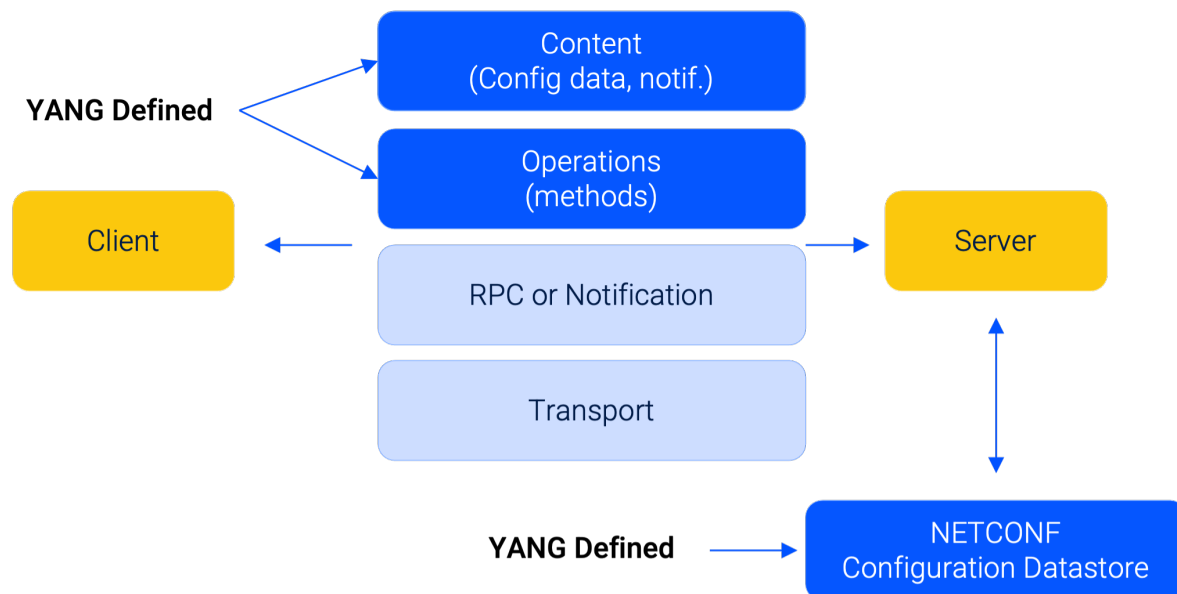


Figure 7 - Overview of YANG and NETCONF

2.3.2. Processes

The processes to deploy and maintain the disaggregated solutions are simplified due to multiple factors like less planning efforts and faster installation time.

2.3.2.1. Less planning efforts

IP network planning can be quite complicated and take months. The planning time is heavily influenced and considers the number of components, service requirements, capacity forecasting, site constraints, and more. While traditional chassis solutions are limited in their abilities to simplify the overall process, disaggregated distributed solutions offer one design process and ease of expansion by relying on only two standard data plane building blocks: Packet Forwarders and Fabric Forwarders, as compared to five to ten integrated router solutions with multiple line

card models which must be compatible with fabric modules and different software revisions for different features on a line card.

2.3.2.2. *Installation Time Considerations*

In this high-paced, modern era of networking, field installations and maintenance are some of the most comprehensive and costly operational activities. They include the engagement of multiple teams, materials, and dedicated planning to adapt to each physical location's constraints like space, power, and cooling. In the disaggregated model, installation process is identical across all node sizes, eliminating the need to constantly train staff and avoid costly installation errors. This cookie-cutter method of installation can be used for the deployments across all fields of use in the network (peering exchanges, cloud exchanges, edge cloud etc.). However, there is a lot of backend fabric fiber and ethernet connections (fondly referred to as "spaghetti-wiring") since it looks like strands of spaghetti that must be installed to connect the nodes as a cluster as opposed to simply sliding in a linecard in traditional chassis case where only the service ports must be connected as the fabric is hidden in the backplane.

There is a lot of up-front cost and effort in a DDC/DDBR architecture in performing the physical install work initially to obtain a longer roadmap, but the benefit is savings in port migration to the new chassis. When forecasting growth and installing with the mindset of room to grow, this eliminates the need to come back and install any additional fabric forwarders and disrupt the system later. Also, unlike scale up in a single chassis router, when the system grows beyond its size, all the client ports ought to be replumbed to another chassis whereas in a disaggregated model scale out consists of only adding another PF or two. This positions the network to consider any un-forecasted growth surges, such as the one a lot of networks experienced during the COVID 19 pandemic. This follows the "set it and forget it" ideals.

3. Operational Considerations and Challenges

DDC/DDBR is an optimal solution that overcomes the most relevant challenges that ISPs/CSPs are facing today when deploying and scaling their IP backbone/aggregation networks. The role of these networks is to route the mobile, voice, broadband & commercial traffic between different network segments at a national and regional level while providing connectivity with external networks such as other service and cloud providers, content data networks, Internet exchange peers and IP transit providers.

The IP backbone networks must regularly scale to support the internet traffic growth, to improve resiliency and reliability, and to meet the expectations of mission-critical types of communications. The essential objective is to lower the cost per bit and improve the overall customer experience and satisfaction.

In the following sections, we will dig deep into the key challenges that exist in the IP backbone networking space, how DDC/DDBR architecture can help in overcoming them, and comprehensive comparison of DDC/DDBR and Spine-Leaf architecture.

3.1. Challenges with traditional router architectures

3.1.1. Supply chain concerns

Akin to all sectors of the telco and cableco networks (Transport, Core, Aggregation, Access), the continuous mergers and acquisitions in the core routers market have resulted in the following limitations, no less important of which is a lack of diversity:

- Notable reliance on a very limited number of suppliers.
- Market which is extremely difficult to enter and compete in leading to an ever-increasing risk from soaring costs.
- Insubstantial innovation and time-to-market speed.
- Unsatisfactory interoperability across different hardware components. For instance, line cards and fabric modules compatibility issues.

3.1.2. Traditional nature of the components

Historically the cablecos have been deploying traditional IP backbone/aggregation routers which are based on proprietary components as demonstrated in Figure-8:

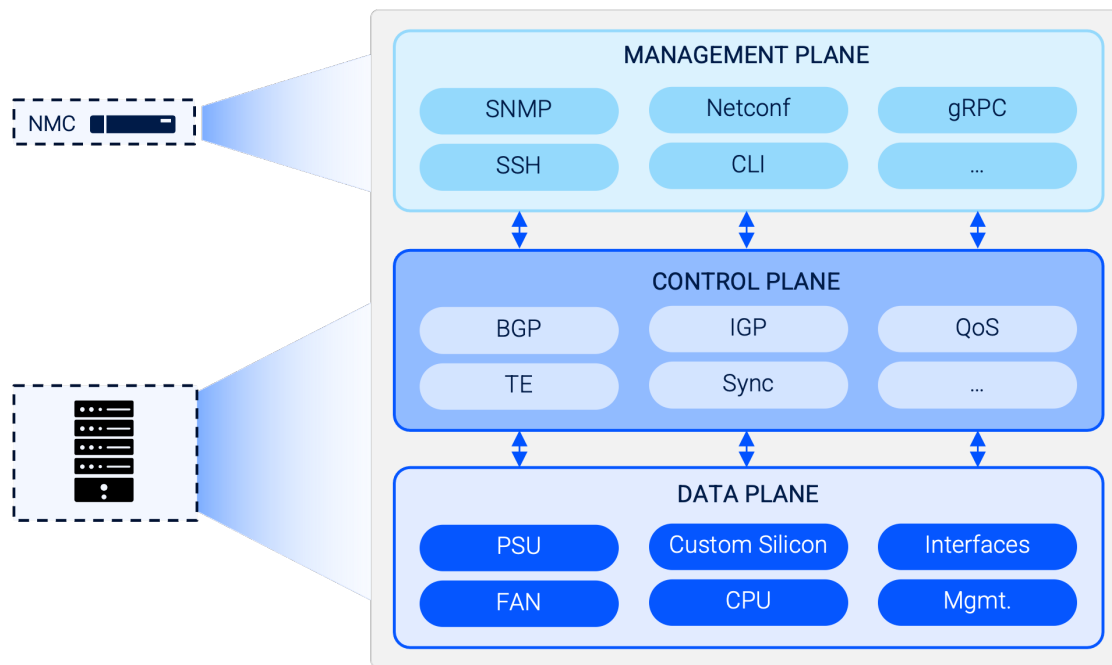


Figure 8 - Traditional Routers

The data plane hardware which contains line-cards is where custom silicon chips are implemented to manipulate the packet processing, traffic management and forwarding. Multiple line cards serve different purposes and have differing applications depending on the place in the network which introduces additional operational complexity and overhead. For example, there

are availability concerns when replacing or upgrading a line card inside the chassis. In the live production environment, an issue with a fabric module or a line card can have an adverse impact on the whole system rather than being limited only to that component consequently increasing the blast radius.

The disaggregated model has a cloud-like pay-as-you grow approach where fixed form-factor PF whitebox servers can be purchased as needed from multiple ODMs for a significantly lower price. Packet forwarders enable operators to utilize any port on the whitebox for any service regardless of the implementation area. This adds tremendous flexibility to enable multiple services and reduce their time to market. The ports on these whiteboxes come with variable native interface speeds which can be reconfigured to operate at different speeds. For example, native 400GE port can be reconfigured to be utilized as 100GE port or can be broken down into 4x100GE via a breakout cable. For the access use case, 100GE can equivalently be adjusted to operate at 10GE speed or, also, broken down into 4x10GE or 4x25GE ports. This allows for gradual growth into the port step functions depending on your deployment budget.

A custom NOS which is specifically designed to efficiently run only on the custom hardware and is comprised of proprietary code and licensing runs the control plane. That includes the drivers which are in control of the hardware components like power, cooling, etc. The firmware is responsible for loading the NOS image when the router boots up including the networking software stack.

In the disaggregated model, multiple NOS vendors in the ecosystem can install their software onto the open networking hardware. The Open Network Install Environment (ONIE) [\[19\]](#) is an open-source initiative that defines an open “install environment”. Before the invention of ONIE, routing equipment was procured with pre-installed NOSs, essentially creating networking devices that locked operators with vendors whose supply chain is integrated and owned completely by them. ONIE runs in the management subsystem of the whiteboxes utilizing capabilities in a Linux kernel. This allows ISPs/CSPs to install target NOS as part of the provisioning process, in the same way the COTS servers are provisioned.

The third component is the management plane, which takes care of the overall platform management: for instance, the interfaces configuration, services provisioning, inventory management, alarm reporting, fault handling, and performance monitoring, all which is tightly coupled with control and data plane; thereby introducing a single point of failure.

This traditional architecture has met the essential needs of operators (capacity, availability, etc.) and has served them for an extended time. Nonetheless, it has impeded them from unleashing the true potential of open networking and has significantly decelerated the innovation in the backbone/aggregation networks.

For instance, having the data and control plane so closely tied together leads to a significant and uncomfortable dependency on the incumbent vendor’s roadmap and prevented them to benefit from features available in a 3rd party NOS supplier. Operators are dependent on the vendor to deliver the HW or SW roadmap on time to upgrade the environment for addressing the growth pressure. The below graph depicts this port growth need versus availability.




Figure 9 - Non-linear growth of ports [4]

Even though multiple software companies have offered powerful Network Management Systems (NMS) and software controllers to manage the traditional routers' management plane, ISPs/CSPs are still exposed to various challenges. One of the major one is the complicated and high integration cost to manage third party products via their existing homegrown NMS. Even though the promising NETCONF protocol is capable to address them, there is a colossal amount of work yet to be done to come up with a vendor-neutral data model for network and device configuration.

3.1.3. Chassis limitations

The current backbone routers are primarily designed based on a chassis structure with front access where the user network interfaces (UNI), network-to-network interfaces (NNI) and the control boards are interconnected into a backplane.

Based on the crucial role of these routers in the network and the huge volume of traffic they pass through, they were required to offer:

- Resiliency to maintain uninterrupted connectivity to the mobile, broadband, and business customers.
- Significant computing and storage capability to store IPv4 and IPv6 global routing tables.
- Port density and capacity enough to support the growth of the customers and services.

Therefore, the chassis had to be equipped with:

- High availability at all levels: control processing, switching fabric, cooling, and power.
- Large TCAM (Ternary Content Addressable Memory), strong computing capability, deep buffers.
- Variety of sophisticated control plane features which includes NSR (Non-Stop-Routing) & ISSU (In-Service Software Upgrade).
- Large port density, number of slots for line cards and backplane switching capacity.

The above requirements have resulted in a platform of considerable size with high cost and huge power consumption, and which necessitates an upfront capital allocation without a guarantee of a reasonable return. This system takes a lot of space and demands a lot of advanced cooling system deployments to ensure adequate operating thermal levels in datacenters and head ends and lacks a capability to adequately grow based on the current capacity needs.

Additionally, due to the limited number of slots in the chassis, running out of ports may turn out to be quite disruptive to an ISPs/CSPs' operational model. This could lead to having to purchase an entirely new chassis resulting in an unnecessarily complex network topology, possible suboptimal traffic flows and a non-linear cost per port model as depicted in Figure-9. This chassis upgrade process is not even remotely agile to empower service providers with a capability to react to unplanned upgrade requests in a timely manner, which results in missing the opportunity to gain more market share by increasing the customer base. This chassis upgrade becomes akin to forklifting the chassis out from the rack and performing “open-heart surgery” to rip-and-replace with an upgraded traditional chassis.

Furthermore, all the NNI & UNI interfaces are centralized on one chassis which creates undesirable operational risks of losing the entire node in the event of a software failure, power issue, executing the Method of Procedure (MoP) document in the wrong order etc. Since telcos and cablecos are dependent on a single chassis in the core of their network, they have the potential to isolate services in case the redundant site goes down increasing the blast radius once again. Even an un-expected fiber cut or an environmental issue that no one has control over can blackhole all services in the network.

3.1.4. Openness and ease of upgrades

According to Yole Développement's “Optical Transceivers for Datacom & Telecom Market 2021” report [\[20\]](#), the optical transceiver market will highly likely grow 14 percent in the next 5 years. This growth is driven by the need of the operators to utilize high data rate modules above 100G.

With the industry trends in the optical pluggable transceivers and the dawn of 400G and 800G QSFP-DD (Quad Small Form Factor Pluggable-Double Density) optics, the ISPs/CSPs need to replace the existing hardware with higher capacity, more compact proportions, elastic thermal management ports which facilitate supporting higher capacity links with advantageous port density per rack unit.

Notably, the 10G/25G/40G user network interfaces need to be upgraded to 100G or even to 400G, and the 100G network to network interfaces or fabric interfaces to 400G and 800G. This means line cards or entire chassis replacement needs to happen to take advantage of the new interfaces and benefit the most from the interface's capacity through the backplane.

To that end, the disaggregated networking approach gives telcos and cablecos an upper hand when replacing or upgrading the installed base to protect their investment in IP backbone/aggregation networks, as the components are based on open hardware. There is no need to use specialized proprietary optics between packet and fabric forwarders because of cell-based packets, keeping with the open-standard theme.

3.2. Comparison of DDC/DDBR and Spine-Leaf architectures

A 3-stage Clos network is the smallest version of a Clos network, and it is relevant to modern scalable carrier and hyperscalers' networks. As its name implies, this network has 3 stages: ingress, middle and egress. Figure-4 from the earlier chapter gives a high-level view of that architecture.

A Spine and Leaf architecture is a derivative from the 3-stage Clos network. Occasionally, it is referred to as a "Folded 3-stage Clos Network", where the ingress and egress points are folded back on top of each other [\[22\]](#) as shown below in Figure-10:

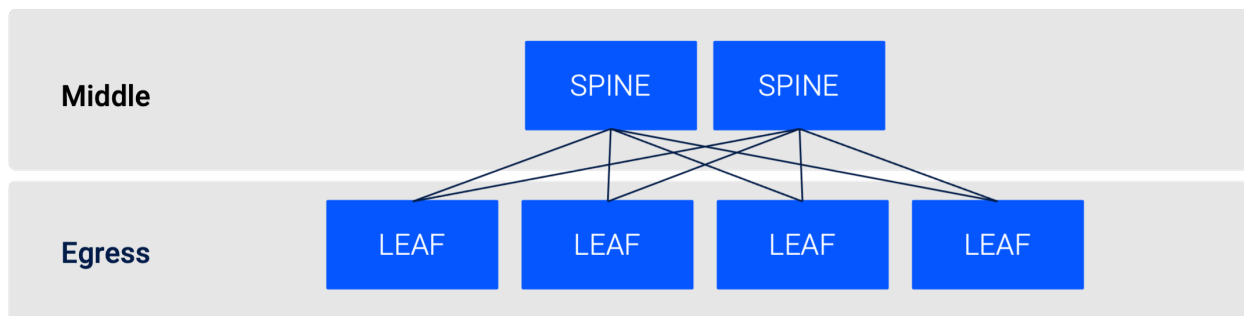


Figure 10 - Folded 3-stage Clos network (also known as Spine-Leaf)

Let's dive into a detailed comparison of DDC/DDBR architecture and Spine/Leaf:

- Over-subscription of Leaf to Spine links in S/L versus equal or over-provisioned fabric to leaf ports bandwidth in the disaggregated solution:

In the S/L architecture oversubscription of links is a common occurrence causing a lot of operational overhead in cases when more bandwidth is needed.

When oversubscription of links ensues (meaning, if there is more traffic on the ingress than the egress on the active link at one time), the procedure for enlarging capacity is complicated. A whole new spine switch or router must be added, and uplinks need to be expanded to every leaf switch thereby adding more cable density within and between the racks. In instances

where device port capacity reaches its limit, a totally new leaf switch/router must be added accruing to the operational complexity introduced by the addition of a spine.

In the disaggregated model on the other hand, clusters are built in a way where bandwidth is equal or over-provisioned to leaf ports [23]. Below is the table with variable cluster sizes ranging from a standalone whitebox with 4Tb capacity up to the large cluster with 192Tb throughput:

Table 1 - DDC/DDBR cluster sizes

	Standalone	Small Cluster	Medium Cluster	Large Cluster
Max Capacity	4Tb	16Tb	96Tb	192Tb
	40x100G	160x100G	960x100G	1920x100G
Port Density	10x400G	40x400G	240x400G	480x400G
	80x10G/25G	320x10G/25G	1920x10G/25G	3840x10G/25G
Packet Forwarder White box	1	4	24	48
Fabric White box	-	2	7	13

- Elephant flows:

In a large Clos network, a typical approach for using the built-in multipath on NNI links is to utilize reliable Layer 3 routing, including on the leaves. Point-to-point links between the spines and leaves are in a specifically allocated subnet and common Equal Cost Multipath (ECMP) routing can be implemented to distribute traffic equally over them. This leads to maintaining large numbers of subnet routes in the tables of the nodes in the topology which results in a requirement to purchase expensive custom hardware with large enough ternary content-addressable memory (TCAM). Although, this architecture is designed using Layer 3 routing, ECMP and provides distinguished bandwidth and latency performance, the preeminent complication here is insensitivity to the workload. Carrier network traffic tends to feature a mix of latency sensitive “mice” flows, and bandwidth intensive longer-lived flows, also known as “elephants”, for which throughput is of higher priority than latency and which constitute only 10% of flows, while accounting for nearly 80% of traffic.

Mice flows are very bursty and short lived, however, elephant flows tend to “pile up” on certain paths even though others are available and ready to be used, filling network buffers, creating congestion events and suppressing mice flows on these links.

This leads to a detrimental degradation of application performance such as online gaming, web requests, multimedia broadcasting and VoIP which are of a particular importance for cable operators who continuously strive to improve quality of experience for their customers [24].

In the DDC/DDBR solution, the NNI links do not rely on ECMP or any sort of Layer 3 routing rather taking advantage of the capabilities enabled by merchant silicon for redistributing traffic equally with remarkable precision.

- Variable sized ethernet frames versus cell-based fabric traffic:

In the Spine-Leaf network architecture packets that are passing through fabric links (NNI links) are of variable size.

Packets transiting the fabric in the DDC/DDBR architecture are proprietary formatted fixed sized cells [\[23\]](#). Thus, monitoring systems must account for differences in the way the payload is transmitted.

- Spreading of traffic over fabric ports:

In the Spine-Leaf network architecture the traffic is spread over disparate NNIs via ECMP and the routing protocols' logic, and the processing occurs in the control-plane by the NOS where the routes are programmed in the tables of the forwarding ASIC.

In the disaggregated solution, cell-based traffic is used on fabric ports and the spreading logic is being executed at the microcode level of the forwarding ASIC and is completely transparent to the NOS. Fairly sophisticated credits/tokens mechanism is involved in the implementation of this feature [\[5, 23\]](#).

- Managing a folded Clos topology of N-independently functioning elements versus all elements are under control of a single NOS:

In the DC Leaf-Spine design, each node ordinarily runs a NOS that is independent. In this case, forwarding decisions are made individually on each node which lack a global view of the network, so the best path computation algorithms are at best locally optimized.

In the DDC-RS, the Packet and Fabric Forwarders function under the oversight of a single NOS. The virtual router needs only a single network management interface and is considered as one big virtual chassis. Subsystems of the operating system can be distributed to the elements of the cluster but the "Brains" of it are centralized on the powerful compute nodes. This also makes monitoring and onboarding automation frameworks easier.

From this point of view, multiple routing algorithms are optimized through a single routing system, The traffic is just traversing through one router – from one port to another port on a separate Packet Forwarder from the end-to-end network perspective [\[5\]](#).

- Easier upgrade process:

Software and firmware upgrades usually take more time and require multiple maintenance windows in the S/L case.

Due to multiple independent components, each of them must be upgraded separately. In many cases, not all traffic can be migrated away from the pod which introduces a need to move traffic away from certain leaves and spines within it to proceed with the upgrade which negatively contributes to the already complex process.

Due to a single point of management, disaggregated backbone/aggregation router upgrades have the potential to upgrade either some or all the elements of the cluster within one maintenance window considerably reducing the operational overhead and the amount of scheduled maintenance windows. This gives the user the flexibility of shifting traffic off the cluster, where in the past with traditional routers or in the S/L network upgrading sets of spines and leaves left the potential to be exposed to being in a hazardous condition.

Furthermore, in the traditional S/L model convergence of the IGP takes more time before the traffic is soaked out of nodes but there is a minimal negative impact on Quality of Experience (QoE) for customers. However, in the DDC/DDBR case due to this being a single virtual chassis traffic can be diverted immediately and converges instantaneously.

There is of course a tradeoff between a graceful traffic soak from S/L with minimal negative impact on the customer Quality of Experience which can take more time and leave the system exposed, whereas in the yank and replace method the protocol must fail the traffic over and traffic reconverges more abruptly but is quicker. There is a certain impact on the customer's QoE in either method.

4. Orchestration, Automation and Analytics

Despite bringing meaningful scale and cost advantages to the table, this new disaggregated network operational model might potentially be considered a risk. This approach has emerged recently in the industry and service providers are weighing the extra complexity related to the orchestration, automation, and management of this open model, specifically in the areas of field installation, upgrades, capacity growth, and troubleshooting.

Cloud orchestration is a solution that is required to deliver the operational simplicity, automation, and visibility to the DDC/DDBR architecture to drive the acceleration of the deployment of this cloud-native networking solution. It should offer detailed visibility into the system's internal architecture including hardware and software components, KPIs related to SLA, alarm, and fault management.

4.1. Automated Operations

This encompasses the lifecycle management of resources and services – from provisioning to decommissioning, which includes:

Zero-touch provisioning (ZTP) – automatically integrates multi-vendor white box hardware and NOS into a working routing platform supporting a secure deployment prone to less errors with limited manual intervention. Ahead of entering the protected and controlled operator network's environment, white boxes go through an extra security measure which is the bootstrap process to ensure system integrity.

Hardware inventory management – grants detailed data on every element within the cluster, including location, model, serial number, firmware version etc. This becomes even more so critical as the cluster scales out and the inventory management gets too complicated to manage manually.

Modular software orchestration – covers the entire stack and can be done selectively per specific software component, counting firmware, base OS, NOS image/container. This capability becomes extremely handy because upgrades can be initiated per component without affecting the overall software stack reducing the risk of failure in cases when the whole cluster is upgraded. Another advantage is the real-time orchestration status which helps engineers to monitor the activity and potentially prevent issues before they occur. Software rollback can also be easily performed in cases when the upgrade was unsuccessful.

4.2. Health Monitoring and Assurance

The orchestration tool automates event and KPI monitoring and ensures availability and performance SLAs such as:

Cluster topology – live view of the cluster's nodes, their states, formation and connectivity across clusters and the entire network, including:

- Hardware components: CPU, memory, PSU, fan, temperature, ports, and interfaces.
- Software components: base OS, firmware, processes, containers, and microservices.

Fault, performance, and alarm management: on top of 3rd party applications that collect data for events management, the cloud orchestration system can provide all these details for the clusters under its management:

- Supports alarms and KPIs at every level of the system – from hardware components to software containers.
- Alarm dashboard to monitor and categorize system alarms.
- Real-time and time-series alarm view.

Tech support integration for in-depth system diagnosis and debugging: being able to retrieve tech-support files when there is a potential software or hardware bug is paramount for network operations teams. The orchestration tool should allow a simple way to retrieve such a bulky file from the clusters for analysis by the NOS and hardware vendors.

4.3. Telemetry

Historically, SNMP has been the de-facto protocol for data collection, however, it may not be suitable for the management of truly large networks because of the performance limitations of polling.

gRPC is Google's project for Remote Procedure Calls between applications. gRPC based telemetry streaming allows to export performance monitor counters and operational-state parameters in a flexible and scalable way which will help operators tremendously as the number of devices in their networks grows exponentially [25]. Unlike traditional performance monitoring (PM) collection methods such as SNMP walk, gRPC based telemetry uses push method for delivering PM data from router to the PM collector. Such an approach gives the following advantages:

- PM Collector does not need to poll each router individually.
- PM Collector can define different set of counters for collection from specific routers.
- Push parameters can be configured to each pushed counter specifically (e.g., sample rate, telemetry packet DSCP value etc.).
- gRPC interface when combined with Protobuf encoding is more efficient in terms of network channel utilization than SNMP and other interfaces such as NetConf and RestConf.

5. Real-life deployment of the DDC/DDBR model

DDC/DDBR architecture is a field-proven concept deployed in the core and aggregation layer of the largest telcos and cablecos in the world [\[26\]](#).

The implementation started with the channel partners wiping the received whiteboxes and installing the ONIE and BaseOS on top of it in preparation for the NOS integration. After this step had been completed, all the elements were shipped to the respective data centers for the site installation and wiring.

Once the cabling was finalized and all elements within the cluster powered up, the route controllers were discovered by the orchestration system which initiated the cluster creation via the Zero-Touch Provisioning process. At this point, the cluster is ready for the enablement of the control plane and for passing traffic.

There are several tools and dashboards needed to monitor production grade Clos networks including the need to monitor the load sharing over ECMP paths, scale of the databases in the IGP domain due to the large number of nodes. On the other hand, with a virtual cluster:

- It became easier for the NOC team to manage the network and respond accordingly to the events within the system.
 - For example, the load sharing on fabric links is astonishingly equal and does not require any monitoring because it is done at the microcode level, as we have mentioned earlier.
- No IGP or any other routing protocol required between the components of the cluster as is needed in a S/L system, thus reducing the number of nodes in the IGP domain.
- There is also no need for a multi-level BGP design within the cluster and consequently no requirement for BGP enhancements and features to improve recovery time during convergence.

Now, an abstraction layer is introduced which hides all the complexity of the Clos network and appears as a single router which is so familiar to network engineers. The next figure demonstrates an example of a large national backbone network:




Figure 11 - Large National Backbone Network

Every backbone site either has a traditional router or a spine-leaf cluster. Let's zoom in into one of the sites that hypothetically contains a spine-leaf cluster:




Figure 12 - Spine-Leaf cluster within a Backbone Site

As can be seen from the above figure, there is an ISIS and BGP mesh that must be maintained within the cluster on top of interconnections between different backbone, aggregation, and peer networks. The next image showcases a similar topology but with a virtual chassis:

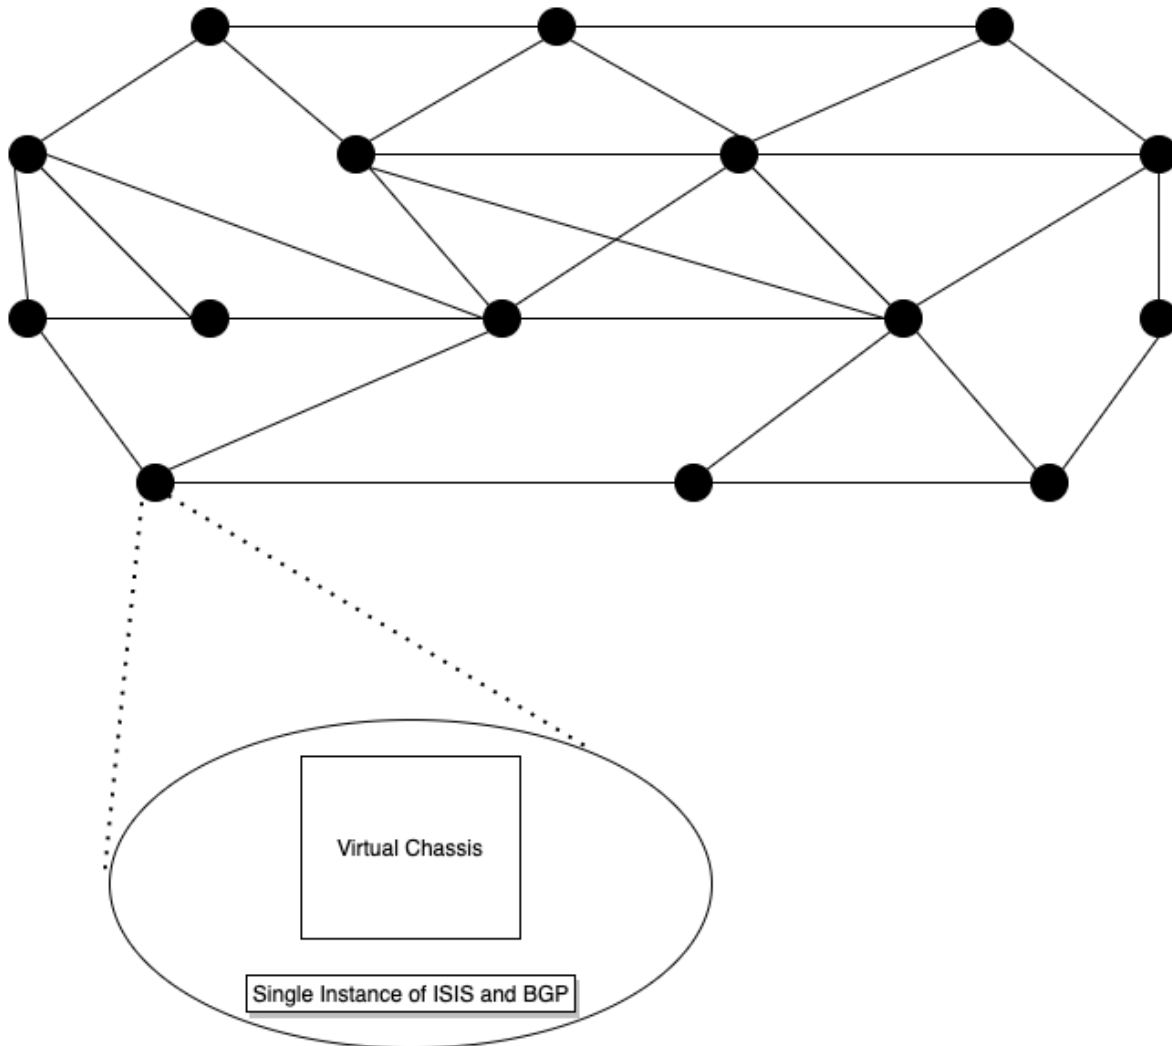


Figure 13 - Virtual Chassis within a Backbone Site

The virtual chassis depicted in Figure-13 eliminates the requirement for intra-cluster control-plane. Now, network architects can focus on mission critical services rather than assuring the SLAs within the cluster.

Any future launches of a virtual chassis in the network will be done by the orchestration system which supports comprehensive validation methods of spaghetti-wiring fabric cabling, management of the code upgrades and device configurations. Another important activity of the engineering and operations teams is auditing the provisioned components count against the actual operational count so that N+M redundancy can be maintained which is already incorporated into the orchestration system alleviating the need for engineers to do it manually.

With a disaggregated approach comes a challenge of the identification of the responsibility domain: hardware and software. Since there are 2 vendors delivering these components, occurrence of the issues requires an understanding of which domain it belongs to, to deliver necessary tech-support files and logs to the respective TAC. There is a need for cooperative collaboration between HW and SW vendors during troubleshooting of such systems so that the root cause of an issue can be reached quickly and amicably.

6. The future of the disaggregated solutions

In the words of the ancient Greek philosopher Heraclitus, “The only constant in life is change”. Nowadays, this most vividly applies to the technology world. The pace of innovation must keep up with the demand and new emerging use cases. Telecommunications companies along with vendors will keep driving innovation in disaggregated network solutions. Essential future developments are multiservice functionality on top of the shared pool of resources, OpenOffload and data center sustainability.

The average number of devices connected to the internet per household and per capita is increasing. Expanding machine-to-machine applications contribute in a major way to device and connection growth and push the expansion of data center infrastructure. With the growth of IoT and “smart” appliances, almost any type of device needs connectivity to the home network. Examples of this are laundry machines, refrigerators, stoves, microwaves, thermostats, garage openers, and even the front door keypad of a home! If one would check the number of devices connected to their home Wi-Fi network, they would be surprised at the number.

One of the promising concepts being developed is utilizing the network as a cloud resource leveraging this model’s High-Availability (HA) capabilities. By applying modern, cloud-based, shared resource methodologies to networks, telcos and cablecos can enhance their network resource utilization by taking advantage of the unified infrastructure that supports multiple network functions as software-based services, such as enterprise, broadband, mobile, firewall, load-balancer etc. [\[27\]](#). Any part of the cluster, designated for a specific function, can be used to enable this service. Sharing the physical infrastructure for various services extensively lowers the physical footprint within data centers and reduces the number of unused ports, resulting in a more efficient employment of compute and networking resources. In the past we had core, edge, and backbone nodes that were separate units. Now they can be aggregated onto a unified cloud-native infrastructure.

Spinning up various network functions is very demanding from the compute resources perspective which calls for original approaches in dealing with processing. There is an initiative in the industry that defines APIs to accelerate network functions and applications by offloading packet processing to the packet forwarders rather than x86 servers which already have their fair share of load. One of the APIs is for applications like Virtual Firewalls and Intrusion Detection Systems called OpenOffload [\[28\]](#), and another one is for functions like VPN Gateways that are designed to offload IPSEC and GENEVE [\[29, 30\]](#) tunnel processing to the hardware. This development will allow ISPs/CSPs to efficiently use the disaggregated clusters by running virtual functions on top without a compromise in throughput.

Network operators are continuously seeking technologies to reduce carbon footprint. One of the sustainable solutions can be found when what was once a sophisticated deployment becomes antiquated for its environment due to the changing conditions on the ground. Decommissioning a disaggregated cluster and re-using its hardware in a developing country where it can be practical for many years to come paves the way for more sustainability. It is the function/capacity separation of the DDC/DDBR architecture that facilitates resiliency in reusing the existing hardware to serve different geographical regions [\[31\]](#).

7. Conclusion

The explosion of the traffic growth over the last two decades has been somewhat unexpected and has urged telecommunications companies to explore alternative routing solutions. The disaggregated approach to the traditional chassis routers has been gaining significant interest from ISPs/CSPs' driven by a variety of motives such as cost reduction, the removal of vendor lock-in and service innovation. Regarding the talent pool, due to the proliferation of the innovative networking solutions, network architects and engineers should master new skills to stay relevant in the industry. Traditional chassis router architectures, having dutifully served for decades, impose multiple challenges like the traditional nature of the components, and the hurdles to overcome upgrade simplicity. DDC/DDBR architecture solves many of the outlined issues while presenting service agility and faster innovation. Spine-Leaf architecture has been deployed in the data centers of hyperscalers and backbone networks of the large telcos and cablecos. Although this approach has multiple drawbacks which are now addressed by the disaggregated solution, this approach is not only valid on paper but there are multiple scalable DDC/DDBR deployments in the backbone and aggregation networks of the largest cable and telecommunications companies in the world. On the automation side, cloud orchestration system is leveraged to streamline cluster management, device configuration and critical files retrieval. In the foreseeable future, development in the network disaggregation domain will focus on multiservice, session offload by leveraging open source OpenOffload API all while trying to remain cognizant of sustainability of the network. It seems as if we are getting closer to a more complete automated, portable, and easily scalable network for the future!

Abbreviations

ASIC	Application Specific Integrated Circuits
API	Application Programming Interface
BGP	Border Gateway Protocol
CLI	Command-line Interface
CNCF	Cloud Native Computing Foundation
COTS	Commercial Off-The-Shelf
CPU	Central Processing Unit
CSP	Communications Service Provider

DC	Data Center
DDBR	Disaggregated Distributed Backbone Router
DDC	Disaggregated Distributed Chassis
DOCSIS	Data Over Cable Service Interface Specification
DSCP	Differentiated Services Field Codepoints
E2E	End-to-end
ECMP	Equal Cost Multipath
FF	Fabric Forwarder
FRU	Field-Replaceable Unit
GE	Gigabit Ethernet
GENEVE	Generic Network Virtualization Encapsulation
GNMI	gRPC Network Management Interface
HA	High-Availability
HTTP	Hypertext Transfer Protocol
HW	Hardware
IGP	Internal Gateway Protocol
IGW	Internet Gateway Router
IoT	Internet of Things
IP	Internet Protocol
IPSEC	Internet Protocol Security
ISIS	Intermediate System - Intermediate System
ISP	Internet Service Provider
ISSU	In-Service Software Upgrade
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
MoP	Method of Procedure
MPLS	Multiprotocol Label Switching
NETCONF	Network Configuration Protocol
NMS	Network Management Systems
NNI	Network-to-Network Interfaces
NOS	Network Operating System
NSR	Non-Stop-Routing

OCP	Open Compute Project
ODM	Original Design Manufacturers
OEM	Original Equipment Manufacturers
OIR	Online Insertion and Removal
ONIE	Open Network Install Environment
OSPF	Open Shortest Path First
PF	Packet Forwarder
PM	Performance Monitoring
Protobuf	Protocol Buffers
PSU	Power Supply Unit
QoE	Quality of Experience
QSFP-DD	Quad Small Form Factor Pluggable Double Density
REST	Representational state transfer
RPC	Remote Procedure Call
RPM	RPM Package Manager
RS	Routing System
S/L	Spine-Leaf
SCTE	Society of Cable Telecommunications Engineers
SLA	Service-level Agreement
SNMP	Simple Network Management Protocol
SW	Software
TAC	Technical Assistance Center
TCAM	Ternary Content Addressable Memory
TCP	Transmission Control Protocol
TIP	Telecom Infra Project
UNI	User Network Interfaces
VPN	Virtual Private Network
WAN	Wide Area Network
WIFI	Family of wireless network protocols/Wireless Fidelity
XML	Extensible Markup Language
YAML	YAML Ain't Markup Language
YANG	data modeling language
ZTP	Zero-touch Provisioning

Bibliography & References

- [1] <https://www.broadbandsearch.net/blog/internet-statistics>
- [2] <https://corporate.comcast.com/stories/announcing-another-10g-milestone-amidst-a-flurry-of-innovation>
- [3] <https://www.broadcom.com/blog/800g-optical-platform-solutions>
- [4] <https://www.youtube.com/watch?v=d90uK4WW0Po>
- [5] *Hardware Specifications and Use Case Description for J2-DDC Routing System*, Tuan Duong; Open Compute Project
- [6] *Distributed Disaggregated Backbone Router (DDBR) Technical Requirements Document*, Eva Rossi, Jose Angel Perez, Kenji Kumaki, Ryuji Matsunaga, Yuji Sonoki, Ahmed Hatem, Diego Marí Moretón; Telecom Infra Project
- [7] <https://www.opencompute.org/>
- [8] <https://telecominfraproject.com/>
- [9] GigaOm Radar for Network Operating Systems: Network Service Providers, https://research.gigaom.com/report/gigaom-radar-for-network-operating-systems-network-service-providers/?utm_content=208592782&utm_medium=social&utm_source=linkedin&hss_channel=lcp-11282464
- [10] <https://www.cncf.io/>
- [11] <https://www.opencompute.org/wiki/Networking/ONL>
- [12] *The Internet of Things: Catching up to an accelerating opportunity*, Michael Chui, Mark Collins, Mark Patel; McKinsey & Company
- [13] <https://restfulapi.net/>
- [14] <https://datatracker.ietf.org/doc/html/rfc6241>
- [15] https://www.w3schools.com/xml/xml_what.asp
- [16] https://www.w3schools.com/js/js_json_intro.asp
- [17] <https://yaml.org/>
- [18] <https://datatracker.ietf.org/doc/html/rfc7950>
- [19] <https://opencomputeproject.github.io/onie/overview/index.html>

[20] *Optical Transceivers for Datacom & Telecom Market 2021: Market and Technology Report 2021*, Dr. Martin Vallo, Pars Mukish, Dr. Eric Mounier; Yole Development

[21] *A Study of Non-Blocking Switching Networks*, Charles Clos; The Bell System Technical Journal, March 1953

[22] *On Nonblocking Folded-Clos Networks in Computer Communication Environments*, Xin Yuan, Department of Computer Science, Florida State University, Tallahassee, FL 32306, 2011

[23] <https://www.ufispace.com/company/blog/what-is-a-distributed-disaggregated-chassis-ddc>

[24] *Engineered Elephant Flows for Boosting Application Performance in Large-Scale CLOS Networks*, Broadcom Corporation, March 2014

[25] <https://grpc.io/>

[26] https://about.att.com/story/2020/open_disaggregated_core_router.html

[27] *Multiservice: Maximizing Infrastructure Utilization*; DriveNets Whitepaper

[28] <https://github.com/att/sessionOffload>

[29] <https://datatracker.ietf.org/doc/html/rfc6071>

[30] <https://datatracker.ietf.org/doc/html/rfc8926>

[31] *Topology and component description of DriveNets Network Cloud as a sustainable solution for multipurpose networking*, Run Almog; Open Compute Project

[32] *Annual outage analysis 2021*, Andy Lawrence; Uptime Institute
https://uptimeinstitute.com/uptime_assets/25ff186d278b32c202fc782e60a0d473bd72bfbcb6d4d65afedfa15dd406c7656-annual-outage-analysis-2021.pdf

The Power of YANG Configuration Templates

A Technical Paper prepared for SCTE by

Pawel Sowinski
VP Technical Excellence
Falcon V Systems
Sudbury, Mass
p.sowinski@falconvsystems.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Flexible MAC Architecture.....	3
3. Addressing the Problem of Scaling Configuration Management Data.....	6
3.1. The Intended Audience	6
4. YANG-based Configuration Templates.....	6
4.1. What are Configuration Templates?	6
4.2. YANG-based Configuration Templates.....	7
4.2.1. Data Payload.....	7
4.2.2. Metadata	9
5. Template-Oriented Configuration Management System.....	10
5.1. Repository	11
5.2. Target Database	11
5.3. Template Management Application	12
5.4. Template Rendering.....	12
6. Discussion of Selected Use Cases	14
6.1. Modular Data Design	14
6.2. Subdivision of Responsibilities within the Organization.....	15
6.3. Configuration Lifecycle Management.....	16
7. Conclusion.....	16
Abbreviations	17
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 - FMA Phase 1 Reference Architecture	4
Figure 2 - YANG-based Configuration Template Framework Scope.....	5
Figure 3 – Yang-Based Configuration Template	8
Figure 4 – YANG-Based Configuration Template Data Payload	8
Figure 5 – YANG-Based Configuration Template Metadata Payload	10
Figure 6 – Template-Oriented Configuration Management System	11
Figure 7 – Template Rendering Example	14
Figure 8 - Libraries of Fundamental Templates	15

1. Introduction

Cable Operators pursuing Flexible MAC Architecture (FMA) deployments are facing the prospect of a N-fold increase in the number of managed devices in their HFC access networks. The increase is a result of a necessary replacement of integrated CCAPs with Distributed Access Architecture (DAA) components including MAC Managers and Remote MAC Devices (RMDs), where RMDs serve a much smaller service-delivery footprint. Typically, RMDs provide services for just one service group. As part of the DAA transition, N-number of RMDs displace each I-CCAP, a large-scale platform supporting N service groups. The number “N” could reach or exceed 128. This dramatic increase in scale requires a shift towards zero-touch, automated provisioning processes. While FMA specifications already provide tangible solutions for this problem at the interface between the MAC-NE and the MAC Manger, the paper examines the impact of this transition on the provisioning tasks within the operators’ back-office.

Two additional key observations lay the technical foundation for this paper. First, a complete configuration dataset of each RMD is unique in that it includes many configuration attributes that are specific to a particular device and its deployment topology. The examples of such device-specific attributes include power levels configured on the RF components, the geographical location, and the inventory tag. Second, most of the device configuration attributes, including service group configurations, are common across large groups of RMDs.

The paper outlines a framework that capitalizes on abstraction of repeatable patterns from the standardized YANG configuration models of RMDs with creation of configuration templates corresponding to the identified patterns. YANG-based configuration templates coupled with unique device attribute values maintained in a relational database allow automated systems to generate complete device configuration datasets enabling effective management of the configuration lifecycle, including software or service upgrades.

The framework provides benefits to Cable Operators by reducing the overall size of fundamental configuration data and by providing structure better suitable for automation. For example, a Cable Operator deploying thousands of RMDs may be able to reduce the footprint of their fundamental configuration dataset to just a few dozen YANG-based configuration templates and use predictable, automated system to dynamically generate full-formed, operational configuration datasets from these templates. The outcomes are in the form of modular, automated, and streamlined provisioning system with increased agility and reduced OPEX.

The paper outlines a YANG-based configuration template framework, based on FMA YANG models developed by CableLabs[®], while providing several real-world examples of the benefits.

2. Flexible MAC Architecture

CableLabs has released a suite of Flexible MAC Architecture specifications and continues to expand the specifications through industry collaboration. The FMA project is a key element of a larger CableLabs Distributed Access Architectures (DAA) program. The DAA program includes additional projects, such as Remote PHY (RPHY) and Coherent Optics as well as an industry-wide YANG model standards library and YANG model development pipeline.

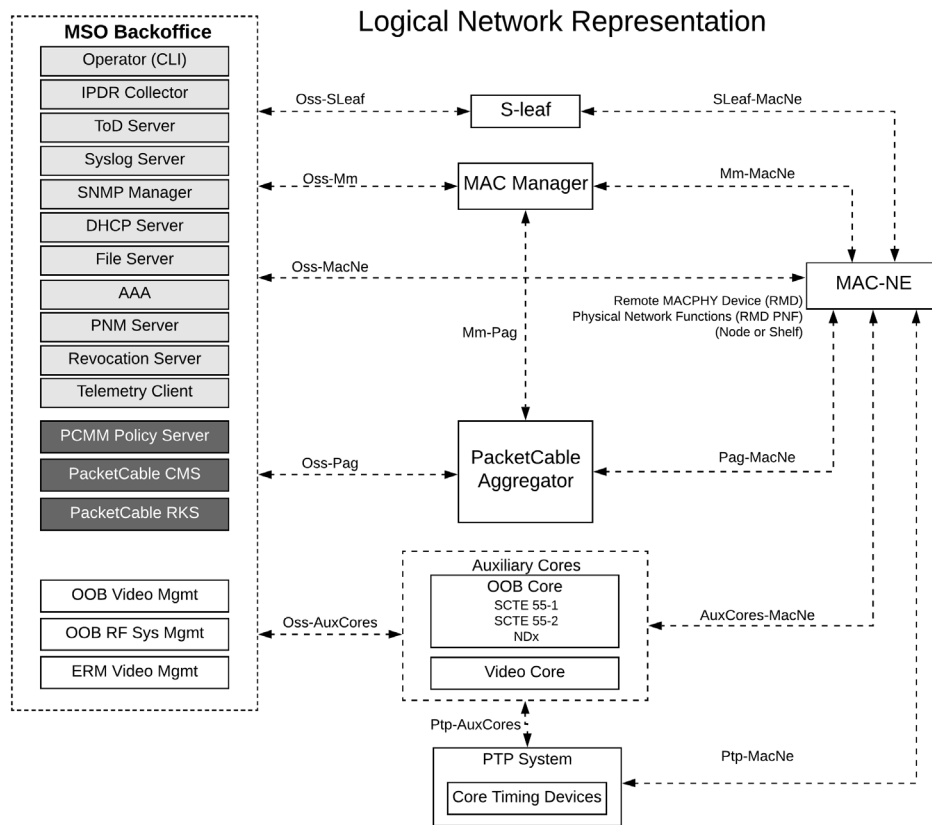


Figure 1 - FMA Phase 1 Reference Architecture

The suite of FMA specifications have achieved a level of stability with the Phase 1 scope officially concluding at the end of CY 2021. The Phase 1 FMA Reference Architecture, as defined in [FMA-SYS], is illustrated on Figure 1. A detailed description for each FMA system component and interface displayed in the Reference Architecture can be found in [FMA-SYS] and [FMA-OSSI].

The following sections of the paper will narrow the focus to those FMA components and interfaces (realized as APIs) which are relevant to the YANG-based configuration template framework. The goal is to align the framework against the FMA Reference Architecture.

Figure 1 depicts the interface between the MAC Manager and the MAC-NE (e.g., RMD). This interface is generally referred to as the MAC Manager to MAC-NE Interface (MMI), labelled as ‘Mm-MacNe’ in the diagram and specified in [FMA-MMI]. The interface’s implementation relies on the RESTCONF protocol, and a library of MAC-NE YANG models developed by CableLabs. The IETF defines the RESTCONF protocol as specified in [RFC 8040].

The MAC Manager fulfills the role of a RESTCONF Client at the northbound of the Mm-MacNe interface, while the MAC-NE operates as the RESTCONF Server at the southbound. Among other primitives, the MAC Manager is responsible for transferring all configuration information to the MAC-NE over the MMI/Mm-MacNe interface. Since the RESTCONF protocol facilitates unicast communication, the MAC-NE YANG models have been designed for transfer of configuration to an individual MAC-NE (RMD).

Another FMA Reference Architecture interface relevant to the subject of the paper, extends between the MAC Manager and the operator's back-office systems, labeled as 'Oss-Mm' in Figure 1. In addition, Figure 1 lists several back-office applications used in the management of an FMA deployment, however it should be pointed out that a network Configuration Management System (CMS) is not included. The CableLabs' FMA specifications do not explain whether configuration management was purposefully omitted as out-of-scope for Phase 1, or under the assumption that configuration of MAC-NEs is an internal domain of the MAC Manager. This paper makes an assumption that a CMS is a mandatory part of the Cable Operators Network Management System (NMS) infrastructure and focuses on selected issues revolving around configuration management functions in FMA deployments.

Figure 2 illustrates an updated reference architecture diagram and outlines the scope of functionality covered by the paper.

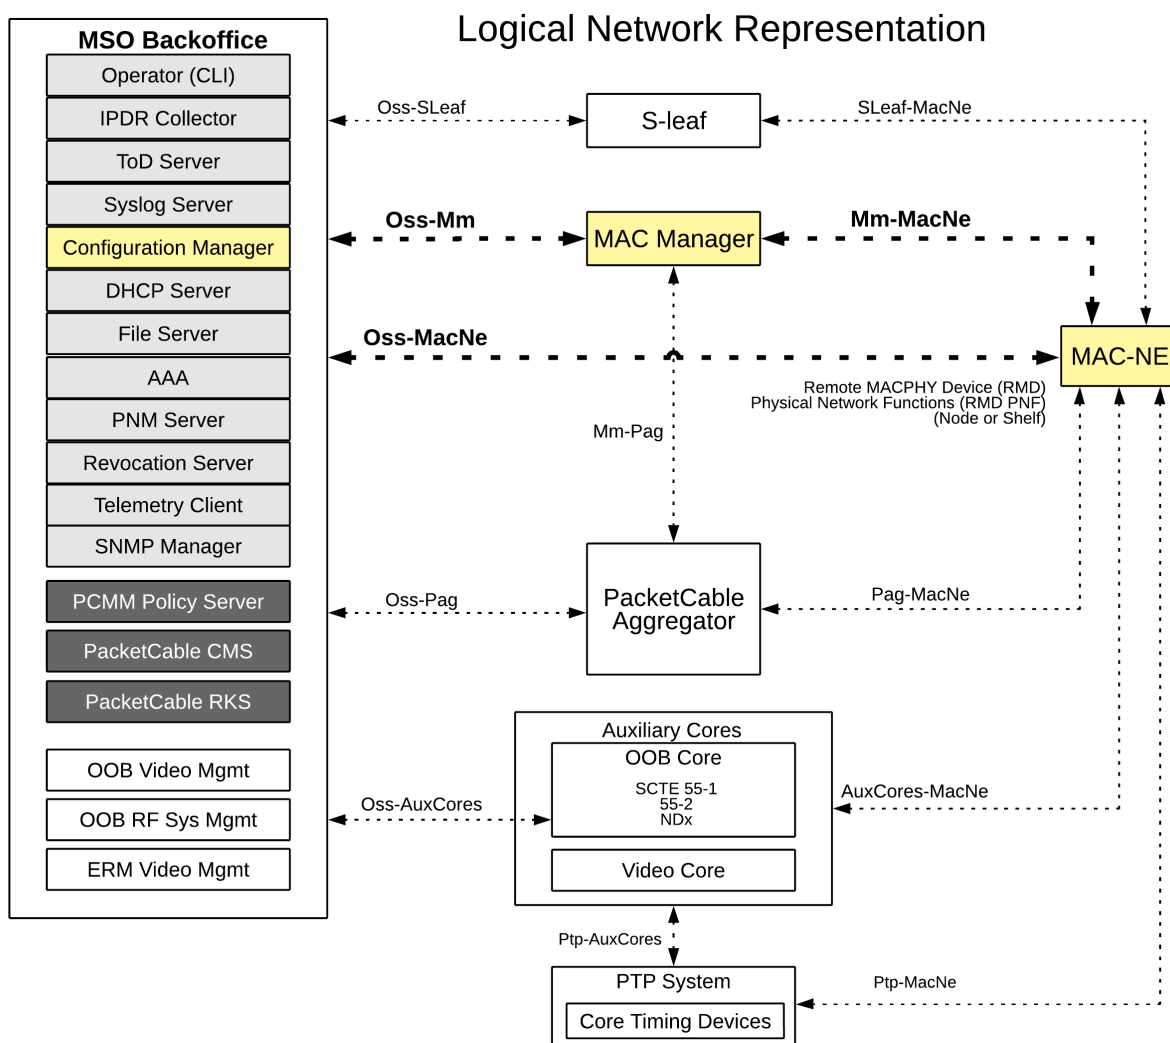


Figure 2 - YANG-based Configuration Template Framework Scope

3. Addressing the Problem of Scaling Configuration Management Data

Let us first briefly illustrate the scale of the provisioning problem faced by Cable Operators deploying FMA through examination of the following example. An HFC network of a medium-sized Cable Operator with a few million paying subscribers may consist of 100,000 service groups. If each service group corresponds to exactly one RMD and considering that an average size of an initial configuration dataset can undoubtedly exceed 500 KB (or 10,000 discrete configuration attributes), and the cumulative size of the initial configuration dataset for such an operator could amount to $500 \text{ KB} * 100,000 = \mathbf{50 \text{ GB}}$ (or 1 billion discrete configuration attributes). This is the volume of the initial configuration information, often referred to as “day-one config”, for all RMDs in this example. This data is encoded and transferred across the ‘Oss-Mm’ interface presented on Figure 2 via the RESTCONF protocol, in a unicast fashion, one-by-one between the FMA MAC Manager and each RMD in the Cable Operator’s network whenever an RMD is initialized or reinitialized.

A data footprint of such scale can very well fit into the storage of an average smartphone. How is that a problem, one might ask?

While storage and transfers of several gigabytes of data cannot constitute a problem for any modern networked computing system, managing “the source of truth” of such sizable footprint of configuration data through the lifecycle of 100,000 independent systems is an entirely different ballgame.

The data abstraction approach presented in the paper takes advantage of the repeatable patterns in the deployment environment, the cookie-cutter like HFC access network design and the standardization of device configuration management interfaces using YANG data models. This methodology relies on breaking down the devices’ configuration datasets into smaller, more manageable modules, extracting individual configuration attributes and their values into a standalone database and creating reusable configuration templates from common, shared datasets. The key enabler of this approach is YANG and its powerful, modular definition of data models.

3.1. The Intended Audience

The intended audience for this paper is the personnel of Cable Operators who are designing or preparing to deploy FMA based systems in their HFC networks, including network provisioning architects, engineers, application developers and business analysts. This paper assumes that readers have some working knowledge of YANG modeling and protocols designed to transport YANG data such as RESTCONF or NETCONF.

4. YANG-based Configuration Templates

4.1. What are Configuration Templates?

This paper adopts the term “Configuration Template” to be commonly understood as **a predefined prototype for instantiated configuration data** that can be further customized and applied to multiple systems in various scenarios.

Configuration Templates are broadly deployed across the software technology industry and distinctly in the networking and cable sectors. For example, some CCAP manufactures provide templates as means to reduce the size of the CLI configuration payloads for network interfaces and selected functions.

4.2. YANG-based Configuration Templates

FMA relies on YANG as a modeling language for MAC-NE status and configuration datastores. YANG was originally defined in [RFC 6020] as a modeling language for the description of data carried in the NETCONF protocol. Since then, YANG has graduated to version 1.1 [RFC 7950] and a wide array of RFCs and other technical standards have been developed to cover adjacent aspects such as encoding, metadata, versioning, translation and handling operational issue with multiple datastores. Today, the networking and cable industry operates with the collateral of thousands of YANG modules, some of which have been published or remain in private distributions.

However, there is no standard or formal definition for YANG-based Configuration Templates (YCTs). No IETF RFC, nor an open-source project or framework provides a formalized definition for a YCT. A YANG-based Configuration Template is a term coined herein. In this context, a YANG-based Configuration Template is a Configuration Template in which the instantiated configuration data is compliant to a particular YANG model and the data encoding generally follows the common encoding rules for YANG data as defined in relevant IETF standards. A YCT can be programmatically validated against the corresponding YANG model.

Each YCT consists of two sections:

1. Data Payload
2. Metadata

The YCT sections are explained in greater detail in the following sections of the paper.

4.2.1. Data Payload

The most essential part of a YCT is the data payload. The data payload contains the instantiated configuration dataset corresponding to a branch of the device's YANG schema. The data payload must not contain any status or operational state data, but partial datasets are permissible. The data payload of one template is typically compliant to a model defined by multiple YANG modules.

The data payloads can be represented in several human and machine-readable formats, such as XML, YAML or JSON. This paper uses JSON for the format of the payload data in the following examples due to its compactness and human readability. The template data in the paper adheres to the standard JSON encoding rules for YANG-defined data per [RFC 7951].

In the simplest form of a YCT, the data payload contains the complete configuration dataset for a branch of the device's YANG schema tree. The example in Figure 3 shows the data payload for configuration of MAC-NE security parameters. Such a template can be incorporated into the target configuration without further changes. The configuration target is defined in section 5.2.

```
"security": {  
  "sav": {  
    "sav-control": {  
      "cm-auth-enable": true  
    }  
  },  
  "tftp-security-config": {
```

```

    "config-file-learning-enabled": true
  },
  "cmts-encrypt": {
    "encrypt-alg-priority": "AES128_CBC_MODE DES56_CBC_MODE DES40_CBC_MODE"
  },
  "certificates": {
    "cert-revocation-method": "NONE",
    "online-cert-status-protocol": {
      "signature-bypass": false
    }
  }
}

```

Figure 3 – Yang-Based Configuration Template

More interesting scenarios ensue when YCT data payloads include variables and/or expressions.

A template variable constitutes a data node value, which is kept within the template as a human and machine-readable name in the form of a character string. Variable names in the paper examples, by convention, are presented as strings that start with the '\$' character (e.g., \$my-variable-name).

The values for the template variables are maintained outside of the template definition, in a Target Database, which is described in a later section of the paper. A variable designates a node value specific to the template target, which is typically an individual MAC-NE device or a group of MAC-NE devices. The Configuration Management System replaces variables with actual values during the process of rendering a template. The template rendering process and the template-oriented configuration management system are described in a later section of the paper.

The example provided in Figure 4 shows a data payload for a configuration of a MAC-NE for Precise Time Protocol (PTP) with a variable for the value of the attribute "ptp-master-addr". The name of the variable in this example is '\$ptp-master-ip-address'.

```

"rdti-cfg": {
  "core-ptp-clk-cfg": {
    "ptp-clk-profile-id": "00:19:a7:02:01:00",
    "ptp-master-addr": "$ptp-master-ip-address",
    "ptp-master-priority": 0,
    "ptp-clk-priority1": 128,
    "ptp-clk-priority2": 255,
    "ptp-clk-domain": 24
  }
}

```

Figure 4 – YANG-Based Configuration Template Data Payload

When such a template is rendered for a selected target, the system retrieves the actual IP address of the PTP master clock from the Target Database entry and inserts it into the configuration dataset. With this approach, the PTP master clock server's IP address can be individually applied for each target and combined with common data from the template to form a complete configuration dataset for the device.

A template expression is a logical or a mathematical formula for algorithmic determination of a value of a data node in the template. An expression can include operations on constants, template variables and values of other data nodes contained within the template, including the key attributes of YANG lists. Similarly, as with variable replacement, expressions are evaluated during the rendering of a template for a selected target.

A familiar example of a template expression could be the formula to calculate the value of frequency of a downstream SC-QAM channel when several such channels are grouped into a single, continuous block. Such a formula can be the sum of the value of a variable representing the center frequency of the lowest channel in a block and channel index multiplied by 6 MHz.

*\$channel-frequency = \$start-frequency + channel-index * 6,000,000*

The syntax of expressions and their usage rules are quite complex. For this reason, and because expressions' encodings are not essential to the purpose of the paper, the detailed definition of template expressions is left out of the scope of the paper.

4.2.2. Metadata

Template metadata consists of a set of attributes that hold miscellaneous properties of a template outside of the configuration data payload. The most important metadata attributes are the **template-name** and the **template-schema-root**.

The template name serves as the primary identifier of the template within the system. The system refers to templates by their names. For this reason, each template name needs to be unique within the system.

The YANG modelled data maintains a tree-like hierarchy where each node can be uniquely identified with a schema path. The template schema root defines the parent node of the target's YANG data schema tree from which the data within the template is in scope. Thus, the template schema root is a YANG schema path. It can include keys and variables so that different templates can be rendered onto specific elements of a single YANG list. A template metadata defines exactly one or zero template schema roots. When the template is rendered for a selected target, the system can override the template schema root with a custom template schema root configured for the target. For example, the template schema root can refer to an individual entry of a list, or to all entries of a list. In the former case a template provides data for the referenced list entry. In both cases, the set of data nodes corresponding to the list can be formed from multiple templates.

Other metadata attributes provide information which can be helpful with administrative tasks. The example of template metadata shown below comprises several other attributes, such as **template-description**, **template-notes** and **template-revision-history**.

The metadata encoding is consistent with the JSON encoding rules for YANG data nodes. All metadata attributes, just like the data payload, are encoded as JSON key-value pairs. Template metadata encoding follows the rules from [RFC 7952].

An example of template metadata is shown in Figure 5.


```

“template-metadata”: {
  “template-name”: “ptp-config-101”,
  “template-schema-root-point”: “/mac-ne/networking/”,
  “template-description”: “This template defines configuration for MAC-NE PTP
operation for G.8275.2 PTP profile”,
  “template-notes”: “This template is incomplete. It is missing a number of
mandatory configuration attributes”,
  “template-revision-history”: [
    {
      “revision”: “2022-01-01”,
      “description”: “Added ptp-clk-domain attribute.”,
      “Author”: “George D.”
    },
    {
      “revision”: “2021-10-01”,
      “description”: “Initial revision”,
      “Author”: “Bob D.”
    }
  ]
}

```

Figure 5 – YANG-Based Configuration Template Metadata Payload

5. Template-Oriented Configuration Management System

This section describes an environment in which all data configuration originates from a template. The difference between a template and configuration payload is the name.

This section provides a rough framework for a system which creates MAC-NE configuration payloads from YANG configuration templates. The paper limits the description of the system to a minimum necessary to illustrate the principles of operation of a system based on Configuration Templates. In the context of this paper, the system is referred to as a Template-Oriented Configuration Management System, or TOC System. The TOC can become an integral part of the Cable Operator’s Data Operations platform architecture.

The TOC System is presented on Figure 6 and consists of three main components:

1. Repository
2. Target Database
3. Template Management Application

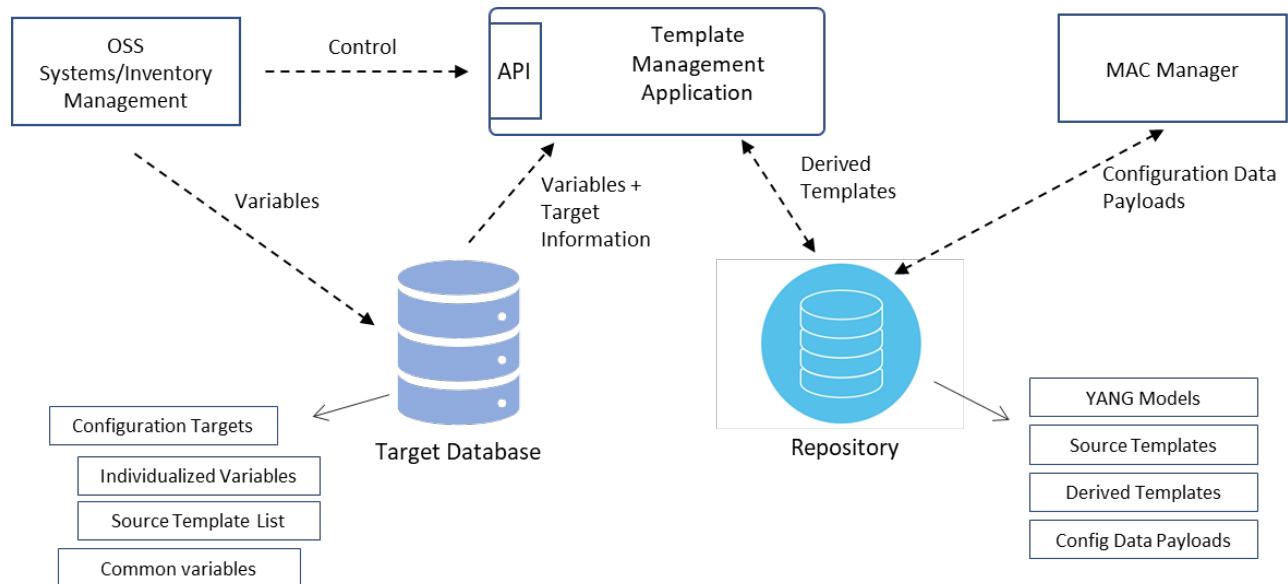


Figure 6 – Template-Oriented Configuration Management System

5.1. Repository

The Repository serves as a storage location for files which house YANG models, source and derived configuration templates, as well as the fully instantiated configuration payloads. The Repository can be implemented on top of a distributed Version Control System (VCS) such as Git or on top of a versioned object storage system such as Amazon S3. However, a distributed VCS provides the features required for a repository of templates. Such features incorporate sophisticated versioning tracking, management and security controls, including user access controls and what is most important, built-in integration for automated CI/CD pipelines. Therefore, the paper considers a VCS such as GitLab or GitHub as a default choice for the Repository.

Fully instantiated configuration payloads from the Repository are communicated to the FMA MAC Manager through a standardized API, e.g., RESTCONF.

5.2. Target Database

The Target Database contains information necessary for rendering configuration targets. The Target Database can be implemented as a relational database, such as an open-sourced Redis database.

This section defines the types of supported configuration targets. The TOC System operates on two categories of configuration targets: Derived Configuration Templates and Configuration Data Payloads.

Derived Configuration Templates are templates which the TOC System generates as the product of rendering of one or more source configuration templates. For example, a derived template representing the configuration of a downstream RF Port of a MAC-NE can be created from multiple source configuration templates where each source template contains configuration payload for one type of channel, for example, a list of downstream SC-QAM and OFDM channels and forward SCTE-55-1 OOB channels.

Another supported type of configuration target, the device Configuration Data Payload represents a portion of, or the entire configuration data instance intended for the target device. The Configuration Data Payload format is essentially the same as the Derived Configuration Template except all variables have been replaced with values, and all expressions have been fully evaluated. Configuration Data Payload must strictly conform to the branch of the YANG schema it represents and its intended use case.

The information maintained in the Target Database includes the following attributes:

- The target identification.
 - The name of the configuration target.
 - The location of the target file within the Repository.
 - Necessary metadata for creation of target payload, including YANG model information.
 - Target revision information.
- A list of source templates from which the target needs to render.
 - A substitution schema root that are optional for each source template. As mentioned earlier, each source template maintains the template-schema-root attribute, which can be overridden by the substitution schema root from the target.
- The YANG schema root for the target, in the form of the YANG path to the root node in the target.
- A list of variable replacement values for the target. When variable values are specific to exactly one target, then the database record of the target contains the values. In other cases, when variable values are shared by multiple targets, the Target Database record of the target provides an indirect reference to the value maintained in a common, shared record. This way the Target Database contains only one record with definition of the value, the single source of truth for multiple targets. The set of variable replacement values can grow quite large, when variables are leaves of multi-dimensional configuration lists.

Most of the information in the Target Database originates from the Cable Operator's BSS/OSS systems, while the information about the relationship between source templates and targets comes from human-assisted YANG schema breakdown into template hierarchy.

5.3. Template Management Application

The third component of the TOC System is the Template Management Application (TMA).

TMA implements all necessary configuration manipulation logic, including the rendering and the validation of target Configuration Data Payloads. The TMA also incorporates machine-to-machine APIs to the BSS/OSS systems through which the processes of rendering and validation of configuration payloads can be automatically initiated, and the status of these processes can be communicated back to the requestor.

5.4. Template Rendering

Template rendering is a procedure for creating target configuration payloads in accordance with their YANG models. It can be summarized as a five-step process:

1. The TMA retrieves the target information from the Target Database.
2. The TMA extracts the payload from all source template payloads identified in step 1.
3. The source payloads are combined in compliance with the applicable YANG model to create the target data payload.

4. The TMA injects the individualized target data into the payload, replacing variables with corresponding values from the Target Database.
5. The data payload of the derived template is written into the Repository.

The product of the rendering procedure is the target payload, which can be in the form of a derived template or a fully instantiated Configuration Data Payload. Next, the TMA validates the product, the target configuration payload against the YANG model and custom validation rules.

The combined source payloads typically form different branches of YANG instance data tree in the target's payload. However, the rendering process can also combine partial datasets from source template payloads that are part of the same branch.

A derived template can inherit selected variables from its source templates. Other variables can be replaced with actual values from the Target Database record associated with the target template. The TOC System decides whether a variable is replaced or inherited by examining the content of the Target Database record for the target. When the Target Database record contains a value for the variable, the variable is exchanged for the value. Otherwise, the variable is inherited by the target in the same form as in the source template.

Similarly, selected expressions can be evaluated during rendering of the source templates while other expressions are inherited by the Derived Template in the same form as present in the source template.

The rendering process can operate in multiple iterations. A Derived Configuration Template created because of rendering, can later become a source template in another iteration of the rendering process. The iterative process is particularly useful in creation of multi-level hierarchies of configuration data from simpler, more manageable, and “flat” templates. For example, the current FMA MAC-NE YANG model places many configuration attributes as high as at the fourth or even the fifth level of the YANG schema tree.

In the example shown on Figure 7, two source templates, named T1 and T2 are rendered into a derived template, named T3. The data payloads of T1 and T2 are combined in compliance with the YANG model. Next, Variable-1 of source template T1 and Variable-2 of source template T2 are replaced with Value-7 and Value-6 respectively. Variable-3 is inherited by derived template T3 from source template T2.

Note, that the payload from source template T1 is inserted at the same level as data in Container Z from source template T2 based on the T1's schema root.

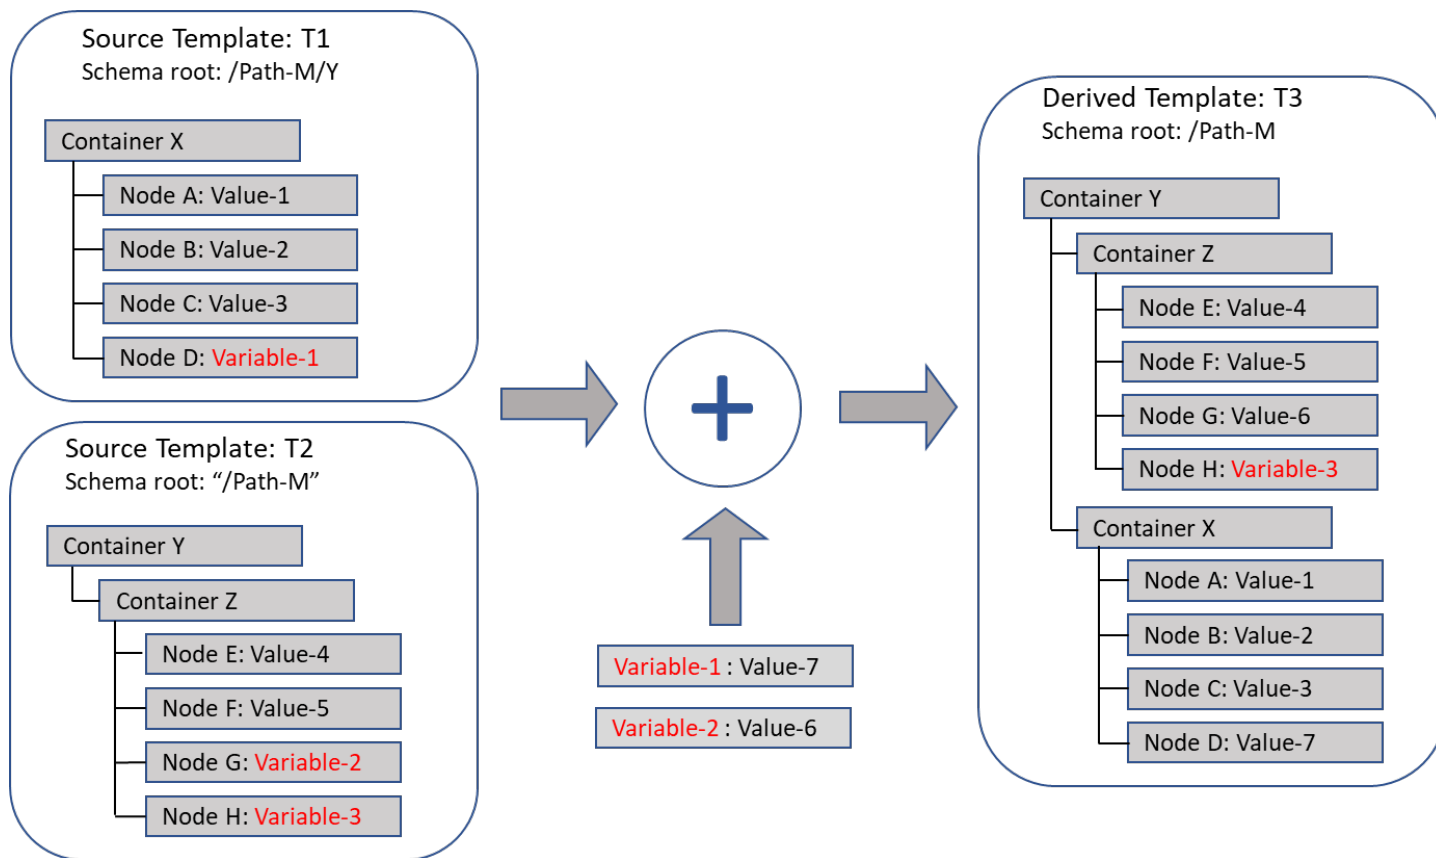


Figure 7 – Template Rendering Example

The instantiated data held in the template payloads, as well as the rules for transforming templates are driven by the YANG model to which the data is compliant. The limitations of the YANG modelling language also limit the TOC System's ability to transform the modelled data in certain, rather rare cases. For example, the data for a user-ordered list cannot be rendered from multiple source templates. The YANG models simply do not contain sufficient information to determine the order of data.

6. Discussion of Selected Use Cases

The following section provides examples of use cases to illustrate and highlight selected aspects of the proposed framework and benefits to Cable Operators.

6.1. Modular Data Design

One of the key advantages offered by the proposed templating methodology is the enablement of modular data design. The YANG-based data model of a MAC-NE can be divided into a set of small-scale segments, where each such segment corresponds to an integral functional area of the MAC-NE. The model break-up naturally follows the hierarchy of the MAC-NE YANG schema tree.

How might this work in practice?

The Cable Operator creates as many templates for each segment of the model as practically necessary to cover all variations of configurations needed for deployment and ongoing management. Several templates

providing configuration prototypes for defined deployment cases in a certain functional area can form a library of fundamental templates. An appropriately crafted set of fundamental template libraries can cover the entire data model of a device.

Figure 8 illustrates how such approach operates with an example based on three libraries of fundamental templates: the RF Subsystem, the Cable Bundle and the PTP Subsystem. Only three fundamental libraries are included in the example for brevity.

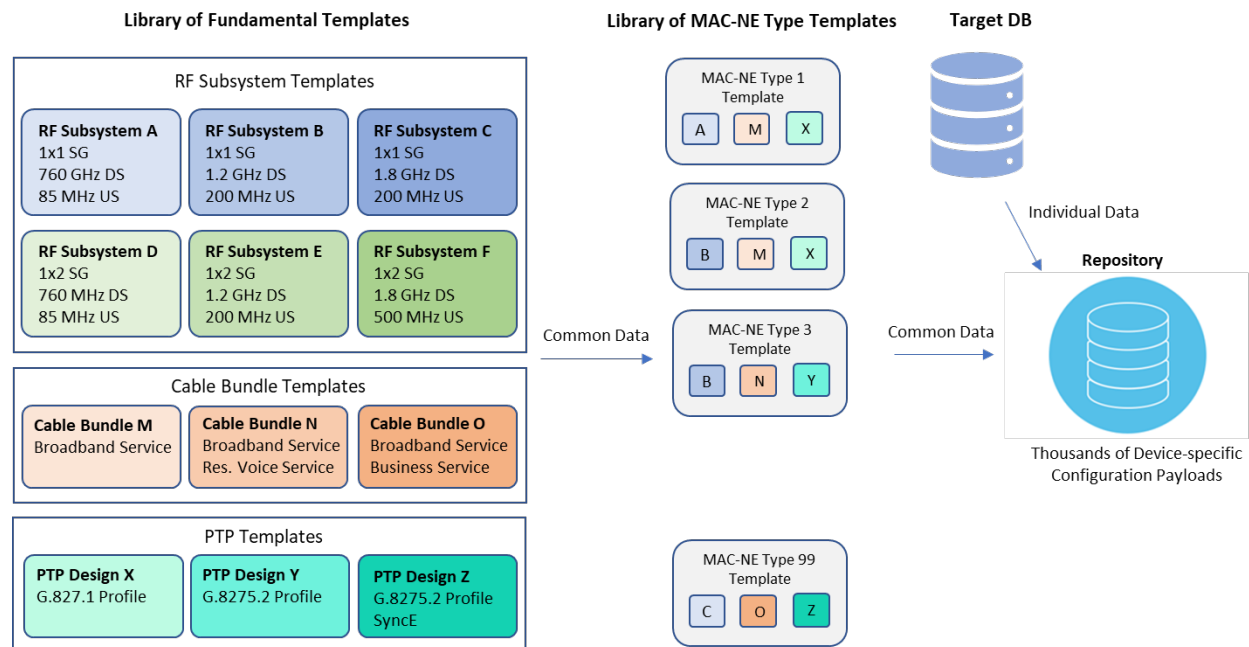


Figure 8 - Libraries of Fundamental Templates

The RF Subsystem template library includes templates for six designs of the HFC Plant. The PTP Subsystem template library includes three templates for commonly deployed PTP profiles. The Cable Bundle template library provides three templates for end-user IP service deployment options.

Next, from the library of fundamental templates, the TOC System renders a larger number of derived templates corresponding to as many variations of MAC-NEs as required. The intermediate library of MAC-NE Type templates is shown in the center of the diagram. Each MAC-NE Type template payload consists of selected payloads from the fundamental libraries.

In the subsequent step, the TOC System injects individual data into a selected MAC-NE Type template to create as many device-specific configuration payloads as needed. The device Configuration Data Payloads are made available to the MAC Manager to load onto the deployed RMDs.

6.2. Subdivision of Responsibilities within the Organization

Dividing work tasks between organizational teams is an essential part of running any business. This is particularly important for Cable Operators' organizations where separate operations sub-departments are often responsible for management of service and equipment lifecycle in functional areas under their jurisdiction. The examples of such functional areas include the DOCSIS[®] Subsystem, the MPEG Video Subsystem, Voice Services, Timing Operations or Subscriber Services.

The proposed framework, through the modularity in configuration data design, also enables the division of responsibilities between sub-departments within the Cable Operator's organization. The personnel of any department can maintain the sole ownership of the set of templates that correspond to the departmental responsibilities. These assignments can be enforced by effective access controls. For example, only the personnel of the department responsible for MPEG video service can be given the ability to modify (write-access to) the templates with configuration of MPEG video channels, DSG and out-of-band services.

6.3. Configuration Lifecycle Management

The flexibility and modularity offered by the Template-Oriented Configuration Management System simplifies many configuration lifecycle management tasks. One specific such task is the process of configuration modification within the template-oriented CMS.

Any changes to the configuration of a device due to network design changes, a software upgrade or service changes are typically contained within a small portion of the devices' configuration set. With proper template design, most configuration modifications typically affect only a single template, or a limited set of target variable values kept in the Target Database.

An administrator can modify a template by one of three methods.

1. The desired modifications are made directly to the payload of the affected template, resulting in a new revision of the template.
2. If the Repository already maintains a template with the desired configuration set, the name of the new template replaces the "old template" name in those Target Database records that need to be affected by the change.
3. The administrator creates the new template with a new name, often by cloning and modifying the previously used template. Then, the administrator modifies the Target Database as in Step 2 above.

After necessary modification to the templates and committing the changes to the Repository, the iterative rendering process predictably propagates the changes to the set of templates derived from the modified template and finally to the set of device Configuration Data Payloads.

7. Conclusion

In recent years the cable industry has embraced Distribute Access Architecture (DAA) including Flexible MAC Architecture (FMA) for innovation and significant investments. The paper augments these cable architecture transition efforts by examining the configuration aspects of DAA deployments and demonstrating how Cable Operators can effectively leverage YANG-based Configuration Template Methodology to scale, simplify and automate system configuration tasks within their back-office systems. The business outcomes are in the form of modular, automated, and streamlined provisioning systems and processes with increased agility and reduced OPEX.

Acknowledgements

We would like to sincerely thank our colleagues who made the writing of this paper possible, especially Brian Hedstrom of OAM Technology Consulting LLC for diligent review and mindful improvement suggestions.

Abbreviations

API	Application Programming Interface
BSS/OSS	Business Support System / Operations Support System
CCAP	Converged Cable Access Platform
CI/CD	Continuous Integration / Continuous Development
CLI	Command Line Interface
CMS	Configuration Management System
DSG	DOCSIS Set-top Gateway
FMA	Flexible MAC Architecture
IETF	Internet Engineering Task Force
JSON	Java Script Object Notation
MAC-NE	media access control network element
MMI	MAC Manager to MAC-NE Interface
MPEG	Motion Pictures Experts Group
OFDM	Orthogonal Frequency Division Multiplexing
OOB	out of band
PTP	Precision Time Protocol
RF	Radio Frequency
RFC	Request For Comments
RMD	Remote MAC Device
SC-QAM	Single Carrier Quadrature Amplitude Modulation
TMA	Template Management Application
TOC	Template-Oriented Configuration Management
VCS	Version Control System
XML	Extensible Markup Language
YAML	Originally: Yet Another Markup Language, recently: YAML Ain't Markup Language
YANG	Yet Another Next Generation
YCT	YANG-based Configuration Template

Bibliography & References

[FMA-MMI]	CableLabs FMA MAC Manager Interface Specification, CM-SP-FMA-MMI-I03-220126.
[FMA-OSSI]	CableLabs FMA OSS Interface Specification, CM-SP-FMA-OSSI-I02-220602.
[FMA-SYS]	CableLabs Flexible MAC Architecture System Specification, CM-SP-FMA-SYS-I03-220126.
[RFC 6020]	IETF RFC 6020, YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF), October 2010.
[RFC 7950]	IETF RFC 7950, The YANG 1.1 Data Modeling Language, August 2016.
[RFC 7951]	IETF RFC 7951, JSON Encoding of Data Modeled with YANG, August 2016.
[RFC 7952]	IETF RFC 7952, Defining and Using Metadata with YANG, August 2016.
[RFC 8040]	IETF RFC 8040, RESTCONF Protocol, January 2017.
[RFC 9195]	IETF RFC 9195, A File Format for YANG Instance Data, February 2022.

The Speed Triangle

speed promises, growing traffic and capacity in balance

An Operational Practice prepared for SCTE by

Robert-Jan van Minnen
Long Range Planning
Liberty Global – Liberty Tech
Boeing Avenue 53
1119 PE Schiphol-Rijk
rvminnen@libertyglobal.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. The 'ks' formula.....	5
3. The expectation of 'speed'	6
4. Utilization in a statistical perspective.....	7
4.1. Definition	7
4.2. The width of the bell curve	8
4.3. The symmetry of the bell curve	8
4.4. The bell curve in reality	9
4.5. The cumulative curve	10
5. Planning capacity	11
5.1. Basic calculation	11
5.2. Calculating a speed upgrade	11
5.3. No standard capacity across the network	12
5.4. Efficiency gain	12
6. Virtual measurement	12
6.1. Customers reaching the median speed	12
6.2. Customers reaching the minimum speed	13
6.3. Median speed room	13
6.4. Median speed in real life	13
7. Congestion	14
7.1. What is traffic congestion	15
7.2. Congestion versus ks.....	15
7.3. Never 100%.....	15
8. Traffic events.....	15
8.1. Many video streams	16
8.2. Big downloads	17
8.3. Speed tests	18
9. Long term outlook.....	19
9.1. XGS-PON.....	19
9.2. Service group size.....	19
10. Conclusions.....	20
Abbreviations	21
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 – theoretical distribution of momentary utilization.....	7
Figure 2 – examples of busy and quiet service groups	8
Figure 3 – examples of asymmetrical distribution	8
Figure 4 – long term utilization distribution of an entire network.....	9
Figure 5 – short-term utilization of a service group.....	10
Figure 6 –cumulative momentary utilization.....	10
Figure 7 –s and the percentage of success	11
Figure 8 –virtually measured speed compared with actual test results	14

Figure 9 –growing traffic surges..... 16
 Figure 10 –impact of video event..... 17
 Figure 11 –impact of simultaneous software download..... 17

List of Tables

Title	Page Number
Table 1 – ks formula definitions	5
Table 2 - overlap and failure of random speed tests	18
Table 3 - DOCSIS4.0 and XGS-PON example calculations.....	19


1. Introduction

The ‘speed’ of internet is typically conceived as the main indicator of the quality of our internet connection. Operators have been using the available top speed as a selling argument. For good reasons because cable could deliver a lot.

As usage grew the race for higher speeds was on. The race could be fueled by adding more and more capacity. But there are limits as spectrum is not free and expansion needs time and investment. Fiber can provide the next order of magnitude but doesn’t come for free as well and over time will also require upgrades.

This paper focuses on the question how we can optimize the capacity we deploy in relation to the growing traffic and speed promise.

A simple formula is proposed, tying capacity together with traffic and speed performance. It is an evolution on the k-factor (re. Tom Cloonan, CommScope). This formula is translated into practical applications to plan into the future and to (virtually) measure the actuals of today.



The aim is to support an optimized business balance of the three corners in a measurable way using only the tools we typically already have.

Possible caveats and the current perception of (avoiding) congestion will be discussed as well as practical proof of the results.

Finally, the formula will be applied to possible futures with extreme speeds on DOCSIS or XGS-PON.

2. The 'ks' formula

The proposed formula is essence a combination of two classical methods to determine the amount of capacity needed in a service group (called 'capacity').

- 50% rule; capacity > the highest speed sold / 50%
- k-factor; capacity > utilization + k * the highest speed sold

The 50% rule is extremely simple to use and can be standardize throughout a network, but it discards the actual traffic. As a result, on busy service groups the speed performance is suboptimal. Even up to the point where congestion occurs. In practice an additional planning rule is needed to avoid this. In service groups that are quiet there is more capacity than needed which means inefficiency.

The k-factor must be applied considering the utilization of each service group individually but only provides a factor to the speed. Analyses of the statistical distribution of traffic revealed however that the chance of having sufficient available capacity for a burst of speed is a function of the utilization and not of the planned speed. This means the formula needed an additional factor to modulate the utilization which has been called the 's'. This stands for service level or safety. The result is the ks formula:

$$\text{capacity} = \text{utilization} * s + \text{speed} * k$$

Table 1 – ks formula definitions

capacity:	the amount of capacity that is minimally needed [Mbps]
utilization:	the average utilization during the interval that the speed should be reached (e.g. 8-10 pm) [Mbps]
s:	the service level factor to use
speed:	the top speed that is desired [Mbps]
k:	the factor that defines how much of the top speed we want to enable

3. The expectation of ‘speed’

The customer perception of speed is influenced by many factors of which some subjective. For network planning a translation into objective parameters is required. In some markets regulators have defined measures to adhere to. Operators may also have defined some themselves. For example:

- *at least 50% of customers at peak time need to achieve headline speed*
- *all customers need to achieve 50% of headline speed at peak*

This in turn requires a definition of the interval for which the rules apply. For example

- Between 8 and 10 pm
- During the busiest hour of the week

It also needs a translation of achievable speed into a network parameter. What a customer would see from a speed test depends on for example:

- The available or unused network capacity during the test interval (typically 10-20 sec)
- Limitations in the customer equipment
- The speed test mechanism

This leads to a basic definition of what ‘speed room’ is from a network perspective:

- **The amount of capacity in a service group that is available for a burst of traffic**
 - At any given time during the busy hours between 8 and 10 pm
 - For the duration of ten seconds

It is recommended to define ‘speed’ in this perspective with care since it directly relates to desired performance.

A limited increase of the factor ‘k’ in the formula can be used to compensate for other factors such as:

- Chance of overlapping speed tests from different users at the same time
- The provisioned speed. Typically, modems are provisioned 5-15% higher than the sold speed. This allows a limited amount of ‘catching up’ when during a moment in the ten second test there is not enough capacity available.

4. Utilization in a statistical perspective

4.1. Definition

The speed room is defined by the capacity of a service group minus the actual utilization. While the capacity is a constant, the actual utilization in a service group is varying from moment to moment. This calls for a deep dive into this fluctuation.

When utilization is measured it is typically:

the total amount of bits transported during a sample and expressed in Megabits per second [Mbps]

In regular measurement tools the sample time varies from five minutes to one week. For an entire network this yields a vast amount of data. Normally this data is aggregated to averages over a certain amount of time and 'peak' levels.

But if the aim is to measure ten second intervals for speed room calculations the challenge will become soon just to handle this data. Since traffic appears to be generated in a semi-random pattern, the working assumption is that it can be described as a bell curve.




Figure 1 – theoretical distribution of momentary utilization

With this assumption only a few parameters are sufficient to determine the utilization during the interval and the chances of having sufficient speed room.

- Average utilization
- Median utilization (50th percentile)
- Width of the curve (measured at 20th, 80th percentiles)

In this example the average utilization is 60%. Because the bell curve is symmetrical, 50% of the samples have a higher utilization, the other 50% lower. The symmetry of the curve implies that:

median utilization = average utilization

4.2. The width of the bell curve

In practice the expected bell curve may be differing between service groups and can also change over time. It may be narrow when many users are active and wider in more relaxed service groups.




Figure 2 – examples of busy and quiet service groups

The duration of the samples has an impact on the width of the curve. Short samples will give a wider curve. The longer the sample are, the narrower the curve will be. However,

for the median speed room, it is not relevant how wide a symmetrical distribution is

4.3. The symmetry of the bell curve

So far it has been assumed that the curve is symmetrical. What if it is not?

These two examples of a skewed distribution visualize a non-symmetric distribution.




Figure 3 – examples of asymmetrical distribution

Calculations on extreme examples have revealed that the median utilization can differ up to 10% from the average utilization. In terms of the formula this means that the median speed is reached when

$$0.9 < s < 1.1$$

This brings in uncertainty in the result. In the left example, the median will be 8% lower than the average. This means the results will be worse than expected if $s=1$ is used. In the right situation we the median is 6% higher than the average. This can be compensated by using $s' = 1.08$ (left) / 0.94 (right).

The uncertainty caused by the translation of average utilization into median can be overcome by:

- Measuring the median utilization in ten second samples which may give additional complexity
- Use a default $s' = 1.1$ and accept a possible overestimation of the traffic
- Calibrate with actual speed test results

4.4. The bell curve in reality

Conversion of actual utilization metrics into a distribution chart can be quite challenging. Part of the challenge is the fact that besides the randomness there are also diurnal, seasonal and long-term trends. On top there are single events such as big downloads from software releases, video events and even the pandemic.

The working assumption has been tested in various set-ups. Regardless of the sample time, network and time of day, a histogram yields a bell-type curve.

In this example of a long period of two years where the busiest hour of each week is shown, the distribution is narrow. The median is very close to the average (0.3% difference).




Figure 4 – long term utilization distribution of an entire network

The following example shows the correlation between ten second- and fifteen-minute samples. Both yield a bell-curve which follows the diurnal and weekly trend and correlate strongly. Because it is measured over two weeks, the results include busy and quiet times.




Figure 5 – short-term utilization of a service group

4.5. The cumulative curve

The distribution from figure 1 can be shown as a cumulative curve. This shows for each level of relative utilization which percentage of time the actuals will be lower. If we depict a speed as a percentage of the service group capacity (e.g. 1 Gbps in a 2Gbps service group is $\frac{1}{2} = 50\%$) we can read the percentage of time the utilization samples will be lower and a speed test would be successful.




Figure 6 –cumulative momentary utilization

In this example the average utilization is 60%. It is a symmetrical distribution, so the median utilization is the same. If this is a 2Gbps service group, the chance of a successful speed test for 1Gbps is only 20%. This is an example that the traditional planning with a 50% rule will give a relatively low performance in a busy service group.

Through mathematical manipulation using the ks formula, any distribution (measured or assumed) can be altered to show which percentage of time success belongs to which s if we measure or calculate it.




Figure 7 –s and the percentage of success

Following the example above, the 1Gbps top speed has an $s = (2 - 1 \text{ Gbps}) / (60\% * 2 \text{ Gbps}) = 0.83$. The corresponding % of time success is 20% as is also found in figure 5.

Note that this is just another representation of the same distribution data. It is useful for:

- Determination of other percentiles than the median
- Sensitivity analyses (what if the growth is higher than expected)
- Reconstruction of the distribution from actual speed test data

5. Planning capacity

5.1. Basic calculation

Calculation of required capacity in a service group to obtain a certain speed room has become simple.

$$\text{capacity} = \text{utilization} * s + \text{speed} * k$$

Example 1: for a service group with an average utilization of 400 Mbps during peak hours and a top speed of 500 Mbps, the minimum capacity = $400 * 1 + 500 * 1 = 900$ Mbps. In that case the median speed room will be the top speed.

Example 2: for a service group with an average utilization of 250 Mbps during peak hours and a top speed of 300 Mbps, 80% of speed tests must be successful. To achieve an 80th percentile of success, $s = 1.25$ must be used. This value is obtained from the cumulative distribution in the previous chapter. The minimum capacity = $250 * 1.25 + 300 * 1 = 625$ Mbps. In that case the 80th percentile of speed tests will be equal to the top speed.

5.2. Calculating a speed upgrade

When the top speed of a network is upgraded the capacity must be expanded accordingly.

As an example, the top speed is upgraded from 1Gbps to 2Gbps. In the traditional calculation with a 50% rule this would require $1\text{Gbps} / 50\% = 2\text{Gbps}$ of additional capacity. The ks approach requires the addition of $k * 1\text{Gbps}$ of capacity. Supposed that $k = 1.2$ this is 1.2Gbps of additional capacity. Since the term $s * \text{utilization}$ remains does not change, the performance of the 2Gbps product after the upgrade will

be the same as the 1Gbps product before. Instead of 2Gbps additional capacity, only 1.2Gbps is added with the same performance as before the upgrade!

5.3. No standard capacity across the network

In essence the required capacity is different for every service group because the utilization is different. This could pose an operational challenge. In practice it will be desirable to have some degree of standardization. A method can be to classify all service groups on their utilization and/or spectrum availability. Then plan for the busiest in each group. This ensures sufficient capacity while avoiding too much idle capacity, unnecessary spectrum occupation or node splits.

5.4. Efficiency gain

Compared to a standardized capacity across all service groups, a deployment based on the ks-formula gives an efficiency gain. The amount depends of course on the distribution of busy and quiet areas and the already used standard capacity. In an example busy network the actuals were compared with the requirements following the ks formula. In the quieter areas 17.5% less capacity was needed. In the busiest areas (4.4% of service groups) the standard capacity was insufficient. To reach the median speed as defined 0.2% of additional capacity would be needed. This would give a net gain of ~ 17%, provided there is spectrum available. If this is not available this should drive segmentations and cost 4.4% more capacity. Still a net gain of ~13% and ensured speed performance across the network.

6. Virtual measurement

The formula has a mathematical form and can be altered to virtually measure both s and k from real life data. This can be used to verify the calculation or even manage the actual operation of a network.

There are two reporting options which are explained here:

- Customers reaching the defined speed
- Median available speed room

6.1. Customers reaching the median speed

In essence this is counting the customers that are on service groups with $s \geq 1$.

In an existing network the capacity, utilization and sold top speed are known. This means if $k=1$ (or another value of k that is desired) s can be calculated for each service group. The ks formula is transformed into:

$$s = \frac{\text{capacity} - \text{speed} * k}{\text{utilization}}$$

The resulting dimensionless s provides information on the propensity of reaching the top speed for a random speed test as explained. For $s = 1$ this is the median result.

When all customers are counted that are connected to a service group that has $s \geq 1$ and divided by the total customer count we have an indication of the percentage of customers that have sufficient speed room.

6.2. Customers reaching the minimum speed

When there is a minimum speed guarantee the calculation can be repeated with a different s . In the example the rule applies for the busiest hour. This means a second s_{\max} must be calculated for the utilization at the busiest hour which of course must be available.

A 100% guarantee is not possible, in the example 95% is chosen. So 95% of random tests must have a minimum of 50% of speed room available. The graph in figure 6 shows that the related $s = 1.7$

Again, the customers are counted. This time from service groups with $s_{\max} \geq 1.7$.

6.3. Median speed room

To calculate the median speed room the s in the formula is fixed: $s = 1$. Another transformation of the k s formula yields:

$$k = \frac{\text{capacity} - \text{utilization} * s}{\text{top speed}}$$

The resulting k gives the speed room / top speed ratio. The actual median speed room is then the found with $k * \text{top speed}$. This value is not directly usable to report on a customer level because CPE are provisioned for a certain top speed. Normally 5-15% more than the sold top speed. This means the median top speed from network perspective must be limited to this value to obtain a value that is achievable for customers. If the overprovisioning is 10%, all results above 110% must be capped at this level before calculating the average for a network.

6.4. Median speed in real life

The virtual measurement of the median speed has been put to the test in a multi-million homes network with several thousands of SamKnows' probes. The virtual measurement theory dictates that if $k \geq 1$, the median speed room equals the top speed. With $k < 1$, the median speed goes proportionally down.

The results from hourly tests from ~1.500 probes are compared here with the results from the virtual measurement with network data as described in section 6.3. The expected relation is clearly visible and the found k to meet the median speed in practice is ~ 1.2 for downstream and ~ 1 for upstream.




Figure 8 –virtually measured speed compared with actual test results

7. Congestion

Network congestion is a thing operators want to avoid. The classical approach is to provide sufficient capacity to keep the utilization under a certain percentage. The knowledge from the ks approach can be used to create a different view on congestion and discover opportunities for improvement.

7.1. What is traffic congestion

First step is to define what congestion is and what should be avoided. For this general road traffic theory is used which defines three stages of traffic flow:

1. Free flow of traffic. There are no obstacles to reduce the speed.
2. Viscous flow of traffic. It is busy and the speed reduces on occasions but still flowing.
3. Jammed traffic. Speeds drop to zero on occasions. This leads to a build-up of even more traffic and longer stand-stills

7.2. Congestion versus ks

The three stages of traffic flow can be defined as speed rules. For example:

1. Free flow: the median speed room is achieved at evening peak hours
2. Viscous: at the busiest hour at least 50% of the top speed must be available in 95% of cases
3. Jammed: the speed room is < 10% for > 1% of time

The cumulative traffic curve(s) in use can be used to find the related s for each stage. These values of s can be put in the formula to add additional rules for planning. At the same time the actual network data can be virtually measured to determine the percentage of customers in each stage.

7.3. Never 100%

The stochastic behavior of traffic predicts that at any time short peaks of traffic will utilize the full capacity, even in quiet service groups. If the samples are short enough, we will 'hit the ceiling' sometimes. This is not a problem as there are buffers and IP has ways of resolving lost packets. Customer impact is expected however when the ceiling is hit too long or often. This is exactly what the cumulative distribution curve can be used for to detect.

8. Traffic events

The ks formula works from the assumption that traffic has a semi-random character. The practical tests have confirmed this but in recent years we have seen an increasing number of events that can change the traffic volume suddenly.

The graph below shows the surges of the combined Liberty Global network. The percentage is relative to the expected volume and compensated for normal growth and seasonal fluctuation. In other words, a sudden spike in traffic for one or more hours.




Figure 9 –growing traffic surges

Qualitative analysis showed that the surges are mainly caused by:

- Video broadcast over IP of popular games
- Releases of large software updates Upstream surges from:
- Work and school from home

Because the events are not always plannable, may last short and data collection is complex there is to date insufficient data to determine the impact on the cumulative distribution. A desk examination on two examples can give some hints on what to expect.

8.1. Many video streams

If many customers are watching a video stream, the traffic per customer is relatively low but the number of users is high. For game events the beginning and end are approximately the same for every user. This means a constant amount of traffic is added during the event with little variation. Note that a significant amount of the video traffic will replace regular traffic so the net addition will be less than calculated by the number of users * traffic per user. This ‘replacement’ is estimated at ~50% of the calculated amount.

If viewed from a ten second sample perspective, the distribution curve shifts to the right during the event. Average and median will also shift with the same amount. To maintain the performance this traffic must be included in the planned capacity as discussed.

If it is not included, the performance will go down. This is shown in the graphs below. The s-curve shows the chance of success according to the initially planned traffic. If we planned for an s=1 and temporarily include 20% of traffic, the percentage of successful speed tests will drop from 50% (median) to 25%.




Figure 10 –impact of video event

8.2. Big downloads

An example is used from an observed surge of 25% linked to a game update of 40 GB. The amount of data corresponds with ~1.3% of all customers downloading the update in the same hour.

A 1Gbps customer would be able to download the update in ~5.3 minutes, a 500Mbps customer 11 minutes and for 100Mbps it would take 53 minutes. If these downloads start at the same time, for the first five minutes the traffic surges +78% after which the surge drops in size. If all customers would have a top speed of 1Gbps, for five minutes the traffic would increase by 281%.

The graph shows the hypothetical traffic increase if all downloads start at the same time and customers have a top speed as they are normally distributed.




Figure 11 –impact of simultaneous software download

In reality traffic increased with 25% for the full hour and beyond.

- downloads do not start at the same time (planned)
- peering or IP network capacity is limited
- access capacity is temporarily at its limit

Chances are that the limit of capacity in terms of this document is reached for some time during the downloading. As a result, the speed room reduces to zero during those times.

Simply adding 25% of traffic in the formula (multiply s with 1.25) may not be sufficient to cater for these events. The reason is that the surge would become higher but still the ceiling would still be hit for some time, only shorter. It is a matter of choice to what extent the speed room must be available during these releases. Managing the release times may be more effective and cheaper option. Or accepting that the performance is not compliant for the short duration of these events.

8.3. Speed tests

Performing actual speed tests from the customer's equipment is of course the ultimate proof of speed performance. For testing the network speed SamKnows probes in modems are available. This test rules out any limitation in the customer equipment and wiring. One could argue to use these systematically to have a reliable indicator of the network performance. But performing tests has an impact on the utilization of the network and in the end influences the performance for regular traffic.

Based on the average provisioned speeds an estimation of the additional traffic at peak hours can be made if all would perform a scheduled hourly test of ten seconds. For downstream the addition would be 19%, for upstream 30%. This means that just because of the proofing of the speed room, the capacity must be increased by these amounts to actually achieve the speed room.

If tests are performed randomly during the hour there is a chance they overlap. In case two top speed tests overlap it is likely both will fail as they compete for the same speed headroom. This chance of failure is the highest for the top speeds but can also happen with other speeds or if more than two overlap. In the example the chances of overlap are:

Table 2 - overlap and failure of random speed tests

speedtest for	overlaps with another for	overlap	failure
1000 Mbps	1000 Mbps	2%	4%
1000 Mbps	500 Mbps	13%	24%
all	all	53%	78%

The percentages indicate that a significant amount of speed tests would be impaired just because of the fact that the tests are performed. Similar percentages occur in the upstream direction.

This self-inflicted performance reduction can be mitigated by scheduling the tests so that they do not overlap. The feasibility of this scheduling is not investigated but assumed possible (SamKnows is supposed to avoid multiple tests at the same time on one node). In the example during 46 minutes of every hour there would be a downstream speed test going on (38 in upstream).

Implementing scheduled speed tests means adding capacity to the network to carry the added amount of traffic. If this is added following the ks formula, the statistical availability of speed room would still be as expected because both consider 10 second samples.

There could be however an impact on the customer perception for normal traffic since this is built up from much shorter bursts. To understand the impact in detail, the utilization distribution of shorter samples would be needed. This is not available to date.

9. Long term outlook

The ks formula works from metrics that are familiar from the DOCSIS world. In essence the metrics relate to properties of the aggregated internet traffic from customers, regardless of the technology. This implies that the same method can be used for other means of internet access.

9.1. XGS-PON

As an exercise the calculations are done for hypothetical DOCSIS4.0 and XGS-PON networks.

Table 3 - DOCSIS4.0 and XGS-PON example calculations

	D4.0 1.8/204		XGS-PON	
	Down	Up	Down	Up
capacity	12280	1228	10000	10000
traffic/customer	10	1	10	1
customers/service group	180		24	
traffic/service group	1800	180	240	24
for median (s = 1)				
speed room (k * top speed)	10480	1048	9760	9976
for k = 1.2				
ensured median top speed	8733	873	8133	8313
for an 8Gbps call out, s =	1.4	1.5	0.7	6.7
speed test success	81%	89%	20%	~100%
utilization	15%	15%	2%	0%

The expected traffic is subtracted from the available capacity to obtain the speed room. When divided by the k of choice (here k = 1.2), the ensured median speed room is obtained. In case the market call-out is 8Gbps (0.8 for upstream in DOCSIS), the resulting s is calculated ($s = (\text{capacity} - k * \text{speed}) / \text{utilization}$). With the sample s-curve from figure 2, the speed test success ratio is determined.

9.2. Service group size

Over time the traffic per customer grows and as a result the number of customers per service group reduces. Metcalfe's law predicts that with less users, the volatility of traffic increases. Investigations have confirmed this in practice. As a result, the cumulative distribution and s-curve will get steeper. This means that smaller variations in traffic or s will result in changed speed room results. For the median it is not relevant how wide the distribution is. But for a guaranteed minimum (95% of time minimum 50%

speed room) a higher s is needed, leading to more capacity. For long-term planning it is recommended to include this in the long term s factor.

10. Conclusions

The capacity that a network needs depends on the traffic and desired speed room. The semi-random behavior of traffic can be modelled with a certain accuracy with a simple formula:

$$\text{capacity} = \text{utilization} * s + \text{speed} * k$$

The formula is most accurate and easiest to implement when the performance is defined for the median speed room. It can be enhanced with multiple planning rules to suit the business or market regulations.

Application of the formula for planning provides the option to diversify the capacity and with that save ~15% of capacity.

The same model provides a measurement tool for speed performance based on regular network metrics. Refinement of the k and s variables can be performed with these virtual measurements or with actual speed test results when available.

Congestion avoidance can be refined because the formula can be used predictively and divide congestion in phases, like in road traffic theory.

Traffic events can disturb the traffic distribution for a short period of time. Video events can be predicted and modelled with the formula or accepted with a reported temporary lower speed room. Big downloads can have a significant and prolonged impact on the performance.

Speed tests are useful to verify the planning and virtual measurement but testing in high volumes should be avoided since the tests influence the result and increase the need for capacity.

The ks formula is a general tool, independent of technology or time. It is recommended for short- and long-term planning and quick speed performance evaluation.

Abbreviations and Definitions

service group	aggregation of channels to serve several customers
Capacity [Mbps]	Total capacity in a service group in Megabits per second
Megabits per second [Mbps]	Megabits per second [Mbps]
GB	Gigabyte
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

k-factor, T. Cloonan; CommScope

Metcalfe's law; Wikipedia

The World Is Changing And So Can You By Using Agile

A Technical Paper prepared for SCTE by

Aparna Srinivasan

Manager Cybersecurity

Comcast Cable

Chennai, India

Aparna_srinivasan@comcast.com

Venkata Ramarao Sanka

Development Engineer 4

Comcast Cable

Chennai, India

venkataramarao_sanka@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Challenges	4
3. What Is Agile Management?	5
3.1. The Pillars of Agile	6
3.2. Agile mindset.....	6
3.3. Agile Values [10]	7
3.4. Agile Principles [9].....	7
4. Planning Agile Implementation.....	8
4.1.1. Understand the culture and needs of the organization.	9
4.1.2. Identify what needs to be changed and what is not changing.	9
5. Plan for Agile Implementation in a non-Agile Organization	10
5.1. Product Vision	11
5.2. Product Roadmap	11
5.3. Create a Release Plan	11
5.3.1. Prioritize Requirements.....	11
5.3.2. Create an Initial Design.....	11
5.3.3. Determine the Success Factor.....	12
5.4. Who Was Involved in the Plan?	12
5.5. Prerequisites for Iteration Planning	12
5.6. Iteration Planning Steps	13
6. How To Implement Agile in An Organization That Is Not Agile	13
6.1. Design Sprint (Sprint 0).....	13
6.2. Project Charter	13
6.3. Scrum Ceremonies	14
6.3.1. Sprint Planning.....	14
6.3.2. Daily Scrum/Stand up	14
6.3.3. Sprint Review	14
6.3.4. Sprint Retrospective.....	15
6.4. Metrics for Agile.....	15
6.4.1. Productivity metrics.....	15
6.4.2. Progress metrics	15
6.4.3. Quality metrics	16
7. The Benefits of Agile Methodology	17
8. Conclusion.....	17
Abbreviations	17
Bibliography & References.....	18

List of Figures

Title	Page Number
Figure 1 – Agile methodology workflow	5
Figure 2 – Pillars of Agile	6
Figure 3 – What is Agile	7
Figure 4 – Difficulty of Agile adoption	8
Figure 5 – Understand organization culture.....	9
Figure 6 – Steps in Agile planning	10
Figure 7 – Create a release plan	11

Figure 8 – Scrum team..... 12
 Figure 9 – Scrum ceremonies..... 14
 Figure 10 – Sprint burndown..... 16

1. Introduction

In this changing business environment, it is crucial to take a proactive stance to be more agile and collaborative. One of the development teams in Comcast Cybersecurity adopted a waterfall methodology for their product development. They committed to understanding requirements, designing to those requirements, and asking teams to develop them once the requirement was signed off.

Agile is a popular methodology for developing products and services, but it hasn't been fully adopted in many organizations. More than ever, Agile has shown to be an effective way to create products and services with proper implementation. It is a culture that changes how teams work together, communicate with each other, and manage their work. It's more than just a new way of doing things. It is a whole new way of thinking about how to do things.

This paper and presentation will talk about the challenges faced by the team, provide a description of Agile, including the planning, implementation and benefits of using this approach. It will show how a transition to Agile delivered significant benefits to our customers.

2. Challenges

The cyber security landscape has been changing rapidly over the last few years, and it has become a constant challenge for organizations to keep up with these changes. There are often delays in the process of analysing customer requirements, often to the point of having them become obsolete before they are addressed. Market conditions often create changes to an initial requirement, which, if not addressed immediately will lead to dissatisfaction of the requesting party. Without an effective communication and feedback cycle, stakeholders are not included in initial discussions or informed of changes leading to a gap in development, stakeholder awareness, and further dissatisfaction. After development and during Testing or User Acceptance Testing can generate minor requirement changes; changes that can take a long time to implement, which might cause the real purpose of the requirement to be lost over time.

3. What Is Agile Management?




Figure 1 – Agile methodology workflow

Agile is an iterative approach to developing software products. It emphasizes the importance of early and continuous delivery of working software, which can be obtained through frequent releases, demonstrating progress early and frequently to stakeholders, and responding quickly to feedback (Figure 1 and Figure 2).

3.1. The Pillars of Agile

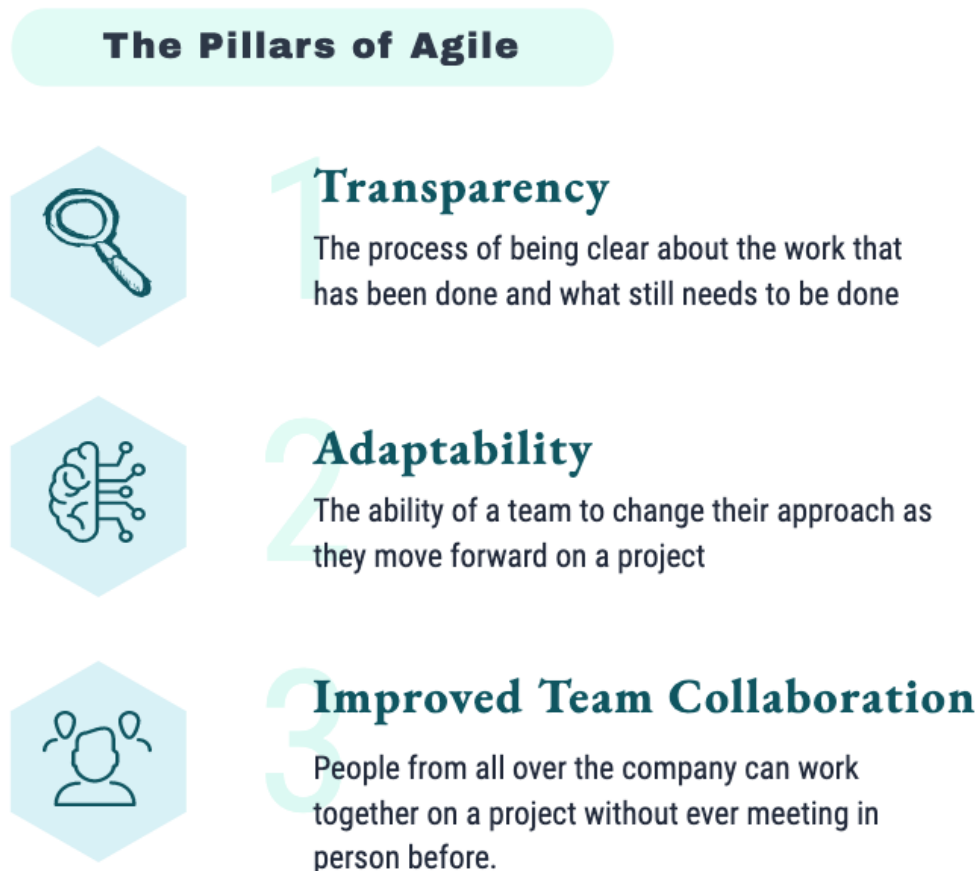


Figure 2 – Pillars of Agile

3.2. Agile mindset

To succeed with Agile methodology, teams must cultivate an Agile mindset. The Agile mindset is a thought process that involves understanding, collaborating, learning, and staying flexible to achieve high-performing results. [3] By combining the Agile mindset with methods and tools, teams can adapt to change and deliver incremental value to their customers. [3] Imagine a concert where each musician plays separate notes without coordination. This is not an impressive performance, and you will be disappointed by the end of the event. Take this case and apply the same scenario to your work environment. Concerts depend on melody, harmony, and rhythm to create a piece of soulful music. Similarly, software teams can successfully deliver the final product when they collaborate, help, share information, and work flexibly.

An Agile mindset (Figure 3) focuses on "being agile" as a foundation for success in "doing Agile." [3] It is defined by the four values and described by the twelve principles of the Agile Manifesto and then manifested through an unlimited number of practices and diverse ways of working. [3]

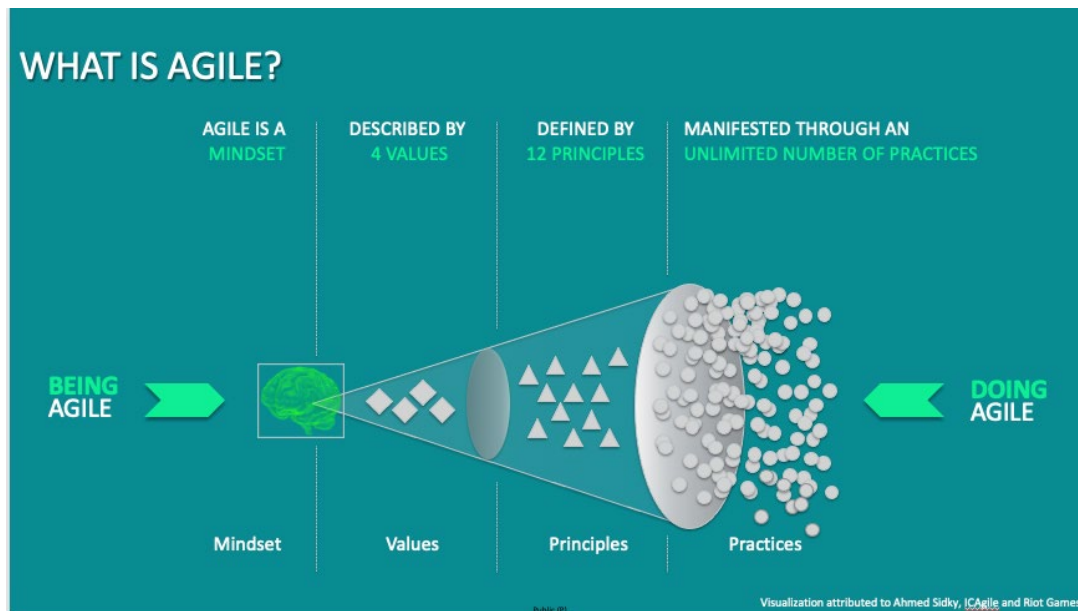


Figure 3 – What is Agile

3.3. Agile Values [10]

We are uncovering better ways of developing software by doing it and helping others do it. Through this work we have come to value:

- **Individuals and interactions** over processes and tools
- **Working software** over comprehensive documentation
- **Customer collaboration** over contract negotiation
- **Responding to change** over following a plan

3.4. Agile Principles [9]

- 1) Satisfy Customers Through Early & Continuous Delivery
- 2) Welcome Changing Requirements Even Late in the Project
- 3) Deliver Value Frequently
- 4) Break the Silos of Your Project
- 5) Build Projects Around Motivated Individuals
- 6) The Most Effective Way of Communication is Face-to-face
- 7) Working Software is the Primary Measure of Progress
- 8) Maintain a Sustainable Working Pace
- 9) Continuous Excellence Enhances Agility
- 10) Simplicity is Essential
- 11) Self-organizing Teams Generate Most Value
- 12) Regularly Reflect and Adjust Your Way of Work to Boost Effectiveness

4. Planning Agile Implementation

“Planning is everything, Plans are nothing”

-Prussian Field Marshal Helmuth Graf von Moltke

Estimating and planning are critical to the success of any software development project of any size or consequences (Figure 4). Plans guide our investment decision. A plan helps us know who needs to be available to work on a project during a given period. [13]

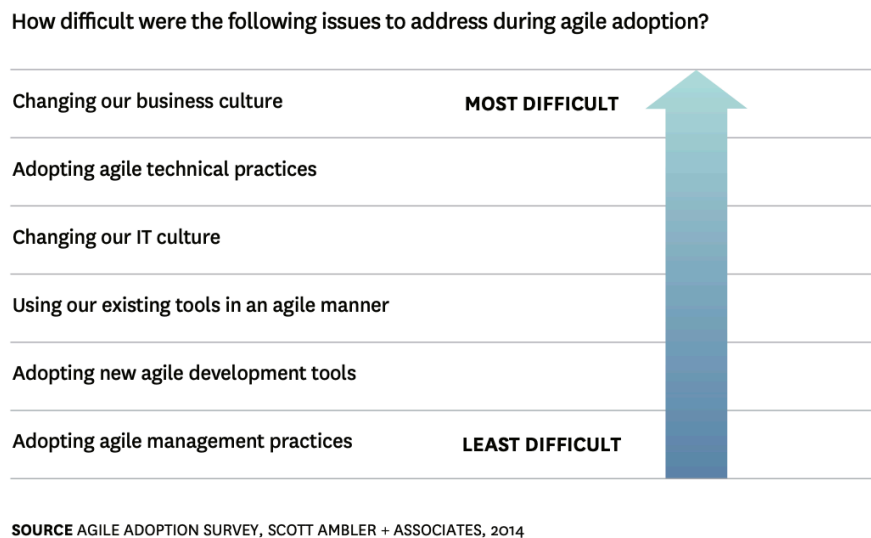


Figure 4 – Difficulty of Agile adoption

Plans are documents or figures; they are snapshots of how we believe a project might unfold over an uncertain future. Planning is an activity. Agile planning shifts the emphasis from the plan to the planning.

Planning for Agile balances work and expense, ensuring we will change the plan throughout our project to improve the outcome. An Agile method is one in which we're not only willing to but are very keen to alter. We want to be flexible because change indicates we've learned something or avoided making an error. We might have discovered that customers want more or less of a specific feature, have determined it is more crucial to have usability than anticipated, or that programming in the new language takes longer than expected. As part of project planning, the team's Agile implementation, we identified and prioritized the requirements, created an initial design, and determined how we'll measure success.

The following are some of the critical steps taken in the planning for Agile implementation in the Cybersecurity team:

4.1.1. Understand the culture and needs of the organization.



Figure 5 – Understand organization culture

Understanding the culture and needs of the team/organization is the key to success in Agile implementation (Figure 5). We realized that change and transformation are not easy. There will be some resistance, but it can be overcome with a proper understanding of the culture, needs, and goals.

The first step we took in understanding the culture of this was to learn about the different modules of their work on an elevated level (what they do), how they do their requirement gathering, design, development, testing (how they do it), identify their end customers, (for whom they do it), and why this development is needed (why do they do it). This helped all parties understand how Agile can fit into their existing framework.

The next step was to listen. Top management was open to hearing other perspectives, and they listened to what others wanted with agility and understood the impact. Once an Agile approach was approved, we identified points of leverage by asking how Agile could serve this project as effectively as possible.

4.1.2. Identify what needs to be changed and what is not changing.

For an organization/team to successfully transition from traditional project management techniques to Agile, it is essential to understand the difference between a change and a non-change, and how these factors affect how Agile implementation should be approached. The answer to this question is dependent on a few distinct factors.

In general, teams like ours which are not Agile need to make fundamental changes to their structure and how they work. We identified what processes need changing and how these processes will be improved.

Some of the key factors we considered when determining if the team needed to change its approach were the goals of the team, the percentage of available time the team spent on projects, and the number of people on the team.

It is essential to understand the different stakeholders' perspectives to communicate with stakeholders about the changes and challenges in Agile implementation. We began by identifying what our stakeholders wanted and expected from us. We then worked specifically on what they wanted, not what they didn't, and then began a process of open and continuous communication so they knew how their ideas were being implemented. This last item, communication, had been identified in our previous development model as an area that could be improved so it was important for us to address.

To ensure that we could identify all the stakeholders and involve them in the process, we created a plan to define the problem/opportunity and to identify the key stakeholders.

There are four main categories of people whom we involved in the change process:

1. those whom the change will directly impact
2. those who need to support the change
3. those who will have a stake in the success of the change
4. those who must manage or implement it

The plan for involving them included getting their buy-in and commitment to participate, identifying what they will do and when they will do it, and identifying what resources they need.

We then developed a strategy for integrating Agile practices into existing organizational processes, policies, and procedures. This allowed us to efficiently integrate these behaviors into all levels of the organization while still preserving those aspects of the existing organizational culture that remain valuable or important to success.

5. Plan for Agile Implementation in a non-Agile Organization

We decided to go with the most popular Agile methodology, which is Scrum. Before planning for an Agile project, we needed the following details (Figure 6).

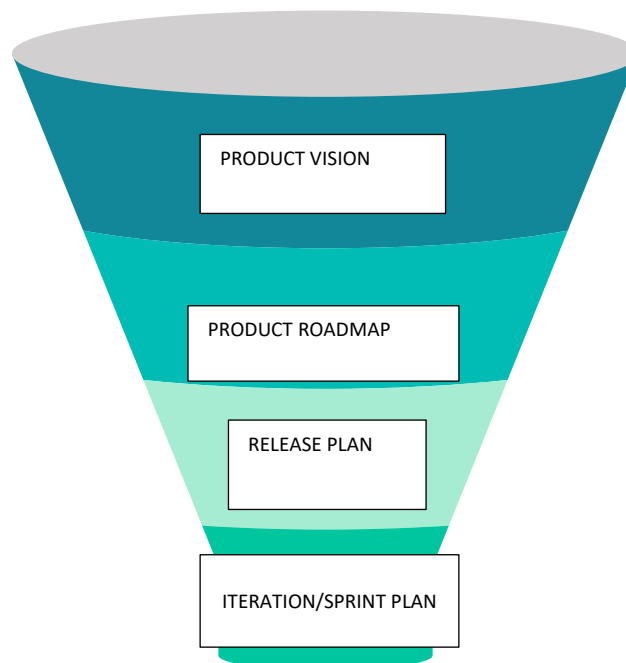


Figure 6 – Steps in Agile planning

5.1. Product Vision

Product vision is the product idea communicated to the team and stakeholders and is the driving force behind a product. We tried to answer the following questions in our product vision document.

1. What is the problem that this product solves?
2. What are we trying to achieve with this product?
3. Who will use this product, and what benefits will they get from it?
4. How does this product make people's lives better or easier?
5. How do we want users to feel when they use our product?
6. What are we going to build, why, and how are we going to do it?

5.2. Product Roadmap

Product roadmaps are a visual representation of a product's future. Creating a product roadmap allowed our teams to visualize the progress of their work and adjust as needed. We used the roadmap to track progress, identify risks, and adjust according to sprint feedback. A product roadmap also helped us set expectations with our customers about what they can expect from our team in the future, which led to increased customer satisfaction and loyalty.

5.3. Create a Release Plan



Figure 7 – Create a release plan

As we did our requirement gathering, we evaluated the time needed for the project along with the constraints of technology, materials, and other considerations. We determined what would be accomplished through technological means or manual labour (Figure 7).

5.3.1. Prioritize Requirements

Agile projects are a way of working that helps teams focus on the priority requirements and deliver the most critical features first. Prioritizing requirements benefited us by allowing our teams to focus on what was essential, allowing us to not get bogged down in less significant details. Prioritization helped us reduce risk as we were able to focus on delivering value quickly, and it encouraged collaboration as we provided a transparent prioritization process.

Prioritizing requirements in our Agile project was a challenging task, and it required our teams to clearly understand the business, product, and product vision. The team also identified the essential requirements to be satisfied first. We prioritized requirements in an agile project using a Moscow (Must Have, Should Have, Could Have, Won't have) analysis.

5.3.2. Create an Initial Design

We developed a design that was easy to understand, aesthetically pleasing, and within the budget. The team worked together to include customers, stakeholders, and Agile team members in building the design.

strategy. We had different perspectives on the problem and did not depend on documentation to express our ideas.

5.3.3. Determine the Success Factor

The success of our project was measured by how well we met our customer's and team's needs during the planning phase. To measure our success, we used criteria such as having a clear understanding of what needs to be delivered to meet customer expectations, a clear understanding of who is going to do each task and how this task will help in achieving the goal. We relied on excellent quality requirements, specification documents, and understood clearly how much time and money it will take to deliver each requirement.

5.4. Who Was Involved in the Plan?

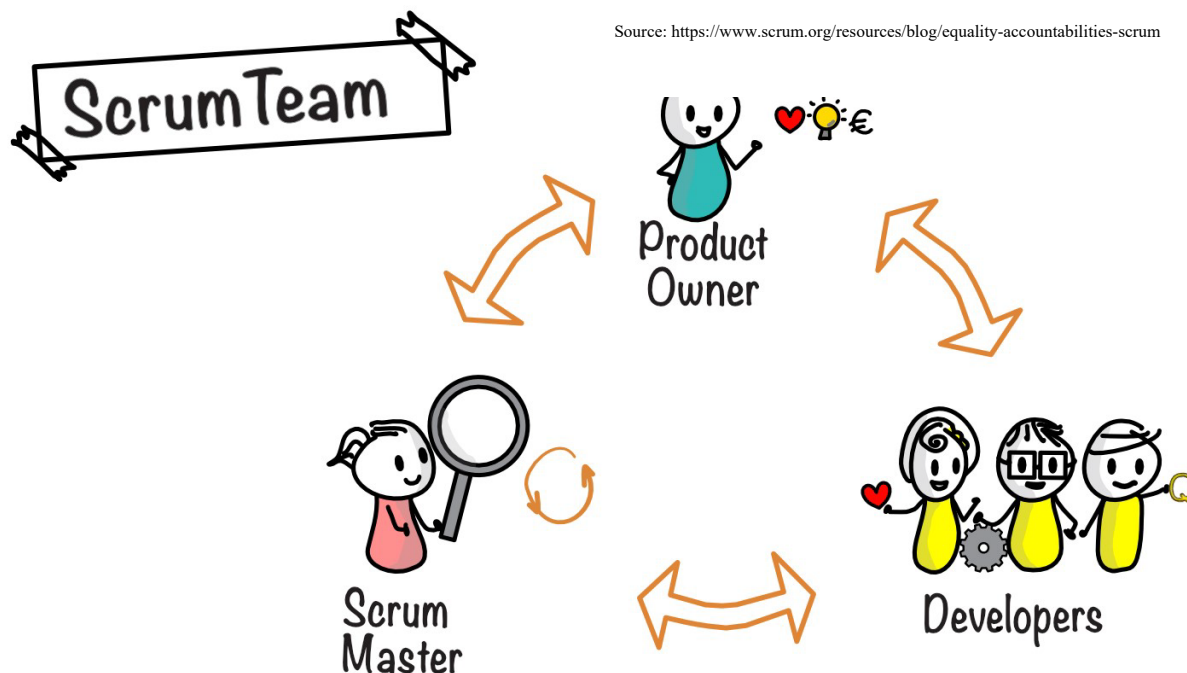


Figure 8 – Scrum team

The Scrum Master is a leader for the team and a facilities provider. He assists the team members in adopting Agile practices to ensure they fulfil their obligations and meet customers' needs. The Product owner provides a full view of the backlog of product requirements and the criteria for acceptance, while the Development Team defines the roles and carries out the work (Figure 8).

5.5. Prerequisites for Iteration Planning

Before we started the Iteration planning, we did the following:

- Ensured that the items in the backlog of the product were sizeable

- Ranked the epics created. Epics are large bodies of work that can be broken down into a number of smaller tasks (called stories) [11]
- Documented the acceptance criteria for each user story. User Stories are short requirements or requests written from the perspective of an end user [11]
- Calculated the number of user stories that could be accommodated into an iteration
- Broke these user stories down into tasks [4] and assigned each to an owner
- Assigned each task an approximate duration

The team members were assigned tasks based on their capacity or speed to ensure that the team member was not overwhelmed.

5.6. Iteration Planning Steps

Iterations are the basic building block of Agile development. Each iteration is a standard, fixed-length timebox, where Agile teams deliver incremental value in the form of working, tested software and systems. [12]

The Product Owner for our product picked up the top item in the existing backlog of items. The team members within explained the tasks needed to complete each user story as they were picked up and were accountable for tasks assigned to them.

The planning poker method was used to determine the complexity of each user story. Those with more story points required more time to develop and test compared to others. Team members were given the flexibility to determine the amount of time it would take them to complete each task.

These steps were then repeated for each item in the iteration process. If any team member was overloaded with work, duties could be re-allocated to others.

6. How To Implement Agile in An Organization That Is Not Agile

The fundamentals of Scrum are simple. To tackle an opportunity, the organization forms and empowers a small team, usually three to nine people, most of whom are assigned full-time. [5] The team is cross-functional and includes all the skills necessary to complete its tasks, and it manages itself and is strictly accountable for every aspect of the work. It is essential to train the team on the Agile framework if they have not previously worked in an Agile environment (Figure 9).

6.1. Design Sprint (Sprint 0)

The design process in an Agile project is iterative, and the process of designing in an Agile project is a continuous cycle of user feedback and design iteration. One of the best practices in Agile Design is to plan for a design sprint.

We started the design sprint by broadly exploring the problem, understanding its scope, and identifying potential solutions. Then, we generated ideas for solutions before focusing on one or two ideas to test initially. Finally, we presented our findings to stakeholders and then chose which idea would be implemented in more detail.

6.2. Project Charter

We kick-started the project by creating a project charter with the team's participation, which included a definition of what "done" meant. This way, when anyone in the team states that they are done with their task, everyone had the same understanding and meaning. We then defined the rules the team would adhere to and be accountable for.

6.3. Scrum Ceremonies

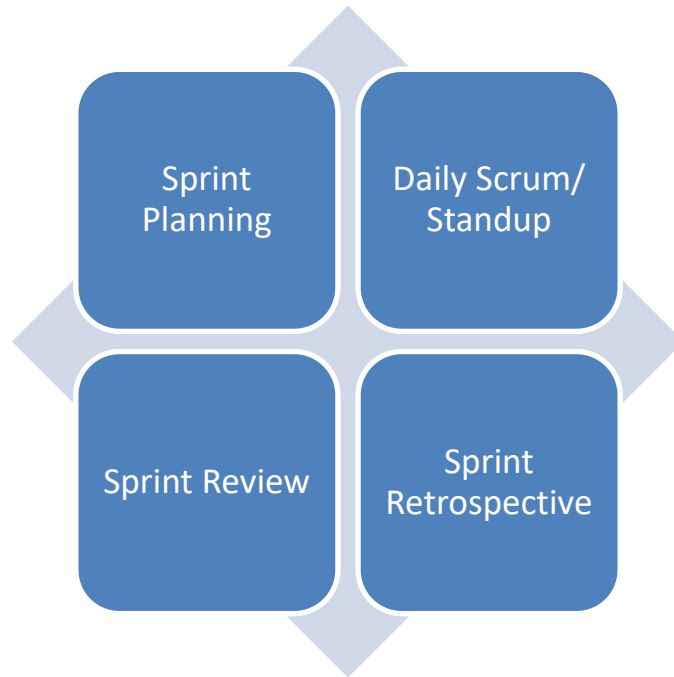


Figure 9 – Scrum ceremonies

6.3.1. Sprint Planning

Stories are committed to developers until they can fill the team's needs and are in line with the priorities of the owner of the product/stakeholders. In the following sprint, the team will utilize average velocity to help them determine the amount they should commit to in any given sprint.

The process began by finding the capacity based on the hours of each team member and subtracting the amount of time required to hold the celebrations during the sprint. For instance, four developers work 40 hours weekly, equal to 80 hours for each 2-week sprint. We subtracted hours for the daily stand-ups and backlog grooming, sprint planning reviews, sprint planning, and the time for sprint retrospectives to calculate the actual hours required to finish stories. Subtract 12 hours for the ceremony from each developer. Thus, their current availability is just 68. Multiplying 68 by 4, the capacity total is 272.

6.3.2. Daily Scrum/Stand up

Four days a week, team members met virtually via a team call and discussed what they accomplished the previous day and the next thing they planned to do the following day. These meetings were less about status and were utilized more for collaboration in the event of dependence and coordination within the team. We also ensured that no one was telling others how to proceed and that all are self-organized to complete the user story at the end of the sprint.

6.3.3. Sprint Review

At the end of the sprint, we conducted a sprint review meeting where the team reviewed/ demonstrated the stories to the product owners. We documented feedback from the product owners and evaluated its

impact on the development process. We used the test environment to highlight the demo where we created records using dummy data instead of actual production data to explain the feature.

6.3.4. Sprint Retrospective

In the sprint retrospective calls, the team reviewed the tasks and efforts they contributed during the sprint. We recorded what went well, what could have gone better, and what did not go well for the duration of the sprint. This allowed us to determine what we would like to put aside, begin doing, and then continue to do. Retrospection at the end of every sprint offered us opportunities for continuous improvement. An enormous improvement in quality, management, and performance happened within just a few sprints (as short as four to six weeks).

As an example, in one of our sprint retrospectives, we discussed that lack of technical documentation led to extended development time as the developers had to spend significant time going through the code to understand and then pick up the enhancement. We started creating technical documentation as, and when, team members had additional time within the sprint which helped us increase the turnaround time. In another example, each team member created their own test framework for testing which led to inconsistencies. We decided to have a common framework on which team members will write their test cases. This helped us have uniformity across testing and test results.

As much feedback as possible was highly appreciated and played a significant role in the overall success of the entire team.

6.4. Metrics for Agile

We categorized Agile metrics into productivity, progress, and quality types.

6.4.1. Productivity metrics

Productivity is a measure of how much work was completed. We used the Agile metrics below to understand how productive the team was.

- **Velocity** measured how much work was completed each time, which can be expressed as story points per day. We calculated sprint velocity based on the story points completed in the first few of sprints.
- **Lead time** measured how long it takes a team to complete their work from start to finish. [6] The team should always aim for short lead times because it's more efficient and cost-effective. [7] <https://kanbanize.com/kanban-resources/kanban-software/kanban-lead-cycle-time>
- We calculated the **average time** it took each team member to start and complete each task over the first four sprints. As an example, we found that Engineer 2 had a slightly longer lead time than Engineer 4 as they spent additional time finding solutions for some challenging problems. We used these metrics and conducted a brainstorming session to identify what we could do as a team to decrease the lead time.
- **Cycle time** is a measure that captures the time taken for the team to complete the sprint, i.e., time taken from In Development to Ready for testing.

6.4.2. Progress metrics

Progress metrics are a measure of how much work is completed.

The sprint burndown (Figure 10) represents actual scrum activities completed compared to the estimates of scrum-related tasks to assess a team's performance throughout a sprint.




Figure 10 – Sprint burndown

Epic and release burndowns are similar to sprint burndowns but epic burndown and release follow the progress across a more significant work body. These metrics are helpful when working with milestones. Also, the burndowns of releases monitor work progress when work is committed to a specific release.

6.4.3. Quality metrics

Quality metrics measure the degree to which the product meets its requirements. There are many quality metrics in an Agile scrum project, but the most important ones include the Quality of Backlog Items which determine how good the items in the backlog are. If they have high quality, they can be easily implemented, and their implementation will not take much effort. On the other hand, if they have low quality, they need more work, or they have a defect or deficiency.

A cumulative flow diagram (CFD) was used to see the status of various tasks - backlog, in-progress review, completed, and in-progress. The metric provided teams with an overview of all the activities in the workflow and the overall progression of the entire project.

The Quality of Team metric determined how good the team was. If they had high performance, the team quickly finished their tasks without much effort. On the other hand, if they had low performance, the team needed more time and effort to finish their tasks.

Escaped Defects allowed us to determine the quality of products received. This measure helped track the number of bugs discovered before a release is put into production.

Throughput measured the team's effectiveness by measuring the amount of work completed and handed over to the client within the period specified. This measure assisted us to determine the duration needed for the creation of software, and it let our end customers evaluate the team's degree of consistency.

The Quality of Product metric determined how good the product was in terms of value delivered (which uses value points to determine the quality of the work delivered to the client).

A control chart was linked to the cycle duration and was used to measure the cycle time of specific issues to assess the method's reliability, predictability, and stability. The control chart is a way to determine the success or failure of various processes within the project and is used to determine if there are any flaws.

7. The Benefits of Agile Methodology

1. Adopting Agile in our project involved teams creating and testing more frequently and delivering improvements to products and services more quickly to customers. Our customer satisfaction increased by 14% in 3 months since implementing this change.
2. The Agile development methodology allowed us to provide top-quality software in shorter release cycles that are more aligned with users' requirements.
3. Agile methodology helped us increase efficiency through more practical communication and was highly flexible to the ever-changing requests of customers. Our overall development-to-testing effort was reduced by over 30%. The number of back-and-forth communications using user acceptance testing decreased from around 5 - 6 rounds to no more than 2.
4. The ever-changing business environment of those customers requires flexibility and innovation on our part. As a team, we wanted to shift toward alignment between IT and business by ensuring that the process could be improved while considering the infrastructure and the technology.
5. We also focused on reducing time, efforts, costs, and expenditures in developing, documenting, testing, and deployment.
6. After Agile implementation in our project, developers and quality assurance (QA) teams were more prepared to work well together, allowing them to create software quickly while working closely with the client.
7. With an Agile approach to software delivery that focused on the set of measurements and metrics, teams were able to organize, plan, and deliver with sufficient certainty and a release-level commitment.
8. Measuring and tracking efficiency improvements cannot be accessed without transparency in the project. Teams responsible for software development and delivery understood the direction they were heading. Further development, testing, and operations teams were made aware of the present state of the project, its performance, and the goals of the project and the company.

8. Conclusion

Agile is a methodology of software development that focuses on the process of iterative and incremental delivery. It is a software development approach that emphasizes collaboration, simplicity, and communication. It also focuses on continuous testing and feedback to achieve higher quality products.

The Agile methodology has been gaining ground in various industries such as healthcare, finance, defence, education, etc. The ability to quickly adapt to changing requirements has made it popular among customers for its flexibility and ability to deliver results faster than traditional methods of project management.

Abbreviations

CFD	Cumulative Flow Diagram
QA	Quality Assurance

Bibliography & References

- [1] <https://www.agilealliance.org/agile101>
- [2] <https://habitation.com/what-is-agile-working>
- [3] <https://thelisting.com/how-to-make-decisions-with-an-agile-mindset>
- [4] <https://dokumen.tips/documents/gullvision-salisbury-university-authors-tim-stlauterburgcosc426finalreportexamplesgull.html>
- [5] <https://www.coursehero.com/file/55754463/5200-AlmetrixDrayton-Assignment1docx>
- [6] <https://industry-era.com/DevOps-Big-Data-and-Social-Media-10080.php>
- [7] https://www.reddit.com/r/transit/comments/v8niy7/whats_the_best_solution_to_rush_hours_in_public
- [8] https://www.reddit.com/r/transit/comments/v8niy7/whats_the_best_solution_to_rush_hours_in_public
- [9] <https://agilemanifesto.org/principles.html>
- [10] <https://agilemanifesto.org/>
- [11] <https://www.atlassian.com/agile/project-management/epics-stories-themes>
- [12] <https://www.scaledagileframework.com/iterations/>
- [13] Agile Estimating and Planning By Mike Cohn

Time is Ripe for D3.0 Farming: Achieving Optimal Spectral Efficiency by Allocating D3.0 Spectrum to OFDM

A Technical Paper prepared for SCTE by

Maher Harb

Director, Data Science
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (215) 990-8376
Maher_harb@comcast.com

Chad Humble

Senior Engineer, Next Generation Access Network
Comcast
4100 E Dry Creek Rd, Centennial, CO 80122
+1 (720) 951-8092
Chad_Humble@comcast.com

Sebnem Ozer

Senior Principal Engineer, Next Generation Access Network
Comcast
1800 Arch St, Philadelphia, PA 19103
+1 (215) 286-8890
SEBNEM_OZER@COMCAST.COM

Dan Rice

VP, Next Generation Access Network
Comcast
4100 E Dry Creek Rd, Centennial, CO 80122
+1 (720) 512-3730
daniel_rice4@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. How Downstream Spectrum is Utilized.....	3
3. D3.0 Spectrum Farming Approaches.....	6
4. View of the State of the Spectrum.....	8
5. 2 Gbps Testing.....	11
6. Conclusion.....	15
Abbreviations	16
Bibliography & References.....	17

List of Figures

Title	Page Number
Figure 1 - Histogram of D3.1 device penetration across the network. Service group counts omitted from the y-axis on purpose as they do not inform the discussion.....	4
Figure 2 - Distributions of (1) number of D3.0 channels per service group, (2) number of devices per service group, (3) D3.0 service group utilization, and (4) D3.0 utilization range across channels within the service group.....	5
Figure 3 - Joint D3.0 and OFDM distribution shown as a scatter plot. The black line is the bisector (45 deg. angle) and the orange line is line-of-best-fit for the scatter.....	6
Figure 4 - Joint D3.0 and OFDM distributions at different number of SC-QAM channels converted to OFDM. (Top left) 0 channels converted—representing the baseline, (top right) 4 channels converted, (bottom left) 8 channels converted, and (bottom right) 12 channels converted.....	7
Figure 5 - Analysis exploring converting up to 16 SC-QAM channels to OFDM. The number of converted channels is specific to each service group and selected such that projected D3.0 service group utilization does not exceed 60%.....	8
Figure 6 - Example of three common downstream spectrum configurations.	9
Figure 7 - The same configuration shows in the previous figure with the outcome of vacancy detection overlaid as dark gray band.	10
Figure 8 - Example of two “conflicts” picked up by the vacancy detection algorithm: (1) a local insert was confirmed at 850 MHz, and (2) ingress due to police repeater was picked up at 885 MHz on a different RPD.....	11
Figure 9 - The data throughput over iCMTS (top 2 panels) and vCMTS (bottom 2 panels) for a single D3.1 cable modem demonstrating speeds in the 2.3-2.4 Gbps range. The different colors indicate different TCP flows utilized by the speed test.	12
Figure 10 - Speed test results for a 2 Gbps speed tier. There are no RF issues observed for this cable modem and corresponding node. However, utilization levels during the speed test have an impact on the results. Purple lines are total octets at service group level while yellow lines are total octets for the OFDM channel. The load telemetry data is displayed with a delay compared to speed test timing displayed in the bottom of each graph and indicated by a red arrow.....	13
Figure 11 - Total data load computed at the traffic generator’s customer premises equipment (CPE) ports connected to the CMs in the serving group. The speed test runs between 7-22 seconds.....	14
Figure 12 - Total average load measured at the iCMTS by summing 30 sec and 1 sec channel loads in the serving group.....	15
Figure 13 - Speed test result from the traffic generator GUI.....	15

1. Introduction

Spectrum in cable networks is often statically configured to support a mix of DOCSIS[®] 3.0 (D3.0), single carrier-quadrature amplitude modulation (SC-QAM), and DOCSIS 3.1 (D3.1) Orthogonal Frequency Division Multiplexing (OFDM) channels. Generally, more spectrum is allocated to D3.0 based on infrequent adjustments to add additional channels, when possible, informed by utilization and modem technology trends as part of long-range capacity and capital planning processes. These capacity management processes consider many business and policy elements including supply chain, procurement, warehousing and inventory, construction permitting, and work order management for node segmentations and cable modem termination system (CMTS) configuration. These important elements can be made even more effective when enabled through software-defined virtual network functions (VNF) that remove many of the manual procedures via automation and deferral of other procedures and capital investment based on real-time channel and spectrum optimization. Automating these solutions also leads to fewer instances of human error, reduces service interruption, and increases availability.

In one illustrative example, the downstream (DS) configuration consists of 44 D3.0 channels (264 MHz) and 1 OFDM (96 MHz) channel in a node that supports 750 MHz of total hybrid fiber coax (HFC) infrastructure spectrum. A configuration that favors D3.0 from a spectrum perspective is counter-intuitive since OFDM supports higher modulation efficiencies (up to 4096-QAM) and a superior low-density parity check (LDPC) error correction algorithm. Using our Profile Management Application (PMA) VNF [1-2], which adjusts physical layer capacity in real time, Comcast is able to improve DS Mbps/MHz by > 44% on average using D3.1 technology. Yet, conventional wisdom states that expansion of the OFDM spectrum is limited by the D3.1 device penetration. Multi-Gbps products, which are enabled only by OFDM/OFDMA, can help drive this device penetration based on consumer demand.

This paper introduces a fresh perspective on optimizing the spectrum to achieve desired outcomes for product speeds, capacity, node segmentation, and cost effectiveness. The goal of the optimization exercise is to maximize spectral efficiency via reallocation of DS D3.0 channels to OFDM while accounting for any increase in utilization, which may lead to costly node augments. We show that the current configuration and CMTS load balancing place a higher burden on the OFDM spectrum, thus making the conditions ripe for spectral reallocation. This reallocation creates a material positive benefit in support of multi-Gbps products. Lastly, we present a VNF concept to farm the D3.0 spectrum and manage load balancing in an automated fashion.

2. How Downstream Spectrum is Utilized

A common spectrum configuration assigns 44 SC-QAM channels for D3.0 and a single 96 MHz D3.1 OFDM channel (equivalent spectrum to 16 SC-QAM channels). Historically, before D3.1 device penetration accelerated, the OFDM channel was thought of as a “bonus” that offers some relief to the D3.0 spectrum and allows the D3.1 cable modem (CM) devices access to exclusive speed tiers (e.g., 1 Gbps and higher). However, D3.1 device penetration has been steadily increasing, mostly by organic means (i.e., new customers opting for the “latest and greatest” cable modems with the best Wi-Fi capability). The distribution of D3.1 CM penetration across all service groups (equivalent to a downstream port on a CMTS) and hardware platforms at Comcast is shown in Figure 1. The median D3.1 penetration has recently crossed a critical level with the consequence that the traffic loads on D3.0 and D3.1 technologies in the spectrum are now comparable on average (but vary per service group). This balance means that a more thoughtful spectrum allocation strategy is now required. Such a strategy may be tailored to the specific characteristics of a service group and considers aspects that include D3.0 and OFDM channel utilization, traffic forecast, projected growth in D3.1 device penetration, and plant

configuration, as well as all the nuances of the current spectrum configuration (location of video, video-on-demand, tones, local inserts, etc.).




Figure 1 - Histogram of D3.1 device penetration across the network. Service group counts omitted from the y-axis on purpose as they do not inform the discussion.

Additional D3.0 distributions of interest are shown in Figure 2. Panel 1(left) confirms that a 44 SC-QAM configuration is the most common, though other variations exist (namely, 40, 36, 32, 24). Panel 2 shows the distribution of the number of DOCSIS devices per service group. Panel 3 is the distribution of service group utilization for the D3.0 part of the spectrum. Service group utilization was calculated by first averaging the point-in-time utilization (sampled at 5-min intervals) across all D3.0 channels in the service group then taking the 98th percentile sample within a 30-day period as representative of the utilization. Typically, service groups that fall above the 80% utilization level are considered to be highly utilized and may require some form of remediation (it could involve, for example, “splitting the node” to create two service groups out of the original one). Finally, Panel 4 (right) shows the D3.0 channel utilization variation range (max – min) within a service group. This metric is a gauge of the CMTS internal load balancing function within the D3.0 spectrum given that each device is allocated a subset of the 44 channels to use (each unique subset is referred to as a downstream bonding group). Though not shown here, we found a large discrepancy between our primary hardware platforms in terms of D3.0 load balancing ability, with some achieving much tighter utilization spread among the SC-QAM channels than others.




Figure 2 - Distributions of (1) number of D3.0 channels per service group, (2) number of devices per service group, (3) D3.0 service group utilization, and (4) D3.0 utilization range across channels within the service group.

What is more important for the concept of spectrum allocation is the load balancing between the D3.0 and the OFDM parts of the spectrum. The CMTS will tend to shift traffic for D3.1-capable devices heavily in favor of OFDM on some of our hardware platforms. For example, on our virtual CMTS (vCMTS) we can influence the load balancing through other VNF and micro-service opportunities. This is beneficial as it allows us to be maximally spectrally efficient, but the higher utilization of D3.1 should be anticipated and accounted for as part of the capacity planning process and policy described in the introduction. To further investigate this feature of CMTS functionality, consider the joint distribution for D3.0 and OFDM service group utilizations shown as scatter plot in Figure 3. In the ideal scenario, with evenly balanced D3.0 and OFDM, the distribution would be centered around the 45° bisector line (shown as black line). Instead, the joint distribution reveals higher OFDM utilization compared to D3.0 (orange line-of-best-fit rotated counterclockwise (CCW) relative to the 45° line). A few points regarding how this data is interpreted:

- At the current median D3.1 penetration level, balancing the D3.0 and OFDM utilizations is challenging for the CMTS owing to the limited spectrum allocated to OFDM. The imbalance is further affected by the fact that the highest speed tiers can only be accessed using D3.1-capable devices.
- More heavily used service groups are driven in large part by the OFDM portion of the spectrum, implying that the D3.0 part has excess bandwidth to yield.
- The joint distribution will continue to rotate counterclockwise (CCW) as D3.1 device penetration continues its organic growth and as CMTS scheduler management is improved, causing more service groups to cross the high utilization threshold due to high OFDM utilization.

Given all of the above, shifting some spectrum from D3.0 to OFDM is a win-win opportunity:

1. It relieves service groups that are heavily utilized in the OFDM spectrum.
2. It adds capacity to the network enabling even higher speed offerings.

Yet, this is not a one-size-fits-all situation and must be targeted within a well-considered spectrum reallocation program.




Figure 3 - Joint D3.0 and OFDM distribution shown as a scatter plot. The black line is the bisector (45 deg. angle) and the orange line is line-of-best-fit for the scatter.

3. D3.0 Spectrum Farming Approaches

This section explores two D3.0 spectrum farming approaches and estimates their impact on service group utilization and capacity. First, presented below are some guiding principles that apply across all analyses:

- It is understood that farming D3.0 spectrum presents a trade-off: overall capacity is increased because of the higher OFDM efficiency compared to SC-QAM, but the D3.0 service group utilization may also increase because of the lost D3.0 spectrum and the fact that not every device in the service groups is D3.1-capable.
- Utilization is assumed to scale linearly with the amount of added/removed spectrum, which is a reasonable assumption that applies to both SC-QAM and OFDM. Thus, this assumption is used as a recipe to project the impact of adding/removing spectrum on utilization.
- For the time being, we ignore second order effects. For example, the fact that increased capacity will enable higher speed tiers, which will in turn lead to higher traffic and increased service group utilization. We also ignore that the load balancing function may adjust given the re-allocation of available capacity between technologies. Note that such changes in speed offerings and customer behavior are tackled within the traffic forecasting and capacity planning functions.

The first approach considers farming the same number of SC-QAM channels across the entire footprint. While such an approach is not favorable from the perspective of the ultimate optimization, it is appealing from the perspective of its simplicity: maintaining a single standard configuration and avoiding development of what could be a complex spectrum management system. Estimating the impact of network wide reallocation of spectrum is straightforward. At any given number of converted SC-QAM channels, the utilizations of the D3.0 and OFDM components of a given service group are scaled proportionally to the lost or added spectrum. The efficiency assumptions of the D3.1 channel are based on years of PMA [1,2] experience across tens of millions of modems.

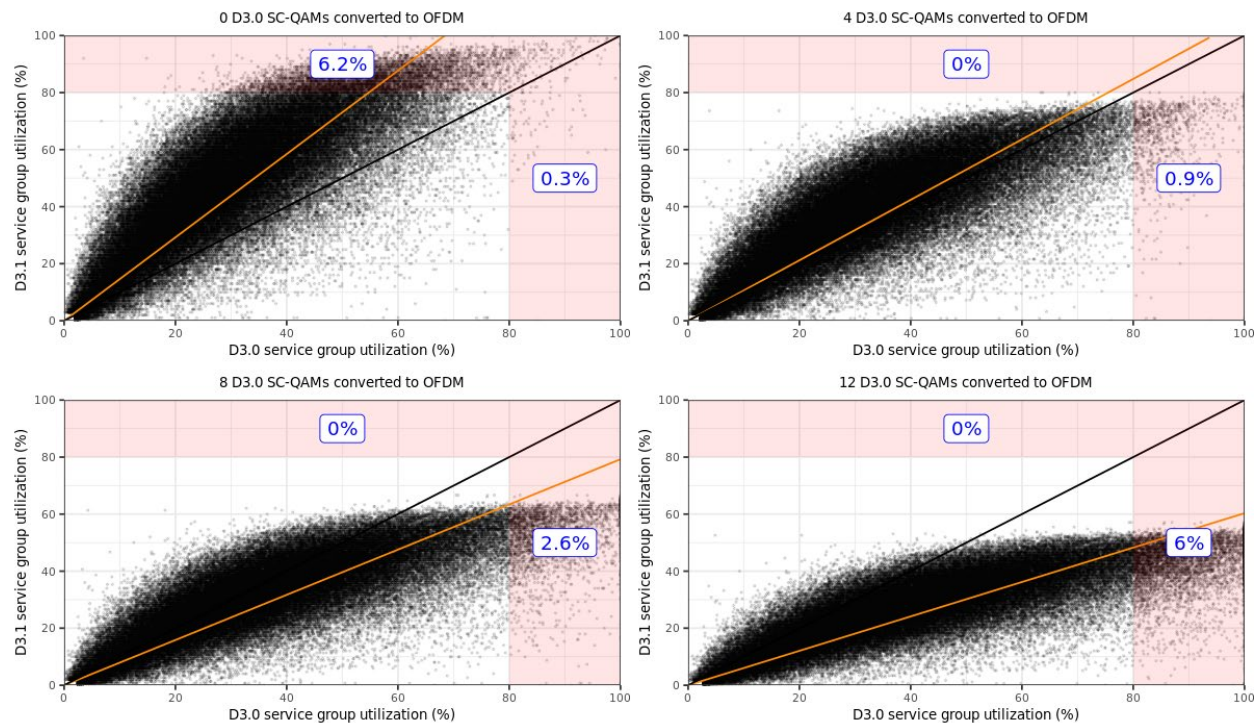


Figure 4 - Joint D3.0 and OFDM distributions at different number of SC-QAM channels converted to OFDM. (Top left) 0 channels converted—representing the baseline, (top right) 4 channels converted, (bottom left) 8 channels converted, and (bottom right) 12 channels converted.

Results of the impact analysis are shown in Figure 4. Each of the panels explores the impact of converting n SC-QAM channels to OFDM, with n being 0 (baseline), 4, 8, and 12. Highlighted on the plots are service groups that are heavily utilized (>80% utilization) falling in the shaded pink area with the percent figure annotated. It is evident that as more SC-QAM channels are converted, the joint distribution rotates clockwise (CW) as expected. At the highest level of converting 12 channels, we end up at roughly the same ratio of service groups that are heavily utilized (~6%) yet these are now driven by the D3.0 utilization rather than OFDM. Under this scenario, no relief is gained in terms of the number of service groups that require remediation as a proxy for equipment and labor capital for adding network capacity. However, ~200 Mbps of capacity was added to each service group in the network (OFDM is estimated to be 16.6 Mbps more efficient than D3.0 per converted 6 MHz channel—this estimate is based on current PMA performance where OFDM is deployed). Also, based on the results in Figure 4, one can argue that converting 4 SC-QAM channels is an easy decision: it alleviates almost all utilization issues while adding a modest ~66 Mbps of capacity to each service group in the network.

To support multi-Gbps downstream product speeds, additional Mbps are needed than can be accommodated with the baseline configuration scenario. This speed requirement is particularly challenging with HFC design spectrum limits such as 750 MHz or 860 MHz node deployments. The only way to accommodate the capacity needed by a multi-Gbps product speed is to re-allocate more OFDM spectrum by reducing D3.0 spectrum in these nodes or via a labor-intensive spectrum upgrade.

The second approach manages spectrum on an individualized service group level. The idea is to convert any number of channels such that:

- The SC-QAM block is not reduced beyond 24 channels as this number corresponds to the maximum number of D3.0 channels to which a D3.0 device is able to bond. Thus, a 44 SC-QAM block will have a maximum of 16 channels to convert.
- The projected D3.0 service group utilization does not exceed 60%, which is a level that allows a safety buffer of several years compounded annual growth rate (CAGR) of 20% before exceeding 85%. Similar models have been developed based on peak period available bandwidth targets.

The analysis shown in Figure 5 explores this approach when applied to service groups that are currently configured with 44 SC-QAM channels. It can be seen from Panel 1 (left) that most service groups are able to release 16 channels for conversion without driving D3.0 utilization above the 60% level. For those service groups that are already utilized above 60%, no channels are converted. Other service groups fall between the 2 extreme cases and are able to convert some number of channels between 0 and 16. Comparison between the current D3.0 utilization (Panel 2) and the projected D3.0 utilization (Panel 3) shows a stark difference. The projected distribution does not exhibit a normal shape as it is artificially constructed from different sub-populations; most notable, the sub-population that was driven to the 60% limit. Lastly, Panel 4 reveals the extent of added capacity: most service groups increased their capacity from just under 2.5 Gbps to above 2.8 Gbps.




Figure 5 - Analysis exploring converting up to 16 SC-QAM channels to OFDM. The number of converted channels is specific to each service group and selected such that projected D3.0 service group utilization does not exceed 60%.

4. View of the State of the Spectrum

The concepts discussed so far present the problem of spectrum farming in an idealized fashion whether the approach is based on a completely dynamic policy optimizing capital investment lifetime or is based on policy constraints. It sufficiently covers the core dimensions pertaining to D3.1 CM penetration, service group utilization, and service group capacity. However, it ignores some important practical aspects to spectrum farming. One of these foundational pragmatic requirements is the challenge around having high confidence visibility into current spectrum allocations on every node. For instance, the Comcast network includes a variety of different hardware platforms, different plant configurations, and different spectrum allocations across localities – e.g., placement and number of video quadrature amplitude modulation (QAM) channels, DOCSIS channels, local video insertions, amplifier automatic gain control (AGC) signals, leakage markers, etc. In a 1 GHz plant, it might not be required to convert SC-QAM channels as the configuration is expected to contain enough vacant spectrum to deploy multiple increased bandwidth OFDM channels. Yet, even in this scenario, it is necessary to understand how the spectrum is currently utilized and the kind of housekeeping that is needed to free up contiguous spectrum blocks. More importantly, striving for a standardized spectrum configuration is a goal in its own because standardization simplifies both day-to-day operations before software management can take over and

deployment of future technologies (e.g., full duplex (FDX), and high split). Given these dynamics, it becomes obvious that automation is a prerequisite to any effort directed towards:

- Optimizing capacity in HFC network segments with limited spectrum by converting D3.0 channels to OFDM.
- Freeing up contiguous blocks of spectrum to deploy OFDM where plant configuration allows it.
- Relocating specific channels in order to move towards a steady state with a few standardized configurations across the network.

Presented in this section are early explorations into concepts that are foundational to the spectrum management tool. The first exercise involved gaining a complete picture of the state of the spectrum configuration across the network. The initial phase of data collection was limited to the vCMTS platform. Data from various sources were merged to provide a unified view on the downstream configuration for every remote PHY device (RPD), which is used in this paper as synonymous to a node, managed by the vCMTS platform. This data analysis task confirmed the lack of standardized configuration as it revealed that ~700 distinct permutations exist for the tens of thousands deployed RPDs. Figure 6 shows three example configurations that represent a small part of the diversity of what is deployed in the field. These examples fall within the top 25 configurations, which covers most of the nodes, creating a “long tail” in the distribution of custom DS channel configurations. Notice the sharp contrast between the three configurations. The first (Rank 1 in RPD count) has a block of vacant spectrum above 750 MHz ready to be used for OFDM expansion. The second (Rank 2) has a lot of vacant spectra but in a more fragmented way, with 2 SC-QAM blocks surrounding the OFDM channel. The third (Rank 21) has vacant spectrum in the expansion region as well but except for a 6 MHz “local insert” channel at ~850 MHz. The local insert could be thought of as a video channel combined locally with the cable signal to serve a specific use (e.g., a security camera system within a hospital or a housing development). This type of spectrum use presents a problem for OFDM expansion. Yet it is often the case that these types of exceptions are not easy to trace since the configuration found in the data is registered for an entire site (large set of RPDs) even if the insert is physically available only on a single RPD or on a node leg. This situation highlights the importance of adding a validation layer on top of the picture that was constructed solely from various configuration data sources.




Figure 6 - Example of three common downstream spectrum configurations.

A validation methodology was adopted that is based on using the full band spectral capture from CM gateways to infer whether a range of spectrum is occupied or vacant. Briefly:

- Raw spectra (117 KHz resolution bandwidth) are captured from every DOCSIS device on the to-be-validated nodes.
- Channel power is calculated from raw spectra for the standardized 6 MHz EIA channel plan.

- A threshold is adopted for determining whether the channel is in use or vacant. For the results presented here, the threshold was set to -40 dBmV.
- Device-channel level data is aggregated to the RPD-channel level and to the configuration-channel level.

The outcome of the vacancy detection process is shown in Figure 7 in the form of a dark gray band overlaid on the configuration map. It is seen that the vacancies detection algorithm performs well as the spectrum gaps neatly correspond with documented vacancies. There are few exceptions. Tones for amplifier AGC, local modulator, or leakage usually overlap with inferred vacancies since these are much narrower than the 6 MHz channel. Furthermore, leakage tones are allowed to be placed within a DOCSIS block, so they should be ignored. What is more interesting is the overlap between the local insert at 850 MHz (Configuration 21) and the inferred vacancy. Since this configuration is shared by 378 RPDs, the local insert may be present in only a small subset of these or even just a single smaller multi-dwelling unit (MDU) on a node that also serves other customers. A deep dive into the device and RPD level data reveals 2 interesting findings shown in Figure 8:

- The local insert was confirmed on one RPD out of the 378 that share the same configuration. 100% of devices on the RPD exhibited energy above the noise level on that RPD.
- Energy was picked up by 50% of devices on a different RPD on a channel that does not correspond to any known use in the configuration (at 885 MHz). Further investigation revealed that the signal represented ingress from a recently installed police repeater.

The first case requires field work to move the inserted channel to a different location and certify that the RPD can proceed with OFDM expansion. The second case does not require any further action since ingress is more appropriately handled by customized OFDM profiles within the PMA system along with plant maintenance dispatch for downstream ingress remediation.




Figure 7 - The same configuration shows in the previous figure with the outcome of vacancy detection overlaid as dark gray band.

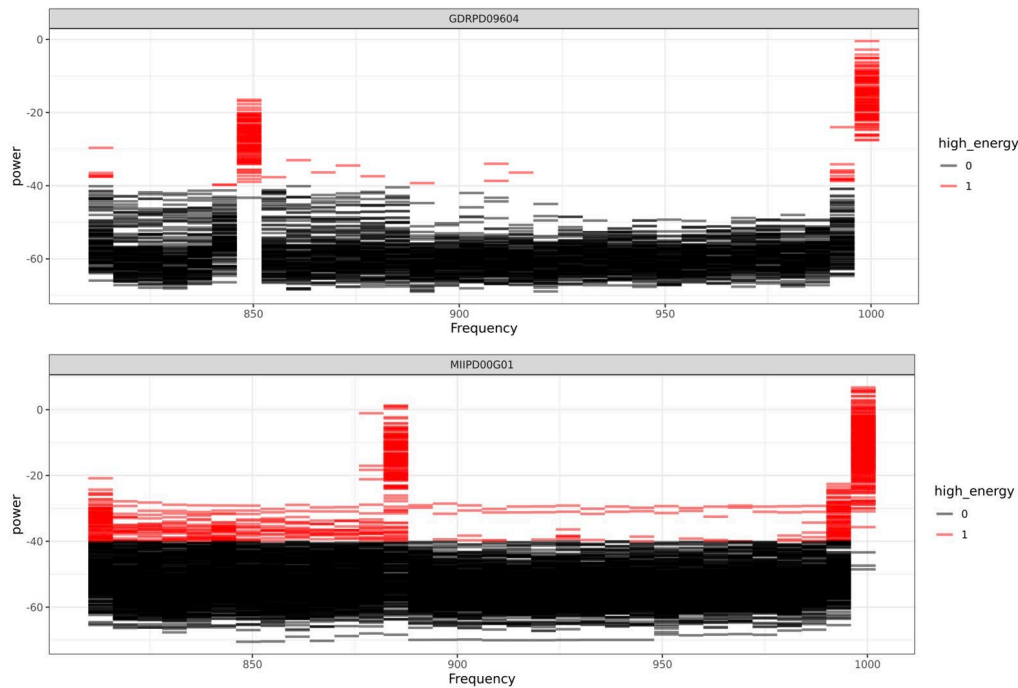


Figure 8 - Example of two “conflicts” picked up by the vacancy detection algorithm: (1) a local insert was confirmed at 850 MHz, and (2) ingress due to police repeater was picked up at 885 MHz on a different RPD.

These early explorations show that we now have a methodology for capturing and validating the state of the spectrum configuration. The subsequent steps involve:

1. Automating the pipeline for building and validating the spectrum configuration for every node in the network (both integrated CMTS (iCMTS) and vCMTS platforms).
2. Implementing an algorithm for moving spectrum around to free up contiguous spectrum blocks for OFDM expansion.
3. Implementing the algorithm for farming D3.0 spectrum where plant configuration restricts use above 750 MHz.
4. Developing a layer that translates the output from both algorithms into the CMTS-specific configuration.
5. Ensuring visibility into real-time spectrum allocations and the technology used in that spectrum is available to all stakeholders. This includes the capacity planning tools and process to ensure optimized investment in network expansion only where and when it is needed.

5. 2 Gpbs Testing

In this section, we introduce some of the preliminary speed test results for a 2 Gbps downstream service to demonstrate how additional OFDM spectrum could be used for higher broadband product speeds. The potential of higher speeds is one of the primary motivations behind the D3.0 farming development. The downstream spectrum configuration for these experiments was 28 SC-QAM channels & a single 192 MHz OFDM on the iCMTS and 44 SC-QAM channels and a single 192 MHz OFDM on the vCMTS.

Figure 9 shows that a single D3.1 cable modem can reach 2.3-2.4 Gbps data throughput over iCMTS and vCMTS systems when there are no congestion or radio frequency (RF) noise issues.




Figure 9 - The data throughput over iCMTS (top 2 panels) and vCMTS (bottom 2 panels) for a single D3.1 cable modem demonstrating speeds in the 2.3-2.4 Gbps range. The different colors indicate different TCP flows utilized by the speed test.

The serving group utilization and RF noise issues may affect the maximum achievable rate at any given time. Figure 10 displays an example for a D3.1 cable modem in a service group with different utilization levels measured within a 15 second window. The results are used to analyze the average and peak utilization levels of the service group to confirm the service group availability for the 2 Gbps speed tier and to drive service group upgrade strategy accordingly. The graphs on the left correspond to relatively high utilization levels and correspondingly impact 2 Gbps speed test results compared to the graphs on the right in which the service group utilization is low. This is analyzed further in this section with more examples to show the impact of utilization burstiness and transmission control protocol (TCP) behavior on supporting higher downstream rates.




Figure 10 - Speed test results for a 2 Gbps speed tier. There are no RF issues observed for this cable modem and corresponding node. However, utilization levels during the speed test have an impact on the results. Purple lines are total octets at service group level while yellow lines are total octets for the OFDM channel. The load telemetry data is displayed with a delay compared to speed test timing displayed in the bottom of each graph and indicated by a red arrow.

With the increased speed tiers and emerging applications, new characteristics have been observed for the ratio of peak-to-average traffic load and peak duration per subscriber and per serving group. As discussed in [3], the higher service tiers have higher burst rates. Although traffic load averaged over longer time periods such as 5 minutes may be low, bursts observed within smaller time windows in the order of seconds can be 10-20 times higher for the current pre-FDX rates. For Gbps rates, microbursts can be observed within windows on the order of microseconds. Microbursts may happen due to service characteristics, network segments with different rates and classic TCP protocol response during congestion [4,5].

Traditionally, utilization in each channel at a CMTS is reported per 5-minute interval. New telemetry models can provide utilization values at 15 second intervals. Measuring traffic utilization and burstiness at smaller window sizes may not be feasible for large scale systems. However, this information is crucial in network analysis. As discussed in [6], bursts and their overlap by multiple users can be infrequent and these short-term bursts may have little impact on aggregate values, but the spikes may have a significant impact on services like gaming, videoconferencing, and augmented reality/ virtual

reality (AR/VR) applications. Internet service providers (ISPs) must develop effective measurement and prediction systems to project granular measurements at small scales to aggregated scales.

An example is shown for an iCMTS with the same settings described for Figure 9. A speed test is performed for a subscriber with 2 Gbps DS speed tier rate, between 7 and 22 seconds as shown in Figure 11. Initially, the serving group load is close to 0.9 Gbps before the speed test starts. After the speed test starts, a set of bursts occurs between 12 and 18 seconds and affects the speed test outcome as the service group is highly utilized during this timeframe. Figure 12 displays the total average load measured at iCMTS. The measurements are collected via a command line interface (CLI) command to get the bit rates for each DS channel. One set of the measurements provides the load averaged over a 30 second window while the other set provides the load averaged over a 1 second window, which is a more accurate representation of the burst that impacts the speed test as shown in Figure 13.




Figure 11 - Total data load computed at the traffic generator's customer premises equipment (CPE) ports connected to the CMs in the serving group. The speed test runs between 7-22 seconds.




Figure 12 - Total average load measured at the iCMTS by summing 30 sec and 1 sec channel loads in the serving group.




Figure 13 - Speed test result from the traffic generator GUI.

This example shows that utilization and other metrics must be analyzed at more granular levels with increased speed tier rates. Burst level analysis is crucial in network planning, cost analysis, quality of experience (QoE), and speed assurance as well as optimization of network functionalities for efficiency and fairness.

6. Conclusion

We presented our perspective on reallocating D3.0 spectrum to OFDM, motivated by the need to increase service group capacity to support higher speed tier offerings. We conclude that:

- The growth in D3.1 device penetration has now reached a level where reconsidering how downstream spectrum is allocated between D3.0 and OFDM is necessary.
- The objective of reallocation is twofold: increase service group capacity to enable higher speed offering and provide relief to a currently highly utilized OFDM spectrum.

- To this end, two approaches were presented: a one-size fits all farming of D3.0 channels, vs. a customized recommendation that considers service group characteristics.
- The latter, due to its dynamic nature, requires building a spectrum management tool; a cornerstone of such a tool is building a unified source of truth for the current state of the spectrum configuration.
- Such a tool also feeds into additional spectrum housekeeping efforts that are targeted towards freeing up space for expanded OFDM region in support of enabling 2 Gbps and higher downstream speeds.
- These spectrum management and farming VNFs will be critical to managing spectrum effectively as the future evolves towards DOCSIS 4.0 and multi-Gbps symmetrical services.
- Initial speed test results reveal the importance of increasing service group capacity via an approach that considers the speed test result sensitivity to utilization and the impact on TCP flow control behavior when testing for these very high speed tiers.

Abbreviations

AGC	Automatic gain control
AR	Augmented reality
CCW	counterclockwise
CLI	Command line interface
CM	cable modem
CW	clockwise
CMTS	cable modem termination system
CPE	customer premise equipment
D3.0	DOCSIS 3.0
D3.1	DOCSIS 3.1
dBmV	Decibels relative to 1 milliVolt
DOCSIS	Data Over Cable Service Interface Specifications
DS	downstream
EIA	Energy Information Administration
FDX	full duplex
Gbps	Gigabits per second
GHz	GigaHertz
HFC	hybrid fiber coax
iCMTS	integrated cable modem termination system
ISP	Internet Service Provider
KHz	KiloHertz
LDPC	Low density parity check
Mbps	Megabits per second
MDU	multi-dwelling unit
MHz	MegaHertz
OFDM	orthogonal frequency division multiplexing
OFDMA	Orthogonal frequency division multiple access
PMA	profile management application
QAM	quadrature amplitude modulation
QoE	quality of experience
QoS	quality of service
RF	radio frequency
RPD	remote PHY device

SC-QAM	single carrier-quadrature amplitude modulation
TCP	Transmission Control Protocol
vCMTS	virtual cable modem termination system
VNF	Virtual network function
VR	Virtual reality

Bibliography & References

1. *A Machine Learning Pipeline for D3.1 Profile Management*, M. Harb, J. Ferreira, D. Rice, B. Santangelo, and R. Spanbauer, NCTA technical paper, 2019.
2. *Full Scale Deployment of PMA*, M. Harb, J. Ferreira, D. Rice, B. Santangelo, NCTA technical paper, 2020.
3. *Traffic Engineering in a Fiber Deep Gigabit World*, John Ulm & Tom Cloonan, SCTE-NCTA Workshop, 2021.
4. <https://support.huawei.com/enterprise/en/doc/EDOC1100086962>
5. IETF L4S Architecture Draft: <https://datatracker.ietf.org/doc/draft-ietf-tsvwg-l4s-arch/>
6. *Roaring into the '20s with 10G*, Robert Howald, Robert Thompson, Sebnem Ozer, Daniel Rice, Larry Wolcott, Tom Cloonan, Ruth Cloonan, John Ulm, Jan Ariesen, SCTE-NCTA Workshop, 2020.
7. <https://github.com/BroadbandForum/obudpst>

Top-5 Things You Should Know About Operating a Virtual CMTS

A Technical Paper prepared for SCTE by

Brady Volpe

Founder

NimbleThis and The VolpeFirm

3000 Old Alabama Rd. Suite 119-434, Alpharetta, GA 30022

404.954.1233

brady.volpe@nimble-this.com :: brady.volpe@volpefirm.com

Thuy Nguyen

Cable Segment Lead

Intel Corporation

2200 Mission College Blvd, Santa Clara, CA 95054

thuy.nguyen@intel.com

Muhammad Siddiqui

Platform Solution Architect

Intel Corporation

2200 Mission College Blvd, Santa Clara, CA 95054

muhammad.a.siddiqui@intel.com

Michael Lafser

Technical Solution Engineer

World Wide Technology

1 World Wide Way, Maryland Heights, MO 63043

Mike.Lafser@wwt.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. DOCSIS Primer.....	4
3. CMTS Overview.....	4
4. How do I deploy?.....	5
4.1. Legacy CMTS Deployment.....	6
4.2. vCMTS Deployment.....	7
5. How Do I Scale and Add capacity?.....	13
5.1. Capacity Expansion Legacy CMTS.....	14
5.2. Capacity Expansion vCMTS.....	15
6. How do I Troubleshoot the CMTS?.....	16
6.1. Troubleshooting a Legacy CMTS.....	17
6.1.1. Hardware Failures.....	17
6.1.2. Software Failures.....	18
6.1.3. Subscriber Impairments.....	18
6.2. Troubleshooting the vCMTS.....	19
7. How do I Multi-task with a vCMTS?.....	23
7.1. What is required to run the DOCSIS workload in software on a general purpose x86 server?.....	23
8. Does a vCMTS Consume More or Less Power Than a Legacy CMTS?.....	25
9. Conclusion.....	26
Abbreviations.....	28
Bibliography & References.....	29

List of Figures

Title	Page Number
Figure 1 - Cable Access Network Diagram.....	5
Figure 2 - QAM profile configuration.....	6
Figure 3 - DS controller configuration.....	7
Figure 4 - GUI configuration example of interfaces.....	8
Figure 5 - Profile List (Left) and Sample Downstream GUI profile.....	9
Figure 6: RPD/SG provisioning process.....	10
Figure 7 - Routing (Left) and interfaces (Right) created.....	11
Figure 8 - Dashboard navigation.....	12
Figure 9 - System Cluster Summary.....	12
Figure 10 - Service group summary.....	13
Figure 11 - Forecast Downstream Bandwidth Demand 2018-2030 (source: Strategy Analytics, Inc., CommScope).....	14
Figure 13 - Power Overview of Legacy CMTS (4.388 kW used).....	17
Figure 14 - Example of Detailed Powering and Temperature Statistics on Legacy CMTS.....	18
Figure 16 - Telemetry and insights data flow.....	20
Figure 17 -Grafana dashboard visualizing two memory reports.....	21
Figure 18 - Show Cable Modem PHY dashboard.....	21
Figure 19 - Anatomy of a Virtual CMTS deployment.....	24

List of Tables

Title	Page Number
Table 1 - Physical attribute comparison of different CMTS systems	6
Table 2 - Power consumption comparison.....	26

1. Introduction

This paper will discuss the top-5 things an operator should know about operating a virtual cable modem termination system (vCMTS). Specifically, five key aspects will be discussed including deployment, scalability, capacity, troubleshooting, multitasking, power, and space needed for operating the vCMTS. The reader will understand that legacy CMTSs and vCMTSs are similar in most cases, especially from an operational and deployment standpoint. This means that if a user is familiar with operating a legacy CMTS then they should be comfortable migrating to a vCMTS. Further, the paper will explore some advantages that may make vCMTSs attractive to cable operators over legacy CMTSs. By the end of the paper the reader should have confidence in understanding that vCMTSs are quite like legacy CMTSs from a features and performance standpoint, while adding other valuable benefits and functionality that the reader may not have considered prior to reading this paper. Sit back and enjoy the read.

2. DOCSIS Primer

Data-Over-Cable Service Interface Specifications (DOCSIS[®]) technology is effectively a transparent Ethernet bridge over a hybrid fiber/coax (HFC) network. There are two functional components in a DOCSIS network, the cable modem (CM) on the subscriber side and the CMTS in the headend or hub site. The CMTS communicates with the CMs on one or more single carrier quadrature amplitude modulation (SC-QAM) channel(s) and/or orthogonal frequency division multiplexed (OFDM) channel(s). Data on these channels is digitally encoded on radio frequency (RF) signals on the downstream path of an HFC network between 108 and 1.2 gigahertz (GHz). The CMs communicate with the CMTS using one or more SC-QAM and/or orthogonal frequency division multiple access (OFDMA) digitally encoded RF channels, transmitted on an upstream HFC frequency between 5 to 204 MHz (note these frequencies will change with DOCSIS 4.0). The digital data contains DOCSIS management information in addition to subscriber traffic. The CMTS is the system scheduler which coordinates the power level, frequency, transmit time, and pre-equalization of all CM signals on the DOCSIS network.

3. CMTS Overview

In any CMTS you have the hardware, the operating system software, the application software, and switching equipment. A vCMTS is made up of these same components by integrating the operating system, management software, and application software on a commercially available server.

What is a legacy CMTS as defined in the introduction of this paper? A legacy CMTS is a chassis based CMTS containing downstream and upstream RF cards utilizing on-board software and/or firmware for system operation. In DOCSIS 3.1, legacy CMTSs may be augmented to support distributed access architecture (DAA). In this scenario, the legacy CMTS may not have on board RF cards, but instead the RF and physical layer (PHY) portion will extend to a remote PHY (R-PHY) node, but the medium access control (MAC) processing is still performed in the legacy CMTS.

A virtual CMTS or vCMTS on the other hand, is software that is purpose built to run on commodity servers. Proprietary vCMTS software is installed on the server that converts the server into a vCMTS platform. A network interface card (NIC) attached to a switch connects the server to an R-PHY node or shelf and a DAA architecture is created, like the DAA architecture described with the legacy CMTS.

Figure 1 shows a high-level architecture of the cable access network showing both a converged cable access platform (CCAP) and a vCMTS with DAA deployments.

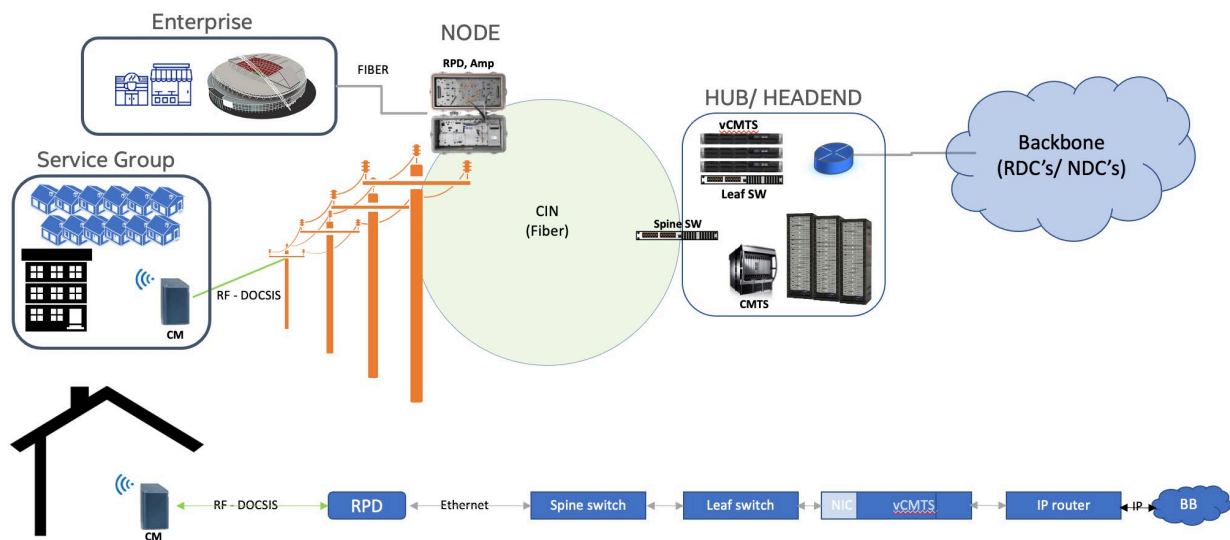


Figure 1 - Cable Access Network Diagram

4. How do I deploy?

The first aspect of turning up any DOCSIS HFC plant, be it legacy CMTS or vCMTS will be the initial deployment. While some aspects differ, particularly when comparing a legacy CMTS to a vCMTS initial deployment, many of the same setup procedures will be the same, just taking different forms. This makes sense because the functionality of a CMTS is just reproduced through virtualization of the CMTS. One area that greatly differs, is the amount of physical setup required for each device since legacy CMTS products require a large amount of either coax (analog) or fiber (digital) for the data ports, along with a much larger physical weight needing to be racked, as can be seen in Table 1.

Table 1 - Physical attribute comparison of different CMTS systems

	I-CCAP	D-CCAP	vCMTS
Data Cables	186 (RF) + 4 (Fiber)	68 (Fiber)	12 (Fiber)
Weight	429 lbs	429 lbs	36 lbs/server (3 servers = 108 lbs)
Power Cables	20 terminal cables	20 terminal cables	6 pluggable cables

Note:

- Weight is inclusive of the equipment such as chassis, line cards, power supplies, etc. Items such as optics, combiners, cables, are excluded from the weight calculations.
- D-CCAP may be several pounds lighter due to PHY modules being removed.

4.1. Legacy CMTS Deployment

Initial deployment of a legacy CMTS begins after the gear has been rack mounted, all linecards and physical NICs installed, and the devices are cabled up. Once the device is up and running, a console terminal is opened, and the management interface, security, host information, and secure shell host (SSH) access are configured through Command Line Interface (CLI). The software and firmware are then upgraded to the proper versions. Networking is then configured with wide area network (WAN) interfaces, loopback interfaces, and routing all being configured for the CMTS to communicate with the network elements (NE). Connectivity, routing, and networking are then verified, completing the initial setup.

DOCSIS configuration can now begin after the initial deployment is completed. There are multiple items that will need to be configured through CLI for the system before service group (SG) provisioning can occur. The first item to be configured is the quadrature amplitude modulation (QAM) profiles; an example of this is shown in Figure 2.

```
Router(config)#cable downstream qam-profile 13
Router(config-qam-prof)# annex A
Router(config-qam-prof)# modulation 256
Router(config-qam-prof)# interleaver-depth I12-J17
Router(config-qam-prof)# symbol-rate 6952
Router(config-qam-prof)# spectrum-inversion off
Router(config-qam-prof)#description new-256-qam
Router(config-qam-prof)#
```

Figure 2 - QAM profile configuration

These profiles contain the physical layer information for the downstream (DS) channels in the system. OFDM and frequency profiles are configured, providing the information required for DS channel configuration for the DS ports, as seen below in Figure 3.

```
Router(config)#controller Integrated-Cable 3/0/0
Router(config-controller)# max-ofdm-spectrum 192000000
Router(config-controller)# max-carrier 32
Router(config-controller)# base-channel-power 36
Router(config-controller)# rf-chan 0 31
Router(config-rf-chan)# type DOCSIS
Router(config-rf-chan)# frequency 801000000
Router(config-rf-chan)# rf-output NORMAL
Router(config-rf-chan)# power-adjust 0
Router(config-rf-chan)# qam-profile 1
Router(config-rf-chan)# docsis-channel-id 1
Router(config-rf-chan)# !
Router(config-rf-chan)#
Router(config-rf-chan)#rf-chan 158
Router(config-rf-chan)# power-adjust 0
Router(config-rf-chan)# docsis-channel-id 159
Router(config-rf-chan)# ofdm channel-profile 50 start-frequency 837000000 width 192000000 plc 930000000
```

Figure 3 - DS controller configuration

This process is repeated for the upstream (US) side with modulation and OFDMA profiles being configured.

With the profiles now created, the controllers can then be configured by applying the profiles to them. With the controllers configured, the Cable, DS Cable, and US Cable interfaces are configured by associating the controllers to them, configuring the MAC-Domains as the channels get bound to it. Configuring SGs continues by creating a Fiber Node by assigning the US and DS cables to it, which creates and associates the US SG and DS SG to the MAC Domain. The CMTS's RF output is checked at the head end to ensure it falls within specified values, and the system is connected to the RF passive equipment. To monitor the system, simple network management protocol (SNMP) is configured and connected to an external software source.

4.2. vCMTS Deployment¹

Initial deployments of a vCMTS cluster consist of a few steps. The server hardware, networking equipment (routers, switches) need to be racked and cabled up, with their management and baseboard management controller (BMC) IP addresses configured, along with the required network configurations so that access is allowed. From there, an external computer, or virtual machine with network access to the devices can be connected and loaded with an automation tool such as Ansible to act as a deployer. The vCMTS software images, operating systems and Ansible playbooks get loaded on to the deployer. A configuration yet another markup language (YAML)² file is created containing all the host information and passwords that are required for the vCMTS cluster deployment. A playbook is then executed, deploying the operating systems, vCMTS software packages, and containers which are automatically deployed and synced for the cluster. After the cluster is deployed, they are connected but unconfigured to talk to the rest of

¹ Much of this section, and its figures, are specific to Cisco's cloud Native Broadband Router (cNBR) but has been generalized where available.

² YAML is a human-readable data-serialization language. It is commonly used for configuration files and in applications where data is being stored or transmitted. YAML targets many of the same communications applications as Extensible Markup Language but has a minimal syntax which intentionally differs from SGML. (source: <https://en.wikipedia.org/wiki/YAML>)

the network elements. This configuration can occur through applying a pre-made JavaScript Object Notation (JSON) configuration files, or directly configuring the NE and networking addresses in a graphical user interface (GUI). These will include all the core addressing, precision time protocol (PTP) information, routing information, etc. for all interface elements, an example is shown below in Figure 4.

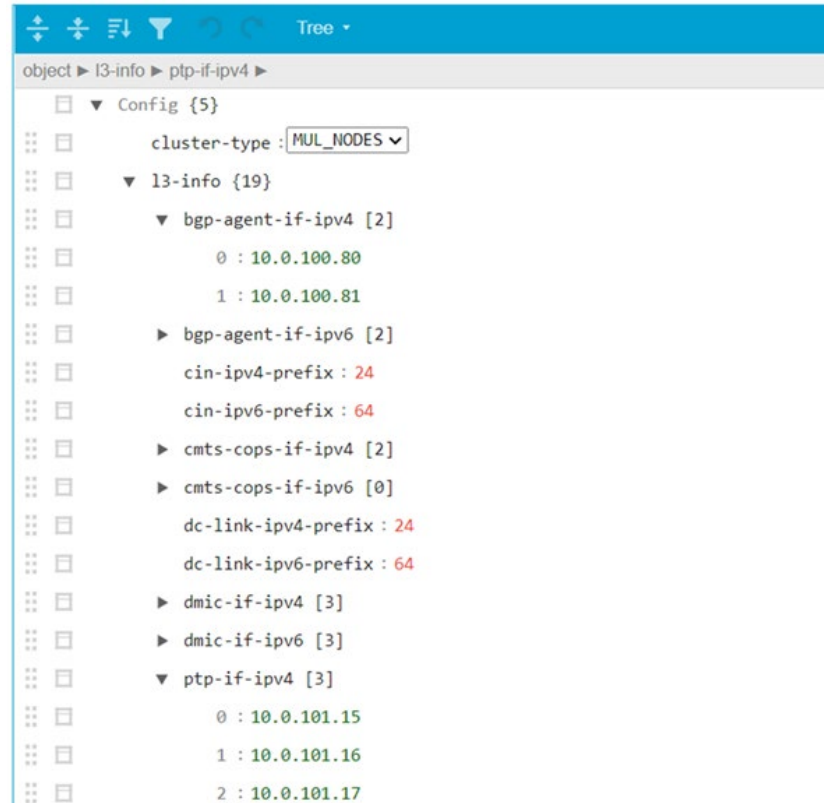


Figure 4 - GUI configuration example of interfaces

Once all the prerequisite day zero configurations are completed, and the vCMTS core cluster is talking to the NE, the basic network setup is completed, and it should be verified that all the NEs can talk to each other, and the systems' PTP clock is locked. This can be accomplished through various means, from ensuring ping capability, checking border gateway protocol (BGP) neighbors and routing tables.

DOCSIS configuration can begin after verification of connectivity. This is once again very similar to a legacy CMTS deployment, and all the elements that would go into configuring a CMTS SG. Profiles for the various elements, the DS channels, US channels, remote PHY device (RPD) PTP information, etc. will all need to be created. These will contain the same information as in a legacy CMTS, as Figure 5 illustrates. This can again be done through importing JSON files into the system, allowing for automation in the configuration process, or through use of a GUI interface.

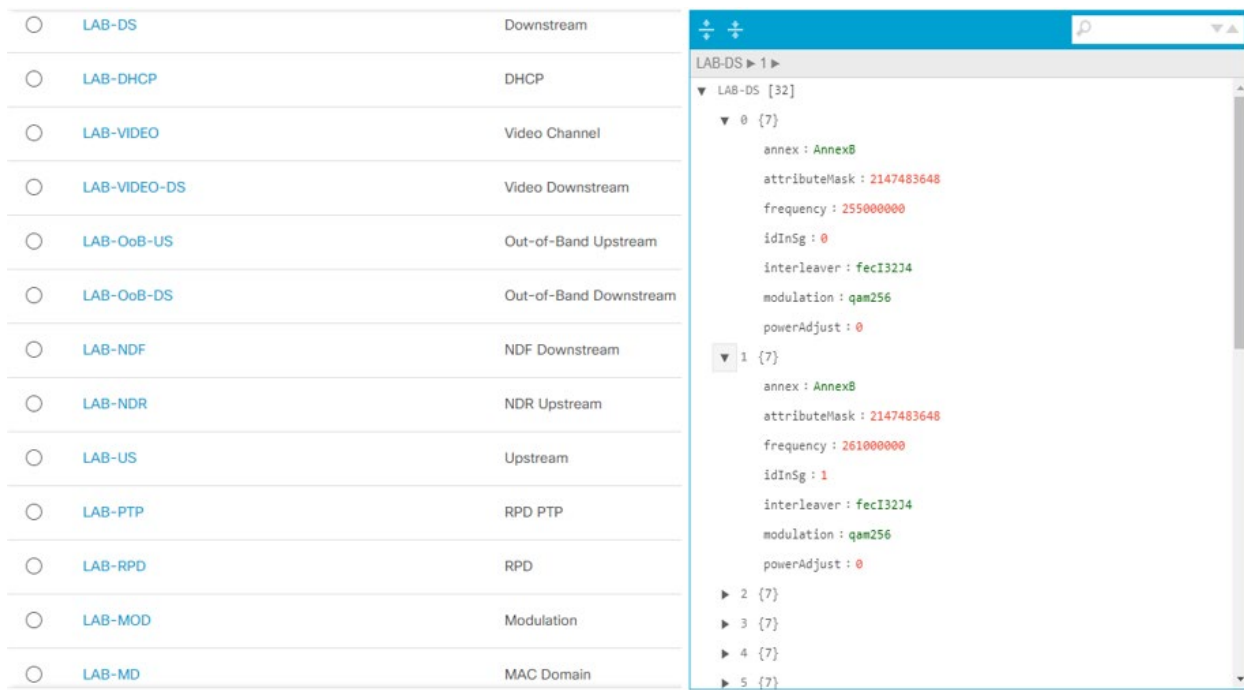


Figure 5 - Profile List (Left) and Sample Downstream GUI profile

Now that all the prerequisite DOCSIS and dynamic host configuration protocol (DHCP) profiles are in the system, SG templates can be created from the profiles, likewise, Layer 3 templates, acting as a cable helper configuration, containing the information for DHCP relay, and BGP peer information, can be created. These are, once again, like a legacy CMTS system and contain all the information required for provisioning.

Unlike previous sections, provisioning SGs in a vCMTS differs from a legacy CMTS. In this case, a separate RPD configuration containing the core information, and DS/US cable information is not needed, as these are configured automatically once a RPD (and consequently a SG) is created. Utilizing a GUI based interface, provisioning is also simplified. Here, the RPD MAC information, names, and the previously created templates provide the necessary information required to provision a SG. As shown in Figure 6, the vCMTS software then checks the configuration, creates the required cabling and routing information, then checks the online status before posting a status report.

RPD MANAGEMENT
Inventory
Add RPD
Edit RPD
Replace RPD
Delete RPD
Image Management
Secure Software Download
Code Validation Check
Cutover RPD

1
Target-Setting
2
Pre-Check
3
RPD-Add
4
Post-Check
5
Report

RPD Details
cnBR Cluster* WWT-CNBR
RPD MAC* 7018.A766.EC20
RPD Name* RPD0
SG Name* SG0
SG Template* n338
Layer 3 Template* L3-Template
Max wait time 12 min
RPD Location
Region Midwest
City St. Louis
Neighborhood WWT
Address 60 Weldon Parkway
Latitude 38.7079857681959
Longitude -90.44010645046485
Next Step

RPD MANAGEMENT
Inventory
Add RPD
Edit RPD
Replace RPD
Delete RPD
Image Management
Secure Software Download
Code Validation Check
Cutover RPD

1
Target-Setting
2
Pre-Check
3
RPD-Add
4
Post-Check
5
Report

Pre-RPD-Add Checklist:
✓ 1. Check RPD MAC valid.
✗ 2. Check RPD Name valid.
✓ 3. Check SG Name valid.
4. Please connect RPD physically.
☒ Please confirm RPD has been connected physically and start RPD config adding.
Next Step

RPD MANAGEMENT
Inventory
Add RPD
Edit RPD
Replace RPD
Delete RPD
Image Management
Secure Software Download
Code Validation Check
Cutover RPD

1
Target-Setting
2
Pre-Check
3
RPD-Add
4
Post-Check
5
Report

100%
RPD Add Progress:
Start RPD adding...
Checking transaction state
Checking SG configuration
RPD configuration in CnBR check passed
Next Step

RPD MANAGEMENT
Inventory
Add RPD
Edit RPD
Replace RPD
Delete RPD
Image Management
Secure Software Download
Code Validation Check
Cutover RPD

1
Target-Setting
2
Pre-Check
3
RPD-Add
4
Post-Check
5
Report

25%
Post-check Progress:
Start polling RPD state...
The maximum waiting time is 12 minutes
RPD Init(gcp)
Next Step

RPD MANAGEMENT
Inventory
Add RPD
Edit RPD
Replace RPD
Delete RPD
Image Management
Secure Software Download
Code Validation Check
Cutover RPD

1
Target-Setting
2
Pre-Check
3
RPD-Add
4
Post-Check
5
Report

Summary Report
Task: Add RPD
Start time: 2:21:35 AM; End time: 2:23:41 AM

cnBR Cluster	WWT-CNBR
RPD MAC	7018.A766.EC20
Service Group List	SG0
RPD State	online
RPD Version	v1.3
Result	Success

Figure 6: RPD/SG provisioning process

While this is a convenient and easy way to provision a SG, for ongoing deployments this would be quite tedious. However, since this is a vCMTS, it is an application programming interface (API) driven setup, provisioning a large amount of RPDs can be accomplished using API calls directly to import the configuration files into the system. As was previously stated, the routing and interfaces are configured automatically by the software, these interfaces are the same ones that would be used in legacy CMTS systems, while the routing is the information for the traffic and DHCP relay, and an example is shown in Figure 7.

SG Name	SG ID	VRF	Prefix	NextHop
SG0	0	default	10.0.122.1/32	10.0.100.103
SG0	0	default	10.0.122.1/24	10.0.100.2
SG0	0	default	10.0.121.1/32	10.0.100.103
SG0	0	default	10.0.121.1/24	10.0.100.2
SG0	0	default	10.0.120.59/32	10.0.100.2
SG0	0	default	10.0.120.1/32	10.0.100.103
SG0	0	default	10.0.120.1/24	10.0.100.2


```

{
  "cluster-id": "10.254.154.173.nip.io",
  "interface-list": [
    {
      "if-name": "10.254.154.173.nip.io_SG0_Cable0",
      "if-type": "DocsCableMacLayer",
      "if-mac": "84:b2:61:59:00:01",
      "if-stack-status": "Up",
      "if-in-octets": 0,
      "if-out-octets": 0
    },
    {
      "if-name": "10.254.154.173.nip.io_SG0_Upstream0",
      "if-type": "DocsCableUpstream",
      "if-mac": "",
      "if-stack-status": "Up",
      "if-in-octets": 0,
      "if-out-octets": 0
    }
  ]
}

```

Figure 7 - Routing (Left) and interfaces (Right) created

After the SG for the system has been provisioned, the configured elements need to be verified working as intended. As with legacy CMTS systems, the RF output of the node (RPD) in the field should be checked with probes. With a vCMTS system, however, many of the monitoring systems have been simplified due to the information being pulled from containers and direct information gathering by streaming telemetry viewable through available dashboards, shown in Figure 8.




Figure 8 - Dashboard navigation

A few examples of this can be seen below, with Figure 9 showing the overall summary of the system, and Figure 10 showing information of an individual SG.

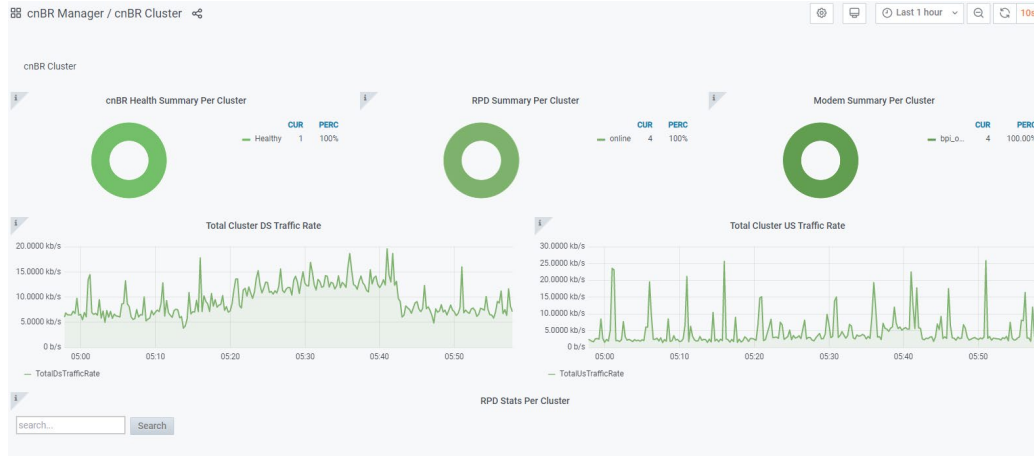


Figure 9 - System Cluster Summary

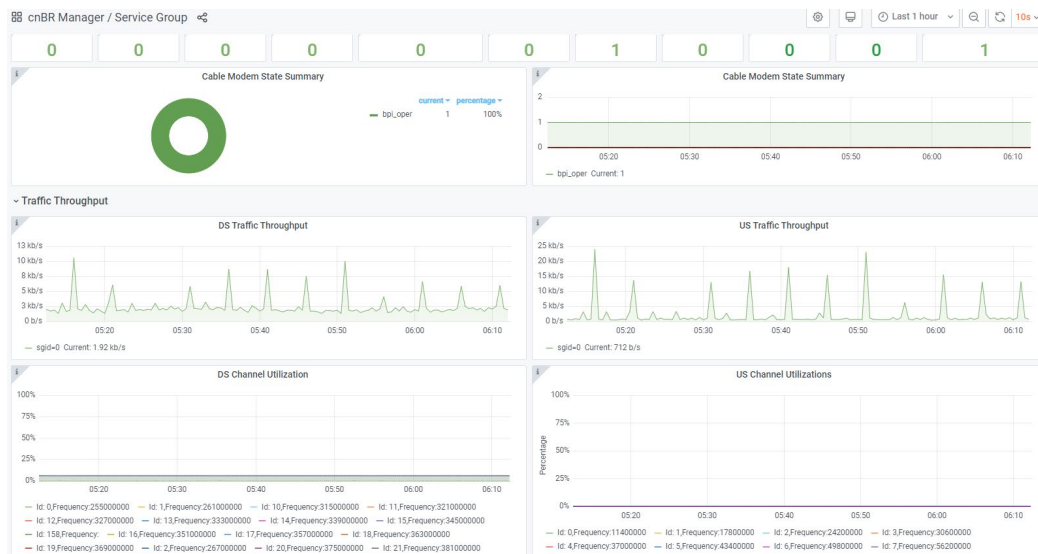


Figure 10 - Service group summary

Virtual CMTS deployments are heavily containerized architectures, with each container acting as an individual application. These containers are orchestrated using the open-source API Kubernetes (often referred to as K8s), or other container management platforms. These containers are applications running the individual components of the overall vCMTS architecture, such as packet processing, packet cable, telemetry, OFDM, SG management, MAC scheduler, etc. While all these containers are running in the background on a vCMTS cluster, they are not needed to be touched or managed by an operator as Kubernetes manages them and their services. So, while a vCMTS does run as all these individual pieces, it operates fundamentally the same as a legacy CMTS. After all, DOCSIS is DOCSIS.

5. How Do I Scale and Add capacity?

A common scenario with any CMTS deployment is increasing capacity. Traditionally, year-over-year growth has been roughly 20% in the downstream and upstream as can be seen in Figure 11. This traffic growth is projected to continue for the foreseeable future. The impact on any CMTS, legacy or virtual, is that the CMTS must be scalable and capable of supporting the continued traffic growth.




Figure 11 - Forecast Downstream Bandwidth Demand 2018-2030 (Source: Strategy Analytics, Inc., CommScope)

5.1. Capacity Expansion Legacy CMTS

Capacity expansion in a legacy CMTS is dependent on the CMTS vendor, though ultimately a well-defined process. Legacy CMTSs will have a finite number of downstream and upstream SC-QAM and OFDM/OFDMA channels. These channels may be combined to create service groups. A SG is defined as a set of downstream channels and upstream channels which are further associated logically in the CMTS as a MAC domain. Downstream and upstream communication over the downstream and upstream RF channels to a specific fiber node will be associated together as traffic. One can think of this as a shared data pipe in the downstream and shared data pipe in the upstream. Similarly, on the CMTS, each individual channel in the downstream is shared or “utilized”, provided the CMTS is properly configured. The same holds true for the upstream.

Consider a downstream with 32 SC-QAM channels at 256-QAM. Assuming a 256-QAM channel supports approximately 38 Mbps after overhead, the total downstream capacity would be:

$$32 \times 38 \text{ Mbps} = 1.216 \text{ Gbps}$$

This is a common CMTS downstream configuration prior to DOCSIS 3.1 OFDM. When subscribers on this fiber node saturate the fiber node in the downstream with continuous traffic during peak hours (which is about 1.2 Gbps) chaos will break out and the CSR phone lines will

start to ring. Subscribers will experience poor quality of experience due to network congestion. This will eventually generate a work order for a node split on that fiber node.

Node splits are expensive and time consuming. One or more new fiber nodes will be installed to reduce the number of subscribers per SG. The CMTS must then be updated to add a new SG(s) for the new fiber node(s). This is a common and ongoing process in most DOCSIS networks. DOCSIS 3.1 has enabled cable operators to delay fiber splits by adding DOCSIS 3.1 OFDM channels in the downstream. These channels can add up to 1.89 Gbps per 192 MHz OFDM block. In either case, the CMTS must support the addition of licensing for more SC-QAM and/or OFDM channels for the SG. It is possible that the CMTS will need additional hardware such as line card(s). The line card(s) provide the conversion of data to RF signals which are referred to as SC-QAM and OFDM channels. If the CMTS is already completely full and no additional slots are available for new line cards, then the cable operator would need to buy a new CMTS chassis and perform the steps mentioned in section 4.1.

Once the CMTS has been licensed and upgraded with hardware, someone with knowledge of the CMTS and required SG update must program the CMTS. The programming is similar to that in section 4.1. CLI programming is not consistent across CMTS vendors. But in general, the concept is the same, driven by DOCSIS specifications. Some cable operators have automated the process to the point where a SG can be created by a script. These scripts may need to be updated as networks evolve and CMTS code changes.

5.2. Capacity Expansion vCMTS

One of the benefits of migrating from a traditional hardware-based legacy CMTS appliance to a software-based containerized vCMTS solution is the agility and time required to add capacity. Also, a legacy CMTS requires more rack space than a vCMTS. If a legacy CMTS is fully populated, then to add capacity, it may require an additional legacy CMTS chassis. A vCMTS on the other hand, that runs on common off the shelf (COTS) servers and may need as little as 1 rack unit (RU) of rack space to add new SGs. This is a significant space savings.

Kubernetes has gained popularity over the years and has now arguably become the de-facto standard for container orchestration and its lifecycle management. Most, if not all, vCMTS core offerings from independent software vendors (ISVs) are orchestrated and managed using Kubernetes APIs. As the existing Kubernetes cluster starts to run at maximum capacity and resources start to run out, to add more SGs, new K8s worker nodes can be added to an existing cluster that runs vCMTS applications. A single K8s cluster can run several hundred worker nodes though officially K8s claim to support up to 5000 worker nodes.

Adding new worker nodes in the headend is relatively easier if the management network and simple IT infrastructure are properly set up for remote access. When a new server with desired hardware specifications is acquired, it is mounted in the rack and connected to the management and data plane networks. An out-of-band management interface can be used to update basic input/output system (BIOS) and enable required hardware features (e.g., Hyperthreading), single root input/output (I/O) virtualization (SRIOV) etc. The new server can either have a pre-installed operating system (OS) with custom configurations (IP, ssh keys, etc.) before it's racked up, or the

OS can be installed remotely after the server is mounted in the rack. This paper presents one example of how a new server with no pre-installed OS can be provisioned in the virtualized cable headend which has a DHCP and preboot execution environment (PXE) boot server with desired OS image already setup. As the new server boots up for the first time it can be booted from the network using a PXE enabled NIC. This will let the OS image be downloaded from the PXE boot server (trivial file transfer protocol [TFTP] server) and install the OS on the newly provisioned server. The OS image can be customized beforehand to enable remote ssh access before its uploaded on to the TFTP server. Once the initial OS is installed using PXE, the server can then be accessed remotely and ready to be added as a new worker node to an existing vCMTS K8s cluster.

There could be multiple ways to add new nodes to a Kubernetes cluster but broadly it can be divided into two different approaches 1) manual 2) automated.

A manual installation approach would require all K8s node software components to be installed and configured individually. This could be a tedious task and prone to errors and doesn't reap the real benefits of virtualizing a cable headend.

A better approach is to adopt automation which significantly reduces time and effort to deploy and scale, while improving quality by preventing human errors. One popular automation tool adopted widely today is Ansible. Ansible is an agentless automation tool that runs Ansible “playbooks” from a remote Ansible controller to perform cluster provisioning, configuration management and performs various actions on remote machines. Ansible controller communicates with remote/host machines over SSH protocol. Ansible playbooks are configuration files that define tasks to be performed on host machines and let users customize multiple parameters to fit their installation needs. Playbooks can be used to deploy new K8s clusters and add new K8s nodes to a running cluster. The whole process of adding new vCMTS nodes takes a few hours if playbooks are configured properly.

Another automation tool widely used today with K8s is Helm. Helm is a package manager tool that runs on top of K8s to automate the installation process of plugins and K8s capabilities and used to deploy and manage applications like vCMTS applications. As the new K8s node is added, Helm can be used to deploy new vCMTS service groups to the newly provisioned K8s node. Scaling out new SGs to a new node consists of running a few K8s CLI commands that spins up new vCMTS SGs. The rest of the SG configurations are the same as previously explained in section 4.2.

6. How do I Troubleshoot the CMTS?

An important aspect of any CMTS is ensuring that subscribers continue to send and receive high speed data without interruption. Being able to identify if the CMTS is having a hardware and/or software issue or if there is an impairment in the plant impacting subscribers is critical. Troubleshooting the CMTS is a feature which must not be overlooked in any purchasing decision of a CMTS. If one is unable to quickly troubleshoot the CMTS and network, then the CMTS is failing at its main mission of delivering data to and from the subscriber.

6.1. Troubleshooting a Legacy CMTS

Legacy CMTSs have had over 20 years for vendors to harden the CMTS, identify internal bugs and resolve them. However, over the past 20 years the DOCSIS specification has continued to evolve meaning more complexity, features, higher density, and increased power use. Problems can and do occur in the CMTS hardware as hardware degrades over time. Legacy CMTS vendors include CLI and SNMP commands that allow one to continuously monitor all aspects of the CMTS for hardware failures.

6.1.1. Hardware Failures

As an example, a typical scenario may be the failure of a power supply. Using simple CLI commands one can see the status of all power supplies and the total power consumed by the CMTS. Figure 12 shows that the current CMTS supporting roughly 5k subscribers with 42 service groups is consuming about 4.4kW of power.

```

cbr8#show environment power
=====
Slot      Controller      Value
-----
P0        PEM Power       771 W
P1        PEM Power       790 W
P2        PEM Power       782 W
P3        PEM Power       771 W
P4        PEM Power       785 W
P5        PEM Power       760 W
-----
Input Power Summary:  4659 W
=====
1         FRU Power       380 W
7         FRU Power       390 W
3         FRU Power       380 W
9         FRU Power       370 W
2         FRU Power       380 W
6         FRU Power       380 W
8         FRU Power       370 W
R1        FRU Power       693 W
R0        FRU Power       725 W
0         FRU Power       320 W
-----
Power Consumed Summary: 4388 W
=====
More Cards can be supported:
-----
LC:                               0
=====

```

Figure 12 - Power Overview of Legacy CMTS (4.388 kW used)

Figure 12 is only a high-level summary. One can drill down to detailed diagnostics of the overall powering of the system. Figure 13 shows a sample of another typical CMTS command where details of temperature and voltages can be observed on the legacy CMTS.

```

cbr8#show environment all
Sensor List: Environmental Monitoring

```

Sensor	Location	State	Reading
I2_CUR: Sens	1	Normal	20 mV
I2_CUR: Vin	1	Normal	12725 mV
I2_CUR: ADin	1	Normal	271 mV
G0_CUR: Sens	1	Normal	73 mV
G0_CUR: Vin	1	Normal	12575 mV
G0_CUR: ADin	1	Normal	0 mV
G1_CUR: Sens	1	Normal	74 mV
G1_CUR: Vin	1	Normal	12625 mV
G1_CUR: ADin	1	Normal	0 mV
LB_CUR: Sens	1	Normal	20 mV
LB_CUR: Vin	1	Normal	12575 mV
LB_CUR: ADin	1	Normal	0 mV
Temp: CAPRICA	1	Normal	54 Celsius
Temp: BASESTAR	1	Normal	58 Celsius
Temp: RAIDER	1	Normal	55 Celsius
Temp: CPU	1	Normal	34 Celsius
Temp: INLET	1	Normal	25 Celsius
Temp: OUTLET	1	Normal	46 Celsius
Temp: DIGITAL	1	Normal	35 Celsius
Temp: UPX	1	Normal	45 Celsius
Temp: LEOBEN1	1	Normal	48 Celsius
Temp: LEOBEN2	1	Normal	50 Celsius

Figure 13 - Example of Detailed Powering and Temperature Statistics on Legacy CMTS

In Figure 13 it is evident that much greater detail is available on legacy CMTSs. This data can be extracted from the CMTS and stored in a database, then plotted over time. Trends which would indicate the CMTS is failing giving the cable operator time to replace the failing element or the entire CMTS if necessary. Further, nearly every element of a legacy CMTS is redundant. These include power supplies, processing cards, NICs, and even the RF cards (if desired). It is standard for legacy CMTSs to operate in high availability (HA) mode, such that if any single component or element fails, there is a backup to take over. This generally makes legacy CMTSs very reliable.

6.1.2. Software Failures

Because legacy CMTSs runs on proprietary code, which for most vendors has been built on years of development, software failures are uncommon. This does not mean that there are not compatibility issues between CMTSs and cable modems. It does happen and more frequently it has occurred in new standards, such as with DOCSIS 3.1 OFDMA. These are not so much software failures as they are software “issues”.

6.1.3. Subscriber Impairments

Troubleshooting subscriber impairments is a key feature for any legacy CMTS. Because legacy CMTSs are standards-based, they support standard management information base (MIBs) providing rich access to SNMP data. This SNMP data is used by reactive and proactive monitoring systems which provide the health and key performance indicator (KPI) metrics of each subscriber as well as the aggregate metrics of the network. Reactive and proactive monitoring systems give cable operators the visibility to know when subscribers are having problems or will have problems in the future. The aggregated view also enables cable operators to have visibility into the total traffic consumption on any given fiber node or an entire CMTS. This is important information to monitor and trend over time so that capacity expansion may be planned, as discussed in section 5.1.

SNMP can create a load on legacy CMTSs. Reactive and proactive systems need a lot of data from CMTSs to do their job. At the same time, legacy CMTSs have finite processing power, of which most is dedicated towards managing subscriber traffic. SNMP queries to the CMTS will impact the legacy CMTS processor. If there is too much strain on the CMTS from subscriber traffic, excessive SNMP queries may result in no response to the SNMP queries or even worse, SNMP queries may impact subscriber traffic. If the former, no SNMP data is returned, then the monitoring system will return false data, possibly causing the cable operator to react to bad data. In the latter case, excessive SNMP queries have been known to impede subscriber traffic and even crash a CMTS. This is quite serious and legacy CMTS vendors, monitoring software vendors and cable operators must be aware of these concerns and always ensure their CMTS are not overloaded by either subscriber traffic or SNMP utilization.

6.2. Troubleshooting the vCMTS

One of the benefits of virtualizing a cable headend is that it enables cable operators to collect huge amount of telemetry - much more than what was traditionally exposed by hardware appliance network functions - including hardware, vCMTS software, and environmental equipment. If the telemetry from vCMTS infrastructure is used efficiently it can provide a holistic understanding of infrastructure and services running on top of it. Telemetry makes it possible to achieve closed-loop automation, streamline root cause analysis, perform timely reactive maintenance, effectively plan for proactive maintenance, and much more.

The vCMTS and the platform telemetries are available through open and industry-standardized interfaces, so it can be used to feed a wide range of applications and workflows [3]. There are many telemetry data collection open-source software available that use various plugins to gather metrics from a variety of sources, including COTS servers and software applications like vCMTS. Some common metrics collection software widely used today are Collectd [4], Telegraf [5] and cAdvisor[6]. This telemetry can be integrated with various monitoring solutions that can help remediate platform and vCMTS-related issues quickly. These metrics collection daemon is typically run on non-dataplane central processing unit (CPU) cores so the metric collection processes doesn't interfere with data plane performance. The metrics are collected periodically and stored in a time-series databases like Prometheus [7], InfluxdB [8], etc. These metrics from the hardware and the vCMTS application can be visualized on a single or multiple Grafana dashboards. Grafana alerts can be created in a single and consolidated view based on certain pre-set conditions, and application and platform metrics thresholds, that make headend management easier, enhance recovery time, and reduce service downtime significantly.

Like legacy CMTSs, vCMTSs can also have hardware failure from time to time. Since a vCMTS runs on commodity servers the common hardware failures are related to power supply, memory dual in-line memory modules (DIMMs), NIC port, disk corruption etc. One great benefit of virtualization is that you can have redundancy available at almost all levels with properly configured HW and closed-loop automation. Intel servers provide a great deal of telemetry from different server components and if monitored properly can indicate hardware failures beforehand. This lets cable operators plan maintenance proactively and isolate bad hardware to mitigate downtime significantly. The hardware telemetry can also be used to enhance scheduling decisions in K8s which enables intelligent placement of vCMTS SGs pods based on up-to-date

platform telemetry and vCMTS workload requirements. If one of the K8s nodes reports any bad hardware metrics, the K8s scheduler will automatically cordon off those nodes and de-schedule SGs from that node and move them to a healthy node.

One key consideration for cable operators is to use automation to manage their virtualized headend's network operations. Automation helps to manage growing and changing networks, fix problems faster, and helps adhere to customer service level agreements (SLAs). To perform automation effectively, it requires end-to-end monitoring of software, services, and the hardware on which these services are running within the network. Intel server telemetry spans a vast number of domains including utilization, power consumption, fault detection, and performance. To offer meaningful insights from this information, Intel has created a portfolio of telemetry reports that provides actionable data about the current status of the server [9]. Combining these insights with vCMTS performance data allows for a more holistic view of the vCMTS network function. These reports can be used to help automate orchestration, self-healing, and energy optimization. There are four telemetry reports currently being developed:

- Platform health; covers overall platform health covering compute, memory, storage, and network interfaces
- Utilization; covers platform utilization and capacity indicators
- Congestion; covers CPU overload or network congestion scenarios
- Configuration; covers platform misconfigurations

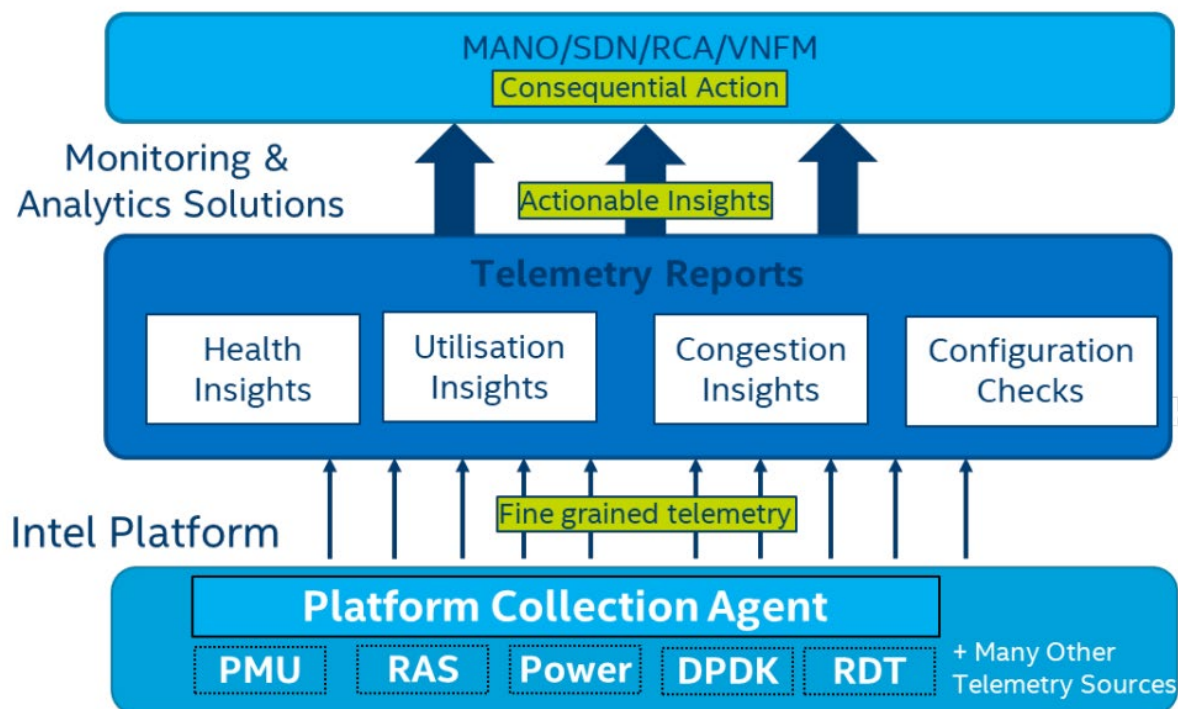


Figure 14 - Telemetry and insights data flow

Figure 14 shows how the telemetry reports use Intel server telemetry provided by Collectd and Telegraf via intel plugins to a monitoring solution that forms the input for the reports. The output

of the reports can be read by an operator or consumed by other management systems. The cable operators can feed these insights generated by the reports into their monitoring systems, which are then processed by online or offline automated systems and can also be used for visualization purposes. Figure 15 (below) illustrates a simple Grafana dashboard visualizing two of the memory health reports, availability and errored seconds. Errors can be documented in this instance and clear indications of these errors can be visualized. [9]




Figure 15 -Grafana dashboard visualizing two memory reports

In addition to the platform metrics in a vCMTS, Grafana dashboards can provide metrics to diagnose and examine subscriber issues. This can be seen in Figure 16 below from a cnBR show cable modem PHY dashboard. This displays the PHY output of the cable modems, just like running the command in CLI on a legacy CMTS.

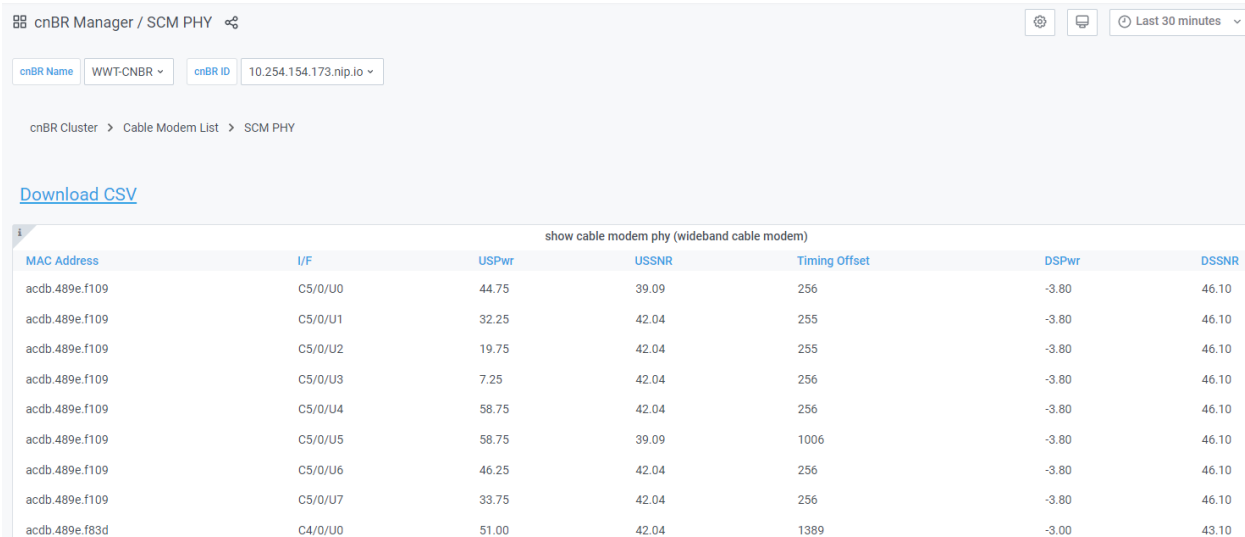


Figure 16 - Show Cable Modem PHY dashboard

Using the metrics collected by the vCMTS software, troubleshooting issues follows the same paths as legacy CMTS systems. The information is provided from the various containers and displaying them in easily navigable outputs. Subscriber impairments can be investigated by examining an individual subscriber and the reported information from the modem. The way this information is presented can vary by vendor, with the subscriber information in Figure 19 being displayed by Harmonic's CableOS, while Figure 20 displaying similar information as reported by Cisco's cnBR. While these software packages come with default dashboards, since these are using Grafana, the operator is able to modify, or completely create new dashboards as needed.

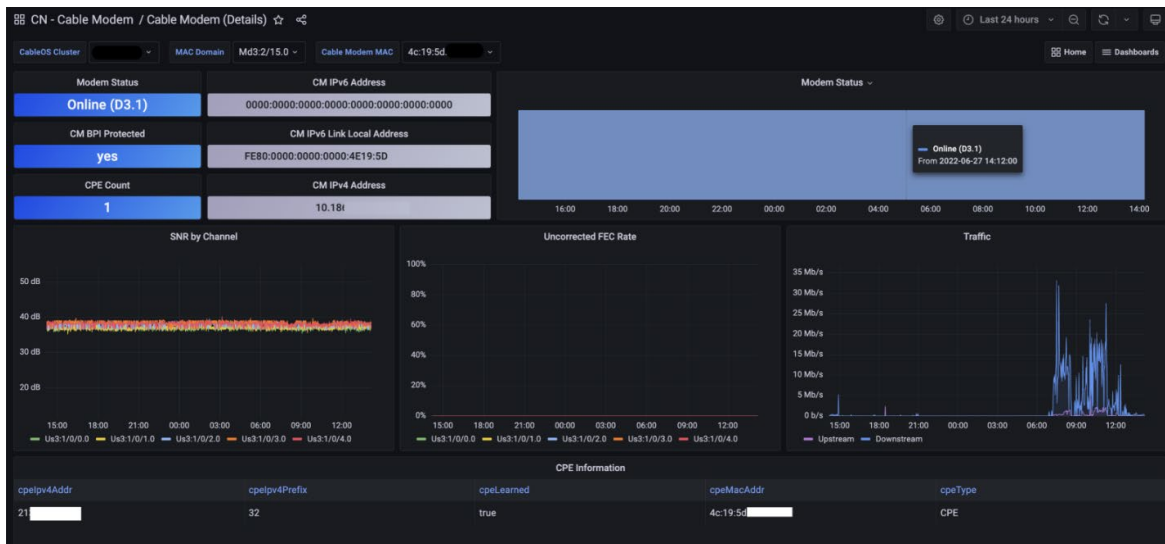


Figure 19 - CableOS dashboard displaying modem detail

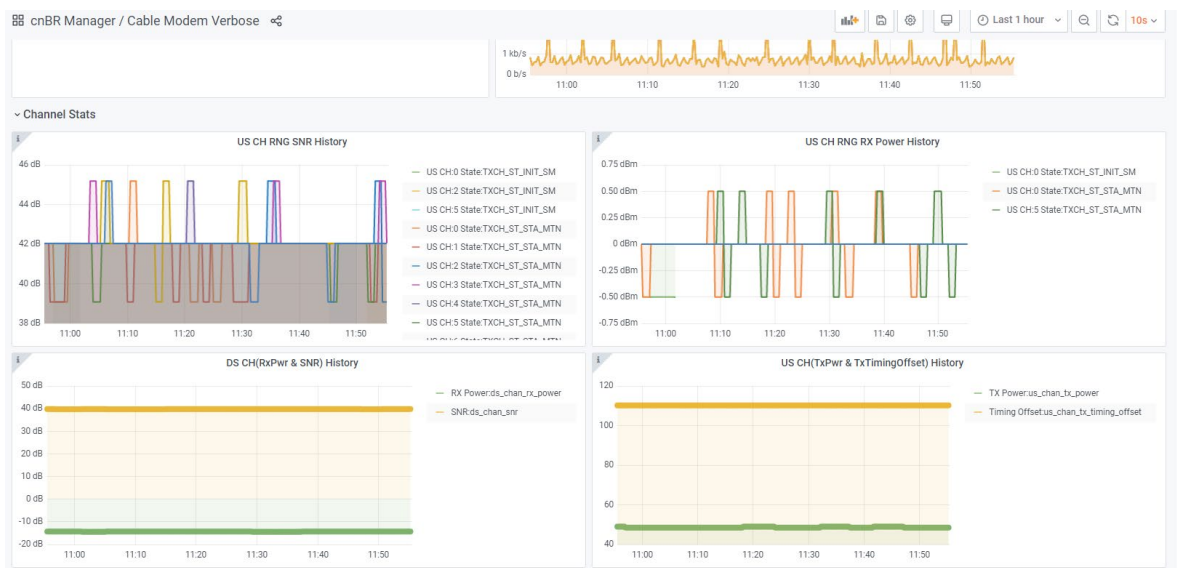


Figure 20 - cnBR dashboard displaying modem detail

A vCMTS ultimately provides the same information that a legacy CMTS provides. As such when diagnosing DOCSIS issues the troubleshooting techniques likewise follow the same steps that are performed in section 6.1. The difference being that the information is provided in an easier to digest form using GUI information, and not requiring SNMP or CLI intervention to obtain the data, thus simplifying operations. However, new telemetry-based tools, such as a Telegraf, InfluxDB and Grafana (TIG) stack, along with a new set of skills and training may be required to take advantage of the improved reporting features available in the vCMTS.

7. How do I Multi-task with a vCMTS?

The disaggregation of the PHY layer-1, MAC layer-2, and IP layer-3 of the CMTS makes it possible for cable operators to scale their access networks - virtually infinitely. As we have seen with the adoption and deployments of DAA, cable operators are reaping the benefits of running the DOCSIS PHY in the field by deploying RPDs. Upstream RF signals are demodulated at the RPD. This means RF signals are transported back digitally to the CMTS. Combining RF signals in racks of RF-combining equipment is no longer required. This has benefits for cable operators looking to reclaim headend space, reduce power and cooling costs in the headend.

The standardization of DAA for DOCSIS and further advancements in the flexible MAC architecture (FMA) standard have enabled the transition to a software-centric cable network infrastructure. [2] One option that the DAA and FMA architectures allows for the DOCSIS MAC software to be deployed as a virtual network function (VNF) on general purpose x86 servers in a cable operator headend as a vCMTS, while the DOCSIS PHY layer is housed in an HFC node near subscribers and business customers.

7.1. What is required to run the DOCSIS workload in software on a general purpose x86 server?

In Figure 17 below, we illustrate the typical components that make up a virtual CMTS deployment. There are hardware and software elements needed to deliver the DOCSIS workload to cable modem users.

The **hardware** required is COTS hardware that is typically found in a data center or server room of companies across all sectors of today's information age such as telecom, media & entertainment, manufacturing, medical & health services, automotive, and financial services. The servers are sold by original equipment manufacturers (OEM) such as HPE, Dell, Cisco, Lenovo, SuperMicro, Intel, and others. The NIC are made by various OEM suppliers such as Intel, Mellanox, Broadcom, Marvell, Cisco, and others. The data center switches are made by OEM vendors such as Arista, Cisco, Dell, Extreme, HPE, Juniper, and others. The hardware comes in various sizes and configurations to meet the compute and networking requirements to serve the workload in that location of the network topology. The benefit of running the DOCSIS workload in software on general purpose servers and switches is the economies-of-scale.

The **software** elements listed in the Figure 17 below vary depending on the vCMTS software selected by the cable operator. Typically, different software is needed for the OS, the container orchestration, management, automation, monitoring, analytics, and the vCMTS VNF. Some of the software to run the server platform and container network function is available in opensource and has been proven to scale to billions of end-users.




Figure 17 - Anatomy of a Virtual CMTS deployment

The advent of SDN and NFV has accelerated innovation for the separation of the data plane function, the control plane function, and the management functions in the delivery of networking applications. Containerization allows further disaggregation of compute, networking, and storage resources needed to move, process, and store packets and data at a much more granular level than previously possible with purpose-built, fixed-function, hardware. An example of this is a CMTS that is built to serve the function of modulating and de-modulating DOCSIS encapsulated packets, processing the packets, and moving them between end-users and the Internet. At the turn of the century, the legacy CMTS platforms were built with shared compute resources to service pre-defined network segments such as an upstream port or a downstream port. In the new world of container networking running on commodity servers, more control and freedom is given to the cable operator to allocate (i.e. orchestrate) compute and networking resources required to service the customer-defined segments. The cable operator can allocate the compute resources down to the granular level of threads in the CPU cores, and the amount of IO on the NIC logical ports at the service group, the upstream, and the downstream level to maximize resource allocation needed per service element. Not all service groups are the same, therefore, the compute and IO resources to serve them can be customizable.

Servers can be purchased as single socket or dual socket systems. Meaning, the server can be configured to have 1 CPU or 2 CPUs in the same system that houses the power supplies, fans, memory, and NIC cards. Each CPU will have multiple CPU cores. A CPU core is a virtual CPU that is a self-contained unit of compute within a CPU. Each CPU core can perform different tasks. Today's servers have CPU sockets, CPU cores and NIC interfaces that do the compute processing and network movements of the data packets. The vCMTS software solution typically has default allocations for CPU sockets, CPU core, threads, and NIC partitions that are pre-defined for ease of operation of the DOCSIS function that cable operators adopt "out of the box."

If desired, an operator could potentially orchestrate the container functions and land a DOCSIS workload on one of the two CPUs in a dual-socket server that services 10 DOCSIS service groups in a small town, while orchestrating the other CPU socket to run a cloud Digital Video Recorder (cDVR) function with some cloud storage resources. Even more granular, an operator may choose to mix the CPU cores in the CPU socket to do different functions. For example, if the operator was running 2 service groups with a few of the CPU cores and decided that there are CPU cores available for running the virtual optical line terminal (OLT) function of the passive optical network (PON) network, or, for terminating the Soft Generic Routing Encapsulation (Soft-GRE) tunnels for the community WiFi service, or for running another service, the operator can do so with flexibility. The x86 CPU cores and NIC cards and ports, coupled with container networking software allows operators much more control, flexibility, and optionality to maximize usage of the computing resources. This flexibility to allocate compute, IO, and storage resources at a much more granular level allows cable operators to scale their networks much more cost-effectively while converging multiple services on commercially available servers.

8. Does a vCMTS Consume More or Less Power Than a Legacy CMTS?

When migrating from legacy CMTS to vCMTS, cable operators will realize significant reductions in power consumption. Depending on the configuration of the vCMTS server clusters, Cable operators will see reductions in power that are 3 to 5 times less than legacy CMTS. Furthermore, coupled with the reduction in vCMTS power consumption and the associated heat being generated, cable operators will see significant reductions in power consumption of the heating, ventilation and air conditioning (HVAC) cooling systems. The reduction in rack space to house vCMTS servers to service the same, or more, service groups will also result in significant cost reductions.

When comparing power consumption of legacy CMTS versus vCMTS, the units of measurements will be normalized down to a cable service group. Below is a chart comparing a legacy CMTS operating in I-CCAP mode, a legacy CMTS performing the DOCSIS MAC functions connected to a remote PHY device, and a vCMTS doing the DOCSIS MAC processing connected to a remote PHY device.

On a per service group per watt comparison, without including the HVAC cooling reductions, a vCMTS uses 2.5 times less power than a distributed converged cable access platform (D-CCAP) and 3.8 times less power than an integrated converged cable access platform (I-CCAP). When comparing the amount of DOCSIS throughput delivered per megabit per second (Mbps) per watt, a vCMTS delivers nearly 10 times as much bandwidth as an I-CCAP.

There are additional power conservation possibilities when running vCMTS software on Intel servers. One example that cable operators can employ is to turn down the CPU clock frequency during off-peak network periods to reduce the power consumption of the vCMTS servers.

Table 2 - Power consumption comparison

	Scenario 1	Scenario 2	Scenario 3
	I-CCAP CBR-8	D-CCAP CBR-8	vCMTS 3 servers + switch + timing server
Power consumed	5,200 watts	5,000 watts	2,635 watts
Service Groups served	56	84	108
Watts per Service Group	92.8	59.5	24.4
Rack space	14 RU	14 RU	8 RU
Aggregate Throughput	100 Gbps	100 Gbps	480 Gbps
Mbps per watt	19.23	20.0	182.2

Notes:

- For this comparison, speed of 3 Gbps downstream and 200 Mbps upstream configurations for each service group were used.
- The RPD power consumption is not included as it will be the same for scenarios 2 and 3. Scenario 1 would have multiple racks of required RF cable combining equipment that are not included in this power consumption analysis of the DOCSIS MAC processing.
- The vCMTS configuration in scenario 3 includes three (3) dual-socket servers powered by Intel 3rd Generation Scalable Processor CPUs, one (1) Cisco Nexus 3232c leaf switch, and one (1) timing server. Each server consumed 635 watts while producing 160 Gbps of aggregate active throughput and powering 36 active service groups and 12 standby service groups.
- Aggregate throughput of I-CCAP and D-CCAP are limited to 100 Gbps due to physical limitations of the I/O cards. Similarly, the aggregate throughput limit of the vCMTS server is the NIC capacity, which is typically 100 Gbps, 200 Gbps, or 400 Gbps per server, as procured today.

9. Conclusion

In conclusion, we hope the readers will appreciate that operating a vCMTS is not much different than operating a legacy CMTS. Some would argue that the web interfaces for managing a vCMTS are easier to navigate than the CLI of legacy CMTSs, similar to the transition from MS DOS to Windows in 1995, which made using a personal computer (PC) much easier.

As the CMTS workload is performed in software, cable operators can run the vCMTS software on commodity servers. These general-purpose servers are used for running a plethora of applications such as DNS, email, video streaming, web services, chat, user authentication services, firewalls, security, billing, monitoring, video encoding, video transcoding, and so many other applications. Thus, millions of servers are purchased and deployed annually by companies, big and small, to run their business services. By running the DOCSIS workload on commodity servers and switches, cable operators can buy the hardware at a fraction of the capital cost of purpose-built hardware. Furthermore, if no longer needed for running DOCSIS, cable operators can repurpose the commodity servers to run other functions such as email, video streaming,

proxy caching, or SD-WAN, or any other software application on the same CPU cores in the servers located at that location in the network.

CMTS hardware has served cable operators well to win market share in the burgeoning broadband market over the past 25 years. Compared to legacy CMTS equipment, vCMTS software running on servers can do the job at a fraction of the hardware cost, consume a fraction of the rack space, use a fraction of the power, and allow cable operators more flexibility to quickly scale and remain competitive.

Abbreviations

API	application programming interface
bps	bits per second
BIOS	basic input output system
BGP	border gateway protocol
BMC	baseboard management controller
CDVR	cloud digital video recorder
CLI	command line interface
CM	cable modem
CMTS	cable modem termination system
CNF	container network function
COTS	commercial off-the-shelf
CPU	central processing unit
DAA	distributed access architecture
DHCP	dynamic host configuration protocol
DIMM	dual in-line memory module
DOCSIS	Data-Over-Cable Service Interface Specification
DPDK	data plane development kit
FEC	forward error correction
FMA	flexible MAC architecture
GHz	gigahertz
GUI	graphical user interface
HD	high definition
HE	headend
HFC	hybrid fiber coax
HVAC	heating ventilation and air conditioning
Hz	hertz
IO	input-output
ISV	independent software vendor
JSON	JavaScript object notation
K8s	Kubernetes
KPI	key performance indicator
MAC	media access control
MANO	management and orchestration
MIB	management information base
MSO	multiple system operation
NE	network element
NFV	network function virtualization
NFVO	network functions virtualization orchestration
NIC	network interface card
OEM	original equipment manufacturer
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OLT	optical line terminal
OS	operating system

PHY	physical layer
PMU	performance monitoring units
PON	passive optical network
PTP	precision timing protocol
PXE	preboot execution environment
QAM	quadrature amplitude modulation
RAS	reliability, availability, serviceability
RCA	root cause analysis
RDT	Resource Director Technology
RF	Radio Frequency
RMD	Remote MACPHY Device
RPD	Remote PHY Device
R-PHY	Remote PHY
RU	Rack Unit
S-CDMA	Synchronous Code Division Multiple Access
SC-QAM	Single Carrier QAM
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networking
SGs	Service Groups
SLA	Service Level Agreement
SNMP	simple network management protocol
SoftGRE	soft generic routing encapsulation
SR-IOV	Single Root I/O Virtualization
SSH	Secure Shell
TDMA	Time Division Multiple Access
TFTP	Trivial File Transfer Protocol
TIG	Telegraf, InfluxDB, and Grafana
US	upstream
vCMTS	Virtual CMTS
VNF	Virtual Network Function
VNFM	virtual network functions manager
YAML	yet another markup language

Bibliography & References


[1] DOCSIS 4.0 - A Key Ingredient of the 2030's Broadband Pie, A Technical Paper prepared for SCTE by Maricevic, Andis, Cloonan, Ulm, 2021

[2], B. Ryan, M. O'Hanlon, D. Coyle, R. Sexton and S. Ravisundar, "Maximizing vCMTS Data Plane Performance with 3rd Gen Intel® Xeon® Scalable Processor Architecture," [Online]. Available: <https://networkbuilders.intel.com/solutionslibrary/maximizing-vcmts-data-plane-performance-with-3rd-gen-intel-xeon-scalable-processor-architecture>

[3] <https://www.intel.com/content/dam/www/public/us/en/documents/guides/nt-telemetry-eguide-for-web-team.pdf>

[4] <https://collectd.org/>

- [5] <https://www.influxdata.com/time-series-platform/telegraf/>
- [6] <https://prometheus.io/docs/guides/cadvisor/>
- [7] <https://prometheus.io/>
- [8] <https://www.influxdata.com/>
- [9] <https://builders.intel.com/docs/networkbuilders/telemetry-reporting-for-network-infrastructure-solution-brief.pdf>



Towards Predictable Low Latency DOCSIS Services

A Technical Paper prepared for SCTE by

Dan Rice

VP Access Architecture and Technology, CONNECT
Comcast
4100 E Dry Creek Road, Centennial CO
daniel_rice4@comcast.com

Sebnem Ozer, Ph.D.

Senior Principal Architect, CONNECT
Comcast
sebnem_ozero@comcast.com

James Martin, Ph.D.

Professor Emeritus
Clemson University
Jmarty@clemson.edu

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Low Latency DOCSIS: Promising Results	3
3. Predictable End-To-End Latency	12
3.1. E2E Marking.....	12
3.2. End-to-End Measurements	12
3.3. End-to-End Optimizations	13
3.4. Network Segments with Classic ECN	14
3.5. Standards Compliance and Larger Ecosystem Support.....	14
4. Conclusion.....	15
Acknowledgments	16
Abbreviations	16
Bibliography & References.....	16

List of Figures

Title	Page Number
Figure 1 – LLD Lab Setup	4
Figure 2 – Idle Latency Analysis	5
Figure 3 – US LLD example showing the impact of QB traffic on NQB traffic latency	6
Figure 4 – US LLD example with two marked NQB traffic flows of the same rates in the CM	7
Figure 5 – Latency statistics for US LLD example with two marked NQB traffic flows of the same rates in the CM	8
Figure 6 – US LLD example with two marked NQB traffic flows of different rates in the CM	9
Figure 7 – Latency statistics for US LLD example with two marked NQB traffic flows of different rates in the CM	10
Figure 8 – Latency statistics for US LLD example with PGS	11
Figure 9 – Utilization statistics reported with 15 sec and 5 min windows	13
Figure 10 – LL Service and Network Architecture	15

1. Introduction

The Internet is based on fundamental design goals that emphasize simplicity and reliability. Early on, the idea was that instead of a centrally controlled network that was a gatekeeper on new features (such as telephony ‘star’ codes) most of the intelligence was decentralized into the edge systems to help enable a tremendous wave of permissionless innovation at the application layer. Any complex network-oriented processing was to be performed by the host devices. A fundamental difference of today’s Internet is ‘predictable service qualities’, which can be achieved only by ‘smart’ components in an end-to-end path. Over the years, applications have adapted as well. A classic success story is Internet streaming. Through maintaining an appropriate size playback buffer and aided with adaptive bitrate control, HTTP-based adaptive streaming (and UDP equivalents) is the dominant application.

As society moves from the information age to the age of the M2M communications, we anticipate further significant changes to the Internet as well as significantly different applications. Machines are making decisions based on data, many times in real-time. There are wired and wireless scenarios to consider. Many of these applications have yet to be invented, but based on early examples such as cloud gaming, enhanced videoconferencing, coordinated autonomous vehicles, drone swarms working in a coordinated manner to carry out missions, multiuser virtual-reality gaming, the Internet needs to be re-examined to determine what must be addressed so that technological breakthroughs in devices and systems are not held-back by Internet performance limitations. These concerns have motivated the key standards bodies to pursue broad initiatives that will enhance their respective technology’s ability to support emerging applications systems that require predictable service qualities.

Cable operators are evaluating Low Latency DOCSIS (LLD) to determine how the new technology might improve subscriber’s perceived quality and how an operator might leverage the technology for new services [1, 2, 3].

In this paper, we first present LLD lab analysis that shows promising results. We then discuss end-to-end factors that are critical in providing predictable end-to-end LL services. We conclude our paper with an architecture including network and service components to deploy an effective LL system.

2. Low Latency DOCSIS: Promising Results

The term Low Latency DOCSIS (LLD) reflects the general Internet and cable community’s efforts to offer an architecture that satisfies the requirements for both non-queue building (NQB) traffic and queue-building (QB) traffic over the same physical network [4, 5, 6]. The core ideas behind LLD are novel and show promising results based on preliminary studies and trials. The goal is to enhance the end-user experience and to allow operators or over-the-top service providers to develop new services.

LLD is compatible and aligned with a new generation of Internet flow control, with Low queuing Latency, Low Loss, and Scalable throughput (L4S) technology [4] which supports end-to-end

latency that manages queuing delays with congestion notifications that will result in flow rate adjustments, thus avoiding the packet drops and retransmission as much as possible for the NQB flows. The goal is to transition to substantially lower queuing delays across the network while coexisting with ‘classic’ congestion controls used widely in the Internet today. This will become important not only to the development of new applications as described here, but also to improve customer experience as sufficient network speeds are now becoming widely available to all Internet users and higher speeds may no longer improve the Internet experience, especially for residential consumers.

LLD aims to support non-queue-building applications that send data quickly and don’t cause latency in the presence of queue-building applications from and to the same subscriber’s modem. The main components are coupled dual-queue with weighted scheduler, proactive grant scheduling and optimized MAP timing and channel settings [6]. The aim is to support ~10ms DOCSIS round-trip-time (RTT) latency for the 99th percentile of LL flows, a ~10X improvement to today’s DOCSIS deployments.

The initial tests show promising results. The following upstream examples are discussed to describe the benefits of LLD in the presence of high home network utilization. The lab tests are carried over an iCMTS with deployment configuration and D3.1 CMs supporting LLD features. A simplified version of the setup is shown in Figure 1.

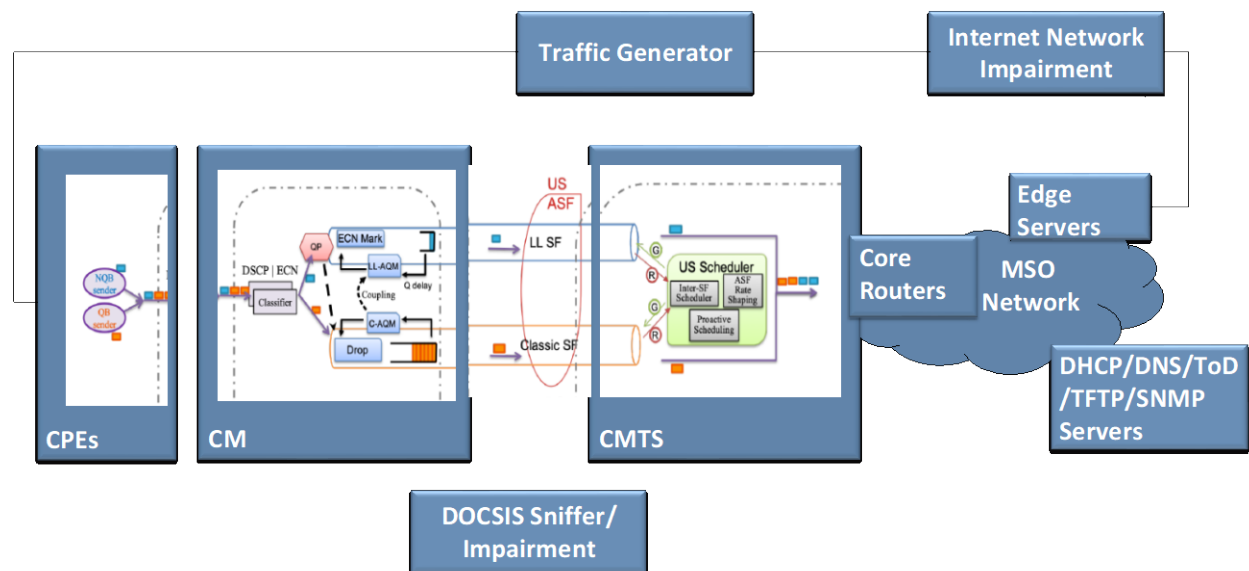


Figure 1 – LLD Lab Setup

The queuing and media access are the major DOCSIS latency sources while propagation, serialization, encoding, and switching also contribute to the latency. Figure 2 displays example idle latency values for a VoIP type traffic in the upstream direction. Idle latency refers to latency when there are no other simultaneous home traffic flows competing with the test traffic for the shared resources. Delays due to propagation, FEC, ATDMA Interleaving/OFDMA Framing,

CM/CMTS MAP processing, queuing (\sim MAP timing/2) and other processing are computed for a total latency estimation as shown in blue line in Figure 2. A traffic generator is used to create VoIP type traffic and measure latency as shown with circles in the same figure. As expected, propagation latency, described by the data point annotation, is the major factor depending on the distances between the CM and the US scheduler as a function of CMTS. Today, most headend-to-CM distances are smaller than 160 km, deployed with typical DOCSIS timings, which provides good latency performance. If better latency bounds are targeted, US scheduler can be deployed closer to the subscriber as a microservice or proactive grant scheduling techniques can be integrated.




Figure 2 – Idle Latency Analysis

The idle latency provides important information on the achievable minimum latency and typical values when LL traffic is not affected by other traffic in the CM. Today in many homes there are many users or sources of traffic that use the network simultaneously with work and school from home that increased dramatically in recent years. As a result, while setting a lower bound on

latency performance, is a less effective metric for user experience. However, LL services require consistent performance values such as jitter.

Latency under load or working latency computed during high home network congestion levels provide upper bounds for latency and jitter a LL service may be subject to, depending on the speed tier rate and conditions of the shared resources. Figure 3 displays an example use case where two audio flows are generated for the same CM, one marked with LL DSCP value and the other one marked with DSCP 0x00. The dual queue LLD features are tested. The top graph displays the latency over time (green line) for unmarked audio flow while the middle graph displays the latency over time for an audio flow marked per LLD specifications. The concurrent TCP flows 'throughput, representing other traffic flows in the home, is displayed in the bottom graph. As the TCP flows fill up the speed tier bandwidth, both audio flows' latency values are affected proportionally to TCP flows' throughput fluctuations as seen in the figure. However, the unmarked audio flow's latency increases proportional to the TCP traffic RTT as it shares the same queuing and scheduling weight while the marked audio flow's latency is bounded with a ~6X improvement. Not only is the absolute delay much lower, but the variation in delay is within a much smaller window which is a critical key improvement for most applications. Note that the MAP time and other physical layer channel settings are not yet optimized for these results set, and will lower latency even further.



Figure 3 – US LLD example showing the impact of QB traffic on NQB traffic latency

In Figure 4, two marked NQB traffic flows share the same LL queue while QB TCP traffic and unmarked NQB traffic is transmitted through the classic queue for the dual queue approach. In [5], suggested unresponsive NQB traffic load is 1 Mbps or a couple of packets per RTT. The marked NQB flows share the resources fairly with bounded latency results (top two graphs), while unmarked NQB traffic (third graph) has inconsistent latency fluctuating depending on the TCP traffic (bottom graph). The latency histograms and Complement Cumulative Distribution Function (CCDF) are shown in Figure 5, with top two graphs for marked NQB flows, third graph for unmarked NQB flow and last graph for CCDF of all three NQB flows. In this example, there were a few packets with latency close to 19 ms and 99%ile of packets have latency smaller than 18 ms. Further optimization of LLD feature parameters and DOCSIS settings can reduce this latency range further. For the same settings, the latency improvement is close to 5X for 99%ile range. The flows in this simulation are longer than those of Figure 3 where TCP flows are ramping up for most of the simulation time. The latency variation for the unmarked NQB traffic can still be observed in Figure 4 during the steady-state TCP conditions as well.

**Consistent
latency up
to ~18ms
for 99%ile
(with LL
marking)**

**Inconsistent
latency up to
~86ms for
99%ile
(without LL
marking)**

**Concurrent
TCP traffic**

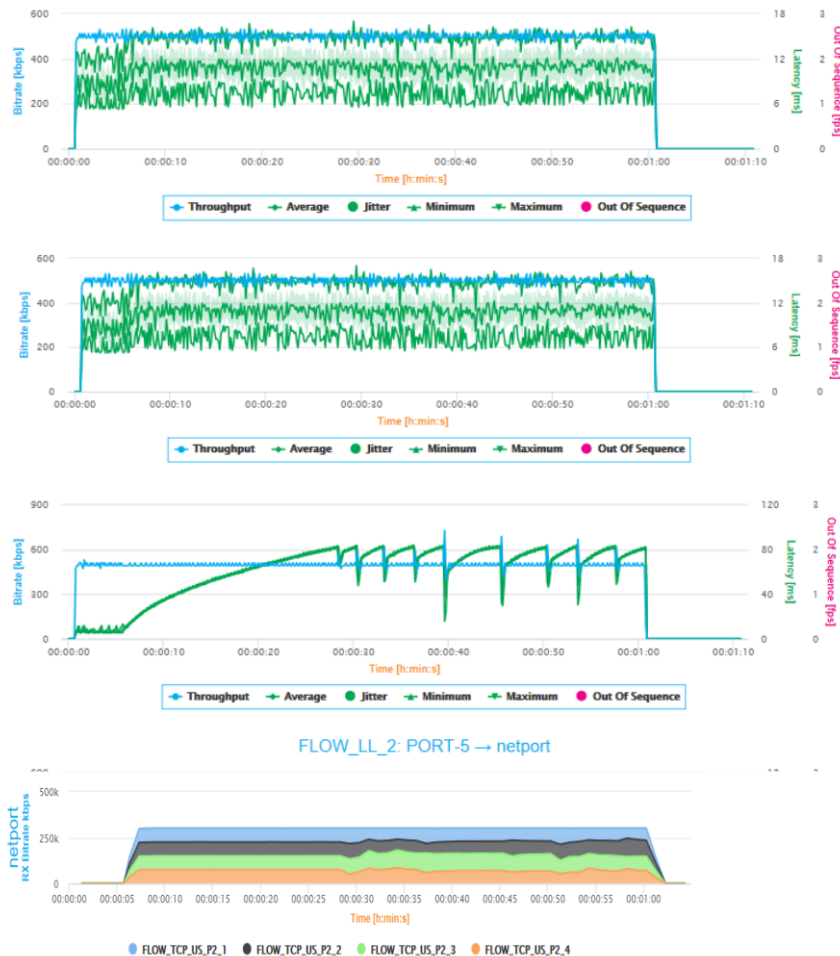


Figure 4 – US LLD example with two marked NQB traffic flows of the same rates in the CM

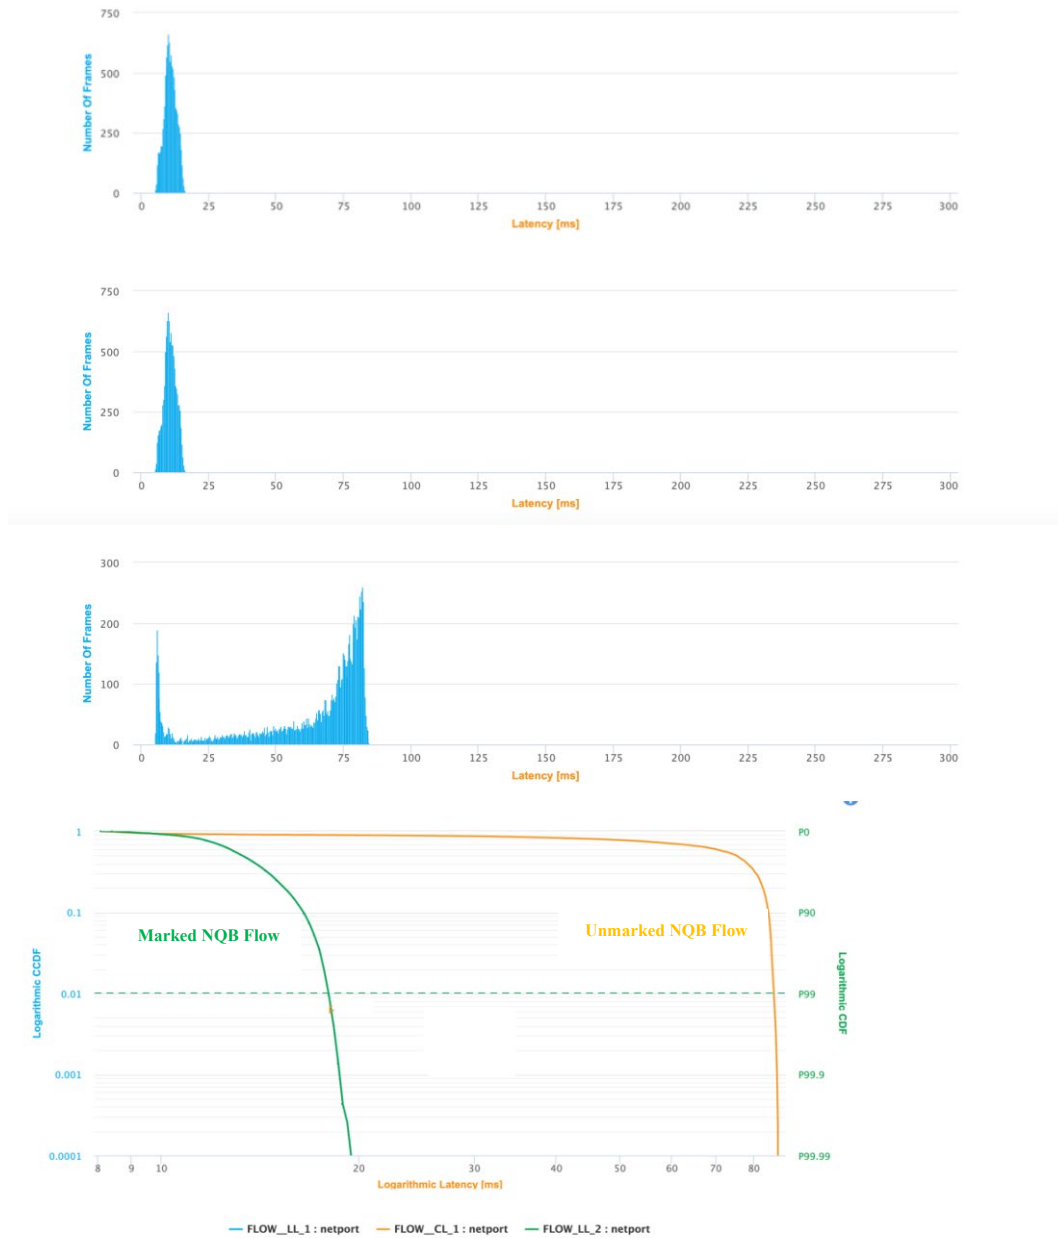


Figure 5 – Latency statistics for US LLD example with two marked NQB traffic flows of the same rates in the CM

A similar use case with two marked NQB traffic flows of different rates are displayed in Figure 6 and Figure 7 to show that the flows are served fairly in terms of latency, jitter and throughput. While improving the 99th percentile latency by ~6x.

Consistent
latency up to
~14.5ms for
99%ile of flows
with 500kbps
and 300kbps
rates (with LL
marking)

Inconsistent
latency up to
~86ms for
99%ile
(without LL
marking)

Concurrent
TCP traffic



Figure 6 – US LLD example with two marked NQB traffic flows of different rates in the CM




Figure 7 – Latency statistics for US LLD example with two marked NQB traffic flows of different rates in the CM

For applications requiring tighter latency bounds or less variations due to use cases such as serving group congestions, Proactive Grant Scheduling as defined in [6] may be preferred. Although PGS eliminates the latency due to request-grant cycle in the upstream direction, it may

create network inefficiency if the grants timing and NQB traffic characteristics are not matched well. In this case, unused grants would decrease network efficiency while less grant allocations might increase the latency. However, as shown in Figure 8, the 99th percentile latency and jitter are further improved by an incremental ~65% with this technique increasing the overall improvement to > 13X. Cable operators may apply PGS for certain services or extend the implementation with traffic activity detection.



Figure 8 – Latency statistics for US LLD example with PGS

3. Predictable End-To-End Latency

As discussed in the previous section, LLD improves the latency and jitter for the NQB traffic. Since NQB traffic doesn't share the same queue as QB traffic, packet loss due to aggressive QB traffic is avoided. While DOCSIS network segment can benefit from LL features, the subscriber's QoE depends on the end-to-end latency. In this section, different factors that cable operators must consider for a predictable latency are discussed. Although latency is mentioned as the main metric, in fact, jitter, packet loss and throughput are analyzed concurrently. Reliability, security and service availability are other metrics that operators consider for a complete service assurance.

3.1. E2E Marking

LLD and LL implementations in other network segments depend on marking unresponsive NQB and L4S traffic with LL DSCP and ECT(1) [4,5]. In traditional systems so far, operators have bleached or remarked marked packets entering their network. Wi-Fi gateways and routers in the upstream direction and interconnect routers and CMTSs in the downstream direction bleach or remark these packets.

Operators must assure their network components do not bleach or remark (except CE marking as explained in [4] and DSCP conversions as explained in [5]) LL packets in their network. Queue protection schemes [4,6] and new security features may be applied to detect NQB traffic violation. Negotiations with peering partners and edge servers may be used to avoid the issue with network components in the segments outside of operator's control.

Although, techniques such as mirroring the marking in the downstream per the received marked packet in the upstream may be used, symmetry is not always guaranteed. Application detection should be used only if subscriber has full control and/or a paid service is deployed.

3.2. End-to-End Measurements

Cable operators have been deploying latency, jitter, packet loss and throughput/speed measurements in their network for idle, LUL/working and real customer traffic [1]. Techniques that can measure latency in the access network as well as latency outside of operator's network has been also deployed [1]. These platforms must be upgraded to measure these metrics per LL marked and classic packets and flows.

These measurement techniques must be analyzed with other monitoring features such as utilization levels, device conditions, routing options, channel conditions and many more resource conditions that can affect e2e latency. As latency targets get lower and/or speed tier rates increase, traditional monitoring windows may not be adequate for this analysis. For example, 5 min averages for channel utilization have been widely used in cable operators' networks.

However, bursts (even microbursts) at smaller windows may have a big impact not only on flow control for very high speed tiers, but also on LL services. Traditional monitoring settings may not catch the correlations and causations. An example is shown in Figure 9 for a serving group with highly variable utilization bursts. 5 min values correspond to 98% of utilization measured within 15 sec and reported every 5 min. 15 sec values correspond to 98% of utilization measured and reported every 15 sec. It can be seen that 15 sec values report high burst rates while 5 min values are smoother. The same behavior may be observed for home network utilization and other channel and network conditions.




Figure 9 – Utilization statistics reported with 15 sec and 5 min windows

LL services QoE may correspond to different QoS levels. Although it is not a straightforward process, operators developing platforms to map QoS levels to LL service QoE will be able to program their networks in an optimized way.

3.3. End-to-End Optimizations

A cable operator's network may be approximated as a tandem queuing and servicing. As shown in the literature, in tandem queuing, the arrival rate and distribution in a network segment depends on the previous network's arrival and service processes. Therefore, analyzing the network segments with interdependency can help the operator to estimate the end-to-end latency and find ways to optimize the network features in an end-to-end way.

Mathematical models for complex tandem queuing system with general arrival and service processes may not exist. Operators may use measurement and machine learning techniques to detect the dependencies of network segments.

Edge servers may be deployed for cases where dependency outside of the operator's network creates an obstacle for certain LL services, depending on the cost and latency analysis.

Another factor to consider is how marking is passed based on the next segment's LL support. For example [5] proposes remarking LL DSCP to a lower DSCP value if the next network segment is not ready to support LL services and a higher DSCP may create unfairness or starvation for other services.

3.4. Network Segments with Classic ECN

Although a rare use case, network bottlenecks employing a shared queue that implements an AQM algorithm that provides Explicit Congestion Notification signaling according to RFC3168 may create issues for the L4S deployment [7]. ECT(1) usage for L4S traffic classification are defined in RFC8311 to update RFC3168. If end-to-end path includes network segments with RFC3168 AQMs, having L4S compliant congestion controlled flows instead of all classic congestion controlled flows may cause a lower throughput of classic congestion controlled flows in the same queue. As explained in [7], this problem is rare and significant only if RTTs are long and rates are high for long running flows. [7] proposes preferred and other options, including using edge servers, upgrading AQM implementations, fairness improvement techniques etc. Operators may use end-to-end monitoring and QoE assessment techniques to discover such cases and deploy options defined in [7]. IETF specifications have recommendations for other cases that may be incompatible for L4S deployments during transitioning to a larger ecosystem support.

3.5. Standards Compliance and Larger Ecosystem Support

Although IETF and CableLabs specifications are the main standards cable operators rely on for LL feature implementations, most cable operators support other technologies in their networks, such as PON, 4G/5G, Wi-Fi and wireless IOT technologies. For a predictable E2E latency values, the standards' LL approaches such as marking and congestion notification techniques must converge. For example, IEEE 802.11be (Wi-Fi 7) specifications introduces new Access Categories and mapping for LL services, that comply with LLD features.

In addition, applications (e.g. online and cloud gaming, videoconferencing, interactive real-time streaming, VR/AR), OS (e.g. Windows, MAC, Linux) and API protocols (e.g. WebRTC, Element) providers must also support LL marking, accurate congestion notification and scalable congestion control techniques. LL services may be supported seamlessly and in a predictable way with larger ecosystem support as every party will win with a common goal.

At the IETF 114, sponsored by Comcast, and held in Philadelphia the week beginning July 23, a Hackathon to test interoperability of L4S, ECN and LLD took place.[9]. There were over 30 participants including cable operators and large Internet application and technology providers

including Google, Apple, Meta, Netflix and NVIDIA among others. During this hackathon network infrastructure for DOCSIS, Wi-Fi and 5G were all tested. Results for LLD traffic with an AppleQu8ic implementation are summarized in Table 1.

Table 1: IETF 114 Hackathon results with very compelling improvements in Latency and Jitter.

Upstream	Classic AQM Network	L4S + LLD Network
Upstream Flows	P99 – 30 msec P99.9 - 125 msec	P99 – 9 msec P99.9 - 10 msec
Downstream Flows	P99 – 56 msec P99.9 - 96 msec	P99 – 1.1 msec P99.9 – 7.8 msec

4. Conclusion

LLD promises predictable low latency service support if end-to-end architecture and deployment strategies are well defined. It requires many service and network components to interact in a harmonious way as shown in Figure 10. To support such an architecture requires disruptive changes in traditional cable operators' networks. Operators should support Software Defined Networking (SDN), Network Functionality Virtualization (NFV), Data-driven automation and service agility for the best LL service support.

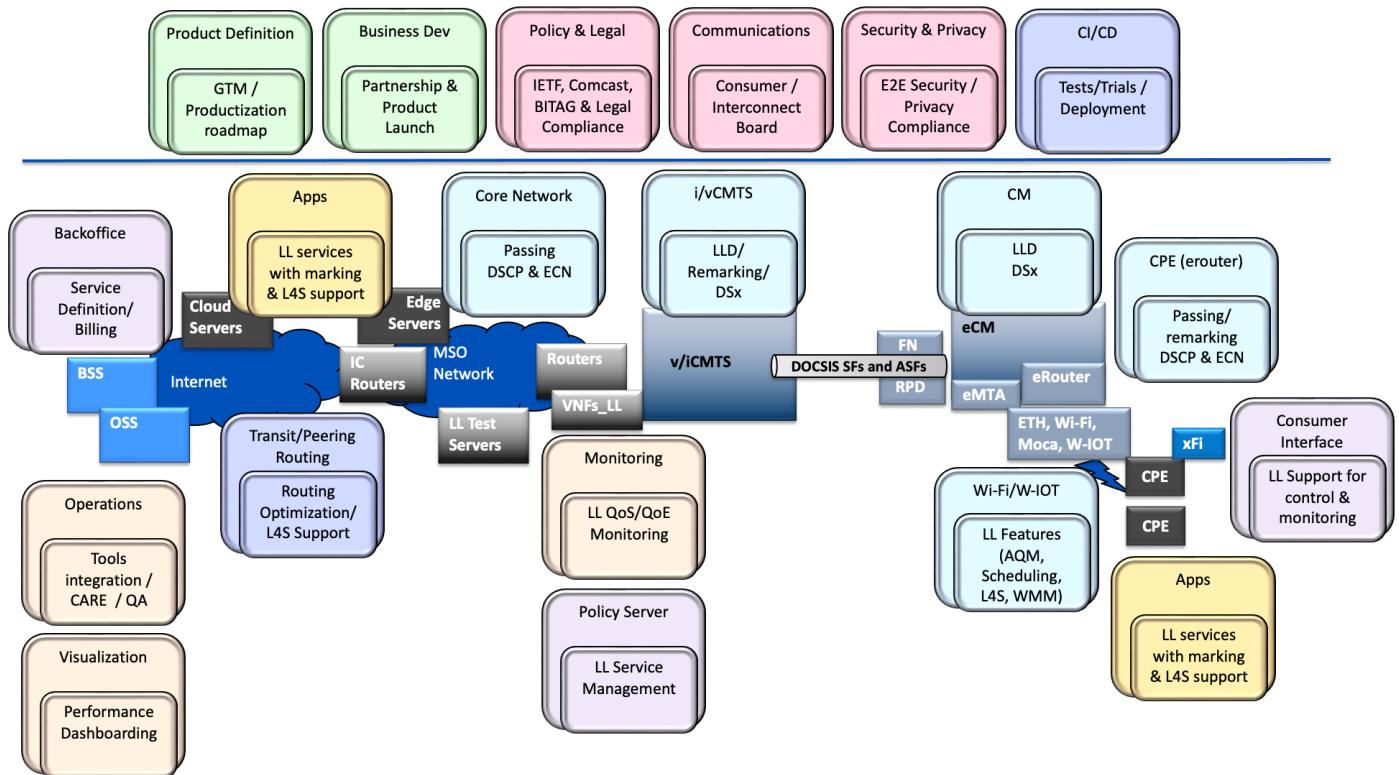


Figure 10 – LL Service and Network Architecture

Acknowledgments

The authors would like to thank Joe McHale, Sumi Chandrashekar, Garrett Miller, Chad Humble for their LLD lab support and Maher Harb for his utilization statistics analysis.

Abbreviations

API	Application Programming Interface
BSS	Business Support System
CM	Cable Modem
CMTS	Cable Modem Termination System
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
eMTA	Embedded Media Terminal Adaptor
HFC	Hybrid Fiber Coaxial
HSD	High Speed Data
MAC	Medium Access Control
MSO	Multiple System Operators
OSS	Operations Support System
PON	Passive Optical Network
QoS	Quality of Service
TFTP	Trivial File Transfer Protocol
ToD	Time of Day

Bibliography & References

- [1] *Fastest Path to Low Latency Services: How Can Cable Operators Deliver Consistent Latency by Following an Efficient and Future-Proof Design Path?*, by S. Ozer, A. Tunstall, C. Klatsky, D. Rice, J. Livingood, J. Chrostowski, J. Raezer, J. Gerson, M. Dolley, P. Sarathy, S. Subbaraj, S. Cho, T. Graffa SCTE Expo 2021
- [2] *Configuring and Deploying Low Latency DOCSIS Networks*, by Greg White & Karthik Sundaresan, SCTE Expo 2021
- [3] *Low Latency Docsis: Concepts and Experiments*, by Tushar Mathur, Ram Ranganathan, Greg Gohman, Bob Zhang, SCTE Expo 2020
- [4] <https://datatracker.ietf.org/doc/draft-ietf-tsvwg-l4s-arch/>
- [5] <https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/>
- [6] *CableLabs DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv3.1, I23, 2022
- [7] <https://datatracker.ietf.org/doc/pdf/draft-ietf-tsvwg-l4sops-03>
- [8] *A Latency Measurement System Using STAMP*, by Karthik Sundaresan, SCTE Expo 2021
- [9] <https://datatracker.ietf.org/meeting/114/materials/slides-114-tsvwg-update-on-l4s-work-in-ietf-114-hackathon-00>

Translating Customer & Employee Experience with Shaw's Data Journey

A Technical Paper prepared for SCTE by

Greg Bone

Principal Architect
Shaw Communications
2400 32 Ave NE, Calgary, AB T2E 6T4
greg.bone@sjrb.ca

Goutam Agarwal

Principal Enterprise Architect
Shaw Communications
2400 32 Ave NE, Calgary, AB T2E 6T4
goutam.agarwal@sjrb.ca

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Key Drivers.....	7
3. Overview of Unified Customer Platform (UCP).....	8
4. Guiding Design Principles.....	10
5. Minimum Viable Product Scope.....	11
6. Components.....	12
6.1. Customer Search.....	12
6.2. Search Results.....	12
6.3. Customer Details View.....	13
6.4. Data API Considerations.....	14
6.5. Snowflake Cloud Data Warehouse.....	15
6.6. API Data Stores.....	15
6.7. Customer Account Linking.....	16
7. Challenges.....	17
8. What's next for Unified Customer Portal.....	27
8.1. Search Improvements.....	27
8.2. Customer Device Details.....	31
8.3. Real-Time Data Loading Process.....	32
9. Conclusion.....	33
Abbreviations.....	34
Bibliography & References.....	34

List of Figures

Title	Page Number
Figure 1 - Pivot in Customer Expectations.....	4
Figure 2 – Shaw's Customer Experience Strategy.....	5
Figure 3 - Shaw's Agent Experience Strategy.....	5
Figure 4 – Key Drivers for UCP.....	7
Figure 5 – High Level System Overview.....	9
Figure 6 – UCP - Minimum Viable Product Scope.....	11
Figure 7 – Top Search Bar.....	12
Figure 8 – Search Results page variation where "Service Type" is shown with text in the Account # column.....	13
Figure 9 – Wireline, no mobile service (in-network wireless quality & eligible for bundle).....	14
Figure 10 – UCP – Challenges.....	17
Figure 11 – Union of Customer Types.....	17
Figure 12 – Example of Getting In-Network Status for a Single Account Number.....	18
Figure 13 – Multiple Account Numbers as Input; Aggregate by Service Provider.....	19
Figure 14 – Keyword Field Tokenization Keeps Josh Smith as a Single Token.....	20
Figure 15 – Text Field Tokenization Separates Name into Two Tokens.....	20
Figure 16 – Removing Leading Zeros from Account Number.....	22
Figure 17 – Account Number Search Query (Leading Zero not Specified).....	22

Figure 18 – Search Results for Query “123456789”	23
Figure 19 – Phone Number Mapping Example.....	24
Figure 20 – Phone Number Analyzer.....	24
Figure 21 – Examples of How to Test Different Phone Number Formats	25
Figure 22 – Results for Analyzing Phone Numbers	26
Figure 23 – Example Synonym File Provides Alternate Names	27
Figure 24 – Mapping Settings that Applies the Name_Synonyms Analyzer	28
Figure 25 – Testing the Original Analyzer without Synonyms	28
Figure 26 – Results of the Original Analyzer	29
Figure 27 – Testing the Name_Synonyms Analyzer	29
Figure 28 – Tokenization Results of the Name_Synonyms Analyzer.....	29
Figure 29 – Searching Names without Synonyms	30
Figure 30 – Both Benjamin Smith and Josh Smith have the Same _Score	30
Figure 31 – Searching Names with Synonyms	31
Figure 32 – Josh Smith Now has a Higher Score with the Synonyms File.....	31
Figure 33 – Device Data Wireframe.....	32

1. Introduction

Through industry partnerships and customer connects, we have pursued insights on how the pandemic has fundamentally shifted consumer and worker behavior, exploring the increased importance of engagement in a world of remote work, and touchless experiences. We discovered a significant shift in customer expectations and behaviors: customers have moved beyond speed and are looking for a unified and simple experience with channel of choice. In a post-pandemic world, preferences have shifted to more personalized, convenient, and connected experiences.

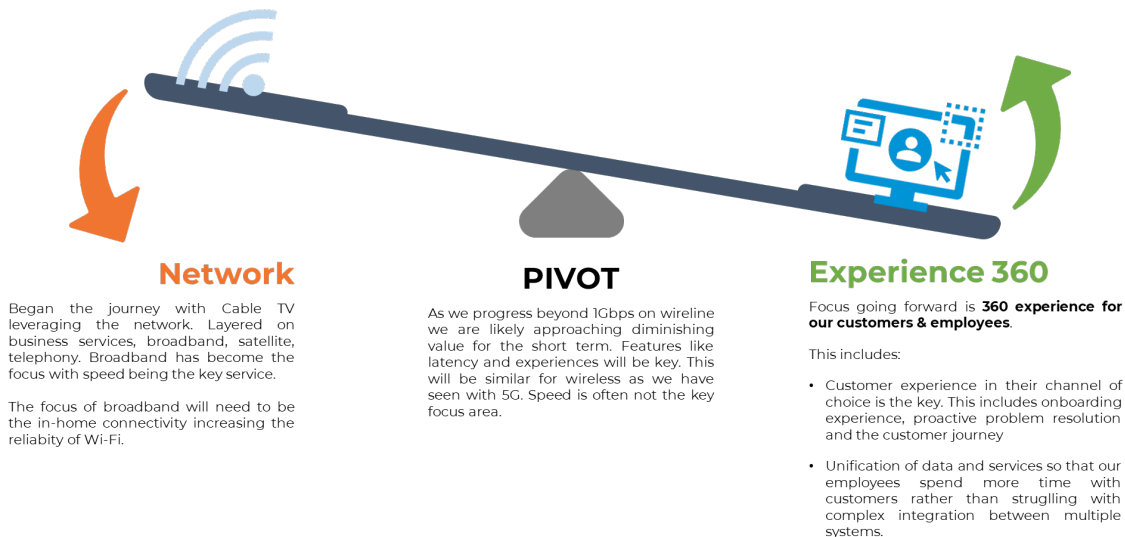


Figure 1 - Pivot in Customer Expectations

A unified channel touchpoint and agent experience is central to providing a connected customer experience. Because agents are behind the various customer interaction channels, their experience is the most crucial factor in meeting these quickly evolving customer expectations. To obtain a 360-degree view of the customer, agents have traditionally needed to access information from multiple systems and applications. However, this “swivel chair” experience was inefficient and impacted call times, wait times and general customer experience. Something had to change.

A Unified Customer Platform (UCP) has since been designed to empower agents through seamless visibility to any Customer in one place – no more swivels. UCP provides an accurate and centralized look at all the services a customer subscribes to, across all lines of business. Available 24/7, UCP visualizes customer and service information (using simple but intuitive UIs and modern search capabilities), allowing sales and support teams to quickly understand the customer is subscribed services. It also serves as the foundation to create customer communications and digital journeys, automated marketing and sales campaigns, enriched customer segmentation and enhanced reporting.

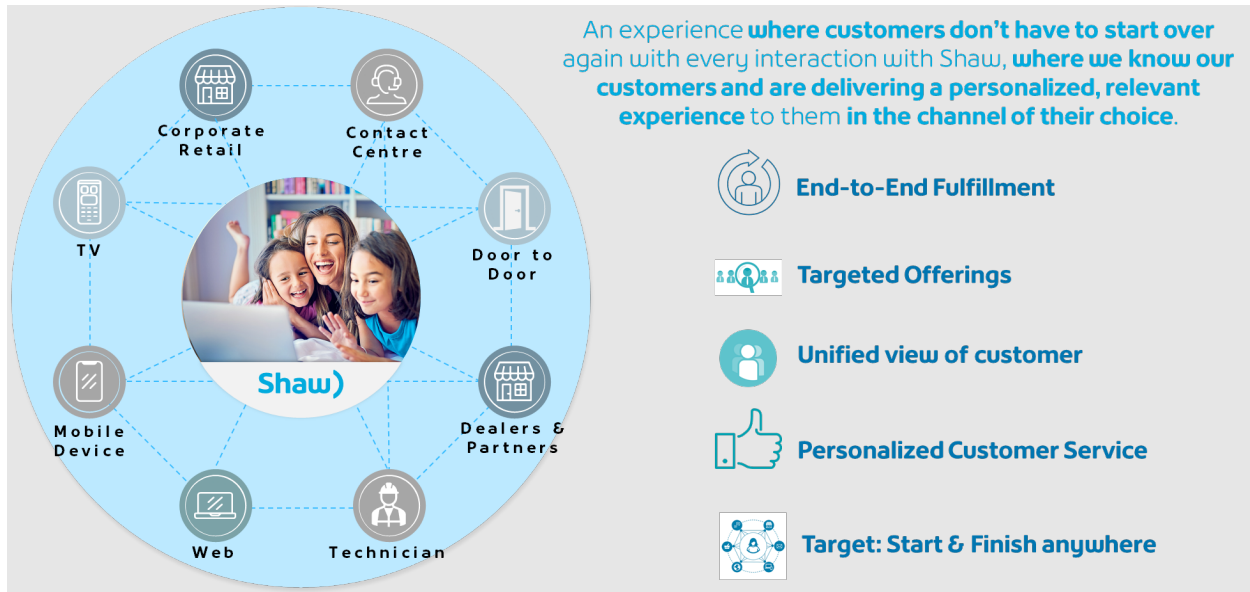


Figure 2 – Shaw's Customer Experience Strategy

In this paper we introduce the Unified Customer Platform (UCP), a modern platform that helps Shaw identify customers who meet eligibility requirements for new product offerings.



Figure 3 - Shaw's Agent Experience Strategy

Eligibility requirements often require taking a 360-degree view of the customer across multiple lines of business to highlight the products and services that are a good fit. This analysis leads to cross sell/upsell opportunities and would typically take place outside of a source system with the data analysis occurring in siloed data sets. Centralizing customer information (account information, subscribed services, history) in

UCP creates a single-entry point for reviewing cross-sell/upsell opportunities. Accessing customer information in UCP via an application programming interfaces (API) or a Google-like search interface was an important design decision. This will be the first time Shaw has had access to a single platform capable of retrieving customer information for all our customers across all lines of business.

Overall, the care teams using UCP are incredibly happy with the platform. The consolidated customer information helps reduce the “swivel-chair” previously required to jump between different billing systems. It provides them with the consolidated account details to best help support the customer base. The key findings that emerged from building and deploying UCP are:

- 1) Customer care teams benefit from platforms that aggregate account data from multiple source billing systems. We saw good adoption from Wireless Care teams who needed to find authorized users listed on a Wireline Internet services account.
- 2) Searching for customer accounts using the Elasticsearch search engine is a comfortable and time saving feature. Expanding customer search to include multiple service addresses and other business-class? options has been collected.
- 3) Modern technology stacks inside public clouds like AWS help reduce the overhead of building a new, common data platform and can deliver extreme performance at reasonable costs.
- 4) The lack of real-time data feeds into UCP was a constraint placed on us because of how data gets batched and scheduled upstream. For key data attributes like upgrades to new services, care teams would like to see an indicator that something is pending rather than no data.

Our intent is to continue to develop the UCP platform as an aggregation platform for a variety of new use cases that reach beyond care teams. Also adding additional account details like pending upgrades and customer interactions would help paint a better picture of our customers. Migrating some of the data loads into event-based streams could also help reduce the data load lag times caused by batched daily loads of data.

This paper will be organized as follows: we begin by introducing some of the drivers that helped ground us on why Shaw pursued building UCP in the first place. We will then present an overview of UCP and how it is being used to improve customer experiences. This is followed by the platform architecture and design challenges. We conclude the paper with a look at UCP’s journey so far, and likely future states.

2. Key Drivers

The key drivers to provide a personalized, convenient, and connected experiences to our customers started with the rebranding of the Shaw wireless product suite. If the Unified Customer Platform helps with selling Shaw Mobile, it could be used for other customer offers in the future. The relatively short project timelines had us looking at ways to modernize the technology stack so that we can focus on building business value without the incumbrance of managing infrastructure at scale.

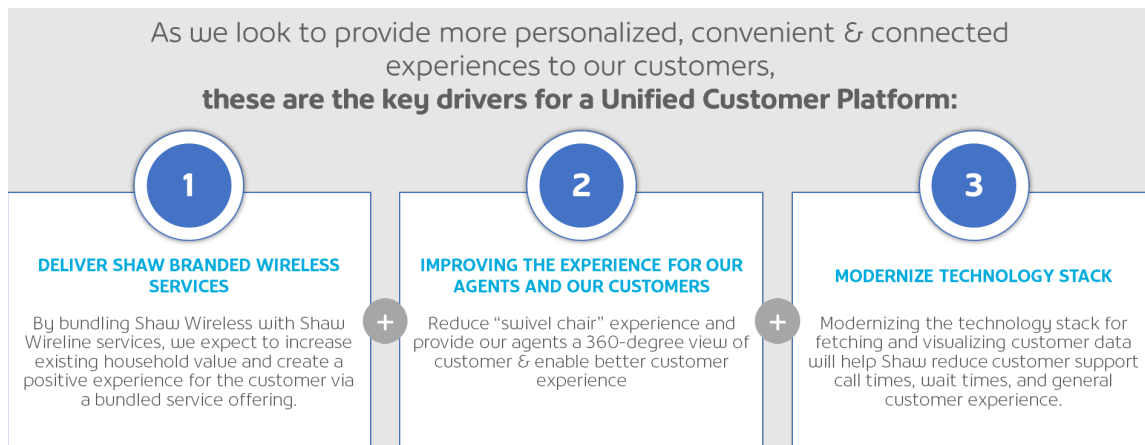


Figure 4 – Key Drivers for UCP

- Shaw wants to deliver a new Shaw-branded wireless product in Western Canada. By bundling Shaw Wireless with Shaw Wireline services, Shaw expects to increase existing household value and create a positive experience for the customer via a bundled service offering.
- Having a complete picture of the customer will provide a better agent experience and improve Shaw’s ability to support the current customer base, improve upsell recommendations, and maximize the effectiveness of marketing campaigns.
- Modernizing the technology stack for fetching and visualizing customer data will help Shaw reduce customer support call times, wait times, and general customer experience. Specifically, we were interested in how a query language for your API (GraphQL, graphql.org) can potentially create a more consumable API when compared to traditional RESTful (REpresentational State Transfer) API’s.

3. Overview of Unified Customer Platform (UCP)

The Unified Customer Platform enables searching for both residential and business customers across all lines of business using a dedicated search engine. Work was done to build data pipelines that aggregated this fragmented data into a single search index. Searching against this index matches your search terms against customer names, account numbers, phone numbers, and email addresses. The search results contain a list of customers, ordered by relevance, that match or partially match the search criteria along with all the associated child accounts for each customer. Why does this matter? UCP provides a single view of the customer instead of an account-centric view that was previously spread across five different source systems. Knowing the customer from an aggregated data perspective provides an enriched customer interaction, creates cross/up sell opportunities, exposes data insights to better serve the customer and allows the right offers to be recommended at the right time.

We felt it necessary to build this complete picture of the customer across three distinct layers. 1) a relational data set that can support reporting and marketing campaigns. 2) An API layer that functions as a data fetching API and gives clients the flexibility to select what data they need with minimal transformations. 3) And because agents need to interact with this data while assisting customers, it was necessary for us to build a new front-end to search and see all the customers data on one screen.

The Unified Customer Platform can be divided up into the following components:

- A front-end user interface that supports customer search and viewing customer details.
- An application programming interface (API) layer that uses the GraphQL to control exactly what data you get back from the API.
- Local data stores that optimize data fetching and retrieval.
- A data transfer layer that handles loading of the local data stores.
- A cloud data warehouse that aggregates customer data.
- A master data management (MDM) platform that uses a set of business rules for linking customer accounts across 5 incongruent billing systems.

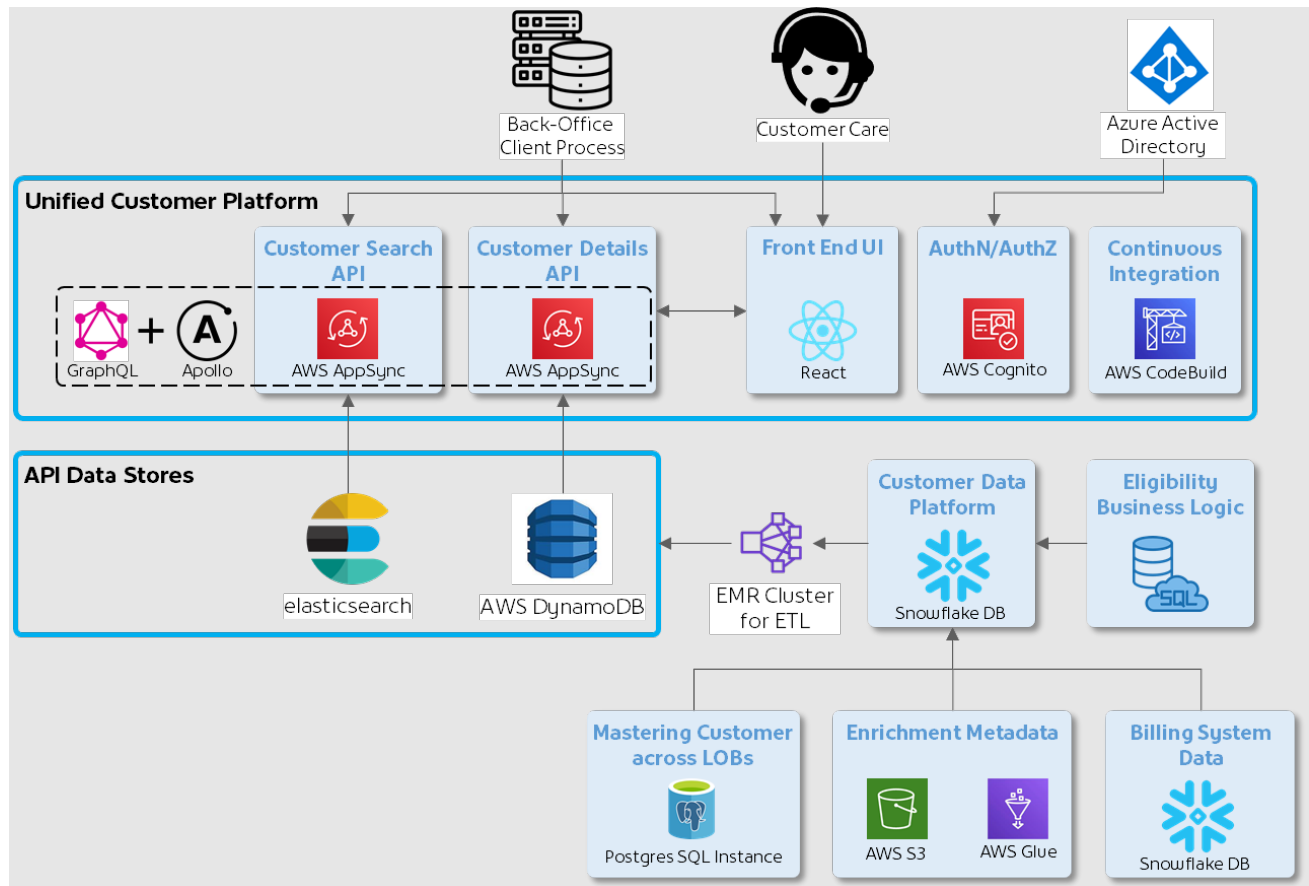


Figure 5 – High Level System Overview

4. Guiding Design Principles

When we sat down to develop this new platform, we were getting several different requirements from different teams. The scope of potential end-users was broad and included Back Office teams, Care teams, Retail support, Marketing, Finance, Revenue Assurance, Technical Operations, and corporate stores across multiple lines of business.

Some teams don't typically interface directly with the customer and have more time to perform troubleshooting steps. Teams like the Back Office, Marketing, and Finance are good examples of teams disconnected from the customer experience. Other teams like the technical operations and retail support provide the necessary support but operate under different expectations for a timelier resolution of issues. And then there are the customer facing teams like the care teams that operate the call centers and the retail store reps. The customer facing teams often need to navigate systems and applications while talking to a customer.

To avoid feature bloat and build a usable product, the design and development teams grounded ourselves on a few key principles. These include:

Basecamp Like Delivery

- Structure work and teams into cycles that last six weeks. We experimented with three, two-week sprints and two, three-week sprints. The number of actual sprints and length of sprints could vary depending on the type of work. But roughly every six weeks we wanted to finish a batch of product work and start preparing for a new batch of work. We were largely influenced by how Basecamp (Cisco, 2018)[1] structures their work.

Right Sizing Teams

- Team sizes are kept small. A team is two or three people that is dedicated to a portion of the product. We had a 2-person team for the UX design, a 2-person team for the AWS infrastructure, a 2-person team for the front-end, and a 2-person team for data engineering work. In a 6-week cycle, team sizes could flex up or down depending on the scope of work that needed to be done.

Prioritize Based on Value

- We could not do everything that we wanted to do and do it well. We did not have the time, resources, people, etc., so we prioritized the features to make the minimal viable product that executes on a few things and does them extremely well. The top requested feature was the ability to view account details for wireless and wireline to see if the account owner or authorized users on the account are eligible for promotions or offers.

API-First Approach

- Employ an API-first approach to building products. An API-first approach means that for any given development project, the API's are treated as "first-class" citizens. The API's can be consumed by both the client applications, other system platforms, or other development teams. Therefore, the API's need to be designed in an intuitive and reusable manner.

5. Minimum Viable Product Scope

The initial set of users for UCP was reduced to the care teams, marketing teams and back-office teams. Having a targeted user base helped us get both timely feedback and a reduced set of requirements plus allowed the development teams to course correct between sprint cycles.

The search feature was limited to customer names, account numbers, phone numbers, and email addresses. Searching by customer name needed to allow for partial matches on either the first name or last name.

Account details needed to include total monthly revenue totals by line of business, identification of the different product subscriptions, how long has this person been a customer, account status, payment status, and any authorized users on the account. Address information for billing and service addresses was also needed.

The unified customer channel needed to include three distinct access patterns for fetching customer data:

- A relational data warehouse that functions as the foundation for data aggregation and enrichment of customer data that can be accessed using Structured Query Language (SQL)
- A GraphQL API that provides some flexibility for the client to choose what data gets returned
- A React Web Application that uses the GraphQL API to efficiently fetch customer data

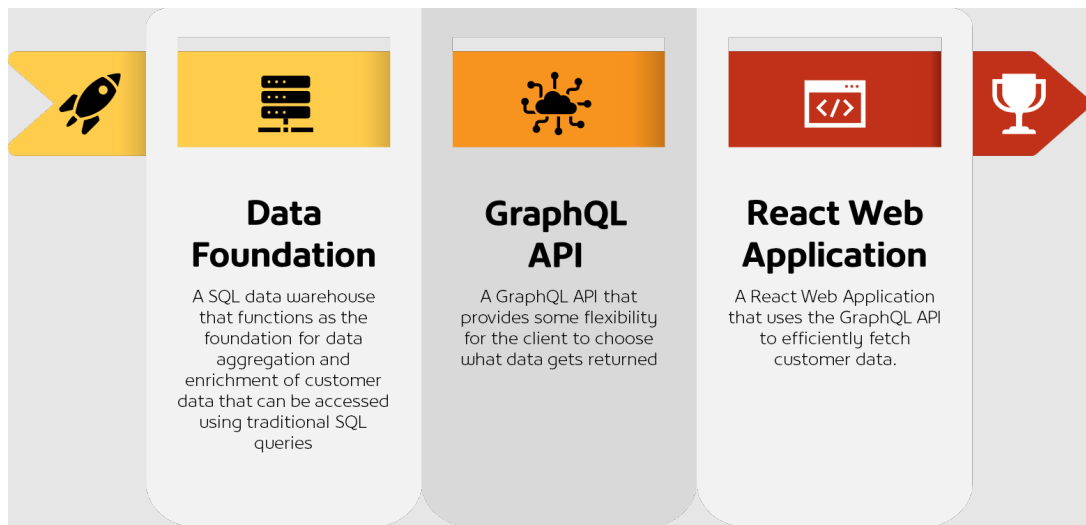


Figure 6 – UCP - Minimum Viable Product Scope

6. Components

6.1. Customer Search

The goal was to develop a single customer search endpoint that would search across the entire customer base in a performant manner. To optimize search performance, we are using Elasticsearch, a distributed search and analytics engine.

Initially, we considered using traditional relational database engines for search, but the search performance that we got from Elasticsearch was far superior to that of a Postgres database. In addition, we also had plans to take advantage of Elasticsearch features for including alternate spellings and nicknames for a customer's first name and addresses. Overall Elasticsearch felt like the best tool for the job.

We wanted searching to be simple, intuitive and with minimal navigation. We evaluated different design options and preferred a single search bar component with some place holder text that tells the user what to enter for search terms.




Figure 7 – Top Search Bar

Searching across residential and business accounts is the default without a need to provide any extra context. We also wanted the search engine to support partial matching of customer names and email addresses.

6.2. Search Results

Search results needs to show the customer, and all their associated linked accounts. (The linking of accounts will be covered later under the Customer Account Linking section.) When accounts are linked, there is one account that gets designated as the primary account. The rest of the linked accounts are considered child accounts. Child accounts can and will span multiple lines of business and be sourced from different billing systems.

4 RESULTS FOR

Customer Name: (John Smith)

Expand all | Collapse all

ACCOUNT NAME	ACCOUNT #	AUTHORIZED USERS	SERVICE ADDRESS	PHONE #	EMAIL	STATUS	
John Smith John Smith	1234567890 Wireline	Ann Smith Dan Smith	81 Chinook Drive Calgary, AB, T2V 2P8	403 555 6789	jsmith@gmail.com	Active ✓	>
John Smith Child 1 John Smith	0000000000 Wireline	Ann Smith Stan Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Active ✓	>
John Smith Child 2 John Smith	0000000000 Shaw Mobile	Tom Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Active ✓	>
John Smith Child 3 John Smith	0000000000 Shaw Mobile	Stan Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Active ✓	>
John Smith Jr John Smith	0000000000 Freedom Mobile	Jack Smith Mary Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Suspended !	>
Ann Smith Ann Smith	0000000000 Wireline	John Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Deactivated ✗	>
Jerry Smith Jerry Smith	0000000000 Wireline	John Smith Sally Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Active ✓	>
Jerry Smith Child 1 Jerry Smith	0000000000 Wireline	Ross Smith Sally Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Active ✓	>
Jerry Smith Child 2 Account owner	0000000000 Shaw Mobile	Lisa Smith Sally Smith	Civic address Municipality, province, postal code	000 000 0000	email@email.com	Active ✓	>

Figure 8 – Search Results page variation where “Service Type” is shown with text in the Account # column

By clicking on one of the arrows on the right side of the page, the user will navigate to a customer details view.

6.3. Customer Details View

The customer details page provides a summary of accounts by line of business. Each line of business is reflected in a separate tab and has a summary of services, date connected, total number of accounts, the total revenue for all services, and whether the customer is in the mobile network and is eligible for a wireless bundle.

Below the tabs, there is an account summary section consisting of the account number, authorized users, revenue per account, source billing system, and account status/payment status. Below the summary are the account details that includes billing and service addresses and the details for the distinct types of services.

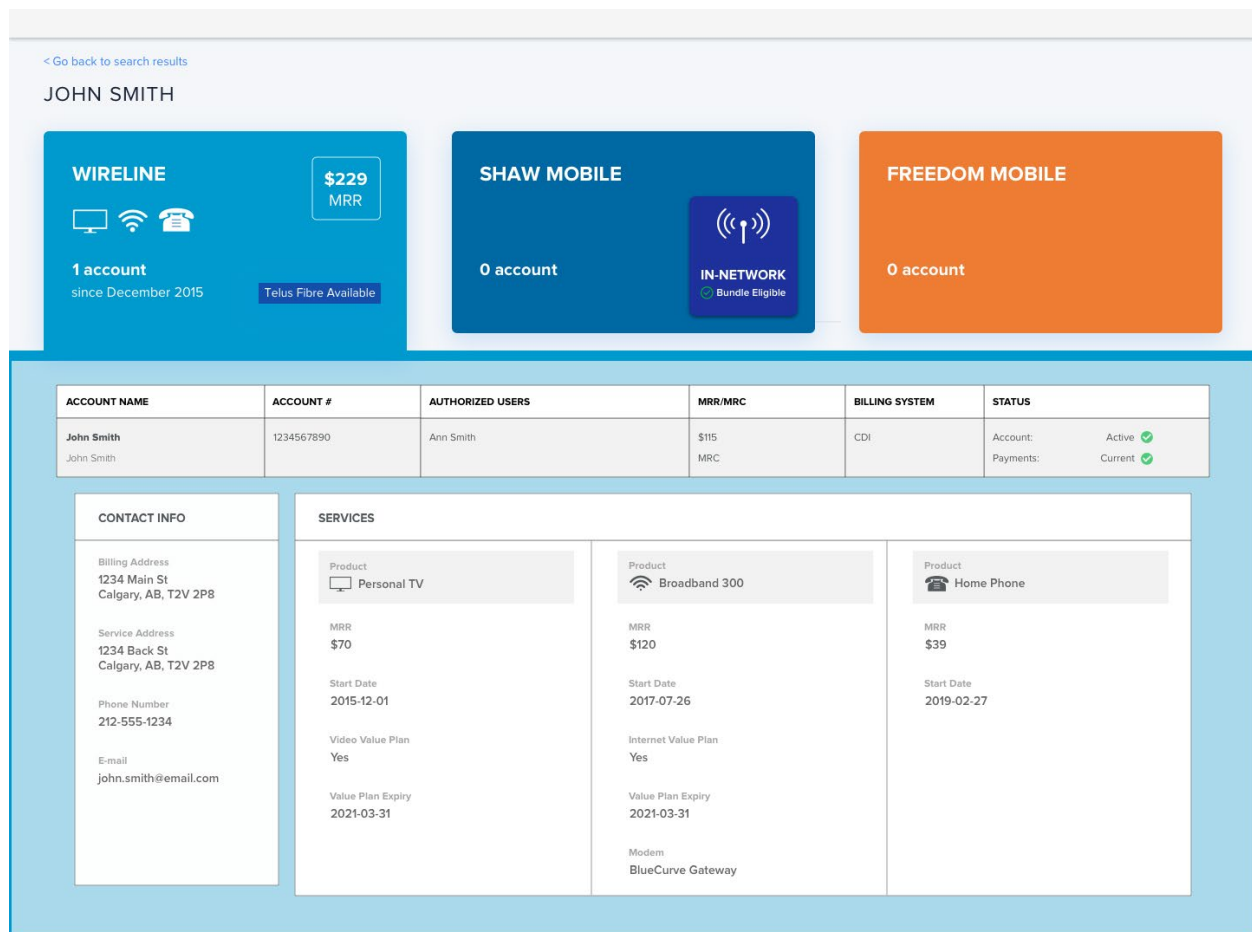


Figure 9 – Wireline, no mobile service (in-network wireless quality & eligible for bundle)

6.4. Data API Considerations

A customer data API needs to be designed in a reusable way to support multiple use cases. The primary use case is to serve as the back-end data fetching API for the new UCP front-end. In addition, we received requests from other platform owners to use this API to enrich customer data by making a simple HTTP request to our API. This type of integration is attractive because we eliminate the need for replicating customer data sets between various systems. Replicating large data sets across disparate systems is expensive, slow, and difficult to maintain.

The decision to use GraphQL over REST was an easy decision from the perspective that by 2020 GraphQL was a proven technology offering many advantages over some of the challenges that faced traditional REST APIs.

For example,, retrieving an aggregated view of accounts using a RESTful API might require making multiple calls to the API to get all the required data for each account. Or, if the data has been aggregated to a centralized landing area, the REST API would most likely be broken up into individual resources that return a set of data specific to a use case and still require a client to make multiple API calls to retrieve all

the required data. Also, the data returned from a REST API call has the potential to return a lot of unnecessary data to the client.

GraphQL is a more modern approach to building API's that gives clients more control over the data that they want to get back from an API request. This quote is from the graphql.org web site [2] which does an excellent job summarizing GraphQL:

“GraphQL is a query language for APIs and a runtime for fulfilling those queries with your existing data. GraphQL provides a complete and understandable description of the data in your API, gives clients the power to ask for exactly what they need and nothing more, makes it easier to evolve APIs over time, and enables powerful developer tools.”

A query language for APIs means that clients can change the shape of data returned by a GraphQL server. Clients will have control to specify different return fields, run multiple queries in a single request, and add aliases to fields to accommodate naming differences across front-end and back-end code. This all helps reduce the amount of data transformations needed on the client to process the data.

In summary, GraphQL is typically served over HTTP/s via a single endpoint with a data schema that informs clients about the shape and types of data that can be returned from the API. This contrasts with REST APIs over HTTP which typically expose a suite of URLs (multiple endpoints), with each URL endpoint exposing a single resource that defines a data format for the return data.

6.5. Snowflake Cloud Data Warehouse

One of the core requirements was to provide a unified customer channel that can be accessed via SQL as well as an API. It seemed logical that the cloud data warehouse would be a landing zone for customer data, so we built a customer data mart for this purpose. This customer data mart was designed as the primary source of data for UCP.

The primary consideration was that we wanted to use a modern data platform that enabled separation of compute and storage so that we can scale compute and storage separately. This requirement had us considering a cloud data warehouse (e.g., Snowflake) or a modern cloud data lake in AWS. We ultimately chose Snowflake because it was an SQL-based cloud data warehouse that allowed us to leverage existing skill sets with the minimum amount of training required to adopt the new platform. Snowflake's user administration was also simpler and more integrated with the product compared to the AWS data lake.

6.6. API Data Stores

The API data stores were needed to meet our performance needs. For the customer search API, we used Elasticsearch in AWS. There is nothing specific about the Elasticsearch configuration that required it to be hosted in AWS, so this decision was made mostly out of convenience.

For the customer details lookup API, we used DynamoDB, a fast, NoSQL Key-Value database. The main consideration was that we needed a database that could manage semi-structured customer data in JSON format. Other NoSQL datastores like MongoDB could be drop-in replacements for DynamoDB.

6.7. Customer Account Linking

Customer account linking is the process of matching accounts by different attributes like name, service addresses, phone numbers, etc. Once a grouping of accounts is known for a customer, one account gets assigned as the primary account and other accounts become child accounts and are linked via a data relationship.

We took a conservative approach to linking customers by building a recommendation engine that only recommends accounts to be linked. Administrators need to accept the recommendation before one account gets assigned the primary account and all the other accounts get linked as child accounts. The recommendation engine factors in a set of rules that helps prioritize which accounts are considered the primary.

This process runs on a separate platform and the results for account linking get replicated into Snowflake for consumption by UCP.

7. Challenges

Building a GraphQL API that returns data for all customer types across all lines of business required defining a customer data type in a GraphQL schema document. Bringing residential and business customers together in a single data type requires some flexibility in the design so that we do not create dependencies for things that are independent.

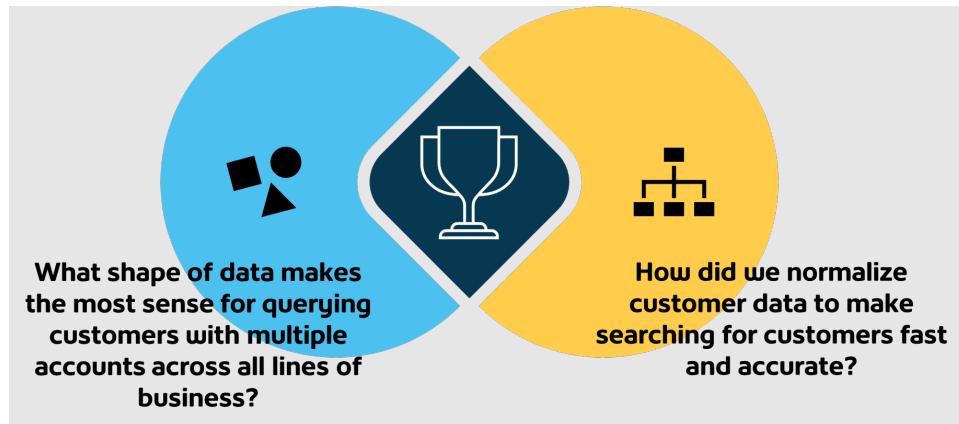


Figure 10 – UCP – Challenges

What shape of data makes the most sense for querying customers with multiple accounts across all lines of business?

The customer API must be flexible and return a range of possible data types depending on the subscribed products and services for each line of business. A customer could be a combination of wireline and wireless services with data coming from different source billing systems. We wanted to future proof the API by not adding any dependencies across billing systems. Billing systems can make changes that are independent of other billing systems and the customer API needs to be able to handle any type of data.

The types of data for a residential customer could differ dramatically from a business customer and the API needs to dynamically handle this. We wanted to remain consistent with GraphQL best practices by having a single API endpoint for all the distinct types of customers and to be able to return all the customer data in one HTTP request. If you only are interested in detail for one type of customer, for example, wireline customer details, the request payload can change to reflect the exact details for a specific customer or even multiple customers.

For the API to handle a range of possible field types, there is a union type that's part of the GraphQL specification.

```
union ServiceProvider = ShawWirelineServiceProvider | ShawMobileServiceProvider | FreedomMobileServiceProvider | AccountNotFound
```

Figure 11 – Union of Customer Types

A query that returns the ServiceProvider type is designed to receive more than one billing account number that could be any one of the ServiceProvider types. If the account lookup fails, the AccountNotFound type gets populated.

In cases where a client is passing in a single account number and already knows the type of service provider, they can build an extremely specific query that reduces the API response to the pieces of data that they care about.

```
query GetInNetworkStatus {
  getAccountsByServiceProvider(getDashboardInput: [{billingAccountId: "TESTESI-11111111"}]) {
    serviceProvider {
      __typename
      ... on ShawWirelineServiceProvider {
        accounts {
          accountName
          billingAccountId
          detailType
        }
        inNetwork
        wirelessCoverageQuality
      }
    }
  }
}
```

Figure 12 – Example of Getting In-Network Status for a Single Account Number

In this example, the client is only interested in finding out whether the customer is in network and has good wireless coverage quality to be able to offer them a Shaw Mobile bundle.

In other cases, we could enter multiple billing accounts for a single customer. In this example, the \$getDashboardInput is an array that accepts multiple billing accounts of type getAccountInput. The account numbers getting passed in could be one of the three different service providers or a failure condition when account number is not found.

```
query getAccountsByServiceProvider($getDashboardInput: [getAccountInput!]!) {
  __typename
  getAccountsByServiceProvider(getDashboardInput: $getDashboardInput) {
    serviceProvider {
      __typename
      ... on ShawWirelineServiceProvider {
        accounts {↔}
        telusFibreAvailable
        inNetwork
        bundleEligible
        minStartDate
      }
      ... on ShawMobileServiceProvider {
        minStartDate
        accounts {↔}
      }
      ... on AccountNotFound {
        __typename
        missingaccounts
      }
      ... on FreedomMobileServiceProvider {
        minStartDate
        accounts {↔}
      }
    }
  }
}
```

Figure 13 – Multiple Account Numbers as Input; Aggregate by Service Provider

How did we normalize customer data to make searching for customers fast and accurate?

As mentioned above, one of the design considerations for search was to use Elasticsearch as the search engine for finding customers by customer name, account number, email address or phone number. Elasticsearch, or one of its derivatives like AWS OpenSearch, typically ranks extremely high in popularity for enterprise search capabilities. DB-Engines ranks Elasticsearch #1 in popularity as of June 2022. [3]

Today, a lot of the undifferentiated heavy lifting for provisioning and scaling an Elasticsearch cluster is handled by third-party cloud providers. We will not go into tremendous detail on the physical infrastructure that is needed to search across millions of customer records. Instead, we will spend more time discussing how the customer data is indexed to achieve fast and efficient searches.

The physical infrastructure for an Elasticsearch cluster in AWS that will support searching across millions of customers could look like the following.

- Use Amazon OpenSearch Service for managing a six-node search cluster running either an Elasticsearch 6.8 cluster or an OpenSearch 1.2 cluster.
- Typically, there would be three master nodes and three data nodes all running on instance types of either c5.xlarge.search (Elasticsearch) or m6g.xlarge.search (OpenSearch).

Indexing data in Elasticsearch is done using analyzers that translate data into tokens that are more suited for search. A token in Elasticsearch contains the string, type, and some positional offsets. In large blocks of unstructured text, the same token could appear in multiple positions. For the examples here, we are dealing with structured data as text or keywords and therefore the positional offsets are not as important.

In Elasticsearch there is a difference between a text field and a keyword field. The text fields get analyzed and broken down into different tokens, whereas a keyword field is typically represented as a single token.

The tokenization of a customer's full name as a keyword field could look like this:

```
1 {
2   "tokens" : [
3     {
4       "token" : "josh smith",
5       "start_offset" : 0,
6       "end_offset" : 10,
7       "type" : "word",
8       "position" : 0
9     }
10  ]
11 }
12
```

Figure 14 – Keyword Field Tokenization Keeps Josh Smith as a Single Token

Whereas the tokenization of a text field for a customer's full name as a text field would look more like this:

```
1 {
2   "tokens" : [
3     {
4       "token" : "josh",
5       "start_offset" : 0,
6       "end_offset" : 4,
7       "type" : "<ALPHANUM>",
8       "position" : 0
9     },
10    {
11      "token" : "smith",
12      "start_offset" : 5,
13      "end_offset" : 10,
14      "type" : "<ALPHANUM>",
15      "position" : 1
16    }
17  ]
18 }
19
```

Figure 15 – Text Field Tokenization Separates Name into Two Tokens

Only text fields can be sent through an analysis process that structures data into tokenized formats that are optimized for search. Elasticsearch ships with built-in [4] and custom analyzers. Custom analyzers will be needed to improve search hits for account numbers and phone numbers.

Normalizers are like analyzers but are performed on keyword fields and only produce a single token. Keyword fields are usually structured data types that are recognizable fields like email addresses, phone numbers, and account numbers that are typically used for filtering and sorting.

Now that there is a better understanding of the difference between keywords and text fields, we can jump into some of the things that we had to do to improve searching by account numbers and phone numbers.

Account numbers defined as keywords will get normalized to make searching against them simpler. Because we are dealing with a handful of different billing systems, we could see variations with how account numbers are formatted. Some accounts are all digits; sometimes accounts include leading zeros, other times they do not; and some accounts have characters that prefix a set of digits that could be capitalized or lower case.

The first step is to deal with the leading zeros, so we do not require this as part of search. The `char_filter` will replace any of the leading zeros with an empty string. For this example, we will just focus on the `customer_account_number.keyword` field. This field uses the keyword normalizer with a custom filter (`char_filter`) that removes the leading zeros and forces the final token to be lowercase.

```
PUT scte_index
{
  "settings": {
    "analysis": {
      "char_filter": {
        "no_leading_zeros": {
          "type": "pattern_replace",
          "pattern": "^(0*)",
          "replacement": ""
        }
      },
      "filter": {},
      "normalizer": {
        "keyword_normalizer": {
          "type": "custom",
          "char_filter": ["no_leading_zeros"],
          "filter": [
            "lowercase"
          ]
        }
      },
      "analyzer": {}
    },
    "mappings": {
      "properties": {
        "customer_account_number": {
          "type": "text",
          "analyzer": "account_number",
          "fields": {
            "keyword": {"type": "keyword", "normalizer": "keyword_normalizer"}
          }
        },
        "customer_name": {}
      }
    }
  }
}
```

Figure 16 – Removing Leading Zeros from Account Number

Using the above settings, searching for a customer using an account number no longer requires specifying leading zeros. The query below uses account number “123456789” as the search term.

```
GET scte_index/_search
{
  "query": {
    "match": {
      "customer_account_number.keyword": {
        "query": "123456789"
      }
    }
  }
}
```

Figure 17 – Account Number Search Query (Leading Zero not Specified)

The query string of “123456789” normally will not match an account number that is indexed with “0123456780” because of the zero prefix. However, if you look at the results below you will notice that the document returned from the query “123456789” matches a document that contains a zero prefix.

```
{
  "took" : 3,
  "timed_out" : false,
  "_shards" : {
    "total" : 5,
    "successful" : 5,
    "skipped" : 0,
    "failed" : 0
  },
  "hits" : {
    "total" : {
      "value" : 1,
      "relation" : "eq"
    },
    "max_score" : 0.2876821,
    "hits" : [
      {
        "_index" : "scte_index",
        "_type" : "_doc",
        "_id" : "-EXVc4EB6ViCQ5dG6p-7",
        "_score" : 0.2876821,
        "_source" : {
          "customer_account_number" : "0123456789",
          "customer_name" : "john smith"
        }
      }
    ]
  }
}
```

Figure 18 – Search Results for Query “123456789”

Normalizing phone numbers required a little more work to get satisfactory results. Phone numbers needed to be only numerical digits, 10 digits minimum, not empty, required removal of all zero prefixes, so that different formats for phone numbers could be handled. In the example below, we use a custom phone number analyzer and search analyzer.

```
{
  "scte_index_03" : {
    "mappings" : {
      "properties" : {
        "customer_phone_numbers" : {
          "type" : "text",
          "analyzer" : "phone_number",
          "search_analyzer" : "phone_number_search"
        }
      }
    }
  }
}
```

Figure 19 – Phone Number Mapping Example

And below you can see how the phone number analyzer is configured. There is a `char_filter` for removal of all non-digit characters from the phone number. The `us_phone_number` is a custom tokenizer that removes any leading 1's from the phone number and preserves the original number.

```
{
  "analyzer": {
    "phone_number": {
      "char_filter": "digits_only",
      "tokenizer": "keyword",
      "filter": [
        "us_phone_number",
        "ten_digits_min"
      ]
    },
    "phone_number_search": {
      "char_filter": "digits_only",
      "tokenizer": "keyword",
      "filter": [
        "not_empty"
      ]
    }
  }
}
```

Figure 20 – Phone Number Analyzer

We can test how the `phone_number` analyzer works by using the `_analyze` feature of Elasticsearch. The `_analyze` feature runs an analyzer with a text string and outputs the generated tokens.

```
GET scte_index_03/_analyze
{
  "analyzer": "phone_number",
  "text": "18005551111"
}

GET scte_index_03/_analyze
{
  "analyzer": "phone_number",
  "text": "1 (800) 555-1111"
}

GET scte_index_03/_analyze
{
  "analyzer": "phone_number",
  "text": "800.555.1111"
}
```

Figure 21 – Examples of How to Test Different Phone Number Formats

The first two _analyze examples resolve to these two tokens [18005551111, 8005551111] while the last example resolves to [8005551111]. The important thing to notice that each example has the 8005551111 as one of its tokens. Actual output from the _analyze command is shown below.

```
{  
  "tokens" : [  
    {  
      "token" : "18005551111",  
      "start_offset" : 0,  
      "end_offset" : 11,  
      "type" : "word",  
      "position" : 0  
    },  
    {  
      "token" : "8005551111",  
      "start_offset" : 0,  
      "end_offset" : 11,  
      "type" : "word",  
      "position" : 0  
    }  
  ]  
}
```

Figure 22 – Results for Analyzing Phone Numbers

Because account numbers, email addresses, and phone numbers are stored in different fields we needed to use what Elasticsearch calls a multi-match query. In a multi-match query, we are running multiple queries with the same search term. Because the search term gets applied to different search fields, we rely on the relevance score to build the search results page.

Search algorithms apply rules and mathematical calculations to derive a relevance score so we then can do a little manipulation of the scores. For example, we boost complete matches against full name and last name higher than a partial match that uses a wild card search. Elasticsearch uses Lucene under the hood so by default it uses the Practical Scoring Function [5]. The details of the scoring function go beyond this paper, but it is worth noting that Elasticsearch does give you the ability to tune the scoring algorithm and change relevance scores.

8. What's next for Unified Customer Portal

There are a lot of opportunities for us to continue development of UCP so that Care agents continue to have better interactions with the customers. Below are three features that we are currently pitching to the UCP product owners.

8.1. Search Improvements

For customer search we would like to add synonyms for both customer first names and alternate spellings for address searches. To accomplish this in Elasticsearch you need to start with a synonyms file that correlates between a name and different nicknames.

```
aaron => erin, ron, Ronnie  
alan => al  
...  
benjamin => ben  
...  
josh => joshua  
...  
william => bela, bell, bill, billy, will, willie, willy
```

Figure 23 – Example Synonym File Provides Alternate Names

Maintaining a synonym list is difficult to do manually so we would like to explore some novel ways on how to automatically push updates to this file. The plan will be to start with an initial list of nicknames that we can eventually compare against actual search logs so that we can measure the synonym list for completeness.

Once a synonym list is added to an Elasticsearch index you can map the synonyms to an analyzer. The example below uses a synonyms field instead of a file, but it is sufficient to prove out how we can take advantage of nicknames. The filter object defined next to number 1 creates the synonyms. The number 2 section tells Elasticsearch to use the name_synonyms analyzer for the name field. And Section 3 configures the analyzer to tokenize text into lowercase and add the synonyms data.

```
PUT scte_index_02
{
  "settings": {
    "analysis": {
      "filter": {
        "name_synonym_filter": {
          "type": "synonym",
          "synonyms": [
            "josh,joshua",
            "ben,benjamin"
          ]
        }
      },
      "analyzer": {
        "original": {
          "type": "custom",
          "tokenizer": "keyword",
          "filter": [ "lowercase" ]
        },
        "partial": {
          "type": "custom",
          "tokenizer": "standard",
          "filter": [ "lowercase" ]
        },
        "name_synonyms": {
          "type": "custom",
          "tokenizer": "standard",
          "filter": [ "lowercase", "name_synonym_filter" ]
        }
      },
      "mappings": {
        "properties": {
          "name": {
            "type": "text",
            "analyzer": "name_synonyms"
          }
        }
      }
    }
  }
}
```

Figure 24 – Mapping Settings that Applies the Name_Synonyms Analyzer

We created three analyzers in the example above and we can use the `_analyze` function to show how text gets tokenized and synonyms are incorporated.

```
GET scte_index_02/_analyze
{
  "analyzer": "original",
  "text": "Josh Smith"
}
```

Figure 25 – Testing the Original Analyzer without Synonyms

The tokenizer named original keeps the text “Josh Smith” as a single token, but it does make it lower case. Take note that these synonyms are not included in this example when using the original tokenizer.

```
1 {
2   "tokens" : [
3     {
4       "token" : "josh smith",
5       "start_offset" : 0,
6       "end_offset" : 10,
7       "type" : "word",
8       "position" : 0
9     }
10  ]
11 }
```

Figure 26 – Results of the Original Analyzer

By changing the analyzer from 'original' to 'name_synonyms' you will see the tokenization of “Josh Smith” look different from the above example.

```
GET scte_index_02/_analyze
{
  "analyzer" : "name_synonyms",
  "text" : "josh Smith"
}
```

Figure 27 – Testing the Name_Synonyms Analyzer

There are now three separate tokens for “Josh Smith” that can be used for searching [“josh”, “joshua”, “smith”]

```
1 {
2   "tokens" : [
3     {
4       "token" : "josh",
5       "start_offset" : 0,
6       "end_offset" : 4,
7       "type" : "<ALPHANUM>",
8       "position" : 0
9     },
10    {
11      "token" : "joshua",
12      "start_offset" : 0,
13      "end_offset" : 4,
14      "type" : "SYNONYM",
15      "position" : 0
16    },
17    {
18      "token" : "smith",
19      "start_offset" : 5,
20      "end_offset" : 10,
21      "type" : "<ALPHANUM>",
22      "position" : 1
23    }
24  ]
25 }
26 }
```

Figure 28 – Tokenization Results of the Name_Synonyms Analyzer

Searching for “Joshua Smith” without synonyms will get a match on “smith” but both “josh smith” and “ben smith” are scored the same.

```
GET scte_index_02/_search
{
  "query": {
    "match": {
      "name": {
        "query": "joshua smith"
      }
    }
  }
}
```

Figure 29 – Searching Names without Synonyms

```
1 {
2   "took" : 2,
3   "timed_out" : false,
4   "_shards" : {
5     "total" : 5,
6     "successful" : 5,
7     "skipped" : 0,
8     "failed" : 0
9   },
10  "hits" : {
11    "total" : {
12      "value" : 2,
13      "relation" : "eq"
14    },
15    "max_score" : 0.2876821,
16    "hits" : [
17      {
18        "_index" : "scte_index_02",
19        "_type" : "_doc",
20        "_id" : "90W8coEB6ViCQ5dGSZ9S",
21        "_score" : 0.2876821,
22        "_source" : {
23          "name" : "benjamin smith"
24        }
25      },
26      {
27        "_index" : "scte_index_02",
28        "_type" : "_doc",
29        "_id" : "Ys08coEBexGa4gSBP-bi",
30        "_score" : 0.2876821,
31        "_source" : {
32          "name" : "josh smith"
33        }
34      }
35    ]
36  }
37 }
38
```

Figure 30 – Both Benjamin Smith and Josh Smith have the Same _Score

After adding synonyms to the search query, you get a good match for a “Joshua Smith” query which will be ranked higher than “ben smith.”

```
GET scte_index_02/_search
{
  "query": {
    "match": {
      "name": {
        "query": "joshua smith",
        "analyzer": "name_synonyms"
      }
    }
  }
}
```

Figure 31 – Searching Names with Synonyms

```
1- {
2  "took" : 17,
3  "timed_out" : false,
4-  "_shards" : {
5    "total" : 5,
6    "successful" : 5,
7    "skipped" : 0,
8    "failed" : 0
9-  },
10- "hits" : {
11-   "total" : {
12     "value" : 2,
13     "relation" : "eq"
14-   },
15   "max_score" : 0.5753642,
16-   "hits" : [
17-    {
18     "_index" : "scte_index_02",
19     "_type" : "_doc",
20     "_id" : "Ys08coEBexGa4gSBP-bi",
21     "_score" : 0.5753642,
22-    "_source" : {
23     "name" : "josh smith"
24-    }
25-   },
26-   {
27     "_index" : "scte_index_02",
28     "_type" : "_doc",
29     "_id" : "90W8coEB6ViCQ5dGSZ9S",
30     "_score" : 0.2876821,
31-    "_source" : {
32     "name" : "benjamin smith"
33-    }
34-   }
35-  ]
36- }
37- }
38
```

Figure 32 – Josh Smith Now has a Higher Score with the Synonyms File

8.2. Customer Device Details

We would like to enhance the customer details by adding device metrics. Device data that is associated with the service line can provide a Care agent more context on how well their mobile plan is working and if changes are needed.

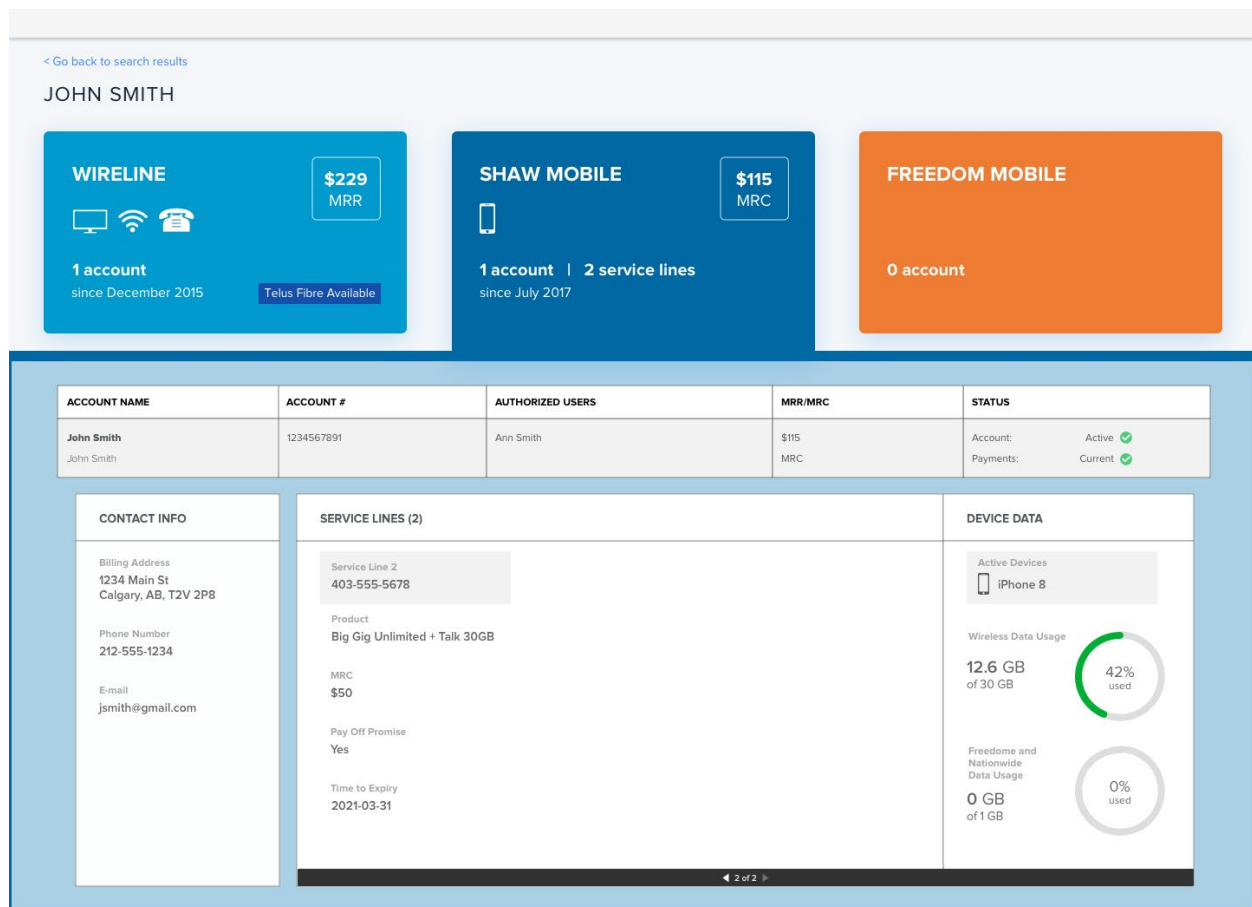


Figure 33 – Device Data Wireframe

8.3. Real-Time Data Loading Process

Our goal has always been to move to more of a real-time data load process for customer account linking and customer search. The challenge is that we are currently dealing with a lot of upstream legacy systems that only support batch exports of data.

There is still some room to optimize the schedules being used for batch loads. Transitioning to event-based schedules will help us load data as soon as it is ready. Snowflake recently announced its Unistore workload which expands the capabilities of Snowflake to support modern transactional data and analytical data in one platform. Keeping data in snowflake could eliminate the Snowflake-to-DynamoDB data load which moves the freshness of data up by 30 minutes.

Real-time data processing will require source systems to support event-based data streams. The UCP platform would subscribe to these event streams and use the data to build additional metrics about a customer.

9. Conclusion

Prior to the introduction of the Unified Customer Platform, our front-line agents would need to access multiple back-office systems across different lines of business. This “swivel chair” experience was time consuming and very inefficient while serving customers that lead to customer frustration and longer call times.

UCP simplifies this for our agents by allowing them to quickly search for a customer by name, account number or phone number. Once the customer is found, UCP displays the accounts and service attributes for all services that the customer subscribes to, eliminating the need to find these accounts across multiple interfaces. UCP provides all the information needed for an agent to understand the current customer engagement with Shaw across different LOBs. In addition, our agents can use the information in the tool to determine if a customer is eligible for any current promotions. UCP serves the intent that we were after: to utilize the data to best serve our customer base and allow our agents to talk intelligently with them.

Over time, we will continue to build on the foundation that we have built and extend the platform by including the ability to link customer accounts in a more formal and verified manner in real time. Our value-based iterative delivery and ability to unlock the true value of enterprise data will help us meet the goals that we have set for ourselves.

Abbreviations

5G	5th Generation
API	Application Programming Interface
AWS	Amazon Web Services
DB	Database
EMR	Elastic Map Reduce
ETL	Extract Transform Load
Gbps	Gigabit per second
GraphQL	Graph Query Language
HTTP	Hypertext Transfer Protocol
JSON	JavaScript Object Notation
LOB	Line of Business
MDM	Master Data Management
MRC	Monthly Recurring Charge
MRR	Monthly Recurring Revenue
MVP	Minimum Viable Product
NoSQL	Not Only SQL
REST	REpresentational State Transfer
S3	Simple Storage Service
SQL	Structured Query Language
UCP	Unified Customer Platform
URL	Uniform Resource Locator
UX	User Experience

Bibliography & References

- [1] <https://m.signalvnoise.com/how-we-structure-our-work-and-teams-at-basecamp/>
- [2] <https://graphql.org/>
- [3] <https://db-engines.com/en/ranking/search+engine>
- [4] <https://www.elastic.co/guide/en/elasticsearch/reference/current/analysis-analyzers.html>
- [5] <https://www.elastic.co/guide/en/elasticsearch/guide/current/practical-scoring-function.html>

Truth/Devil is in the Details: The Fusion of Design and Telemetry Information in Access Networks

A Technical Paper prepared for SCTE by

Matt Wichman
Director
Comcast
+1 (425) 754-8388
Matthew_Wichman@Comcast.com

Venk Mutalik, Fellow

Joann Shumard, Vice President

Kathy Fox, Vice President

Table of Contents

Title	Page Number
1. Abstract	3
2. Introduction.....	3
3. Digitization of Data	3
4. Operational Logic, XMF-R.....	4
5. Acting on Real-time Data	5
5.1. Process Changes	5
5.2. Tool Changes	6
6. The Access Network	6
6.1. Coherent Optics in the Access Domain	6
7. Putting it All Together: Role of Orchestration.....	7
7.1. Deployment and Implementation	7
7.2. Operational Integration.....	8
7.3 An Informed Fiber Design	8
7.4 Proper Change Implementation	8
8. Conclusion.....	8
Abbreviations	9
Bibliography & References.....	9

List of Figures

Title	Page Number
Figure 1 - Operational Logic.....	5
Figure 2 - Access Network Convergence	7

1. Abstract

‘The Truth/Devil is in the Details’ is a common saying referring to the mysterious or hidden elements knowable only with infinite parsing of all obtainable information. And truth be told there are many places where valuable information about a network as vast as Comcast’s resides. To obtain, visualize and correlate all these pieces of information requires innovation in hardware and software – in intersecting cycles that virtuously enhance each other.

Comcast serves tens of millions of residential customers and businesses. In our efforts to increase capacity and enhance reliability, Comcast has a host of hardware and software telemetry tools that monitor various factors that affect the performance of our network.

This paper will focus on ways we are taking advantage of our highly digitized network to automate wavelength management, speed up fiber event resolution, and enhance fiber construction. We will detail how real-time fiber management and auto-calibration of new fiber links are essential to maintaining complex networks and delivering exceptionally reliable services. We will also review how we are addressing fiber repairs. We all know that fiber outage events and wavelength management are significant industry challenges because we rely on multiple data sources and integration of the data for accurate records and expedient repair. We also rely on the identification of the location and distance of events to speed the arrival of repair agents. With our solution, learn how utilizing telemetry, multiple data sources, and innovative technology, a fiber event can be detected, assessed, and dispatched to an accurate location.

2. Introduction

Our networks are becoming more fiber-rich every day. We have smaller node sizes as we are driving fiber closer to the customer in both residential and commercial solutions and the amount of fiber connecting our facilities has only increased. In the past when there was a suspected fiber cut, our operations teams would notify the necessary employees or contractors to deploy to troubleshoot the fiber to determine if there was a fiber cut and if so, the associated footage so the fix agents could get to the repair location as quickly as possible. With XMF, (Xfinity Meter Fiber), a solution that automates traditionally manual fiber impairment analysis work, we have saved considerable time because there is no need to dispatch someone to perform the OTDR (Optical Time Domain Reflectometer) measurement of the fiber, assess the distance, then review the prints to identify the location. Instead, this solution provides the necessary details in minutes, allowing fix agents to be deployed more quickly and reducing MTTR (Mean Time To Repair), and improving the customer experience. But it is not just the fiber that is driving the MTTR enhancements, it is the combination of fiber-intelligent network devices, tooling, process sameness, and automation that drives improved network reliability.

3. Digitization of Data

The documentation of fiber has varied throughout the industry and each group, even within organizations, may have different methodologies to track information. The key areas of fiber documentation are (1) the physical design location (plant maps), (2) the fiber amount or size, and (3) the assignments of what is traversing each fiber. When focusing on repair efforts, the digital analysis and resulting determination of the physical location is crucial to minimizing outage repair time and the digital analytics are the catalyst for making this possible. These analytics define the digital transformation of fiber outage analysis. Digital transformation relies on 4 key business competencies, technology, data, process, and organizational change capability [5].

Technology must contribute not complicate the transformation process and be adaptable to the needs of the business [5]. The first aspect of fiber technology is the fiber itself and how it is utilized within the system architecture. The technology architecture is key to reducing latency and driving resiliency. The fiber location documentation must be captured in a centralized and consistent system that becomes the source of the truth repository. In the past, much of this documentation was maintained in multiple formats, sometimes even as a drawing on the wall of fiber engineers. Digital transformation of the fiber requires that the information be accurately captured in a tool that can then be utilized to provide key location information in the event of a fiber outage.

Data integrity standards may vary, and transformation requires understanding unstructured or invalidated data [5]. Many know that data is important, but sometimes data quality is lacking. Understanding the integrity of fiber information data is an important investment in ensuring the expected impact of digital fiber cut detection is realized. Implementing methodologies to define fiber metadata, automate the capturing of such data, and wavelength management is key to accomplishing data proficiency.

Process transformation requires innovative thinking to envision a new way of utilizing data to evolve manual processes into automation [5]. In many areas, fiber triage requires multiple manual reviews and physical confirmation of the stance of the breakage. Horizontal process management breaks down traditional hierarchal triage methodologies and the results are the implementation of a process-driven by digital and technological capabilities. For XMFR, the remote version of the XMF solution, this requires a detailed review of current processes and alignment with the updated processes to support integration.

Organizational change capability is the final key talent that requires courage to change what is historically rooted in current technology, data, and process [5]. Building new technology solutions is key to digital transformation, but without organizational change, adoption and adaptability will not happen within the business. Committing to understanding the human factor is critical to adoption and success.

All four of these elements are key to the integration of XMFR. The technology platform is the engine that delivers the capability. Data integrity is what fuels the engine to provide key accurate information to fix agents. The process is how we guide the organization to achieve consistent positive results. And organizational change is how we bring all these things together for a successful implementation that is recognized for its value to the business and impact on positive customer experience.

4. Operational Logic, XMF-R

Once the organization is aligned on the digital transformation plan, the operational logic reviews the high-level flow of how an outage is managed and identifies key automation opportunities. In many cases today, fiber outage events are not easily discernable from another outage event. The first step in the operational logic is to identify how to utilize the XMFR capability to create a specific alarm type that indicates a fiber outage event which reduces triage times for the operations center. The next step in the logic is to determine how to eliminate the need to dispatch a headend tech to OTDR the fiber to identify the distance to the fiber event. Direct dispatch of this automated information to the fix agents/persons creates speed to resolution and utilizes the technology and data information to confirm the location of the fiber breakage. The fiber logic flow in Figure 1 illustrates the value of reducing the total number of steps required to drive restoration.




Figure 1 - Operational Logic

Within the operational logic, the reduction of steps optimizes performance, but data and payload provided to the fix agents are essential for optimized restoration. The fix agents require clarity of the event, breakage distance, and location accuracy. These three elements are guides to building the technology, data, and process for operational integration. This one event message then includes:

- Fiber Outage Event
 - Clarity for the fix agent this is a fiber event and not another type of system impairment
- Distance Payload (XMFR OTDR)
 - The XMFR performs an OTDR measurement, and that information is integrated into the event messaging
- Fiber Design Reference
 - Latitude and Longitude information that is referenced to design systems to provide a specific location for the fiber outage event

5. Acting on Real-time Data

5.1. Process Changes

Our operations centers react to many diverse types of alarms and alerts. When a fiber cut is suspected, our normal response for many years was to contact the appropriate teams, based on location, and dispatch them to the appropriate location to troubleshoot the fiber, determine if there was a cut or major impairment, providing the footage details once they were obtained. The Operations Centers would then update the ticket with the footage and dispatch the appropriate fix agents to the site to make the repairs.

With the introduction of XMF, we focused first on deploying the technology between facilities, headend to the headend. This had an immediate operational impact in saving us the time we used to spend sending someone to the facility. This is especially valuable in the instance of an unmanned facility, where, if an issue occurred at night, the time to dispatch, troubleshoot the fiber, and obtain footage could easily be more than an hour. With the XMF solution, the alarm is received in seconds and within a minute or two the operations team receives the distances electronically, updates the ticket accordingly, and dispatches the fix agent.

The second focus included deploying the XMF technology at the node level (XMF-R), in the access network, as we upgrade and shift our nodes from analog to digital. As each node upgrade is implemented, the XMF-R solution is also introduced, allowing for the same benefit across all digital nodes.

The third element that enhances the solution is the end-to-end automation of the ticket and the dispatch with no manual intervention from an operation's center technician. When the XMF solution is installed and configured, many critical details are captured. (The information captured is detailed in the Digitization of our Network section of this paper) This information allows for the autodetection of fiber impairments when our tools see variances based on details captured during configuration (which creates a baseline), versus real-time performance (which either matches the baseline or finds an anomaly to escalate). As tools detect the impairments, tickets are auto-created with all the necessary information and auto-dispatch occurs with no manual, technician intervention. As noted, prior, the result is faster detection and dispatch, and a decrease in fiber troubleshooting/triage and ticketing data entry for the operations center technicians.

5.2. Tool Changes

If the XMF solution only allowed for the time savings associated with removing the need to dispatch and manually shoot the fiber, it would still be worth installing and using. However, we did not stop there! Why capture information automatically and not take advantage of removing manual steps?

Our operations teams utilize National Watch Tower (NWT) for alarming and dispatch activities associated with the Access Network for both demand and planned Maintenance Work. Our NWT teams partnered with engineering to understand the XMF solution and devised a plan to allow NWT to ingest the XMF alerts and payloads, developing the necessary routing solutions to automatically send the demand ticket to the correct fix agent for fiber cuts. Once the fiber cut solution was in place, the NWT team began to develop a similar solution for various types of fiber impairments. Both advancements share the necessary payload details to ensure fix agents have all the needed information to enable the speediest repair.

6. The Access Network

We refer to the access network as the portion of the network from the fiber node to the customer premise. As expected from a top tier service provider, we have been very focused on the reliability of our network. Reliability does not simply mean that the network is 'up', and devices are online. It means the network is operating optimally, delivering expected performance and devices are online. To achieve superior network reliability, we are deploying digital nodes with much more telemetry than traditional analog nodes, utilizing our HFC fiber to deliver expected speeds and capacity, and then complimenting that with intelligent tooling to quickly let us know when the network is not operating optimally, and where to go to fix it quickly.

6.1. Coherent Optics in the Access Domain

A practical deployment strategy of fibers is essential at Comcast. The ability to use a portfolio of fibers for some of the emerging low latency markets, while also driving fiber technology deeper into the network, is the essence of our approach in building the industry initiative of 10G. This technology convergence is an essential balance that optimizes the system architecture.

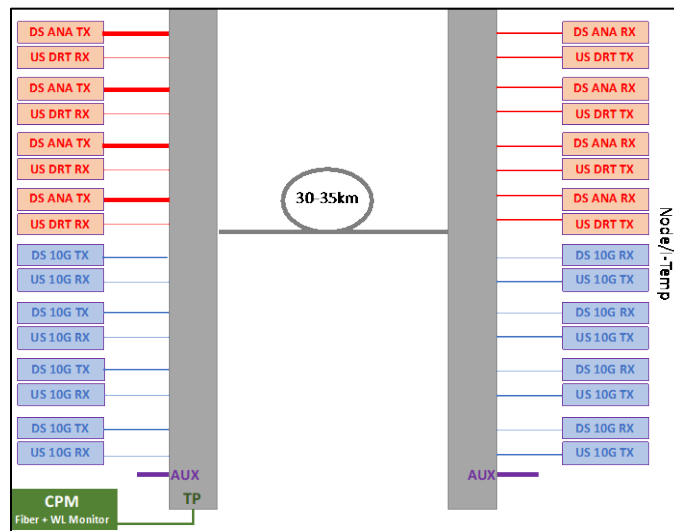


Figure 2 - Access Network Convergence

As the largest broadband company in the US, Comcast serves millions of customers and businesses with a reach that stretches coast to coast. All of this is the result of a large optical network that spans core, metro, and access layers with multiple intersecting points, all intended to increase capacity, reduce latency, and enhance reliability. Operational integration ensures the integrity of the system is maintained with a heightened customer focus.

Photons flood into the Comcast backbone network from the Internet and peer-to-peer traffic through mass-scale routers connected to optical consolidation equipment and reach the various metro centers. At the metro center, the data is reincarnated into photons on a massive IP network and the photons courses thru the highly meshed Converged Regional Area Network infrastructure that terminates in the thousands of our headends. At the various headends, data becomes photons again, traverse access fibers, and light up the many homes, businesses, and fiber nodes eventually completing their journey downstream. A similar process ensures the upstream where photons transmigrate through the access system to the various headends and mesh metro circuits before making their way back through to the internet.

7. Putting it All Together: Role of Orchestration

Deploying an advanced fiber solution and adding intelligence to the network does not immediately deliver the desired benefits. There is a great deal of orchestration across engineering, deployment, tooling/systems development, operations, and, of course, proper people management, required to bring the solution to fruition and maximize the benefit for our employees and our customers. The fiber technology must work together with the operationalization that optimizes event response.

7.1. Deployment and Implementation

Once a new network solution has been developed and tested, it becomes GA (Generally Available) across the network. Engineering teams cannot simply throw it over the fence, expecting adoption by the masses. Engineering partners with the deployment teams to update purchasing plans to acquire the necessary materials and to update our installation plans so the correct materials are ordered, installed, and configured as specified. The solution in this paper includes the fiber purchases as well as the XMF/XMF-R equipment and new policies/procedures for installation and configuration. The operational logic

illustrates how the payload of the OTDR distance and geolocation supports the ability to resolve any post-implementation fiber events quickly and effectively.

7.2. Operational Integration

Once the deployment team has completed their installation and configuration work, they turn it over to operations where validation that the equipment is present within the necessary systems and tools. It is also critical to share these changes and expected benefits with our field teams. We want to make certain everyone is aware of the change and the impact it has on their processes and the customer experience. When deploying the XMF, we communicated the expected reduction in an after-hours dispatch to OTDR fiber, and the fiber benefits regarding reduced latency and increased speed, all of which were well received by the field operations teams

7.3 An Informed Fiber Design

Live fiber information from fiber monitoring tools can improve fiber designs. When teams make design changes to the fiber network, they must have access to real time documentation. Many times, this means a request is made to the Headend or Network Team to validate the fibers or open channels. Utilizing live fiber information, a fiber designer can build a link to a new endpoint knowing what has been provisioned on a mux. This added information improves the speed of design and reduces rework and fiber assignment changes due to the use of outdated fiber system designs.

Use existing fiber information to improve assumptions, resulting in better use of the fiber resource. Fiber designers can utilize real time measured distance and loss to the mux. This new information allows teams to potentially reach further into the network, opposed to using conservative loss or storage assumptions. Fiber teams can also compare system assumptions to better understand the amount of storage or splice loss, from system to system. This new view of the network can ensure that future changes made are based on good information.

7.4 Proper Change Implementation

When implementing process, technological, or organizational changes, it is important to properly lead our people through the change. Detailing what is changing, why it is changing, and understanding the impact on our network and our people is critical to successful implementation and adoption. Recognizing the human factor of change is how organizations optimize evolutionary transformation.

8. Conclusion

Given the general focus on tooling and automation today, it is critical that we analyze all areas of our network to see how we can create better customer experiences and efficient workflows through automation. Given increasingly intelligent alerts and data access, this task becomes increasingly realistic within traditionally manpower intensive areas of the network. This paper discusses the impact of data digitization and the operational logic that is foundational to the technology needed to run advanced fiber networks. The operational logic transforms the manual fiber anomaly analysis procedure from four manual steps to automated processes that reduce troubleshooting time for the business. The process updates drive the development of tool innovation to support automation goals.

The advanced Comcast 400G network described herein, together with the optical monitoring, provisioning, and visualization, we are creating unprecedented robustness and enhancing our customer experience.

At the heart of the Comcast technology evolution is our commitment to our customers and our business goals.

Abbreviations

AP	access point
bps	bits per second
FEC	forward error correction
Hz	hertz
K	kelvin
MTTR	Mean Time To Repair
OTDR	Optical Time Domain Reflectometer
SCTE	Society of Cable Telecommunications Engineers
XMF	Xfinity Meter Fiber
XMF-R	Xfinity Meter Fiber-Remote

Bibliography & References

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

[2] *The ARRL Antenna Book, 20th Ed.*; American Radio Relay League

[3] Code of Federal Regulations, Title 47, Part 76

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

[5] <https://hbr.org/2020/05/digital-transformation-comes-down-to-talent-in-4-key-areas>

Turbocharging DOCSIS 3.1 Technology: An Incremental Step on the Way to DOCSIS 4.0

A Technical Paper prepared for SCTE by

Colin Howlett

Chief Technology Officer
Vecima
colin.howlett@vecima.com

Jay Rolls

Chief Technology Officer
Broadband Success Partners
jrolls@broadbandsuccess.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Network Evolution	4
3. DOCSIS 3.1 Minimum Capability	5
4. DOCSIS 4.0 Outlook and Considerations	5
5. Turbocharging DOCSIS 3.1 Technology	7
5.1. Distributed Access Architecture (DAA)	8
5.2. Profile Management Application (PMA)	9
5.3. Low Latency DOCSIS (LLD)	10
5.4. DOCSIS 4.0 Cable Modems in 1.2 GHz Plant.....	13
6. Operational Considerations.....	14
6.1. Spectrum Reclamation – Sunset QAM Video	14
6.2. Overlapping OFDMA Channels.....	15
6.3. Leakage Detection	17
7. Upgrade Costs	18
8. Conclusion.....	19
Abbreviations	20
Bibliography & References.....	21

List of Figures

Title	Page Number
Figure 1 - DOCSIS 4.0 Frequency Split Options	7
Figure 2 - Distributed Access Architecture Examples.....	9
Figure 3 - PMA Deployment Architecture Example from [Karthik1].....	10
Figure 4 - Bandwidth Usage for Capacity Seeking Applications with TCP	11
Figure 5 - Dual-Queue Approach for LLD from [White1]	12
Figure 6 – Round-Trip Latency with LLD for NQB-Marked Traffic from [White1]	13
Figure 7 - Capacity by Upstream Split and Plant Maximum DS Frequency with D4.0 CM	14
Figure 8 – Need for Overlapping OFDMA Channels	16
Figure 9 – Overlapping OFDMA Channel Arrangement	16
Figure 10 – Overlapping OFDMA Channel Solution – UCD and MAPs	17
Figure 11 - Leakage Detection Test Signals from [Coldren1].....	18

List of Tables

Title	Page Number
Table 1 – DOCSIS 3.1 HFC Outside Plant Minimum Capabilities	5
Table 2 – DOCSIS 4.0 Capabilities.....	7
Table 3 – DOCSIS Cable Modem Minimum Capabilities	13
Table 4 – Relative Upgrade Costs	19

1. Introduction

In North America, cable operators, delivering broadband over their hybrid fiber-coax (HFC) infrastructure, have become leading providers of broadband in the communities that they serve. It's not uncommon for a multi-system operator (MSO) to enjoy 65% market share. And subscriber additions accelerated during the first 18 months of the pandemic.

However, MSOs are now facing growing competition, in multiple forms across their service areas.

These threats mainly originate in three areas. First, fiber competition, both from large telco incumbents such as AT&T, Verizon, and Lumen, and Frontier and small upstart fiber operators. Second, the Mobile Network Operators (MNOs) have been pushing broadband over their mobile platforms very heavily over the past 24 months. (A sub-component also arises from Fixed Wireless Access providers but is a much smaller percentage of wireless broadband). And lastly, a not insignificant industry has materialized that specializes in just servicing multi-dwelling units (MDUs), and that is normally also served via fiber. MDUs can often represent 30-40% of an operator's customers.

Amidst this burgeoning competition, MSOs and their HFC architecture can point to an admirable record of remaining resilient and competitive. In fact, the limitations of the evolving DOCSIS standards have always been discussed, and those limitations have always been shattered.

Up to this point, none of the previous versions of DOCSIS[®] (1.0, 2.0, 3.0, and 3.1) specifications required any major changes to an operator's Outside Plant¹. Implementations have mainly relied on changes to electronics on either end of the 'wire'; either in the headends and hub sites (Inside Plant), or in the home (the cable modem).

Outside Plant changes have been operator-specific and generally driven by overall increases in frequency capability operating within the limits of existing active device (amplifier) housings. These changes have included:

- DS frequency migrations from 550 MHz to 750 MHz, 860 MHz, 1 GHz or even 1.2 GHz by updating node and amplifier modules
- US frequency extension from 42 MHz to 85 MHz or in rare cases to date, 204 MHz, by updating amplifier modules with new frequency diplexers

However, that is about to change.

DOCSIS 4.0², which will be introduced between 2023 and 2025, will require changes to the outside plant. Introducing this new version of DOCSIS, which will activate frequencies all the way up to 1.8GHz, is going to require an overhaul to today's HFC cable plant. This will be a far more invasive and expensive undertaking. Every device in the outside plant will need to be

¹ Perhaps this overlooks the initial 2-way activation of the cable plant in the 1990's, which was a precursor to offering broadband services.

² When discussing DOCSIS 4.0 specifications, this paper will focus on the part that many operators are planning to implement: ESD, or Extended Spectrum DOCSIS, which bumps the cable plant spectral capabilities up to 1.8GHz.

replaced (or in more rare occasions, upgraded). To put this in perspective, using an example of a rather large “fiber node” with a 350 to 400 home serving area, this equates to replacing between 10 and 20 amplifiers, between 10 and 25 splitters, and over 100 coaxial cable taps. The cost to undertake such a retrofit will start to be driven more by the labor involved than the equipment itself.

In this paper, we reference “turbocharged” DOCSIS 3.1 technology as a collection of specific improvements that have been added to DOCSIS 3.1 and DOCSIS 4.0 specifications since the original DOCSIS 3.1 specification was published:

- New deployment architectures (Distributed Access Architecture or DAA)
- Operation with varying impairments at each customer (Profile Management Application or PMA)
- Application latency improvements (Low Latency DOCSIS or LLD)
- Use of enhanced DOCSIS bonding group channel counts in DOCSIS 4.0 cable modems along with higher speed Ethernet interfaces

The “turbocharged” collection of improvements is not a separate specification, but a useful way to categorize an incremental step possible for operators before they deploy networks capable of realizing all parts of the DOCSIS 4.0 specifications. This step is a valuable way to achieve the next step in DOCSIS performance at lower cost than both full DOCSIS 4.0 network deployment and far lower cost than full fiber to the home rollouts.

2. Network Evolution

Broadly speaking, there are a handful of potential technology paths for evolving an HFC network. In between and amongst those options, many smaller choices will need to be made.

The two most significant factors to consider in meeting growing demands on a broadband network are capacity (scale) and speeds. Capacity is the ability to meet the growth in the number of subscribers on the network, as well as the growing consumption needs of each subscriber. Speeds relate to speed tier offerings, and in particular the top speed offered on the network.

To meet the demands outlined above, operators have these options to consider, as they evolve their network:

- Create capacity by segmenting the shared nature of an HFC network into smaller parts. Most often this is facilitated by splitting neighborhoods into smaller service groups, via node segmentations and node splitting.
- Expand the capacity of the network by increasing the breadth of frequencies supported. Choices here include expanding today’s 1GHz plant to 1.2GHz and/or to 1.8GHz.
- Leveraging technology to get more out of the network at hand. This includes running more of the signaling with advanced modulation (orthogonal frequency division multiplexing or OFDM), running higher modulation 2k and 4k QAM, and leveraging systems that can operate on more channels.
- Lastly, changing to a different network altogether; in this case fiber to the premises (FTTP).

3. DOCSIS 3.1 Minimum Capability

DOCSIS 3.1, introduced to field deployment in ~2015 (see [MULPIv3.1] and [PHYv3.1]), has been able to keep pace with advancements needed in capacity and speed to serve both increasing customer usage (widespread uptake in streaming video platforms, proliferation of consumer devices to name a couple) and competition from operators widely deploying FTTP in the same service area.

The current HFC outside plant is faced with several constraints. Today's 1GHz and below plant will eventually create a capacity constraint. And the current low split (~42MHz in North America) configuration, leads to high bandwidth asymmetry, with far more capacity in the downstream than in the upstream. Loosely this is along the lines of a 10:1 ratio.

Upstream speed for subscribers is limited based on frequency split, the condition of the plant (achievable US modulation) and need to support legacy DOCSIS 3.0 and earlier Advanced Time Division Multiple Access (ATDMA) channels and is summarized in Table 1. Operators looking to move to support gigabit upstream (~1 Gbps) speeds are now actively testing, trialing, or deploying High Split networks to move the upstream upper band edge frequency to 204MHz and the downstream lower band edge frequency to 258 MHz.

In the downstream direction, DOCSIS 3.1 cable modems have a minimum bonding group of 32 channels of single carrier QAM (SC-QAM) and two 192 MHz orthogonal frequency division multiplexing (OFDM) channels. With 256QAM modulation for SC-QAM and 4096QAM modulation for DS, the maximum throughput of this bonding group is 32 x ~35 Mbps + 2 x ~1.7 Gbps or 4.5-4.6 Gbps. The availability of spectrum for this total capacity depends on maximum DS frequency limit and the amount of spectrum used to delivery QAM video service, both of which are operator- and location-specific.

Commercially available DOCSIS 3.1 cable modems today are limited to supporting the minimum bonding group noted above and are also limited to maximum Ethernet interfaces using 2.5 Gigabit Ethernet (2.5GE). The combination of these commercially available minimum implementations results in service offerings with the values summarized in Table 1:

Table 1 – DOCSIS 3.1 HFC Outside Plant Minimum Capabilities

Upstream Frequency Upper Band Edge	Upstream Capacity	Maximum Upstream Service Tier	Downstream Bonding Group Capacity	Maximum Downstream Service Tier
Low Split (42MHz)	100-250 Mbps	50-100 Mbps	4.5-4.6 Gbps	3 Gbps
Mid Split (85MHz)	~400-550 Mbps	250 Mbps	4.5-4.6 Gbps	3 Gbps
High Split (204MHz)	~1.3-1.6 Gbps	1 Gbps	4.5-4.6 Gbps	3 Gbps

4. DOCSIS 4.0 Outlook and Considerations

The next version of DOCSIS, 4.0, is due to see equipment released between 2023 and 2025.

The primary objectives in DOCSIS 4.0 as the next version include:

- Support for customer upstream service tiers beyond 1 Gbps
- Support for customer downstream service tiers beyond 2 Gbps

The DOCSIS 4.0 standard (see [MULPIv4.0] and [PHYv4.0]) presents two different implementation options depending on an operator's objectives. The options are:

1) FDX - 1.2 GHz Full Duplex (FDX)

FDX is generally intended for Node+0 HFC with no active amplifiers but industry initiatives are underway to build FDX repeater amplifiers to stretch operation into networks including amplifiers. FDX targets maximum DS frequency with the same 1218 MHz limit as DOCSIS 3.1.

2) FDD/ESD - 1.8 GHz Frequency Division Duplex (FDD)

FDD/ESD is generally intended as an incremental upgrade to a traditional HFC network of actives and passives with an arbitrary number of amplifiers. This second option is commonly referred to as ESD, or Extended Spectrum DOCSIS. Industry demonstrations in 2022 have shown ESD operation to be compatible with traditional amplifier cascades and plant designs. This paper specifically focuses on ESD as an incremental path from traditional DOCSIS 3.1 to DOCSIS 4.0 technology.

Upgrading an HFC cable plant to 1.8GHz will be challenging. The signal attenuation on a cable plant to signals above 1.2GHz are significant. The passive components currently deployed are designed for 1GHz or 1.2GHz and may extend slightly higher, but many of these components will need to be changed out to support operation to the 1.8GHz maximum capability of the DOCSIS 4.0 FDD specification:

- Taps required to direct a portion of the signal energy to/from specific subscribers
- Power inserters for feeding AC into the plant
- Directional couplers and splitters used within the trunk portions of the coaxial

Likely some design changes will be required to other parts of the HFC cable plant. Some areas may require smaller booster amplifiers, estimated to be in the range of 10-20% in early design studies from multiple North American operators. In some cases (less than 5%), older deployed coax cable itself might not support 1.8GHz, requiring new coax to be deployed. These are areas of study that will be necessary for an operator to perform as each company assesses their readiness for DOCSIS 4.0, and their ability to meet the challenges that will arise.

Since the outside plant is changing and frequency is being pushed higher than what DOCSIS has done before in the HFC network, the implementation of DOCSIS 4.0 will require more extensive lab and field testing than previous versions. Many operators are likely to take a "wait and see" approach to DOCSIS 4.0 to see how early field tests and deployments work before making decisions on the future of the HFC plant.

Even when deciding to move forward, with DOCSIS 4.0 deployment, expectations are that operators will progress through multiple incremental phases of DOCSIS 4.0 ultra-high frequency

splits (UHS) as capacity and speed requirements increase. These frequency split options are shown in Figure 1:




Figure 1 - DOCSIS 4.0 Frequency Split Options

Note that the DOCSIS 4.0 standard also includes other splits but these are not included based on operator feedback on probable deployments:

- UHS-300 (300/372MHz) offers only a small improvement above 204 MHz so operators are much more likely to switch to UHS-396 or UHS-492 as the next increment
- UHS-684 (684/834MHz) has no additional downstream bandwidth while requiring more spectrum and providing very high peak upstream rates that are unlikely to be used by service offerings

DOCSIS 4.0 capabilities for these tiers are summarized in Table 2 based on up to 900 Mbps per 96 MHz OFDMA (700 Mbps in low band) and 1.8 Gbps per 192 MHz OFDM channel:

Table 2 – DOCSIS 4.0 Capabilities

Upstream Frequency Upper Band Edge	Upstream Capacity	Maximum Upstream Service Tier	Downstream Bonding Group Capacity	Maximum Downstream Service Tier
High Split (204MHz)	1.6 Gbps	1 Gbps	14.4 Gbps	10 Gbps
UHS-300 (300MHz)	2.6 Gbps	2 Gbps in great plant	13.3 Gbps	10 Gbps
UHS-396 (396MHz)	3.5 Gbps	2 Gbps	12.2 Gbps	10 Gbps
UHS-492 (492MHz)	4.4 Gbps	3 Gbps	11.1 Gbps	8 Gbps
UHS-684 (684MHz)	6.2 Gbps	5 Gbps in great plant	9.0 Gbps	6 Gbps

5. Turbocharging DOCSIS 3.1 Technology

Since the first field deployment of DOCSIS 3.1 networks in ~2015, cable operators and vendors have continued to innovate to improve the performance and operation of DOCSIS-based systems. Bundled together, these enhancements on top of the original DOCSIS 3.1 turbocharge DOCSIS 3.1 technology as a concept of a noticeably improved DOCSIS implementation suited to long term deployment, and as an incremental step towards full DOCSIS 4.0 deployment.

These options available in the DOCSIS 3.1 specifications, if leveraged efficiently, can instantly turbocharge any DOCSIS 3.1 deployment:

- High Split frequency split – full support of the upstream capability specified in DOCSIS 3.1 specifications to enable gigabit upstream service tiers

- Distributed Access Architecture (DAA) – extend all-digital Ethernet fiber deep into the network for reduction in hub space and power needed for an increased number of service group, along with improvements in RF performance
- Profile Management Application (PMA) – optimized subcarrier modulation in downstream and upstream to improve channel capacity for impaired modems or when overall RF plant conditions are non-optimal
- Low Latency DOCSIS (LLD) – reducing overall end-end packet delay through optimized queuing and standardized packet marking, especially for applications such as gaming
- DOCSIS 4.0 cable modems in 1.2 GHz plant – taking advantage of extra OFDM and OFDMA channels likely to come in DOCSIS 4.0 CMs to increase overall bonding group capacity in both DS and US while also supporting service tiers above 2 Gbps in the DS

5.1. Distributed Access Architecture (DAA)

The cable industry has long relied on a centralized architecture in configuring traditional HFC networks. Signals that were broadcast on the cable plant originated in centralized headend and hub site facilities. This practice continued with the introduction of broadband services delivered via the DOCSIS standard. However, the industry is in the middle of a major transition away from this approach, moving electronics from centralized to more distributed locations, with key elements (electronics) being placed in the outside plant fiber nodes. The broad acronym in use to describe this new approach is DAA – Distributed Access Architecture.

A primary initial use case for DAA was to reduce hub space and power by moving the modulation and demodulation of physical layer (PHY) signals (QAM and OFDM) from the hub to the node location. This reduction is necessary to avoid hub expansions as the number of subscribers per service group continues to shrink to provide more capacity and speed to individual subscribers so more DOCSIS and RF equipment is needed. DAA allows the removal of RF equipment including RF combiners, analog/digital optical RF systems, and the extensive wiring needed to connect everything together. Instead, all hub DOCSIS-related equipment is now digital with simplified Ethernet/IP interconnects or removed altogether in the case of Remote-MACPHY architectures.

Along with the reduction in hub space and power, one of the most important advantages of this new approach is that it allows operators to move from “analog optics” to “digital optics” – thus simplifying how signals are transmitted on the fiber portion of an HFC network, and greatly increasing the fidelity of the signals that are transmitted. Field experience from a number of DAA deployments has shown 5 to 7 dB improvement in C/N (carrier to noise ratio), which in essence represents a doubling in RF performance. Such an improvement allows for higher modulation (more bits per hertz), thus increasing overall system capacity.

The advent of DAA also holds promise in addressing some operational problems. Some areas that might be operating on the margin today (ie. Signal quality), could see improved performance from a DAA implementation and reduced trouble calls.

DAA implementations come in two primary configurations specified by CableLabs under the umbrella of Distributed CCAP Architecture (DCA) (see [DCA]). These configurations are

known as Remote-PHY, and Remote-MACPHY. Both have the advantage of moving to digital optics. Both can serve to allow for better RF performance out of the cable plant. Entire technical papers have been devoted to the merits of one over the other, with details being beyond the scope of this analysis. Figure 2 shows a high level view of the DAA architectures.

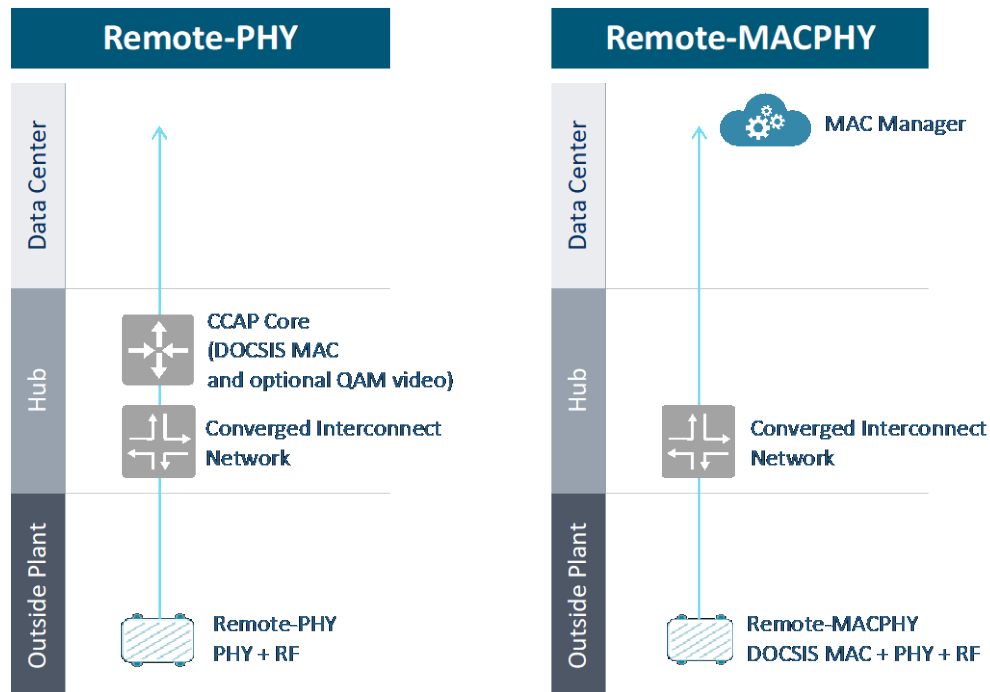


Figure 2 – Distributed Access Architecture Examples

One extremely important consideration, accepted by the industry while developing the specifications, is that DAA architecture will be required to implement DOCSIS 4.0. DAA will be paramount in addressing challenges with supporting spectrum as high as 1.8GHz. As an incremental step towards DOCSIS 4.0, turbocharging DOCSIS 3.1 technology assumes the use of DAA as well.

5.2. Profile Management Application (PMA)

OFDM DS and OFDMA US channels as specified in [MULPIv3.1] and [PHYv3.1] support the use of multiple “profiles” which allow for different data modulations (256QAM, 1024QAM, 4096QAM, etc.) to be configured for each modulated subcarrier. The use of these profiles allow the HFC system performance, in both robustness and overall capacity, to be optimized to the current conditions in the network.

The Profile Management Application (PMA) (see [PMA-TR]) as shown in Figure 3 is an external software solution which uses data on receiver MER and codeword errors from cable modems and CMTS, along with sophisticated algorithms and significant server processing power to optimize the set of available profiles to match current conditions. Without the dynamic changes in available profiles through PMA, modems that are temporarily impaired or operating outside normal design targets would generally use a default modulation of 256QAM.




Figure 3 - PMA Deployment Architecture Example from [Karthik1]

Since the cable plant channel model is quasi-static with slow variations over time due to temperature effects or specific impairment events, PMA solutions update profiles on a periodic interval measured in minutes.

The types of impairments or plant conditions that benefit from PMA include:

- Ingress from wireless communications within the cable band
- Operation in the rolloff region above the current plant frequency. This is particularly interesting as part of turbocharging DOCSIS 3.1 technology since operators can use PMA to extend into a rolloff region (say 1.1 GHz for 1 GHz tap and gain several hundred Mbps of additional capacity without the labor costs of swapping taps in the field for 1.2 GHz.
- Standing waves from impedance mismatches, grounding issues, or other non-ideal operation

Gains achievable, as seen in [Karthik1] are based on a significant set of operator data across several geographies. Capacity gains in excess of 30% improvement over a default 256QAM profile are achievable using PMA techniques.

5.3. Low Latency DOCSIS (LLD)

Many applications such as gaming, interactive videoconferencing, and web browsing do not require significant capacity, but instead require timely responses from remote servers to provide the best user experience. Low Latency DOCSIS technology (LLD), as part of [MULPIv3.1], has been developed by CableLabs, cable operators, and vendors to improve overall application latency, especially for those applications that are not using significant bandwidth.

The primary problem to solve from an end-to-end perspective relates to applications and transport protocols (e.g. TCP) which are capacity seeking. High bandwidth applications like file transfers and streaming IP video run over TCP, using up available network bandwidth until they

experience congestion. Congestion is recognized by way of dropped packets, at which time the application backs off until congestion ceases and then ramps back up until congestion is experienced again, at which time they once again back off. This cycle is repeated, creating a sawtooth pattern of bandwidth usage as shown in Figure 4.



Figure 4 - Bandwidth Usage for Capacity Seeking Applications with TCP

The problem this creates for applications that need low latency is that their packets get queued up during periods of congestion and experience latency and jitter patterns that follow the sawtooth bandwidth usage pattern of the capacity seeking application protocols. The industry has settled on the terms "queue building", to describe capacity seeking applications and transports, and "non-queue building" (NQB), to describe the transport that is needed by low latency applications.

Early methods for dealing with queue building applications included creating larger queues in network devices to absorb bursts without loss of packets. However, these methods were easily defeated by applications which just continued to ramp up bandwidth demands until packet loss was experienced. This made delay and jitter even worse and was referred to unkindly as "buffer bloat".

The solutions included in the overall LLD ecosystem include the following. Many of these solutions are just now progressing to readiness in the broader consumer ecosystem including gaming platforms and major device operating systems:

- A new NQB packet marking that applications can use to indicate that they are non-queue building so that their traffic can be treated differently in the network (e.g. classified into a Low Latency Service Flow), see [NQB1]
- Recommended use of Explicit Congestion Notification (ECN) and ECN Capable Transport (ECT) markings in the Traffic Class field of the IP header by low latency applications and networking equipment (e.g. routers), see [ECN1]
- Support for a new congestion control scheme called Low Latency, Low Loss, Scalable Throughput (L4S) (see [L4S1] and [L4S2]), which leverages Active Queue Management (AQM) techniques and which is intended to be applied in next generation transport

protocols supporting low latency applications. L4S congestion controls are leveraged in low latency applications and in network equipment as follows:

- Low latency applications will mark ECT in their packets
 - Network equipment experiencing congestion involving that ECT traffic will mark ECN Congestion Experienced (CE) in packets before forwarding rather than dropping packets
 - Low latency applications supporting ECT and receiving ECN CE will respond by marking ECN CE in traffic toward the originating application
 - The originating application will respond by reducing its transmission rate
- DOCSIS queuing improvements address queueing delay by allowing applications to avoid waiting behind the delays caused by the current TCP or its variants. At a high level, the low-latency architecture consists of a dual-path approach that treats both queues as a single pool of bandwidth as shown in Figure 5:




Figure 5 - Dual-Queue Approach for LLD from [White1]

- DOCSIS scheduler improvements address media acquisition delays by:
 - Lowering the request-grant delay with a shorter MAP Interval broadcast by the CMTS and associated MAP Processing Time in the cable modem
 - Adding a new scheduler service known as Proactive Grant Service (PGS) to proactively grant to a service flow without incurring the request-grant delay

Simulations published in [White1] show the dramatic improvements in round-trip latency and especially consistency of round-trip latency possible for NQB-Marked Traffic. These results show the 99th percentile of traffic having round-trip latency (DOCSIS part of network) well below 10ms without PGS and below 1ms with PGS, results that are 2-3 orders of magnitude better than standard DOCSIS 3.1.




Figure 6 – Round-Trip Latency with LLD for NQB-Marked Traffic from [White1]

5.4. DOCSIS 4.0 Cable Modems in 1.2 GHz Plant

One key change in the DOCSIS 4.0 specifications is an increase in the minimum number of OFDM DS and OFDMA US channels that can be bonded by the cable modem as shown in Table 3.

Table 3 – DOCSIS Cable Modem Minimum Capabilities

Modem Generation	Upstream OFDMA Channels	Downstream Channels
DOCSIS 3.1	2 x 96 MHz	32 x SC-QAM + 2 x 192 MHz OFDM
DOCSIS 4.0	7 x 96 MHz	32 x SC-QAM + 5 x 192 MHz

In addition, commercial DOCSIS 4.0 cable modems are expected to shift from 2.5GE Ethernet interfaces to 10GE Ethernet interfaces to take advantage of the bonding group capacity to deliver very high single cable modem speeds and enable new service tiers not possible today.

As shown in Figure 7, leveraging this capability without incurring the extra costs of full 1.8 GHz outside plant changes allows an operator to provide a premium service tier of 5 Gbps or higher instead of the 3 Gbps limits resulting from maximum bonding in DOCSIS 3.1.

The UHS-396 with 1.2 GHz plant and UHS-492 with 1.2 GHz plant options in Figure 7 are unlikely to be interesting for deployment unless upstream usage and customer need changes significantly from current projections. The overall DS capacity is limited in 1.2 GHz plant. In addition, amplifiers capable of supporting UHS-396 and UHS-492 diplexers are likely 1.8 GHz capable.




Figure 7 - Capacity by Upstream Split and Plant Maximum DS Frequency with D4.0 CM

6. Operational Considerations

Several operational considerations need careful attention from operators to ensure the largest benefit can be achieved from turbocharging DOCSIS 3.1 technology. These include:

- Spectrum Reclamation – transitioning DS spectrum from QAM video to DOCSIS data use
- Overlapping OFDM Channels – ensuring Mid Split and High Split modems can utilize their full available bandwidth when mixed in the same service group
- Leakage Detection – ensuring regulatory compliance when sensitive aeronautical bands in the 130 MHz range move from being DS signals to US signals as part of High Split transitions

6.1. Spectrum Reclamation – Sunset QAM Video

A key service still provided by most operator networks today is QAM digital video delivery to traditional set-top boxes. The number of channels used is dependent on market and QAM video technology (i.e. broadcast vs. switched digital video (SDV)) but can occupy 30-60 or even 80 channels.

Concerns that arise from keeping QAM video available in the network:

- Insufficient spectrum to transition to High Split

In this case, the transition from low or mid split to high split requires freeing up 150 MHz (108-258 MHz for mid to high split) or 204 MHz (54-258 MHz) of downstream spectrum. If downstream spectrum is relatively full already due to multiple OFDM channels or outside plant DS frequency limits, this spectrum reclamation may not be possible.

- Insufficient DOCSIS spectrum for service tier goals

This case is a concern when pushing service tiers to 2 Gbps and above, especially if using DOCSIS 4.0 cable modems to increase DS bonding group size. QAM video spectrum needs to be reclaimed to achieve maximum capacity.

- Legacy set top box carriers – SCTE 55-1 or SCTE 55-2

Support for operation of the legacy set top box DS carriers above 130 MHz for SCTE 55-1 or 55-2 is a possibility with some set-top boxes but this varies by the specific set top box. Adaptation solutions do exist to deal with these issues: downconverters at the STB from ~250 MHz to 70-130 MHz for seamless operation or use of DOCSIS Set Top Gateway (DSG) to out-of-band (OOB) carrier generators can help.

The primary solution for Spectrum Reclamation is transition to IP video service while sunsetting QAM video except possibly for niche use cases. This transition, already done at many operators, with the full digital QAM video lineup available in IP format to stream as data over DOCSIS channels, also helps operators deliver video service to FTTP customers without the complications associated with RF over the fiber.

6.2. Overlapping OFDMA Channels

The need for Overlapping OFDMA Channels (OOC) comes from DOCSIS 3.1 requirements and common implementations which limit the number of available OFDMA channels in an Upstream Service Group to two.

Deployment of High Split ideally has one of the two OFDMA channels placed at 108 to 204 MHz, and the other OFDMA channel placed immediately below that (e.g., 12 to 108 MHz). The problem is that many operators have fielded Mid-Split DOCSIS 3.1 CMs, which have a diplexer at 85 MHz and which are unable to make use of any OFDMA channel that spans 85 to 108 MHz (or which goes above 85 MHz, to be precise). See Figure 8 for further details on the concern.




Figure 8 – Need for Overlapping OFDMA Channels

As a result, in moving to High-Split operators may be forced to downgrade service levels for subscribers with Mid-Split CMs as these CMs would no longer be able to use any OFDMA. Conversely, the operator would be forced to place one of the OFDMA channels completely below 85 MHz (e.g. 12 to 85 MHz), which leaves spectrum from 85 to 108 MHz vacant.

To make matters worse, most operators are reluctant to deploy time and frequency division multiplexing (TaFDM) which overlaps ATDMA and OFDMA channels and ensures that there are no overlaps in the scheduler. If avoiding using TaFDM, these operators would carve out explicit spectrum for ATDMA channels below 85 MHz and explicit spectrum for the OFDMA channel below 85 MHz, making the OFDMA channel that exists there very small.

DOCSIS 4.0 CMs are also not required to be able to make use of 85 to 108 MHz. This leads to an analogous problem for those CMs when Ultra-High Split is in play (i.e., they cannot make use of an OFDMA channel that spans 85 to 108 MHz).

The solution to this issue is shown in Figure 9 and Figure 10 and now included in [MULPIv3.1]. Each CM (Mid-Split or High-Split) sees the channel that it is capable of transmitting, while the CMTS generates different long and short Upstream Channel Descriptors (UCDs) and MAPs. No schedule conflicts or channel descriptor conflicts exist since Mid-Split just sees a shortened version of the same messages.




Figure 9 – Overlapping OFDMA Channel Arrangement




Figure 10 – Overlapping OFDMA Channel Solution – UCD and MAPs

Implementation of OOC for a plant transitioning from Mid-Split to High-Split can maintain an additional ~200 Mbps of capacity while allowing both Mid-Split and High-Split CMs to be simultaneously deployed.

6.3. Leakage Detection

System leakage monitoring and detection is required to ensure regulatory compliance and avoid interfering with the sensitive aeronautical band in the 108 to 137 MHz range. In traditional Low-Split and Mid-Split deployments, legacy methods for accomplishing this using downstream spectrum and portable field meters detecting specialized CW carriers have been in place for many years and are well-understood.

With High-Split and DOCSIS 4.0 Ultra-High Split frequency plans, the aeronautical band is no longer in the downstream spectrum and instead falls within the upstream spectrum. With these splits, leakage test signals must instead be generated by the CM just above the aeronautical band as shown in Figure 11. The CM uses OFDM Upstream Data Profile (OUDP) test probes, normally assisting with upstream profile validation, to create specific pilot patterns which are detected by the field meters. [Coldren1] provides a summary of the approach.




Figure 11 - Leakage Detection Test Signals from [Coldren1]

Since the standard updates were published in 2021, significant progress has been made towards industry interoperability between operators, field meter vendors, and CMTS/node vendors. Field results has proven the ability to generate the signals in the CM, detect the signals with the field meter, and effectively troubleshoot leakage situations in the outside plant.

7. Upgrade Costs

Absolute costs for upgrading a particular deployed network are location-specific and depend on many factors such as existing installed equipment and cables, operator service offerings, regional regulatory requirements, labor rates, and operational practices.

In assessing turbocharged DOCSIS 3.1 technology as a step in HFC deployment, four upgrade points are considered for relative cost per home passed:

- 1) Turbocharged Baseline: This is an upgrade of the HFC actives (amplifiers and nodes) to high split with 1.2 GHz maximum DS. “Turbocharged Baseline” Assumptions:
 - Passives are assumed to support >1 GHz operation and are not deliberately a part of the upgrade
 - No significant changes to plant design or spacing
 - Amplifiers are upgraded in place to 1.2 GHz – only obsolete or unsuitable amplifiers are replaced
 - Nodes are transitioned from analog and/or digital return nodes to DOCSIS 3.1 DAA nodes
 - Cable modems are DOCSIS 3.1
- 2) Turbocharged Baseline with D4.0 Cable Modem
 - Above and beyond the Baseline scenario, customers are provided with DOCSIS 4.0 cable modems to enable higher DS tiers above 2 Gbps
- 3) Full DOCSIS 4.0 Upgrade
 - All passives are proactively replaced in the network to support 1.8 GHz or greater operation
 - Amplifiers are upgraded in place or swapped for 1.8 GHz

- Booster amplifiers are added in 10-20% of locations
- Some cable replacement (1-3%) is required
- Labor is a very significant portion of the overall cost in this scenario due to the amount of plant touch

4) FTTP Upgrade

- Service group is swapped over to fully fiber to the premise
- Average/median costs across a wide range of geographies and types are used for this analysis

Based on industry and operator feedback for each of these items, we estimate the relative upgrade costs as shown in Table 4:

Table 4 – Relative Upgrade Costs

Upgrade Scenario	Relative Upgrade Cost
Turbocharged Baseline	1.0
Turbocharged Baseline + D4.0 CMs	1.2-1.5
Full D4.0 Upgrade	1.8-2.4
FTTP	10

8. Conclusion

Operators have many tools available to increase capacity and performance of their existing HFC networks to address evolving customer experience targets and competition. DOCSIS 4.0 technology is clearly the long term future of HFC to provide multi-gigabit downstream and upstream capability, but may require significant plant investment to roll out 1.8 GHz capability across the plant to take full advantage of all that is available in the specification.

A more incremental approach with turbocharged DOCSIS 3.1 technology is possible – take advantage of the changes that have been added in the DOCSIS 3.1 specifications with recent industry performance enhancements and include the higher multi-channel capability of DOCSIS 4.0 modems if needed. This turbocharged approach allows an operator to gradually roll out changes to the outside plant while enabling DS tiers to 5 Gbps and significant improvements in latency.

HFC will continue to deliver the service needed by customers for many years – and turbocharging DOCSIS 3.1 technology provides yet another way to get there incrementally.

Abbreviations

AC	alternating current
AQM	Active Queue Management
ASF	Aggregate Service Flow
ATDMA	Advanced Time Division Multiple Access
bps	bits per second
CCAP	Converged Cable Access Platform
CMTS	Cable Modem Termination System
DAA	Distributed Access Architecture
dB	Decibel
DCA	Distributed CCAP Architecture
DOCSIS	Data-over-Cable Service Interface Specifications
DS	Downstream
ECN	Explicit Congestion Notification
ECN CE	ECN Congestion Experienced
ECT	ECN Capable Transport
ESD	Extended Spectrum DOCSIS
FDD	Frequency Division Duplexing
FDX	Full Duplex
FTTP	fiber-to-the-premises
HFC	hybrid fiber-coax
Hz	Hertz
IP	Internet Protocol
L4S	Low Latency, Low Loss, Scalable Throughput (L4S)
LLD	Low Latency DOCSIS
MAC	media access control layer
MDU	multi-dwelling unit
MNO	mobile network operator
MSO	multi-system operator
NQB	non-queue building
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PGS	Proactive Grant Service
PHY	physical layer
PMA	Profile Management Application
QAM	quadrature amplitude modulation
RF	radio frequency
SCTE	Society of Cable Telecommunications Engineers
SDV	switched digital video
TCP	Transmission Control Protocol
UHS	ultra-high frequency split

Bibliography & References

[Coldren1] – *Leakage Detection in a High Split World: Industry Progress Toward a Viable Solution*, Rex Coldren et al., SCTE Expo 2021

[DCA] – *Distributed CCAP Architectures Overview Technical Report*, CM-TR-DCA-V01-150908, CableLabs

[ECN1] - *Explicit Congestion Notification (ECN) Protocol for Very Low Queuing Delay (L4S)*, draft-ietf-tsvwg-ecn-l4s-id-28, K. De Schepper et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-ecn-l4s-id-28>

[Karthik1] – *Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)*, Karthik Sundaresan et al., SCTE Expo 2019

[L4S1] - *Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service Architecture*, draft-ietf-tsvwg-l4s-arch-19, Bob Briscoe et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-l4s-arch/>

[L4S2] - *DualQ Coupled AQMs for Low Latency, Low Loss and Scalable Throughput (L4S)*, draft-ietf-tsvwg-aqm-dualq-coupled-24, K. De Schepper et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-aqm-dualq-coupled-24>

[MULPIv3.1] – *DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv3.1-I23-220328, CableLabs

[MULPIv4.0] – *DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv4.0-I05-220328, CableLabs

[NQB1] - *A Non-Queue-Building Per-Hop Behavior (NQB PHB) for Differentiated Services*, draft-ietf-tsvwg-nqb-10, Greg White et al., <https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-nqb-10>

[PHYv3.1] – *DOCSIS 3.1 Physical Layer Specification*, CM-SP-PHYv3.1-I19-211110, CableLabs

[PHYv4.0] – *DOCSIS 4.0 Physical Layer Specification*, CM-SP-PHYv4.0-I05-220328, CableLabs

[PMA-TR] – *DOCSIS 3.1 Profile Management Application Technical Report*, CM-TR-PMA-V01-180530, CableLabs

[White1] – *Low Latency DOCSIS, Overview and Performance Characteristics*, Greg White et al., SCTE Expo 2019

Understanding Latency across PON systems

(With few comparisons to DOCSIS systems)

A Technical Paper prepared for SCTE by

Karthik Sundaresan

Distinguished Technologist
CableLabs
858 Coal Creek Circle, Louisville
3036613895
k.sundaresan@cablelabs.com

Evariste Some

R&D Wired Intern, PhD Student,
CableLabs, & University of Colorado, Boulder
858 Coal Creek Circle, Louisville
3036619100
e.some@cablelabs.com

Special Acknowledgments for Lab support to:

Doug Jones, Jay Zhu, Sheldon Webster, Aaron Quinto, CableLabs

James Lin, Radhouan Allani, Kyrio

Randy Wiggins, Tibit

Table of Contents

Title	Page Number
1. Introduction.....	4
2. PON Technologies Background.....	4
3. Background of EPON and GPON networks.....	6
3.1. EPON Latency background.....	7
3.2. GPON Latency Background.....	8
4. Results from Latency Measurements.....	9
4.1. Lab experimental Setup	9
4.2. Test Scenarios	10
4.3. PON Throughput Test Results	11
4.3.1. XGSPON Throughput	11
4.3.1. 10GEPON Throughput.....	12
4.4. Baseline Latency (RTT) Test Results	13
4.4.1. Baseline XGSPON Latency	14
4.4.2. Baseline 10GEPON Latency.....	14
4.4.3. Baseline DOCSIS Latency.....	15
4.4.4. Baseline DOCSIS (with LLD) Latency	15
4.5. UDP Load Test Results.....	16
4.5.1. GPON UDP 1Gbps Test Results	16
4.5.2. GPON UDP 9Gbps Test Results	17
4.5.3. EPON UDP 1Gbps Test Results.....	18
4.5.4. EPON UDP 9Gbps Test Results.....	19
4.5.5. DOCSIS UDP Test Results.....	20
4.6. TCP Load Test Results	22
4.6.1. XGSPON TCP Test results	22
4.6.2. 10GEPON TCP Test results	23
4.6.3. DOCSIS TCP Test results	24
4.7. GPON/ EPON /DOCSIS Test Results summary.....	25
4.8. Multiple LLIDs Test Results (EPON).....	27
4.9. Challenges and Further testing.....	29
5. Conclusions.....	29
Abbreviations	30
Bibliography & References.....	30

List of Figures

Title	Page Number
Figure 1 – Evolution of PON Acces technology	5
Figure 2 – Simplified EPON Network.....	6
Figure 3 – Simplified Media access in EPON: Gate-Report-data	7
Figure 4 – GPON Upstream Transmission mechanisms	9
Figure 5 – PON Network Lab Test Setup	10
Figure 6 – XGSPON Throughput & one-way Latency(UDP)	12
Figure 7 – XGSPON Throughput & one-way Latency in Numbers (UDP).....	12
Figure 8 – 10GEPON Throughput & one-way Latency(UDP).....	13
Figure 9 – 10GEPON Throughput & one-way Latency in Numbers (UDP).....	13

Figure 10 – Baseline XGSPON latency (no load).....	14
Figure 11 – Baseline 10GEPON latency (no load)	14
Figure 12 – Baseline DOCSIS (with No LLD enabled) and no load	15
Figure 13 – Baseline DOCSIS with LLD latency and no load.....	16
Figure 14 – XGSPON latency – with 1 Gbps UDP background traffic, non-load ONU	16
Figure 15 – XGSPON latency – w 1 Gbps background traffic, Load ONU	17
Figure 16 – XGSPON latency – with 9 Gbps background traffic on non-load ONU	17
Figure 17 – XGSPON latency – with 9 Gbps background traffic on load ONU	18
Figure 18 – 10GEPON latency – no load.....	18
Figure 19 – 10GEPON latency – w 1 Gbps UDP background traffic- load ONU.....	19
Figure 20 – 10GEPON latency – w 9 Gbps background traffic - non-load ONU	19
Figure 21 – 10GEPON latency – w 9 Gbps background traffic - load ONU	20
Figure 22 – D3.1 (with AQM) latency – (4.7 and 1.7 Gbps UDP Load).....	21
Figure 23 – D3.1 with LLD, UDP test at 10% above service flow limit	21
Figure 24 – XGSPON latency – w TCP(1 Gbps, byteblower) - non-load ONU	22
Figure 25 – XGSPON latency – w TCP(9 Gbps,byteblower) - load ONU	22
Figure 26 – XGSPON latency – with Iperf TCP on non-load ONU	23
Figure 27 – XGSPON latency – with Iperf TCP on Load ONU	23
Figure 28 – 10GEPON latency – with Iperf TCP on non-load ONU	23
Figure 29 – 10GEPON latency – with Iperf TCP on load ONU.....	24
Figure 30 – DOCSIS no AQM config– with 1 Gbps TCP.....	24
Figure 31 – DOCSIS AQM Enabled Config, no LLD and Iperf TCP load	25
Figure 32 – DOCSIS LLD Enabled, Iperf TCP load	25
Figure 33 – GPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU	27
Figure 34 – EPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU	27
Figure 35 –10GEPON throughput (5 LLIDs 2Gbps per LLID)	28
Figure 36 –10GEPON Latency (5 LLIDs 2Gbps per LLID).....	28
Figure 37 –10GEPON visual of Throughput & Latency (for one US and DS LLID).....	28

List of Tables

Title	Page Number
Table 1 – PON Configuration settings	10
Table 2 – Latency Tests.....	11
Table 3 – GPON latency Test results , P0,P50,P99	26
Table 4 – EPON latency Test results , P0,P50,P99	26
Table 5 – DOCSIS latency Test results , P0,P50,P99.....	26

1. Introduction

HFC/DOCSIS networks are the most widely deployed technology for delivering Internet data services to the consumers. A passive optical network is another option for access technology deployed by operators using EPON, or GPON technology.

One of the performance questions which many in the operator community want to understand is the latency performance of each type of technology and how they compare. We setup a 10GEPON and XGSPON systems in our labs at CableLabs. We ran speed and latency benchmarks on these types of networks in a lab environment and compared the latencies that we get under a variety of conditions. DOCSIS has various options and configurations (D3.0 SC-QAM channels, D3.1 OFDM/OFDMA channels, AQM, LLD etc., while PON has various flavors: 10G-EPON, XGSPON etc. and configurations as well). This paper will give us a better understanding of the efficiencies and the latencies of each of these technology options and the tradeoffs this entails. Based on time constraints during lab testing, this paper is focused primarily on the PON latencies, with some limited testing on DOCSIS systems and a more complete testing and analysis will be part of future work. This paper will give a brief overview of each technology and its evolution and bring out a theoretical comparison of the latency characteristics of each type of technology and mainly reports on the lab testing results for each technology under different operating and load conditions.

2. PON Technologies Background

The continuous growth in bandwidth usage has put demands on the core, metro, and access networks. Gigabit access speeds are now the norm with a large share of US population having an option. Per NCTA, 88% of US Homes have access to gigabit Internet speeds available to them.

Passive optical networks (PON) have been used to serve the end users with high data rate, high split ratio, and high bandwidths. PON provides this to a large number of customers at a low cost. In PON, the Time Division Multiplexing (TDM) method raises the number of end users by using power splitters that split the power of the channel into multiple segments.

P2MP PON based architectures has proven to be the most popular among operators for FTTx deployments because it typically results in lower infrastructure costs. With P2MP, a shared fiber carries traffic to a passive splitting complex close to the end-customers, which can be single stage or cascaded. Consequently, there is no need to deploy a single dedicated fiber all the way to the customer premises from the central office/Headend.

Optical fiber access systems based on PON are currently being deployed on a mass market scale by numerous network operators worldwide. These systems typically exploit Gigabit-class PON systems such as G-PON or EPON as standardized by the ITU-T and IEEE respectively. Both these standards bodies have also defined 10 Gigabit-class PON technologies XGS-PON and 10G-EPON and are widely available today. So, we focus our energy on understanding the latency performance of the XGS-PON and 10G-EPON systems.

The figure below shows the evolution of PON technologies as defined by the IEEE and the ITU-T. both the EPON and GPON standards have been evolving over the last 20 years going from 1 Gbps to now 50 Gbps and beyond.

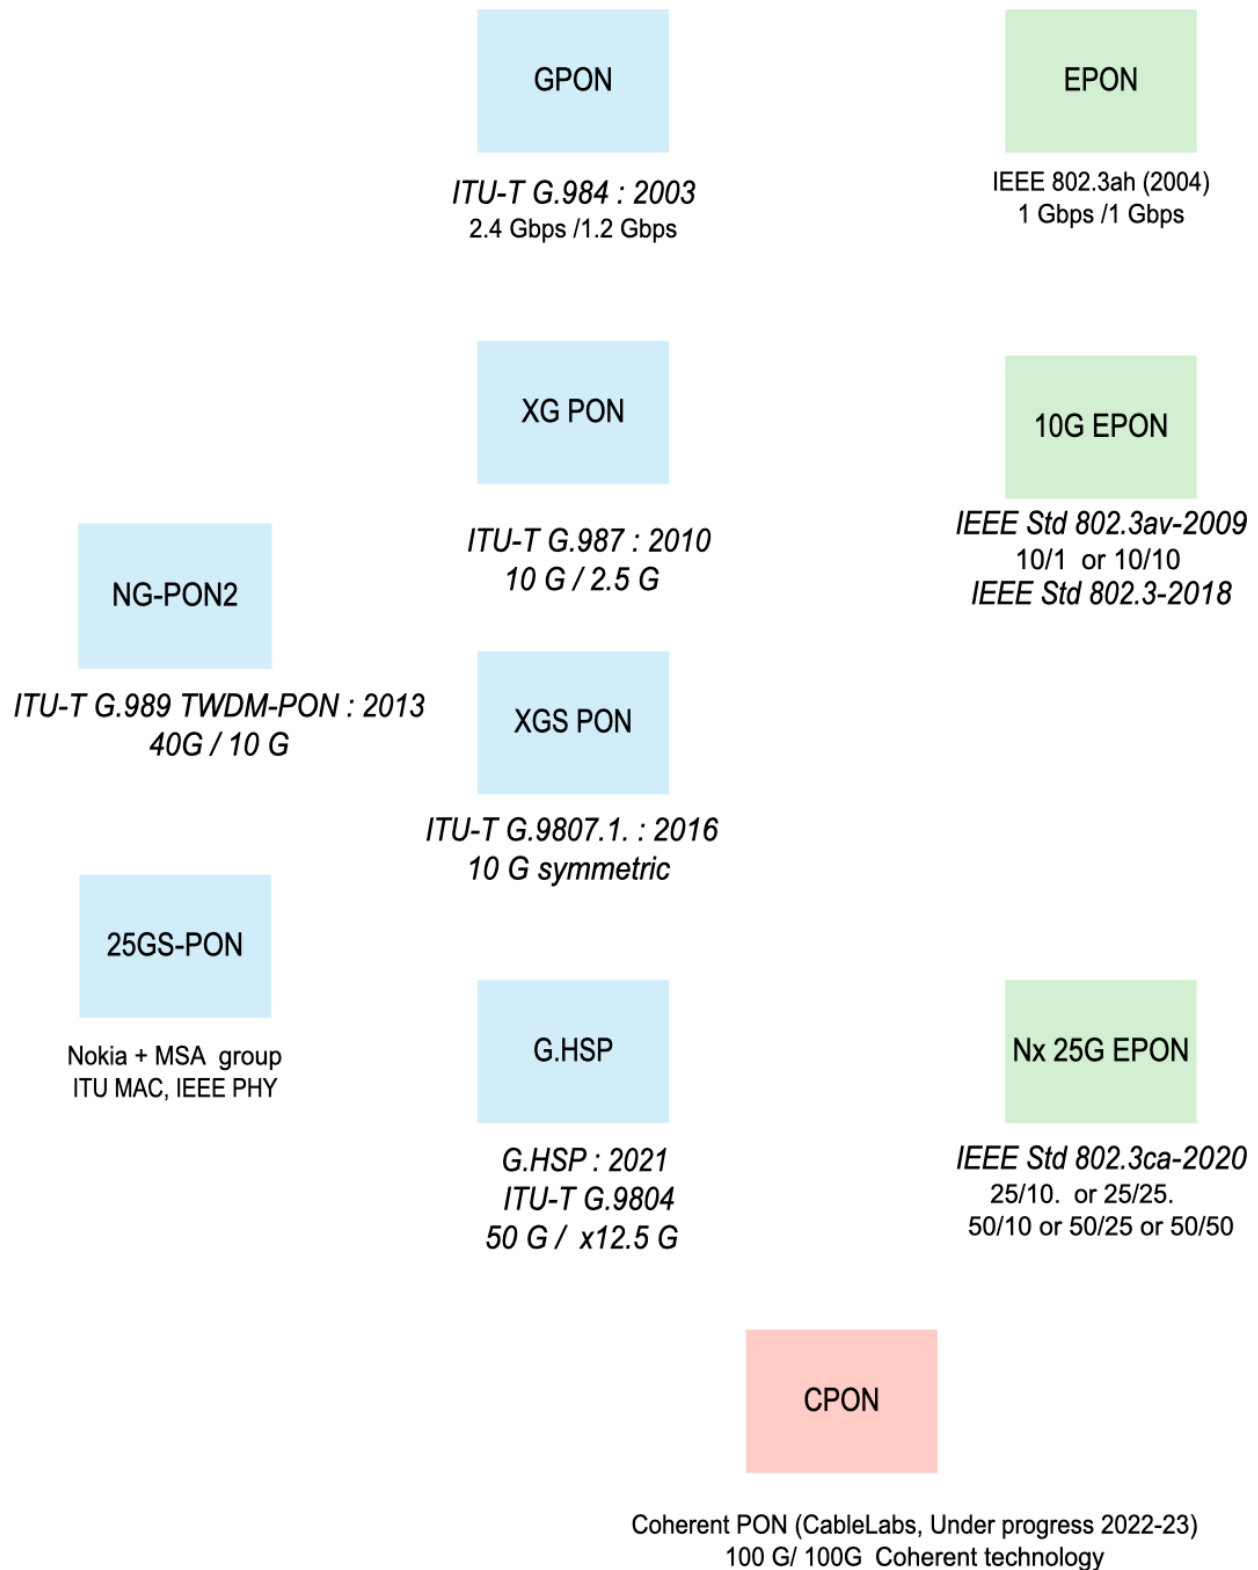


Figure 1 – Evolution of PON Acces technology

Over the last two decades, PON based technology evolved greatly and rapidly, and is one of the most attractive access network solutions for delivering high-speed data and video services. After initial 1 Gbps speeds, PON standards evolved into 10 Gbit/s Ethernet PON and 10-Gigabit-capable PON (XGS-PON). Over the last few decades, the access speed of PON over single wavelength has increased almost 100 times to support the ever-growing bandwidth demand from the emerging services. The 25/50G next-generation Ethernet PON (NG-EPON) specification has been approved by IEEE in 2020, based on 25 Gbps per single wavelength. Most recently, the ITU-T 50G PON standard, based on 50G single-wavelength is close to being a reality as consent has been achieved on multiple part of the series in 2021. Figure 1 above shows a summary of the existing PON standards and the supported downstream/upstream transmission rates.

3. Background of EPON and GPON networks

Optical Access Networks have typically been deployed using point-to-multipoint PON based architectures. In PON networks, Optical networking units (ONUs) communicate with the Optical Line Terminal (OLT) via a passive optical distribution network consisting of a feeder fiber (connecting the OLT and the passive splitter), a passive splitter, and numbers of distribution fibers (connecting each ONU to the passive splitter), as shown in Figure 2.

In the downstream direction (from OLT to an ONU), signals transmitted by the OLT pass through a 1: N passive splitter (or cascade of splitters) and reach each of the ONUs. In the upstream direction (from ONUs to OLT), the signal transmitted by an ONU will only reach the OLT, and not any other ONUs. To avoid data collisions and increase the efficiency of the subscriber access network, ONU's transmissions are arbitrated by the OLT.

This arbitration is achieved by allocating a transmission window (grant) to each ONU. An ONU defers transmission until its grant arrives. When the grant arrives, the ONU transmits frames at wire speed during its assigned time slot. A simplified P2MP topology example is shown in the figure below.

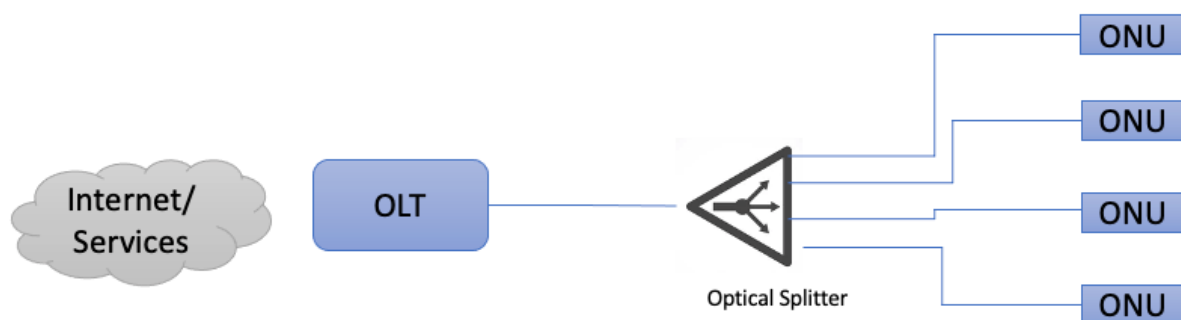


Figure 2 – Simplified EPON Network

The packet delay in a PON network is the time elapsed between a frame's generation at the ONU and its arrival at the OLT. In the most general case, the delay components experienced by a frame include the Delay between frame generation and the transmission of the next REPORT message to the OLT, Propagation of the REPORT message from a given ONU to the OLT, (different for each ONU), processing delay at the OLT between the arrival of the REPORT and the transmission of the GATE message, Propagation of the GATE message to the ONU, delay between the arrival of the GATE message at the ONU and the beginning of the granted transmission window, propagation of the frame to the OLT and the duration of the frame transmission itself.

3.1. EPON Latency background

In EPON, the Multipoint MAC Control Protocol (MPCP) defines the messages and timers, to control access to the P2MP ODN topology in 10G-EPON. Every P2MP ODN topology consists of one Optical Line Terminal (OLT) plus one or more ONUs, as shown in the Figure.

The MPCP and underlying PHY allows an underlying P2MP network to appear as a collection of point-to-point links to the higher protocol layers (at and above the MAC Client). The MPCP achieves this by providing a Logical Link Identification (LLID) to each MAC which is dynamically assigned by the registration process.

The MPCP has the ability to arbitrate the transmitters of a number of ONUs. The OLT controls an ONU's transmission by assigning grants. The transmitting window of an ONU is indicated in gate message where start time and length are specified. An ONU will begin transmission when its local time counter matches start time value indicated in the gate message. An ONU will conclude its transmission with sufficient margin to ensure that the laser is turned off before the grant length interval has elapsed.

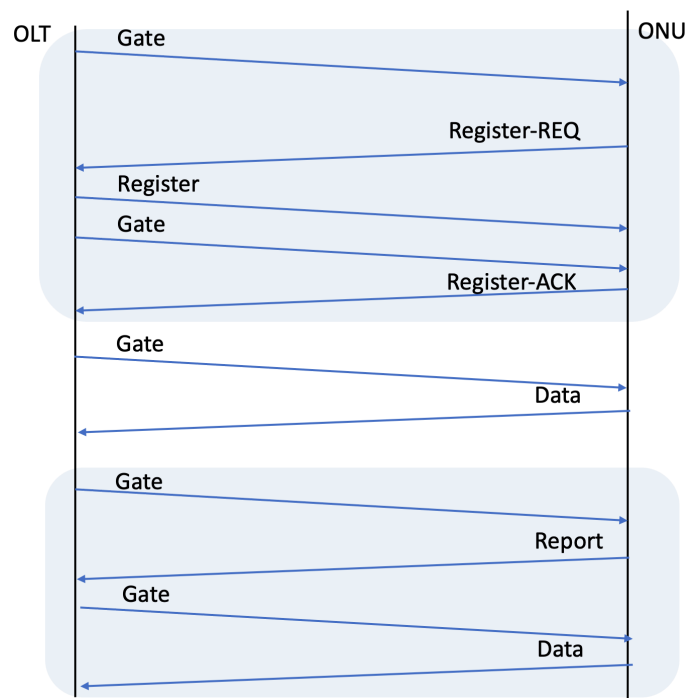


Figure 3 – Simplified Media access in EPON: Gate-Report-data

Multiple outstanding grants may be issued to each ONU. The OLT does not issue more than the maximal outstanding grants as advertised by the ONU during registration. In order to maintain a watchdog timer at the ONU, grants are periodically generated. For this purpose, empty gate messages may be issued periodically.

The purpose of gate message is to grant transmission windows to ONUs for both discovery messages and normal transmission. Up to four grants can be included in a single gate message. The length of the signaled grant, this is a 16-bit unsigned field. The length is counted in 16-bit time increments. The laser-on time, sync time, and laser-off time are included in the grant.

In the [IEEE 802.3] EPON technology, the MPCP protocol relies on strict timing based on distribution of timestamps. The implementations guarantee a constant delay through the MAC and PHY in order to maintain the correctness of the timestamping mechanism. The actual delay is implementation dependent; however, a complying implementation maintains a delay variation of no more than 16-bit times through the MAC stack.

The OLT does not grant less than 1024-time quanta into the future, in order to allow the ONU processing time when it receives a gate message. The ONU shall process all messages in less than this period. The OLT shall not issue more than one message every 1024-time quanta to a single ONU. The unit of time quantum is defined as 16 ns.

A report message has several functionalities. Time stamp in each report message is used for roundtrip time calculation. In the report messages ONUs indicate the upstream bandwidth needs they request per IEEE 802.1Q priority queue. Report messages are also used as keep-alives from ONU to OLT. ONUs issue report messages periodically in order to maintain link health at the OLT. In addition, the OLT can specifically request a report message. Status reports are used to signal bandwidth needs as well as for arming the OLT watchdog timer. Reports shall be generated periodically, even when no request for bandwidth is being made. This keeps a watchdog timer in the OLT from expiring and deregistering the ONU. For proper operation of this mechanism the OLT shall grant the ONU periodically.

EPON implementations typically adopt the traditional Interleaved Polling with Adaptive Cycle Time (IPACT) scheme (with refinements) as the method to allocate bandwidth across multiple ONUs and try to minimize this Gate-Report-data cycle.

3.2. GPON Latency Background

Recommendation ITU-T G.9807.1 describes a 10-Gigabit-capable symmetric passive optical network (XGS-PON) system in an optical access network for residential, business, mobile backhaul and other applications. This system operates over a point-to-multipoint optical access infrastructure at the nominal data rate of 10 Gbit/s both in the downstream and the upstream directions.

A GPON network will be used to carry different types of services, the delay requirements are designed and standardized according to the characteristics of these services.

TDM PON systems use the TDMA mechanism for communicating with multiple ONUs connected to the same OLT PON port. Therefore, Dynamic bandwidth Allocation (DBA) mechanisms and ONU activation mechanisms are needed in the convergence layer, which periodically introduce additional delay to the data transmission in the upstream direction. In real implementations, the data transmission delay of the PON link is considered to fluctuate within about 1.5 ms between service node interface (SNI) and user-network interface (UNI). When the OLT creates a quiet window as part of the ONU activation mechanism, the data transmission delay is significantly increased as a consequence

The XGS-PON OLT support DBA for the efficient sharing of upstream bandwidth among the connected ONUs and the traffic-bearing entities within the individual ONUs based on the dynamic indication of their activity. The dynamic activity indication can be based on the following two methods: (1) status reporting DBA employs the explicit buffer occupancy reports that are solicited by the OLT and submitted by the ONUs in response; (2) traffic monitoring DBA employs OLT's observation of the actual traffic amount in comparison with the allocated upstream transmission opportunities. To guarantee multi-vendor interoperability, the standard specifies the formats of the SR DBA status enquiries and buffer occupancy reports and the associated protocol. An ONU supports DBA status reporting, and transmit upstream reports as instructed by the OLT.




Figure 4 – GPON Upstream Transmission mechanisms

4. Results from Latency Measurements

4.1. Lab experimental Setup

We built a PON plant with a 10GEPON and XGSPON system at CableLabs, as shown in Figure 5. These systems consisted of SFP+ modules of 10GEPON and XGSPON OLTs. These SFP+ OLTs were connected to a 48-port 100Gps switch (Arista 7280R). The fiber cable connection coming out of OLTs were attenuated with a series of optical attenuators and fed into a 5-port optical splitter. An alternate setup was 10 Km of optical fiber to the ONUs. Coming off the splitter we had the two ONUs connected. We had two 10GEPON ONUs and two XGSPON ONUs connected to each of the respective OLTs.

The latency measurement system used here includes a Measurement agent (a STAMP Session sender) and a STAMP session reflector. Additionally, a Controller/Collector was used for visualization of results from the Measurement agent. See [STAMP SCTE21] for further details. These prototype software components are available at [C3 CableLabs] for operators to use.

A hardware traffic generator (ByteBlower with two 10 Gbps optical ports) was used to load the system with traffic. An iPerf server and client were also installed on dedicated servers with 10G NIC cards.

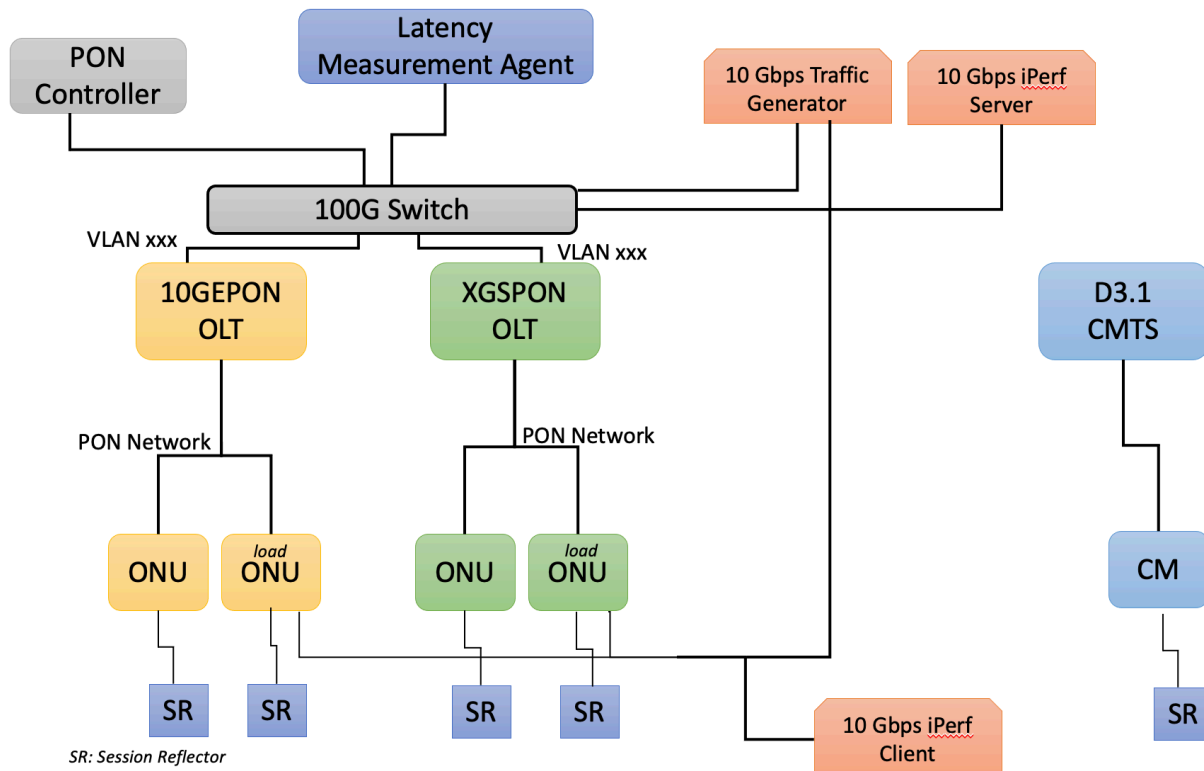


Figure 5 – PON Network Lab Test Setup

The goal of the lab testing was to understand the practical latency on the PON networks by sending different types of load traffic and measuring the latency through the network. Table below shows the default PON configuration on the OLTs

Table 1 – PON Configuration settings

XGSPON		10GE PON	
Discovery Period [ms]:	3000	Discovery Period [ms]:	3000
Encryption:	Bidirectional	Grant Spacing [16ns]:	16
Encryption Key Time [s]:	600	Encryption:	Downstream
Max Frame Size [bytes]:	9600	Encryption Key Time [s]:	900
Downstream FEC:	true	Sync Time [16ns]:	16
Upstream FEC:	true	Laser OFF [16ns]:	32
Guard Time [12.8ns]:	64	Laser ON [16ns]:	32
Preamble Length [12.8ns]:	64	Max Frame Size [bytes]:	9600
Error Det Max Ratio:	20%	FEC:	true
Error Det Min Sample [bursts]:	100	Fiber Reach:	Standard (0..20 km)
Fiber Reach:	Standard (0..20 km)		
Downstream & Upstream SLA Guaranteed: 512 Kbps, rest Best Effort		Downstream and Upstream SLA Guaranteed: 1.28 Mbps, rest is Best Effort	

4.2. Test Scenarios

The main test scenarios and traffic definitions for the tests were as follows.

- Latency Measurement Test Traffic (round trip time - RTT tests)
 - o Raw measurement: 10 pps, 256-byte packets for 2 mins = 20Kbps
 - o Percentile test: 30 pps, 256-byte packets, for 2 mins = 62Kbps
 - o Histogram test : 60 pps, 256-byte packets, for 2 mins = 122Kbps
- UDP Background traffic
 - o Send UDP traffic at different rates as load / background traffic and then measure the round-trip latency of the test traffic
 - o The idea is to send UDP traffic in steps of 1, 3, 5, 7, and up to 9 Gbps (just above the capacity of the two PON networks)
- TCP Background traffic
 - o Send TCP traffic and then measure the round-trip latency of the test traffic
 - o The idea was to have the TCP flows capped at certain rates (1, 3, 5, 7) and also unlimited (capped only by the network capacity ~9 Gbps)

Table 2 – Latency Tests

Scenario name	Load traffic Downstream Data rate	Load traffic Upstream Data Rate
UDP		
S0 Baseline	No load	No load
S1 UDP	1 Gbps UDP	1 Gbps UDP
S3 UDP	3 Gbps UDP	3 Gbps UDP
S5 UDP	5 Gbps UDP	5 Gbps UDP
S7 UDP	7 Gbps UDP	7 Gbps UDP
S9 UDP	9 Gbps UDP	9 Gbps UDP
TCP		
S1 TCP	1 Gbps TCP	1 Gbps TCP
S3 TCP	3 Gbps TCP	3 Gbps TCP
S5 TCP	5 Gbps TCP	5 Gbps TCP
S7 TCP	7 Gbps TCP	7 Gbps TCP
SU TCP	Unlimited TCP	Unlimited TCP

4.3. PON Throughput Test Results

4.3.1. XGSPON Throughput

We ran a baseline throughput test, through the XGSPON system (OLT +ONU), just to understand what the maximum throughput is and what the OLT to ONU one-way latency is. A first port on the traffic generator device is directly connected via the ethernet switch to the OLT and a second port is connected to the ONU as a CPE device. Traffic is being sent in both upstream and downstream directions. The XGSPON system, shows a maximum throughput off 8.564 Gbps in the downstream direction and 8.416 Gbps in the upstream direction. Based on the latency numbers as calculated by the traffic generator device, the average one-way latency is 0.277 milliseconds on the downstream and 0.708 milliseconds in the upstream. The 10GEPON configuration was in the default state as detailed in Table-1.

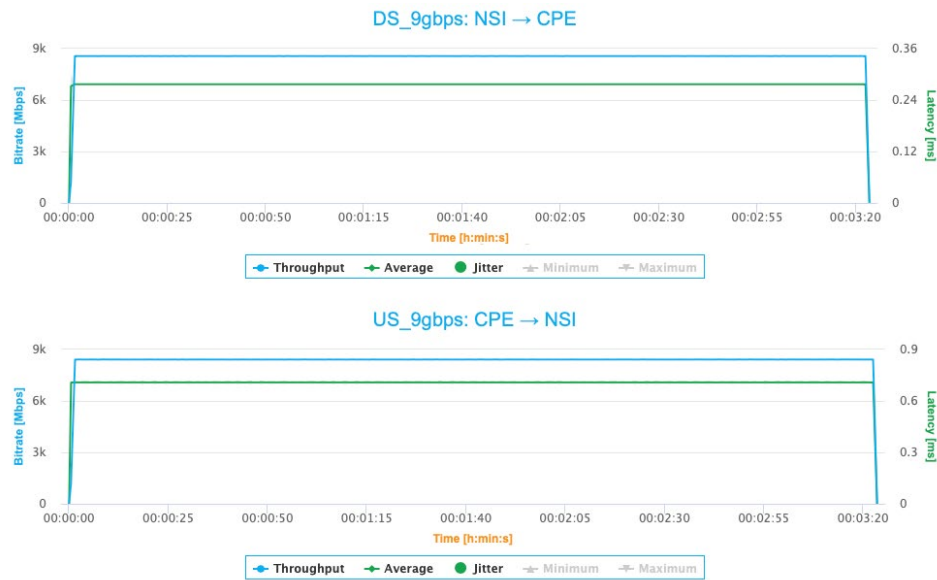


Figure 6 – XGSPON Throughput & one-way Latency(UDP)

Frame Blasting Flows : Throughput

Flow	Source	Destination	TX Frames	Rx Frames	Frame Loss	TX Bytes (+VLAN)	RX Bytes (+VLAN)	Byte Loss	Duration	Average Throughput (Mbps)
DS_9gbps	NSI	CPE	151,822,917	144,442,983	4.86%	227,734,375,500 (+607,291,668)	216,664,474,500	4.86%	3m, 22s 379ms 305µs 757ns	8,564.69
US_9gbps	CPE	NSI	151,823,300	141,561,578	6.76%	227,734,950,000	212,342,367,000 (+566,246,312)	6.76%	3m, 22s 378ms 708µs 613ns	8,416.25

Frame Blasting Flows : Latency

Flow	Source	Destination	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)	Jitter (ms)
DS_9gbps	NSI	CPE	0.024	0.277	0.281	0.001
US_9gbps	CPE	NSI	0.532	0.708	1.091	0.007

Figure 7 – XGSPON Throughput & one-way Latency in Numbers (UDP)

4.3.1. 10GEPON Throughput

We ran a similar baseline throughput test, through the 10GEPON system (OLT +ONU), again to understand what the maximum throughput and the OLT to ONU one-way latency is. The traffic generator device has the same configuration as the above test. The 10GEPON system, shows a maximum throughput of 8.596 Gbps in the downstream direction and 8.436 Gbps in the upstream direction. Based on the latency numbers as calculated by the traffic generator device, the average one-way latency is 0.636 ms on the downstream and 1.739 ms in the upstream. The XGSPON configuration was in the default configuration as state in Table-1.

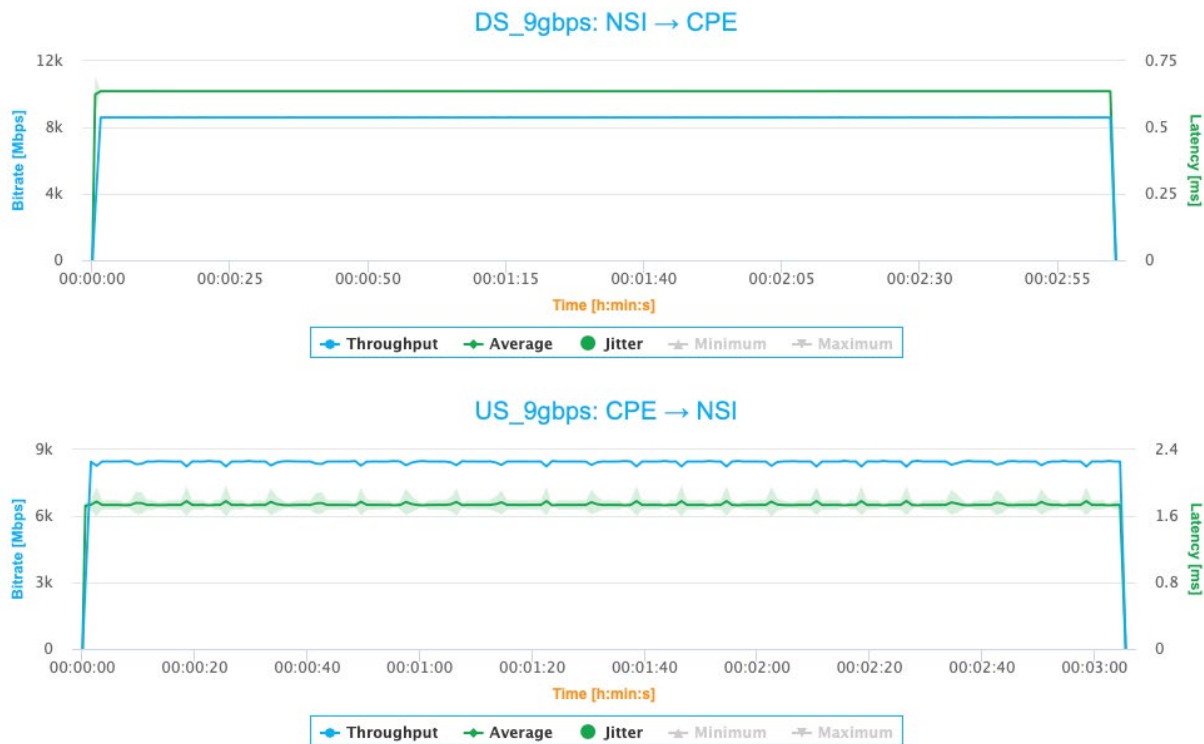


Figure 8 – 10GEPON Throughput & one-way Latency(UDP)

Frame Blasting Flows : Throughput

Flow	Source	Destination	Tx Frames	Rx Frames	Frame Loss	Tx Bytes (+VLAN)	Rx Bytes (+VLAN)	Byte Loss	Duration	Average Throughput (Mbps)
DS_9gbps	NSI	CPE	138,613,303	132,370,647	4.50%	207,919,954,500 (+554,453,212)	198,555,970,500	4.50%	3m, 4s 771ms 334µs 308ns	8,596.83
US_9gbps	CPE	NSI	138,613,615	129,561,132	6.53%	207,920,422,500	194,341,698,000 (+518,244,528)	6.53%	3m, 4s 769ms 526µs 608ns	8,436.89

Frame Blasting Flows : Latency

Flow	Source	Destination	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)	Jitter (ms)
DS_9gbps	NSI	CPE	0.024	0.636	0.641	0.003
US_9gbps	CPE	NSI	1.336	1.739	3.604	0.080

Figure 9 – 10GEPON Throughput & one-way Latency in Numbers (UDP)

4.4. Baseline Latency (RTT) Test Results

The next set of tests that we ran is a set of standardized latency measurement tests as defined in section 5.1. Here the basic idea is to have a measurement agent (on the NSI side of the OLT), which sends STAMP (UDP) packets to a session-reflector device (located behind an ONU as the CPE). The session-reflector essentially sends the packet back to the measurement agent who then computes the roundtrip time. The sequence of 1200 packet latency measurements will be used to show a time series of the latency, a histogram of the latency and a latency CCDF as defined in [SCTE21 Latency Measurement]. The figures below also show any packet drops during the test. We also ran a couple of additional tests at the same time, which compute only the histogram of latency and the second one computes only the percentile values though we are not sharing the results from those tests, as for now the raw latency test

has all the information we need. These baseline tests are run without any load on the PON network, i.e., there is no background load traffic

4.4.1. Baseline XGSPON Latency

As shown in the figure below, we see the latency of the XGSPON system to be between 1 to 2.1 ms, with a few outliers going all the way from 2.5 to 3.5 ms.

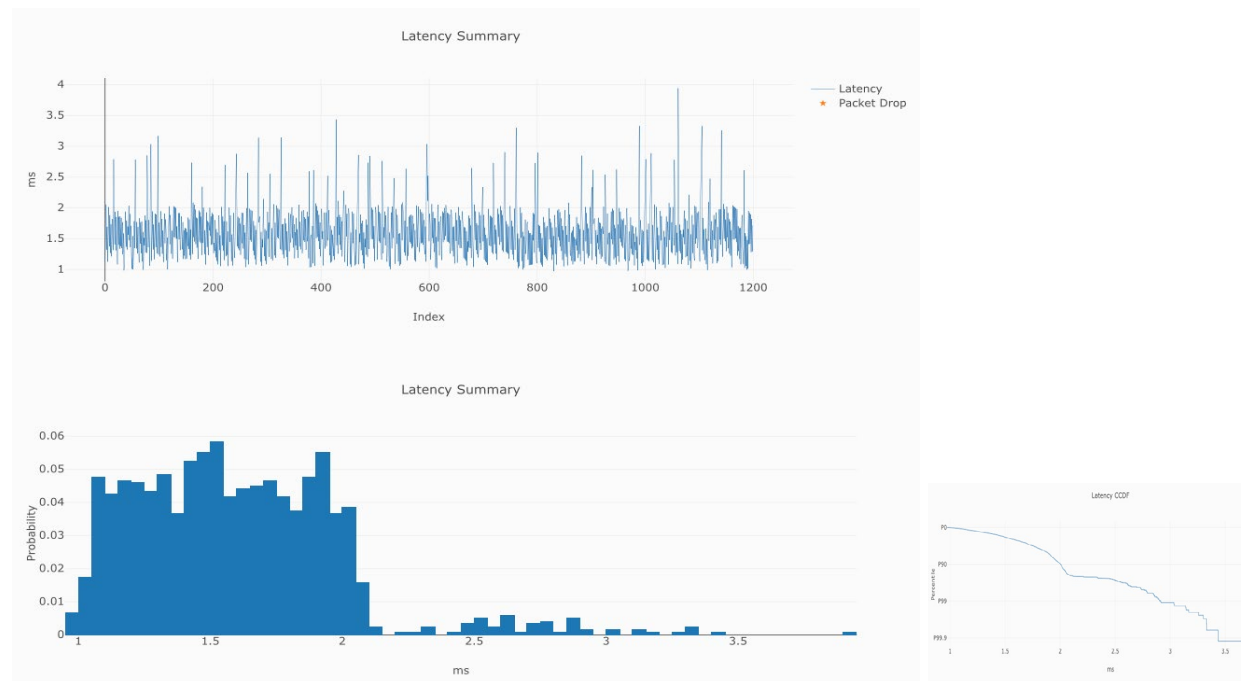


Figure 10 – Baseline XGSPON latency (no load)

4.4.2. Baseline 10GEPON Latency

As seen in the figure below on a 10GEPON system, we see the latency values between 0.8 and 1.9 milliseconds




Figure 11 – Baseline 10GEPON latency (no load)

4.4.3. Baseline DOCSIS Latency

We ran the same baseline latency test on a DOCSIS CMTS (no Low Latency DOCSIS (LLD) configured), to observe a baseline latency between 6.2-9 ms, with outliers in the 9-11 ms range.




Figure 12 – Baseline DOCSIS (with No LLD enabled) and no load

4.4.4. Baseline DOCSIS (with LLD) Latency

We ran the same baseline latency test on a DOCSIS CMTS which is configured with the new Low Latency DOCSIS features (LLD) and observe a baseline latency between **2 and 3 ms**, with some outliers in the 3-4 ms range. This was just out-of-the-box configuration with LLD enabled on both the upstream and downstream. The latency measurement traffic in this test was marked with an IP DSCP Marking of 46 (expedited forwarding). With this we can see that enabling LLD features on a DOCSIS 3.1 CMTS definitely puts the latency in the same level as the EPON and GPON systems




Figure 13 – Baseline DOCSIS with LLD latency and no load

4.5. UDP Load Test Results

Now we have a set of UDP tests where we run different levels of UDP traffic in the background. The load traffic is sent through one ONU, while the latency RTT measurements are being done on both the reflectors that are behind the load ONU, as well as the second ONU which doesn't have any load traffic passing through it.

4.5.1. GPON UDP 1Gbps Test Results

When we run background UDP load traffic of 1 Gbps, we see that for the ONU that is not carrying the load, the latency distribution does not change much and the histogram and the CCDF curves look similar to the baseline case.

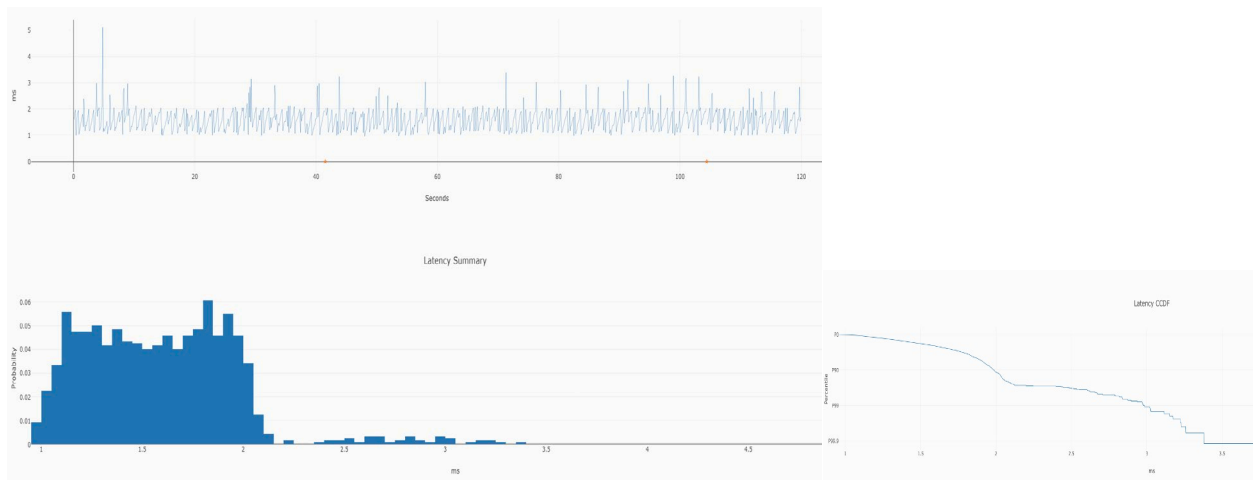


Figure 14 – XGSPON latency – with 1 Gbps UDP background traffic, non-load ONU

When we run background UDP load traffic of 1 Gbps, we see that for the ONU that is carrying the load, the latency distribution does change and the histogram and the CCDF curves actually get better. From the

percentiles test we can tell that the latency for the P50 and P99 numbers actually improved. The 50th percentile latency improved from 1.557 to 1.278 ms and the 99th percentile latency from 2.78 ms to 1.534 ms. With the increased load on the PON network we can see that the polling-request-grant delay cycle is reduced, with a ONU buffer status report being a part of the upstream data frame header itself.

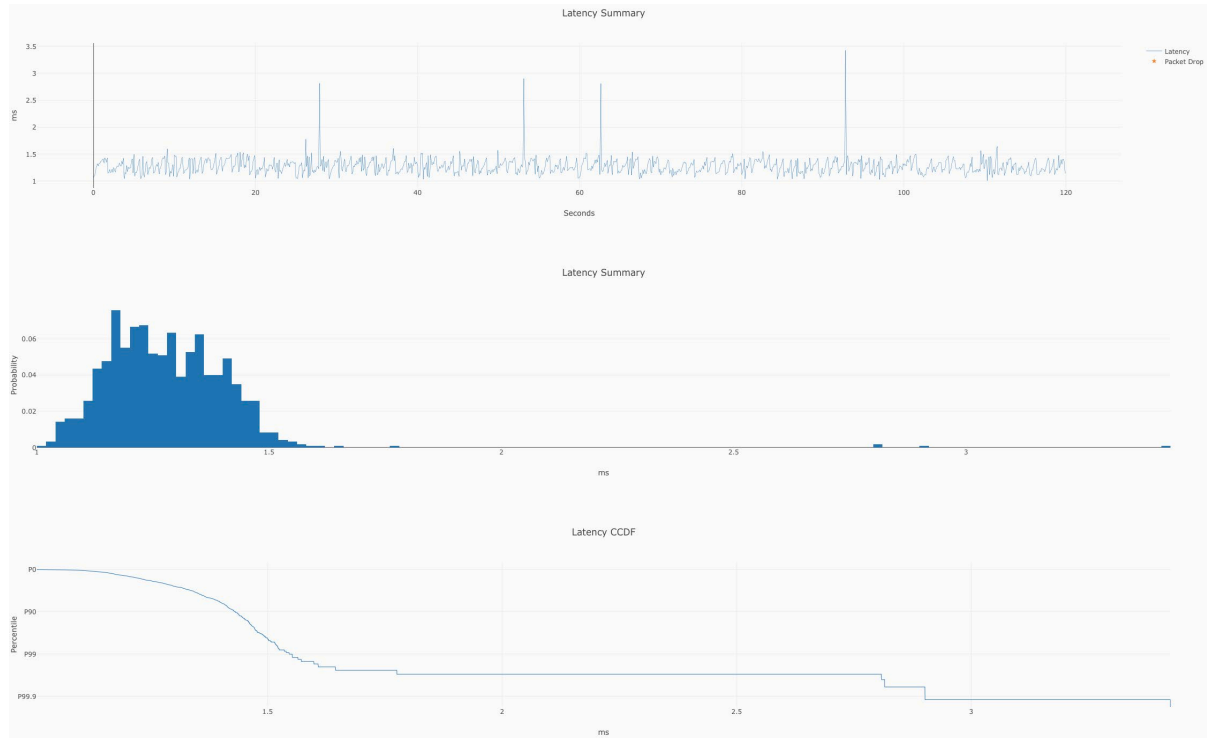


Figure 15 – XGSPON latency – w 1 Gbps background traffic, Load ONU

4.5.2. GPON UDP 9Gbps Test Results

When we run background UDP load traffic of 9 Gbps (above the capacity of the network), we see that for the ONU that is not carrying the load, the latency distribution does not change much and the histogram and the CCDF curves look similar to the baseline case.




Figure 16 – XGSPON latency – with 9 Gbps background traffic on non-load ONU

When we run background UDP load traffic of 9 Gbps, we see that for the ONU that is carrying the load, the latency distribution does change much and the histogram and the CCDF curves get worse as expected. From the percentiles test we can tell that the latency for the 50th percentile latency (P50) increased from 1.557 to 2.423, and the 99th percentile latency changes from 2.78 ms to 3.482 ms. With the fully loaded PON network we can see that there is packet loss and the increased latency of the packets that did make through. This also points to a relatively short buffer on the ONUs and OLTs.




Figure 17 – XGSPON latency – with 9 Gbps background traffic on load ONU

4.5.3. EPON UDP 1Gbps Test Results

When we run background UDP load traffic of 1 Gbps, we see that for the ONU that is carrying the load, the latency distribution shifts to the left (decreases) by 0.4 ms. For the ONU that is not carrying the load the latency remains the same as the baseline case.




Figure 18 – 10GEPON latency – no load.




Figure 19 – 10GEPON latency – w 1 Gbps UDP background traffic- load ONU.

4.5.4. EPON UDP 9Gbps Test Results

When we run background UDP load traffic of 9 Gbps (above the capacity of the network), we see that for the ONU that is not carrying the load, the latency distribution does not change much and the histogram and the CCDF curves look similar to the baseline case. This behavior for this 10GEPON system, is similar to what we saw with the XGSPON system.




Figure 20 – 10GEPON latency – w 9 Gbps background traffic - non-load ONU

When we run background UDP load traffic of 9 Gbps, we see that for the EPON ONU that is carrying the load, the latency distribution does change and the histogram and the CCDF curves get worse as expected. From the percentiles test we can tell that the latency for the 50th percentile latency (P50) increased from 2.132 to 2.423 ms as expected, though the 99th percentile latency improved from 2.769 ms to 2.548 ms. With the fully loaded PON network we can see that there is some packet loss and the increased latency of the packets that did make through. The packet loss on these test flows is much less than what we saw on the XGSPON network, this implies that the EPON implementations have a per-flow queuing mechanism. This also points to a relatively short buffer on the EPON ONUs and OLTs.




Figure 21 – 10GEPON latency – w 9 Gbps background traffic - load ONU

4.5.5. DOCSIS UDP Test Results

For this test we used a high-split CMTS, which has a downstream capacity of 4.8 Gbps downstream and 1.7 Gbps upstream. When we run background UDP load traffic of 4.7 Gbps downstream and 1.65 Gbps (just fully loading the DOCSIS link capacity configured for this setup), we see that for the CM, the latency distribution does change and the histogram and the CCDF curves shift to the right as expected. From the percentiles test we can tell that the latency for the 50th percentile latency (P50) increased from 7.2 to 15.5 ms, and the 99th percentile (P99) latency changes from 11.27 to 20.4ms. This test essentially shows the power of the AQM feature which all operators must ensure is turned on in their network.

As seen in the figure below, we can see that the baseline latency with fully loaded network is at the 7-8 ms mark. When we start the UDP load at about 16 seconds into the test we see that the latency spikes up to 30 milliseconds and then once the AQM algorithm kicks in, the latency stabilizes at about 15 ms.




Figure 22 – D3.1 (with AQM) latency – (4.7 and 1.7 Gbps UDP Load)

In the graph below we have an experiment with Low Latency DOCSIS (LLD) is turned on and the latency measurements are between 2 to 5 milliseconds even when we are sending traffic at 10% above the configured data rate. (In this configuration, service flows are limited to 200 Mbps down/20 Mbps up).




Figure 23 – D3.1 with LLD, UDP test at 10% above service flow limit

The test itself runs traffic at 10% above the 200 Mbps downstream rate limit and 20 Mbps upstream rate limit.

4.6. TCP Load Test Results

The next set of tests includes TCP traffic as load or background traffic either on the ONU under test or on a neighboring ONU on the same PON.

4.6.1. XGSPON TCP Test results

The next three figures show three different TCP Load test scenarios. In the first case we rate limit the TCP flows to be 1 Gbps, in the second we rate limit the TCP flows to 9 Gbps and in the third we run TCP flows using iPerf which are limited only by the network capacity. In all three cases we see that the XGSPON system handles bidirectional TCP flows without a significant change to the latency numbers.




Figure 24 – XGSPON latency – w TCP(1 Gbps, byteblower) - non-load ONU




Figure 25 – XGSPON latency – w TCP(9 Gbps, byteblower) - load ONU

Again, as before when the latency measurement is being done on the ONU which is carrying the load traffic, we see that the latency numbers are actually lower, compared to the ONU which is not carrying any load traffic.




Figure 26 – XGSPON latency – with Iperf TCP on non-load ONU

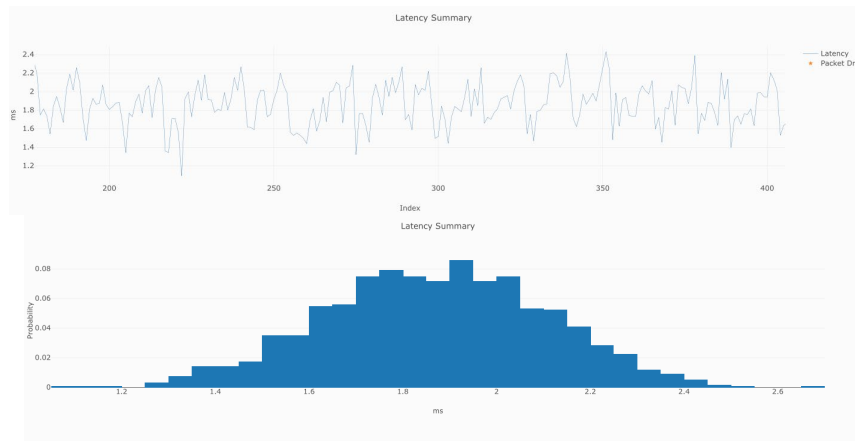


Figure 27 – XGSPON latency – with Iperf TCP on Load ONU

4.6.2. 10GEPON TCP Test results

Similarly, for the 10GEPON system, we see that the system handles bidirectional TCP flows without a significant change to the latency numbers.




Figure 28 – 10GEPON latency – with Iperf TCP on non-load ONU

When the latency measurement is being done on the EPON ONU which is carrying the load traffic we see that the latency numbers get worse, compared to the ONU which is not carrying any load traffic. This behavior is similar to the results we observed for the UDP tests on the EPON network




Figure 29 – 10GEPON latency – with Iperf TCP on load ONU

4.6.3. DOCSIS TCP Test results

We ran a few of the TCP load set of test cases on a DOCSIS CMTS. Now due to various networking issues and corporate VLAN switching issues to connect to the DOCSIS CMTS in a different lab the number of test cases were limited and are left for a future testing effort.

In an older DOCSIS system configuration (without AQM configured or AQM Latency target configured to a high value) we see the latency with a TCP load test increases dramatically compared to the baseline DOCSIS latency seen in the previous sections. This is **no longer an acceptable configuration** and operators should ensure that AQM is enabled on all of their CMTS /CM devices.




Figure 30 – DOCSIS no AQM config– with 1 Gbps TCP

In a D3.1 system correctly configured with AQM, (See figure below), we see the latency with a TCP load test does increase (the histogram shifts to the right by a millisecond and also extends a bit more) to 7-12 ms compared to the baseline DOCSIS latency (6-8 ms) seen in the section: Baseline DOCSIS Latency.




Figure 31 – DOCSIS AQM Enabled Config, no LLD and Iperf TCP load

In a D3.1 system configured with LLD, we see the latency with a TCP load test does increase, the histogram shifts to the right and is a bit more spread (1.9-5ms) compared to the baseline DOCSIS LLD latency (1.8-3 ms) seen in the previous section Baseline DOCSIS (with LLD) Latency.

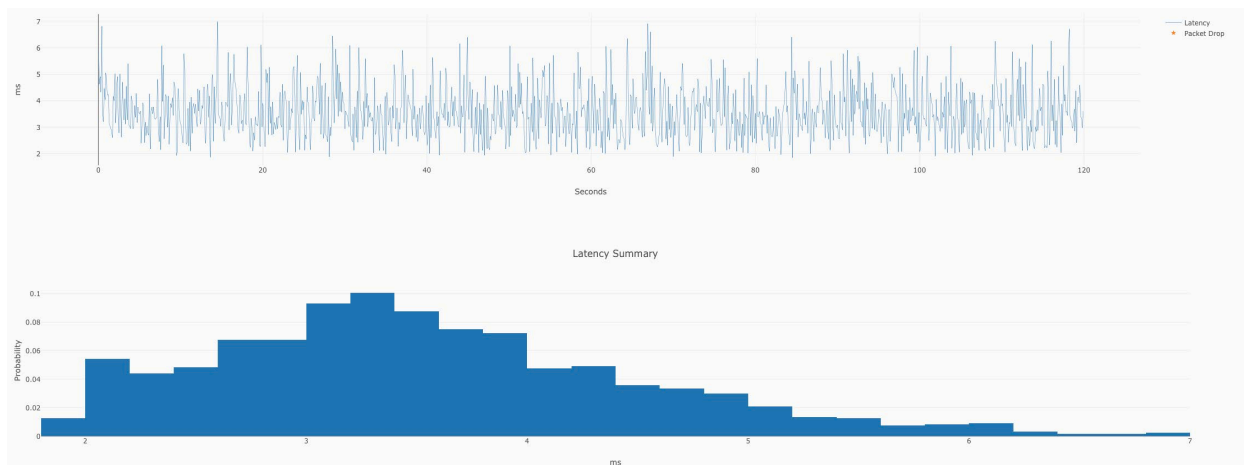


Figure 32 – DOCSIS LLD Enabled, Iperf TCP load

4.7. GPON/ EPON /DOCSIS Test Results summary

Looking at each of these above graphs and CDF's is very enlightening. The below table tries to summarize each of the tests that we ran by documenting three percentile values (0th, 50th, 99th), designated as P0 (minimum), P50 and P99. Please see the next three figures for a visual representation of the three tables. Please note that the DOCSIS testing was done on 3 different CMTS's/CMs, and as such these results from one DOCSIS CMTS cannot be compared with another due to change in the CMTS/CM hardware/Software versions and different plant configurations.

Table 3 – GPON latency Test results , P0,P50,P99

Scenario name	Reflector 4 (load ONU)	Reflector 3
		All units are ms
S0 Baseline	0.960, 1.557, 2.78	0.960, 1.550, 2.91
S1 UDP	1.019, 1.278, 1.534	0.948, 1.558, 2.93
S3 UDP	1.000, 1.177, 1.417	0.942, 1.550, 2.92
S5 UDP	1.000, 1.140, 1.340	0.961, 1.580, 2.99
S7 UDP	0.998, 1.150, 1.340	0.941, 1.580, 3.03
S9 UDP	2.840, 2.950, 3.482	1.670, 2.330, 4.023
S1 TCP	0.920, 1.280, 1.54	0.920, 1.569, 2.97
S3 TCP	0.961, 1.150, 1.36	0.945, 1.550, 2.90
S5 TCP	0.991, 1.140, 1.32	0.956, 1.580, 2.89
S7 TCP	1.005, 1.167, 1.35	0.960, 1.580, 2.98
S9 TCP	1.027, 1.227, 1.45	1.007, 1.640, 3.05
SU TCP (iPerf)	1.043, 1.667, 2.137	1.089, 1.900, 3.289

Table 4 – EPON latency Test results , P0,P50,P99

Scenario name	Reflector 2 (load ONU)	Reflector 1
		All units are ms
S0 Baseline	0.569, 2.132, 2.769	0.792, 1.384, 1.900
S1 UDP	0.547, 1.906, 2.125	0.673, 1.397, 1.932
S3 UDP	0.545, 1.925, 2.115	0.641, 1.419, 1.962
S5 UDP	0.522, 1.859, 2.041	0.656, 1.421, 1.957
S7 UDP	0.505, 1.824, 1.975	0.690, 1.400, 1.970
S9 UDP	1.118, 2.423, 2.548	0.990, 1.770, 2.341
		All units are ms
S1 TCP	0.549, 1.901, 2.126	0.633, 1.393, 1.927
S3 TCP	0.559, 1.899, 2.095	0.655, 1.389, 1.951
S5 TCP	0.421, 1.959, 2.724	0.625, 1.390, 1.947
S7 TCP	0.471, 1.950, 2.717	0.708, 1.393, 1.946
SU TCP (iPerf)	0.545, 2.122, 2.807	0.679, 1.403, 1.971

Table 5 – DOCSIS latency Test results , P0,P50,P99

Scenario	Reflector 5 (Different CMTS/CM ***)
S0 Baseline	2.80, 7.28, 11.27
S0 Baseline w LLD	1.76, 2.47, 3.834
UDP: No AQM **	89.3, 91.1, 91.97
UDP: with AQM	6.28, 15.5, 20.43
UDP with LLD	1.84, 3.42, 6.14
TCP: With AQM	3.11, 7.57, 10.64
TCP : With LLD*	1.94, 8.37, 12.625

** Prototype LLDCMTS software, not mature implementation. ** Recommend MSOs to always enable AQM. This test was just for historical purposes. *** 3 different CMTS/CM, so results cannot be compared.*




Figure 33 – GPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU




Figure 34 – EPON Latency Summary (P0,P50,P90) for UDP/TCP , load vs unloaded ONU

4.8. Multiple LLIDs Test Results (EPON)

The next set of tests that we wanted to do on the 10GEPON system was the use of multiple LLIDs. We created some DPOE (DOCSIS Provisioning of EPON) configuration, which essentially classifies packets based on the incoming VLAN tags and forwards each of those packets on to a different LLID on the PON system. These LLIDs were configured on both the upstream and downstream direction.

The load test we ran was sending five different flows, 2Gbps each, in both the upstream and downstream directions which matched those VLANs. This essentially meant we were pushing the limits of the EPON system by sending 10 Gbps in each direction. Given that the capacity of the PON network was 8.5 Gigabit per second on the downstream and 8.4 Gigabit per second on the upstream we expect to see that the scheduler on the OLT equally shares the bandwidth between the five LLIDs, in both the upstream and downstream direction. In the test results we see that the average throughput on the downstream is about 1.712 Gbps for each of the 5 downstream flows. On the upstream we notice the first LLID almost has no drops with a rate of 1.968 Gbps and the remaining 4 flows share the remaining bandwidth with 1.585 Gbps each.

Flow	Source	Destination	Tx Frames	Rx Frames	Frame Loss	Tx Bytes (+VLAN)	Rx Bytes (+VLAN)	Byte Loss	Duration	Average Throughput (Mbps)
DS_DSCP_80_f7a0	NSI_4001	CPE_4001	20,000,000	17,398,331	13.01%	30,000,000,000 (+80,000,000)	26,097,496,500 (+69,593,324)	13.01%	2m, 2s, 243ms, 90µs, 737ns	1,712.46
DS_DSCP_112_f7a0	NSI_4002	CPE_4002	20,000,000	17,398,076	13.01%	30,000,000,000 (+80,000,000)	26,097,114,000 (+69,592,304)	13.01%	2m, 2s, 252ms, 591µs, 819ns	1,712.30
DS_0_f7a0	NSI_4000	CPE_4000	20,000,000	17,397,410	13.01%	30,000,000,000 (+80,000,000)	26,096,115,000 (+69,589,640)	13.01%	2m, 2s, 252ms, 591µs, 819ns	1,712.24
US_0_f7a0	CPE_4000	NSI_4000	20,000,000	19,996,993	0.02%	30,000,000,000 (+80,000,000)	29,995,489,500 (+79,987,972)	0.02%	2m, 2s, 251ms, 376µs, 496ns	1,968.11
UP_DSCP_80_f7a0	CPE_4001	NSI_4001	20,000,000	16,107,957	19.46%	30,000,000,000 (+80,000,000)	24,161,935,500 (+64,431,828)	19.46%	2m, 2s, 246ms, 146µs, 871ns	1,585.42
UP_DSCP_112_f7a0	CPE_4002	NSI_4002	20,000,000	16,106,642	19.47%	30,000,000,000 (+80,000,000)	24,159,963,000 (+64,426,568)	19.47%	2m, 2s, 248ms, 709µs, 251ns	1,585.25
DS_DSCP_160_f7a0	NSI_4003	CPE_4003	20,000,000	17,397,824	13.01%	30,000,000,000 (+80,000,000)	26,096,736,000 (+69,591,296)	13.01%	2m, 2s, 252ms, 591µs, 819ns	1,712.28
UP_DSCP_160_f7a0	CPE_4003	NSI_4003	20,000,000	16,106,557	19.47%	30,000,000,000 (+80,000,000)	24,159,835,500 (+64,426,228)	19.47%	2m, 2s, 244ms, 13µs, 4ns	1,585.31
DS_DSCP_104_f7a0	NSI_4004	CPE_4004	20,000,000	17,397,827	13.01%	30,000,000,000 (+80,000,000)	26,096,740,500 (+69,591,308)	13.01%	2m, 2s, 237ms, 765µs, 881ns	1,712.49
UP_DSCP_104_f7a0	CPE_4004	NSI_4004	20,000,000	16,105,553	19.47%	30,000,000,000 (+80,000,000)	24,158,329,500 (+64,422,212)	19.47%	2m, 2s, 250ms, 278µs, 316ns	1,585.13

Figure 35 –10GEPON throughput (5 LLIDs 2Gbps per LLID)

Using the measurements done by the traffic generator we observe that the average one-way latency on the downstream for each of the LLIDs is 0.291 ms, while in the upstream direction the one-way latency is 0.781 ms with the first Upstream LLID with a latency of 0.830ms.

Flow	Source	Destination	Min Latency (ms)	Avg Latency (ms)	Max Latency (ms)	Jitter (ms)
DS_DSCP_80_f7a0	NSI_4001	CPE_4001	0.026	0.291	0.307	0.005
DS_DSCP_112_f7a0	NSI_4002	CPE_4002	0.025	0.291	0.307	0.005
DS_0_f7a0	NSI_4000	CPE_4000	0.026	0.291	0.307	0.005
US_0_f7a0	CPE_4000	NSI_4000	0.328	0.830	1.814	0.153
UP_DSCP_80_f7a0	CPE_4001	NSI_4001	0.315	0.781	1.850	0.153
UP_DSCP_112_f7a0	CPE_4002	NSI_4002	0.320	0.781	1.789	0.153
DS_DSCP_160_f7a0	NSI_4003	CPE_4003	0.027	0.291	0.307	0.005
UP_DSCP_160_f7a0	CPE_4003	NSI_4003	0.307	0.781	1.837	0.153
DS_DSCP_104_f7a0	NSI_4004	CPE_4004	0.025	0.291	0.307	0.005
UP_DSCP_104_f7a0	CPE_4004	NSI_4004	0.326	0.781	1.793	0.153

Figure 36 –10GEPON Latency (5 LLIDs 2Gbps per LLID)




Figure 37 –10GEPON visual of Throughput & Latency (for one US and DS LLID)

4.9. Challenges and Further testing

This was the first time we were setting up an 10GEPON and XGSPON in our labs. So as expected there was a very steep learning curve in all aspects of the setup configuration and operation of the PON network. Once we got past setting up the physical network using the needed attenuators (10 to 15 dB), the main challenge was related to VLAN configuration. These SFP+ OLTs were connected to a switch which needed careful configuration of VLANs, to make sure we were bridging the traffic to the devices it needed to go. We also needed to build up the lab capabilities with optical to copper 10 Gbps port converters, 10 Gbps NIC cards on servers and clients etc. Learning the provisioning of services on these GPON and EPON systems and getting the right configurations in was another step in the learning process. Also getting packets to flow through the corporate network and reach the DOCSIS CMTS in a different lab had its own set of challenges and limited the testing to some extent.

In terms of future testing, we would like to study in more detail the performance of TCP flows on these PON systems. For an XGSPON system, we would like to enable and test multiple flows (AllocIDs in XGSPON) in each ONU. We would like to do more detailed testing of the DOCSIS system with LLD configured and varying the different configuration parameters. We would also like to measure the performance of TCP flows and the latency, when the measurement agents are located much farther in the network (i.e., adding additional round trip time or delay in the network path). We would also like to enable testing of PON systems with a larger number of ONUs (more than the two that we had currently on each system). We would also like to mix and match different on ONUs with different OLTs to understand the performance. There are many scheduling parameters which can be tweaked on the OLT scheduler, this includes parameters like polling interval, minimum and maximum granting period etc., and we would like to understand the differences in performance by varying each of these parameters.

5. Conclusions

10GEPON and XGSPON networks show remarkably good performance with 8.6/8.5 Gbps symmetrical throughput. Both PON systems show consistent latency performance on an unloaded or loaded PON network. The majority of the latency numbers fall within 1 to 2 milliseconds on a lightly loaded PON network, while we observe latencies up to 2 to 3.5 milliseconds under heavily loaded conditions. (Note that this loaded condition is for the home network i.e. this ONU's load, whereas when the whole serving group is heavily loaded and each ONU is transmitting above the minimum guaranteed rate, the latencies will likely increase.) At least in the configurations we tested it looks like 10GEPON has slightly lower latencies than the XGSPON setup we tested, but given they are all in the 1-3 ms range they both are very similar latency ranges. (The 10GEPON did have more guaranteed bandwidth than the XGSPON setup). We tested the network using UDP load traffic and TCP load traffic and the latency numbers stayed relatively consistent and in the same range. We also compared the latency in a DOCSIS network and see that while without Low Latency DOCSIS(LLD) features, the latency in a DOCSIS network can vary from 6-7 ms (baseline) to ~15 to 30 ms (under load). Enabling LLD brings the latency numbers to be at comparable levels, 2-3 ms (baseline) to 2-6 ms (under load), to that of a PON network. As latencies drop to the sub-5ms levels, the connection to content servers (outside of access network) through the Internet can become bigger contributors to overall end-to-end latency.

All testing was done with limited variations from default out-of-the-box configurations, varying the configuration parameters(for each of EPON, GPON, DOCSIS) may provide more variation in the results, leading to further analysis and understanding.

Abbreviations

10GEPON	IEEE 10 Gigabit EPON (Ethernet Passive Optical Network) technology
XGSPON	ITU-T 10 Gigabit Symmetric PON (Passive Optical Network) technology
RTT	round trip time
CDF	cumulative distribution function
DSCP	Diff Serv Code Point
ms	millisecond
LLD	Low Latency DOCSIS
STAMP	Simple Two-Way Active Measurement Protocol
Gbps	Gigabits per second
OLT	Optical Line Terminal
ONU	Optical Networking Unit
CMTS	Cable Modem Termination System
CM	Cable Modem
Pps	Packets per second
NSI	Network Side interface (i.e., Interface north of the CMTS or OLT)
CPE	Customer Premise Equipment (i.e., device behind CM or ONU

Bibliography & References

[ITU-T G.9807.1] 10-Gigabit-capable symmetric passive optical network (XGS-PON).

[IEEE 802.3] IEEE Std 802.3™-2018 IEEE Standard for Ethernet, Section 5 (10GEPON)

[LM SCTE 20] Latency Measurement: What is latency and how do we measure it? Karthik Sundaresan, Greg White, Steve Glennon, SCTE 2020

[STAMP SCTE21] A Latency Measurement System Using STAMP, Karthik Sundaresan, SCTE 2021

[C3 CableLabs] CableLabs Common code community, <https://community.cablelabs.com/wiki/display/C3>

Unified Optical Architecture to Support Wavelength and High Bandwidth Ethernet Services

A Technical Paper prepared for SCTE by

Stephen Ruppa
Sr Principal Engineer
Comcast
2001 York Rd, Oak Brook, IL 60523
847-489-7953
Stephen_ruppa@comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Network evolution.....	3
2.1. Many small CATV providers.....	3
2.2. Acquisitions and HFC	4
2.3. Networks merge, advanced services, cell backhaul and ROADMS	5
2.4. All ROADMs, Flex Technology and one vendor	5
3. Products	6
3.1. Wavelength Services	6
3.2. High Bandwidth Ethernet Services	7
4. The Building Blocks.....	7
4.1. Network Terminating Equipment.....	7
4.1.1. NTE Remote Management	8
4.2. Wave Integration Shelf.....	9
4.2.1. Capacity Planning	10
5. Leaveraging Access and Metro together: OTN Tails	11
5.1. Access MUXs to support 400 Gbps	12
6. Complete end to end architecture	13
7. Benefits	14
7.1. Engineering and Service Delivery	14
7.2. Procurement and Deployment Engineering	14
7.3. Operations and Service assurance	15
8. Conclusion.....	15
Abbreviations	15

List of Figures

Title	Page Number
Figure 1- Three independent CATV providers.....	4
Figure 2- Two neighboring networks continue to grow	4
Figure 3- Single network with CBH, redundant hubs, some ROADMs and various transport vendors.....	5
Figure 4 - Meshed Metro Optical Networks with ROADMs and Flex Technology	6
Figure 5 - LR to DWDM conversion	8
Figure 6 - Management IP for NTE being distributed using GCC0.....	9
Figure 7 - Muxponder. All circuits share A and Z termination	10
Figure 8 - OTN aggregation allowing different A and Z termination	10
Figure 9- Recolor with back-to-back cards in WIS.....	12
Figure 10- Spectrum consumption by 200 Gbps, 400 Gbps and 800 Gbps Wavelengths	13
Figure 11 - Unified Optical Architecture	14

1. Introduction

The optical networks present today for cable television (CATV) systems look nothing like their predecessors first deployed with the birth of fiber nodes. Over the years, as technology evolved, these networks were augmented and upgraded as customer expectations also grew requiring a more reliable, higher-quality experience. As many smaller CATV companies were operating with various technologies and architectures, mergers and acquisitions continued to take place creating larger multiple system operators (MSOs). Having this variety was not inherently an issue with the residential services the networks were designed to support, which were primarily one-way plants providing linear video. With the advent of two-way plants and high-speed Internet (HSI) access, it was still not an issue to have separate optical networks. Other than backbone, services either originated or terminated on devices in a local headend. There was not a requirement for an optical service to continue beyond this termination. Even as commercial services were originally productized to support cell backhaul (CBH), and soon after retail metro Ethernet (ME) customers, the same access fiber infrastructure used to support hybrid fiber/coax (HFC) architectures fit the bill. Once the fiber circuit reached a headend, if it needed to continue to another location it became an Ethernet service and leveraged the routed Internet Protocol (IP) network. The introduction of two new commercial offerings, wavelength and high bandwidth Ethernet services, has made using the existing architecture a challenge.

This paper demonstrates how the metro optical network, optical transport network (OTN) tails over the access fiber, small optical shelves at customer sites, and dedicated commercial shelves at hub sites can be combined to support multiple types of commercial business. This enables new features like remote management, performance monitoring (PM) data, alarming, and a full end-to-end circuit view including the customer site. In addition to these operational benefits there are other efficiencies seen by using the same hardware and software as the rest of the core network.

2. Network evolution

To understand what the new architecture means it is important to look in a little more detail as to where the metro and access networks started, how the current state was reached, and where they are seen going.

2.1. Many small CATV providers

Some of the earliest deployments of fiber in CATV were to support one-way, video-only nodes, typically using low-count fiber cables feeding large pockets of customers from a headend, and areas that were beyond the reach of long coaxial trunk runs. These were single-threaded and there were some, but not many, hub sites. At this time it was not uncommon to have several small CATV operators with shared borders. Figure 1 represents this with three independent providers and a limited amount of fiber, hubs, and fiber nodes.




Figure 1- Three independent CATV providers

2.2. Acquisitions and HFC

With the upgrades to two-way plant and HFC architectures many new nodes were added as operators began creating hub sites with diverse routes and transport. These were mainly passive, point to point systems with amplifiers (AMP) and multiplexers (MUX) but no reconfigurable optical add / drop multiplexers (ROADM). Mergers and acquisitions continued over the years resulting in larger cable companies. This is shown in Figure 2. The three previously independent CATV networks illustrated in Figure 1 have grown and merged and are now owned by a single operator.




Figure 2- Two neighboring networks continue to grow

2.3. Networks merge, advanced services, cell backhaul and ROADMS

With advanced service deployments like HSI and video on demand (VOD) in full swing, the idea of placing routers in the hubs with redundant feeds to regional aggregate routers and video content servers became common. Systems continued to merge and some of the headends became hub sites. Each previously separate network was connected with fiber and active transport. ROADMs become more prevalent. This is also when the first commercial services over fiber were introduced, consisting primarily of CBH connecting cell towers to carrier's mobile switch centers (MSC). This made sense as the access fiber placed during HFC builds passed close to many towers. The routed infrastructure was getting more robust and resilient. At this point, there were many transport vendors and technology types. While the geographic borders between networks are blurring, transport networks that do not interop stop in the same hub keeping the optical networks still mostly segregated.




Figure 3- Single network with CBH, redundant hubs, some ROADMs and various transport vendors

2.4. All ROADMs, Flex Technology and one vendor

Over time, bandwidth needs became much greater for residential and commercial needs. Many systems within a region became part of the same network. The same technology type and vendor of transport are being deployed and meshed in the metro optical network. The access still stops at the metro hub.

From here forward all networks are being deployed with flex technology with an initial focus on supporting 400 gigabits per second (Gbps) wavelengths.¹ This paves the path for the IP infrastructure to use 400 gigabit Ethernet (GbE) native interfaces. Flex technology is also known as a colorless

¹ To better understand sections in the remainder of this paper there are important definitions when referring to wavelengths that are dependent on context. Wavelengths can be defined by the center frequency that aligns with an International Telecommunication Union (ITU) standard channel. The width of the wavelength can also be referenced in gigahertz (GHz). And lastly while referring to a network facing wavelength the transmission rate can be reflected in Gbps.

system. Instead of a traditional fixed wavelength grid add and drop structure with either 50 GHz or 100 GHz channels, flex technology allows 6.25 GHz slices of spectrum to be combined into any width wavelength. Doing this allows more efficient use of spectrum and future proofs the network against new technology that may require wider wavelengths than the fixed grid systems can support. At the same time, Comcast has decided to add two additional optical design requirements: Layer 0 (L0) control plane capable which allows protection and restoration of optical paths through wavelength switching and L band. These both make the network more robust by adding resiliency and capacity.




Figure 4 - Meshed Metro Optical Networks with ROADMs and Flex Technology

3. Products

Equally important to understanding the underlying fiber and transport infrastructure is knowing the requirements and challenges of the two products that drove the creation of the unified architecture.

3.1. Wavelength Services

Wavelength services is a product that delivers a single or diverse point-to-point, optical (Layer 1) 10 Gbps or 100 Gbps circuit over the Comcast dense wavelength division multiplexing (DWDM) network. There is no Layer 2 (L2) switching, Layer 3 (L3) routing, or any IP customer premise equipment (CPE) provided by Comcast.

Simply stated, wavelength services are point-to-point, transparent pipes over long distances with extremely low latency. The customer can connect their own IP devices between sites and as far as they are concerned, it is like a fiber jumper, is transparent, and lower latency than Ethernet services.

It is different than dark fiber in that PMs are available to observe for degradation or other issues. It can also go much farther because it uses optical transport.

3.2. High Bandwidth Ethernet Services

This product differs from wavelength services in that there is a Comcast owned and managed CPE on the end of the circuit and can be either point-to-point between customer sites or Internet access, referred to as Ethernet dedicated Internet (EDI). To achieve this at 10 GbE and lower rates, the CPE are connected over the access network to Comcast routers using coarse wavelength division multiplexing (CWDM) or DWDM optics, and leverage the IP infrastructure to transit to other customer sites or the Internet. With customers now requesting much larger bandwidth (up to 100 GbE) this cannot always be supported without extensive augments to the IP network and platform upgrades to support 100 GbE interfaces. This can lead to long deployment times and be cost prohibitive.

To solve this issue the access and metro optical networks can be used together with the same unified optical architecture as wavelength services to backhaul the high bandwidth Ethernet services to locations that the IP infrastructure can support.

4. The Building Blocks

The unified optical architecture started as a solution for wavelength services. This product needed to be launched in several markets that had different transport vendors and technology types. All tools for order processing, design, service delivery, and service assurance were centralized and needed standards applicable across the board to successfully provide the experience customers expect. It would not be possible to handhold each service and use development cycles on several iterations of the same product.

As seen above in the network evolution, the metro optical network is starting to be well-defined and positioned to carry any line of business (LOB) to any location. While the access is also maturing to support distributed access architecture (DAA) technology, it does remain segregated from the metro. This is not necessarily a bad thing provided the two networks can work together when needed.

With the metro being able to carry any LOB agnostically there needed to be a way to provide demarcations for engineering, service delivery, and service assurance. This was done by subtening independent devices dedicated to commercial services under the metro shelves. These needed to have the same role and functions regardless of the vendor. Thus, the network terminating equipment (NTE) and wave integration shelf (WIS) were born.

4.1. Network Terminating Equipment

The NTE is a small optical transport shelf that today can support 10 GbE and 100 GbE services over up to 400 Gbps wavelengths. It resides at a single customer site or data center to serve multiple customers and assumes the role similar to a CPE or edge gateway (EG) would for Ethernet-only services. A challenge when developing wavelength services was deciding how to hand off to the customer. Being an all-optical product, a point of demarcation was required to be the handoff. When only a demarcation was needed, a fiber patch panel could be used. A fiber patch panel directly connecting a customer to the network was not an option because there also needed to be a way for the wavelength and transmit power to be controlled.

A major roadblock for using high bandwidth Ethernet services was that anything over 10 GbE of committed information rate (CIR) would require the use of a 100 GbE interface. The use of 100 GbE

optics in IP platforms is not new. This has been widely adopted by the industry for years. 100 GbE has become the new 10 GbE. However, these are all 1310 nm, long reach (LR) optics. The only way a CPE could be connected to the IP core with these optics would be through a dedicated pair of fiber with no MUXs assuming the customer was within ~10 km of the Comcast router. At the time of creating this architecture, coherent DWDM optics at 100 Gbps or greater for the routers and CPE were in their infancy and were not within the project timeline. Even as this paper is being written these optics and their host platforms are still being developed and are not ready for production at scale.

The solution for this is to use optical transport cards that support DWDM optics at 100 Gbps (or greater). This allows taking a 1310 nm, LR signal from the router and CPE, and converting it to a DWDM wavelength for use over the access MUXs as shown in Figure 5.

To deploy this LR to DWDM conversion a place was needed to house the transport cards. The NTE fit this role perfectly at the customer site.

An additional benefit to placing the NTE is now the device can be remotely managed and enrolled in the transport network management system (NMS). This allows alarms to be monitored, PM data is available, loopbacks can be placed remotely for troubleshooting, provides an end-to-end circuit view and other features useful for supporting the service.

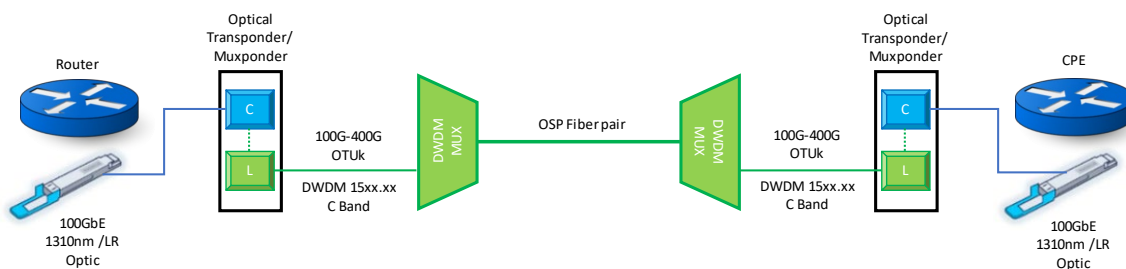


Figure 5 - LR to DWDM conversion

4.1.1. NTE Remote Management

With many proposed solutions come new challenges. To enroll the NTE in the transport NMS and be able to reach it remotely it needs a management IP address. In most cases the IP for transport devices is provided by directly connecting to a co-located switch. With the NTE being in a customer site there is no switch present. The closest switch is typically in the hub at which the fiber from the customer site terminates. The solution is to use the general communications channel (GCC). The GCC is bytes embedded in the overhead of transport links and can be used for small amounts of information. There is GCC0, which is two bytes within the optical transport unit (OTU) overhead, and GCC1, which is two bytes in the optical data unit (ODU) overhead. GCC0 was selected because the OTU is on the line (or network facing) interface. GCC1 is associated with ODU which is tied to the client payload. While both are transparent to the customer it was decided to keep the management on the network facing interface partially due to more hardware supporting GCC0.

One caveat when considering hardware in this architecture is that not all interfaces support GCC0 or GCC1. Due diligence was required to select the right transport cards and shelf for the NTE role to ensure it supported this feature.

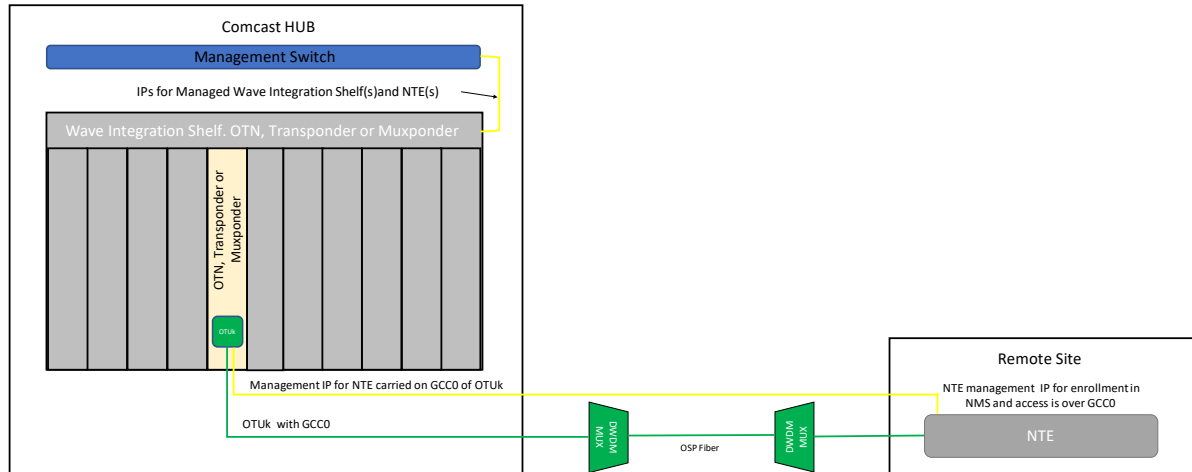


Figure 6 - Management IP for NTE being distributed using GCC0

4.2. Wave Integration Shelf

In addition to needing a device in customer sites another role was needed in the hubs and headends. It serves a similar point of demarcation function as the NTE. It also functions as an optical aggregation point. Other than more effectively using wavelengths through aggregation this also solves another issue. As optical networks are upgraded to flex technology they are optimized for coherent channels. In doing this the dispersion compensation modules (DCM) are removed. Without dispersion compensation 10 Gbps non-coherent channels will not work on most spans. By making the WIS an aggregation point for 10 G channels they can be combined onto 100 Gbps to 400 Gbps coherent wavelengths for transit across the flex optical network. There was a decision to be made here regarding the type of aggregation technology to use.

Traditional muxponders could be used to combine the nx10 Gbps circuits to a single coherent trunk. As seen in Figure 7, all traffic must have the same A and Z ends. Analysis showed it would not be common for multiple 10 Gbps circuits to be going to same end point except in the case of dense data centers.

A better solution was to use OTN switches. This allows traffic to be assigned to a slot on the trunk that can be picked out at a location along the path between A and Z without impacting the pass-through services. This is particularly effective when you expect a common A end with different Z ends along a path as is often seen with single customer sites that need to reach a common data center.




Figure 7 - Muxponder. All circuits share A and Z termination

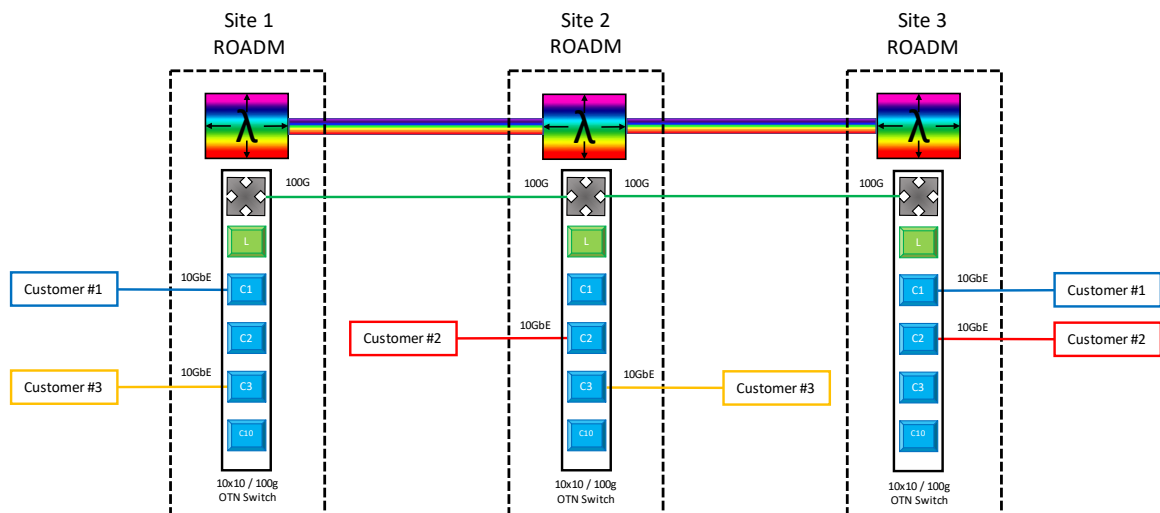


Figure 8 - OTN aggregation allowing different A and Z termination

4.2.1. Capacity Planning

Much of the focus on creating the unified optical architecture has been on engineering and operations but another important piece is planning, specifically capacity planning. Commercial services have sales forecasts, but they are very fluid. Residential networks have well-known, proven compound annual growth rates (CAGR) and can be planned with good accuracy. Since the metro optical network supports all LOBs to effectively manage slot capacity, each demand, or group of demands, benefits from having its own client shelf to support the cards needed. The WIS fits this role perfectly. It provides a shelf that will only support commercial services. Here is an example to illustrate the benefit.

Assume a shelf in a headend that supports both residential and commercial had five open slots, and the CAGR for residential traffic dictated that over a year one slot per quarter would be consumed.

That leaves one slot to support a commercial sale. Initially, that is not a significant issue. With commercial being so fluid though, sales can come in at any time. If more than one was to come in before the planned residential circuits were turned up, the commercial asks would consume the slots planned for residential. Everything could be oversized to make sure there is enough slot capacity but with space and power being a premium it does not make sense to do this “just in case.”

The WIS solves this issue by giving the commercial cards a place to reside other than the well-known static needs of the residential network.

This benefit proved itself with the capacity explosion during COVID-19. With no risk of a commercial demand unexpectedly consuming a slot in the residential shelf, engineers were able to execute augments to keep up with bandwidth needs on the residential network without delay while at the same time being able to provide high bandwidth commercial services to business customers to support remote employees needing to reach their infrastructure.

5. Leaveraging Access and Metro together: OTN Tails

Now that the hardware and roles have been defined to solve those challenges, we can address the separation of the access and metro networks and how they can be merged when needed to create a single end-to-end circuit.

As seen in the above fiber evolution maps, the metro transport has become more meshed and better able to support an any-to-any traffic pattern. The access has also gone through evolutions but what has stayed the same is the fact that it stops at the hub site. It has been found that most customers can connect to the access and be brought back to a hub site. But the remaining question is how can they be connected to each other if their fiber terminates at different hubs? Or in the case of EDI, how do they get to the router that provides internet Access? This is where the access and metro can work together by using OTN tails.

OTN is taking an Ethernet (or other protocol) payload and encapsulating it in a wrapper to be carried across an optical transport network. It is completely transparent, and a standard governed by ITU G 7.09.

An OTN tail is transmitting that wrapped payload over dark fiber or a channel on a MUX over fiber.

Using this tail and transponders (or muxponders) in the NTE and WIS facing each other can extend the metro core over the access without using ROADMs or even AMPs.

This tail is then transmitted over the core by a second card in the WIS with an OTU client-to-client connection to the card facing the NTE. This back-to-back configuration allows a DWDM OTN signal to face the NTE and the add / drop structure of the metro. Client-to-client connections have been made in the past as Ethernet where a similar design was used for regens. In this case the decision was made to keep them OTN so that the circuit remains transparent from end-to-end. There is no conversion to Ethernet and back to OTN. One advantage of this is that the trail trace identifier (TTI) trail trace remains intact from end-to-end even if the circuit crosses to another network or even to another vendor. You can transmit a trace on the A end and read it on the Z end. This is particularly useful for troubleshooting and topology verification.

The signal flow described above is applicable to 100 GbE services using a 100 Gbps to 400 Gbps trunk.

10 Gbps services differ slightly at the WIS. They still use an OTN tail over access from the NTE to the WIS but instead of back-to-back cards the 10 Gbps is aggregated on the OTN switch card.

Using the back-to-back cards or OTN switch card has also solved another issue with using MUXs in the access. By having the ability to tune to any wavelength facing the NTE or the metro it is not required to maintain the same wavelength end to end from NTE to the metro, across the metro, and to the other NTE. An effective way to illustrate this is to imagine many MUXs being deployed in the access with ITU channels (Ch) 20-59. The metro uses the same channels and, in some cases, adds another set of 50 GHz spaced channels the access does not use. If Ch 20 is used on the A end without this recoloring (that is, changing wavelength) in the WIS you would need to have Ch 20 open on your metro and the other end of the circuit over the access. Once you use Ch 20 for one circuit in the access or metro you cannot use it again. Back-to-back cards allow any channel that can reach the hub site where the WIS is located to be recolored to any open channel across the metro.




Figure 9- Recolor with back-to-back cards in WIS

5.1. Access MUXs to support 400 Gbps

With the metro being a flex system, it can combine slices of spectrum in 6.25 GHz increments to create the width needed per channel. Legacy fixed grid transport used two sets of 44 channel 50 GHz MUXs. This provided up to 88 channels but with a maximum per channel width of 50 GHz. This would support up to 200 G channels that are typically ~37.5 GHz wide.

The Access followed the same fixed grid pattern with one important distinction. The MUXs used were 100 GHz spaced. Why is this important? The last major technology advance that has been heavily adopted is the use of 400 G line rates. Depending on vendor and specific technology type, the width of this is ~75 GHz. With the access having up to 100 GHz of width available per channel, 400 G wavelengths can be deployed on a system that was first created and optimized to run 10 G. This is a fantastic use of the available spectrum; however, it does appear to be the end of the road. The next technology being adopted by the industry, including Comcast, is 800 Gbps wavelengths. As can be seen in Figure 10, an 800 Gbps wavelength occupies ~112 GHz of spectrum and will not fit in the 100 GHz MUXs used in the access.




Figure 10- Spectrum consumption by 200 Gbps, 400 Gbps and 800 Gbps Wavelengths

There is one particularly important caveat when using 100 Gbps and greater wavelengths over the access. Today's coherent optics by nature require both the transmit and receive frequency (wavelength) to be the same. This requires the use of a two-fiber MUX. A single-fiber MUX uses different frequencies for the transmit and receive.

Being that single fiber MUXs are the standard deployment at Comcast for all new DAA builds this puts the unified optical architecture at odds with DAA. However, work is being done in the industry to create bidirectional coherent optics to use for this application that would operate on a single fiber MUX.

Comcast Fellow and industry leader in this effort Venk Mutalik says, "Many businesses are within a short distance of optical fiber nodes, and with the expansion of DAA, we are going to have more of them out in the field. So, a really great way to use fiber assets already available is to converge business services on the same fiber as the residential services. This is done rather easily with 10 Gbps services, but with the ability to do the same with 100 Gbps and even 400 Gbps coherent services on the same fiber, access convergence is a game changer for the industry!"

6. Complete end to end architecture

With the building blocks established and the access and metro working together, a complete end-to-end architecture can be established. In Figure 11, each colored line represents a different commercial service type and rate.

All services, regardless of type, originate or terminate on an NTE at the customer site, use an OTN tail over access and are added into the metro optical network via WIS.

In Figure 11 the following types of services are shown.

- Blue – 100 GbE EDI
- Red – 10 Gbps wavelength service, customer site to data center
- Green – 100 Gbps wavelength service, customer site to customer site
- Orange – 100 GbE EDI

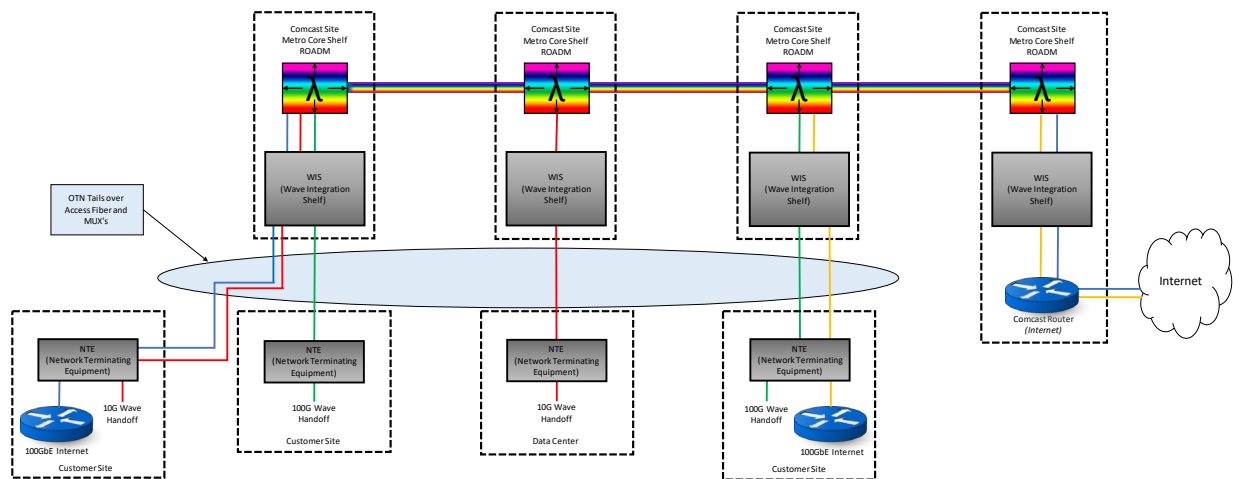


Figure 11 - Unified Optical Architecture

7. Benefits

All the work is now done. Building blocks have been assembled, technology challenges have been overcome, and a unified optical architecture to support multiple lines of commercial services has been created. The benefits are simple and clear. By developing this vendor-agnostic, unified optical architecture, commercial circuits can be deployed and supported the same way every time regardless of the type of service or traffic pattern, while increasing efficiency and maintaining or improving customer satisfaction.

The overarching theme is sameness. As explained by Shane Portfolio, Comcast Senior Vice President of Network Services, and leader of the Comcast One Network Initiative, “The value of sameness should not be underestimated. It allows for speed to market, greater scale, increased customer felt reliability, costs reductions, and enables next generation technology like machine learning, automation, and artificial intelligence to come to life because it is a homogeneous system that can be built to leverage these technologies in ways doing things differently simply cannot. It’s a game changer.”

7.1. Engineering and Service Delivery

By using the same architecture regardless of vendor, engineers can streamline processes and become deeper subject matter experts. The designs will always include the basic building blocks of the NTE, WIS, access and metro core. They will always follow the same device and port naming standards.

7.2. Procurement and Deployment Engineering

Using the same equipment as the metro networks benefits procurement by reducing the number of unique items in the ordering systems. A basic tenant of economics, volume buying power, also applies here.

Deployment engineering benefits from the ability to shuffle hardware to the highest priority project. Using the same hardware for multiple demands makes this possible. This again was a valuable option to meet demand and bandwidth shifts during COVID-19.

7.3. Operations and Service assurance

Possibly the biggest benefactors are the operations and service assurance teams. They are tasked with keeping the circuits and network healthy and code up to date. The ability to see PM data like optical levels and errors on an interface as well as place loopbacks at the customer site remotely is invaluable. By using the access to create OTN tails and merging them into the metro, the power of the NMS can be used to see a circuit from end to end including a simple green is good, red is bad, graphical representation. Alarms are also collected from all the optical devices and centrally managed in the NMS. This can dramatically cut the time to resolve an issue. Using the same hardware as the metro for NTEs and WIS also streamlines code certifications, deployment, and development efforts by reducing the variety of equipment deployed.

One of the greatest accomplishments this architecture has made possible is the development of software that merges customer account information, provisioning details, and correlates L2 and L3 ports to their connected transport clients. This also shows a hop-by-hop trace of the circuit from the NTE over the access and metro. Work was even done to stitch together circuits on multiple networks provided by different vendors. This creates a true end-to-end view with all the information an engineer needs to understand the circuit in one place. Comcast and a partner company co-developed this software.

8. Conclusion

All the work that was done to create a unified optical architecture may seem obvious. Why wouldn't the same equipment and standards already be used everywhere? The reality is that with the speed at which networks have merged and evolved they have been in a perpetual brownfield state with a variety of architectures. If everything were a greenfield build with no existing network to support, this would certainly be a lighter lift. In addition, optical technology has gone from 10 Gbps to 1.2 terabit per second (Tbps) wavelengths over the last decade. Simply keeping up has required the full attention of all teams. To put everything together has taken several years and a dedicated effort beyond keeping the lights on. In the end, though, by creating this unified optical architecture the goals to continue to deliver or improve the experience customers are accustomed to, offer new technology, and increase efficiency have been achieved. By not only solving the challenges present today but looking ahead to the future while creating this architecture, the networks are well-positioned to deliver the next generation of 400 Gbps wavelengths and 400 GbE to the customers with no changes.

Abbreviations

AMP	amplifier
CAGR	compound annual growth rate
CATV	cable television
CBH	cell backhaul
Ch	channel
CIR	committed information rate
CPE	customer premise equipment
CWDM	coarse wavelength division multiplexing
DAA	distributed access architecture

DCM	dispersion compensation modules
DWDM	dense wavelength division multiplexing
EDI	Ethernet dedicated Internet
EG	edge gateway
Gbps	gigabits per second
GbE	gigabit Ethernet
GCC	general communications channel
GHz	gigahertz
HFC	hybrid fiber/coax
HSI	high-speed Internet
IP	Internet Protocol
ITU	International Telecommunication Union
L0	Layer 0
L2	Layer 2
L3	Layer3
LOB	line of business
LR	long reach
ME	Metro Ethernet
MSC	mobile switch center
MSO	multiple system operator
MUX	multiplexer
NMS	network management system
NTE	network terminating equipment
ODU	optical data unit
OTN	optical transport network
OTU	optical transport unit
PM	performance monitoring
ROADM	reconfigurable optical add drop multiplexer
SCTE	Society of Cable Telecommunications Engineers
Tbps	terabit per second
TTI	trail trace identifier
VOD	video on demand
WIS	wavelength integration shelf

Up Your Up-Time with Automation: Outage Pre-Verify

A Technical Paper prepared for SCTE by

Kathy Fox

VP Program Management-Operations Support
609 Odin Rd, Coudersport, PA 16915
(856) 912-7850
Kathy_Fox@cable.comcast.com

Joann Shumard

VP Engineering Operations-Network Services
1800 Arch St, Philadelphia, PA
(770) 652-3836
Joann_Shumard@cable.comcast.com

Table of Contents

Title	Page Number
1. Abstract	3
2. Introduction.....	3
3. Process Alignment	3
4. Technical Development.....	4
4.1. Design Methodology	4
4.2. Design Details	5
4.3. Challenges	6
4.4. Delivery	6
5. Trusting The Technology	6
6. Change Impact.....	7
6.1. Planning the Change.....	7
6.2. Implementing the Change	8
6.3. Challenges	8
7. Value & Results.....	8
8. Conclusion.....	9
Abbreviations	10
Bibliography & References.....	10

List of Figures

Title	Page Number
Figure 1: Manual to Automated Transformation	4
Figure 2: Agile Methodology [1]	5

List of Tables

Title	Page Number
Table 1: Average savings, in minutes, for triage activities with automation	9
Table 2: Quarterly Time Savings Potential	9

1. Abstract

The goal of tooling automation in operations centers is to speed the identification of issues and translate them into the swiftest fix agent response. Efficient management of outage events relies on multiple layers of detection tools that provide critical information for operations centers to coordinate event management and resolution. Numerous operation centers rely on Lines of Questioning (LOQs) to walk agents through key troubleshooting points to create consistent outcomes for identifying the root cause of system events.

The opportunity to reduce triage time comes from integrating key event knowledge points and automating LOQs to streamline assessment and dispatch tactics. Outage Pre-Verify (OPV) is a tooling automation strategy that creates technology solutions to automate current operation center processes. Process automation requires in-depth linkage of available system status data, understanding outcomes expected from LOQs, tooling development strategy, process optimization, and the capability to transition people through system changes. OPV takes the technology outage telemetry and enhances operational center processes by reducing triage time, improving event response time, and optimizing fix agent resources.

This paper and presentation will review strategies for identifying automation opportunities in outage triage. It will discuss how to optimize effectively transitioning people and business processes to integrate new tooling automation. Authors Kathy Fox, VP of Product Management, Excellence in Operations Centers (XOC), and Joann Shumard, a VP of Engineering Operations at Comcast, are leaders in Shane Portfolio's organization focusing on operational technology integration. They will detail the tooling transformation process and the impact on the user community. Emphasis will be given to effectively integrating automation change while simplifying the user experience and transforming operational processes.

2. Introduction

For several years now, teams at Comcast have focused on finding best practices and implementing them across all three Comcast divisions. As part of this best practice identification, process sameness opportunities have been discovered and implemented in the design, engineering, construction, and operations. While finding and implementing a best practice across three different divisions often takes a good bit of analysis, debate, and planning, once process sameness is in place, we then look to identify opportunities for simplification and automation. In this paper, we will review the automation of Comcast's OPV work where we focused on implementing process sameness related to both digital and analog node outages, measured the results, and then worked closely with these teams to identify specific improvements we could implement. We will review how we worked through data analysis, use case development, tooling requirements development, and finally tooling development, testing, and implementation to deliver automation that resulted in time savings for operations center technicians and truck roll reduction for our plant maintenance technicians both of which ultimately lead to a better customer experience.

3. Process Alignment

One area we previously gained alignment on across our three division operations centers was OPV. After much analysis and discussion, the teams had agreed upon a single process which included the same LOQs (Lines of Questioning) and timing for node outages to "soak", a term we use to indicate the period during which node components self-clear any alarms. It is important to note that nodes go into soak when a certain percentage of modems served by that node goes through a registration state change (to offline). The cost savings that resulted across our divisions with this process were noteworthy with a significant reduction in No Trouble Found (NTF) or Power Outage-related incidents as well as overtime savings.

However, the additional troubleshooting did create many more tasks (which means more time!) for our access network staff in the operations center. On a few occasions, the LOQs added enough triage time that the operations technicians were not able to keep up with all the triage activities, especially when node outages increased during inclement weather events, resulting in increased plant maintenance truck rolls.

The team began to investigate what efficiencies or automation we might deliver to the business to lighten the load for our operation center teams. There is a great deal of information that can be obtained from our network and the project team began to investigate what information might be available to help us more quickly determine if there was a power outage in the area. Any information that would tell us that nodes were in soak because a lot of customers had lost power would save our operations technicians a good bit of triage time either checking power providers' outage notification tools or calling the appropriate power company.

4. Technical Development

Aligned processes pave the way for automation. When processes and procedures are different across groups, the tooling is maintained in a customized and typically inefficient way. Once there is full cooperation in the way work is completed, the LOQs can be evaluated for automation opportunities. For OPV, building automation is focused on accuracy in identifying dispatchable outage events while reducing triage events for the operations centers. When considering the major steps of event detection—triage (LOQs), and Dispatch of a fix agent—automation provides the opportunity to bypass the manual triage and move the event to an auto dispatch state.

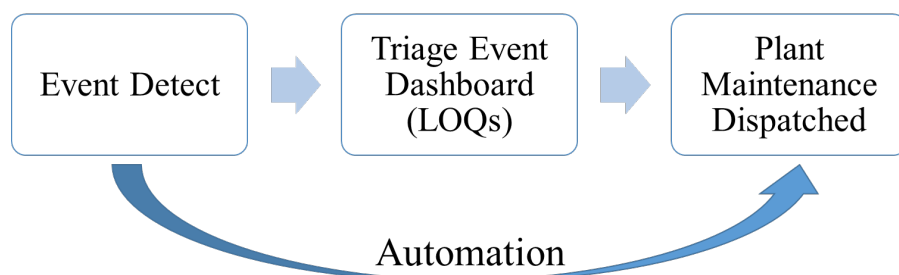


Figure 1: Manual to Automated Transformation

4.1. Design Methodology

In agile development methodology, the needs of the users drive the development requirements [1]. Sprints are used to develop iterations and progress towards initial implementation. For OPV automation, the requirements were captured from the initial manual implementation of standardized LOQs. This required a full review of which aspects of the LOQ could be translated into automated machine language.

Another aspect of agile development is to ensure there is clarity in task prioritization [1]. This is important to reduce the potential of rework and lost development. For OPV, this required full

alignment will all key users to collaborate to maintain the standard from manual to automation. This ultimately translated into the initial user story used for the technology development.

During development, continuous user feedback was required to make final adjustments or address issues [1]. This required consistent connections with the teams currently managing outage events using the manual implementation of OPV. It also provided an opportunity to ensure the process was collaboratively reviewed for improvement opportunities.



Figure 2: Agile Methodology [1]

4.2. Design Details

The logic of recreating the LOQs and drive automation capability was built to align with reduced field truck roll requirements while also reducing transactions for the operation centers completing the manual review. The summary of the key logic points of the LOQs and automation creation are:

- Upstream (US) Bandwidth Utilization.
 - Identification of a significant drop in US traffic;
- Upstream In Bps.;
 - Identification of a significant drop in US traffic.
- Power Supply Telemetry;
 - Identification of power current changes;
- Node Device Offline %;
 - The entire node was determined as a whole node outage determined by % of devices offline;
- Interactive Voice Response (IVR) Call Threshold; and
- Storm Mode Status.

These logic points are based on leading outage-based plant characteristics and time-boxed to provide assessment considerations for the algorithm. The target was to ensure that the logic was constructed so that it could be configured and adjusted as a part of continuous agile improvement. When considering the build of this type of automation, these characteristics and thresholds may fluctuate based on system

architecture variations. The storm status is a specific consideration that considers how the algorithm adjusts as a system transitions from plant outages created by typical event types (degradation, plant damage, etc.). to the more rapid pace of storm creation where dispatching may be held until it is safe for fix agent technicians to engage.

4.3. Challenges

Through the agile process, consistent feedback from the users was key to preparing for the adoption of the technology. One challenge was ensuring consistency in the thresholds across the organization. There was considerable discussion about the variety of plant characteristics and local considerations that had to be integrated into the final algorithm. Overall, this challenge was met through the partnership commitment to consistency and the agility of adjustments throughout the build process.

The integration of power outage information is a continuing challenge that the team continues to evaluate. When the power supply telemetry identifies the absolute loss of commercial power, it combines into the logic cleanly. The issue is that power grid boundaries do not consistently line up with node boundaries which makes the automation more challenging. The logic requires telemetry of every system power supply to be reporting consistently but due to maintenance issues, a number of units may not be fully reporting. This ambiguity in the power status supporting the node can cause the event to transition from an automated event to a manual review triage requirement.

Obtaining a unified buy-in of the working logic of the resource teams required a commitment to integrating continuous feedback while utilizing data modeling to confirm impact assessment. The logic build required that consistency is maintained through collaboration and strategic focus on a common goal. With multiple division teams, opinions on the solution path varied at times but having a change strategy supported the success of the implementation.

4.4. Delivery

This logic sits in a user interface tool that is used to process the logic and translate it into actionable alerts and jobs. When the automation logic triggers a pre-verified outage event through automation, the logic flows to an automated dispatch of a plant maintenance technician to resolve. This ultimately reduces triage time by operations team members. Events that do not meet the logic parameters or are unclear become a dashboard task for triage and validation. The introduction of automation creates an immediate reduction as well as creates a platform for continuous process improvement.

5. Trusting The Technology

Research tells us that people will inherently trust technology to be accurate, but it relies on the simplicity of the data and how well they know the technology [2]. Research shows that people are more suspicious of accuracy when they tend to confirm information validity using their expertise and skills. This algorithm complexity also drives the teams who have historically reviewed the work directly to want to see incremental validation before trusting the value.

The trust-skepticism balance can be overcome through targeted training and providing cross-reference data that illustrates confirmation of event accuracy [2]. Balance can be achieved by implementing system verification and reporting to work in parallel with the new technology algorithm. The skepticism is created from the psychological bias that some may have against trusting the technology to replace the work they are completing. The vetting of information through the agile process development is a key component of the feedback loop.

Since Outage Pre-Verify (OPV) was an embedded manual process being conducted by industry experts, building trust in the information was key to preparing to adopt the new technology. This was accomplished through the trial process where events were confirmed valid when tagged for automated dispatching. Another key aspect was the creation of reports that illustrated that the volumes were not significantly changing while also measuring the value to the business.

Understanding the impact of the change is a key component of preparing the teams to trust the innovative technology. This logic analysis is designed to reduce LOQ manual tasks. The trust in the automation requires the team to build confidence in the accuracy of what is dispatched without manual validation.

6. Change Impact

With any technological change, we must care for the tools, the processes, the documentation, and most importantly, the people. Organizational change management has become part of our standard project practice and has been well-received by the teams we serve. Our teams have been utilizing the Prosci ADKAR Model (Awareness, Desire, Knowledge, Ability, Reinforcement) for a few years now with impressive results. This powerful model is based on the understanding that organizational change can only happen when individuals change. The ADKAR Model focuses on individual change—guiding individuals through a particular change and addressing any roadblocks or barrier points along the way [3]. We have found that investing time upfront to understand the impact on our people and to develop plans that care for those impacts, enables an improved implementation experience and is time well spent.

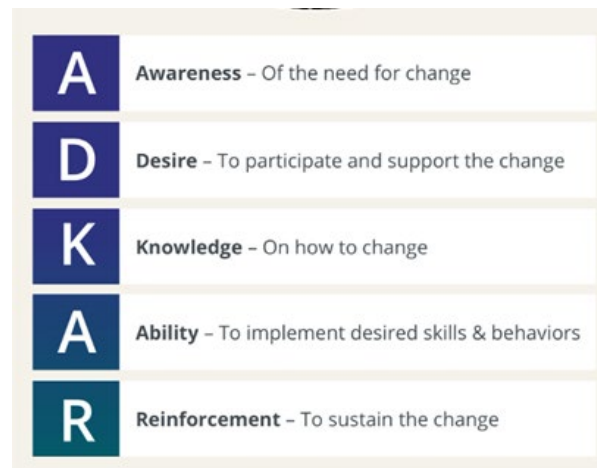


Figure 3: The Prosci ADKAR Model [3]

6.1. Planning the Change

During the initial project planning stages, we perform a change impact assessment. This involves project team members, change agents, and representatives from each of the expected impact areas. During the assessment, we look at each impacted role, discuss the expected impact in detail and determine if the change impact is high (a ‘5’) or low (a ‘0’ or ‘1’) or somewhere in between. That information is then used in the development of the organizational change management plan.

For OPV, our operations center teams had been investigating node outages for several years, so they knew the time needed to determine if a node is in soak because of a commercial power outage. Our three division operations centers have somewhat different staffing models. Although the three divisions had agreed upon a single set of LOQs for OPV, one of our divisions had difficulty executing all the questions

for the node soak triage because of staffing. As a result, they omitted some of the questions in the LOQ, expecting that future development would introduce time savings and negate the need for these questions. One of our other divisions had added the full set of LOQs but had not increased staffing for this effort so the team was feeling the impact on their workload. There was a desire on the part of the divisions to reduce the workload for the operations teams while maintaining the benefit of reduced plant maintenance truck rolls. The solution developed by the technical team was expected to reduce triage time by three minutes which is a significant impact given the thousands of triages performed by the operations centers each week. For the one division that had the abbreviated list of LOQs, the automation aligned them with the other two divisions.

For this technological change, the team quickly understood the resulting work benefit and considerable time savings, but there was also the potential that this advancement could introduce feelings of job insecurity or there could be a lack of trust in the solution being implemented. Insecurity or lack of trust in the solution could lead to resistance by the individuals impacted by this change. We wanted to proactively address the concerns and make certain the operations teams had Awareness of the change and a Desire to adopt it and the Knowledge to trust the solution (the “A” Awareness, “D” Desire, and “K” Knowledge in ADKAR).

6.2. Implementing the Change

When we can introduce an operational change as a trial, allowing the end-users to execute testing, it is a significant benefit. The first benefit is those end-users know the different scenarios they encounter every day, and they usually assign their most skilled users to conduct testing, validate the solution, or find issues that need to be addressed. The second benefit is the testers are part of the solution, they take ownership and often become champions of the change being implemented, helping other team members to adopt and accept the change once we move to full operationalization, supporting the “A” for Ability in the ADKAR Model.

Implementation plans will vary based on the Change Assessment results. For high-impact changes, we are certain to include formalized training, extensive communications, both written and verbal, and often, testimonials from testers or early adopters. We have created a detailed list of tactics that should be used based on the degree of impact. For the automation related to OPV, we considered it a medium-level impact and focused our communication on the efficiencies being provided to the teams while also sharing the new areas of focus for the teams. We were removing busy work to free them up to focus on more complex activities.

6.3. Challenges

Innovative technology implementations typically produce challenges. With OPV, we were fortunate to have the operations teams take part in the trial/testing phase which helped identify several technical improvement opportunities before implementation. We also found a number of process variances related to multi-node outages but were also able to address those, allowing for the solution development to be built one way.

7. Value & Results

The development process first determined what the initial goal was to create business value for the organization. The expected results were used to identify what level of reduction could be accomplished to reduce manual tasks. The reduction is calculated based on the total number of Hybrid Fiber Coax (HFC) outage events created then calculating the number of outages that were captured in the OPV engine

and sent directly to the fix agent technicians. The total time per task for the manual process is calculated for this series of events. To calculate an estimate of cost savings, an industry estimated rate of \$40 per hour is used for illustration. The launch of the new algorithm was implemented in mid-January 2022. The year-to-date average of reduced manual triage handling is 24% which is now the basis to increase automation in future iterations.

The figure below shows how significant the time savings are. The top line shows OPV performance for the three months before implementation of the OPV Automation. The bottom line shows the improvement for the most recent three months' performance with automation. The number of outages to be triaged is significantly less and the average completion minutes has dropped by more than the expected three minutes. Proper change planning ensured that the teams were part of both the journey and the solution, knowing what other tasks they would be working on once this automation was delivered. This type of reporting is shared with the teams to reinforce the benefits and highlight their contributions; the "R" in ADKAR is for reinforcement.

Table 1: Average savings, in minutes, for triage activities with automation

	#Triage	Avg Completion Min	Avg Min on Dash	Max Min on Dash
OPV Prior To Automation	220,917	9.8	18.6	797
OPV With Automation	167,958	5.1	12.2	594
Reduction/Improvement	52,959	4.7	6.4	203

Susan Bean, Sr. Director, XOC & Excellence in Plant Maintenance (XPM) summed up the OPV benefit this way. "Last year, the division XOCs set out on a journey to showcase processes and collectively solve for areas of opportunity. One area of focus was finding efficiencies for node outage triage with the goal of quicker restoration times for our customers. The collaboration of the operations and development teams was critical to finding a common approach. We were able to identify a successful solution, removed non-complex tasks for our operations teams, and redirected their efforts to more complex work."

When calculating the value of the business implementation there are two considerations. The first is the overall task reduction which can be multiplied by the average manual task time to get working time saved for the quarter view. The secondary improvement is generated in the reduction of task triage time through the streamlining of the LOQs and reducing handle time. This led to an opportunity to save approximately 250k minutes in the quarter that could then be used to support additional events.

Table 2: Quarterly Time Savings Potential

	# Jobs Triage	Average Completion Minutes	Total Time Minutes
OPV Prior To Automation	220,917	9.8	2,164,987
OPV With Automation	167,958	5.1	856,586
Time Saved	52,959	4.7	248,907

8. Conclusion

This evolution of automation followed a path from the historical manual triaging process of plant-related events but was transformed into a path of optimizing performance through automation. Aligning processes is foundational to tool development. When processes vary, the implementation of the tool is inconsistent and inefficient thus restricting the ability to automate. The focus of this initiative was to

capitalize on the efficiencies gained through the consistent process and LOQs and utilize technology to reduce transactions and time to complete tasks.

Having an agile and inclusive development plant paves the way for the adoption of these new changes. Through collaboration, the continuous feedback loop in the development lifecycle aided in change acceptance and building trust in these new ways of completing work tasks. This first phase of automation has reduced total transactions by 24% while improving LOQ time efficiency by 48%. These strong results illustrate the reinvestment opportunity of these resources completing manual work by providing them the opportunity to focus more on complex tasks.

As the network continues to become more intelligent, there will be more information available, allowing our tools to self-diagnose and self-service for repair. This will free up our people resources to focus on more complex issues that need investigation as well as occasional problem escalations for resolution. To be successful, we must utilize agile development methodologies and exercise proper organizational change management practices to advance the network, systems, tools, and people collectively.

Abbreviations

ADKAR	Awareness, Desire, Knowledge, Ability, Reinforcement
LOQ	Line Of Question
OPV	Outage Pre-Verify
US	Upstream
IVR	Interactive Voice Response
HFC	Hybrid Fiber Coax
XOC	Excellence In Operations Centers
XPM	Excellence In Plant Maintenance
NTF	No Trouble Found
SCTE	Society of Cable Telecommunications Engineers

Bibliography & References

[1] “What Is Agile Methodology? Benefits of Using Agile.” *Nvisia*, <https://www.nvisia.com/insights/agile-methodology>.

[2] Thibodeaux, Wanda. “People Trust Technology, to Be Honest, Study Finds. Here Are the Big Implications of That.” *Inc.com*, Inc., 8 Nov. 2019, <https://www.inc.com/wanda-thibodeaux/people-trust-technology-to-be-honest-study-finds-here-are-big-implications-of-that.html>.

[3] Inc., Prosci. “The PROSCI ADKAR[®] Model.” *Prosci*, <https://www.prosci.com/methodology/adkar>.

Using Profile Management to operationalize the roll-off region in DOCSIS 3.1 HFC networks

A Technical Paper prepared for SCTE by

Diana Linton

Network Engineer III
Charter Communications
14810 Grasslands Drive, Englewood, CO 80112
720-536-1027
Diana.Linton@charter.com

Justin Stiles

Principal Engineer I
Charter Communications
14810 Grasslands Drive, Englewood, CO 80112
720-629-4424
Justin.Stiles@charter.com

Esteban Sandino, Distinguished Engineer / Charter Communications

Roger Stafford, Principal Engineer III / Charter Communications

Vinod Dani, Principal Engineer I / Charter Communications

Jay Liew, Advanced Analytics Architect / Charter Communications

Keith Auzenne, Principal Engineer II / Charter Communications

Table of Contents

Title	Page Number
1. Introduction.....	4
2. HFC Lab Network Description.....	5
3. Deploying OFDM in the Roll-Off Region	6
3.1. Baseline testing: characterizing roll-off and MER impact using flat profiles	8
3.2. Understanding CM operation in the roll-off	12
3.2.1. Results for CM performance when the OFDM channel width is varied	12
4. Optimizing OFDM Performance for Roll-off Operation	14
4.1. Pre-PMA implementation baseline.....	15
5. PMA vs No-PMA	20
6. Conclusions.....	23
Abbreviations	25
Bibliography & References.....	25

List of Figures

Title	Page Number
Figure 1. Typical HFC plant using 1.2 GHz actives and 1.0 GHz passives.....	5
Figure 2. Cumulative Insertion Loss after four 1.0 GHz passives - two different vendors.....	6
Figure 3. Cumulative Return Loss after four 1.0 GHz passives - two different vendors.....	7
Figure 4. Roll-Off in the 1.0 – 1.2 GHz region at each of five selected tap locations across the network	7
Figure 5. Reported CM RxMER levels across a 192 MHz OFDM after three 1.0 GHz passives	10
Figure 6. Reported CM RxMER levels across a 192 MHz OFDM after seven 1.0 GHz passives.....	10
Figure 7. Reported CM RxMER levels across a 192 MHz OFDM after 11 1.0 GHz passives	11
Figure 8. Reported CM RxMER levels across a 192 MHz OFDM after 13 1.0 GHz passives	11
Figure 9. RxMER values across a 192 MHz OFDM after 20 1.0 GHz passives.....	12
Figure 10. D3.1 Avg. RxMER values per CM at three & seven 1.0 GHz passives for a 192 MHz OFDM channel.....	14
Figure 11.D3.1 Avg. RxMER values per CM at 3,7,11,13 & 20 passives for 192 MHz OFDM channel.....	16
Figure 12. RxMER per sub-carrier for CM 1 located at three 1.0 GHz passives.....	16
Figure 13. RxMER per sub-carrier for CM 3 located at 20 1.0 GHz passives	17
Figure 14. Reported RxMER per four CM vendors and used profiles after three 1.0 GHz passives	17
Figure 15. Reported RxMER per four CM vendors and used profiles after seven 1.0 GHz passives.....	18
Figure 16. Reported RxMER per four CM vendors and used profiles at 11 1.0 GHz passives.....	18
Figure 17. Reported RxMER per four CM vendors and used profiles at 13 1.0 GHz passives.....	19
Figure 18. Reported RxMER per four CM vendors and used profile at 20 1.0 GHz passives	19
Figure 19. Typical Throughput of OFDM at 1.0 GHz for Full Coverage of All Taps.....	20
Figure 20. Pre-PMA Baseline 4 profiles.....	21
Figure 21. CableLabs PMA tool – Four variable modulation profiles	22
Figure 22. Baseline vs PMA overall throughput.....	22
Figure 23. Median vs Standard Deviation for CM across the HFC plant.....	23

List of Tables

Title	Page Number
Table 1. Initial OFDM settings.....	9
Table 2. Configured Flat MER Profile Thresholds	13
Table 3. Optimized OFDM Settings	15
Table 4. PMA engine MER profile thresholds.....	21

1. Introduction

As high split deployments begin to ramp up, demand for broadband services keeps driving the need for higher network capacities. Operators are required to do plant hardware reconfigurations and upgrades across their Hybrid-Fiber Coaxial (HFC) networks. Plant hardware reconfigurations and upgrades are expensive. This is why implementation of techniques and strategies that minimize physical changes in the HFC networks is crucial to allow cable operators to optimize their HFC networks for the long-term, and to remain competitive.

DOCSIS 3.1 introduced features that not only allow operators to increase network capacity through their HFC networks, but also to optimize downstream and upstream transmissions in a more granular manner to obtain better performance. One of those features is Orthogonal Frequency-Division Multiplexing (OFDM), which is associated with higher modulation orders (up to 4096-QAM) and can be deployed over existing HFC plants to significantly increase network capacities, while deferring the need for immediate RF spectrum expansions. However, higher modulation orders do require higher Modulation Error Ratio (MER) and distortion performance, and not every HFC network will support the highest orders ubiquitously. Flat OFDM profiles only support one order of modulation across the entire OFDM channel and can be deployed initially to take advantage of higher MER conditions in specific areas, but do not necessarily result in optimal transport capacity. This is especially the case when plant impairments in specific areas of the HFC spectrum force the use of lower modulation orders across every single OFDM sub-carrier. Optimizing bandwidth utilization and transport capacity requires a dynamic mechanism to track MER performance for every OFDM sub-carrier, identify HFC plant impairments, and assign the most efficient modulation orders to one or more OFDM sub-carriers for a given set of plant conditions.

The Profile Management Application (PMA) is a cost-effective, software-based implementation to address the demand for more network capacity through data collection, analysis, and efficient selection of OFDM profiles. The selected profiles for each OFDM channel can assign different modulation orders to groups of subcarriers within the channel, known as segments, based on plant conditions. PMA not only will prove to be a valuable tool in evolving the network, but also will help MSOs improve network efficiency by maximizing modulation rates to and from each modem in the network, enabling higher user throughput overall.

This paper will present the results of the practical application of PMA in an HFC plant. The primary focus will be on OFDM operation in HFC plant roll-off areas as a practical mechanism to gain additional network capacity by increasing the modulation order and bit loading capacity of individual subcarriers for those CMs reporting higher RxMER values. The paper will discuss first the observed roll-off characteristics between 1.0 GHz and 1.2 GHz for a representative HFC plant segment configured with a mix of 1.2 GHz amplifiers and 1.0 GHz passives, and the worsening conditions impacting this portion of the HFC spectrum as the number of cascaded 1.0 GHz passives increases. It will discuss next the maximum expected MER and modulation orders at various tap locations throughout the cascade, and how an initial optimization of OFDM parameters for operation in roll-off areas using only flat modulation profiles can deliver limited capacity gains prior to a PMA deployment. Finally, the paper outlines how a full PMA approach is used to correctly identify physical channel conditions as seen by individual customer modems, and how this data is used to dynamically select the most appropriate OFDM profiles for optimal roll-off operation of modems connected throughout the cascade. Expected overall capacity gains when applying a PMA approach that enables efficient operation of OFDM carriers in roll-off areas will be discussed.

2. HFC Lab Network Description

Testing was performed in a typical node-based HFC network segment designed with a remote fiber optical node feeding a cascade of RF amplifiers, taps, and express, distribution, and drop cables. Total cable footage was 4,800 feet of RG-11 cable.

As seen in Figure 1, the express cables, illustrated in green, connect the optical node to the first (amp 1) and second (amp 2) 1.2 GHz system amplifiers. These express actives are Mini-Bridger (MB) type radio frequency (RF) amplifiers, designed to overcome cable attenuation and passive losses over sections of the network where there are few splitters and directional taps. The distribution portion of the HFC network, represented in blue, includes a combination of two line extenders (amp 3 and amp 5) and one MB amplifier (amp 4) to boost the RF levels to provide adequate signal level to multiple taps. In between the distribution amplifiers there are 1.0 GHz passives including: directional couplers (DC), power inserters (PI) and taps.

The drop portion of the network consisted of a 100-foot span of RG-6 cable from the tap port to the input of a two-way splitter, and 50 feet of RG-6 cable from the output of the two-way splitter to the Cable Modem (CM) input.

For the worst case scenario at the farthest end of line, there are a total of 20 1.0 GHz passives between the optical node and the CMs connected to the last tap.

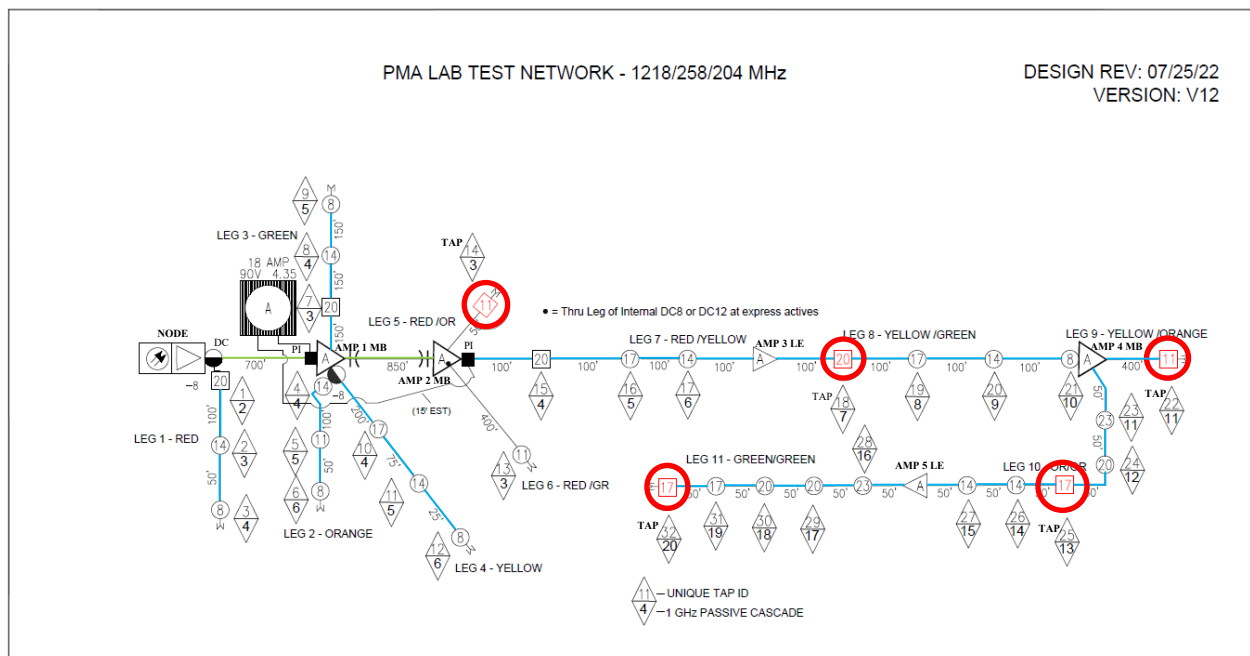


Figure 1. Typical HFC plant using 1.2 GHz actives and 1.0 GHz passives

The series of tests described below were performed using a high split Cable Modem Termination System (CMTS) that supports two downstream OFDM channels and two upstream OFDMA channels. Five sets of four different D3.1 CMs were connected to tap ports at five different tap locations, circled in red, across the network (Tap ID 14, 18, 22, 25 and 32). The 1.2 GHz actives in the test-bed were aligned to meet the designed RF output levels and tilt up to 1.2 GHz, but only loaded with a partial 1.0 GHz channel

line-up. This ensured that after completing the alignment, the addition of a 192 MHz OFDM channel in the 1.0 to 1.2 GHz range, required for the first round of testing, would not affect amplifier performance.

3. Deploying OFDM in the Roll-Off Region

The advantages of OFDM operation over traditional Single Carrier Quadrature Amplitude Modulation (SC-QAM) channels have been widely documented. OFDM has allowed cable operators to increase downstream efficiency and to transmit data more robustly. OFDM transmits data over a combination of orthogonal narrowband sub-carriers. Each sub-carrier can have a different modulation order (bit loading) within the OFDM channel as determined by the different segments configured within a modulation profile. The number of segments that can be configured within a profile will vary by CMTS vendor.

The ability to configure multiple profiles allows the OFDM channel to operate higher modulation orders in parts of the spectrum with minimal impairments. Conversely, OFDM sub-carriers can operate with lower modulation orders or profiles that are more robust in the presence of impairments such as frequency roll-off in the 1.0 to 1.2 GHz range, which is introduced by the accumulated non-flat signature of cascaded 1.0 GHz passives.

The frequency response curve of 1.0 GHz passives beyond their defined passband has different peaks and valleys. The non-flat frequency response of a single 1.0 GHz passive in the 1.0 – 1.2 GHz range can be considered fixed, but it does vary in severity by manufacturer. The combination of poor return loss and higher insertion loss creates a non-flat roll-off signature that increases in severity as the number of 1.0 GHz passives increases. The cumulative effects of higher insertion loss and poor return loss in the 1.0 to 1.2 GHz range after four passives in cascade, using passives from two different vendors, are illustrated below in Figure 2 and Figure 3.




Figure 2. Cumulative Insertion Loss after four 1.0 GHz passives - two different vendors




Figure 3. Cumulative Return Loss after four 1.0 GHz passives - two different vendors

As shown in Figure 4 the non-flat frequency response signature creates a roll-off that worsens as the number of 1.0 GHz passives increases. Eventually, the distortion of the OFDM carrier is such that CM operation may not be possible. RxMER data from individual CMs was no longer accessible after 11 passives, CMs stopped receiving data on the lowest modulation profile due to the amount of uncorrectable codeword errors. Ultimately CMs were unable to bond to the OFDM channel. As a result, the sequence of plots of the 192 MHz-wide OFDM channel across the five selected tap locations was obtained using a spectrum analyzer.




Figure 4. Roll-Off in the 1.0 – 1.2 GHz region at each of five selected tap locations across the network

Consequently, it is important to understand the overall impact of the roll-off on CM operation, and to select optimal OFDM parameters prior to enabling PMA in order to minimize issues such as those listed below and described in [PMA-D3.1-CL]:

- Profile flapping, which is defined as the process of a CM switching from a higher modulation profile to a lower modulation profile and back. It can be a result of the RxMER for the OFDM channel being marginal, falling slightly below the configured threshold for the CM to decode a certain modulation order. Therefore the CM starts to experience uncorrectable codeword errors on the assigned profile. The CM reacts by sending a CM-STATUS profile failure message to the CMTS, the CMTS response is to send data on a lower profile to the CM. If the RxMER increases over the configured threshold, then the CM will send a recovery CM-STATUS message to the CMTS, informing the CMTS that it can use a higher-modulation profile.
- Impaired channel operation due to loss of lock on the OFDM channel caused by high uncorrectable codeword errors on the assigned modulation profile. When loss of lock is

detected, the CM sends a loss-of-lock event message to the CMTS. The CMTS can either reprogram the OFDM channel or the profile.

- Partial channel mode occurs when, despite detecting lock loss, the OFDM channel is still partially functional for data reception, but uncorrectable codeword errors are high either in the PLC, NCP or profile A, and the CM tries to maintain use of the OFDM channel. When a CM enters partial channel mode the CMTS can either switch the CM to another profile on the OFDM channel or stop sending data on the reported profile.

HFC networks have different cascade depths that vary from node to node, as does the number of passives and type of passives that have been deployed throughout the years. Consequently the cumulative frequency response across different HFC plants will also vary. This is why it is important for MSOs to implement tools that help characterize and adjust to their network's performance. This ensures that subscriber quality of experience (QoE) is not impacted due to high volume packet loss and loss-of-lock events while operating an OFDM channel in the roll-off region. In addition, it is important to implement profiles with different modulation orders that account for the frequency response of the HFC plant to avoid significant network capacity reduction and/or intermittent connectivity issues caused by profile flapping. Further it will benefit MSOs to leverage extra network capacity by assigning a set of profiles based on the RxMER levels reported by different groups of CMs across different areas of the HFC spectrum.

Additional variables cable operators should consider when deploying and operating DOCSIS 3.1 (OFDM) channels in an HFC network with roll-off, due to the non-flat frequency response of 1.0 GHz passives in the 1.0 to 1.2 GHz range, are the configuration settings of the OFDM channel. A clear understanding of DOCSIS 3.1 configuration parameters is crucial to leverage stable CM operation of OFDM channels in the roll-off region. Parameters such as location of the OFDM channel, location of the PLC and NCP subcarriers within the channel and CM-STATUS messaging frequency can all affect the overall efficiency of the OFDM channel [PMA-D3.1-CL]. This impact will be discussed in more detail in the following sections.

3.1. Baseline testing: characterizing roll-off and MER impact using flat profiles

A first round of baseline measurements was taken using a DOCSIS 3.1 field meter. A 192-MHz OFDM channel was configured using default OFDM settings as shown in Table 1. No attempt was made to optimize these setting as the primary objective was to identify the impact on RxMER due to the non-flat frequency response above 1.0 GHz when mixing 1.0 GHz passives and 1.2 GHz actives across the HFC network.

The first round of testing was performed using the default OFDM settings as shown in Table 1. These parameters were re-configured as shown later in Table 3 to optimize CM operation in the roll-off region pre-PMA implementation. The latter approach is discussed in more detail in section 4.

Table 1. Initial OFDM settings

OFDM SETTINGS	PARAMETER
OFDM Channel	996-1188 MHz
PLC location	1004 MHz
Primary capable	No
Pilot-scale-factor	48
NCP Profile	16-QAM Flat
Max-event-hold-off	20 minutes

As mentioned in [OFDM RxMER], RxMER plots per sub-carrier are a valuable tool to characterize impairments such as roll-off. CMs will report via SNMP MIBS different RxMER values based on tap location and RF spectrum conditions across the network. Therefore an imperative first step in our process was to measure and record RxMER values using a typical DOCSIS 3.1 field meter. Not only did these measurements, collected manually, help quantify and map the impact of plant roll-off across a full 192-MHz channel spectrum and at various points in the cascade, but they could also be used to validate the accuracy of CM reported RxMER values.

For the initial measurements, the test bed was configured as follows:

1. Channel loading: 4 SC-QAMs below 1.0 GHz and a 192 MHz OFDM starting at 1.0 GHz;
2. OFDM settings were as described in Table 1;
3. Four flat profiles were configured, each supporting a single modulation order for all subcarriers across the entire OFDM channel.

RxMER values at different tap locations through the HFC network were recorded to identify RF signal degradation from 1.0 to 1.2 GHz due to the non-flat frequency response of the 1.0 GHz passives. As expected, the overall result is increasingly degraded CM performance (RxMER) as the number of cascaded actives and passives increases. In practical terms, this means that while CMs closer to the node can support 4096-QAM modulation across all sub-carriers within the OFDM block, CMs that are progressively farther away can only support 4096-QAM order of modulation over an increasingly reduced set of sub-carriers within the same block.

Recorded RxMER values at increasingly deeper locations within the cascade can be grouped by frequency blocks of varying widths across the entire 192 MHz OFDM channel. Each colored block then represents the maximum reported RxMER value within that frequency block that can be maintained by the CMs at a given tap location. The colored frequency blocks for the last tap in the cascade would illustrate the maximum RxMER values that could be maintained by the farthest end-of-line CMs connected to that last tap.

As seen in Figure 5, the RxMER values reported by four different CMs across a 192 MHz channel at the third 1.0 GHz passive (Tap ID 14) in cascade is within tenths of a dB and it is enough to support 4096-QAM operation across the full OFDM channel.




Figure 5. Reported CM RxMER levels across a 192 MHz OFDM after three 1.0 GHz passives

As seen in Figure 6, when the amount of 1.0 GHz passives in cascade is at seven, the roll-off region starts to affect a minor portion of the upper frequency edge of the 192 MHz OFDM channel. It was observed that CMs started to experience a few correctable codeword errors due to the lower RxMER values in the last 48 MHz of the OFDM channel, but the RxMER degradation of the affected sub-carriers did not have a meaningful impact on overall performance. RxMER was enough to support the highest modulation profile (4096-QAM).




Figure 6. Reported CM RxMER levels across a 192 MHz OFDM after seven 1.0 GHz passives

Referring to Figure 7, as the 192 MHz OFDM carrier travels even deeper into the cascade and the number of 1.0 GHz passives increased to 11 (Tap ID 22), the first 48 MHz block within the OFDM channel remains capable of supporting a MER of 40 dB for all sub-carriers within that block. As the roll-off starts to affect a higher percentage of sub-carriers in the OFDM channel RxMER progressively drops to 38 dB for the next 76 MHz block, to 36 dB for the next 20 MHz block, to 30 dB and to 26 dB for the next two 10 MHz blocks respectively. Ultimately RxMER for the sub-carriers of the last 28 MHz portion of the 192 MHz OFDM channel sank to 18 dB. At this 11th tap, CMs were not able to bond to the OFDM channel and stopped receiving data on the lowest modulation profile due to the amount of uncorrectable codeword errors.




Figure 7. Reported CM RxMER levels across a 192 MHz OFDM after 11 1.0 GHz passives

At 13, 1.0 GHz passives (Tap ID 25) as visually illustrated in Figure 8, the first 48 MHz block within the OFDM channel remain capable of supporting an MER of 40 dB for all sub-carriers within that block, which should enable modulation orders of up to 4096-QAM . The second 26 MHz block was not significantly affected by the roll-off and can support an MER of 38 dB. However, as a higher number of sub-carriers are affected by the roll-off the MER values started to progressively drop to 32 dB and 30 dB for the third and fourth 40 MHz blocks. Eventually the fifth 10 MHz block can only support an MER of 20 dB and the last 28 MHz block of the OFDM channel can only support an MER of 13 dB across the entire network. As a result, none of the CMs were able to bond to the OFDM channel.




Figure 8. Reported CM RxMER levels across a 192 MHz OFDM after 13 1.0 GHz passives

At 20, 1.0 GHz passives in cascade at the EOL tap (Tap ID 32), Figure 9 shows only the first 24 MHz block still able to support a MER of 39 dB, but the severity of the roll-off significantly reduced the usable bandwidth of the OFDM channel. As a result the MER values decreased at a higher rate, reducing to 37 dB for the next 24 MHz then to 33 dB and 27 dB for the next 6 and 10 MHz blocks respectively. After MER values sank rapidly to 17 dB for the next 10 MHz block. Ultimately the MER for the last 118 MHz started to drop below 10 dB.




Figure 9. RxMER values across a 192 MHz OFDM after 20 1.0 GHz passives

This distribution of MER values across frequency blocks within a 192 MHz OFDM channel across different tap locations begins to uncover the value of a PMA approach. Without PMA, only flat modulation profiles are possible across the RF spectrum. Referencing Figure 5 through Figure 9, a flat modulation profile based on 4096-QAM operation that leverages the MER value distribution of Figure 5 can only be used by those few CMs closer to the node (Tap labeled 14). Likewise at seven passives in cascade (Tap labeled 18), 4096-QAM could be supported even though some sub-carriers started to be affected by the roll-off as shown on Figure 6. However CMs at the farthest end of line and reporting a distribution of MER values similar to that shown in Figure 9 reported a maximum RxMER that is considerably below the threshold for 4096-QAM. These CMs failed to operate with a flat 4096-QAM profile even though the subcarriers in the first 48 MHz block within the OFDM channel would support it.

3.2. Understanding CM operation in the roll-off

There are different ways for a CMTS to check the CM performance on a downstream OFDM channel. In the case of this paper, the CMTS sends an OFDM Downstream Profile Test Request (OPT-REQ). The OPT-REQ is used to test the CM ability to receive the specified OFDM profile and query the CM RxMER statistics. The CMTS uses the data from the OFDM Downstream Profile Test Response (OPT-RSP) to decide which OFDM profile is better for the CM at the time of collection. After this, the CM can still fail any available OFDM profile for other reasons such as the ones mentioned on [MULPIv3.1].

3.2.1. Results for CM performance when the OFDM channel width is varied

The objective of this test was to determine the tap location and RxMER threshold at which CM would start downgrading OFDM profiles. In addition, a throughput test was performed at the selected tap locations to quantify the additional throughput obtained above 1.0 GHz when using OFDM channels of varying width. All CMs were monitored as the roll-off severity increased across the test plant.

1. Channel loading: 4 SC-QAMs below 1.0 GHz and a 48 MHz wide OFDM starting at 1.0 GHz;
2. The same four flat profiles from the previous section were configured across the entire OFDM channel;
3. Tap location, profile, and the value of RxMER were recorded;
4. Throughput measurements were taken;
5. The OFDM bandwidth was increased in 6 MHz steps from 48 MHz to 192 MHz, and steps 3 & 4 were repeated.

As the OFDM bandwidth increased, CMs at tap locations further from the node were not able to achieve stable operation on the reported profile when the reported RxMER was between the threshold values showed in Table 2.

Table 2. Configured Flat MER Profile Thresholds

MODULATION RATE (QAM)	FLAT MER [dB] PROFILE THRESHOLDS
64	20.5
128	23.5
256	31
512	29.5
1024	33
2048	36
4096	38

For instance, when the RxMER was slightly below or above a configured RxMER threshold value (Table 2), as the upper end of the OFDM channel increased to about 1145 MHz, at 11 1.0 GHz passives in cascade (tap 22), all four CMs locked to the OFDM channel but CMs started to experience uncorrectable codeword errors. When CMs experienced enough codeword errors, the CMs sent a CM-STATUS message to the CMTS to communicate a profile failure on profile 2048-QAM due to the amount of uncorrectable codeword errors in the profile. The CM was reporting an RxMER slightly below 36 dB, which is the configured threshold for profile 2048-QAM. It was observed that CMs took a long time to drop to the next profile. After analyzing the CM's debug log it was found that the CM-STATUS message was not sent in a timely manner. The CMTS default setting for the max-event-hold-off was set to twenty minutes. See Table 1. This parameter was lowered to five minutes to allow the CMs with a marginal RxMER to switch to a lower profile faster.

Similarly, for the four CMs located farther from the node 13 passives into the cascade (tap 25), as the OFDM bandwidth increased beyond 1110 MHz some CMs reported RxMER below the 1024-QAM threshold. Other CMs were able to operate on 1024-QAM with some correctable errors, but when the upper end of the OFDM channel was increased up to 1150 MHz CMs were not able to bond to the OFDM channel due to the amount of uncorrectable codeword errors reported on the 1024-QAM profile. When the OFDM bandwidth was increased in 6 MHz steps between 1110 to 1150 MHz, CMs began to experience profile flapping between profiles 2048-QAM and 1024-QAM.

At 20, 1.0 GHz passives in cascade (tap 32) the cumulative effect of the 1.0 GHz passives worsened, and with a 48 MHz-wide OFDM channel some CMs started to experience uncorrectable codeword errors. These CMs notified the CMTS of a profile failure on 4096-QAM due to the RxMER being below the 38 dB MER configured threshold. This behavior was not as expected because we predicted that CMs connected at 20 1.0 GHz passives in cascade would work with a 48 MHz OFDM channel on the highest modulation profile as shown in Figure 9. In addition, when the OFDM channel width was increased beyond 66 MHz, all CMs went into partial channel mode or partial service mode due to high FEC errors in the used profile (1024-QAM) but continued trying to use the OFDM channel. These CMs sent a CM-STATUS message to inform the CMTS of a loss-of-lock event. Some CMs were able to fall back to the

lowest modulation profile 256-QAM, but due to the amount of uncorrectable errors were not able to lock to the OFDM channel.

During this test, the main observation was that only CMs closer to the node (less or equal than seven passives deep, tap 18) were able to operate using a full 192 MHz OFDM channel, and did not experience uncorrectable codeword errors because the severity of the roll-off was not affecting as many sub-carriers in the OFDM channel. As shown in Figure 10, all four CMs connected to the third and the seventh 1.0 GHz passive (tap 14 and tap 18 in Figure 1) reported a similar RxMER of approximately 40 dB. The CMTS sent traffic on the highest modulation profile and the overall throughput was not impacted. Due to the CM operational issues just described, CM RxMER data was not measured at 11, 13 and 20 passives in cascade.




Figure 10. D3.1 Avg. RxMER values per CM at three & seven 1.0 GHz passives for a 192 MHz OFDM channel

4. Optimizing OFDM Performance for Roll-off Operation

It can be seen from section 3 that it is crucial to revise certain OFDM configuration parameters before PMA implementation in the roll-off region to ensure stable operation. The OFDM settings mentioned previously in Table 1 were optimized in Table 3. Optimized settings can help minimize Partial Channel Mode, impaired OFDM channel, and profile flapping events similar to those reported in this paper. Optimization will also help ensure that a CM is assigned the most appropriate modulation profile for specific plant conditions such as severe roll-off.

Table 3. Optimized OFDM Settings

OFDM Setting	Optimized Parameter
OFDM Channel	996-1188 MHz
PLC Location	1004 MHz
Pilot-scale-factor	120
NCP Profile	16-QAM from 996 to 1100 MHz and 0 QAM from 1100 MHz to 1188 MHz.
Max-event-hold-off	5 minutes

After analyzing the flat OFDM profile configuration and CM behavior it was determined that NCP-profile should be dynamically adjusted based on the roll-off characteristics of the plant. Also the pilot scale factor was set to the highest value (120). As described in [PMA-D3.1-CL] the CMTS defines modulated sub-carriers with a particular modulation pattern as pilots in the downstream. All the CMs in the system know this to allow interoperability. The pilot scale factor increases the amount of continuous pilots that occur at fixed frequencies in every symbol, as shown in the following formula:

$$\text{Number of Continuous Pilots} = \min \left[\max \left(8, \text{ceil} \left(M * \left(\frac{F_{\max} - F_{\min}}{192^6} \right) \right) \right), 120 \right]$$

The value of M is the pilot scale factor in the equation and can be adjusted at the CMTS between ($120 \geq M \geq 48$). The pilot factor was adjusted from $M = 48$, which results on a total of 56 pilots (48 + 8 PLC pilots) for a 192 MHz channel to $M = 120$, which equals to 128 pilots (120 + 8 PLC pilots) for 192 MHz to improve OFDM downstream channel estimation.

4.1. Pre-PMA implementation baseline

The objective of this test was to establish a new baseline after adjusting OFDM settings to ensure the maximum amount of CMs were able to use the OFDM channel in the roll-off region across the entire HFC network pre-PMA implementation:

1. Channel loading: 4 SC-QAMs below 1.0 GHz and a 192 MHz OFDM at 1.0 GHz;
2. The following 4 profiles were configured across a 192 MHz OFDM channel:
 - a. Profile 0: 256-QAM from 996 MHz to 1050 MHz and 0 bit loading from 1050 MHz to 1188 MHz
 - b. Profile 1: 256-QAM flat
 - c. Profile 2: 1024-QAM flat
 - d. Profile 3: 4096-QAM flat
3. OFDM was configured as described in Table 3;
4. RxMER at different tap locations throughout the HFC network was recorded to identify RF signal degradation from 1.0 to 1.2 GHz due to the non-flat frequency response of the 1.0 GHz passives.

As seen in Figure 11, after adjusting OFDM settings to the values in Table 3, CMs were able to operate in the roll-off region with a 192 MHz OFDM beyond tap 18 (seven 1.0 GHz passives in cascade). Consequently, all CMs located at 11 (tap 22) and 13 (tap 25) passives in cascade were able to bond to the OFDM channel after initialization based on the configured CMTS MER thresholds in Table 2. At the end of line (EOL) tap 32 only CM 3 was able to bond to the OFDM channel. CM 1, CM 2 and CM 4 were not able to bond to the OFDM.

Figure 11 quantifies the positive impact the changes made in the OFDM settings had in CM operation throughout the network. Most CMs were able to successfully bond to a 192 MHz channel after the changes, and only two CMs did not report RxMER at the EOL tap. The root cause for the failure to bond for these two CMs is still under investigation.




Figure 11.D3.1 Avg. RxMER values per CM at 3,7,11,13 & 20 passives for 192 MHz OFDM channel

At three (tap 14) 1.0 GHz passives in cascade, the average (mean of all supported sub-carriers) RxMER value of CM 1, represented in blue in Figure 11, was obtained from the expanded RxMER plot per sub-carrier shown in Figure 12. Correspondingly, the RxMER of CM 3 located at 20 (tap 32) 1.0 GHz passives, represented in green in Figure 11, was calculated from the RxMER plot per sub-carrier illustrated in Figure 13.




Figure 12. RxMER per sub-carrier for CM 1 located at three 1.0 GHz passives




Figure 13. RxMER per sub-carrier for CM 3 located at 20 1.0 GHz passives

At three (tap 14) and seven (tap 18) 1.0 GHz passives in cascade all CMs reported an average RxMER close to 40 dB allowing the CMs to use all profiles and operate with the highest profile and the highest modulation order (4096-QAM) as shown in Figure 14 and Figure 15.




Figure 14. Reported RxMER per four CM vendors and used profiles after three 1.0 GHz passives




Figure 15. Reported RxMER per four CM vendors and used profiles after seven 1.0 GHz passives

After configuring 0 bit loading on the NCP-profile in the severe roll-off region, the ability of CMs to bond to the OFDM channel was greatly improved. For the Pre-PMA baseline, as seen in Figure 16 all CMs located at 11 (tap 22) 1.0 GHz passives were able to bond to the OFDM. An interesting observation during this test was that different D3.1 CMs reported slightly different RxMER values. CM 3 and CM 4 stopped receiving data on the highest profile three due to uncorrectable codeword errors. The CMTS responded to the CM-STATUS message by downgrading the profile to a lower profile in this case profile 2.




Figure 16. Reported RxMER per four CM vendors and used profiles at 11 1.0 GHz passives

Even though the cumulative effect of the non-flat frequency response of the 1.0 GHz passives was worse at 13 1.0 GHz passives, all CMs still bonded to the OFDM channel. However, higher differences in the

reported RxMER values from all four CMs were observed. CM 1 and CM 3 were able to receive data on profile one, but CM 2 and CM 4 were able to only use profile zero due to the amount of uncorrectable errors on profile 1. See Figure 17. The differences in reported RxMER values among CM manufacturers have not yet been explained at the time of this writing. Further testing needs to be performed to research possible causes.




Figure 17. Reported RxMER per four CM vendors and used profiles at 13 1.0 GHz passives

As the non-flat frequency response of the 1.0 GHz passives accumulated the severity of the roll-off worsened. As a result CM 1, CM 2 and CM 4 located at tap 32, after 20 1.0 GHz passives, were not able to use any of the configured profiles. See Figure 18. However CM 3 was able to receive data on profile 0, a major observation is that it appears that certain D3.1 CM manufacturers do not report RxMER values if the channel is in Partial Service/Channel Mode. This can potentially be of concern for a PMA implementation since profile generation is based on CM RxMER reported data.




Figure 18. Reported RxMER per four CM vendors and used profile at 20 1.0 GHz passives

In summary, as shown in Figure 19, some CM manufacturers were able to achieve a maximum throughput of approximately 1.569 Gbps after up to 11 passives in cascade, while other CMs were able to achieve a maximum throughput of 1.343 Gbps at 11 passives. With optimized OFDM settings more CMs were able to operate with higher flat modulation profiles at the same tap location in the network compared to typical OFDM settings. At the cascade EOL, at 20 1.0 GHz passives in cascade only one CM was able to operate on profile 0 while the other CM were impaired on the OFDM channel.




Figure 19. Typical Throughput of OFDM at 1.0 GHz for Full Coverage of All Taps

5. PMA vs No-PMA

PMA alone does not allow operation in the roll-off region, particularly for CMs further away from the node that are severely affected by the cascading effect of the non-flat response of the 1.0 GHz passives. We must use a combination of optimized OFDM settings to get as many CMs online as possible before PMA is implemented.

CM performance testing has proven that CMs can operate modulation orders with MER thresholds below the specified D3.1 MER thresholds [MULPIv3.1]. Hence the thresholds used in the PMA engine to generate the variable modulation profiles were lowered and were less conservative compared to the MER thresholds used during the baseline testing described in section 3.2. As a result, more CMs were able to operate in higher order modulations at the selected tap locations. The lowered MER thresholds are listed in Table 4.

Table 4. PMA engine MER profile thresholds

MODULATION RATE (QAM)	PMA MER [dB] THRESHOLDS
64	19
128	22
256	25
512	28
1024	31
2048	34
4096	37

Figure 20 illustrates the four modulation profiles used during the Pre-PMA baseline testing performed in section 4.1. Profile 0 is deliberately non-flat because, as seen in Figure 9, the severity of the roll-off is such that 256-QAM can only be supported up to 1050 MHz. Since profile 0 must be usable by all CMs, zero-bit loading was configured for the rest of the subcarriers in the OFDM channel.




Figure 20. Pre-PMA Baseline 4 profiles

Figure 21 shows the four variable modulation profiles generated using the modified CableLabs PMA tool based on the reported RxMER data from 18 out of 20 CMs connected to the five tap locations mentioned previously in Figure 1. The MER thresholds used to generate the four variable modulation profiles in Figure 21 are listed in Table 4. These profiles were converted to CMTS commands and applied to the CMTS.




Figure 21. CableLabs PMA tool – Four variable modulation profiles

PMA allowed us to create custom-made profiles consisting of segments of different modulation orders based on plant conditions over the frequency range of the OFDM channel. In the case of impairments such as roll-off, the profiles are adjusted to follow the non-flat frequency response of the 1.0 GHz devices. As a result most CMs were able to bond to a 192 MHz OFDM channel.

Figure 22 contrasts the difference in overall throughput by applying four flat modulation profiles (Figure 20) vs. the four variable modulation profiles obtained from the PMA tool (Figure 21). By using the four variable modulation profiles generated by the modified CableLabs PMA engine, some CMs were able to operate at higher speeds across the network than when using flat profiles. For instance, as shown in Figure 22, at 11 passives there was a 4% increase in average throughput across four CMs. At 20 1.0 GHz passives, there was a 35% increase in average throughput.




Figure 22. Baseline vs PMA overall throughput

There is a close correlation between the mean RxMER reported by the five four-CM groups at five tap locations across the HFC network and the results previously presented in sections 3 and 4. As the cumulative effect of the non-flat frequency of the 1.0 GHz passives worsens, an increasing number of sub-carriers in the OFDM channel are affected and mean RxMER decreases creating five distinct clusters that correspond to the five selected tap locations in the HFC network as shown in Figure 23. The RxMER data reported by the four CMs located at tap 14 (three 1.0 GHz passives), circled in yellow, resembles the mean RxMER values of Figure 14. Similarly, the average RxMER values

circled in lilac, correspond to the average RxMER values reported by the four CMs located at tap 32 (20 1.0 GHz passives in cascade) and illustrated in Figure 18.



Figure 23. Median vs Standard Deviation for CM across the HFC plant

PMA is a valuable tool that generates customize profiles with different modulations (segments) based on impairments CMs experience across different HFC networks. PMA not only helps MSOs increase the reliability of their networks, but it can leverage the operation of an OFDM channel in HFC networks with severe roll-off. PMA implementation provides extra network capacity that can be used to alleviate traffic during peak hours. Further testing will be performed to continuously implement the benefits that PMA can provide.

6. Conclusions

- When operating in the roll-off region, there will be significant differences between the average reported RxMER values across the HFC network at different tap locations. As the number of 1.0 GHz passives increased, the severity of the roll-off increased, and usable modulation orders decrease due to the lower RxMER values. This was expected since the non-flat frequency response of the taps worsened as the number of passives in cascade increases.
- The amount of uncorrectable codeword errors on the OFDM channel operating in the roll-off region increased as the severity of the roll-off affected a higher number of sub-carriers across the network. As a result, CMs farther from the node started experiencing loss of lock to the OFDM channel. This was alleviated by changing OFDM configuration parameters such as NCP-profile and pilot-scale-factor, allowing more modems to be online at the EOL tap. It also helped some modems to use higher modulation profiles across the network. This also reinforces the need for optimization of OFDM parameters prior to PMA deployments.
- When using flat profiles, the non-flat frequency response of legacy 1.0 GHz passives in the 1.0 to 1.2 GHz range creates challenges as a result of profile flapping and partial service for D3.1 CM operating in the roll-off. Uncorrectable codeword errors were recorded for the OFDM channel as the severity of the roll-off increased across the network.
- The severity of the frequency response (peaks and dips) will vary by tap manufacturer, and will worsen as the number of passives in a cascade increases. As a result, successful reception of an OFDM channel will be limited to a maximum cascade depth depending on deployed passive

manufacturer. Since cascade depth varies from node to node, as does the number of passives and the nature of the induced roll-off, attempting to create custom OFDM modulation profiles manually can be time consuming and yield unpredictable results. A PMA approach will greatly simplify operation in the roll-off regions.

- PMA provides operators the ability to implement proactive and adaptive network operations in a network with inevitable impairments such as roll-off. The benefits include reduced trouble calls, a higher throughput, and the ability to scale and deploy incrementally. However, should an operator decide to operate OFDM in roll-off regions, the OFDM carrier should be configured/optimized first to ensure the highest possible performance prior to PMA implementation. Once optimal performance is achieved pre-PMA, higher and stable capacity gains can be obtained through the implementation of a fully automated PMA solution.
- When using PMA, CMs were ranked by RxMER quality, they naturally fell into clusters or groups that closely aligned with the tap locations selected and the severity of the experienced impairment, as shown in Figure 23. This natural clustering helps ensure that the set of profiles created by the PMA tool are optimized and targeted to the specific degree of impairment severity as seen by the CMs across the HFC network.
- Cable modems experienced the effect of the roll-off differently across the network. The impact of the roll-off region on cable modem operations was observed as a degradation of the RxMER values of the sub-carriers operating in the roll-off region. But by configuring the profiles generated by the PMA engine, the CMTS assigns customized profiles based on the type of collected RxMER per groups of cable modems, thus minimizing transmission errors on the network and maximizing the overall network capacity with up to a 35% gain at EOL tap compared to a non-PMA solution.
- The maximum throughput of approximately 1.569 Gbps over the OFDM and SC-QAM channel combination used for this testing was not achieved consistently, unless the CM is in close proximity to the node and within a limited number of passives in cascade (depends on the profile signature of the passive) and the CM manufacturer.
- Higher network capacity gains may be possible if the number of modulation profiles increases beyond four. Additional testing is planned to explore the overall network capacity gains in the near future with six and eight dynamic modulation profiles.

Abbreviations

CM	Cable Modem
Gbps	Giga bits per second
FEC	forward error correction
HS	high split
MHz	Mega-Hertz
PMA	Profile Management Application
SCTE	Society of Cable Telecommunications Engineers
RxMER	Received Modulation Error Ratio

Bibliography & References

[MULPIv3.1] DOCSIS 3.1, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I17-190121, January 21, 2019, Cable Television Laboratories, Inc.

[PMA-D3.1-CL] Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)

[OFDM RxMER] Characterizing Network Problems Using DOCSIS 3.1 OFDM RxMER Per Sub-carrier Data.

[D3.1 PMA-CC] DOCSIS 3.1 Profile Management Application

[BL-PMA] Making the Most of Your HFC Network, Bit by Bit

[Full-Scale-PMA] Full Scale Deployment of PMA

Validating Access Network Maps for Plant Extensions and Capacity Upgrades

A Technical Paper prepared for SCTE by

John Schlack

Distinguished Engineer
Comcast Cable
1800 Arch St
Philadelphia, PA 19103
+1 267-592-8063
John_Schlack2@comcast.com

Rina Hayashi

Senior Director, Product Management
Comcast Cable
1800 Arch St
Philadelphia, PA 19103
+1 267-562-9437
Rina_Hayashi@comcast.com

Table of Contents

Title	Page Number
Table of Contents	2
1. Introduction.....	4
2. Pre-Vet Data Collection and Analysis	5
2.1. System Overview	5
2.2. RF Network Graph	8
2.3. Linking Addresses to the Graph.....	10
2.3.1. Addresses Inside Node Boundary	11
2.3.2. Addresses for RF Taps Outside Node Boundary	13
2.3.3. Addresses Associated with Node Segments	13
2.3.4. Linking Addresses to RF Network.....	15
2.4. Linking Devices to RF Network Graph.....	18
2.5. Data Aggregation	19
2.6. Bus Leg to Node Segment Identification.....	22
2.7. Opportunities for Map Improvement	23
2.7.1. Unconnected Subscriber.....	23
2.7.2. Address Not Serviceable	23
2.7.3. Unconnected RF Tap.....	23
2.7.4. Incorrect Node Segment.....	23
2.7.5. Node Segment to Bus Leg Mismatch	23
2.7.6. Outside Node Boundary.....	24
2.8. Pre-Vet Analysis Results.....	24
3. Maps and Data Output	24
4. Conclusion.....	28
Abbreviations	29
Bibliography & References.....	Error! Bookmark not defined.

List of Figures

Title	Page Number
Figure 1 - Simplified Plant Map	4
Figure 2 - Adding Serviceable Addresses and Device Counts	5
Figure 3 - Using Device Locations to Identify Map Improvement Opportunities.....	6
Figure 4 - Pre-Vet System Diagram.....	7
Figure 5 - Example RF Network Graph.....	9
Figure 6 - Correlating RF Network Graph with Drafted Network	10
Figure 7 - Node Segments to Bus Leg Association	11
Figure 8 - Address Locations within Node Boundaries.....	12
Figure 9 - Drafted Supports and RF Taps Relative to Service Addresses	15
Figure 10 - Address to Support/Tap Assignment by Proximity and House Count.....	16
Figure 11 - Example Data Model to Associate Addresses to Supports/Taps.....	18
Figure 12 - Address Locations in RF Network Graph	18
Figure 13 - Aggregating Device Counts at the RF Output Port of Amps and Splitters.....	19
Figure 14 - Depth First Search to Calculate Device and Homes Passed Counts	20

Figure 15 - Map of Plant with Bus Legs Highlighted and Address Errors.....	25
Figure 16 – Device and Homes Passed Counts at Amp in Node Boundary View.....	26
Figure 17 – Device and Homes Passed Count at Amp (Focused View)	27
Figure 18 - Bus leg table	27

List of Tables

Title	Page Number
Table 1 - Example Port to Port Connectivity from RF Flood Trace.....	8
Table 2 - Example Service Address Information ¹	13
Table 3 - Example Node Segments for HSD Devices ²	14
Table 4 – Example Device and Homes Passed Count by Aggregator ³	22

1. Introduction

Operators of Hybrid Fiber Coaxial (HFC) have a long history of delivering continuous advancements in speed, capacity and performance. With the rise of DOCSIS 4.0 and 10G, operators are poised to take the next huge leap in delivering faster speeds, greater capacity, advanced security, and improved reliability. Anytime we update the HFC plant, we must conduct a thorough analysis of the drafted plant to decide whether the existing maps are accurate and if so, where and how plant upgrades should be made. This document will refer to this analysis as “pre-vet.” While HFC plant maps are a reliable source for beginning a pre-vet, changes to the area (new streets, buildings, plant repair, etc.) may not always make it into the plant map in time for the analysis. Layering additional data onto the plant maps can further enrich the data and increase the level of accuracy for the designer to make an informed decision on where capacity upgrades and plant extensions should take place.

The pre-vet process begins by capturing the node boundary of the existing plant map that has been documented. **Figure 1** shows a very simplified example of an HFC plant map. Information such as street addresses, parcel size, and where the node and other HFC equipment is located is represented on the map. While this helps show where the equipment is and the type of equipment used, it does not address capacity concerns on its own.

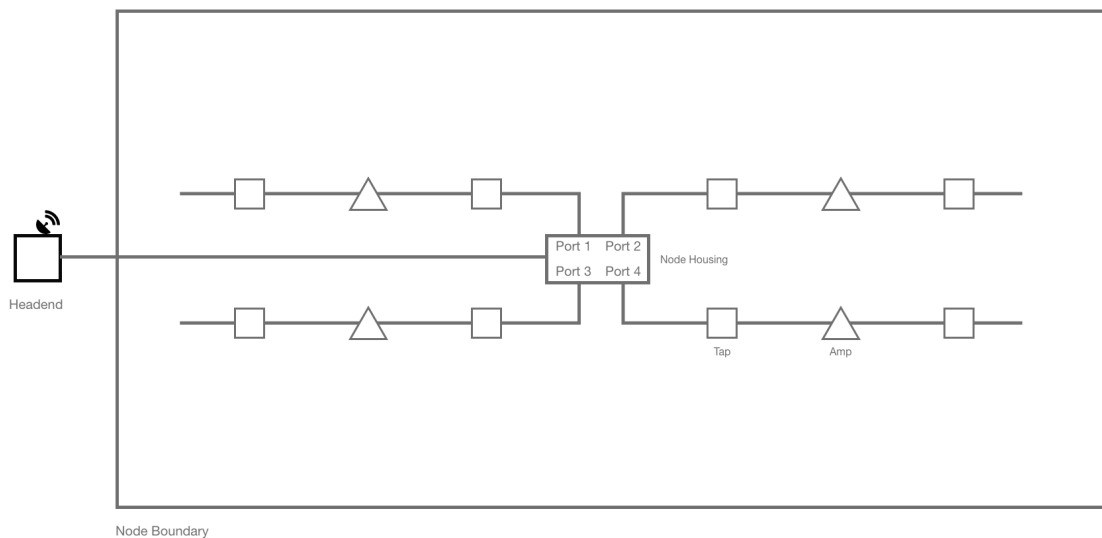


Figure 1 - Simplified Plant Map

Next, as shown in **Figure 2**, serviceable addresses and device counts connected to the CMTS are added to the plant map. The combination of the two data points indicates to the designer where the strain is on the network. By aggregating this information at the amplifier (amp) level, the designer can pinpoint where on the bus leg to place a new node to alleviate the strain on the old node. In addition, this level of data can identify gaps in the plant maps. For example, each leg should serve around 50-500 homes. If the aggregated data shows more than that it is a clear indication that either the plant map requires adjustment, or the serviceable addresses associated with the plant requires investigation regarding accuracy.

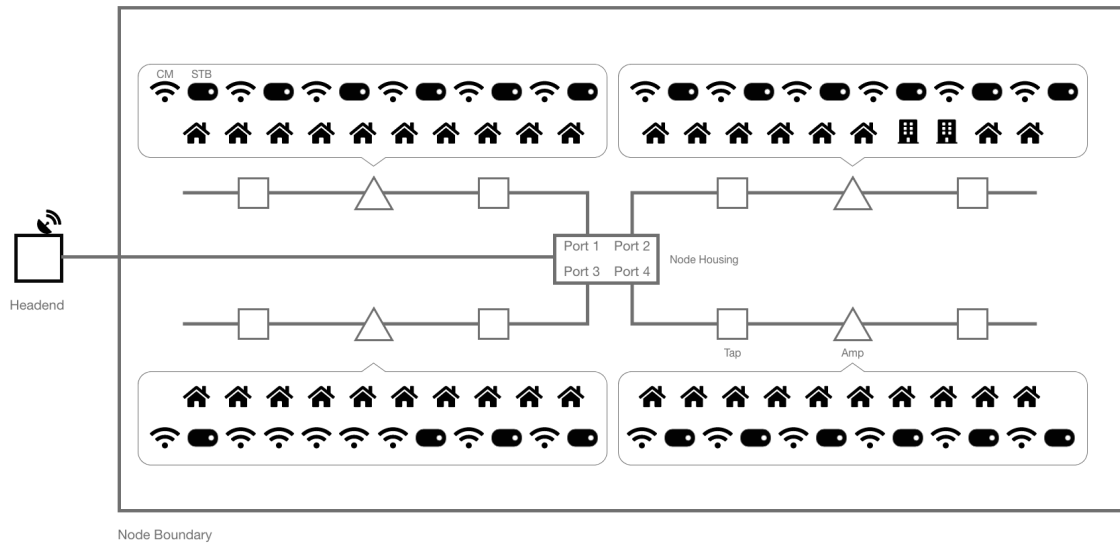


Figure 2 - Adding Serviceable Addresses and Device Counts

Historically, creating a pre-vet has been a manual, labor-intensive process, taking around 2-8 hours to create the package for a single node housing. As the need increases for upgrading the HFC plant, it will be pertinent to make this process more efficient by automating the aggregation of the serviceable addresses and device counts and visualizing the data on the plant map. This paper will discuss the data collection, analysis, and automation of creating a pre-vet package.

Note that the network maps and data shown in this paper do not include any names or personal data about subscribers.

2. Pre-Vet Data Collection and Analysis

2.0. System Overview

Subscriber service location (address) and device network communication location correlated against the drafted RF network map provides clues about the accuracy of the maps. Clusters of service locations that cannot be connected to the drafted network indicate opportunities for improving the drafted map. The example network shown below (which is NOT actual data) shows a cluster of service locations that cannot connect to the drafted network. An orange circle highlights these service locations. Notice that some service locations below the circled addresses reside outside the node boundary, but within acceptable distance to the drafted network.

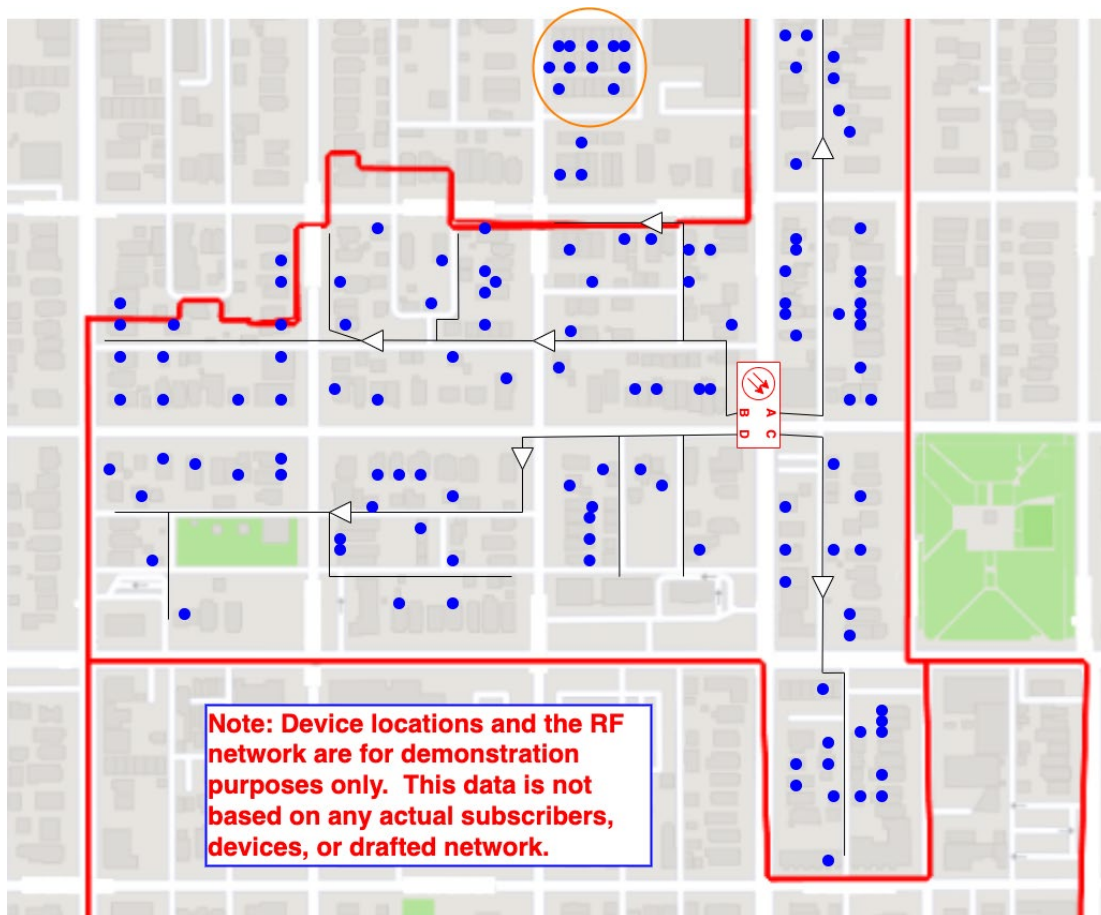


Figure 3 - Using Device Locations to Identify Map Improvement Opportunities

Once the drafted RF network map has been validated, the pre-vet process calculates homes passed and actual device counts at key locations throughout the RF network map to aid the designers in node segmentation or node split analysis.

To perform an automated pre-vet analysis of a node housing, data must be collected, correlated, and analyzed from several sources including:

- Design and drafting platform;
- USPS address database;
- Subscriber account database; and
- Device telemetry.

The figure below shows an example system for performing the pre-vet operation.

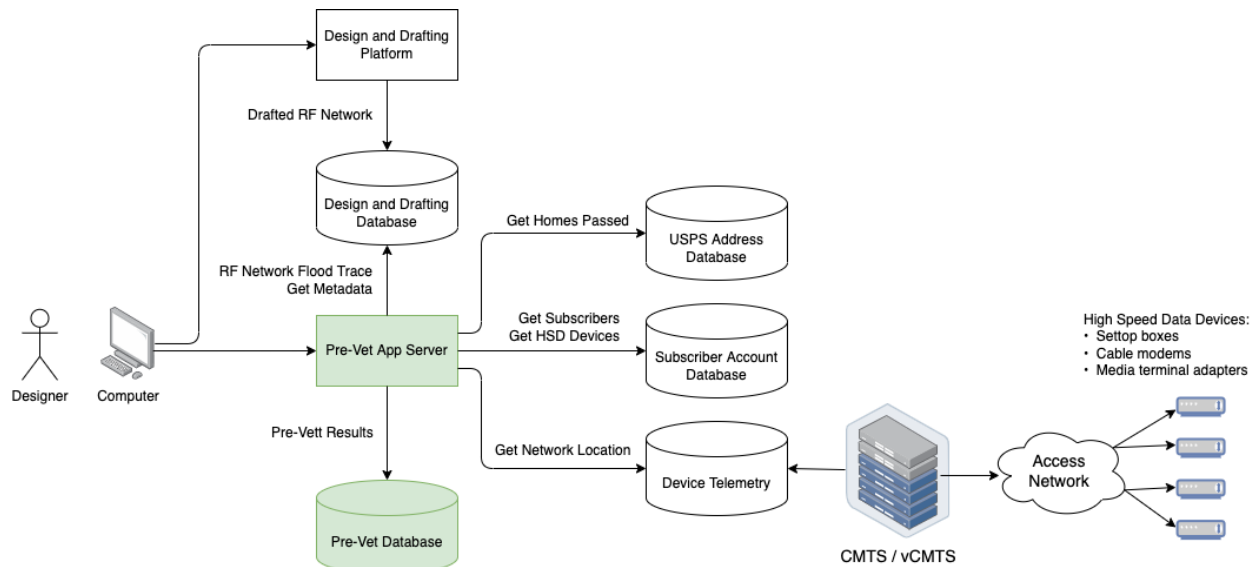


Figure 4 - Pre-Vet System Diagram

The design and drafting platform maintains the physical details of the support structure, cables, and equipment for the access network. This platform maintains connectivity information between elements, provides maps of the access network, and performs power and loss calculations. The drafted data may link the access network equipment to service address locations, but this linking is inconsistent and may be outdated.

The USPS address database contains the official USPS addresses and their geolocations (latitude and longitude). This database minimally must contain addresses for premises passed by the MSO access network.

The account database stores the subscriber account information, including service address and list of devices provisioned to the account. Each subscriber account links to a service address.

Device telemetry identifies the presence of communication with high-speed data (HSD) capable devices, such as cable modems (CM), set-top boxes (STB), and media terminal adapters (MTA). Telemetry data provides the set of ports on the cable modem termination system (CMTS) or remote PHY device (RPD) where messages from the HSD devices are received. The set of subscribers that communicate upstream to a specific CMTS or RPD port are described in this paper as a "node segment."

The steps to collect and correlate this data include:

- Build a graph representation of the RF network for a specific node housing.
- Retrieve the service location addresses based on the node boundary.
- Determine node segments if not already known.
- Retrieve service location addresses associated with node segments.
- Associate service locations to RF taps.
- Count homes passed and devices at amplifiers and splitters.
- Identify service address or network gaps.

Details for these steps appear in the following sections.

2.1. RF Network Graph

To analyze the RF network for a given node housing, one needs to understand the relationship between the equipment and cables downstream of the node housing. The RF network shall be represented as a graph data structure, which is a series of vertices and edges. A graph provides efficient mechanisms for determining connectivity between elements and also provides searching capabilities. Some example uses for the RF network graph include counting devices serviced by each amplifier and determining max amplifier cascade by bus leg.

Vertices in the graph represent the cables and equipment (node housing, RF cables, taps, amplifiers, splitters, etc.) while edges represent connectivity between the vertices. Notice that the edges do NOT represent cables. Edges are connectivity. A vertex associated with an RF cable will have two edges, representing other RF equipment or cables connected at each end of the cable in question. The vertex for a node housing will contain 4 edges, one per bus leg (RF output port). For the RF network analysis, the fiber network connecting to the node housing does not need to be represented in the graph. The RF network graph will have a root node that is the node housing.

Building the RF network graph is dependent on the design and drafting platform used by the MSO. In general terms, performing an RF flood trace on each bus leg of a node housing will specify a list of elements and connected ports. An example of port connectivity is shown below.

Table 1 - Example Port to Port Connectivity from RF Flood Trace

Source			Destination		
Type	ID	Port	Type	ID	Port
CoaxCable	DMTARCMC1853065972850	N/A	Splitter	DMTARPMC18536598	Input
Splitter	DMTARPMC18536598	Output-1	CoaxCable	DMTARCMC1853065991823	N/A
Splitter	DMTARPMC18536598	Output-2	CoaxCable	DMTARCMC1853065964851	N/A
CoaxCable	DMTARCMC1853065991823	N/A	Tap	DMTARTMC18536621	Input

Notice that port level connectivity means that most RF equipment or cables will have multiple entries in the RF flood trace details – one for each input and output port. The method to build the graph must account for this duplication so that only a single vertex is created for each RF element. Ports identified in the RF flood trace shall be annotated in the edges of the graph.

The graph building algorithm will also create a hash map that maintains the equipment ID to vertex relationship. This will be called the equipment map. The equipment map provides a method to quickly find the vertex of any RF element without requiring a graph search.

The algorithm to create the RF network graph and equipment map is specified below:

- Create vertex for the node housing and add the vertex as the root of the graph.
- Load the node housing metadata from the design and drafting system.
- Add node housing metadata to vertex.
- Add the ID and vertex of the node housing to the equipment map.
- For each RF output port on node housing:
 - Perform RF flood trace on RF output port; and
 - For each connection pair in RF flood trace:
 - Search equipment map using the equipment source ID and destination ID from the connection pair to find the equipment vertices.

- For each vertex not found:
 - Create vertex for the equipment and add to graph;
 - Load equipment metadata from the design and drafting system;
 - Add equipment metadata to vertex; and
 - Add the equipment ID and vertex to the equipment map.
- Create an edge between equipment vertices.
- Add the port IDs and port types of the equipment to the edge.

Port types can include coax ports, power ports, and drops. Tracking the port type is important since the pre-vet analysis will not analyze power and thus power ports can be ignored. Although shown in the figure below, the power port edges can be pruned from the graph.

An example RF network graph is shown below for node housing “NJHB00750”. Each of the four bus legs can be seen connected to the root node of the graph, which is the node housing.




Figure 5 - Example RF Network Graph

To better visualize the correlation of the RF network graph with the drafted plant data, the figure below focuses on bus leg A from the figure above. The below figure shows a side-by-side comparison of the RF network graph and the drafted network. Each RF cable or RF equipment element is represented by a

vertex in the graph. Notice that the power supply is missing from the graph. This may be a drafting issue indicating that the power cable running to the power supply is not actually connected.

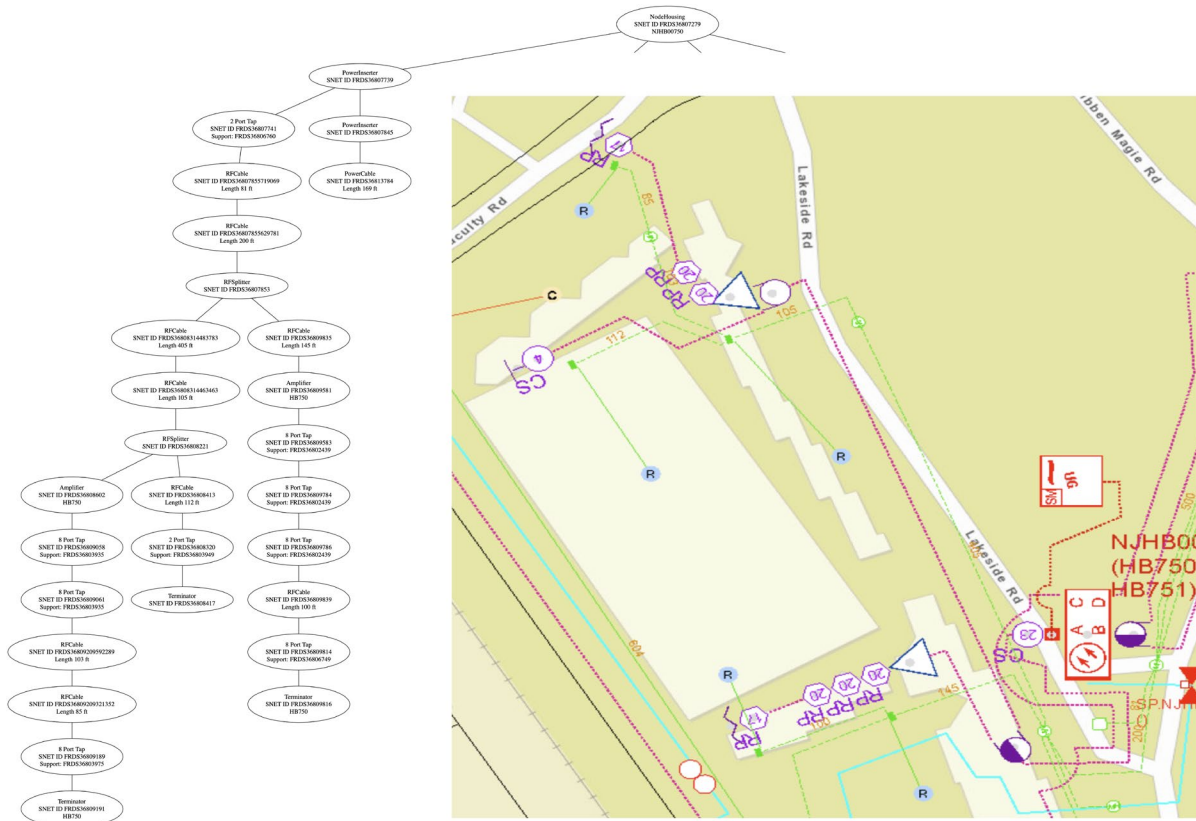


Figure 6 - Correlating RF Network Graph with Drafted Network

2.2. Linking Addresses to the Graph

An important method to evaluate the accuracy of the drafted RF network is to examine the subscriber locations and the high-speed data (HSD) capable devices on the network. HSD devices communicate with the cable modem termination system (CMTS). The upstream communication for a set of subscribers – the “node segment” – is typically on a specific CMTS port for a physical CMTS or remote PHY device (RPD) port in the case of a virtual CMTS (vCMTS). A node segment may be comprised of one or more bus legs from the node housing. The figure below shows an analog 2x2 node housing, with two upstream node segments.

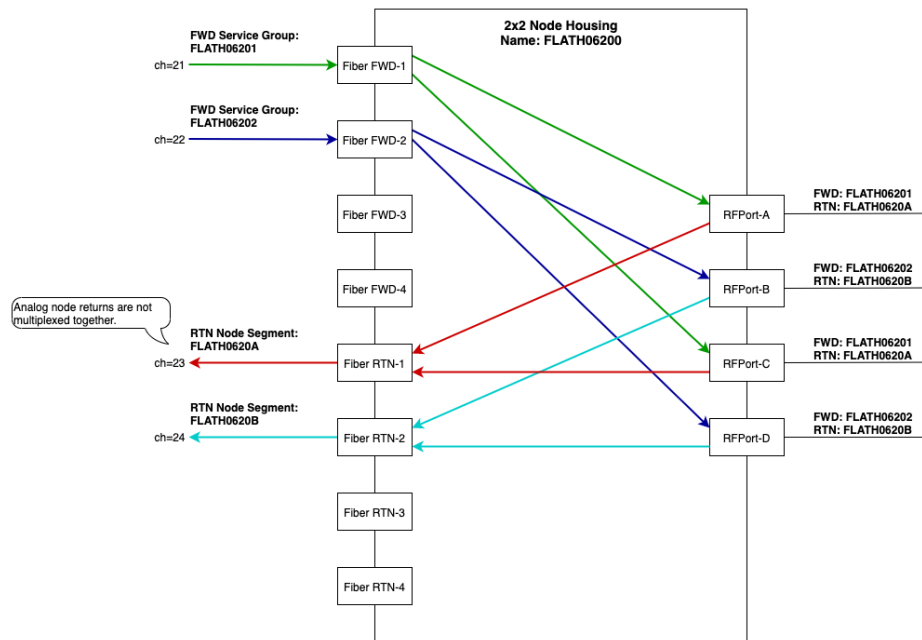


Figure 7 - Node Segments to Bus Leg Association

The node segment for an HSD device can be determined when the device communicates with the CMTS or RPD. Through the provisioning process, HSD devices are linked to subscriber accounts. These accounts have service address locations. Under ideal conditions, the drafted network includes drops that link the RF taps to drafted address locations. These drafted address locations exactly match the service locations assigned to subscribers and their associated provisioned devices. Thus, a direct link can be made from the drafted RF tap to the subscriber HSD devices and the associated node segments. Unfortunately, there are several factors that could prevent that direct association:

- Constructed network has changed since the design was drafted (e.g. plant extension or capacity upgrade).
- No drop was drafted to link the RF tap to the address location.
- Address has changed since the network was originally drafted (e.g. street renamed).
- Originally drafted address was not the official USPS address, but a place holder for new construction (e.g. drafted network has lot numbers instead of USPS addresses).

The pre-vet analysis will perform several steps to overcome these conditions. These steps include querying address locations based on the node boundary, querying addresses associated with node segments of devices, and linking address locations to RF taps by proximity.

2.2.1. Addresses Inside Node Boundary

Designers draft boundaries to indicate the coverage area for each node housing. Most address locations serviced by the RF network of a node housing reside within the drafted boundary. Therefore, querying the drafted boundary provides most of the service addresses needed for the network analysis. Note that depending on the address database that the MSO queries, the address locations may be homes passed (e.g.

USPS address database) by the RF network or subscribers' service addresses (e.g. subscriber account database).

Designers may draft distinct types of node boundaries including node housing, node segment, and bus leg boundaries. One or more bus leg boundaries should comprise the node segment boundary. Node segment boundaries comprise the node housing boundary. Only one type of node boundary needs to be queried to retrieve the service addresses associated with the node housing. Querying the service address locations within the node's drafted boundaries and associating those addresses with active devices will allow discovery of the node segments associated with the node housing. The figure below shows address locations of homes passed (not subscribers) for a set of node housing boundaries.




Figure 8 - Address Locations within Node Boundaries

An example Oracle query to retrieve the service address locations of homes passed within a node boundary is shown below:

```
SELECT A.ID, A.STREET_NUM, A.STREETNAME, A.UNIT_TYPE, A.UNIT_VALUE, A.CITY,
       A.STATE, A.POSTALCODE, A.LATITUDE, A.LONGITUDE
FROM ADDRESSES A
WHERE
SDO_RELATE (
  A.LOCATION,
  SDO_GEOMETRY(
    2003,
    4326, -- WGS84 coordinate system
    NULL,
    SDO_ELEM_INFO_ARRAY(1,1003,1), -- one polygon (exterior polygon ring)
    SDO_ORDINATE_ARRAY(
      <LIST OF COORDINATES>
    ),
    ),
    'mask=anyinteract' ) = 'TRUE'
```

where <LIST OF COORDINATES> are the node boundary polygon coordinates. Coordinates are specified as longitude and latitude pairs.

Example:

```
SDO_ORDINATE_ARRAY(
  -87.28482973277373, 37.25299509730186,
  -87.27747974176278, 37.25299509730186,
  -87.27747974176278, 37.24625992935859,
  -87.28482973277373, 37.24625992935859,
  -87.28482973277373, 37.25299509730186 )
```

Note that some addresses in this query may not actually be serviced by the node housing in question. Drafted node boundaries are not precise and RF networks from different node housings may overlap such that it is extremely difficult to draft boundaries that are completely distinct. Additional processing of the data, discussed below, will isolate and remove address locations associated with other node housings.

Fields returned in the query include a unique ID assigned to the address, street number, street name, unit type and value (for multi dwelling units), city, state, postal code, latitude, and longitude. The latitude and longitude are based on the geocoder for the address database, not the drafted location.

Table 2 - Example Service Address Information¹

ID	Number	Street	Unit Type	Unit Value	City	State	Zip	Latitude	Longitude
2198101	101	Main St.	Apt.	A	Philadelphia	PA	19103	39.9547	75.1685
36198354	222	First Ave.			Philadelphia	PA	19103	39.9621	75.0981

2.2.2. Addresses for RF Taps Outside Node Boundary

The design and drafting platform will have the latitude and longitude of the RF taps (and the location of the support structure where the RF tap connects). Query all RF taps within the node boundary used to query the service addresses. In an ideal case, all RF taps from the RF network graph are in the query results. However, it is possible that some RF taps in the graph are outside the drafted node boundary. It is also possible that some RF taps from the RF network of other node housings are in the query results. Save these “extra” RF taps for use in the proximity-based address to RF tap correlation step below.

On the graph, RF taps that are significantly outside the node boundary will require an additional address query to capture the associated service addresses. Drops that connect the RF tap to an address location are typically limited to about 300 feet. To account for white space in the drafting process, if the location of the support structure where the RF taps are attached is more than 500 feet outside the drafted node boundary, then an additional address query needs to be executed to capture the locations.

For each RF tap outside the node boundary, query all addresses within a 500-foot circle of the support structure on which the RF tap is connected. Add these addresses to the addresses from the node boundary query.

2.2.3. Addresses Associated with Node Segments

To identify missing plant extensions from the drafted map, it is not sufficient only to examine the addresses in the drafted node boundary or addresses associated with drafted RF taps since the data

¹We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

returned can only be associated with the drafted access network. The address search needs to expand to include any address that has devices in the node segments associated with this node housing.

Using the assumption that the drafted network is at least partially accurate, one can use the service addresses queried from the map to obtain the set of node segments associated with the node housing. Some of the queried addresses passed by the drafted network will be actual subscribers. These subscriber accounts link to provisioned devices, whose telemetry includes the node segment on which the device reports.

Join the addresses retrieved above with the subscriber account database to identify subscriber accounts among these addresses. Find the device IDs of the provisioned HSD devices for the accounts and use that to retrieve the node segment from telemetry data. These queries are MSO specific. Below is an example table with the joined account to device to node segment data. Note that the data below is not actual customer data and was created only for demonstration purposes.

Table 3 - Example Node Segments for HSD Devices²

Account ID	Device ID	Device Type	Node Segment
335056384	A3:98:31:C9:ED:71	CM	FLD0010A
242012225	22:3E:C1:B5:CC:8E	CM	FLD0010C
263370322	B7:C0:39:D2:AC:4F	STB	FLD0010D
619042765	FA:45:22:13:D2:1C	CM	FLD0010A
390653365	6F:4A:70:D6:E8:44	MTA	FLD0010B
133703447	E5:7B:B2:84:0A:FC	CM	FLD0009D
423544086	9A:16:8C:CB:0E:B5	STB	FLD0010C
195661518	37:7E:08:47:DC:F4	STB	FLD0010B
444325235	0E:03:81:4C:60:D5	CM	FLD0010B
386919575	AA:6C:D7:AE:61:FA	MTA	FLD0010D
339519483	FD:00:9A:C0:04:9E	CM	FLD0010A
110301259	9F:F7:FD:E6:0B:5C	STB	FLD0010A

In an ideal case, the result should be between 1 and 4 node segments. There cannot be more than 4 node segments since the node housing has a maximum of 4 bus legs. Note that some of the address locations may be associated with other node segments and may have been retrieved since node boundaries of other node housings are adjacent or overlap. If more than 4 node segments are found, use the 4 instances that occur most frequently. In the example table above, the node segments FLD0010A, FLD0010B, FLD0010C, and FLD0010D appear most frequently.

It is still possible to identify more node segments than are associated to the node housing. For example, the node housing may actually contain only 2 node segments, but 4 node segments were identified due to node boundary overlap with adjacent nodes. These extra node segments will be removed when individual bus leg to node segment association is calculated below.

²We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

Given the set of node segments, query the device telemetry data for the full set of HSD devices that communicate on these node segments. Then use the device provisioning data to determine the service addresses for these accounts.

Add these addresses to the previously collected addresses.

2.2.4. Linking Addresses to RF Network

Addresses must be associated to RF Taps in the Graph to find map inconsistencies, calculate homes passed and device counts, and find the node segment per bus leg. In an ideal case, the service addresses from the USPS database or subscriber account database can be connected to the addresses specified in the design and drafting platform, which are then linked to the RF taps by drops on the drafted RF network. However, in many cases, drop details are missing or the address specified in the drafted network does not match USPS addresses or subscriber account addresses. Another method must be used to link addresses to the RF taps in the RF network.

The algorithm outlined below uses proximity of the address geolocation to the support structure on which the RF tap is attached. Support structure locations are typically more accurate than the RF equipment locations, which may be offset for white space management purposes. The figure below shows a drafted set of support structures and RF equipment relative to the parcels of service addresses.




Figure 9 - Drafted Supports and RF Taps Relative to Service Addresses

Each address location shall be linked to the closest support structure location that has capacity for the address if the distance does not exceed 500 feet. Note that long drop lengths are typically limited to 300 feet. However, allowances must be made for the geocoder assigned location of the address not matching the actual drop location of the building or the drafted position of support structures not being exact.

The figure below shows proximity-based assignment of addresses to support structures associated with RF taps. Each parcel contains a cross with an “R,” which is the geocoder assigned location of the address. For comparison, the estimated drop location of the building is the plain cross. The pink circle, square, and hexagon connected by a pink dashed line are RF Taps, representing taps with 2, 4, and 8 ports, respectively. The small red circles with “X” connected by solid red lines are support structures (poles and strand).

The RF taps are physically located on the adjacent support structure. Notice that the two-port tap in the center of the image has two connections. There is an address below it which should connect to the two-

port tap based solely on shortest distance. However, since the RF Tap already has both ports occupied, the address must connect to the more distant RF tap to its "northeast."




Figure 10 - Address to Support/Tap Assignment by Proximity and House Count

The proximity algorithm requires a sorted list of distances for each combination of address and support structure with an RF tap. The algorithm uses this list to assign the next address based on shortest distance. The pseudo-code to track distances is shown below:

- Create an empty distance list that will store objects with the triplet of distance, address ID, and support ID.
- For each combination of address and support:
 - Calculate the distance from the address lat/lon to the support lat/lon; and
 - If the distance is less than max drop distance (e.g. 500'), add the triplet of distance, address ID, and support ID to the distance list.
- Sort the distance list by distance value in ascending order.

Support structures may contain multiple RF taps. It does not matter how the addresses are distributed across these RF taps if the RF taps are on the same bus leg. This is because the goals are determining address assignment per bus leg and determining the device and homes passed counts aggregated at amplifiers and splitters.

Each support will have a maximum number of addresses that can be connected. This will be the smaller of the total output ports of the RF taps on the support and the house count associated with the support. For example, a support may have 2 RF taps - one with 2 ports and one with 4 ports. The support may also have a house count of 5. The maximum number of addresses that can be connected is 5, since the total RF tap port count is 6, but the house count is 5. The converse may also occur. If the house count was 7, then the maximum number of addresses would be 6 since there are only 6 ports available between the 2 RF taps on that support.

The algorithm to assign addresses to supports by proximity will iterate over the sorted list of distance calculations. If the address has not yet been assigned to a support, and the referenced support has space for another address, the address shall be assigned to the support. The available capacity for the support is decreased by one.

There is no actual pruning of addresses from the distance list when assigning the address to an RF Tap. Each address appears multiple times in the list. Removing each address would require sequential traversal of the list to find all address instances. It is more efficient to traverse the distance list once and keep track of assigned addresses in the hash map of assigned IDs.

The following pseudo-code assigns addresses to RF taps based on support distance and capacity:

- Create an assigned hash map to store the list of assigned addresses. It is simply a lookup table whose key is the address ID but contains no data.
- For each element in the distance list (distance-address-support triplet) traversed in order of ascending distance:
 - Search assigned hash map. If a match is found, go to the next element in the distance list.
 - If support does not have capacity (assigned address count \geq max address count), go to next element in the distance list.
 - Find an RF tap for this support with available ports (connected address count less than port count):
 - Assign the address ID to this RF tap;
 - Decrease the support's available capacity by one; and
 - Add the address ID to the assigned hash map.

It is possible that some addresses cannot be assigned when there is insufficient RF tap port or house count capacity. It is equally possible that some RF tap ports do not have assigned addresses.

Below is an example of the data model used by the algorithm:

- **Distance List** - List of distances between an address and a support. List is sorted by distance in ascending order. Every combination of address and support distance is included in this list if the distance is less than the maximum drop distance.
- **Address Hash Map** – Stores the service addresses retrieved in the boundary and node segment queries above. The hash map uses the Address ID as a key. It contains the geolocation of each service address and may also store the node segment (if determined by device association).
- **Support Hash Map** – Stores the support structures on which the RF taps reside. The hash map uses the support ID as a key. Each support structure entry contains the list of RF taps on this support, maximum address count, and assigned address count.
- **RF Tap Hash Map** - RF tap lookup table using RF tap ID as a key. Contains all RF taps downstream of a node housing. Each RF tap entry stores the port count and the addresses associated with its support based on distance. The assigned address count must not exceed the port count.

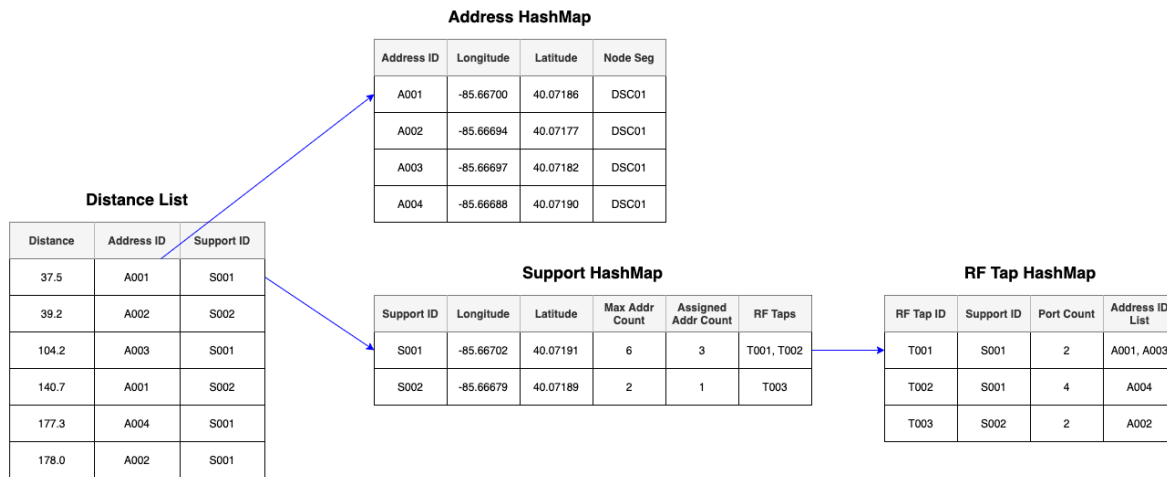


Figure 11 - Example Data Model to Associate Addresses to Supports/Taps

Use the results in the RF tap hash map to update the RF network graph with the RF tap to service Address association. Create a new vertex for each assigned address and create an edge between the address Vertex and the RF tap vertex.

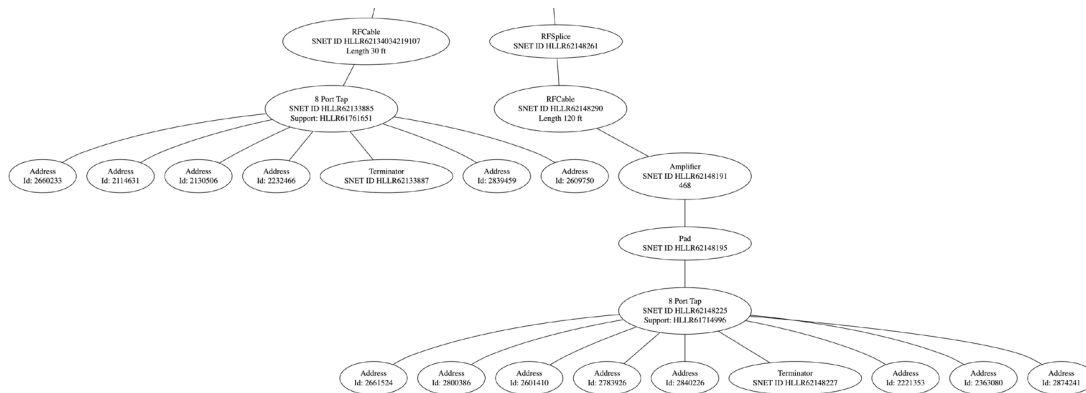


Figure 12 - Address Locations in RF Network Graph

2.3. Linking Devices to RF Network Graph

With the service address locations linked to the appropriate RF taps in the RF network graph, it should be possible to link the service locations to subscriber accounts and obtain the provisioned HSD devices for those accounts. A similar operation occurred above when identifying the possible node segments for the node housing. Remember that not all homes passed are subscribers and therefore only a subset of the service locations will have associated HSD devices.

The provisioned HSD devices may be added to the RF network graph as vertices connected to the address vertices. However, for the Comcast implementation, the devices were simply added to the address vertex itself as additional data.

Note that the graph contents are never provided to the designer since it identifies individual subscriber service locations. The device count data will be aggregated and anonymized before being presented to the user.

2.4. Data Aggregation

HSD device counts and homes passed counts are aggregated at each of the node housing RF ports as well as output RF ports for amplifiers and splitters. The equipment where the data is aggregated shall be referred to as “aggregators” in this document. The device and homes passed counts provide the designer a view into the distribution of devices and service locations across the network, which will assist in the node segmentation or node split operation. Each of these elements will track both “local” counts and the “cumulative” counts. The local counts are the device and homes passed count from the specified aggregator to the next downstream aggregator for each of the output ports. The cumulative counts are the device and homes passed count from the specified aggregator to all downstream components and branches.

The figure below shows the port level data aggregation. Focusing on amplifier A1, there are two output ports: port-1 and port-2. The local counts are based on the device and homes passed counts on all RF taps from output port to the next aggregator (amplifier A2 and Splitter S1 for A1 port-1 and A1 port-2, respectively). The cumulative counts for a given output port is the summation of device and homes passed counts of all RF taps downstream of the port.

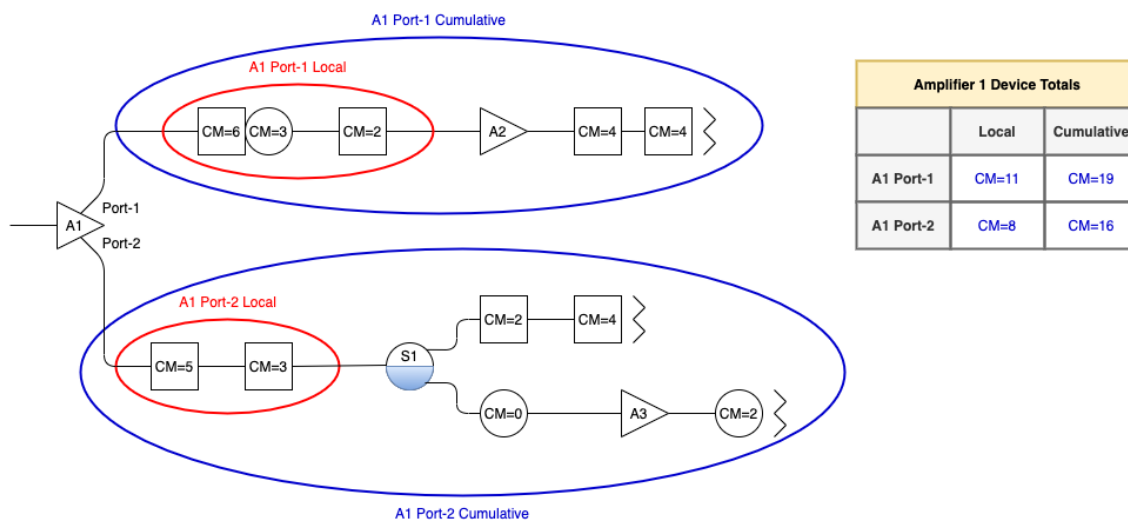


Figure 13 - Aggregating Device Counts at the RF Output Port of Amps and Splitters

MSOs can divide HSD device counts into multiple categories. Comcast uses 3 categories for device counts: cable modem (CM), set-top box (STB), and media terminal adapter (MTA). Homes passed can also be categorized as well, such as residential or commercial properties.

The algorithm to calculate the device and homes passed counts of each aggregator in the RF network uses a depth first search (DFS) to visit each of the vertices of the graph. A stack data structure keeps track of

the visited aggregators during the search. The last aggregator placed on the stack will be used for updating local counts. All aggregators on the stack will be used for calculating cumulative counts.

The DFS executes for each bus leg on the node housing. It visits each vertex in the graph, traversing as deeply as possible along one path before backtracking to visit alternate branches in the graph. When a DFS passes in the downstream direction through the RF output port of an aggregator, the RF output port of that aggregator gets pushed onto the stack. Conversely, when the DFS passes in the upstream direction through an RF output port on an aggregator, that RF output port pops off the stack.

When the DFS reaches an RF tap, it will traverse into the connected addresses. Each address contributes to the aggregated homes passed. If the address has associated provisioned HSD devices, these devices contribute to the aggregated device counts. The aggregator at the top of the stack will have its local counts updated. All aggregators in the stack will have the cumulative counts updated.

The figure below shows an example bus leg for a node housing. It includes a set of amplifiers, splitters, RF taps, and addresses. The RF equipment highlighted in red shows the elements traversed in a depth first search and the resulting stack at a point in the DFS when visiting vertex T11.

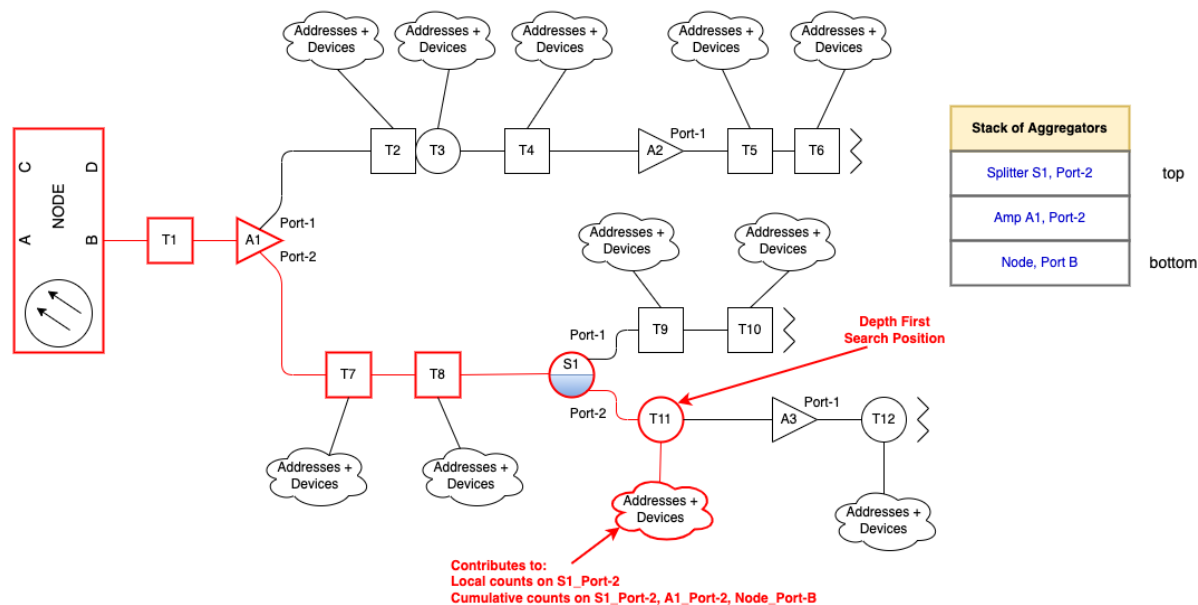


Figure 14 - Depth First Search to Calculate Device and Homes Passed Counts

The DFS began at Bus Leg B on the node housing. This RF output port is the first entry on the stack (bottom of the stack). The DFS proceeds to visit vertices T1 and A1, following the edge connected to Amp A1 Port-2. Amp A1 Port-2 shall be added to the stack since amplifiers are aggregators. The DFS proceeded to vertices T7, T8, and through Splitter S1 Port-2. Splitter S1 Port-2 shall be added to the stack since splitters are aggregators.

When the DFS traverses the address vertices connected to T11, each address contributes to the local homes passed count of S1 Port-2. Each address contributes to the cumulative homes passed counts for all elements on the stack, which include S1 Port-2, A1 Port-2, and Node Port B. Any provisioned HSD devices associated with the addresses will contribute similarly to the elements on the stack.

The DFS can also be used to calculate total amplifier counts and maximum amplifier cascade per bus leg. Each time an amplifier is encountered in the DFS, the amplifier total for that bus leg is incremented and the number of amplifiers on the stack are counted to update the maximum cascade count.

The pseudo-code for the local and cumulative device and homes passed count calculations are shown below. There will be an initialization function to kick off the processing of each bus leg. Aggregation begins at the node housing's RF output ports, which define each bus leg.

- For each node housing RF output port (bus leg), perform the following steps:
 - Create a "visited" hash map that tracks each of the vertices that have been visited.
 - Add node housing vertex to the "visited" hash map.
 - Create a stack that tracks RF port of aggregators visited as part of the DFS.
 - Push the node housing's RF output port onto the stack.
 - Set max cascade to 0.
 - Set amplifier count to 0.
 - Call ProcessVertex function for the vertex connected to this RF output port of the node housing.
 - Save max cascade and amplifier count data for this bus leg

Depth first search is performed by a recursive function that (1) skips visited nodes, (2) aggregates data from addresses connected to RF taps, and (3) manages the stack contents based on RF ports of aggregators that are visited.

- ProcessVertex function (SourceVertex)
 - Add SourceVertex to "visited" hash map.
 - If SourceVertex is a power supply, return without processing its connections.
 - If SourceVertex is an address:
 - For tracking local counts, add device and homes passed count from this address to the top element on the stack; and
 - For tracking cumulative counts, add device and homes passed count from this address to all elements in the stack.
 - If SourceVertex is an amplifier:
 - Increment amplifier count.
 - For each vertex connected to the SourceVertex:
 - If connected vertex is in the "visited" hash map, go to next connected vertex.
 - If SourceVertex is an aggregator:
 - Get RF output port from edge between this vertex and SourceVertex;
 - Add RF output port to aggregator stack;
 - Count the number of amplifier RF output ports on the stack;
 - If amplifier count on stack > max cascade,
 - max cascade = amplifier count on stack.
 - Call ProcessVertex function for this vertex.
 - If SourceVertex is an aggregator,
 - Pop stack.

The “visited” hash map is required so that searches do not loop back on themselves. Although this should not happen, it is good to have logic to prevent that case.

After executing the algorithm, the bus legs of the node housing and the RF output ports of amplifiers and splitters will have local and cumulative counts of devices and homes passed. The maximum number of amplifiers and maximum amplifier cascade will be known by bus leg.

The table below shows example device and homes passed counts for the aggregators in on Bus Leg D of a node housing:

Table 4 – Example Device and Homes Passed Count by Aggregator³

Type	Device ID	Port	Local				Cumulative			
			Homes	CM	STB	MTA	Homes	CM	STB	MTA
Node	UTAARMC18121634	D	0	0	0	0	96	39	43	69
Amp	UTAARAMC18121825	1	4	2	1	3	55	23	21	39
Amp	UTAARAMC18121825	2	3	1	2	2	41	16	22	30
Split	UTAARPMC18121923	1	2	1	1	1	21	8	8	17
Split	UTAARPMC18121923	2	0	0	0	0	30	11	12	19

Local device counts and homes passed counts may be zero if there are no RF taps between aggregators.

2.5. Bus Leg to Node Segment Identification

The set of node segments for the node housing has been discovered based on the service address locations linked to the RF network graph and the association of provisioned devices to accounts. The next step is to determine the actual node segment assignment by bus leg. All devices on the same bus leg must communicate with the same CMTS or RPD port. Therefore, they will share the same node segment.

Proximity based assignment of address locations to RF Taps is not deterministic, providing the chance for errors. Therefore, the node segment assignment to each bus leg shall be determined based on the most common node segment of the devices on the bus leg. This can be accomplished using a depth first search of each bus leg, counting each occurrence of a node segment for the devices associated with accounts on that bus leg. Details of the depth first search algorithm are provided above. In fact, this calculation can occur at the same time as the device aggregation.

Device telemetry indicates the node segment on which each device communicates. With devices associated with accounts, and the service address of the accounts linked to the RF taps in the RF network graph, one can calculate the node segment for the bus leg.

The pseudo-code to calculate the node segment assigned to a bus leg is shown below. Perform these steps for each bus leg of the node housing:

- Create hash map for node segments; key is the node segment, value is the count.
- Perform a depth first search on the RF output port of the node housing:
 - If vertex is an address, then for each device:
 - Use the node segment of the device to search the hash map;
 - If a match is found, increment the count for that node segment;
 - Else (no match found) add the node segment to the hash map with a count of 1.

³We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

- Set the node segment of the bus leg to the node segment with the highest count in the hash map

The node segment assignment to the bus leg shall be used to identify areas of inconsistencies on the map.

2.6. Opportunities for Map Improvement

After assigning node segments to each bus leg and linking homes passed or service addresses to the RF taps in the access network, one can examine each of the addresses to determine whether there is any mismatched data. Common exception conditions are outlined below.

It is possible that multiple exception conditions may occur per service address location. All exceptions that occur for each service address shall be recorded for display purposes to allow the designer to interpret the quality and accuracy of the map.

2.6.1. Unconnected Subscriber

The address location is associated with a subscriber account and/or active device. However, there is no RF tap with which the address location is associated. This condition may be caused by a missed plant extension.

2.6.2. Address Not Serviceable

The address location is not connected to an RF tap. Additionally, the address location is not associated with an active account and does not have a provisioned device. This may not be an error, since sometimes homes passed by the RF network are simply too far to connect. However, large numbers of non-serviceable addresses may indicate a map deficiency.

2.6.3. Unconnected RF Tap

An address location connects to an RF tap that does not connect to the RF network of any node housing. The RF tap does exist and can be plotted on the map. However, tracing the RF cable that connected to its import port does not lead back to a node housing. This typically represents a drafting error where the RF tap input port was not connected to an RF cable.

2.6.4. Incorrect Node Segment

The address location contains one or more provisioned devices whose node segment as reported by the telemetry data indicate it is on a different node segment than any of the node segments assigned to this node housing. If there are other RF taps from adjacent node housings within a typical drop distance, then the issue is either a drafting error or the proximity assignment algorithm.

2.6.5. Node Segment to Bus Leg Mismatch

The address location contains one or more provisioned devices whose node segment as reported by the telemetry data indicate it is on a node segment for a different bus leg of the node housing. This may be a drafting error or error in the proximity-based address to tap assignment.

2.6.6. Outside Node Boundary

Address locations connected to RF taps in the graph reside outside the node boundary. This condition is minor and would just require an update to the drafted node boundary.

2.7. Pre-Vet Analysis Results

Pre-vet data collection, correlation, and analysis, generates a set of data for use by the designers to determine where drafted maps have opportunities for improvement and how to segment or split the RF network for a capacity upgrade.

- Graph representation of all elements in the RF network.
- Service locations of homes passed and the connection points into the RF network
- Local and cumulative homes passed counts and device counts for each bus leg, amplifier, and splitter.
- Number of amplifiers and max amplifier cascade per bus leg.
- Service location exceptions including incorrect node segment, no connection to the RF network, and address not serviceable.

This data will be overlaid on the drafted maps to visualize the analysis results and any inconsistencies that have been found.

3. Maps and Data Output

Visualization of the pre-vet analysis results is overlaid onto the drafted access network map. The pre-vet analysis shall be presented to the designer using multiple layers that can be individually toggled on and off. This capability permits the designer to view the overall set of results or focus on a specific aspect of the analysis.

In Figure 15, service addresses and address errors are color coded based on the associated bus leg. This helps designers identify which addresses are tied to which bus leg. It also provides a visual indication of how many addresses are on a single bus leg and identifies the opportunities for map improvements.

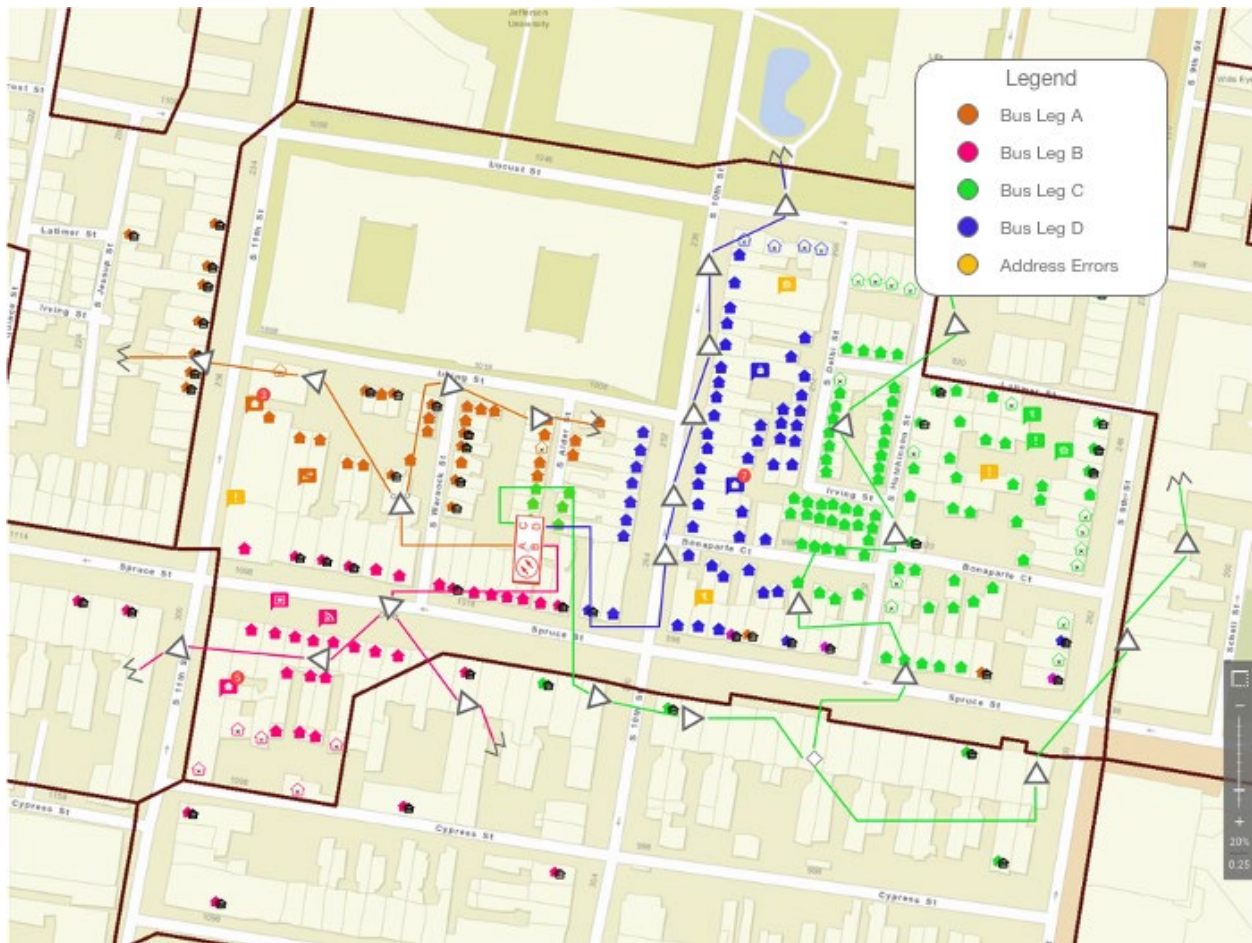


Figure 15 - Map of Plant with Bus Legs Highlighted and Address Errors

In the next two figures, users can see the device counts (aggregation of STB, CM, and MTAs) and homes passed at the amp level. The designer typically chooses an amp location to create a node split; Figure 16 shows how this view helps the user identify where there is pressure on the HFC plant.



Figure 16 – Device and Homes Passed Counts at Amp in Node Boundary View⁴

Figure 17 is a focused version of the previous figure. The image is centered on the end of a bus leg to view the cumulative and local homes passed and device counts at the amp level. Starting at the RF terminator, the cumulative counts add up the devices towards the node housing. The total number of devices and homes passed on a bus leg is shown at the node housing.

⁴ We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.



Figure 17 – Device and Homes Passed Count at Amp (Focused View)

In addition to the map view, the user interface provides a summary table of the bus leg. Figure 18 does not reflect real data, but a representation of how the data is displayed for users to easily digest.





Bus Leg Info				
Legs	 A	 B	 C	 D
# of Amps	12	0	18	0
Max Cascade	5	0	5	0
Serviceable Passings	78	1	103	3
Not Serviceable Passings	7	1	8	0
Device Count	66	0	110	2
Active Subscribers	46	0	72	1

Figure 18 - Bus leg table

The simplification and visualization of the data helps users understand the current state of the HFC plant in an efficient manner.

4. Conclusion

The pre-vet process aggregates and correlates data from multiple sources; the results are used to analyze the drafted RF network and determine the quality and accuracy of drafted maps. If these maps are deemed accurate, the pre-vet process provides information to the designer regarding where and how capacity upgrades should be applied to the RF network.

The process to automate these steps involves building a graph model of the RF network for a given node housing based on data in the design and drafting platform. The set of homes passed by the network are determined using both USPS addresses in the node boundary of the drafted network and HSD device telemetry data. These service locations are correlated against the drafted network to identify deficiencies in the drafted map.

A depth first search algorithm is employed on the graph to aggregate homes passed and device counts at the bus legs, amplifiers, and splitters. A view of the distribution of homes passed and devices across the network aids the designer when performing the capacity upgrade design. Additional data such as amplifier count and maximum amplifier cascade is also collected. The automation system presents the results to the designer through both a map-based interface and as tabular data.

The manual pre-vet process requires 2-8 hours to complete for a single node housing. Automating many of the pre-vet steps significantly reduces the amount of time a designer dedicates to this process. Time savings of 50% or more are possible, freeing the designer to focus more on the data analysis and design, rather than data collection and correlation.

Looking to the future, additional opportunities exist to enhance the pre-vet process. One example would be to collect upstream and downstream historical utilization data by node segment; this could be presented to the designer to augment the existing pre-vet's homes passed and device count data when designing the capacity upgrades. Beyond pre-vet automation can be applied to steps in the node capacity upgrade design process. Mid split designs, node segmentation, and even node splits share many common design steps, such as node housing swaps, amplifier upgrades, and power analysis. Automating these steps accelerates the design process, permitting faster delivery of network capacity upgrade designs.

Abbreviations

Amp	amplifier
CM	cable modem
CMTS	cable modem termination system
DFS	depth first search
HFC	hybrid fiber coaxial
HSD	high-speed data
MTA	media terminal adapter
RF	radio frequency
RPD	remote PHY device
STB	set-top box
USPS	United States Postal Service
vCMTS	virtual cable modem termination system

Voice Control of Set-Top Box for Customers with Non-Standard Speech

A Technical Paper prepared for SCTE by

Adina Halter

Sr. Principal Software Architect
Comcast

1701 John F Kennedy Blvd, Philadelphia, PA 19085
267-658-0261
adina_halter@cable.comcast.com

Sara Smolley

Co-Founder, Head of Partnerships
Voiceitt

700 Canal Street, Stamford, CT 06902
716-348-8229
sara@voiceitt.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Non-standard Speech and the Set-top Box	3
2.1. What is non-standard speech	3
2.2. Voice control and the set-top box	3
2.3. Current solutions for non-standard speech	4
2.4. Our approach	5
3. ASR Technology for Non-Standard Speech	5
4. Adaptive Remote Technology	7
5. Companion App.....	8
6. Integrating the Technologies.....	10
6.1. Customer authentication	10
6.2. Voice imprint.....	10
6.3. Set-top box control	11
6.4. Compatibility with iPhone accessibility options	11
7. Evaluating Impact of Integrated Solution with Customers	11
8. Conclusion.....	12
Definitions and Abbreviations	13
Bibliography & References.....	14

List of Figures

Title	Page Number
Figure 1 – customized model downloaded onto the edge device.....	6
Figure 2 – feedback loop.....	6
Figure 3 – adaptive remote on iPad (top and bottom scroll).....	7
Figure 4 – adaptive remote architecture diagram	8
Figure 5 – companion app architecture	9
Figure 6 – authentication model.....	10
Figure 7 – voice imprint model.....	10
Figure 8 – set-top box control model	11

1. Introduction

Voice control of set-top boxes is becoming the norm. But voice technology needs to be able to understand non-standard speakers as well. Non-standard speech can be a factor for people affected by deafness, disabilities, medical disorders, or even foreign language speakers.

Comcast and Voiceitt, an Israel-based voice technology startup, have collaborated to explore a solution using a mobile application paired to a set-top box. In this paper, we describe how we are applying machine learning and artificial intelligence to create unique voice command models for individuals with speech disabilities to access the set-top box. In this paper, we will cover:

- Non-standard speech and how this translates to customer set-top-box control
- Solution method and architecture
- Description of our ongoing customer trial.

Offering accessible voice control for non-standard speakers can open the opportunity for all customers to experience the joy and convenience of a voice-enabled home entertainment system.

2. Non-standard Speech and the Set-top Box

This solution can be adapted to other devices besides set-top boxes (mobile, TV, computer, streaming adapter) using similar approaches and methods described throughout this paper.

2.1. What is non-standard speech

Non-standard speech is speech that is not readily understood by others or by standard speech recognition. This could be because the speaker has an accent, has deaf speech intelligibility issues, uses speech synthesis, or has a physical or neurological disability such as dysarthria. Other types of non-standard speech disabilities such as Wernicke's Aphasia create a disconnect between thought and utterance. The technological solution described in this paper will work if there is a regular, repeatable connection between the thought and the utterance.

2.2. Voice control and the set-top box

Voice-driven technologies are proliferating rapidly. Growing adoption of smart speakers and smart assistants is likely to make speech recognition a primary means to interact with the technological world around us, including home entertainment.

In 2015, the Emmy-award winning Xfinity Voice Remote Control introduced the ability to control a set-top box with one's voice. Using machine learning, Comcast's Natural Language Processing platform ensures that the remote delivers precise results. "The platform leverages machine learning to understand what customers mean when they say certain words or phrases and deliver highly relevant results."

The speech recognition engine requires understandable speech to perform speech-to-text (STT) conversion. Thus, people with non-standard speech cannot access mainstream SST voice technologies.

"Switch the channel to HBO"

For many of us, this simple, familiar voice command spoken aloud in our voice remote control is a convenient way to instruct our home entertainment system. For people with speech and motor disabilities, being able to use their voices would give them an opportunity to take ownership of their TV, increase their independence, and decrease their dependence on people around them to perform tasks such as changing the channel, recording a show, or browsing content.

In her report "Xfinity Adaptive Remote for Accessibility Audiences" Theresa Murzyn Ph.D., Lead UX Researcher at Comcast states "Not being understood takes a mental and emotional toll on users with motor/sensory challenges. The Adaptive Remote must enable users to feel understood regardless of their input method."

Customers' physical challenges, multi-sensory personas, and video mindsets inform what they expect from set-top box control technology, namely:

- Independence
- Fewer steps
- Speedier navigation through the TV/cable interface
- Less mental and physical effort
- Being valued

2.3. Current solutions for non-standard speech

"It doesn't understand me. I don't know why." (giggle)
— Comcast customer with amyotrophic lateral sclerosis (ALS)

Many of the people who can benefit from "voice first" technologies cannot access those technologies because they do not have the standard speech patterns that are recognizable to commercial automatic speech recognition (ASR) algorithms.

So what options are currently available to them to navigate this "voice-first" world?

"I have the capabilities of doing streaming if someone else is here pushing the buttons for me. But I can't do it myself. And so, I never do it."
— Comcast customer with spinal injury

Non-standard speakers may rely on friends, family, or caregivers for basic tasks, including controlling their devices for everyday tasks. Voice control for their set-top boxes can provide independence in these everyday routines.

"I use my iPad for a lot of YouTube. And I used to use a [sic] voice activation to get to it, but now it doesn't understand me any longer. So, I've kinda' lost the use of it."

— Comcast customer with ALS

Technologies such as the Xfinity Adaptive Remote are beginning to give touch and text alternatives to those who cannot speak into a physical remote control. Customers can pair the Adaptive Remote with assistive technologies (ATs) such as eye control, mouth sticks, gross-motor options for swiping and large-target tapping. This pairing gives these customers an opportunity to trigger set-top box actions or submit a text string version of the "voice" command they are interested in. Submitting this text string bypasses the ASR algorithm.

And yet, this technology often requires many steps to complete a simple task. This makes content foraging slow, tiring, and often exasperating. As one of our customers with ALS remarked, "Every click is time."

"They are already managing many challenges. let's not add more to them."

— Theresa Murzyn, PhD.

2.4. Our approach

In this project, Voiceitt's non-standard speech recognition was integrated with the Xfinity Adaptive Remote's video code (vcode) and string input capability, using Comcast's Companion App architecture as the technology bridge. This enables customers with speech disabilities to access and control their set-top boxes (and wider home and entertainment platforms) by voice.

The solution presented here was the innovation of two companies, Comcast and Voiceitt. Please note that while we will often refer to our different technologies by company or product name in this paper, this is done to keep our two companies' individual contributions and solutions clearly differentiated. Similar solutions can be developed by your organization's product team as well.

3. ASR Technology for Non-Standard Speech

Voiceitt's ASR technology is designed to recognize the speech of people with speech disabilities. The technology includes both discrete and continuous ASR for non-standard speech. Discrete ASR offers the ability to recognize a predefined list of phrases which the user with speech disabilities can customize. For example, if the speaker trains the software with the vocal pattern "uhwuh o uhah" and its meaning is, "I want to go outside," the software learns to recognize this pattern and associate it with its meaning, which it can then produce through digital speech.

Continuous ASR, now in Beta, extends this functionality to recognize the user's speech more flexibly. With the continuous ASR, there are no longer constraints to use predefined phrases from a phrase bank, thus allowing the user to speak more spontaneously and freely.

Both the continuous and discrete ASR technologies are customized solutions tailored to the individual user. As such, they rely upon enrollment data (training data): samples of the user's speech. This enrollment data is used to adapt the acoustic model to provide a more accurate representation of the individual's speech.

Further, hands-free activation is supported using the Voiceitt wake word technology, extending further accessibility for users with disabilities.

In the case of the discrete ASR solution, this customized model is downloaded onto the edge device as illustrated below.

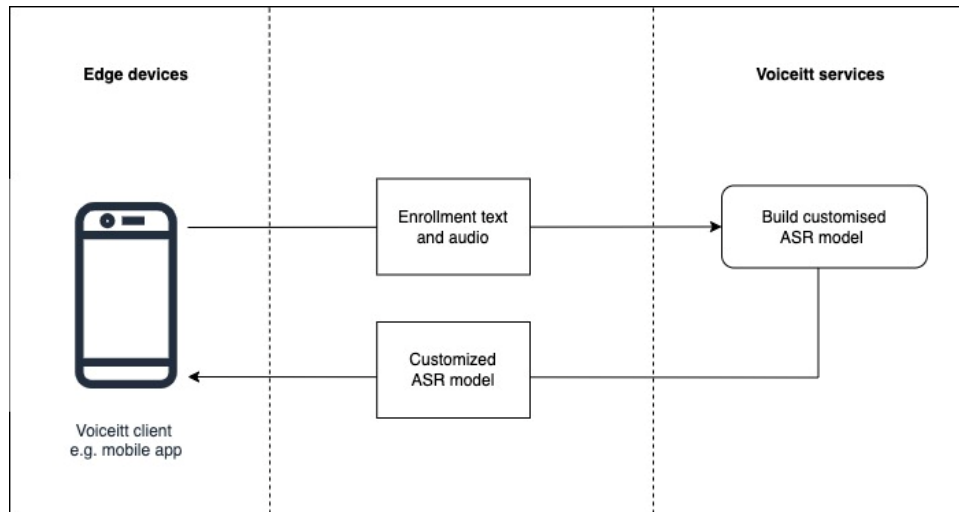


Figure 1 – customized model downloaded onto the edge device

Once the enrollment phase is complete, the user may use the edge device to recognize his/her speech. This recognition takes place on the device in the case of the discrete ASR solution. Importantly, a feedback loop is implemented which continuously improves the accuracy of the solution. This feedback loop is illustrated below.

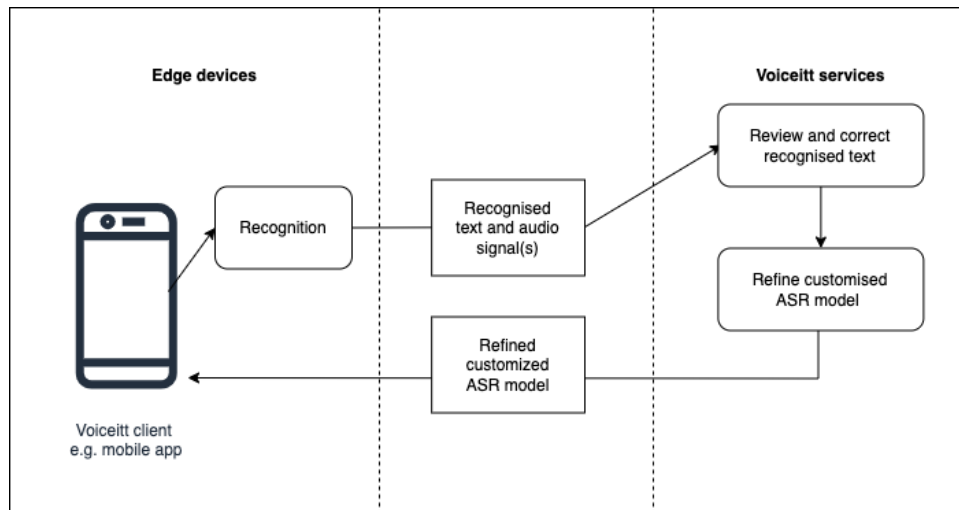


Figure 2 – feedback loop

Recognized text and its associated audio signals are sent to the backend database. The text and signals are reviewed and potentially corrected by a team of skilled annotators to provide additional training material to the model customization procedure. A further customized model is then delivered to the edge device for future recognition. This process iterates while the user engages with the technology, creating a virtuous machine learning cycle which delivers improved ASR accuracy.

A very similar workflow of enrollment, recognition and feedback is deployed in the case of continuous ASR. The primary difference is that continuous ASR uses a combination of an acoustic model and a language model to add the ability to recognize free speech using words and phrases that were not pre-trained, as well as phrases that are not in the pre-defined order.

This method could also be integrated with speech-based technologies such as those found on voice-controlled interfaces such as set-top boxes and smart home devices, effectively enabling users to access such technologies.

4. Adaptive Remote Technology

The Xfinity Adaptive Remote (<https://remote.xfinity.com/>) is a web application written in NodeJS which allows users to control their set-top boxes with various ATs such as the Tobii Eye Gaze solution.

The original project motivation was to provide remote tuning capability for our customers with ALS (also known as Lou Gehrig's disease) using an eye-tracking device such as the Tobii Eye Gaze. Later, features such as support for voice commands and voice-as-text commands were added.

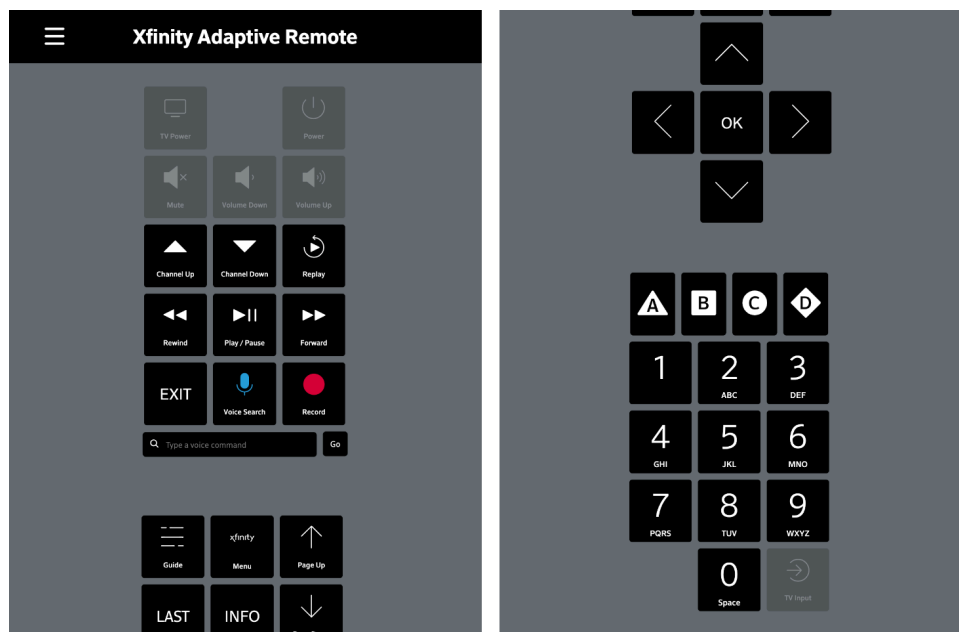


Figure 3 – adaptive remote on iPad (top and bottom scroll)

Tapping buttons on the adaptive remote simulates the press of one of the keys on the physical remote by sending the same vcode to the set-top box that the physical remote would send.

As an alternative to voice control, the adaptive remote has a field to enter a text string that would mimic the desired spoken command. Examples of voice commands include: “NBC”, “Peacock”, “Show me comedy movies”, “Guide”, “Channel up”.

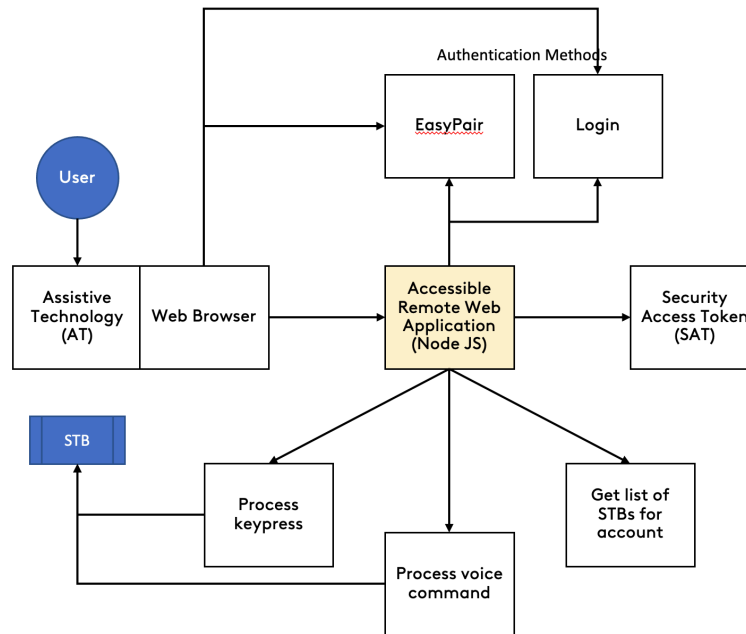


Figure 4 – adaptive remote architecture diagram

5. Companion App

The Adaptive Remote (AccRem) architecture supports “companion applications”. These companion apps use the existing AccRem app for login and TV Box selection and then use an AccRem web service application programming interface (API) developed specifically to support these companion applications. This AccRem architecture enabled the Voiceitt companion application to control the Xfinity set-top box.

Cross-Origin Resource Sharing (CORS)

CORS is an HTTP-header based mechanism that allows a server to indicate any other origins (domain, scheme, or port) than its own from which a browser should permit loading of resources.

Companion IDs

Each run / instance of a companion app must uniquely identify itself by a universally unique identifier (UUID). This is the value that will be implicitly passed via the adaptive remote app when submitting key presses, sending custom text commands, etc.

User Flow

The overall flow for an end user with speech disabilities will typically go something like this:

1. User visits a companion web app, hardware solution, or “fat client” app.
2. The companion app generates a new companion UUID value.
3. The companion app puts up a login button/link with a URL that contains the UUID.

4. Immediately after the user clicks on the button (which opens the AccRem app in another tab or mobile web view), the companion app puts up a “please wait” screen and begins polling using an API pairing endpoint.
5. Once the user logs in and choose her set top box, the AccRem app goes to the companion-success page which shows a message such as “Go back to your Companion App” and the pairing endpoint returns a token value in its response, which causes the companion app to stop polling and move the user to the companion app’s “main screen”, showing buttons, an input field for entering custom commands to control the set-top box.

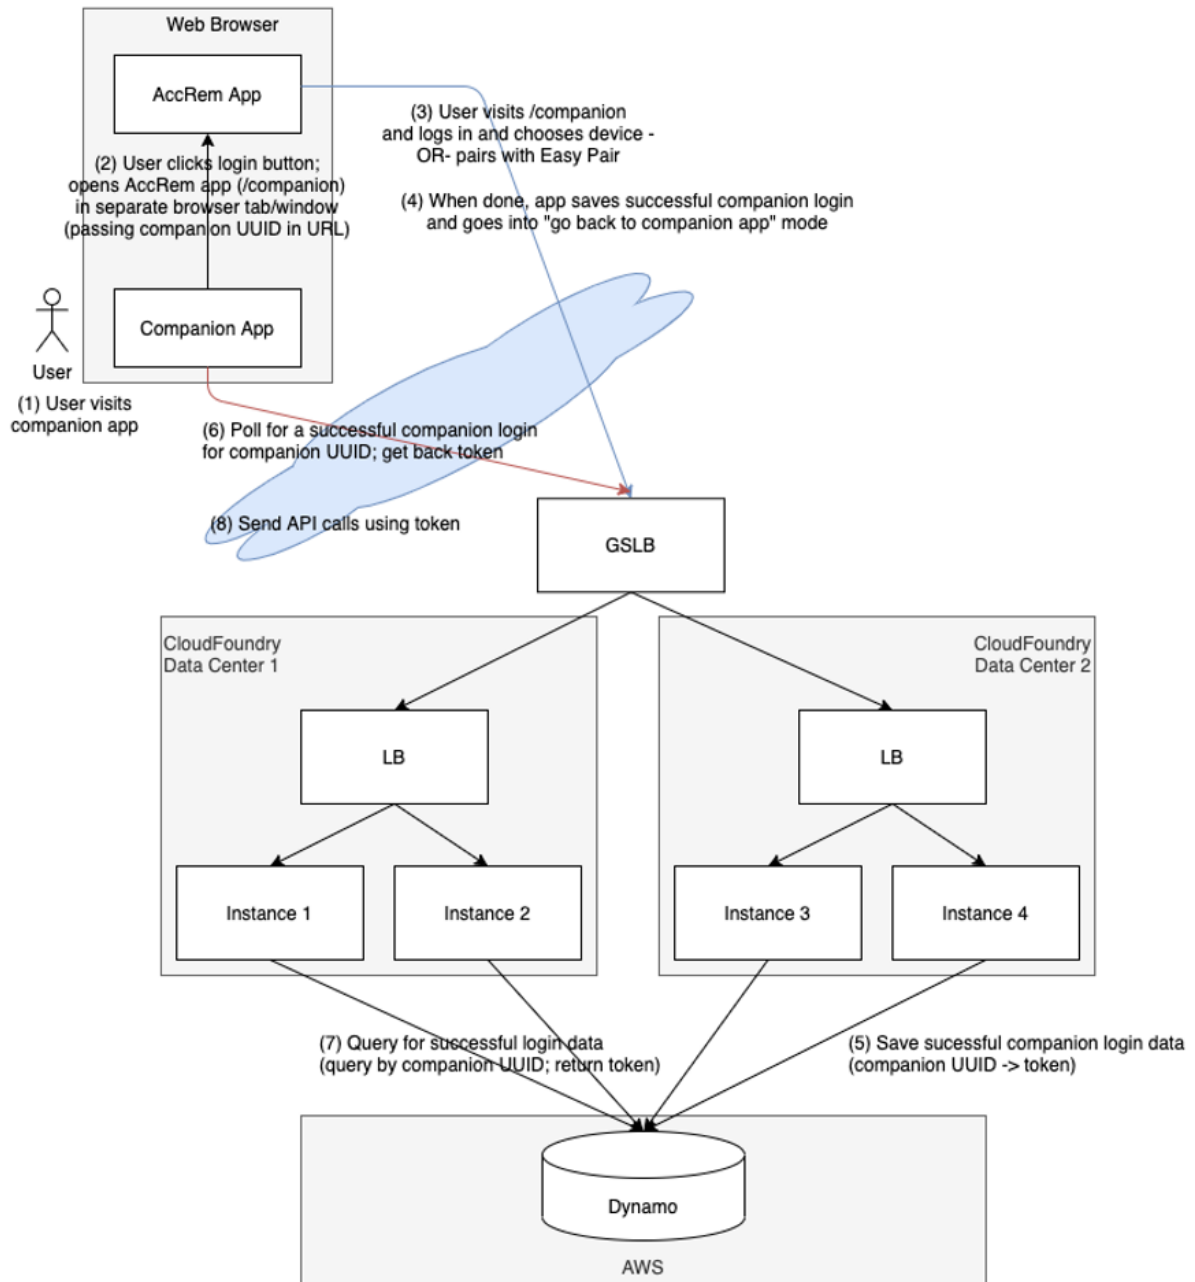


Figure 5 – companion app architecture

6. Integrating the Technologies

Comcast's Companion App was written in NodeJS, therefore Voiceitt translated this to C++ to work with their iOS mobile app. Personal identifiable information (PII) privacy was key in our joint solution.

(We are using our company/product names here to illustrate how we integrated the different technologies while ensuring the privacy of each company's customers. Again, similar solutions can be developed by your organization's product team as well.)

6.1. Customer authentication

To ensure PII privacy for Comcast customers, authentication is done on Xfinity-domain interfaces rendered in the Voiceitt app's web view. No authentication is done through the app itself.

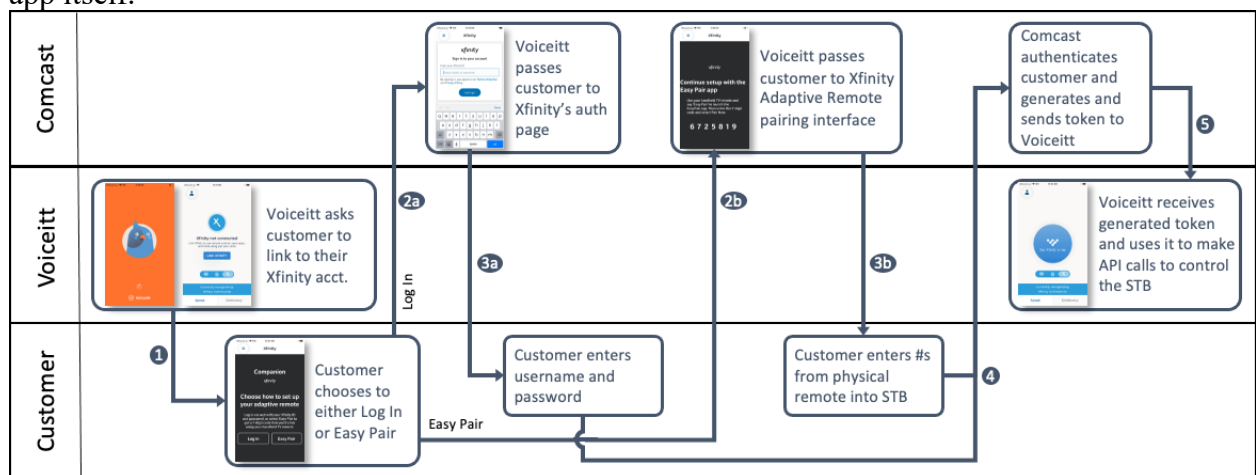


Figure 6 – authentication model

6.2. Voice imprint

To ensure PII privacy for Voiceitt customers, no voice recording is ever shared with Comcast, and moreover is compliant with international data privacy protocols.

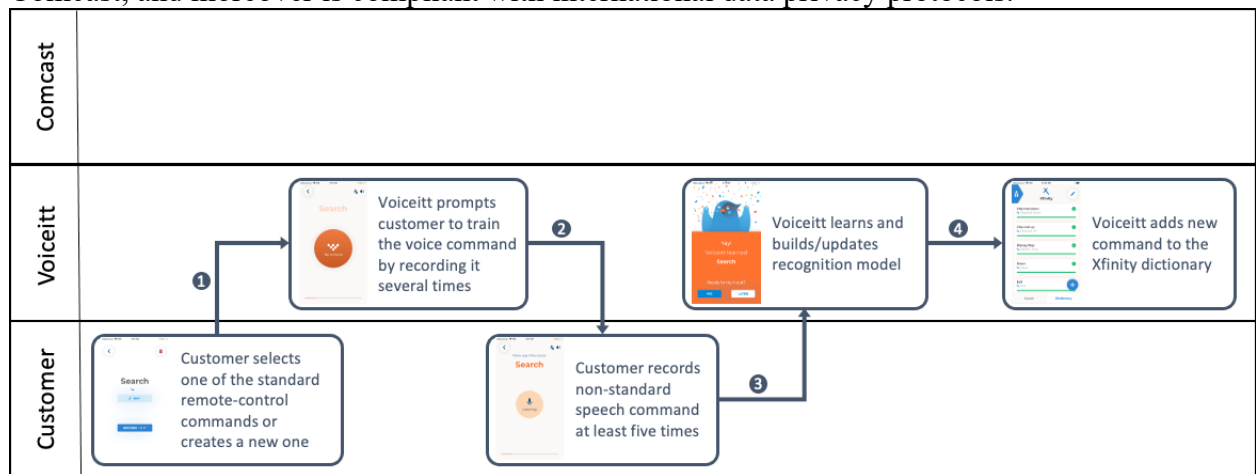


Figure 7 – voice imprint model

6.3. Set-top box control

The cable customer with non-standard speech is now able to control their set-top box with their voice.

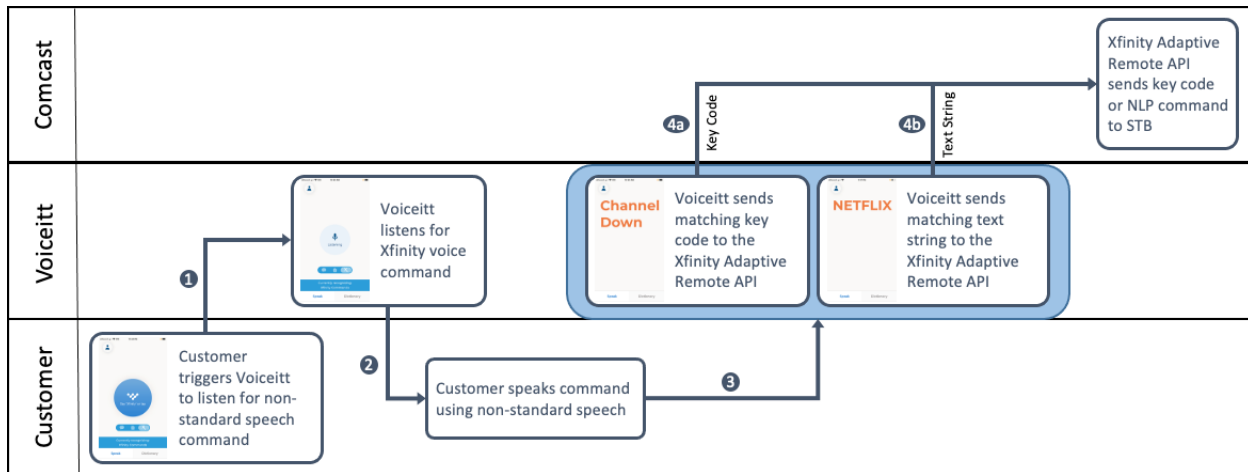


Figure 8 – set-top box control model

6.4. Compatibility with iPhone accessibility options

Any app should be compatible with OS accessibility settings and ATs so that those with disabilities can use it. We ensured all content was readable and in order when using a screen reader. We ensured all actionable items could be reached via finger swipe or AT and were labeled properly. We made sure that all [World Wide Web Consortium \(W3C\) Web Content Accessibility Guidelines \(WCAG\)](#) were followed at the AA level.

The companion mobile app is also designed and developed with innovative accessible user experience and design so that individuals with motor control impairments, cognitive, and dexterity challenges may access it as independently as possible.

7. Evaluating Impact of Integrated Solution with Customers

Comcast and Voiceitt have collaborated with a specialty nursing care facility in Philadelphia to evaluate the integrated solution described in this paper with end users with dysarthric speech.

The objective of the pilot is to evaluate how the Voiceitt app, which has integrated Xfinity's Adaptive Remote technology, improves independence and quality of life for individuals with dysarthric speech.

In an ongoing pilot, participants with highly atypical speech patterns correlated with cerebral palsy use Voiceitt's customizable speech recognition engine to activate a series of voice commands to their Xfinity X1 set-top box via Voiceitt's consumer application. The participants may not have had prior experience with voice devices or speech recognition; or, they may have previously tried to use these devices but without success. The available voice commands are chosen by each participant, sometimes with the help of a caregiver. The user calibrates the system by recording their voice, following prompts on the screen of their mobile device.

The pilot, now ongoing, will include input from participants, their caregivers and support professionals. Recognition accuracy and daily usage is measured through the companion

application. Impact on customer satisfaction, engagement, and usage, as well as quality of life and independence will be evaluated through a series of interviews and questionnaires with participants, facility administrators, and their daily support professionals.

8. Conclusion

The opportunity to give customers with non-standard speech (especially those whose speech is impaired due to neurological or physical disability) the ability to use their natural voices to control their entertainment system returns to them a sense of independence that offers fewer steps, speedier navigation, less mental and physical effort, and greater overall enjoyment of these offerings.

By integrating two solutions via a companion app bridge, we may serve not only those in the disability community with non-standard speech, but also those with accents, age-related tonal changes, etc. In short, while accessible solutions are necessary for some, they can be helpful for everyone.

We would like to acknowledge Theresa Murzyn and Mike Fine at Comcast. Theresa's user research on how those with ALS and Spinal Injury use media and home entertainment has been invaluable. Mike Fine's assistance in understanding the architecture behind the adaptive remote and companion app APIs has been vital to this technical paper. We would like to give special thanks to our partners at the facilities who provide help in recruiting and supporting individuals with speech disabilities participating in this collaborative pilot.

As our joint pilot progresses, further input from customers with disabilities will inform refinements to the technological approach described here, which will make our solution even more impactful and effective.

Definitions and Abbreviations

ALS	amyotrophic lateral sclerosis, also known as Lou Gehrig's disease, is a progressive neuro-degenerative disease that affects the brain and spinal cord.
Aphasia	The inability to understand what is being said, find the necessary word for something, or formulate sentences due to damage in the brain, often from a stroke or accident.
API	Application Programming Interface
ASR	Automatic Speech Recognition
Dysarthria	A speech disorder caused by either muscle weakness or the inability to control speech muscles due to brain damage.
Content foraging	entertainment system navigation and searching techniques to find content via direct retrieval or orienteering
Easy Pair	A method to connect a remote control to a set-top box by typing the numbers shown on the set-top box interface using the keypad on the remote to be paired.
IoT	Internet of Things
ML	Machine Learning
NLP	Natural language processing
PII	personally identifiable information such as name, address, streaming content history, etc.
RDK	Reference Design Kit (https://rdkcentral.com). RDK is a fully modular, portable, and customizable open-source software solution that standardizes core functions used in video, broadband, and IoT devices.
RDK-V	Reference Design Kit for Video.
STB	Set-top box
Speech synthesis	Artificial production of human speech by computer or speech synthesizer.
UUID	universally unique identifier. A 128-bit alpha-numeric to identify a person, peripheral, etc. without PII (personally identifiable information).
Vcode	Video code. The code sent to the set-top box when a remote-control button is pressed.
Voiceitt	Voice technology startup that has developed automatic speech recognition for non-standard speech.
W3C	World Wide Web Consortium
WCAG	Web Content Accessibility Guidelines. Guidelines written by the W3C's Web Accessibility Initiative directing designers and developers of web applications on accessibility requirements and standards. There are three levels of compliance: A, AA, AAA.
Wernicke's Aphasia	Seemingly fluent speech which is made up of unrelated words (schizophasia, often termed "word salad") or even non-words.

Bibliography & References

Accessible Remote Technical Information, 2021. Mike Fine, Principal Software Architect, Entertainment Experiences, Comcast

Companion App Developer's Guide, 2021. Mike Fine, Principal Software Architect, Entertainment Experiences, Comcast; Adina Halter, Sr. Principal Software Architect, Accessibility Innovations, Comcast

[*Introducing the New X1 Voice Remote*](#), Dec 11, 2017. Jonathan Palmatier, Comcast

[*Comcast Wins Emmy Award for X1 Voice Remote Technology*](#), Aug 29, 2017. Comcast

[*Voiceitt Makes Alexa Accessible for People with Disabilities*](#), PR Newswire

Xfinity Adaptive Remote for Accessibility Audiences, May 19, 2022. Theresa Murzyn, Ph.D., Lead UX Researcher, User Research, Comcast

WAN And LAN Speed and Service Matching – Are We Engineering It Correctly for Consumer Services Growth In The Next 5+ Years

A Technical Paper prepared for SCTE by

J.R. Flesch
Director, Advanced Technology
CommScope
3871 Lakefield Drive, Suwanee, GA 30024
Jr.flesch@commscope.com

Kurt Lumbatis, CommScope

Ian Wheelock, CommScope

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Bitrate and Latency, Averages and Peaks.....	4
2.1. The Wi-Fi (LAN) environment.....	4
2.1.1. Wi-Fi 6, Wi-Fi 7 and Standard Power.....	6
2.2. The DOCSIS (WAN) Environment.....	9
2.2.1. DOCSIS 4.0 Enhanced Spectrum DOCSIS (ESD).....	9
2.2.2. DOCSIS 4.0 FDX.....	10
3. Managing E2E Packet Latency Across WAN and LAN.....	11
3.1. Low Latency DOCSIS.....	11
4. Extant and Emerging Premise Services Landscape.....	14
4.1. Consumer Impressions of Wi-Fi Performance and Feature Value.....	15
4.2. Premises Services Data Payloads.....	25
4.3. Data Demands of a Model Smart Home in 2027.....	28
5. Performance of the Adjoined Networks.....	31
5.1. Stressing the Wi-Fi 7 Home Network.....	31
5.2. Taxing the DOCSIS Wireline Network.....	33
6. An OTT Sidebar.....	34
7. Conclusion.....	35
Abbreviations.....	36
Bibliography & References.....	38

List of Figures

Title	Page Number
Figure 1 – Standard Power Benefit over LPI at 6 GHz in Arris Wi-Fi House.....	7
Figure 2 – Effect of MLO on Goodput at Various Bonding Options.....	8
Figure 3 – DOCSIS 4.0 ESD Variants.....	10
Figure 4 – DOCSIS 4.0 FDX Variants.....	11
Figure 5 – Evolution of DOCSIS pre-4.0 Roundtrip Time Latency.....	12
Figure 6 – 2-Component Flows in DOCSIS Pipe, Low Latency and Classic.....	12
Figure 7 – Differential Flow Performance Upstream for Latency and Jitter.....	13
Figure 8 – Candidate Service Mounts for Low Latency Grooming.....	14
Figure 9 – Regional Basis for User Network Impressions.....	15
Figure 10 – Users’ Broadband Speed Tiers.....	16
Figure 11 – Budgeted Monthly Cost for Network BB Data.....	16
Figure 12 – Upstream Bitrate Rating.....	17
Figure 13 – The Upstream Value Challenge.....	18
Figure 14 – The Downstream Value Challenge.....	19
Figure 15 – Where Latency Fits for Users VS DS Bitrate.....	20
Figure 16 – Where Jitter Fits for Users VS DS Bitrate.....	21
Figure 17 – BB Connection Rate Impetus to Shopping.....	22
Figure 18 – VR impact to BB Connectivity Requirements.....	23
Figure 19 – Most Important Broadband Network Attribute.....	23

Figure 20 – Broadband Network Enabled Home Services	24
Figure 21 – Simultaneous User Leverage of WAN.....	25
Figure 22 – Stacked Bitrate Increases Driving VR Simulation Immersion and Super-Resolution Screens	26
Figure 23 – Wi-Fi House Basement Device Outfitting	29
Figure 24 – Wi-Fi House Main Level Device Outfitting	30
Figure 25 – Wi-Fi House Top Level Device Outfitting.....	30
Figure 26 – In-Home Wi-Fi 7/Standard Power (with Extender) Network Performance	31
Figure 27 – In-Home Wi-Fi 7/LPI only (with Extender) Network Performance	32
Figure 28 – In-Home Wi-Fi 7/Standard Power (no Extender) Network Performance.....	32
Figure 29 – Mid-day, one Wi-Fi hop OTT modem performance	35

1. Introduction

In the monotonic drive for ever-greater bitrate consumption posed by the increased variability and sheer number of home dweller-facing digital network services (particularly immersive experiences and IoT-anchored emergency services), a review of the balancing act between WAN and LAN capacity and responsiveness begs execution. Such review seems particularly warranted at this point in the cable distribution and home Wi-Fi development epochs, given that the recalculation of transmit and receive capacities posed by the DOCSIS 4.0 WAN upgrade nearly coincides with the adoption of in-home Wi-Fi 7 LANs (themselves also potentially enhanced by higher EIRP /AFC-enabled Standard Power).

This paper will examine the implications of emerging network demands in a post-pandemic home services environment and attempt to evaluate the impact to WAN and LAN infrastructure; identifying behavioral, economic, or capacity mismatches which may alter the scope and rate of proposed network evolutions. What seems solicited is an advisory to network stakeholders on the scope of WAN and LAN futureproofing and a weighting of appropriate investment required to adequately meet anticipated mounted service connection needs in a moderately near-term adoption horizon.

Note that the exclusion of OTT WAN inroad is an intentional artifact of the constrained scope of this paper, but such competitive dismissal assumes minimal impression (consumer exploitation) in the 5-year consideration window. A brief sidebar is included for future reference and study, however, given the threat's higher global accessibility than fiber (PON), modest capability and cheaper entry ticket.

2. Bitrate and Latency, Averages and Peaks

Both the network WAN and the in-home LAN exhibit effects of mixed-flow data stochastics, though with the obvious contrast of WAN versus LAN data endpoint counts (clients and upstream sources). The WAN's high end client count and dedicated multicast video channel assets guarantee that its carry-to-capacity during peak use exhibits a higher average ratio than the bursty (though still largely video-based ABR data consumption) associated with adaptable-rate Wi-Fi traffic bound for on-premises distribution. In fact, in terms of raw bitrate capacity, Wi-Fi airtime capacity is typically leveraged to a much smaller fraction of its upper limit in the average US household (which is obviously not the case with the wireline WAN, especially during peak use periods). This begs questions around the required number of contending Wi-Fi clients (or competing in-range, interfering APs) which would render the LAN incapable of meeting its QoS aspirations – especially given next generation Wi-Fi MAC market penetration.

From the wireline WAN standpoint, DOCSIS 4.0 (in one or more of its species) promises improved duplex signaling capacity via modified exploit of the wireline's duplex split (including dynamic management of same in FDX instantiations), jettisoning video's SC-QAM in favor of full OFDM leverage and ESD (initially to 1.8 GHz) to accompany higher spectral efficiency modulation schemes. All these improvements promise higher bitrate support in both directions over the WAN.

2.1. The Wi-Fi (LAN) environment

The summary standard of on-premises network performance is Wi-Fi bitrate, and its singular comparative assessment is “the higher, the better”. And while raw delivery speed makes up for many signal distribution sins, it is worthwhile to recognize that data movements between Wi-Fi endpoints lever orchestrated solicitations of clear airtime in one or more radio bands to guarantee transfer of packets, the integrated view of which amounts to flows of varying bitrate and waypoint latencies. Contention in the airtime space – especially from native-mesh-asynchronous radio sources – compounds the access issue and confounds the negotiating parties; this serves to reduce delivery determinism and inject noise into the

packet goodput in the form of varying latency (jitter). In pathological cases, it can result in intermittent service interruption (if not outright failure) for any given link.

Fundamentally, Wi-Fi has evolved to include methods for framed exchange (RTS, CTS) and priority handling (WMM with its varying backoff counts) and now with the advent of MAC 7, features much broader channel BWs, higher spectral efficiency (denser QAM), cross-band channel aggregation, bi-directional MIMO and inherits OFDMA, BSS coloring and TWT from MAC 6, all of which serve to give its scheduler more tools to rapidly – and with low latency and higher dependability – transfer packets across endpoints even under challenging interference conditions. The goal, of course, is “responsive immersion” for end users in cases where content is being consumed or interacted with; and immediate cloud network connectivity for those services which operate on emergency care or home security tiers within a given residence.

Wi-Fi responsiveness has two components – bitrate and latency. The twitch gaming community (and simulation scenarios in general) have drawn attention to roundtrip command processing latency and the specific requirement of predictable, low jitter – given that now the issue of raw data pumping capacity has moved well past a 25x surfeit of capability for any mounted service in the home wireless domain. In other words, since we seem to be able to burst data to/from endpoints with low enough airtime duty cycle, it now becomes more critically important to remove queuing shims from the transfer process and more appropriately address tiers of delivery priority for the packets themselves.

On the Layer 2 front, the typical resolution for latency issues comes in the form of different priority queues (with delivery deferral of, and pre-empted delivery for, lower versus higher priorities respectively), clever multi-receiver packing of right-sized data payloads (OFDMA) and hold timeouts (buffer truncations) capping the maximum dwell time for lined-up packets. The Layer 1 mitigation involves identifying surplus airtime in underutilized Wi-Fi channels (in bands mutually supported by the attaching endpoints). Further, these types of solutions may be reactive (algorithmic responses to particular link telemetric triggers like airtime capacity or buffer lengths) or anticipatory (via managed pre-emptive adjustments given TOD and stored prior successes). Condensed, the Layer 1 consideration involves the capacity calculus of the given band and channel (so, typically, MCS x BW x SS) versus any steering options on the band/channel front for the affected client(s). If better spectrum efficiency (shorter burned airtime) is available in an alternate channel or band, the scheduler chooses that more optimal solution. And by leaning on MCS as one of the calculus components, the EIRP benefits of alternate spectrum choice will automatically receive an accounting (more power meaning better received C/N and hence higher MCS). A key leverage here is the 6 GHz spectrum, where even simple single user round-robin (non-OFDMA) scheduling can achieve 2 msec link latencies provided the served spectrum client count does not exceed a half-dozen or so participants.

The initial step in effectively handling latency expectations is to either infer requirements based upon some period of traffic observation or have the user explicitly designate a particular client/service binding which sets priority expectations for its data needs. The inferential determination can be done autonomously (and aligns with “fingerprint” methodology necessary when MAC randomization becomes a fixture in home wireless LANs – certainly likely during our study horizon). To mitigate startup jitter – and during the period when client traffic inference is still pending, the default priority handling should be set to highest (least latency and jitter). Over some initial transfers, the traffic from AP to client will establish its stochastic metrics (fingerprint) and “graduate” to the appropriate priority and latency setting. And while WMM provides a Wi-Fi mechanism which suggests leverage of 4 different priority queues, tremendous improvement for jitter comes with just the first parsing of traffic from “first-come, first served” to “warrants special handling”.

Once the traffic is parsed for latency priority, the scheduler can then employ various buffer, payload packing and spectral leverages to manage the latent aspects of data connectivity for the impacted services. In the case of a low latency need, for example, a PHY optimization would consider the immense spectral width of the 6 GHz band. This large greenfield spectrum makes it straightforward to select a wide BW channel (to enhance burst speed) and locate it such that CCI is not measurably present (to guarantee a C/N which maximizes MCS and hence, bitrate); the Achilles' heel is the obvious requirement that both ends of the wireless link must support the 6 GHz band. The upside is that Standard Power and Wi-Fi 7 – both due within a year's time (or so) -- will both serve to catalyze adoption of devices which can lever that band.

2.1.1. Wi-Fi 6, Wi-Fi 7 and Standard Power

Briefly, Standard Power 6E allows the AP to paint the home expanse with up to 9 dB additional EIRP (over LPI at 160 MHz BW in the 6 GHz band) and essentially walks all 6 GHz clients up the MCS bitrate cascade (for downlink delivery), reducing consumed transmit airtime for a given packet size.

Wi-Fi 6E

LPI to SP Benefits



Figure 1 – Standard Power Benefit over LPI at 6 GHz in Arris Wi-Fi House

For its part, Wi-Fi 7 allows the AP to bind data delivery over multiple bands, with more granular BW leverage (think: spectrum puncturing, fractional RU use and multiband exploit), than any prior Wi-Fi MAC – and so can more cleverly leverage gaps in contested spectrum (in up to 3 bands, at that) to increase delivery rate and reliability over that of a single monolithic channel/single band at best-case C/N.




Figure 2 – Effect of MLO on Goodput at Various Bonding Options

The Wi-Fi 6E stranglehold on operations in the 6 GHz band will not even be fully two years old before Wi-Fi 7 devices will debut, beginning late this calendar year/early next. And it appears this landmark will also (roughly) coincide with the availability of Standard Power for AP devices in that same band. The significance of these coincident events manifests in outstanding bitrate reach throughout even large footprint estates – and couples this with multiband robustness and unprecedented levels of predictable, low latency signal distribution. As previously alluded, APs with Wi-Fi 7 will be able to implement schedulers which have at their disposal broader channel bandwidths than those offered by Wi-Fi 6 (320 vs 160 MHz), spectrum puncturing, MLO, OFDMA, bidirectional Mu-MIMO and TWT operations. Orchestrated data exchanges to the heterogeneous MAC limitations of respective client devices should be possible. This bonanza in Wi-Fi 7 additions to 6E – and their so-proximate release to the first 6E-capable devices – has largely rendered 6E a moot (OBE) waypoint proposition (succeeding only in fragmenting the 6 GHz Wi-Fi market in its early days, unfortunately). There exist two exceptions to this proclamation on 6E early obsolescence: “virtual wireline”, low-latency dedicated 6 GHz Wi-Fi channels between AP

and gaming endpoint and AP:AP (non-backed-off power and maximum BW on a dedicated channel) trunk links for wireless extension within the premises.

Regarding the former, addition of Ethernet (or USB)-to-6E dongles to AP and client on opposing ends of a dedicated-channel 6E link provide an instantaneous means of running a (virtual) Ethernet cable between those devices and establishing near-wireline delivery latency on said link. Regarding the latter, for large-footprint homes, Wi-Fi extension via a dedicated 4SS 6E link between gateway device and mid-home can provide Wi-Fi extension with trunk capacity easily exceeding 2 Gbps – again with low single digit millisecond latency over the link. And by doubling down on this arrangement via leverage of AFC-enabled standard power, 4W EIRP at both ends of the link will be allowed; this trumps the power footprint for APs in any other Wi-Fi band. Boutique solutions, certainly – but also worthy of note in the time horizon under study.

2.2. The DOCSIS (WAN) Environment

Though (for some) glacier-paced, the onset of a DOCSIS 4.0 era has nonetheless proceeded and now finds itself well into commercial chip implementation phase within the development domains of the two most pedigreed commercial participants in the DOCSIS legacy. However, the end game – as regards deployment particulars and timing – is very much undetermined. There exist multiple permutations of DOCSIS wireline epoch which could conceivably meet network needs in the coming few years with rather steep implementation costs for some weighing on the solution calculus. Relief for the restrictive legacy return path is the only identifiable common element at this juncture, with the trunk diplex split, continued exploit of SC-QAM for video and final forward bandwidth still under evaluation – as well as the MAC subtype(s) to be exploited across the full band (such impacting the effective spectral density and hence capacity of the total spectrum).

Regarding the latter, a “budgetary” approach to upgrading the DOCSIS network might be to move to DOCSIS 3.1, slide the diplex frequency up to 204 MHz and set the forward band limit to at least 1.2 GHz. Such a conservative approach might attract some attention from (and provide temporary relief for) capex-bound systems, but for the purposes of analysis we will focus on the DOCSIS 4.0 variants matrixed with more future-proofed options on diplexer frequency and upper band edge (in anticipation of a higher adoption rate which better describes the coming WAN environment for the greatest population of NA subscribers). A brief summary of DOCSIS 4.0 options follows, in order of increasing implementation complexity.

2.2.1. DOCSIS 4.0 Enhanced Spectrum DOCSIS (ESD)

The following pictogram illustrates diplex split operations for DOCSIS 4.0 ESD in either 5x5 or 5x7 formats (referring to number of OFDM 96 MHz blocks of upstream traffic):




Figure 3 – DOCSIS 4.0 ESD Variants

These are but two examples; the spec lists five different diplexer options in up/downstream operation. Note the corresponding shift in the diplex corner frequencies in these two cases and the option to shift the full band forward cutoff up to 1.8 GHz if desired. If we conjure full exploit of the 1.8 GHz bandwidth with only OFDM carriage (no video SC-QAM) and maximum upstream bandwidth (so closest to most symmetric wireline operation), the EOL performance of a cascade would be capable of a downstream PHY rate exceeding 10 Gbps for all but the highest diplexer splits.

2.2.2. DOCSIS 4.0 FDX

ESD pushes the forward spectrum out and eliminates less efficient SC-QAM, but FDX goes a step further (and more complex) – shiftable duplex signaling bands (adjustable diplex frequencies). The CableLabs visualization of this looks as follows:




Figure 4 – DOCSIS 4.0 FDX Variants

The clear benefits on the flexibility front fall into the category of adaptable duplex ratio. Post-pandemic in particular, it is not entirely clear whether the traditional 10:1 split (in the US) for downstream:upstream capacity will continue to make sense. (The SOHO, IoT and data-sourcing aspects may see relatively higher growth rates than traditional lean-back video consumption, for example.)

3. Managing E2E Packet Latency Across WAN and LAN

Proper stewardship of packet delivery across the concatenated DOCSIS and Wi-Fi networks involves a synonymous mapping scheme for the various packet priorities within the two domains; anything less promotes thrash in the queuing arrangements on one side or the other, resulting in jitter increases as the packet queues are spasmodically serviced. Fundamentally, synonymous mapping involves marrying the wireline LLD (low latency DOCSIS) – where appropriate -- with Wi-Fi's WMM and assigning sensible, staggered priorities for the client services drawing data connectivity support from the two networks.

3.1. Low Latency DOCSIS

Pre-DOCSIS 4.0 already features an unevenly implemented technique for minimizing forwarding delays for certain priorities of packets called Low Latency DOCSIS (LLD). This is an evolutionary endpoint along a vector of buffer management trials which produced the following signatures:

(CM/CMTS RTT)			
	<i>When Idle</i>	<i>Under Load</i>	<i>99th Percentile</i>
<i>DOCSIS 3.0 Early Equipment</i>	<i>~10ms</i>	<i>~1000ms</i>	<i>~1000ms</i>
<i>DOCSIS 3.0 w/ Buffer Control</i>	<i>~10ms</i>	<i>~100ms</i>	<i>~100ms</i>
<i>DOCSIS 3.1 Active Queue Management</i>	<i>~10ms</i>	<i>~10ms</i>	<i>~100ms</i>
Low Latency DOCSIS 3.1	~1m	~1ms	~1ms

Figure 5 – Evolution of DOCSIS pre-4.0 Roundtrip Time Latency

Fundamentally, LLD calls for a parsing of duplex service flows into LL and Classic – with a management function to apply service classifier membership to discharge data flows identified as not requiring rapid forwarding from the LL flow into the Classic mix. It also monitors queue health along both flows to make sensible, weighted exploit of the low latency path under dynamic loading conditions.





Figure 6 – 2-Component Flows in DOCSIS Pipe, Low Latency and Classic

With this simple expedient of a second, shunt buffer queue for high priority/low latency traffic, the following remarkable results are obtained in the DOCSIS domain:

High Classic Traffic

p99 Latency vs USSG Utilization
(90% classic Traffic)



Tmax: 20 Mbps

USSG Capacity: 200 Mbps

High Classic Traffic

Jitter vs USSG Utilization (90% classic Traffic)




Figure 7 – Differential Flow Performance Upstream for Latency and Jitter

As can easily be identified, both latency and its deviation (jitter) are remarkably constrained once LL discipline is applied in a second service flow (typically 10 msec RTT latency with 1 msec jitter, across a service group loading from 25% to 90%).

So what services require low latency assistance? The following attempts a parsing of typical home data services into “LL” and “non-LL”:

Flow Mapping for Typical Data Services

Low Latency	Bulk/Background
Web	Software Updates
Videoconferencing	Dropbox
Audio Streaming	Email
VOIP	VPN*
Gaming	Video

** may also require LLD treatment*

Figure 8 – Candidate Service Mounts for Low Latency Grooming

Essentially, then, rate tiers provide a means of selectively monetizing progressively more aggressive user bitrate demands and this type of selectivity can now be applied to the latency domain and enhance those services requiring preferential, low latency handling.

4. Extant and Emerging Premise Services Landscape

Any well-founded attempt at aligning network data delivery to, and scavenging from, home premises necessarily must consider both present (established) service mounts and those calculated to emerge and mature – or wither and go dormant -- over the five-year consideration horizon posed in this paper to supply foundation for modification of network behaviors. A general expectation seems to be that data flows will increase (and on a duplex basis) between premise and cloud; but as always there needs to be sober justification for the timing and magnitude of capex investment to enhance data flows versus the life cycles of legacy and emerging services and the migration of opex balance sheet operating points (typically subscription revenues against licensing or other service origination/maintenance costs). This

section will attempt to assign relative values to services which are expected to populate the network application service palette in the postulated time window.

4.1. Consumer Impressions of Wi-Fi Performance and Feature Value

A proper sampling of present network performance service levels and features – as perceived by the end users themselves – seems an appropriate and necessary first consideration in establishing any motivation for alteration of the current state of network data operations. This can also provide the rational basis for critique of the relative end perceived value of various modification strategies. And providing a voting mechanism on alternate value strategies would help establish success vectors (from the users' standpoints) regarding where to steer the performance of the combined cloud-to-ground (and back) environment experienced by these end users. Such a survey was constructed and the details of the canvassed topics – and the recovered results from 546 global participants – follow.

The regional breakout for the data – which sources from the CommScope associates database – exhibited the following detail:

Asia-Pacific	250
North America	187
Carib./Latin America	21
Europe	80
Africa	2

Figure 9 – Regional Basis for User Network Impressions

The survey attempted to identify users' appreciation of the common metrics used to described network data attachment and assess those parameters and values which resonated from a value standpoint. To orient some of the subscription realities, users were asked their broadband speed tier:




Figure 10 – Users' Broadband Speed Tiers

Note the gross preponderance of sub-Gbps subscription, with the plurality lying in the 200-1000 Mbps bucket. Cost for this privilege was solicited:




Figure 11 – Budgeted Monthly Cost for Network BB Data

Note that globally, caps appear at 1 Gbps speed and \$100/month cost. In terms of satisfaction with this service exchange rate, nearly 79% of respondents indicated they were happy with the value proposition – perhaps underpinned by the fact that just over ¾ of the users had an alternate supplier handy (so,

presumably, might have done a “best offer” type of alternative analysis). To add to the inertia, nearly 87% of those answering defined themselves as not actively looking for replacement broadband services.

However, in terms of total network support, just over half (56%) of respondents felt that their home Wi-Fi matched the performance of their WAN attachment – and 62% felt that, were investment in their total networking were to be made, that investment would fall to the Wi-Fi side. (There appears to be some legacy ill will or suspicion associated with Wi-Fi – perhaps created in bygone single-band days, or with a preponderance of 2.4 GHz based, single antenna clients. Certainly, Wi-Fi beyond dual-band MAC 5 – 802.11ac – has no problem with disposition of packets at the less than 10 Mbps throttling rate which represents average WAN attachment pace for dwellings!) And it is true that, wherever network issues arise, the human proximity of Wi-Fi CPE lends itself to that nearby equipment shouldering blame for connectivity issues (whether or not they originate there).

In the next section of the survey, users were tested as to specific opinions on upstream and downstream data – the latter both regarding speed and latency/jitter. First, as regards upstream connection speed:




Figure 12 – Upstream Bitrate Rating

Note that nearly half the results peg the upstream as being either “good” or “excellent”; whilst an uplifting opinion, one wonders if users understand how little of their connectivity traffic is actually carried upstream – or how to quantify its “goodness”. Still, their resolution on this point was tested by the alternative to either stay with their current provider at a 10% price cut or move to another provider for twice the speed at the same cost they currently experience. The result was telling:




Figure 13 – The Upstream Value Challenge

Only about a third of respondents could see the worth in switching up suppliers for twice the speed at identical cost to them; still, this shows some fair sensitivity to upstream performance value – so it is noteworthy the majority push would be for less cost. (This opposite view holds in high connectivity speed regimes, so clearly even when satisfied with downlink speeds, there is opportunity to mine interest in better upstream capability). Next, the same question was posed regarding downstream performance. The response shift is interesting:




Figure 14 – The Downstream Value Challenge

Note that roughly 10% of respondents opted up for more downstream speed than were prepared to do so for upstream improvements. This suggests that they are aware of some correlation between DS bitrate and in-home services performance (while perhaps not being convinced that it is worth that much more than getting a quick 10% service cost rebate for the extant DS rate). Note that the higher the subscription speed, the more likely the user was to push for cost mitigation than more bandwidth (obviously viewing their connection as at least adequate).

The survey then shifted to probing users' feel for the value of latency and jitter in their network performance; first, as to latency:




Figure 15 – Where Latency Fits for Users VS DS Bitrate

It is difficult to infer anything here: the plurality of users would have latency rank equally with bitrate as regards connection “goodness” but the weighted average of value from this chart produces no other insight – other than high bitrate users REALLY appreciate their high bitrate tier more than latency. Perhaps unsurprisingly, the jitter evaluation by users produces a near-mirror copy of the latency result:




Figure 16 – Where Jitter Fits for Users VS DS Bitrate

If these results were perplexing, perhaps they appear less so when a litmus test for the magic effect of high bitrate labeling is concocted. Users were asked if, rather than receive a guarantee of bitrate, they would instead prefer a promise that all their equipment would be set to operate under the best bitrate, latency and jitter conditions possible, an astounding 2/3 replied “No”. When this offer was sweetened with the inducement of a 10% rate reduction, the “yes” responses did increase – but only about 10 percentage points (to ~44%). So there is obviously a great deal of marketing value to network connectivity distilled to a single rate tier value. Given this reality, the respondents were quizzed as to what speed grades they looked for: the lowest to do the job, the highest offered or something in between. The responses came down as follow:




Figure 17 – BB Connection Rate Impetus to Shopping

To add a layer to this query, the users were asked if new broadband services (specifically like VR) factored into their purchase thinking for network connectivity. The responses came thus:




Figure 18 – VR impact to BB Connectivity Requirements

Nearly 70% of users indicated a high level of sensitivity to high bitrate demanding services emerging in their homes and appeared prepared to include consideration of these in present and future BB purchases. (This is a critical leverage point, as VR developments will very shortly push higher bitrates into the home, as we detail in following sections).

The users were then asked to rank the top attributes they examined when looking for broadband services; the rank order of the number 1 responses was distributed as follows:




Figure 19 – Most Important Broadband Network Attribute

Cost, reliability and topline bitrate were noted – but oddly enough, Wi-Fi performance was seen to count for inclusion in broadband network performance as well. In terms of home uses for WAN services, the following were cited:




Figure 20 – Broadband Network Enabled Home Services

Unsurprisingly, watching streaming video set the bar, but in the post-pandemic world, working from home and teleconferencing received high notice. Not yet too noteworthy, but of interesting acknowledgment are the votes cast for remote healthcare; expect these to rise significantly as AIP takes hold.

A final query had to do with concurrent users in the home. Given the US household average of 2.6, perhaps these numbers are not too far-fetched:




Figure 21 – Simultaneous User Leverage of WAN

The overall impression of users seems to be that they are very satisfied with their broadband investment; they value cost, speed and reliability highly (but are largely clueless on the value of latency and jitter); spend the largest chunk of time ingesting streaming video (still) and are curious about the impact higher-value video services may have on their networking environment (in which they link Wi-Fi performance directly to that of the WAN).

4.2. Premises Services Data Payloads

For the next five years (and in fairness, the conceivable future) video services will continue to represent the largest tranche of downstream service data consumption. This is not one's grandparents' video experience; broadcast channels have morphed to applications, with the appended '+' symbol an indicator of viewing options which exist underneath a brand headliner. But however the content is catalogued and presented, it is clear that lean-back video consumption will remain a lynchpin service for almost every subscriber given that current consumption metrics peg video as burning close to 80% of premise-bound downlink data payloads. And while the present HD resolution norm for premium video meters at ~ 5 Mbps given present codec efficiency and common rendering devices, growth of in-home 8K-capable TVs promises to boost the video data appetite to something in the realm of 50-100 Mbps per screen (at coming codec efficiencies and higher frame rates) in the very near future. (It is, in fact, the rate of consumer uptake of higher quality video services which most challenges the forecast of any premises data consumption budget.) In terms of perceptible wireless impact, video picture quality is also the most relatable quality metric of wireless performance accessible by the average non-technical user and provides an obvious marketing leverage point in side-by-side comparisons of network bitrate value. (Imagine the marketing penalty associated with the realization that network A can mount the latest VR

headset and network B cannot.) Pegging an inflection point for adoption of super-resolution and frame rate rendering devices seems fraught with false starts, but perhaps an understanding of what is implied (in terms of network performance) might serve to frame the coming expectations:

Bandwidth Requirements of Cloud VR

- The estimated per-user bandwidths required by strong-interaction services in the three phases of Cloud VR development are as follows:

[1]

Phase		Fair-experience Phase	Comfortable-experience Phase	Ideal-experience Phase
Typical content resolution		2K (equivalent full-view resolution: 4K)	4K (equivalent full-view resolution: 8K)	8K/16K (equivalent full-view resolution: 12K/24K)
Typical terminal resolution		2K	4K	8K/16K
FOV		90° to 110°	120°	120° to 140°
Color depth (bits)		8	8	10~12
Coding standard		H.264/265	H.265	H.265/266
Compression ratio (I-frame/P-frame)		25/75	38/165	50/255(8K), 83/585 (16K)
Strong-interaction VR service	Typical bitrate	40 Mbit/s	90 Mbit/s	Full-view: 290 Mbit/s (12K) 1090 Mbit/s(24K) FOV: 155 Mbit/s (12K) 580 Mbit/s(24K)
	Typical bandwidth requirement	80 Mbit/s	260 Mbit/s	360 Mbit/s (8K) 1.5 Gbit/s (16K)

 **Note relentless increase in video bitrate**

Figure 22 – Stacked Bitrate Increases Driving VR Simulation Immersion and Super-Resolution Screens

A distant, if otherwise distinct, common second service type which may be reasonably expected to continue (if not grow) its ranking among services is network gaming. (This includes immersive AR/VR equipment such as Oculus – the 80% market share leader – which is projected to exhibit a CAGR between 40 and 60% over the time period 2018 to 2025). This compound growth suggests an exit value from our 5-year window at north of \$20 billion dollars of sales revenue for AR/VR. The desired network attachment profile for these types of gaming services includes a latency component which is two orders of magnitude tighter in tolerance than that required for ABR delivery of video packets for lean-back viewing. The desired jitter descriptor is a further order of magnitude more restrictive than that (so low single-digit milliseconds). Present Oculus data burn mimics that of a moderate resolution planar video

screen but upcoming versions (described above) seem to intend on increasing resolution and pushing frame rates, *a la* 16K TVs. (And recall that VR/AR equipment paints not just frame content straight ahead, but also prospective adjacent frames to the sides, above and below, to reduce rendering stutter in rapidly moving head exercises). As such, it should be expected that near-term-available Oculus II and III headset gear will exhibit an even higher data appetite (from 50 Mbps to perhaps as much as 400-500 Mbps) than those 8Kp120 UHD screens which lie largely unexercised in present home viewing rooms.

AR/VR systems aside, dedicated gaming stations and PCs with legacy planar rendering surfaces will continue to see heavy use in the residences of end users during the projection period. As is the case for immersive (head-worn) rendering systems, the twitch aspects want for predictably low latency with corresponding small amounts of packet delivery jitter. In terms of raw bitrate support, however, network gaming peaks at 150 MB/hour (~ 300 kbps). As regards pipe demand, this is on the same order of magnitude as listening to streaming audio services or conducting an IP-based phone call.

An emerging video option, yet ungraduated from boutique levels of interest and deployment, is holographic imaging. Requiring multiple video planes to construct, it represents the largest potential downlink consumption in the residence (if such is so equipped). Although not yet representing notable consumer interest, it seems an inevitable (if expensive) extension along the vector of greater immersive simulation or learning. Current state-of-the-art requires a data downpipe of nearly 2 Gbps capability – several multiples of present-day “full up” home services network attachment – so when adoption becomes real, such will manifest a huge uptick in network connectivity. However, given the abbreviated window of expectations for this study, the marketing outlook for holographic support in the home may be presumed to be beyond the time scope of this analysis.

Five other, lower bitrate and less stringent latency-invoking, service profiles can easily be identified: website browsing, digital assistant interplay, live streaming, SOHO teleconferencing and IoT telemetry forwarding. Of these five, the latter three represent the largest upstream bitrate generators: camera feeds. Depending on video resolution and codec, these can vary from 500 kbps to 2 Mbps per feed and (constant rate teleconferencing aside) the common use case involves intermittent (triggered) operation, involving transfer of perhaps 7.5 MB of video/audio clip data over a 30-second capture window. From a security standpoint, it would not be unusual to expect camera coverage of all premises ingress/egress points, resulting in perhaps 4 streams worth of randomly triggered data per dwelling. In general, however, the total contribution to network connectivity (and upload at that) might only see 5 or 6 episodes dispatched per 24-hour period (so 40 MB total security payload). The teleconference load depends on meeting times (naturally), but six hours of daily meetings would contribute 5.4 GB/day worth of upstream traffic. Live streaming (Tik Tok and the like) can vary from short bursts to sustained use on the order of a teleconference hit per mounted instance.

Digital assistant interactions process audio clips of typically 48 kbps sampling rate. Assuming a multichannel far-field array capture, one could budget 400 kbps uplink and perhaps 200 kbps downlink traffic, representing around 6 Mb uplink and 1.5 Mb of downlink utilization over a 2-minute, randomly occurring interaction cycle. As with security cameras, assistants will be sprinkled throughout the home (and recall these can be apps on multiple types of CPE, including phones, remote controls, audio players and TVs). A dozen of these “conversations” among all household members might be a reasonable estimate for maximal usage during a day, which yields much less than 10 MB uplink / 2 MB downlink total payload over the course of a day.

Smart home IoT command exchanges and telemetry uploads over Wi-Fi tend to be even less of a network burden than smart assistant interchanges. If we exclude the voice-ordered aspects and consider their impact captured above, telemetry and commands might amount to no more than 30kB (up and down)

worth of daily interactions; though what is critical is timely forwarding of these exchanges into, and out of, the premises, given the potential emergency nature of security or health alerts included in the payloads. (From a behavioral point of view, persistent and rapid-rate-signaled “alarm clamoring” – until such is acknowledged by the cloud receiving entity – would seem to make sense). This also presumes verbose loop-in of a cloud element, though what is expected is that routine executive control of the premises (as regards IoT) would not likely exit the home edge (and such would be hardened for loss of power and wireline WAN connectivity, at that).

4.3. Data Demands of a Model Smart Home in 2027

Combining several of our expectations, expressed above, for data interchange between home LAN and network WAN, allows us to posit the following model for a smart home in 2027. Note that we will be leveraging Arris’ 5300 square foot Wi-Fi house as the representational challenge and will outfit its ~ 2x average home footprint with a full array of CPE Wi-Fi (including sub-mesh Matter support for SED – battery powered -- IoT devices, arranged so that Wi-Fi acts as aggregating long haul for the Thread edge routers commanding these distributed resources). Note that MLO triband capability will be assumed to be implemented in all link endpoints (AP and client alike). We will also liberally equip the home with A/V kiosks (2 per floor, say – though representationally, these might well be phone applications) to support voice control of IoT devices from throughout the premises.

From a human-centric viewpoint, we can suppose nine concurrent clients at peak network utilization, perhaps arranged with the following locations and application demand profiles:

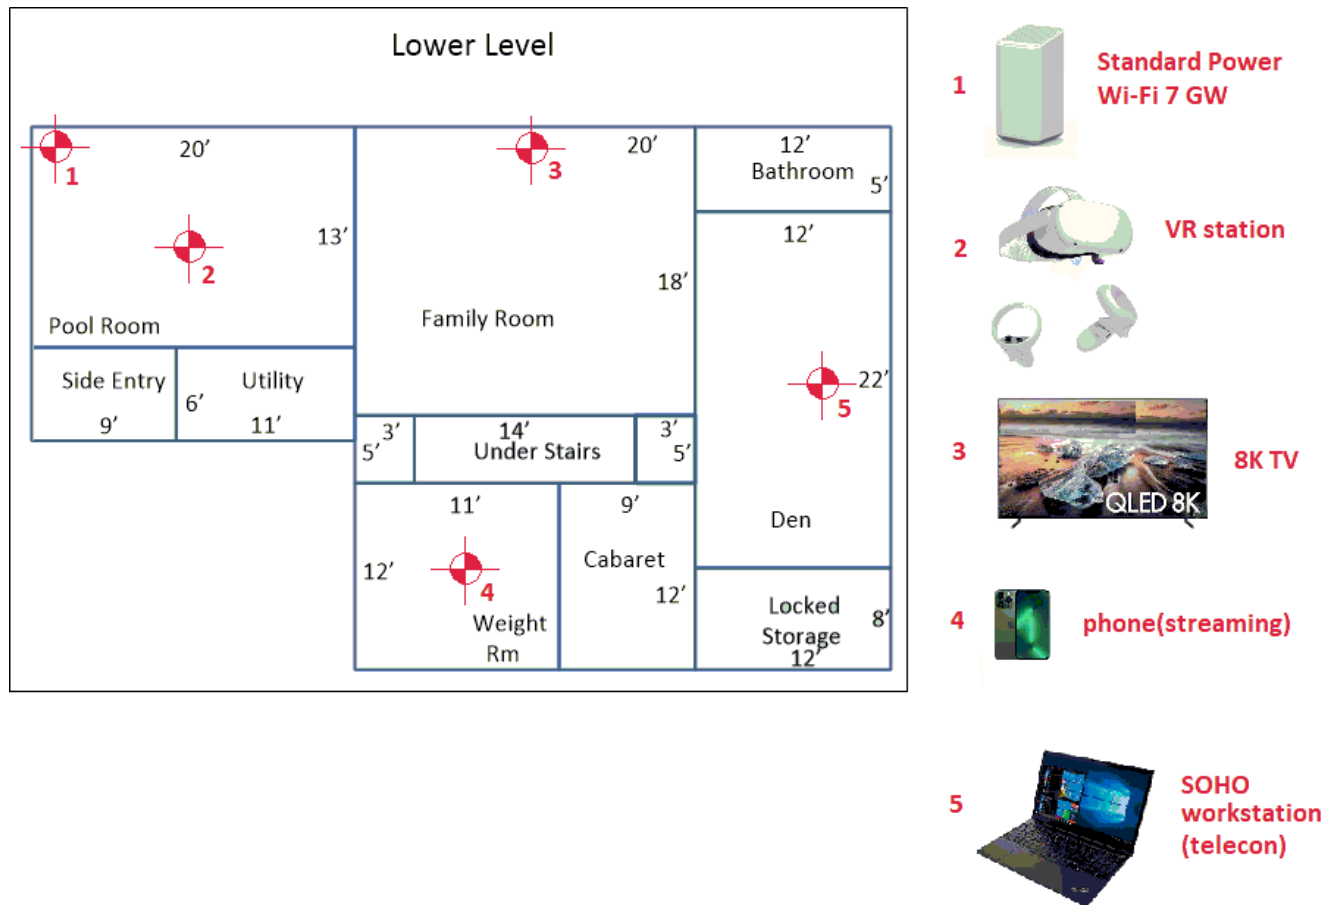
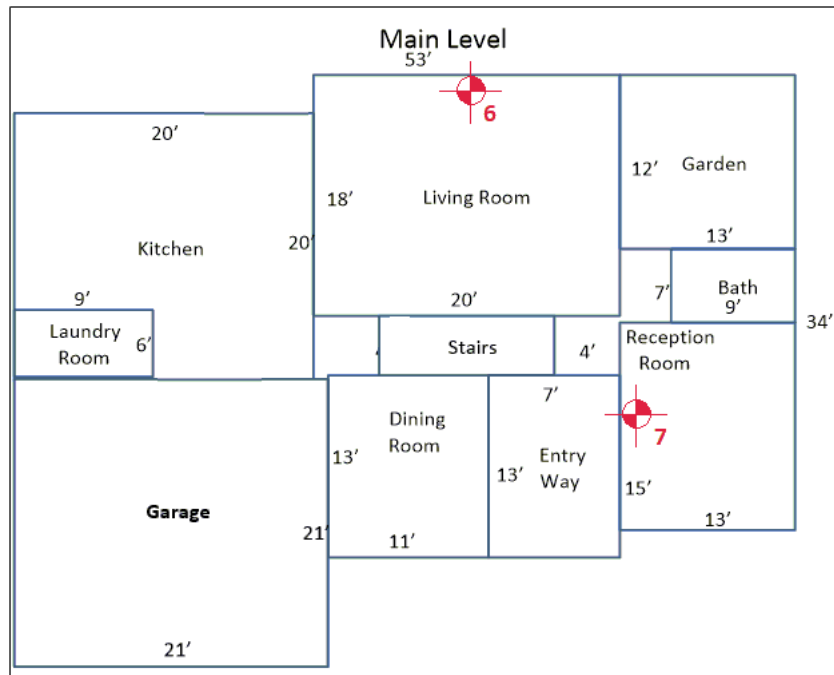


Figure 23 – Wi-Fi House Basement Device Outfitting



6



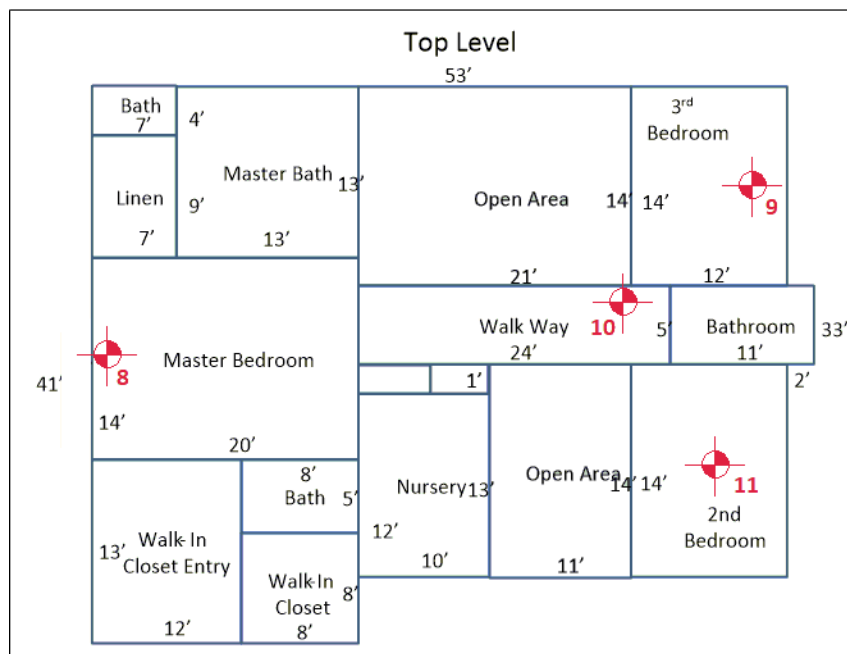
8K TV

7



iPad gaming

Figure 24 – Wi-Fi House Main Level Device Outfitting



8



8K TV

9



Playstation

10



Extender

11



phone

Figure 25 – Wi-Fi House Top Level Device Outfitting

So we can imagine just north of 250 DS Mbps in mounted services, roughly half of which would require the attentive grooming of low latency flow management. In addition to these dynamic client mounts, the home features external security family feeds, internal kiosks for voice command capture and a multi-drop

Thread mesh (for Matter IoT support of automated home security items such as door locks, lighting and garage door control, among others). Note that this latter, alternate-PHY mesh seeks backhaul support via Wi-Fi to the hub (edge) control functions set up in the WAN gateway.

The model home network infrastructure will adopt forward-leaning technologies and so will exhibit a main trunk extension from WAN entry point in the basement game room corner up to the third-floor juncture between the Jock and Jill bedrooms. This will feature a Standard Power and 320 MHz BW channel leverage (at 6 GHz) which will be time-split for attached client MLO/MRU support (essentially dividing client servicing duties with the WAN GW, such split determined by best sustained MCS between AP and client). As all clients will be presumed '11be capable, scheduling behavior will be expected to optimize (minimize) Wi-Fi airtime burn through exploit of appropriate BW in each of the three Wi-Fi bands. Note, however, that for MLO operations, only bonding of 5 and 6 GHz bands will be considered (conservatively capping the available Wi-Fi 7 connections).

5. Performance of the Adjoined Networks

Now we arrive at a discussion of our highly outfitted model home networked to the cloud via its DOCSIS 4.0 WAN wireline attachment. For the purposes of stressing the network attachment, the intention is to array the model with a rather large assortment of Wi-Fi dependent CPE (as described above) and amplify this data aperture by indulging multiple concurrent users with disparate services. We will examine the network behaviors in reverse order, starting with the LAN.

5.1. Stressing the Wi-Fi 7 Home Network

Given the described locations for the types of clients and their respective service mounts, the predicted airtime distribution costs for support of the services array gives us the following:

Device	Location	AP	Path Loss	Link Capacity	Service Bitrate	Service Latency	Service Jitter	Airtime %	Low Latency?
VR station	downstairs pool room	Gateway	10.8 dB	6500 Mbps	100 Mbps	5 msec	2 msec	1.54	yes
8K TV	downstairs family room	Gateway	24.2 dB	5250 Mbps	50 Mbps	250 msec	50 msec	0.95	no
phone (streaming)	downstairs weight room	Gateway	35.5 dB	1800 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.11	no
laptop (SOHO, teleconf)	downstairs den	Gateway	33.3 dB	2850 Mbps*	2 Mbps (up)	5-10 msec	2-3 msec	0.07	yes
8K TV	middle level, living room	Gateway	33.8 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
iPad, gaming	middle level, reception room	Extender	25.5 dB	4950 Mbps	350 kbps	10 msec	3 msec	0.007	yes
8K TV	upstairs, master bedroom	Extender	26.1 dB	4900 Mbps	50 Mbps	250 msec	50 msec	1.02	no
Playstation gaming	upstairs, bedroom 2	Extender	13.1 dB	6500 Mbps	350 kbps	10 msec	3 msec	0.005	yes
phone (streaming)	upstairs, bedroom 3	Extender	18.1 dB	4500 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.044	no
4SS MLO Trunk (5G+6G)	pool room to upper landing	AP/EXT	35.9 dB	8250 Mbps	52.7 Mbps	< 2 msec	< 1 msec	0.64	yes
Total, Gateway:				*client upstream EIRP limited				4.42	
Total, Extender:								1.716	
Client Bitrate Service					256.7 Mbps				

Figure 26 – In-Home Wi-Fi 7/Standard Power (with Extender) Network Performance

As can be seen, even restricting the Wi-Fi 7 MLO to bonding only 5 and 6 GHz channels (albeit, at 160 MHz and 320 MHz bandwidths, respectively), the in-home network is not at all appreciably taxed by the 256 Mbps worth of mounted client services. And it should come as no surprise that the intermittent demands of multiple security camera feeds and periodic polling of the Thread subnetwork do not impinge on operations in the measurable slightest. The leverage of standard power along with a quad-band Wi-Fi 7 extender amounts to extreme futureproofing; the WAN gateway commits to just over 4% airtime serving all its clients and the 4SS trunk to the upper far end of the home, while for its part the extender located there commits less than two percent of its time transmitting. Note that, excepting the trunk peer-to-peer 4SS link, all other links amount to 2SS connections (given client radio restrictions); and the 3

uplink feeds' bitrate capacities are all bound by client EIRP limitations (battery-driven, in the two phone cases).

These distribution costs seem extremely well controlled. What if we restrict the gateway and extender to only LPI levels at Wi-Fi 7? The following results:

Device	Location	AP	Path Loss	Link Capacity	Service Bitrate	Service Latency	Service Jitter	Airtime %	Low Latency?
VR station	downstairs pool room	Gateway	10.8 dB	6500 Mbps	100 Mbps	5 msec	2 msec	1.54	yes
8K TV	downstairs family room	Gateway	24.2 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
phone (streaming)	downstairs weight room	Gateway	35.5 dB	1750 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.065	no
laptop (SOHO, teleconf)	downstairs den	Gateway	33.3 dB	2900 Mbps*	2 Mbps (up)	5-10 msec	2-3 msec	0.055	yes
8K TV	middle level, living room	Gateway	33.8 dB	3425 Mbps	50 Mbps	250 msec	50 msec	1.46	no
iPad, gaming	middle level, reception room	Extender	25.5 dB	4500 Mbps	350 kbps	10 msec	3 msec	0.008	yes
8K TV	upstairs, master bedroom	Extender	26.1 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
Playstation gaming	upstairs, bedroom 2	Extender	13.1 dB	6025 Mbps	350 kbps	10 msec	3 msec	0.006	yes
phone (streaming)	upstairs, bedroom 3	Extender	18.1 dB	2250 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.089	no
4SS MLO Trunk (5G+6G)	pool room to upper landing	AP/EXT	35.9 dB	6000 Mbps	52.7 Mbps	< 2 msec	< 1 msec	0.88	yes
Total, Gateway:				*client EIRP limited				5.11	
Total, Extender:								2.093	
Client Bitrate Service					256.7 Mbps				

Figure 27 – In-Home Wi-Fi 7/LPI only (with Extender) Network Performance

Note that all considerations are still very well met, with only a slight ballooning in airtime for the two AP devices (certainly acceptable, given that the limits at ~ 5% and 2% are still low single digits). So it appears that, with extender, the use of Wi-Fi 7 at the two endpoints of a trunk link limited to LPI EIRP would still work very well. This begs the removal of the extender, coupled with bumping up the single gateway AP (at the WAN injection point to the home) to Standard Power to see if coverage would still be possible:

Device	Location	AP	Path Loss	Link Capacity	Service Bitrate	Service Latency	Service Jitter	Airtime %	Low Latency?
VR station	downstairs pool room	Gateway	10.8 dB	6500 Mbps	100 Mbps	5 msec	2 msec	1.54	yes
8K TV	downstairs family room	Gateway	24.2 dB	5250 Mbps	50 Mbps	250 msec	50 msec	0.95	no
phone (streaming)	downstairs weight room	Gateway	35.5 dB	1750 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.11	no
laptop (SOHO, teleconf)	downstairs den	Gateway	33.3 dB	2900 Mbps*	2 Mbps (up)	5-10 msec	2-3 msec	0.069	yes
8K TV	middle level, living room	Gateway	33.8 dB	4500 Mbps	50 Mbps	250 msec	50 msec	1.11	no
iPad, gaming	middle level, reception room	Gateway	38.0 dB	3700 Mbps	350 kbps	10 msec	3 msec	0.0095	yes
8K TV	upstairs, master bedroom	Gateway	41.5 dB	2575 Mbps	50 Mbps	250 msec	50 msec	1.94	no
Playstation gaming	upstairs, bedroom 2	Gateway	42.3 dB	2450 Mbps	350 kbps	10 msec	3 msec	0.014	yes
phone (streaming)	upstairs, bedroom 3	Gateway	47.6 dB	355 Mbps*	2 Mbps (up)	15-20 msec	5 msec	0.56	no
Total, Gateway:				*client EIRP limited				6.3025	
Client Bitrate Service					256.7 Mbps				

Figure 28 – In-Home Wi-Fi 7/Standard Power (no Extender) Network Performance

As can be seen, whole home coverage is still doable – though note the noticeable burn of airtime just to enfranchise the Tik-Tokker in the far upper bedroom. This type of airtime stress will always be the worst – trying to reach a battery-powered device a large distance from the AP and mounting a service thereon which produces upstream data (instead of consuming downstream) – hence being limited to the link MCS achievable with the relatively weak transmitter of the client device.

Still, in all these solution cases, a relatively pathological mix of concurrent services can be mounted and supported by the in-home wireless network built of Wi-Fi 7 componentry. Now the question shifts: can the DOCSIS 4 WAN adequately support these client mounts in reasonably scaled service groups – and of what particular DOCSIS 4 species must we avail ourselves?

5.2. Taxing the DOCSIS Wireline Network

From historical trends, the appearance of a 250 Mbps persistent (and pervasive) home network demand in a DOCSIS service group (SG) is not remotely projected through 2037, based upon current modeling. (Reference ULM et al 2022). Which is not to say it could not exist, merely that capacity planning using present tools has no historical basis to predict the arrival of multiple CPE with bitrate service demands at, or exceeding, 50 Mbps each. This is cataclysmic disruption, relative to present architected capacity and projected bitrate growth; at issue is whether the weighted take-up of these types of bitrate demands within multiple homes comprising a SG produces statistical impact on network service structure before historical trends suggest it might. The best way to judge the scale of this impact, perhaps, is to note that with our present Tav_g value of 3.5 Mbps, UHD TV video consumption (the major single data sink per home in the network) produces SG sizes of perhaps 250 subscribers (with a forward BW of < 1 GHz). 250 Mbps represents an 80-fold increase in persistent home data consumption which would have to be defrayed by the combination of increased BW, denser spectral modulation leverage and smaller service groups. Succinctly, even the most aggressive deployment of DOCSIS 4 would find this an untenable challenge. So a bit less of an impact needs to be bitten off on first chew. (In fairness, nine concurrent users with four huge concurrent service consumers of DS data in a single HH is simply much-too-much of a service stretch example when it comes to taxing the WAN – especially if we claim prevalence in a SG, so we will be more circumspect from here on out.)

To establish a more realistic perspective (and to suppose that even 25 Mbps represented a reasonable Tav_g DS service goal – note the current value is 3.5 Mbps), we can refer to the Cloonan DOCSIS capacity planning tool for the impact to SG scaling. For reference purposes, we will suppose an ESD solution to 1218 Mhz with a diplex split at 204 Mhz. This gives us roughly 1 GHz of forward BW which, if committed to OFDM parsing only (no SC-QAM) at 9.7 bps/Hz, gets us fairly near the magical 10 GBps downstream capacity.

If we further conjure a top advertised SLA of 5 Gbps, the SG scaling is then set:

N_{sub} (the SG size) $\leq (10^9 - 1.2 \cdot 5^9) / 25^6$, or $SG \leq 160$ (if all are at the premium SLA tier – an admittedly unusual occurrence).

This is on the smallish size (range typically varies from 100-400) but at least reads as a very workable number. The 250 Mbps Tav_g, on the other hand, would produce an SG of only 16 under the same spectrum and SLA assumptions. To drive home the difficulty of this proposition, if we instead applied even a full 1.8 GHz DOCSIS network with only a modest diplex split of around mid-UHF (say 400 MHz), we would end up with forward capacity in the region of $9.7 \cdot 800 \text{ MHz} + 8.5 \cdot 600 \text{ MHz}$, or just shy of 13 Gbps. At the same SLA tiering, this would put us at an SG of roughly 28 – still untenable without reworking expectations. Doing so, however – perhaps looking to support a Tav_g of 50 Mbps and top SLA tier of 10 Gbps (a magical enough marketing icon) -- the SG moves to 60 (certainly a reasonable asymptote given that not all subscribers will select the premium tier, and the actual number will push into perhaps the 120+ range).

Without doubt, however, WAN network capacity becomes quickly tested when high resolution simulation environments become the norm on the client side; and the capacity “fat” at this juncture in the distribution network lies almost wholly within the LAN environment.

6. An OTT Sidebar

A brief acknowledgement of competitive (5G) forays into OTT connectivity for residential CPE in the US seems in order, if for no other reason than to inject some compounding considerations which might resolve teetering wireline network strategy decisions. As of this writing, two cellular ISPs have introduced 5G home gateway solutions, one a sub-6 GHz band version and the other a more urban-centric mmWave band variant. Given the full-USA coverage of these MNO/ISPs, it implies that both urban and suburban customers are to be granted immediate (as opposed to fiber's buildout-paced) access to premise internet coverage (as a 5G/Wi-Fi composition) which also exposes cable video consumption to co-option via streaming services.

Cable MSOs have carefully marked – and reacted to – telecom rollout of fiber. With fiber's pitch of symmetric 10Gbps connectivity (despite a lack of evidential service need for such BW) and further marketing forays into 25 and even 50 Gbps service tiering, cable has dutifully ramped up development of FDX DOCSIS 4.0 to at least place a bookmark in the derby for bragging rights to what – to-date -- may be called an unreasonably thick data pipe to the residential home. But competition for network connectivity bargains (and perhaps more concerning, in lower tiers of QoS) is already everywhere to be found, courtesy of OTT plays.

The appeal in both OTT cases lies in mining interest among these ISP budget-aware shoppers. The sub-6 GHz solution, for example, features a lowest bitrate tier which produces roughly 150 Mbps down/30 Mbps up with 25 msec latency inclusive of one Wi-Fi hop (one floor vertical) off of the modem for around \$50/month (the modem being a GW device which provides dual-band Wi-Fi coverage to the premises). Via contractual agreement for a couple of years, such attachment can produce similar performance at half that subscription rate. A typical midafternoon Ookla SpeedTest sample follows:

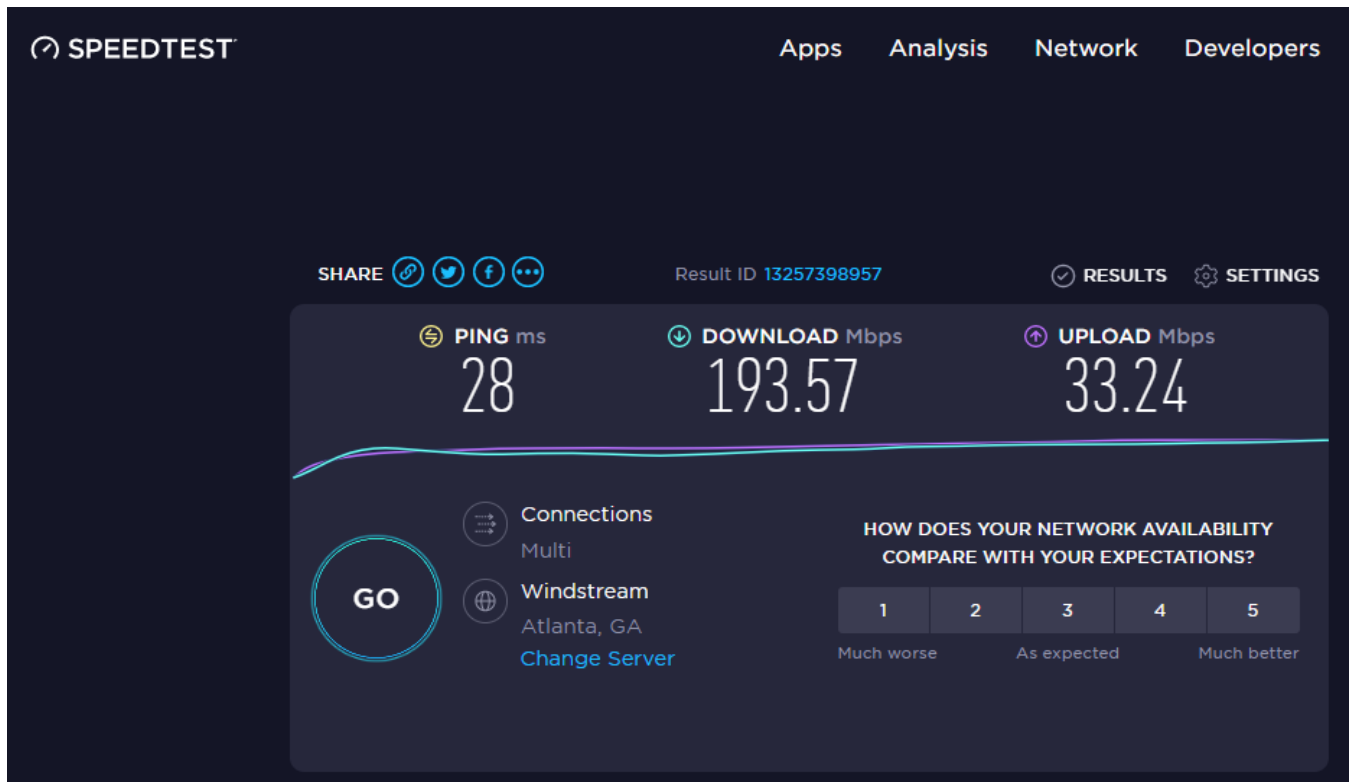


Figure 29 – Mid-day, one Wi-Fi hop OTT modem performance

In MDU and small home scenarios (with perhaps only 2 active Wi-Fi devices concurrently operating), such an arrangement can be seen to provide perfectly adequate internet connectivity, with the benefit of a single point of contact for both phone and network support and logistics. Awareness of this potential for internet fulfillment at less than eye-popping performance numbers ought to serve as a cautionary reminder to not ignore lowest QoS/bitrate markets, as opportunistic competition stands prepared to infiltrate same.

7. Conclusion

As has been shown, Wi-Fi premises bitrate budgets are much more than adequately covered for the near (and coming) term; even pathological, simultaneous adoption cases for extremely high resolution and high frame rate video displays are easily accommodated by home Wi-Fi emerging MACs during the coming five years. With proper leverage of low latency planning and scheduling tools – applied synonymously across the concatenated networks -- immersive simulation environments at modest resolutions can be created which suitably interdict artifacts in the constructed virtual domain(s) and facilitate mounting of rich gaming and instructive services. And there does not yet appear to be widespread adoption of an application or service which swamps the up- or downstream user bitrate flows over home Wi-Fi to the extent that inadequate QoS could be asserted; the home LAN, it seems, is extremely well positioned for all anticipated service mounts.

Packet latency conditioning seems acceptably met in both domains, with LLD capable of dual-piping the WAN and both WMM and buffer congestion management (by virtue of highly flexible scheduling) dissolving any congestion clots within the LAN; married, the E2E performance advertises itself to easily meet sub-20 msec type of E2E latencies just across these networks.

When it comes to raw bitrate, however, WAN prep seems perhaps not congruent with immediate market acceptance of high resolution, high scan rate immersive simulation environments becoming the sudden norm. If this type of precipitous advance in simulation entertainment occurs, then a stepped two-orders-of-magnitude increase (from a *sustained* sub-10 to multi-hundred Mbps) in on-premises average downstream bitrate demand will transpire and network capacity will be challenged as never before -- perhaps tainting even planned 10G/10G FTTP with an “inadequate” QoE estimation.

It's a gambler's paradise, perhaps; truly immersive entertainment options challenge the average home entertainment budget and so might be expected to suppress rapid adoption of the most compelling of lean-back media environments. We have been 8K-capable and pending appropriate source material for several years now, after all – granting WANs some respite in the process. And the CAPEX for system upgrades begs a lag-lag type of adoption curve, certainly. But the penalty for substandard network data support of high-end CPE involves an experiential litmus test which will be difficult to spoof if the long-awaited video nirvana does, in fact, take root. And, to be sure, just as the in-home distribution infrastructure of a very near adoption horizon seems easily capable of underpinning these viewing and gaming experiences, the wireline WAN which most commonly serves it seems, as-yet, resolutely under-capable.

Abbreviations

ABR	Adaptive Bitrate
AFC	Automated Frequency Coordination
AP	access point
AR	Augmented Reality
A/V	Audio/Video
BB	Baseband
bps	bits per second
BW	Bandwidth
CAGR	Compound Annual Growth Rate
CAPEX	Capital Expenses
CCI	Co-Channel Interference
C/N	Carrier-to-Noise
CPE	Consumer Premise Equipment
dB	Decibel
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
E2E	End-to-End
EIRP	Effective Isotropic Radiated Power
ESD	Extended Spectrum DOCSIS
FDX	Full Duplex
FEC	forward error correction
G	Giga
GB	Giga-Bytes
Gbps	Giga Bits Per Second
GW	Gateway
HD	high definition
HH	Household
Hz	hertz

IoT	Internet of Things
ISP	Internet Service Provider
Kbps	Kilo Bits Per Second
LAN	Local Area Network
LL	Low Latency
LLD	Low Latency DOCSIS
LPI	Low Power Indoor
MAC	Medium Access Control
MB	Mega-Bytes
Mbps	Mega Bits Per Second
MCS	Modulation and Coding Scheme
MLO	Multi-link Operation
MNO	Mobile Network Operator
MRU	Multiple Resource Units
Msec	Milliseconds
Mu-MIMO	Multiple User-Multiple In, Multiple Out
OBE	Overcome By Events
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OPEX	Operational Expenses
OTT	Over The Top
PC	Personal Computer
PHY	Physical Layer
PON	Passive Optical Network
QoE	Quality of Experience
QoS	Quality of Service
RTT	Round-Trip Time
RU	Resource Unit
SC-QAM	Single Carrier-Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SG	Service Group
SOHO	Small Office / Home Office
SLA	Service Level Agreement
SS	Spatial Stream(s)
Tavg	Average Bitrate
TOD	Time of Day
TWT	Targeted Wait Time
UHD	Ultra-High Definition
US	Upstream
USB	Universal Serial Bus
VoiP	Voice over Internet Protocol
VPN	Virtual Private Network
VR	Virtual Reality
WAN	Wide Area Network
WMM	Wi-Fi Multimedia

Bibliography & References

Broadband Capacity Growth Models, John Ulm, Zoran Maricevic, Ram Ranganathan, SCTE Technical Paper 2022



Creating Infinite
Possibilities.

Why, How, and Where to Converge Fixed and Mobile Networks

Bob Hallahan
Global Head of Cable Strategy
Nokia

AGENDA:

- ✓ Introduction
 - ✓ Current state of cable fixed and wireless networks
 - ✓ Drivers of convergence
 - ✓ Cable moves to 5G, and problem to be solved
- ✓ Why you should converge fixed, and wireless networks
- ✓ How to look at convergence
- ✓ Where and how to converge to achieve maximum value
 - ✓ Converging MVNO with cable 5G MNO Cores
 - ✓ Using the access gateway function (AGF) to converge broadband
 - ✓ Converging fixed and wireless voice services
 - ✓ Using a 5G adaptive core to achieve access network agnostic service delivery
- ✓ Why converging networks requires a new approach
- ✓ Summary

- Current State of Cable Fixed and Wireless Networks
 - **Network Centric Service Access:**
 - Fixed and wireless networks have historically been built with separate access, cores and OSS, making the job of managing and assuring services and customers experience in near real time exponentially more complicated and costly as services begin to span networks.
 - Services are not easily accessible seamlessly across network access technologies today

- Drivers of Convergence

- **Bundling:** Services and applications are becoming more bundled and agnostic to network access technology
- **Network Agnostic Services :** Network services need to be decoupled from the network access, to enable network agnostic delivery through a multitude of access technologies, devices, and locations with mobility at the center.
- **FMC (Fixed Mobile Convergence) :** The Distinction between fixed and wireless access will be blurred into a seamless blended access capability.
- **Improving Opex :** Initial driver behind network convergence in 2022 and beyond will be the need to offload wireless data from the MVNO 4/5G MNO network on to the MSO's 5G network
- **Market pressures and Digital Divide :** FWA (Fixed Wireless Access) Broadband with OTT (over the top) applications and content to out of network footprint residential and business customers, and Government subsidized network builds (RDOF)

- Cable Moving to 5G problems to be solved
 - **Cable Goes 5G :**
 - Deploying 5G networks is further complicating the need to manage access to Fixed/Wi-Fi, MVNO, 5G, Private, and partner networks.
 - Managing service assurance, reliability, quality, devices, and customers experience as service access becomes networks agnostic and devices are intelligently switching networks.
 - **Reducing complexity and cost:**
 - With the advent of 5G service oriented open architecture, it is now possible to eliminate redundant network and OSS functions and costs, while achieving service and device network access agnosticism

Why you should converge fixed and wireless networks

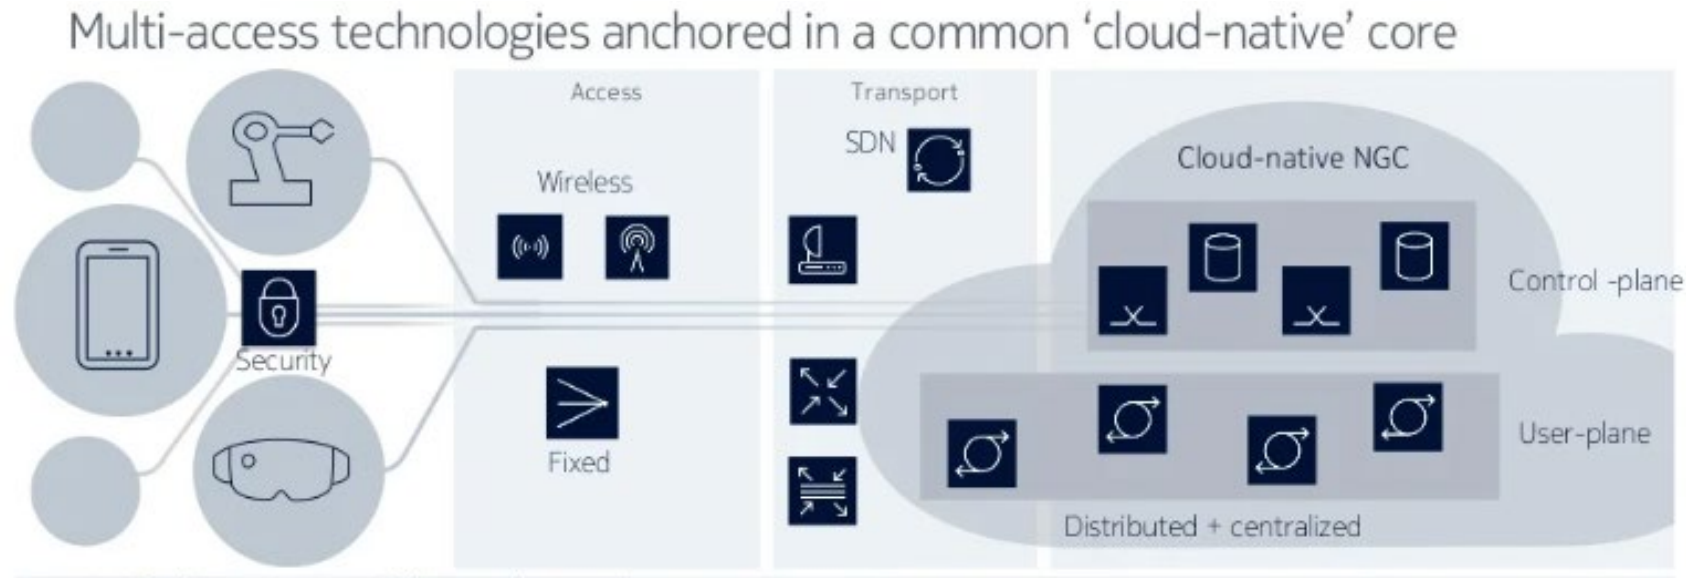
- **Understanding the Benefits of Convergence**
 - **Improve Opex** : Reduce high level of redundancy, cost and complexity in common converged layers
 - **Improve Service Quality** : Seamless service control and resilient continuity for all devices Improving ability to assure service quality and customer experience resulting in reduced churn and increased revenue, loyalty, monetization and retention.
 - **Improve Top/Bottom Line** : New Revenue generating opportunities in the enterprise market for converged private wireless, and the delivery of low latency applications at 5G speeds.
 - **Fast TTM** : In the Residential market it means for example new revenue streams from cloud gaming at the edge accessed through 5G slices, 4/5G MVNO and Wi-Fi seamlessly

How to look at Convergence

- **Agnosticism:** Cable fixed and wireless services and applications should be agnostic from the network access types fixed or wireless
- **The COREs :** Stitching Cores together enables session and management continuity, and moving to a universal adaptive core enables optimal convergence.
- **The Data :** Enabling fixed and wireless network operations to create and measure a universal set of KPIs (key performance indicators) will require a common shared data layer (SDL) containing an aggregated blended set of operational network data from Wi-Fi, 5G SA, MVNO, DOCSIS, PON, for all services under a single management umbrella.
- **B/OSS :** Fixed and Wireless OSS/BSS systems can be normalized via convergence of specific O/BSS functions to support multiple network and access types enabling Network agnostic customer experience management (CEM), Device Management (DSDS, eSIM/iSIM, and common service profile, entitlements, policy and assurance management.

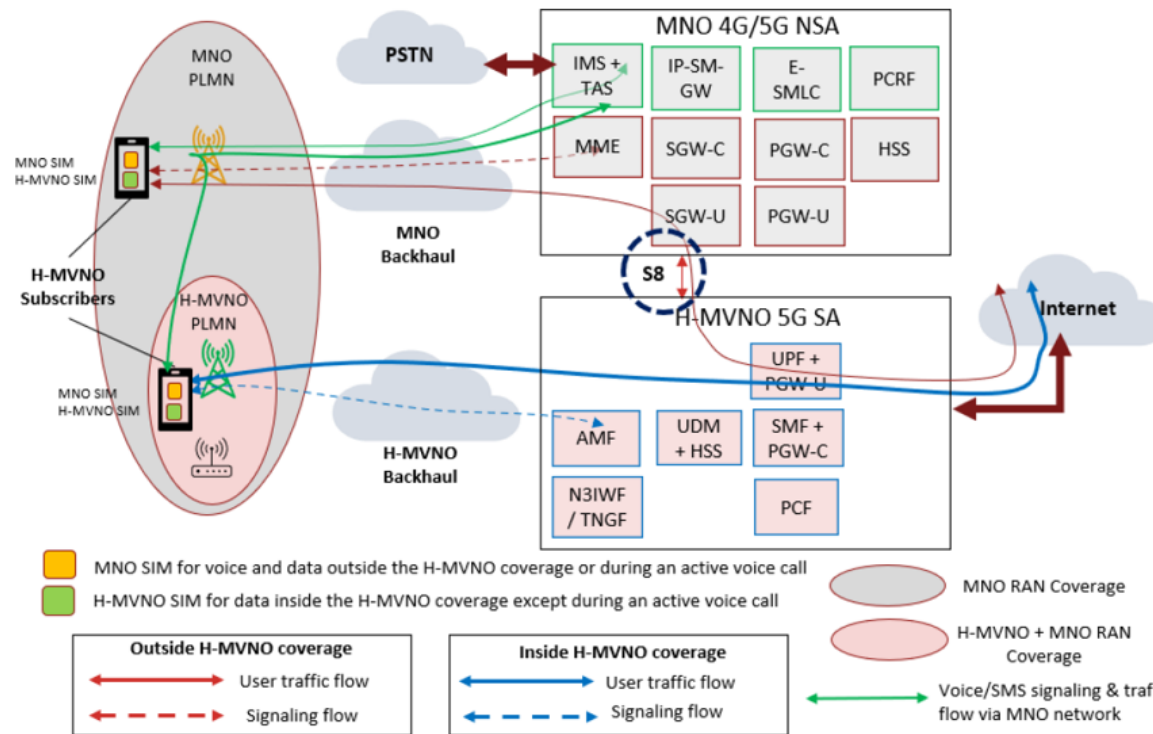
- Where to Converge for maximum Results and Value
 - **Convergence between the Cores**
 - Converging MVNO and MNO 5G Cores
 - Thin vs Thick MVNO models, enable “more knobs to turn”
 - Using AGF to Converge Data services Wi-Fi and 5G
 - Converging Fixed and Wireless Voice at the IMS Core
 - Achieve access network agnostic service delivery via universal adaptive core
 - **Convergence at the Data Layer**
 - Leveraging SDL to optimize network agnostic service access
 - **Convergence at the B/OSS layer**
 - Service Assurance Profiles and entitlements
 - Customer service experience Insights

- **Converging Fixed and wireless network Cores enables..**
 - Unified Session, mobility, and security management
 - Multiple simultaneous connections
 - e2E policy Service Profile management
 - Enhancing service QoS with core automation
 - Converged common multi technology transport layer



Where and how to converge to achieve maximum value

- Converging MVNO with Cable MNO 5G Cores



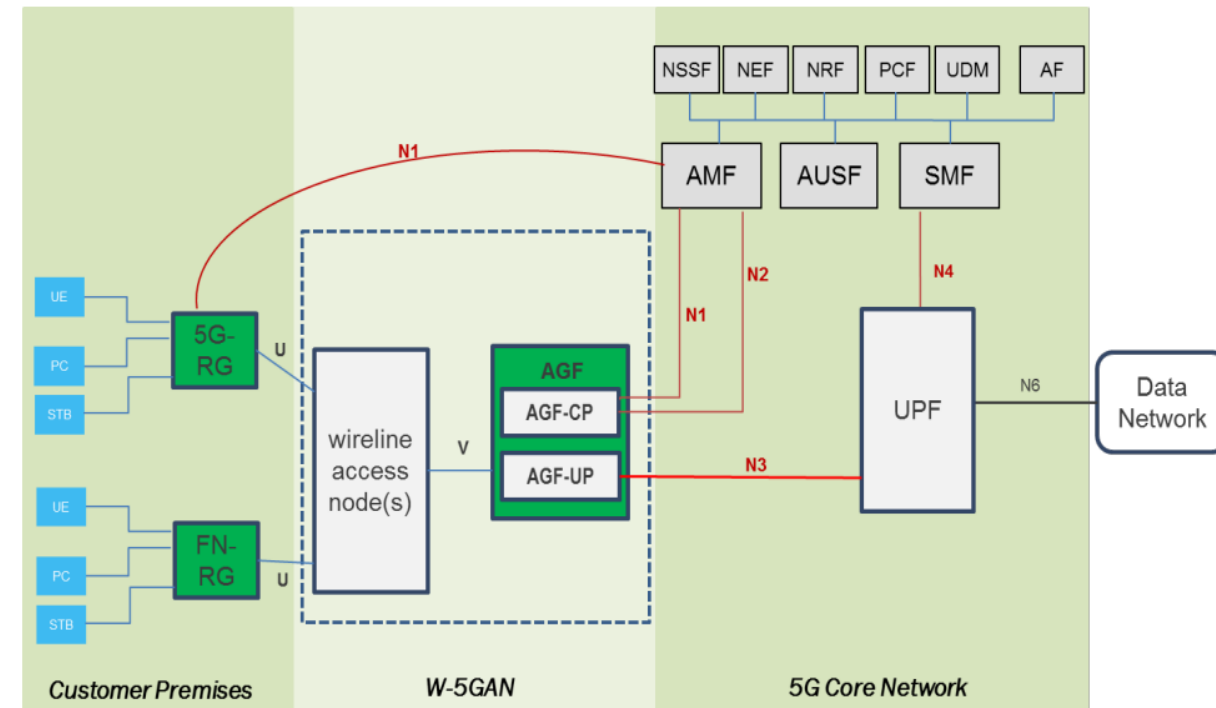
Recommendation from CableLabs using the S8 interface between SGW-U and the UPF/P-GWU to maintain session control and management.

Figure shows MNO network to be a 4G/5G NSA, but the architecture also applies to a scenario where both MNO and H-MVNO networks are 5G SA
The core network elements shown within the MNO and H-MVNO networks will use standardized interfaces

2021 SCTE Fall Technical Forum : Evolved MVNO Architectures for Converged Wireless Deployments

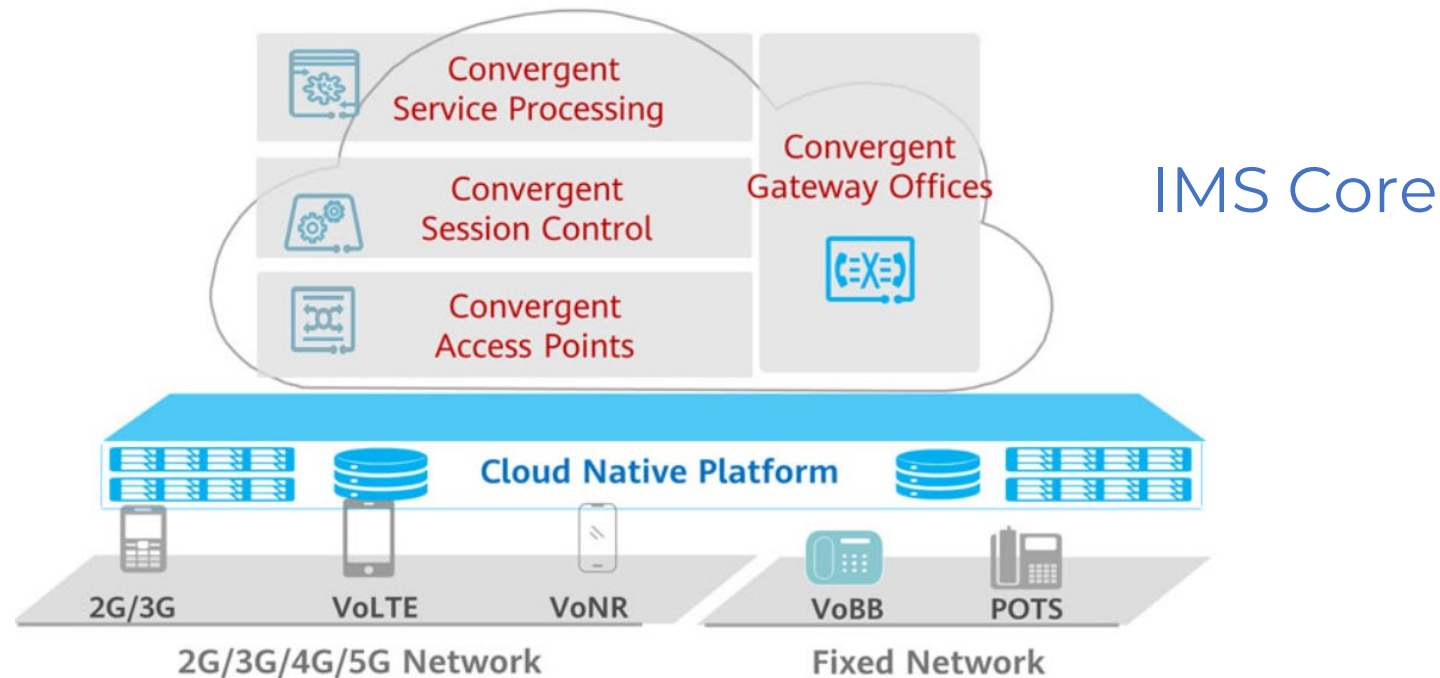
Where and how to converge to achieve maximum value

- Using the access gateway function (AGF) to converge broadband
- One of the main points of convergence is where devices access the networks.
- A key network function which enables fixed and wireless convergence is the access gateway function (AGF),
- The AGF controls network access for fixed networks to the 5G core, and is critical to the success of any network convergence initiative



Where and how to converge to achieve maximum value


- Converging fixed and wireless voice services
 - Convergence should be achieved at the IP Multimedia Subsystem (IMS) core level to bringing together wireline and wireless voice services across all access networks, which will enable centralized management and control of voice services access thru and spanning fixed and wireless networks.




Where and how to Converge to achieve maximum value

- Using an Adaptive Core to enable network access agnostic service delivery

Future Evolution to 5G Universal Adaptive Core
Flexibility to deliver any service over any access



Fixed-Mobile Convergence



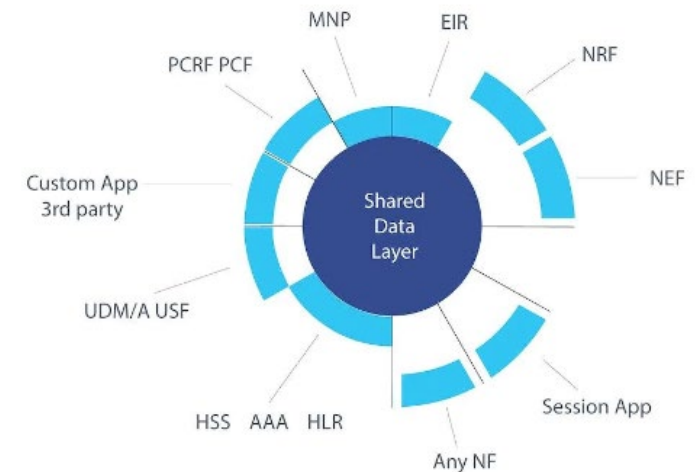

Fixed-Mobile Access Gateway

- Common SMF for fixed-mobile converged interworking
- UPF Packet Forwarding Control based on 3GPP Sx/N4 interface
- Dynamic UPF selection based on APN/DNN
- 5G services over any access!

- New services will require complete “connectivity agnosticism” with subscriber expectation of seemingly infinite service quality and ubiquitous connectivity.
- Enabling a massive scale in access, coupled to an evolved converged core, or new universal adaptive core control function that provides seamless service control and resilient continuity for all devices and associated flows..

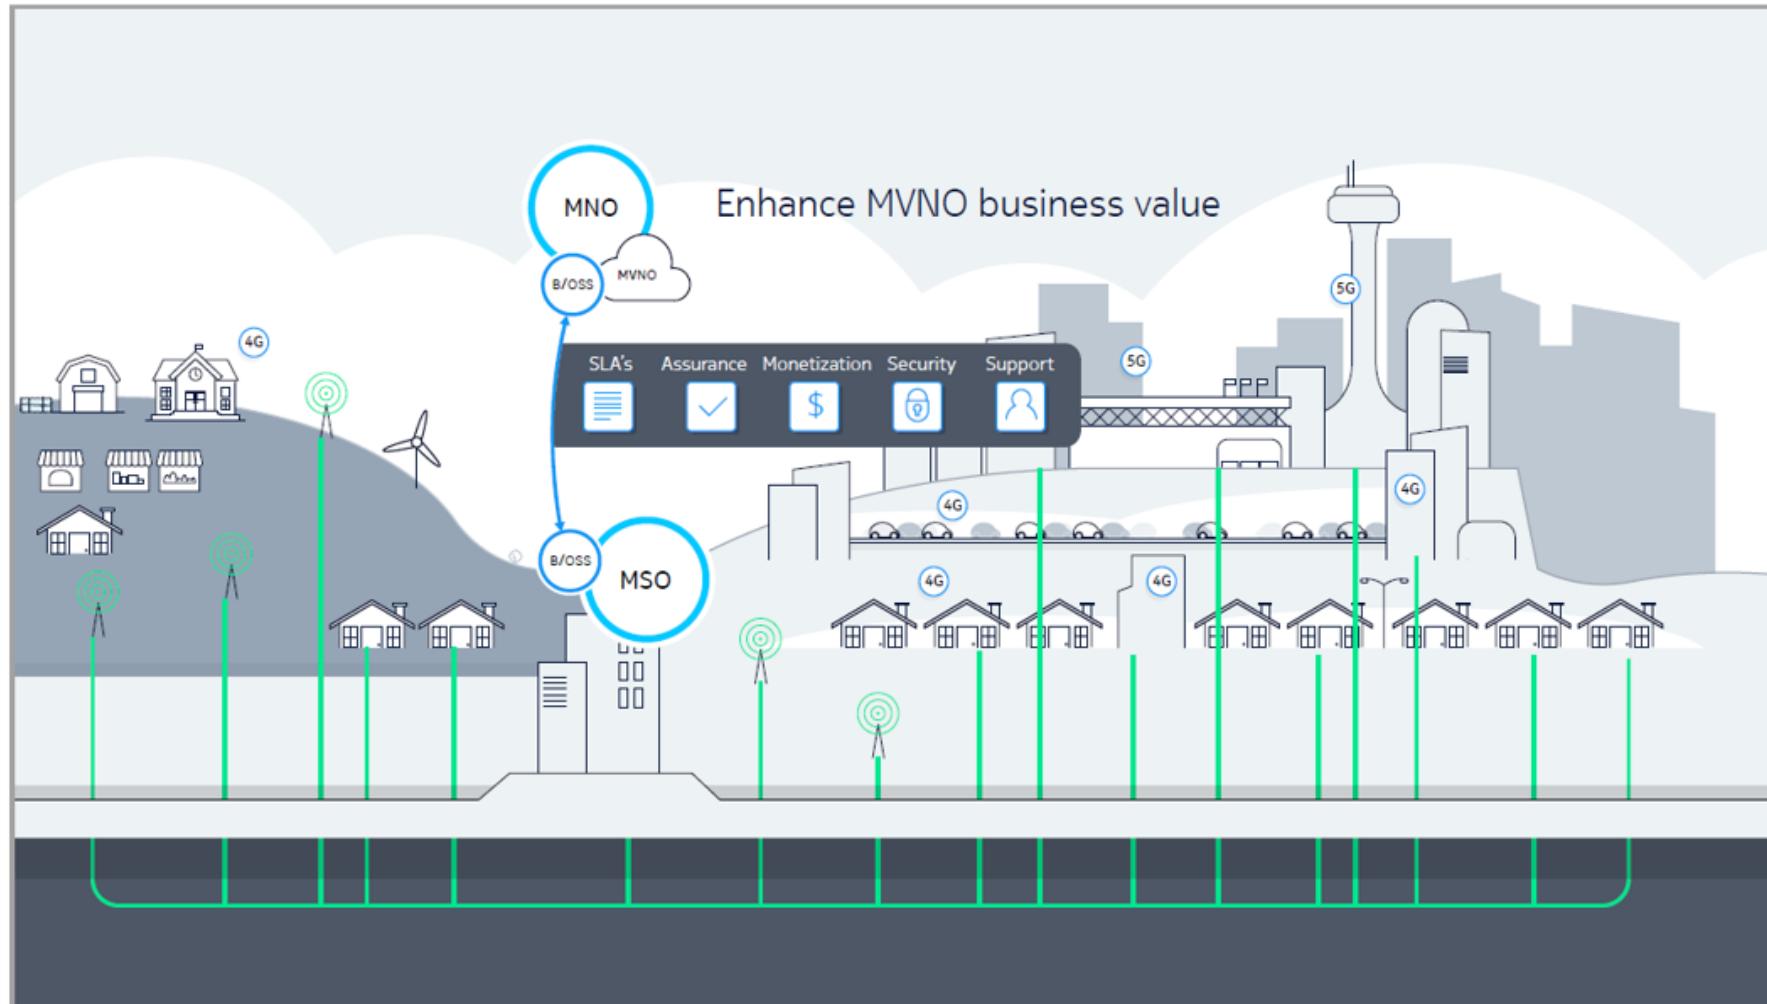
Why Converging networks requires a new approach

- Convergence of network and user data
 - Achieving the right convergence of fixed and mobile networks will mean rethinking the way you currently collect, aggregate, and normalize service and experience data from different access networks
 - Leveraging a User Data Repository (UDR) eliminates data redundancy, error, complexity and operational inefficiencies
 - The shared data layer (SDL) will enable optimal fixed wireless convergence of network and subscriber data



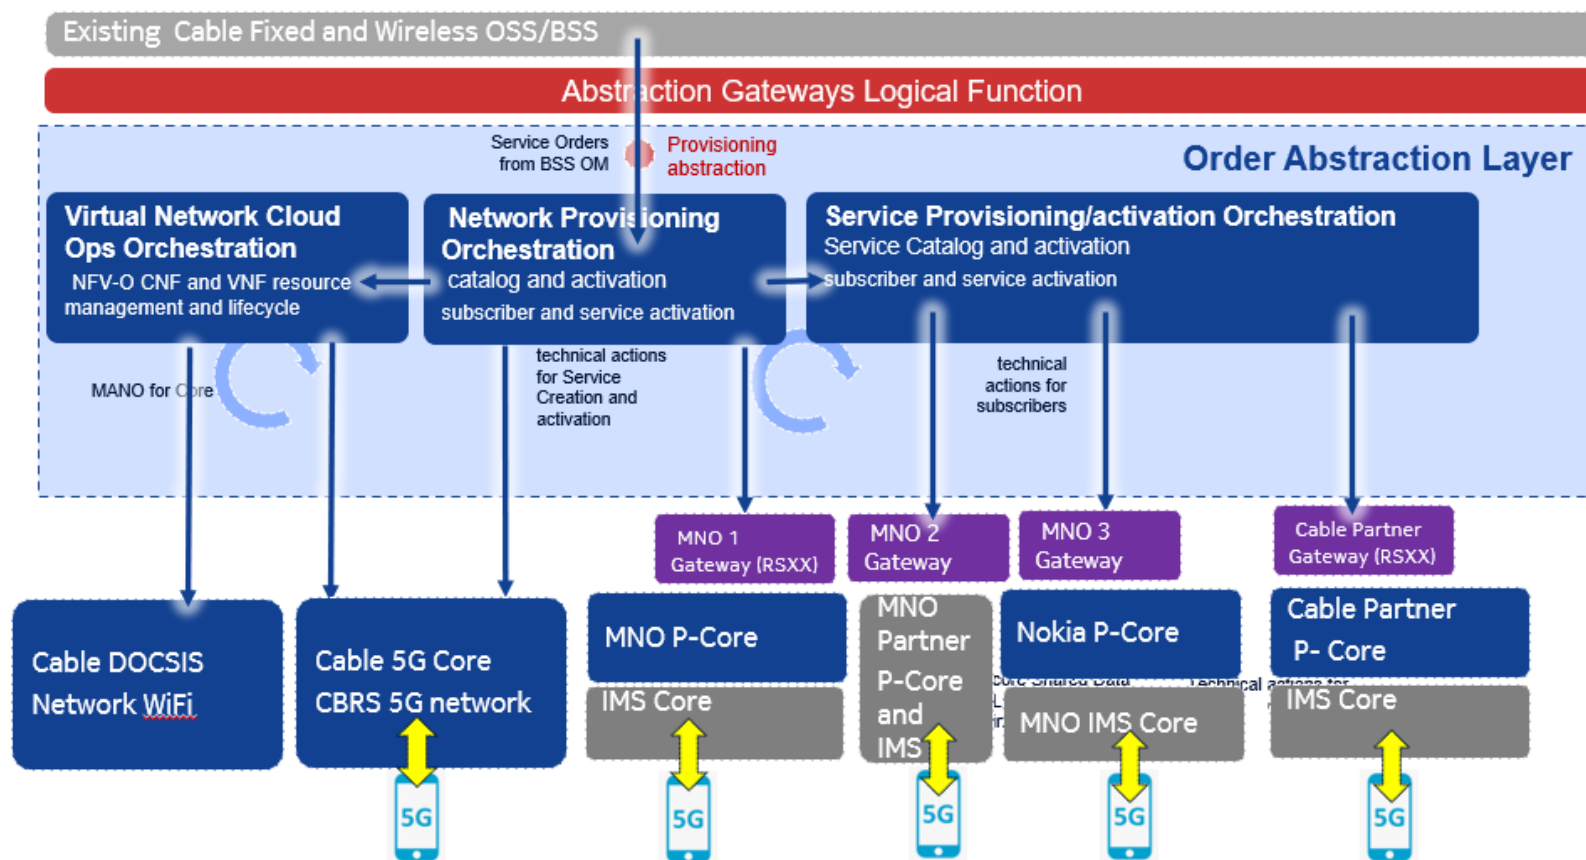
Why Converging networks requires a new approach

- Getting a converged view of subscribers on multiple networks



Why converging networks requires a new approach

- Simplifying Order management for multiple network service access



- **Not all convergence is the same** : Cable operators with multiple Fixed and Wireless networks will need to execute the right convergence architecture to offer, manage and assure network agnostic services
- **True Network Agnostic Services** : Customers will expect to access services and applications regardless of the countless dynamic handoffs which will occur between Wi-Fi, MVNO 4/5G, and CBRS 5G core and RAN owned or partnered.
- **Isolate Current Complexities** : An affective strategy is isolating B/OSS and normalizing network and subscriber data of current fixed and networks from services, devices to enable universal customer experience and situational awareness.
- **Its Convergence at the Network, Data and IT that counts** : Converging at the right points across multiple Cores enables an aggregated umbrella management architecture to provision, activate, assure, control, and measure customer experience via analytics insights, while enforcing entitlements, managing devices, service profiles and policies of services for customers across all fixed and mobile network domains.



Creating Infinite
Possibilities.

Thank You!

Bob Hallahan

Global Head of Cable Strategy

Nokia

Bob.hallahan@nokia.com

Wi-Fi Sensing

Detecting Motion for Security, Aging in Place, and More

A Technical Paper prepared for SCTE by

Josh Redmore

Principal Architect, Wireless Access Technologies | CWNE #376
CableLabs
858 Coal Creek Cir., Louisville, CO 80027
j.redmore@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	3
2. Wi-Fi Sensing Technology Overview	3
2.1. IEEE 802.11bf	4
3. Wi-Fi Sensing Use Cases	5
3.1. Home Security.....	5
3.2. Aging in Place	5
3.3. Advanced Medical Applications	5
3.4. Future Innovations	6
3.4.1. Gesture Sensing	6
3.4.2. XR Integration	6
4. Conclusion.....	6
Abbreviations	6
Bibliography & References.....	7
Acknowledgements	7

List of Figures

Title	Page Number
Figure 1 – RF Shielded, static environment, demonstrating constant CSI	3
Figure 2 – RF absorption of a human walking through the environment causes changes in CSI	4

1. Introduction

Wi-Fi sensing uses functionality already present in Wi-Fi radios to detect motion. The immediate use cases for home security are certainly valuable by themselves, but the future of sensing holds immense promise for elder care, aging in place, advanced medical applications, IoT control, immersive gaming, and more. Imagine being alerted to a change in behavior of a remote relative, or even potentially detecting a fall or other life-threatening condition. Imagine a doctor being able to have real-time information about a patient's breathing or heart rate - all with the patient comfortably at home with no expensive, dedicated medical hardware. Imagine these solutions not requiring a professional installation, relying solely on the devices already deployed in a customer's home. This paper covers the base technology behind Wi-Fi sensing, current applications, deployment best practices, and near-future use cases.

2. Wi-Fi Sensing Technology Overview

Wi-Fi sensing is the ability to detect motion inside an area covered by Wi-Fi. Current implementations use Channel State Information (CSI) to detect this motion.

“CSI characterizes how wireless signals propagate from the transmitter to the receiver at certain carrier frequencies. CSI amplitude and phase are impacted by multi-path effects including amplitude, attenuation, and phase shift. Each CSI entry represents the Channel Frequency Response (CFR)”

$$H(f; t) = \sum_n^N a_n(t) e^{-j2\pi f \tau_n(t)},$$

where $a_i(t)$ is the amplitude attenuation factor, $\tau_i(t)$ is the propagation delay, and f is the carrier frequency” [1].

In short, CSI is the sum of all information needed by an access point (AP) or station (STA) in order to decide how best to transmit a wireless frame. In a completely static environment (e.g., an AP and a STA in a shielded chamber – see Figure 1), CSI should have extremely low variability, as all direct and reflective RF paths are constant and there are no external interferers. In this simplified example, CSI is on a scale from 1 to 10, with 10 being perfect signal conditions.




Figure 1 – RF Shielded, static environment, demonstrating constant CSI

If a person were to walk through this chamber (Figure 2), these RF paths would encounter attenuation, reflection, etc., thus affecting the CSI. By observing the changes in CSI on the receiver, one can infer motion.




Figure 2 – RF absorption of a human walking through the environment causes changes in CSI

Wi-Fi sensing relies on significant Wi-Fi utilization over multiple devices. It is not currently good enough to have a single AP and a mobile device to enact sensing. An “ideal” residential deployment would have several static APs and STAs (e.g., smart speakers / Wi-Fi doorbells / Wi-Fi thermostats) deployed at the boundary of the coverage area. Mobile devices within this area currently cannot participate in sensing, other than to provide a user interface for configuration or motion reporting.

2.1. IEEE 802.11bf

Current Wi-Fi Sensing deployments, while based on 802.11 standards, are all unique in their implementations as the process for turning CSI information into inferred motion is proprietary. The IEEE has formed a task group to standardize sensing. This will be a critical step in helping wide adoption and has the potential to include moving STAs into the sensing cluster. Additionally it will enable a cross-platform mix of devices to participate seamlessly.

“Task Group bf is expected to develop an amendment that defines modifications to the IEEE 802.11 medium access control layer (MAC) and to the Directional Multi Gigabit (DMG) and enhanced DMG (EDMG) PHYs to enhance Wireless Local Area Network (WLAN) sensing (SENS) operation in license-exempt frequency bands between 1 GHz and 7.125 GHz and above 45 GHz.

This amendment enables:

- *Stations to inform other stations of their WLAN sensing capabilities*
- *Request and setup transmissions that enable WLAN sensing measurements to be performed*
- *Exchange of WLAN sensing feedback and information” [3]*

Draft 0.1 for initial comment collection was released in April 2022, and final 802 EC approval is scheduled for July 2024.

3. Wi-Fi Sensing Use Cases

The uses cases listed below are presented in chronological order from where we are today, to where this technology could take us in the future.

3.1. Home Security

Home security is the simplest implementation of Wi-Fi sensing, as it merely detects some level of motion in the coverage area. While it will be the most basic use case described here, it is not without significant challenges – primarily in false signal detection. These include:

- Fans, robotic vacuums, or other household devices that move on their own
- Pets
- Outdoor motion (e.g., a tree on a windy day)

Addressing these false detections is achievable. Fans and devices with consistent repetitive motion can be identified and filtered out. Most pets will have a much lower effect on CSI than a human, so a minimum motion level could be set. Outdoor motion can be an issue but could also be useful in the proper context. For example, outdoor motion on the second floor of a house may be ignored, while outdoor motion on the ground floor could trigger an alert as a possible trespasser.

3.2. Aging in Place

The next evolution of Wi-Fi sensing will address aging in place, which is the ability for people to retain independence and stay in their homes for longer than is currently possible. There are currently 5.8 million people over age 65 in the United States with Alzheimer's or related dementias [2]. These diseases are progressive in nature, with increasing levels of care needed the further along the person is. Technology designed for aging in place can help delay the transition to memory care facilities by offering minimally invasive remote monitoring by a relative or healthcare professional. Several specific uses for Wi-Fi sensing include:

- Fall detection
- Sleep monitoring
- Lack of motion alerts
- Boundary crossing alerts

3.3. Advanced Medical Applications

Wi-Fi sensing is also being developed for dedicated medical applications, beyond the residential deployments covered above. Heart rate and breathing detection have already been shown as achievable, but these have deployment considerations that initially make them only appropriate for controlled environments. For example, envision a hospital bed with an AP on the ceiling above the patient and a STA directly beneath them on the floor. These APs would not be used to service other non-medical STAs and would exist as a separate network from anything providing data services.

While this level of medical application is only possible today in a specific healthcare environment, Wi-Fi sensing technology should advance to the point where these same services are achievable in the home.

3.4. Future Innovations

We've only begun to glimpse the full potential of Wi-Fi sensing, as we're only in the very first years of innovation. And considering there are over 20 Billion Wi-Fi devices deployed today, there's an incredible base upon which to develop new uses for Wi-Fi sensing. The following represents a small sample of what is technically possible.

3.4.1. Gesture Sensing

As seen in section 2.1 above, the 802.11bf task group is working to enable sensing at frequencies above 45 GHz. In this range, gesture recognition becomes possible at very small motion ranges. We've seen advances in IoT control where most home devices can be controlled with a voice command, and sensing will enable an even more natural level of control. If you want all the lights off in a room, just gesture towards the light switch as if you were turning it off.

3.4.2. XR Integration

The natural evolution from gesture recognition is full XR (AR/VR/MR) integration, without the need for on-body sensors or dedicated tracking systems. This will increase immersivity while decreasing hardware costs and battery utilization, making XR systems lighter and more powerful. Potentially, the only hardware needed would be the head-mounted display (HMD).

4. Conclusion

Wi-Fi sensing represents a paradigm shift in how we leverage existing and future Wi-Fi deployments. Home security applications available today are just the very beginning. The biggest impacts will be in home healthcare, granting more independence and a better quality of life for far longer than ever possible before. On top of home and health, the potential value to entertainment markets is enormous, and the field for not-yet imagined sensing innovations is limitless.

Abbreviations

AP	access point
AR	augmented reality
CSI	channel state information
EC	engineering change
GHz	gigahertz
HMD	head-mounted display
IEEE	Institute of Electrical and Electronics Engineers
RF	radio frequency
STA	station
VR	virtual reality
WBA	Wireless Broadband Alliance
XR	extended reality

Bibliography & References

1. Yongsan Ma, Gang Zhou, and Shuangquan Wang. 2019. WiFi Sensing with Channel State Information: A Survey. ACM Comput. Surv. 52, 3, Article 46 (June 2019), 36 pages.
<https://doi.org/10.1145/3310194>
2. <https://www.cdc.gov/aging/publications/features/Alz-Greater-Risk.html>
3. https://www.ieee802.org/11/Reports/tgbf_update.htm
4. WBA Wi-Fi Sensing (<https://wballiance.com/resource/wi-fi-sensing/>)
5. WBA Wi-Fi Sensing – Test Methodology and Performance Metrics
(<https://wballiance.com/resource/wi-fi-sensing-test-methodology-and-performance-metrics/>)
6. WBA Wi-Fi Sensing Deployment Guidelines (<https://wballiance.com/resource/wi-fi-sensing-deployment-guidelines/>)

Acknowledgements

This document could not have been written without the hard work and contributions of the Wireless Broadband Alliance's Wi-Fi Sensing Workgroup and Cognitive Systems' staff, especially Chris Beg.

With Great Power Comes Great Electricity Bills

Reducing grid dependance of the Access Network as it evolves toward 10G and beyond

A Technical Paper prepared for SCTE by

Tobias Peck

Sr. Director of Broadband Product Management

EnerSys

3767 Alpha Way, Bellingham, WA 98226

(360) 392-2247

tobias.peck@enersys.com

Table of Contents

Title	Page Number
1. Introduction.....	4
2. Powering Challenges into the Future.....	4
2.1. Powering 10G Vision.....	5
3. Power Mapping the HFC Grid.....	7
3.1. Where does the power go?.....	8
3.2. Customer Connectivity.....	8
3.3. Conversion Loss.....	9
3.3.1. Converting utility voltage to plant voltage.....	9
3.3.2. AC to DC battery charging or DC to AC battery backup.....	12
3.3.3. Conversion of plant power to DC power for actives.....	12
3.4. Transmission Loss.....	12
3.4.1. Modelling power transmission loss.....	13
3.4.2. Concepts for voltage increase in the HFC.....	15
3.5. Battery Charging and Management Overhead.....	15
3.5.1. Power use from battery charging.....	15
3.5.2. Power use from DOCSIS status monitoring.....	16
4. Modelling the impact of loss factors in the HFC grid.....	16
4.1. Baseline assumptions for modelling.....	18
4.2. Conversion Loss impact scenarios.....	19
4.2.1. Power supply conversion loss.....	19
4.2.2. Conversion loss in active devices.....	19
4.3. Transmission loss impact scenarios.....	20
4.4. Battery Charging and monitoring overhead.....	21
4.5. Cumulative effect of power savings strategies.....	21
5. Other power reduction strategies.....	22
5.1. Solar power augmentation.....	22
5.1.1. – Options for solar implementation in the HFC plant.....	23
5.2. Time-of-use mitigation.....	24
6. Conclusion.....	25
Abbreviations.....	25
Bibliography & References.....	26

List of Figures

Title	Page Number
Figure 1 - Cable Operator Power Consumption Pyramid.....	5
Figure 2 - Diagram of the Access Network in the near future.....	6
Figure 3 - Areas of power consumption in the HFC Grid.....	8
Figure 4 - Simple circuit diagram of a ferroresonant transformer.....	10
Figure 5 - Noise and transient reduction capability of a ferroresonant transformer.....	10
Figure 6 - Current limiting of a ferroresonant transformer during a short-circuit condition.....	11
Figure 7 - Sample network for power loss calculation.....	14
Figure 8 - NESC section 2, excerpt defining allowable voltage in the communication space.....	15
Figure 9 - Pictures of the model plant and the ice cream truck where I lost my assistants.....	17
Figure 10 - Network power diagram for scenario modelling.....	17

Figure 11 - Example of a solar augmented cable power supply 23
 Figure 12 - Basic diagram of a cable power supply augmented by solar 24
 Figure 13 - Example utility time-of-use rate structure 24

List of Tables

Title	Page Number
Table 1 - Quick reference guide for coaxial cable resistance	13
Table 2 - Power draw for modeled plant actives	18
Table 3 - Baseline power breakdown of modeled network	18
Table 4 - Model of power consumption varying power supply efficiency.....	19
Table 5 - Model of improved power conversion efficiency of actives	20
Table 6 - Model of the impact of increased voltage on transmission loss	20
Table 7 - Cumulative effect of power savings strategies	22

1. Introduction

The coaxial network that is the backbone of the Cable Broadband industry was designed with a simple, ingenious powering scheme that allows multiple “actives” to be powered from a single point of connection to the utility grid. This tried-and-true powering strategy allows operators to minimize their points of interaction with an unpredictable utility grid and create a more reliable HFC grid to power network actives that support customer connectivity.

However, unlike the utility grid, the coaxial network’s primary functions are both optimizing the flow of data to customers and transporting energy. As we continue to move toward 10G and beyond, exponential progress has been made in speed and volume of data that can be delivered over HFC, often by developing new architectures that increase dataflow with only minimal increases to power. But, as our industry seeks to be more sustainable and reduce operational energy spending, what should we be doing to improve power usage efficiency? Are there network efficiency gains that can be applied to all architectures to reduce energy usage and further improve the amount of data per dollar that the network can deliver?

This paper will first give a brief overview of network powering as we move toward 10G and beyond to understand the urgency of developing more efficient network powering. From there, the goal will be to map the power consumed by the network from the point of connection to the utility, through every transition point where the energy is transformed, transported, split, or ultimately consumed, to identify where energy is being used without creating value for customers. Finally, with our map in hand, we will discuss where efforts can most effectively be applied to reduce energy consumption and associated cost from the Access Network.

2. Powering Challenges into the Future

One operator recently commented that their outside plant (OSP) coax installed in the 1980s will meet DOCSIS 4.0 performance requirements with only passive component upgrades (splitters and taps). Similarly, OSP powering elements such as enclosures, power inserters and ferro-resonant transformers can function for decades. And while OSP power electronics and batteries will require expected periodic replacement resulting from normal use, the components that currently comprise coaxial network can support network demand into the foreseeable future. The question that needs to be addressed is how the OSP network can deliver power to the network more efficiently and sustainably. To fully understand the impact of the OSP on overall industry power usage, the SCTE Energy 2020 subcommittee recently updated the “Power Pyramid” showing that OSP power usage makes up nearly half of the overall consumption in the Cable Broadband industry.




Figure 1 - Cable Operator Power Consumption Pyramidⁱ

The urgent need to increase power efficiency is inherently connected to not only an understanding of the past and present of the HFC network, but also the future. The trek to 10G requires us to understand what the network of tomorrow will look like in order to plan for optimum power efficiency. Access network evolution has created unique powering challenges that must be addressed to assure that network powering upgrades will endure for decades to come. Let's review some of our powering challenges from the lens of the evolving access network with some help from CableLabs.

2.1. Powering 10G Vision

The CableLabs 10G initiative is the catalyst behind several technology innovations designed to deliver future proof internet speeds up to 100 times faster than most consumers are experiencing today. 10G aims to provide 10Gbps symmetrical, secure, low latency data services. 10G innovations will affect every aspect of the broadband network including headends, the access network and the customer's premises. The access network specifically must undergo enhancements to support new performance levels. Underlying technologies used to move network performance towards and beyond the 10G vision requires power. Assuring the availability of additional, reliable, and intelligent power for the 10G capable network is both essential and challenging since network architectures are evolving and much of the 10G enabling technology is still being developed. To approach this dilemma, let's review the 10G architecture vision as it exists today to ensure we're planning the appropriate powering infrastructure to support this near-future vision. For our 10G powering discussion, refer to this CableLabs 10G network architecture shown in Figure 2.

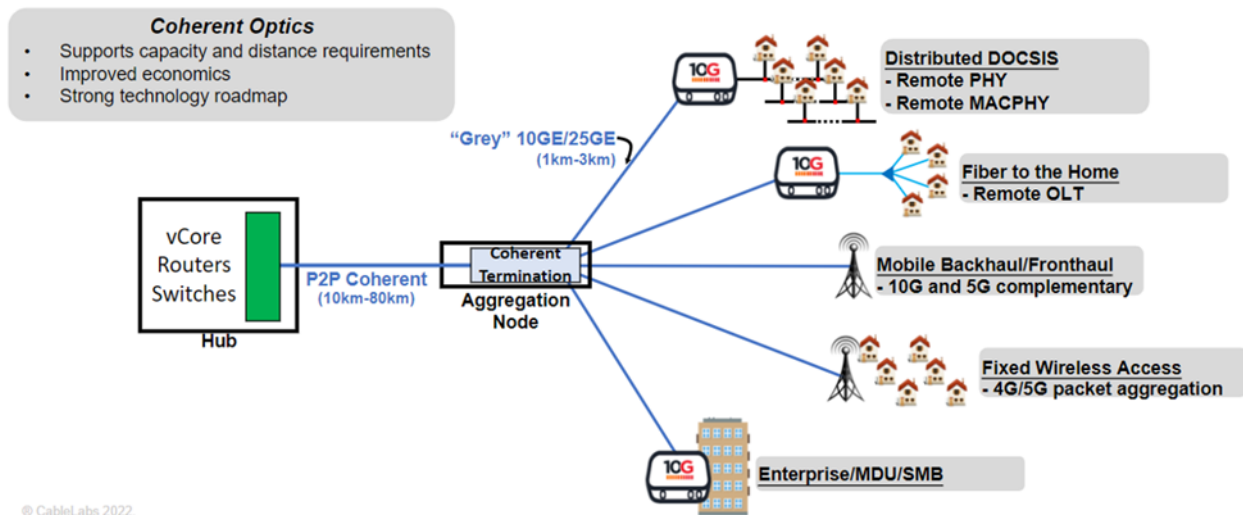


Figure 2 - Diagram of the Access Network in the near future

This near-future network contains several new concepts all combined into a single diagram. Let's review these network elements through the lens of power.

Aggregation Node: Fiber leaving the headend is routed to a new network component, the aggregation node (AN). In this architecture, new high speed fiber communications are pushed deeper into the network by use of high-speed coherent optics between the hub and the new AN. This link could be 50Gb, 100Gb or faster. The AN will act as an optical Ethernet switch providing 10Gb (or higher speed) optical links to downline elements such as DAA nodes and OLTs. The result is higher capacity services deeper into the network and closer to subscribers. From a powering vantage point, we've simply added another remote active element to the mix. Although the AN concept is new, one likely implementation would be a clamshell enclosure with fiber ports for upstream and downstream links, with a traditional node power port connecting to a standard 90VAC HFC power source. Once AN units begin to be deployed, power requirements will need to be analyzed. Given today's state of coherent pluggable modules combined with known high-speed switching elements, we anticipate that the AN may require approximately 200W in a fully configured state. This estimate remains highly speculative until actual architectures using AN prototypes are deployed.

DAA Nodes: Today's DAA nodes are fed from 10G Ethernet backhaul. Future DAA nodes may utilize 50G, 100G or faster backhaul links from an AN. These higher capacity links would be implemented using pluggable coherent optic modules. DAA node outputs support four DOCSIS 4.0 RF QAM channels over coax. Future looking, fully loaded DAA nodes may approach 180W power requirements.

R-OLTs: PON networks are deploying at increasing rates. The RDOF initiative is contributing this increase. Today's strand-mount R-OLTs support up to 512 subscribers using either 10G-EPON or XGS-PON protocol. Higher speed PON architectures are being discussed. 50Gb coherent optics support is planned for next generation PON implementations. Current R-OLT units can consume up to 140W. Future 50Gb PON R-OLTs will likely require up to 165W.

Wireless Backhaul: Wireless backhaul for 5G using the existing HFC infrastructure is expected to grow significantly in the near term as 5G scaled rollouts continue to increase. Coax powered radio units (RU), access points (AP) represent a significant incremental power draw to the coax plant.

Enterprise/MDU: high-capacity enterprise applications as well as large MDUs have historically used P2P dedicated optics for backhaul. DOCSIS 4.0 may be used for many new commercial applications when coax already exists in the region. New ANs providing P2P services or utilizing 10G or 50G PON to commercial and MDU customers are anticipated in the future.

Having this vision in mind of the near-future HFC network and its ability to support the data and backhaul demands for 10G and beyond, let's now discuss how to make the future more energy efficient.

3. Power Mapping the HFC Grid.

Reducing energy use and OpEx spend from the access network starts by understanding where power is used in the plant. More specifically, identifying and reducing or eliminating energy waste is the key to optimizing the effectiveness of the network. By mapping the flow of power from the point of connection to the grid and cataloging how each Watt of energy is consumed or wasted, we can develop data-driven strategies to optimize network sustainability and significantly reduce utility spend.

As we look to understand how to save energy, it is important to keep in mind these simple conversions based on the US average utility cost and CO₂ impact per Watt:

1 Watt of power savings = 8.76 kWh annually

8.76 kWh @ \$0.119 per kWhⁱⁱ = \$1.04 of savings = 7.45 Lbs of CO₂ⁱⁱⁱ

So, every Watt of savings yields a dollar of annual savings on energy bills, with the potential for 2 to 3 times that in regions with higher utility rates. This combined with the fact that there are more than 750,000 power supplies in North America alone, help to emphasize the value that can be gained by reducing wasted power in the OSP.

3.1. Where does the power go?

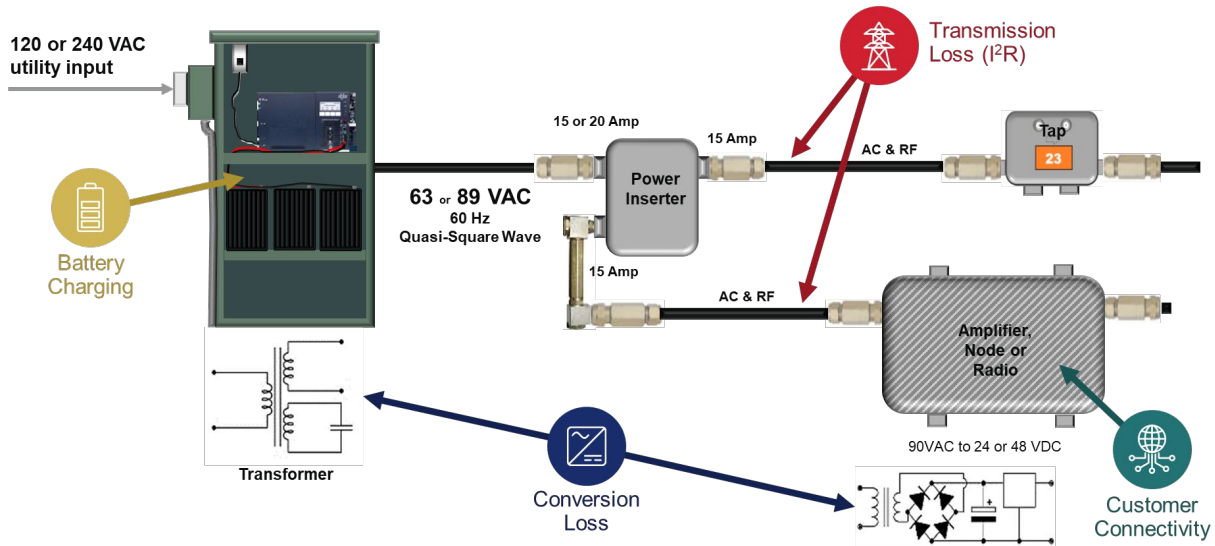


Figure 3 - Areas of power consumption in the HFC Grid

From the point of connection to the utility, power in the HFC network is essentially consumed in 4 ways:

1. **Customer Connectivity** – Powering equipment necessary to transport data through the network and connect customers to the digital world requires energy
2. **Conversion Loss** – Converting electrical power to a different voltage or from AC to DC and back
3. **Transmission Loss** – Power loss inherent from moving power across the network (I^2R loss)
4. **Battery charging and management overhead** – power used for operation of the network power supplies including the cable modem (transponder) and battery charger

Understanding the reasons behind each of these areas of energy consumption and the locations where they happen will allow us to develop strategies to greatly reduce them. Because the plant is a complex electrical circuit, changes to any of the four elements of power consumption above can often have an additive effect. For example, conversion loss in actives can have an impact on transmission loss throughout the plant. We will discuss this in more detail later.

It is also necessary to consider power usage in the network holistically, as in some cases small power sacrifices in the Access Network are made to make the network more robust or intelligent, and often are offset by significant power savings in other areas of plant operations. In order to better quantify the impact of each of the four types of power consumption for comparative analysis against operational impacts we will build a model network and observe the impact of changes to each of them. But first, let's dive a little deeper into each of these four types of power consumption in the network.

3.2. Customer Connectivity

The HFC network is the central nervous system of the connected world. The actives that extend the reach of the network or increase the speed and bandwidth of data that flows through it are the reason that the HFC grid exists. Significant effort is consistently being invested into network actives to enable them to do more with less power. While these efforts have significant impacts in reducing network power consumption, for the purposes of this paper we will view their power draw as a mathematical given within

our power calculations, understanding they are a non-negotiable within the network. We will discuss potential savings from maximized conversion efficiency within actives as a piece of Conversion Loss.

3.3. Conversion Loss

When converting electrical power from one form to another there will always be some power loss. The method by which that power is transformed can have a significant impact on power consumed in the HFC network. However, any conversation around transformer efficiency needs to be had with a more wholistic view considering effects on plant resiliency and operating costs. Currently in the HFC network there are 3 main points of conversion where loss needs to be considered.

3.3.1. Converting utility voltage to plant voltage

The first point of conversion is at the utility input to the power supply that feeds power to the plant. At this point utility power at either 120VAC or 240VAC is converted to plant voltage of 60VAC or 90VAC. In the overwhelming majority of powered coaxial plant in the world this conversion is done by a power supply with a **ferroresonant transformer**.

3.3.1.1. A Brief description of ferroresonant transformers

“Ferroresonance” is a phenomenon associated with the behavior of iron transformer cores, which in the instance of broadband power supplies, have two separate sections. The input core section is designed to prevent saturation at the maximum utility input so it cannot trip the input circuit breaker. The output section is designed to saturate in a controlled manner to provide tight output voltage regulation. The core is designed to optimize the efficiency at full load and allow the saturation loss only at lower load for output voltage regulation.

Normally, a ferroresonant transformer causes distortion of the output sine wave shape at lower load to regulate the output voltage. Ferroresonant transformers also have a magnetic shunt between the input winding and the auxiliary secondary winding, which is paralleled with one or more capacitors, forming a resonant circuit tuned to the power supply frequency. This resonant LC or “tank” circuit amplifies the input voltage and drives the output core to saturation to regulate the output voltage. In addition to providing isolation from input to output windings, this tank circuit provides additional filtering of input noise and resilience against high utility energy surges, such as lightning and industrial surges. The tank circuit also limits the output current during plant short circuit conditions to prevent tripping the input circuit breaker.

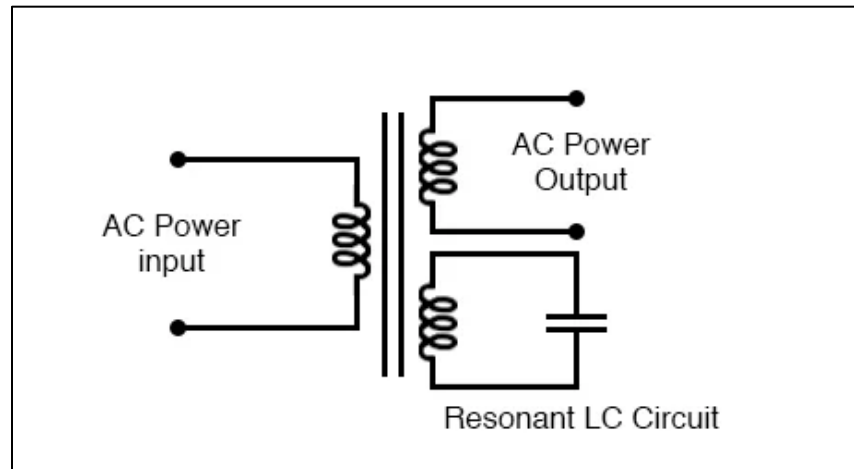


Figure 4 - Simple circuit diagram of a ferroresonant transformer

Because of their unique design, ferroresonant transformers provide several key benefits in powering the OSP network. First, they provide network equipment protection for surges and other transients as well as filtering out input noise that might otherwise be passed on to the plant. Additionally, they provide significant protection from shorts on the plant by current limiting to protect other equipment, and no-touch recovery once the short condition has been repaired. Finally, their relatively simple design, with minimal component count makes them extremely reliable and resilient to the uncontrolled environmental conditions in the majority of OSP locations.

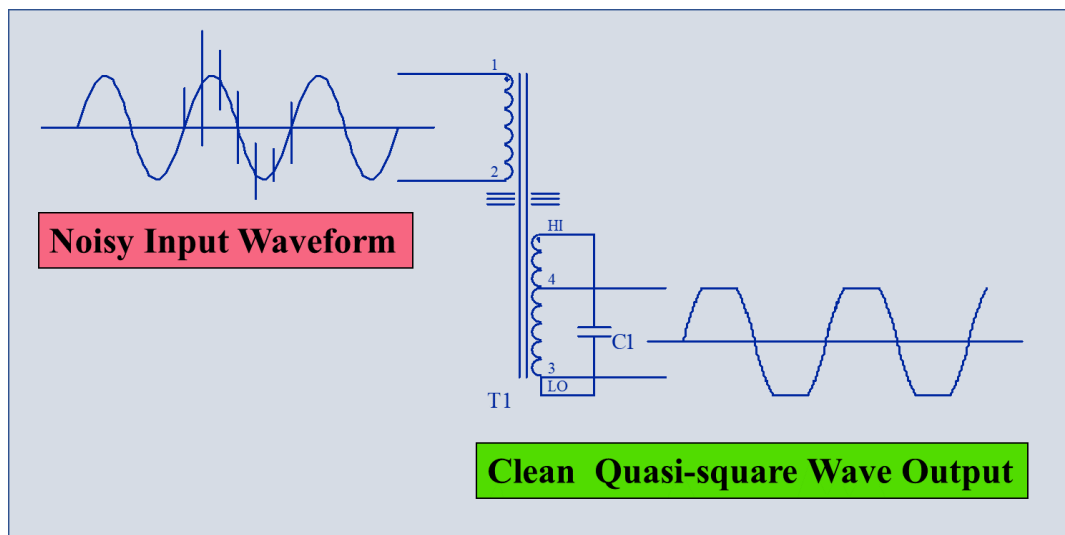


Figure 5 - Noise and transient reduction capability of a ferroresonant transformer




Figure 6 - Current limiting of a ferroresonant transformer during a short-circuit condition

However, with the additional power required for both the tank circuit, and the iron core saturation, the benefits mentioned above come at the cost of some efficiency loss. Since these elements of a ferroresonant transformer are optimized to handle a peak load, there is an inherent overhead to running these transformers that generally makes them significantly less efficient at lower loads. For example, in a power supply that is rated for up to 18A of output, there may be an inherent load of 100W or more to ensure resiliency up to peak load. At 3A of output at 90VAC or 270W of load, an additional 100W of inherent load means the transformer is operating at just under 73% efficiency.

Fortunately, much has been done over the past 30 years to gradually decrease the inherent load required for a ferroresonant transformer and make these devices more efficient. For example, a standard 15A power supply outputting 700W before 1999 would have had an efficiency of around 82.5 %. In the 2000s the standard 15A power supply was up to around 86% efficient at that load and today it is an additional 2% higher. Additionally, as the peak rating of a ferroresonant transformer is lowered so is the inherent load required to support that peak rating. Therefore a transformer that was rated for 5A, for example, would have an inherent load significantly lower than the 18A transformer.

3.3.1.1. Load Matching

It is important to note that when loaded to 50% or more of max rating, ferroresonant transformers can reach efficiencies of 85% to 94%. Knowing this, one strategy that could be deployed to reduce energy usage without losing the operational benefits of a ferroresonant transformer-based power supply is **Load Matching**. This strategy uses power supplies designed to tiered peak capacities in order to lower the inherent load on power supplies at lower amp draws. For example, if power supplies designed to peak loads of 5A, 10A, and 18A are deployed in an operator's network, power supplies can be targeted to unique plant loads between 2.5A and 18A so that all power supplies operate at efficiencies above 85%, with many above 90%. While this strategy comes with some added management operationally, it could be an effective strategy to deploy during regular replacement cycles that shows moderate improvement for little added cost. We will model the potential impact of this strategy later on.

3.3.1.2. Linear and Switch-mode transformer efficiencies

While other transformer architectures, such as linear or switch mode, may be able to achieve conversion efficiencies of 95% or greater, these efficiency gains need to be weighed against either lost operational

benefits, or the energy cost to add intelligence or redundancy into their architecture to achieve desired resiliency targets. For the purposes of this paper, we will model potential energy efficiency gains from increased efficiency of moving to these architectures, but thorough analysis should be done on real-world impacts of switching to a non-ferroresonant transformer.

3.3.2. AC to DC battery charging or DC to AC battery backup

The next conversion to be discussed briefly is the conversion in the standby power supply from AC to DC for battery charging, or DC to AC to power the plant during utility outages. These conversions have very little impact on overall power consumption since they happen very infrequently. And, while there is a constant float charge generally being applied to lead acid batteries in a standby power system to keep them healthy and ready for discharge, even in the worst-case scenario a 20% change in the efficiency of this conversion would yield less than a Watt of savings.

3.3.3. Conversion of plant power to DC power for actives

Another location in the plant where conversion loss is notable is at each active. Plant actives are generally powered internally by 48VDC or 24VDC which means that conversion from 90VAC plant power to the required DC voltage is necessary. Most actives do this conversion through small voltage conversion devices which, generally have low losses with efficiencies generally ranging from 95% to 98%. The presence of the ferroresonant transformer providing power to the network allows these higher efficiency transformers with less resiliency against transients and surges to be safely used in network actives. The efficiency of these, and all transformers, does tend to decrease with age as transformer insulation wears down and more energy is lost to heat. Reducing network impact by ensuring maximized transformer efficiency of actives is key to reducing wasted energy in the network, as inefficiency of actives can also create additional transmission loss in the plant.

3.4. Transmission Loss

Transmission loss refers to the power lost from pushing current down the coax to feed actives, stated mathematically by Ohm's law, and it is one of the most important factors to understand in order to reduce the power consumption of the OSP. Depending on the particular section of plant, transmission loss can account for up to a 25% of the power consumed in the plant.

Since the plant is, from a powering standpoint, an electrical circuit with a series of loads drawing current and a series of resistors between them, in a single span of coax feeding an active we can express the transmission loss by Ohm's law as:

$$P(\text{loss}) = I^2R$$

Where:

$P(\text{loss})$ = power lost from coax line resistance, measured in Watts.

I = current through the cable required to feed an active, measured in amps

R = resistance of the length of cable, measured in Ohms

To understand the total plant load and how and where transmission loss occurs, it is also important to remember that active loads on the plant are constant power devices, or again from Ohm's Law

$$P = IV$$

Where:

P = power required for the active to function

I = current through the cable required to feed that active, measured in amps

V = voltage feeding the active

While these concepts are fairly simple mathematically, the interaction between them and a number of properties of the plant make them more complex to model in the plant.

3.4.1. Modelling power transmission loss

The layout and makeup of the HFC plant varies greatly from section to section. And while the number and type of actives in the plant are the most significant factors in the power draw of the plant, the coax and the passives have an impact on transmission loss.

3.4.1.1. Coax Loss

Everywhere that the plant is not drawing power it is resisting its flow. The resistance of any span between two actives is generally determined by three factors which, again, vary greatly from plant to plant – the length of the coax, the diameter of the center conductor of the coax and the passives in the power path. The coaxial resistance is fairly predictable as long as the cable size is known, which is not always a given especially when reaching toward last-leg actives. For modelling purposes, assuming that we know the length and diameter of coax, each span from active to active can be treated as a resistor with a certain ohmic value calculated from the table below.

Table 1 - Quick reference guide for coaxial cable resistance

P3 Cable Resistance	
Cable Dia	Ohm/Ft
0.5"	0.00172
0.625"	0.0011
0.75"	0.00076
0.875"	0.00055
625 PF	0.0003

3.4.1.2. Passive impact to transmission loss

Passives such as taps, splitters, directional couplers and power inserters are less impactful to plant resistance, but also less understood. Their per length resistance is inherently greater than that of the coax that surrounds them in the plant, but the value of this addition resistance is not well-documented and therefor hard to quantify. A cursory check of resistances for several random passives recently showed resistances from input to output ports of between 10 and 35 milliohms, which based on the table above is the equivalent of an additional 10 to 50 ft of coax depending on diameter. Assuming that approximately one passive is in each 200 ft. span of coax, this means that passives could add up to somewhere between 5% and 10% to the overall resistance of the plant.

To be clear, these results do not come from a well-organized, systematic study of a large sample of various types and manufacturers of passive equipment. They are an initial glimpse at a potential opportunity for plant savings that should be further understood. With the need to update passives to handle higher frequencies required by DOCSIS 4.0, this could be an opportune time to determine a resistance specification for passives that is not cost prohibitive.

3.4.1.3. *Constant power actives*

The vast majority of active equipment in the access network today is of a constant power nature, or simply the input current and voltage can vary as long as their mathematical product provides the requisite power for the device to operate. This is not to say the power draw of actives cannot vary moment to moment due to changes in throughput required. The principles of ‘constant power devices’ dictate that any change in voltage is offset by an inversely proportional change in current. Thus, if a 90W load initially is receiving 2A of current at 45V, and the voltage on the plant is doubled, the current through the lines will be halved. For the HFC network this has the added benefit of reducing the voltage drop through the coax, thereby reducing power loss. As you can see, this process would continue to occur until the changes in voltage and current were immeasurable and the plant voltage “settled.” In the plant this settling of voltage happens almost instantaneously. For a mathematical model of the plant, it can be more time consuming to calculate without the use of some additional calculating power.




Figure 7 - Sample network for power loss calculation

Consider the simple section of plant above. To calculate the total power draw from this section of plant you would sum the power drawn from each active (A-1,2,3) and the total transmission loss from each span (S-1,2,3.) Before calculating the loss in span S-2, you need to calculate the loss in S-3 as the power loss in that span is a factor in determining the amount of current that is drawn through S2. The same goes for S-1 where you need to calculate the loss in S-2 and S-3 before you can calculate that span.

So first:

$$P_{(\text{loss in S-3})} = (P_{(A-3)} / V_{(A-3)})^2 \times R_{(S-3)}$$

Then:

$$P_{(\text{loss in S-2})} = ((P_{(A-3)} + P_{(A-2)} + P_{(S-3)}) / V_{(A-2)})^2 \times R_{(S-2)}$$

Finally:

$$P_{(\text{loss in S-1})} = ((P_{(A-3)} + P_{(A-2)} + P_{(A-1)} + P_{(S-3)} + P_{(S-2)}) / V_{(A-1)})^2 \times R_{(S-1)}$$

For purposes of power modelling the plant later in this paper a computer model was built that sets the end-of-line (EOL) voltage at the last active to a relatively low value then calculates loss back to the point of power. The model then checks the voltage against the known output of the power supply then repeats the process iteratively by raising the EOL voltage incrementally until the calculated voltage at the power supply matches the known output voltage. From there all transmission loss in the plant can be summed and analysis performed on the impact of any power saving strategy.

3.4.2. Concepts for voltage increase in the HFC

The proportional relationship between voltage and current with regard to power opens up interesting opportunities for reducing transmission loss by raising the voltage of the plant to lower the current, and therefor transmission loss, through the coax. First, there is still roughly 25% of plant in North America that is 60V. We will show later in our model that raising the voltage of plant from 60VAC to 90VAC can yield a significant reduction in transmission loss. Second, the National Electric Safety Code (NESC), Section 2 (excerpt in Figure X below) allows up to 150VDC in the communication space and while there are details that need to be better understood with DC powering of the HFC plant, there may be some relatively simple solutions that could allow the voltage to increase to that level. This increase would not only reduce power draw on the plant, but would also increase the reach of the plant and potentially allow for the reduction of operational energy usage by reducing the total number of sites to be maintained. It could also enable easier deployment of wireless technologies powered from the HFC grid. We will model all these scenarios below to better understand their impact on reducing utility power draw.

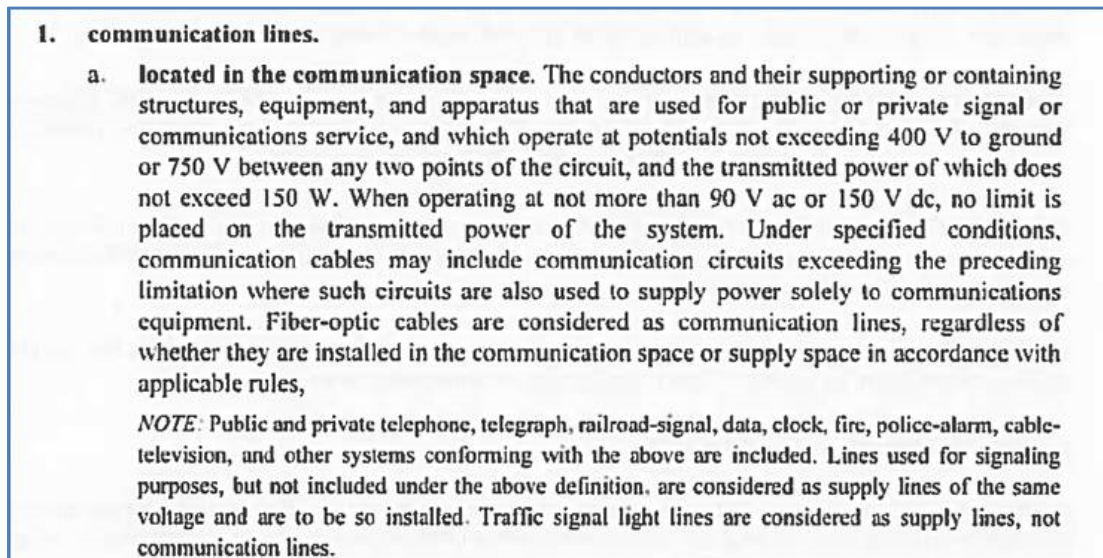


Figure 8 - NESC section 2, excerpt defining allowable voltage in the communication space

3.5. Battery Charging and Management Overhead

Standby powering has made the HFC grid an extremely resilient architecture. There are two key elements required to make standby powering such a robust architecture that draw additional power.

3.5.1. Power use from battery charging

Standby power requires batteries, battery charging can use unnecessary power if not properly managed. Typical gel cell VRLA batteries require a trickle/float charge of up to 3-4W constantly to stay healthy.

This is necessary to negate the effects of self-discharge and ensure batteries are topped off and ready for an outage. As batteries age and internal resistance increases the power used for float-charging these batteries will continue to increase. Therefore, it is important to ensure that batteries are properly maintained so that aging batteries with higher internal resistances, which are often no longer able to provide adequate back up time, are replaced regularly.

Some newer advanced AGM lead-acid batteries present a small opportunity for power savings. With lower internal resistances and self-discharge rates, these batteries have the ability to “rest” without being charged for up to 75% of their life, reducing the power needed to float charge them. Lithium-Ion batteries have extremely low self-discharge rates and, therefore, do not require float charge at all. There are small power gains of around 4 Watts to be had here, but as a part of a broader strategy it could help add up to larger total reduction of power.

3.5.2. Power use from DOCSIS status monitoring

Remote DOCSIS status monitoring of network power supplies has helped to greatly improve power reliability in the OSP by providing insight into potential issues before they become customer impacting. As our ability to extract additional data and perform more advanced analysis on these devices continues to improve, plant operations practices will become increasingly more streamlined. While these devices have become mission-critical to plant operations, they do have some draw that should be noted. Like other cable modems, the majority of devices currently in the OSP have an average draw of 5 to 7 Watts which will increase as networks advance and require higher-powered DOCSIS devices.

4. Modelling the impact of loss factors in the HFC grid

To understand how power is used in the HFC plant and more importantly where it is being consumed without adding value, we can look at an average section of plant and map the power flow. Then we can run mathematical scenarios based on targeting each type of loss and measure the impact of each strategy.

As mentioned earlier, the layout and makeup of the HFC plant varies greatly from section to section. This exercise is designed to show potential impact from loss factors, but the weighting of these factors may vary slightly depending on architectures or location of the plant—rural, suburban, urban—being viewed. The average power supply in North America has an output load of around 650W-750W so for the model it is ideal to find a section of plant with an output load within that range. Fortunately, in order to find this model section of plant all I needed to do was take two aspiring young cable engineers on a walk around my neighborhood to assist me in mapping the power flow from my local power supply. They were very engaged and helpful, until we reached the node which had an ice cream truck parked in front of it.



Figure 9 - Pictures of the model plant and the ice cream truck where I lost my assistants

For the sake of simplified mathematical modeling I used some generalized estimations of power draw for various actives and calculated the output power from my local power supply to be approximately 720W, which just happens to fit nicely into the average load range. I've mapped our model network in **figure 10**, but for the purpose of this exercise I have arranged the elements to more effectively show power flow instead of a more standard optical and RF-path layout. It is also important to note that, while loads were based on data from a cursory review of real actives of varying models and capabilities, the accuracy of the plant active power estimations used for modelling power draw is somewhat irrelevant due to the broad variation of active loads that the plant can support today and will support in the future.

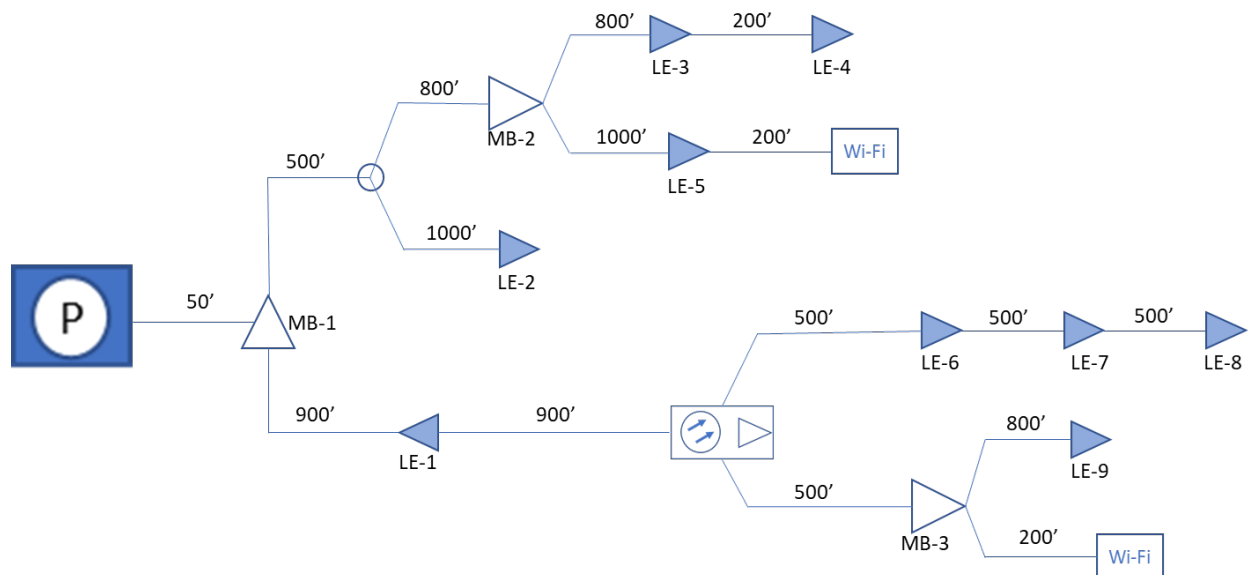


Figure 10 - Network power diagram for scenario modelling

4.1. Baseline assumptions for modelling

Below are the baseline assumptions that were used to build our initial power map.

- 1) Estimates for power draw of plant actives






Table 2 - Power draw for modeled plant actives

Active type	Power Consumed (W)	Quantity	Total Draw (W)
Legacy Node	90	1	90
WiFi Access	55	2	110
Amplifier (LE)	35	9	315
Mini bridger	50	3	150
Total Customer Connectivity Load			665

- 2) Per the chart above the total Customer Connectivity Load is assumed to be 665W.
- 3) All actives are assumed to have a conversion efficiency of 95% which is included in power consumption estimates.
- 4) Based on 2&3 above, the theoretical minimum power to run the network actives at 100% plant efficiency would be 631.75W.
- 5) Power supply conversion efficiency is assumed to be 86% based on a 15A power supply deployed prior to 2010 running at this load. This represents more than 55% of the deployed population of broadband power supplies in North America today.
- 6) Plant passives are assumed to add approximately 10% to the total plant resistance
- 7) Coax carrying power in the plant is assumed to be a 75/25 mix of .75 and .625 P3 cable to represent real-world variation in cable types.
- 8) Battery charging and management overhead is assumed to be a fixed 10W.
- 9) There is assumed to be 1 passive for every span or approximately 200 ft.

Inputting the baseline assumptions into the computer model, here are the baseline results:

Table 3 - Baseline power breakdown of modeled network

Original Baseline	
 Required power for connectivity	631.75 Watts
 Conversion Loss (Actives)	33.25 Watts
 Transmission Loss	59.48 Watts
 Battery Management Overhead	10 Watts
 Conversion Loss (Power Supply)	119.57 Watts
Total Utility Power Draw	854.05 Watts






So, based on the above assumptions, my local network has a customer connectivity load 631.75W. This is the theoretical minimal power that is needed to keep the actives in this section of plant functioning if no other power was consumed in the plant. This is impossible to attain, but the goal is to get as close as possible with reasonable investment. The annual utility usage cost of this power supply based on the US average utility price of \$0.16 per KWh is about \$892.

4.2. Conversion Loss impact scenarios

4.2.1. Power supply conversion loss

Based on our baseline above, this is the most significant loss factor so let's look at the real-world impact of improving power supply efficiencies on conversion loss. We will start by looking at the impact of upgrading to a newer 15A power supply as a part of regular planned replacements. From there we will compare the impact of load-matching a more appropriate 10A power supply to this section of plant. Finally, we will look at the impact of a theoretical power supply that can maintain a maximum efficiency of 94% across a wide range of loads. See the results in the table below.

Table 4 - Model of power consumption varying power supply efficiency






	Original Baseline	Newer 15A PowerSupply	Load-Matched 10A PS	Theoretical High-Efficiency PS
 Required power for connectivity	631.75 Watts	631.75 Watts	631.75 Watts	631.75 Watts
 Conversion Loss (Actives)	33.25 Watts	33.25 Watts	33.25 Watts	33.25 Watts
 Transmission Loss	59.48 Watts	59.48 Watts	59.48 Watts	59.48 Watts
 Battery Management Overhead	10 Watts	10 Watts	10 Watts	10 Watts
 Conversion Loss (Power Supply)	119.57 Watts	100.16 Watts	80.7 Watts	46.88 Watts
Total Utility Power Draw	854.05 Watts	834.64 Watts	815.18 Watts	781.36 Watts

Overall, strategies to increase power supply efficiency seem to have good results in our model network. Moving from an aging power supply to a Load Matched 10A power supply yields a 4.5% reduction in overall power consumption or an annual savings of around \$40 annually. Any additional economically viable efficiency that can be gained toward the maximum theoretical efficiency above should obviously continue to be pursued as well.

4.2.2. Conversion loss in active devices

Maximizing the efficiency of power conversion in plant actives is not an easily deployable strategy as it requires multiple touch points for every section of plant. For example, in our model network there were 15 actives versus just one power supply. That being said, as we transition networks from current architectures to the near-future network discussed earlier, we must remember to be vigilant about ensuring maximum efficiency in plant actives as gains or losses out the plant have a cascading effect that you can see in the table below. This scenario models hypothetical increase in efficiency of all actives from 95% to 98%.

Table 5 - Model of improved power conversion efficiency of actives






	Original Baseline	Plant actives at 98% efficiency
 Required power for connectivity	631.75 Watts	631.75 Watts
 Conversion Loss (Actives)	33.25 Watts	12.89 Watts
 Transmission Loss	59.48 Watts	55.52 Watts
 Battery Management Overhead	10 Watts	10 Watts
 Conversion Loss (Power Supply)	119.57 Watts	115.61 Watts
Total Utility Power Draw	854.05 Watts	825.77 Watts

This scenario is interesting as we get to see just how complex plant power is through the ripple effect that happens by changing the load on the plant. Gains or losses in conversion efficiency at the actives can create carry an additive change of up to 30% by impacting transmission loss and power supply conversion loss. Notice here that by reducing active load by 20W the rest of the plant load was reduced by an additional 7W, or almost an additional percentage point.

4.3. Transmission loss impact scenarios

In the prior scenario we began to see a small impact on transmission loss when adjusting slightly the amount of power drawn by actives. However, the best strategy to reduce transmission losses, as previously mentioned, is to increase the voltage at which power is delivered to actives and thereby reduce current flowing through the plant. In this scenario we will take a small step backwards to show what our model network would look like if it were running at 60VAC and then take a big step forward to the potential future scenario of powering at 150VDC.

Table 6 - Model of the impact of increased voltage on transmission loss

	Plant modeled at 60V	Original Baseline	Theoretical 150VDC
 Required power for connectivity	631.75 Watts	631.75 Watts	631.75 Watts
 Conversion Loss (Actives)	33.25 Watts	33.25 Watts	33.25 Watts
 Transmission Loss	154.67 Watts	59.48 Watts	18.46 Watts
 Battery Management Overhead	10 Watts	10 Watts	10 Watts
 Conversion Loss (Power Supply)	135.06 Watts	119.57 Watts	112.89 Watts
Total Utility Power Draw	964.73 Watts	854.05 Watts	806.35 Watts
Voltage at last active	47.2 Volts	79.8 Volts	145.3 Volts

The biggest thing that jumps out in these scenarios is the power savings going from 60VAC to 90VAC, where we see the impact of transmission loss in older sections of plant and again the additive effect that loss can have on power supply conversion loss. Fortunately, adjusting the output voltage of the plant in most locations can be done in a few minutes without an upgrade to the power supply. Unfortunately, as many operators are already aware, in some of these sections of plant actives would need to be swapped out to handle higher voltages.

The theoretical 150VDC power supply does show significant power savings due to a big reduction in transmission loss. Another key detail to note in the model is the voltage of the last active in each scenario. As mentioned earlier, potentially the biggest gain from going to an increased output voltage like this is the ability to combine sections of plant and reduce power supply locations that need to be managed by plant operations. In 60V plant not only does current need to increase to deliver adequate power, but voltage decreases much more quickly and only needs to drop 18V (from 63VAC to 45 VAC) before it can no longer power actives. In 90V plant, voltage drop toward the minimum voltage for actives is much less pronounced, but would still be impacted significantly if the powering needs of two sections of plant were combined. Additionally, as the next generation of actives with higher power draw are deployed, the additional current required to power these actives will accelerate voltage drop and power loss beyond what we see in the scenarios above. In 150V plant, voltage drops so insignificantly, and the gap between power supply output voltage and minimum voltage for actives is so great, that two sections of plant could easily be combined without adding significant additional losses or risk of dropping actives due to low voltage.

4.4. Battery Charging and monitoring overhead






Assuming that the load of the cable modem used to monitor the health of network power is a given and that these devices are at maximum efficiency, reductions in this can be modeled based on a simple binary scenario. If we can, as a part of near-future network upgrades, leverage an energy storage technology that reduces or eliminates the need to float charge, we can eliminate 4 Watts of draw on the HFC network or the equivalent of 35 KWh annually. No need to create an elaborate model for this, however, there are other strategies to leverage the existing energy storage to create additional energy or utility cost savings that we be discussed in a moment.

4.5. Cumulative effect of power savings strategies

One last scenario to look at is the potential cumulative effect of implementing all of these strategies to have an idea of how much wasted energy we could theoretically remove from the plant. In this scenario, we will optimize all the variables to the limit of their realistic values.

- Increase conversion efficiency in actives to 98%
- Increase power supply conversion efficiency to 94%
- Increase line voltage to 150 VDC
- Reduce the additive resistive load from passives to 5%
- Remove battery float charging of 4W from the battery overhead

Table 7 - Cumulative effect of power savings strategies

	Original Baseline	Theoretical Best Case
 Required power for connectivity	631.75 Watts	631.75 Watts
 Conversion Loss (Actives)	33.25 Watts	12.89 Watts
 Transmission Loss	59.48 Watts	16.48 Watts
 Battery Management Overhead	10 Watts	6 Watts
 Conversion Loss (Power Supply)	119.57 Watts	42.84 Watts
Total Utility Power Draw	854.05 Watts	709.96 Watts

While this model showing a 17% reduction in power draw is only theoretical, it is a realistic view of what is possible. It emphasizes the point that small cumulative gains will add up to significant progress in reducing our industry's footprint and drastically reducing operational costs. Using this average example as a baseline, if we could realize similar savings across all 750,000 power supplies in North America, our industry would be able to save nearly a terawatt hour of energy, 800,000,000 lbs of CO₂, and more than \$110M of utility operational cost. In addition to that, raising plant voltage and increasing reach could save millions of additional dollars in maintenance costs.

The biggest question for many of these opportunities is how to make the solutions economically viable. To achieve many of the gains we've identified, investment is required and the payback period for any of these initiatives individually often makes them difficult to justify. One solution may be to pair these concepts for power savings into network upgrades already planned for the access network of the future, specifying key power saving elements into next-gen actives and passives. Additionally, prioritization of plant upgrades for next gen architectures could factor in opportunity for power savings at each sites so that gains are accelerated.

5. Other power reduction strategies

While significant gains can be made by targeting wasted energy in the plant, that strategy is limited. In our model example our maximized efficiency scenario removes 144W of grid energy and cost but still leaves 710W of dependency. To accelerate sustainability efforts it will be necessary to find additional ways to reduce energy draw from the grid. Here is a brief overview of two other concepts that should be considered to reduce grid dependence and operational spend.

5.1. Solar power augmentation

Renewable sources of energy bring the promise of sustainability and reduction of our carbon footprint. While many sources of renewable energy don't scale well to an OSP powering installation, solar energy not only scales effectively, but also could have the additional benefit of shading installations from excess solar load.

Over the past decade, the cost of an installed renewable energy system has gone down between 60% and 70%^{iv, v, vi}, to the point where the cost of solar energy is around \$2.50 per Watt fully installed. This means that two common 370W panels would have an installed cost of less than \$1500, including permitting. In

most locations in the US these panels would generate around 3.2 KWh of power per day. With the average site drawing 700W, solar production could reduce overall energy usage by an additional 20% or around \$140 of annual savings per site, with potential upside from more intelligent implementations. This could be scaled up to four or more panels where more space is available to cover larger loads.



Figure 11 - Example of a solar augmented cable power supply

5.1.1.– Options for solar implementation in the HFC plant

There are a couple of options for installation of solar in the OSP. The first is to simply leverage the real-estate of the plant to install an AC based grid-tied solar system and sell back energy produced by the panels to the utility company. This option is generally considered the simplest to implement, but it has several drawbacks. First, most utilities buy back at wholesale rates and not the retail rate that users generally pay so the payback model is reduced. Second, the power cannot be used for any other plant benefits like increasing runtime in an outage or charging batteries with free power after an outage. Finally, this type of "unintelligent" solar production has flooded the grid and created uneven demand for utilities. Thus, utilities have begun to discourage this scenario by removing incentives and increasing costs for permitting and metering of these sites.

An alternative, patented method exists to harvest solar energy locally at the point of load. This approach, which is shown below in **Figure 12**, creatively leverages the typical powering elements already existing in the OSP and does not involve any grid interactive utility permitting. By coupling the solar power directly onto the battery (DC) bus, the harvested energy can be utilized automatically for battery charging or extended runtime in an outage. By adding supplemental battery capacity, the system can also be used intelligently for offsetting utility power consumption when it makes the most sense (e.g. higher billing hours). The key element of the system is the master controller which can tailor solar consumption to various system sizes, standby runtime requirements, and utility billing scenarios to optimize the available solar resource and minimize utility billing while still maintaining critical backup capacity.

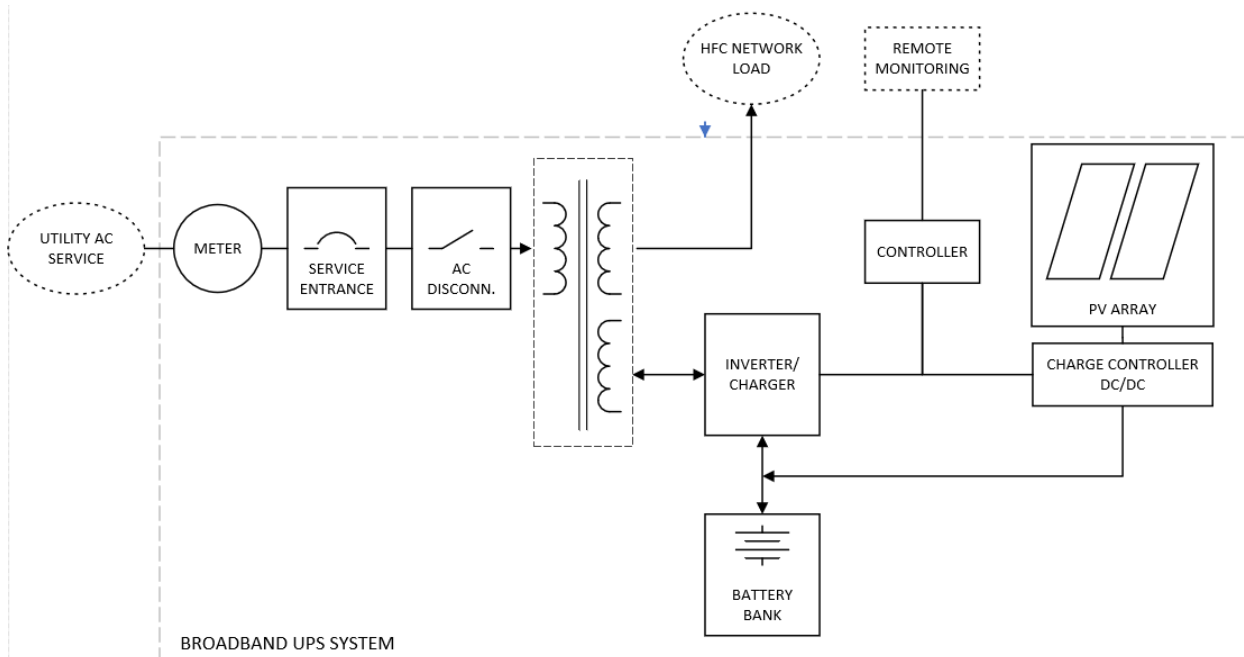


Figure 12 - Basic diagram of a cable power supply augmented by solar

5.2. Time-of-use mitigation

Large scale generation of energy is not inherently scalable, especially with cleaner sources of energy like solar and wind. When customer demand increases past a certain point, utility companies pay premiums to turn up secondary “Peaker Plants” to meet demand. Many large utility companies have rate structures that increase the price of energy during defined peak hours to encourage customers to reduce energy use during these times. This creates another interesting opportunity for reduction of operational energy cost as most OSP installations already have energy storage on site for emergency backup.

Business

Business Time Periods and Delivery Rates

	Peak	Off-Peak
Hours	8 a.m. to 10 p.m.	10 p.m. to 8 a.m. and all day on weekends
TIME-OF-USE DELIVERY RATES		
June 1 to Sept 30	29.38 cents/kWh	1.08 cents/kWh
All other months	14.47 cents/kWh	1.08 cents/kWh

Figure 13 - Example utility time-of-use rate structure

By slightly oversizing the battery array and deploying a battery with the capability for frequent cycling, batteries can be used to power the plant during hours with more expensive rates and recharged later with lower cost energy. For example, using the rates above a power supply drawing 700W from the utility would about \$3 a day to power in the summer. By using batteries to remove it from the grid during peak rates for 4 hours and recharging later in the day when energy is inexpensive you could reduce your energy bill by \$1 each day in the summer.

With major utility companies deploying time of use rate structures with peak rates as high as \$0.57 per KWh this strategy should be explored further. Although, this strategy does not actually reduce total energy use, peak energy usage historically tends to be more carbon-intensive. Thus, this strategy does reduce carbon footprint, if not overall energy consumption. However, by combining time-of-use mitigation with solar augmentation and building intelligence into network power supplies to use the energy source that is most economically viable at any point of time - grid, solar, or batteries – This strategy can yield significant savings in power usage and operational energy spend.

6. Conclusion

As we progress along the path to 10G and beyond it is imperative that we continually seek out new solutions to reduce energy consumption of the network. There is significant opportunity to improve the way we power the access network and specifically reduce the inefficiencies in the plant. The possibility of leveraging renewable energy to reduce grid dependency of the OSP is becoming more economically viable. There is opportunity in front of us, however there is much work to be done.

The power map that we have created in this paper and the mathematical model used to analyze it are an early step to realizing energy savings. The opportunities identified above must each be investigated in depth and economically viable solutions developed to address them. More efficient power conversion, reduced transmission losses, and renewable energy augmentations all can have huge impact on energy usage. Even smaller changes like ensuring minimal resistive impact from network passives and reducing the need to float charge backup batteries can have a cumulative impact that can save millions of dollars and megawatt hours of energy. With significant network upgrades on the horizon to facilitate the path to 10G and beyond, the time is now to understand everything that can be done to reduce power consumption in the access network and ensure that near-future network is as energy efficient as it is fast.

Abbreviations

AC	alternating current
AGM	absorbent glass mat
AN	aggregation node
DAA	distributed access architecture
DC	direct current
HFC	hybrid fiber coax
MDU	multi-dwelling Unit
MSO	multiple-system operator
OLT	optical line terminal
OSP	outside plant
P2P	peer-to-peer
PON	passive optical network

RDOF	Rural Digital Opportunity Fund
UPS	uninterruptible power supply
VRLA	valve regulated lead acid
kWh Google	kilowatt hour

Bibliography & References

R. Anderson, "Network Power Considerations for 10G Enhancements", SCTE- Technical Journal, vol. 1, no. 1, pp.23-37, June 2021. Institute of Electrical and Electronics Engineers, and American National Standards Institute. National Electrical Safety Code C2-2017. 2017th ed., Institute of Electrical and Electronics Engineers, 2017.

ⁱ <https://www.scte.org/standards/energy-2020/energy-2020-powering-cables-success/>

ⁱⁱ <https://www.eia.gov/tools/faqs/faq.php?id=74&t=11#:~:text=In%202020%2C%20total%20U.S.%20electricity,CO2%20emissions%20per%20kWh.>

ⁱⁱⁱ <https://www.eia.gov/electricity/data.php>

^{iv} <https://www.marketwatch.com/picks/guides/home-improvement/solar-panel-costs/>

^v <https://www.solarreviews.com/blog/how-has-the-price-and-efficiency-of-solar-panels-changed-over-time>

^{vi} <https://www.nrel.gov/news/program/2021/documenting-a-decade-of-cost-declines-for-pv-systems.html>

EnerSys is a trademark and the property of EnerSys and/or its affiliates. Other brand and product names may be trademarks or registered trademarks of their respective holder(s). The information contained herein is subject to change without notice. EnerSys shall not be liable for technical or editorial errors or omissions contained herein.

Zero Trust Security Architecture for The Enterprise

A Technical Paper prepared for SCTE by

Sarah Weinstein

Vice President, Engineering
Comcast

Philadelphia, PA
(215) 286-3412

sarah_weinstein@cable.comcast.com

Christopher Zarcone

Distinguished Engineer
Comcast

Moorestown, NJ
(856) 638-4116

christopher_zarcone@cable.comcast.com

Stephen Zevan

Director, Product Management
Comcast

Philadelphia, PA
(215) 286-8275

stephen_zevan@cable.comcast.com

Table of Contents

Title	Page Number
1. Introduction.....	4
1.1. Zero Trust Architecture	4
2. Background	4
2.1. History of Zero Trust	4
2.2. Elements of ZTA	4
2.2.1. Foundational	4
2.2.2. Security Support Structure.....	5
2.3. ZTA Fundamentals.....	6
2.4. ZTA Frameworks.....	8
2.4.1. Google – BeyondCorp	8
2.4.2. NIST – Special Publication 800-207	9
2.4.3. Other Frameworks	9
2.4.4. Using the ZTA Frameworks	10
2.5. Scope of ZTA	10
2.6. Embarking on ZTA	10
3. Brief History of ZTA at Comcast.....	11
3.1. 2019 – The Start of Our ZTA Program.....	11
3.2. 2020 – COVID-19.....	12
3.3. 2021 to Present – Acceleration	13
4. Guiding Principles	13
5. Security Domains	14
5.1. Security Hygiene	14
5.1.1. User Identity & Access Management.....	15
5.1.2. Resource Identity & Access Management.....	16
5.1.3. Asset Ownership	16
5.1.4. Device Identity & Management.....	16
5.1.5. Visibility & Hardening	17
5.2. Microsegmentation.....	18
5.2.1. Definition	18
5.2.2. Example	19
5.3. Application Access	20
5.3.1. Architecture	21
5.3.2. Policy Enforcement Point.....	22
5.3.3. Policy Decision Point	22
5.3.4. Enterprise Risk Engine	22
6. Program Governance.....	23
7. Lessons Learned.....	24
7.1. Engagement.....	24
7.2. Cloud First.....	24
7.1. Buy vs. Build.....	25
7.2. Application Owner Collaboration.....	25
7.3. Business Partner Collaboration.....	26
8. Future Work.....	26
8.1. The Multi-Year Journey	27
8.2. Centralized and Continuous Risk Evaluation.....	27
8.3. Microsegmentation.....	27
8.4. Converged Application Access	28
9. Service Provider Considerations.....	28
9.1. Infrastructure Hardening	28
9.2. Network Visibility	29
10. Conclusions.....	30

Abbreviations	31
Bibliography & References.....	32

List of Figures

Title	Page Number
Figure 1 – Users, Devices & Resources	5
Figure 2 – Security Support Structure.....	6
Figure 3 – Establishment of Trust Via Authentication.....	7
Figure 4 – Establishment of Trust Via Risk & Compliance Assessment	8
Figure 5 – Classic Physical Perimeter Defense Architecture	11
Figure 6 – Shipboard Compartmentalization	19
Figure 7 – An Unsegmented Network and a Microsegmented Network.....	20
Figure 8 – Access Proxy Arbiting Access to Resources.....	21
Figure 9 – Exploded View of Access Proxy Components.....	22
Figure 10 – Converged Access Proxy Architecture	28

1. Introduction

1.1. Zero Trust Architecture

Zero Trust Architecture (ZTA) is an information security model that de-emphasizes computer networks as trust factors, focusing instead on strong user & device authentication and contextual, risk-based authorization. In this paper, we:

- Review the background and history of ZTA (Section 2).
- Summarize the history of our organization's Zero Trust journey (Section 3).
- Discuss highlights of our organization's ZTA program, including its guiding principles, security focus areas, and governance (Sections 4 – 6).
- Reflect on lessons learned (Section 7).
- Consider future directions for our ZTA program (Section 8).
- Evaluate ZTA as it relates to service providers (Section 9).

2. Background

2.1. History of Zero Trust

Zero Trust Architecture is an information security model based on the principle of **Never Trust, Always Verify**. The term “zero trust” can be traced to a 1994 doctoral dissertation by Stephen Paul Marsh of the University of Stirling, and the concept was later popularized in a 2010 whitepaper by John Kindervag of Forrester Research. ZTA argues that the traditional perimeter network security model is obsolete, and that modern security programs should instead focus on authentication of security principals, security policy compliance, and continuous risk assessment.

ZTA assumes that all computer networks are untrusted by default, and that enterprise networks are no different – and hence no more secure – than non-enterprise or public networks. Viewed in this light, Never Trust, Always Verify requires:

- Strongly authenticating, authorizing, and auditing all access to resources and services.
- Measuring user and device compliance against organizational security policy.
- Continuously assessing the security posture of all principals (users, devices, resources).

2.2. Elements of ZTA

2.2.1. Foundational

The foundational elements of a ZTA are:

- **Users.** Users are human actors and include the organization's employees, contractors, business partners, and other third parties.
- **Devices.** A device is a physical computing resource designed to be used interactively by a user. Desktop computers, laptop computers, smartphones and tablets are all examples of user devices.

- **Resources.** Resources are repositories of organizational data, and as such, primary targets of attack. They include both physical and virtual machines, as well as the server operating systems, containers, services, and applications running on such machines.

At the most basic level, Users use Devices to access Resources. All three elements are security principals, in that they have identities and can attempt to perform actions. A ZTA authenticates identities and authorizes actions as required.




Figure 1 – Users, Devices & Resources

Additional dimensions to each element must be considered as well:

- For any given Resource, User audiences might vary (employee, business partner, guest, etc.). Additionally, the Users may be local to one of the organization's offices, or they may be remote.
- Resources might be owned and administered organizationally at one or more premises locations. Alternatively, the Resources might be hosted by a service provider (e.g., **Cloud** services).
- Devices might be owned and administered organizationally. Alternatively, Devices might be owned and administered by a third party, or personally owned by their Users.

2.2.2. Security Support Structure

The Security Support Structure (SSS) is an all-encompassing term used to describe the set of services that enable ZTA. To varying degrees, the foundational elements depend on each of these services. Common SSS components include:

- Application Access Frameworks – on-premise, off-premise
- Configuration Management Databases (CMDB)
- Directory Services (LDAP)
- Endpoint Detection & Response (EDR)
- Identity & Access Management (IAM) & Identity Providers (IdP)
- Multi-Factor Authentication (MFA)
- Mobile Device Management (MDM)
- Public Key Infrastructure (PKI)
- Security Information and Event Management (SIEM)
- User Behavior Analytics (UBA)

Collectively, the SSS bolsters the Zero Trust posture of Users, Devices, and Resources. An example SSS can be depicted as follows (significant integrations illustrated):

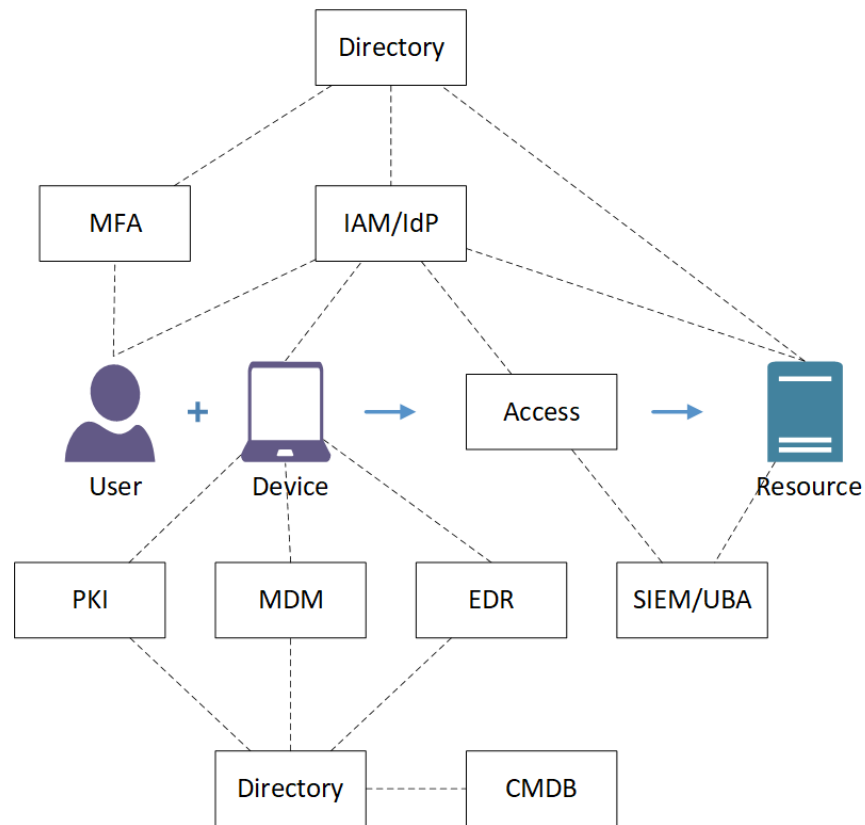


Figure 2 – Security Support Structure

2.3. ZTA Fundamentals

All Users and Devices begin a work session in a state of Zero Trust. That is, absolutely no security relevant attributes are assumed or inferred about a User or a Device. Examples:

- Assertions of Identity
 - “I’m Alice!”
 - “I’m Bob’s computer!”
- Assertions of Network Location
 - “I’m on the Internal network!”
- Assertions of Compliance
 - “I’m fully patched!”
 - “My storage is encrypted!”

Users and Devices must therefore establish trust before obtaining access to Resources. Trust is established in three primary ways:

- **Authentication.** User and Device identities must be confirmed with high confidence. This typically requires strong authentication, which can be defined as:
 - Multi Factor Authentication (MFA) for interactive User access. Examples include knowledge factors (passwords, PINs) combined with physical or logical possession-based factors (time- or event-based One-Time Password (OTP) generators, digital certificates, challenge/response mechanisms, FIDO 2 tokens).
 - Cryptographically sound authentication for Device-level, non-interactive access. Examples include shared symmetric keys, asymmetric key pairs, digital certificates, refresh tokens, and so forth.

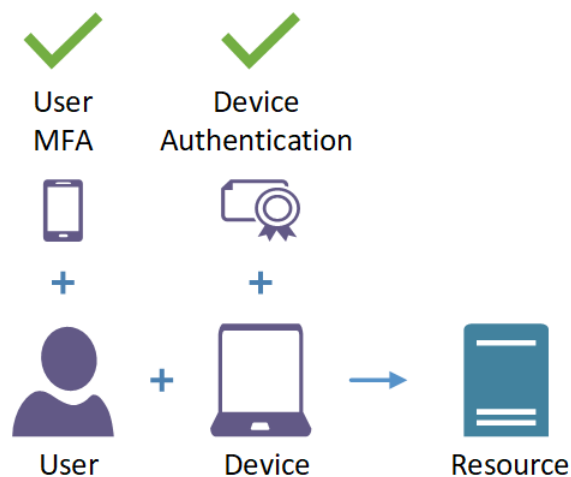


Figure 3 – Establishment of Trust Via Authentication

- **Security Policy Compliance.** Users and Devices must be assessed for compliance with organizational security policy. This could include factors such as:
 - User Authorization Checks
 - There are many possible criteria for user authorization, including directory group memberships and/or other directory attributes, requested Resource, work schedule (time of day, day of week, etc.) and so forth.
 - Device Authorization Checks
 - Determination that device properties and configuration are acceptable (e.g., idle UI timeouts are enabled, storage media is encrypted)
 - Verification that all expected security tooling is installed and operational
 - Confirmation that OS and/or security tool patch levels are suitably recent

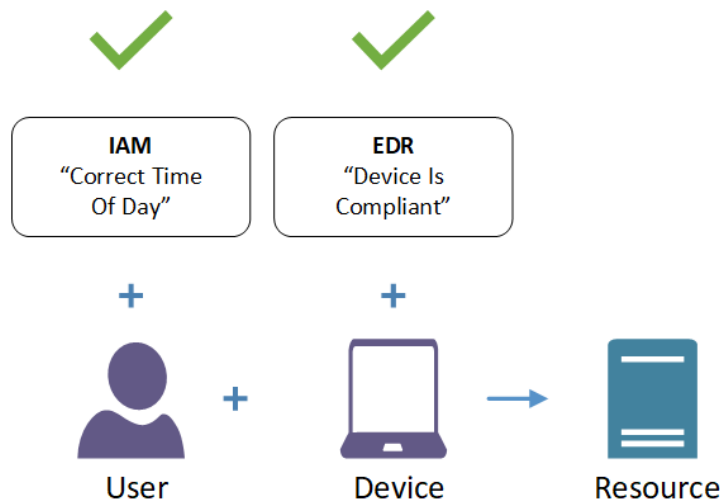


Figure 4 – Establishment of Trust Via Risk & Compliance Assessment

- **Continuous Risk Assessment (Context).** Risk is not a static property to be evaluated only once, at session initiation. Rather, risk is a dynamic property, and the level of risk posed by a User or Device can change mid-session. For example, a Device could move from a trusted state to an untrusted state (e.g., EDR detects a malware infection, or Device IP address reputation changes from benign to malicious). For this reason, User and Device risk should be continuously assessed throughout the duration of a work session. This continuous risk assessment is referred to as Context.

Once a sufficient level of trust in a User and a Device has been established, access may be granted to one of more Resources as specified by policy.

2.4. ZTA Frameworks

The descriptions above provide a generic, high-level understanding of basic ZTA principles. In practice, however, ZTA cuts across many different policy and technology domains, and implementing a full ZTA program can require considerable effort and attention to detail.

Industry, academia, and standards bodies have responded to this challenge by producing several reference frameworks for ZTA, to guide organizations in their approach to Zero Trust. Here, we review a few examples.

2.4.1. Google – BeyondCorp

Starting in 2014, Google published a series of influential whitepapers documenting their own multi-year Zero Trust implementation. Collectively referred to as “BeyondCorp,” the six whitepapers propose Google’s vision of a ZTA, with special emphasis on migrating to a network architecture without traditional perimeter controls.

Google's ZTA features five design objectives:

1. Securely Identifying the Device
2. Securely Identifying the User
3. Removing Trust from the Network
4. Externalizing Applications and Workflows
5. Implementing Inventory-Based Access Control

In general, BeyondCorp shifts access controls from network edges and perimeters to individual User and Device identity, embracing the oft-repeated maxim that “identity is the new perimeter.” This focus allows an organization's employees, contractors, and business partners to work more securely and remotely from any location without the need for a traditional remote access technologies like Virtual Private Networks (VPN). Key to this architecture is a next-generation access model for resources, an Internet-Facing Access Proxy (alternatively called an “Identity-Aware Proxy” or simply an “Access Proxy”). The Access Proxy enforces the security requirements of traditional VPN – strong authentication, strong encryption, etc. – in a centralized manner, without the protocol overhead of IPSec and similar tunnelling approaches.

2.4.2. NIST – Special Publication 800-207

In 2020, National Institute of Standards and Technology (NIST) developed its own framework for ZTA (NIST Special Publication 800-207). The standard does not propose a definitive architecture for Zero Trust; instead, it proposes several architectural variations, emphasizing the advantages and disadvantages of each. Each of these variations, however, reflect a common set of design goals (or “tenets” as they are described in the standard). They are as follows:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.
5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
7. The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications and uses it to improve its security posture.

2.4.3. Other Frameworks

Other Zero Trust frameworks of note:

- Although not exclusively billed as a “Zero Trust” architecture, Gartner's Continuous Adaptive Risk and Trust Assessment (CARTA) framework embodies many basic ZTA principles, such as continuous risk assessment.

- The US Cybersecurity and Infrastructure Security Agency (CISA) has published a “Zero Trust Maturity Model,” intended primarily for consumption by federal agencies. CISA framework focuses on five domains, each with maturity levels ranging from “traditional” to “optimal.”
- The United Kingdom’s National Cyber Security Centre (NCSC) outlines a six-pronged Zero Trust network architecture in its broader published guidance for device security.

2.4.4. Using the ZTA Frameworks

There is no uniform, “one size fits all” approach to ZTA. Organizations will have to evaluate their current IT strengths, capability gaps, and the overall threat landscape to develop their own unique approach to ZTA. To the extent that ZTA frameworks (such as those mentioned above) can guide and assist with this effort, so much the better.

2.5. Scope of ZTA

It should be noted that the scope of a ZTA typically applies to:

- An organization’s enterprise Users, Devices, and Resources.
- Third parties (contractors, business partners, vendors, etc.) performing enterprise duties on behalf of the organization.

Although the same fundamental concepts apply, information security management of an organization’s customers or subscribers should be treated separately from its enterprise needs.

2.6. Embarking on ZTA

It is often stated that Zero Trust is not a specific product or technology. Rather, it is an assemblage of multiple components, all operating in concert, to realize a specific security posture. As such, any approach to ZTA is most fundamentally an effort in systems integration. In terms of project management, organizations would be well advised to approach ZTA from this perspective.

It is also often stated that Zero Trust is not a specific project deliverable or operational practice, but rather a “journey.” This choice of language reflects the reality that – for all but the smallest of organizations – migrating to a Zero Trust posture will require sustained effort over a significant period of time. For example:

- In a 2022 executive memorandum from the US Office of Management and Budget titled “Moving the U.S. Government Toward Zero Trust Cybersecurity Principles,” OMB noted that “Transitioning to a zero trust architecture will not be a quick or easy task for an enterprise as complex and technologically diverse as the Federal Government.” The memo further requires that specific Zero Trust goals be met within a three-year timeframe (end of Fiscal Year 2024).

Accordingly, claims of “overnight” transformations from legacy networks to Zero Trust should be regarded with skepticism. The Zero Trust journey is one of incremental successes – it is a marathon, not a sprint.

3. Brief History of ZTA at Comcast

3.1. 2019 – The Start of Our ZTA Program

Comcast has been on its ZTA journey since 2019. The journey began as most large initiatives do – by building the business case for Zero Trust and why it was important to the organization.

At the time, the need to consider new security strategies was self-evident – the rate of publicly disclosed security incidents was increasing dramatically. To cite just one example of many, in the third quarter of 2019, the RiskBased Data Breach QuickView Report noted that there had been 5,183 breaches during the year, exposing 7.9 billion records. Compared to the 2018 Q3 report, the total number of breaches had increased 33% year over year, and the total number of records exposed had increased 112%.

Zero Trust was quickly identified as the security strategy most worth pursuing. In terms of socializing this new approach to information security, the “Castle and Moat” metaphor really helped paint the picture whenever we needed to explain the concept of Zero Trust to internal stakeholders. For example, a traditional medieval castle may feature high walls made of brick or stone, perhaps further protected by a moat and a drawbridge-style gate. The walls and moat form the “perimeter” by which the residents of the castle are protected from external threats. Entry and egress to the protected castle is governed by the gate. However, if an intruder can get past the gate, they have unfettered access to what is inside the castle. There may be some impedance once past the walls and moat, but the reality is that access can be quite open.




Figure 5 – Classic Physical Perimeter Defense Architecture

"Drawbridge" by j4p4n. Licensed under Creative Commons Zero 1.0 License. openclipart.org

Replace “castle” with “internal and/or enterprise network” and “gate” with “firewall” and one arrives at the traditional network security perimeter model, which was developed in the late 1980s and served the Internet community well for the better part of two decades.

In our new reality, however, with the rise of business partner integrations, public cloud services, and application deployment at network’s edge, it is not enough to rely on the traditional network perimeter model. We need to verify every User, every Device, and every attempt to access Resources (systems, applications and information), every time.

The “Castle and Moat” analogy served to ground technical and non-technical professionals and executives, which in turn was important to gain their support. A tremendous amount of stakeholder engagement was required on how we were approaching Zero Trust because it required ongoing partnership across the organization, with entities including our network engineering team, information technology, and our geographic divisions. Another important partner in this endeavor was Comcast’s Cybersecurity Guild. With over 200 members, the Cybersecurity Guild is a rich and interactive community of employees who are cyber-minded or have day-to-day responsibilities for cybersecurity in their business units.

Although developing the business case involved significant effort, it found a receptive audience, as our executive leadership has long regarded cybersecurity as a critical partner to the business. We didn’t have to convince everyone that it was important to embark on the Zero Trust journey – everyone knew that we just had to do this. The challenge was figuring out how we were going to do this and devising a strategy in a relatively short period of time.

We soon began developing the overall structure of our Zero Trust program. Some of the earliest influences included the BeyondCorp framework (discussed in Section 2.4.1) and Microsoft, which was underway with enterprise solutions and SSS to enable Zero Trust at scale.

Our leadership, architects, engineers, and other stakeholders eventually coalesced around a three-pronged program structure, consisting of:

- **Guiding Principles**
- **Security Domains**
- **Program Governance**

We discuss these in more detail in Sections 4, 5, and 6, respectively.

3.2. 2020 – COVID-19

In early 2020, the COVID-19 pandemic emerged as a worldwide health concern. The pandemic created a seismic shift in how our employees and business partners worked, and everyone was living with the reality of remote work and further diminished network perimeters. Companies like ours could no longer rely on the safeguards that we took for granted before the pandemic – working together, in common facilities, on an enterprise network that provided a basic degree of assurance. Our new way of working made it much more clear why Zero Trust is so important.

Arguably, because of the COVID-19 pandemic, Zero Trust is even more relevant today than it was in 2019.

The threat landscape also changed significantly during COVID-19. The Harvard Business Review highlighted the challenge that with the migration to remote work during COVID-19, cyberattacks increased exponentially. Businesses saw more attacks of every kind, but the headline for 2020 was ransomware attacks, which increased 150% over the previous year. Equally worrisome, the amount of money paid by victims of these attacks increased more than 300% in 2020.

COVID-19, however, did not deter our ZTA efforts. Quite to the contrary, COVID-19 accelerated them. Moving a significant portion of our workforce to remote work re-emphasized the need to move beyond the traditional perimeter security model.

3.3. 2021 to Present – Acceleration

Recently (and particularly within the last one to two years) there has been more industrial awareness of security as a business enabler. Conversations have shifted, with business clients increasingly asking about their service provider's adherence to Zero Trust practices.

In May 2021, we saw the issuance of the White House's "Executive Order on Improving the Nation's Cybersecurity." The order states (emphasis added):

The United States faces persistent and increasingly sophisticated malicious cyber campaigns that threaten the public sector, the private sector, and ultimately the American people's security and privacy. The Federal Government must improve its efforts to identify, deter, protect against, detect, and respond to these actions and actors. The Federal Government must also carefully examine what occurred during any major cyber incident and apply lessons learned. But cybersecurity requires more than government action. Protecting our Nation from malicious cyber actors requires the Federal Government to partner with the private sector. ***The private sector must adapt to the continuously changing threat environment, ensure its products are built and operate securely, and partner with the Federal Government to foster a more secure cyberspace.***

4. Guiding Principles

As previously discussed, the Zero Trust model assumes that every security principal (every User and every Device) is a potential threat until verification. This assumption holds regardless of whether those principals' network location appears to be internal or external to the organization.

We also assume that despite our best efforts to adopt sound information security practices and secure our enterprise, incidents can and will occur. Many factors give rise to this reality, most notably Zero-Day exploits (exploits directed at previously unknown and/or undisclosed vulnerabilities).

These assumptions give rise to the two guiding principles of our ZTA:

1. **Never Trust, Always Verify.** Users and Devices must establish trust before obtaining access to Resources.
2. **Prepare To Be Breached.** Our networks, systems, and applications must be designed to limit the impact of incidents.

Principle #1 was discussed at length in Section 2 but bears repeating. Users/Devices begin a work session in a state of Zero Trust, and Users/Devices must therefore establish trust before obtaining access to Resources. Trust is established through authentication, policy compliance, and risk assessment. Importantly, Trust must *not* be established by the perceived physical or logical network location of a User or Device.

Principle #2 asserts that if incidents cannot be prevented with certainty, the next best strategy is to plan for their eventuality. The goal is to minimize the degree to which Resources are affected by incidents in the broader technical environment. This is often referred to as minimizing the “blast radius” of an incident, preventing spillover effects and lateral network movement.

These two guiding principles informed our work across our entire ZTA program. For every significant architectural decision, we would ask the same two questions:

1. How does this decision improve the trustworthiness of Users and/or Devices?
 - a. Or conversely, how does this decision further reduce our reliance on networks as trust factors?
2. How does this decision help mitigate the effects of an incident?

5. Security Domains

Next, we applied the guiding principles across three broad security domains:

- **Security Hygiene.** Hygiene concerns the security posture of our main ZTA principals – Users, Devices, and Resources.
- **Network Microsegmentation.** Microsegmentation focuses on partitioning our network environments into smaller, workload-specific enclaves.
- **Application Access.** Access governs how our Resources are exposed to Users and Devices, regardless of their physical or logical network locations.

Each of these domains is discussed in turn.

5.1. Security Hygiene

Like many large companies, Comcast has a complex IT environment that was built through years of product and platform growth and acquisitions. There was some foundational work that needed to be addressed to ensure that as a company, we were ready for our Zero Trust journey. This foundational work was called **Security Hygiene** (or simply “Hygiene”), defined as a focus on the security posture of our security principals: User, Devices, and Resources.

We defined the scope of this work across several parallel workstreams, termed **Pillars**. For each Pillar, we defined the problem space that we needed to resolve, and definitions of what success looked like. We also aligned on quantified targets for success wherever possible.

The Pillars were defined as follows:

- 1. User Identity & Access Management**
- 2. Resource Identity & Access Management**
- 3. Asset Ownership**
- 4. Device Identity & Management**
- 5. Visibility & Hardening**

5.1.1. User Identity & Access Management

Section 2.3 notes that strong authentication for Users plays a pivotal role in ZTA. This more broadly requires an examination of how an organization performs Identity & Access Management (IAM) for its Users. IAM refers to an entire suite of tools and services for provisioning, managing, and deprovisioning User identities. IAM is also used to manage User privileges and entitlements, as determined by a User's title, position, and/or organizational roles.

A core component of IAM is the Identity Provider (IdP). The IdP's primary responsibilities are:

- **Authenticating Users.** This can be accomplished using one or more of the mechanisms described in Section 2.3.
- **Asserting User Identities to Resources.** This can be accomplished using protocols such as Kerberos, Security Assertion Markup Language (SAML), and Open Authorization & OpenID Connect (OAuth, OIDC). Proprietary protocols also exist.

Ideally, an organization should implement one, centralized IAM platform with as few IdPs as possible (again, ideally one).

The alternative is multiple IAM platforms with competing/duplicative identity stores that complicates:

- User provisioning and deprovisioning
- Assignment of user authorizations and entitlements
- Compliance
- Reporting
- Far from least important, the User experience (e.g., access to some systems works whereas access to other systems does not, confusion over which credentials or MFA/SSO mechanisms to use for a given system, and so forth).

5.1.2. Resource Identity & Access Management

Resource IAM is essentially the complement to User IAM. If ZTA requires Users to produce strongly authenticated assertions of identity, it becomes necessary for Resources to consume those assertions. Otherwise, the identity assertions serve no purpose.

This generally requires Resources to align with:

- The same Identity Provider (IdP) employed for Users
- The same authentication protocols (SAML, OAuth/OIDC) supported by the IdP

The amount of effort required to integrate a Resource with a modern IdP varies considerably. It can range from the trivial (e.g., simple changes to a configuration file) to the substantial (e.g., installation and integration of third-party authentication libraries, custom code, and extensive testing). The latter is often the case with older, “legacy” Resources, for which modern protocols like SAML/OIDC did not exist at the time they were deployed.

5.1.3. Asset Ownership

Information systems cannot be protected if they are unknown to the organization. This requires an accurate inventory of:

- Devices
- Resources
- The organizational owners of such Devices and Resources

In most organizations, this requires the creation and maintenance of a CMDB or similar inventory/asset management system. As with user identities, the ideal is a single CMDB that provides complete coverage for all Devices and Resources in the organization. Multiple, competing CMDBs exhibit many of the same challenges experienced with multiple, competing IAM platforms – inconsistent views of asset ownership, impaired ability to protect unregistered assets, and so forth.

5.1.4. Device Identity & Management

As with User identity, Device identity is extremely important for purposes of ZTA. It allows organizations to assign specific roles and entitlements to the Device (for example, an administrative workstation, compared to workstations intended for normal workforce use), establish unique accountability for Device activity, and so forth.

Device identity should be established through cryptographically strong mechanisms. Examples include digital certificates, public key pairs (independent of certificates), bearer tokens such as OAuth refresh tokens, and so forth. These credentials are typically created at the time a Device is activated or joined to an administrative framework (such as an Active Directory domain, a Kerberos realm, and so forth).

Device management is equally important as establishment of identity. A Device is said to be “managed” if an organization has administrative control of the Device and can apply security policy to it; otherwise, the Device is considered unmanaged. An insufficiently managed or unmanaged Device poses risk to a ZTA. The device may lack appropriate security tooling like EDR, or may not be configured to require commonplace security controls, such as idle UI session timeouts.

A subset of device management, Mobile Device Management (MDM) refers to the administration of mobile devices, such as laptops, smartphones and tablets. MDM caters to the unique characteristics of mobile devices, in that such devices are typically more limited in their capabilities than server, desktop or workstation computers. Mobile devices are also more easily lost or stolen than their desktop counterparts; as such, MDM is frequently used to enforce confidentiality-preserving security requirements, such as storage encryption.

5.1.5. Visibility & Hardening

Visibility (often also called Telemetry) refers to the monitoring of systems and networks in real time.

Resource Visibility refers to the degree to which system-level properties and activity can be observed on a Resource. Examples include:

- Resource utilization (CPU, memory, storage, network & other input/output)
- Process and thread activity
- File- or filesystem-level artifacts (existence of absence of specific files, etc.)
- (Windows) Existence and/or values of specific registry settings
- Collection and/or examination of system logs

Depending on the level of visibility required, some management frameworks (such as membership in an administrative domain or MDM) may provide sufficient data. In other cases, the deployment and management of specialized agent software may be required.

Resource Visibility also assists with maintaining accurate asset management (Pillar 3).

By contrast, **Network Visibility** refers to the monitoring of network communications in real time. Ideally, such monitoring can be summarized into data flows, or sequences of network packets from sources to destinations.

Communications may be monitored:

- Via the capabilities of standard network infrastructure (both physical and virtual), such as routers and switches.
- Via specialized, media-specific collection devices, often referred to as network taps.

The purpose of network visibility is to detect signals that possibly indicate:

- Lateral movement within an organization
- Data exfiltration

Hardening applies to Devices, Resources, and network infrastructure (collectively, “nodes”) and primarily consists of:

- The disablement of all unnecessary network services on a network-attached node.
- Enabling all practical security controls on the node’s remaining network services.

The main objective is to reduce the **attack surface** of network-attached nodes.

In security taxonomy, “attack surface” is the set of all possible entry points into a Device, a Resource or other information system. A system with a large amount of attack surface is thought to be more vulnerable than a system with less attack surface. The rationale is simple – every entry point represents a potential avenue of attack, either in terms of vulnerability, incorrect configuration, or both. As such, the more entry points an attacker has to work with, the greater the likelihood that one of them can be successfully exploited (and lead to system compromise).

Fortunately, attack surface can be reduced through appropriate system hardening. However, hardening does have its practical limits. Some services must inevitably remain operational for a Resource to perform its intended function – a web server must respond to requests for HTTP content, for example, and a DNS server must respond to name service queries. Other services might need to remain for purposes of monitoring or management, such as SNMP or SSH.

5.2. Microsegmentation

5.2.1. Definition

Even with proper hardening, the overall attack surface of many environments continues to increase, in many cases simply due to the sheer number of Resources and other nodes that are internetworked. The inevitable result is:

- More humans interacting with machines.
- More machines interacting with machines.

In terms of ZTA, one way to address this challenge is through the adoption of **segmented** network architectures. A network is said to be segmented if it is physically or logically separated into smaller, isolated subnetworks. The goals of this approach are twofold:

1. Separate public-facing components from private (e.g., non-public-facing) components.
2. Separate unrelated components from each other.

Taken to its logical end state, this approach results in **microsegmentation**, or the subdivision of networks down to the level of discrete, application-specific workloads. In a microsegmented environment, each distinct application receives its own dedicated, logically independent subnetwork.

Many technologies exist to facilitate microsegmentation. Virtual Local Area Networks (VLANs) are a classic example, and variations of this approach (e.g., Private VLANs) exist as well. More contemporary approaches include:

- Software Defined Networks, which use a centralized component to make forwarding decisions for individual packets. OpenFlow is a classic implementation of SDN.
- Overlay networks (e.g., VxLAN), which are networks layered on top of other networks.
- Hypervisors, which can enforce network isolation between guest virtual machines.
- Host-based firewalls, in some cases complemented with agent-based implementations to facilitate firewall rule management.

It is important to consider the various different approaches to microsegmentation when building and deploying new applications and application components.

5.2.2. Example

A good motivation for microsegmentation comes from the shipbuilding industry. The watertight body of a maritime ship is referred to as the hull. Modern shipbuilding techniques separate the hull into “compartments,” or individual watertight subdivisions. The rationale for this design is that damage to any portion of the hull can (hopefully) be limited to a specific compartment, instead of flooding the entire hull and sinking the ship.

An example of compartmentalization is depicted in Figure 6. Here, one shipboard compartment has sustained damage and subsequently floods with water. Adjacent compartments, however, are sufficiently isolated from the damaged compartment, and hence do not flood. In effect, the severity and extent of the incident has been minimized.



Figure 6 – Shipboard Compartmentalization

This design can be extended into the worlds of computer networks and distributed systems. For example, consider Figure 7:




Figure 7 – An Unsegmented Network and a Microsegmented Network

Here we have an example of an unsegmented network (left). In typical “Castle and Moat” architecture, it has a substantial exterior perimeter but no interior perimeters. Should an adversary make it past the perimeter firewall and compromise a Resource, they can leverage that access as a foothold to further move laterally within the network. In the depicted example, an attacker manages to compromise one system (lower left), then uses that system to launch additional attack against other reachable systems in the same network (lower right, upper right).

Compare this to the microsegmented network depicted on the right, where each system occupies its own individual subnetwork. Each of these “micro” subnetworks is isolated from one another, and each features its own individual perimeter. Breaching the perimeter of one subnetwork (and the systems contained within) does not gain an attacker access to other, unrelated subnetworks. In this manner, and as with the example of the shipboard compartment, breaches and incidents can be isolated to the fewest possible number of systems. Not only does this architecture reduce the “blast radius” of an incident, but it also gives organizations more time to detect the incident and mount a response, before too much additional damage is done.

5.3. Application Access

Before delving into our ZTA access model for application Resources, it would be helpful to review the traditional access model that has accompanied “Castle and Moat” perimeter architectures.

In the traditional model, organizations deploy applications to their internal, trusted enterprise networks. Users access these applications from the safety of those internal, trusted networks, shielded from external abuse and attack by firewalls and other perimeter controls.

Now consider the use case where the Users are remote – other organizational locations, business partner locations, or perhaps the Users’ residences – and hence in untrusted network locations. In such cases, technologies like VPNs are often employed to provide secure connectivity over public networks, like the Internet. In the traditional model, VPNs provide paths to trusted networks over untrusted networks.

However, ZTA asserts that computer networks should be regarded as untrusted. Viewed this way, VPNs provide paths to untrusted networks over *other untrusted* networks, which is nonsensical. ZTA essentially puts all users into a perpetual remote access security posture, but traditional remote access mechanisms (like VPN) are not philosophically aligned with ZTA.

This situation gives rise to the Identity-Aware Proxy/Access Proxy models discussed in Section 2.4. Comcast decided to pursue deployment of Access Proxy infrastructure as the third domain of its broader Zero Trust program.

5.3.1. Architecture

With ZTA, the Access Proxy becomes the primary model by which Resources are exposed to Users and Devices.




Figure 8 – Access Proxy Arbiting Access to Resources

We decided to decompose the Access Proxy into three functional components:

- Policy Enforcement Point (PEP)
- Policy Decision Point (PDP)
- Enterprise Risk Engine (ERE)

This was done primarily so that we could:

- Develop each component independently, on its own release schedule.
- Scale each component horizontally within a given datacenter environment.
- Geographically load balance the components across our enterprise footprint.
- Facilitate an active/active configuration for high availability and redundancy.

An exploded view of these components and their integration can be depicted as follows:

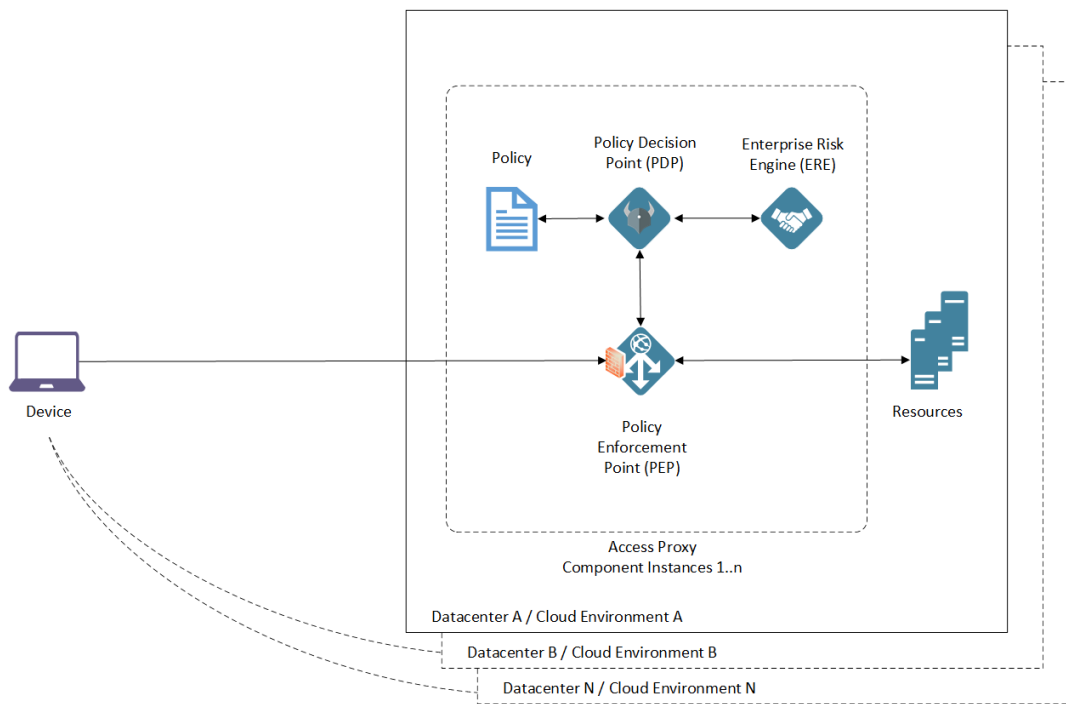


Figure 9 – Exploded View of Access Proxy Components

5.3.2. Policy Enforcement Point

The PEP is essentially an identity-aware reverse proxy. It is responsible for directly mediating connectivity between security principals (Users and Devices) and enterprise Resources. The PEP serves as the ZTA boundary between Resources and their microsegmented networks (which are managed by Resource owners) and all other networks.

5.3.3. Policy Decision Point

The PDP is primarily responsible for authorizing access to enterprise Resources. The PDP uses a combination of static policy (in the form of predefined security rules) and context (in the form of dynamic, near-real-time risk data and other signaling) when making access control decisions. Access control decisions are then enforced by PEPs as necessary.

5.3.4. Enterprise Risk Engine

The ERE is the primary aggregator and provider of context, primarily in the form of risk signals. A risk signal is any kind of data or metadata that provides contextual insight into the security posture of a User, a Device, or a Resource the User/Device is attempting to access. Examples of risk signals include geolocation, velocity, IP address reputation, and Device compliance.

The ERE continuously receives risk signals from sources of truth, including MDM, EDR and IP Reputation. It can then evaluate and share the risk levels of Users and Devices across the security ecosystem.

At of this writing, the PDP is the primary consumer of ERE services, which in turn considers risk signals provided by the ERE when making access control decisions. The ERE is essentially an adjunct used by the PDP to make context-aware access control decisions.

6. Program Governance

We believe that success is tied to how you engage with your enterprise partners. Any effort as large as a ZTA journey requires significant governance and project management. To this end:

- Our ZTA program was championed by our Executive Vice President, who is also our Chief Information Security and Product Privacy Officer.
- Product, Program and Architecture leaders were nominated to guide the overall effort.
- A team was defined to support the overall ZTA program. Each hygiene Pillar was assigned an Executive Leader, a Product Lead and a Program Lead. In most cases, this team was comprised by cybersecurity professionals; in other cases, it included partners from other departments within our organization (IT, Reliability Engineering).
- An Executive Committee was established to provide guidance on a quarterly basis that included leaders from our organization, our business units, Human Resources, Procurement, Legal, Compliance and Finance. We work for a very financially disciplined company. The meetings are a force and function to regularly remind us of the goals that we set, our progress, what's next and whether we need to shift course.
- A Steering Committee was established to provide stakeholder direction between the various organizations contributing development, engineering, and support resources to the program. This committee meets twice a quarter.

We engaged our Executive and Steering Committee members directly very early on in the program to ensure we were properly framing our strategy. They in turn connected us with the critical influencers in our organization who were very supportive and conveyed their commitment to our ZTA journey.

As ZTA started to gain traction, we expanded the formal means of communication to include:

- Cybersecurity Stakeholders – This includes engagement with our Business Information Security Officers (BISOs), Portfolio Leads and Cybersecurity Leads across the company.
- Employee Resource Groups – We closely partnered with technology resource groups to promote, inform and engage partners in our ZTA program:
 - Cybersecurity Guild – Our primary partner is our Cybersecurity Guild with hundreds of members. The Guild hosts a ZTA-themed event at least every quarter that is open to all of our Comcast technical resource groups.
 - We partner with other technical resource groups where they can provide necessary subject matter expertise. For example, we presented our architectural strategy as early concepts to Comcast's Software Ambassador's Guild, and we gained invaluable feedback that influenced our onboarding approach and commitment to automation wherever possible.
- Our ZTA scope is included in annual and quarterly roadmap planning across our technology and business units.

- We are very intentional about how we approach wider employee communications. We have a very tech-savvy workforce and recognized that ZTA would have an impact on how we work. We design our employee experiences and workforce communications around the north star of making it easy for our employees to do the right thing to protect our organization. We partner with our Communications Team to produce campaigns that make the concepts relatable. Cybersecurity Awareness Month is always observed in October, during which time we socialize important security trends. Last year, we introduced ZTA in a game called “Cyber Splash,” a popular smartphone application that is available to all our employees. Our “Trust Titan Superhero” made his big debut in October 2021.

As our ZTA program gained further traction, our BISOs soon became critical partners. The BISOs facilitate a streamlined engagement model for delivering security services to our business units. The BISOs worked with the leaders across Comcast to manage ongoing priority setting for ZTA programs and to continually assess feedback on how those programs were going. The feedback wasn’t always easy to hear. There are times when we had to address technical issues, process issues, and team issues. In addition, our technical and business partners also felt the strain of the convergence of several compliance programs. We view feedback as a gift and used the opportunity to improve our approach and for corporate leadership to stay better aligned.

7. Lessons Learned

7.1. Engagement

It is crucially important to get key stakeholders involved early. An organization’s IT function is the most obvious candidate for collaboration, and our security organization spent much time meeting with IT and aligning efforts. Other key partners included our Office of the CIO, Network Engineering, Procurement, Finance, Human Resources, and Corporate Communications. The governance framework described above ensured ongoing executive engagement and alignment with our key stakeholders.

7.2. Cloud First

Applying elements of ZTA to existing systems and infrastructure – particularly microsegmentation – can be a challenge. Early experimentation and prototypes confirmed this reality. As such, we initiated our ZTA journey with a priority on our new cloud deployments. We focused on enabling microsegmentation and other controls through automated policy. Over the course of one year, we trained our development teams in a “cabin crew” structure that allowed for imbedded expertise within each development organization. Later, we expanded the focus to our legacy applications that require a more hands on, consultative approach.

7.1. Buy vs. Build

In one sense, ZTA is like any other information technology initiative, in that solutions can be realized in many ways, each with its advantages and disadvantages:

- Purchase commercial solutions (licenses, subscriptions, etc.)
- Leverage open-source solutions
- Develop in-house, proprietary solutions

This is often referred to as the “Buy vs. Build” scenario.

More often than anticipated, we had to make buy vs. build decisions throughout our program. These buy/build decisions were driven by our requirements and speed. We had to say “no” to some prototype development efforts in the beginning of our journey, and that’s hard to do when you have talented architecture and engineering resources at your disposal. Like any strategic initiative, ZTA needed to deliver value for the company. In some cases, we built our own ZTA components and SSS; in others, we purchased commercial solutions and services. In the latter case, vendor partnerships and strategic integrations have helped us get to where we need to be.

In our experience, the key to successful ZTA is exploring and finding the optimal mix of all three approaches (commercial off-the-shelf offerings, open-source, custom development) and spending the time and effort to integrate these solutions.

7.2. Application Owner Collaboration

For Application Access, we developed an onboarding process to identify strategic applications and register them for access via our Access Proxy infrastructure. Our initial approach entailed identifying the application owners for such applications. We soon discovered that our CMDBs used for identifying and cataloging application and asset information can become outdated over time. To accommodate this, we employed an intensive “white glove” engagement strategy to help identify application owners. This impeded our initial progress, but we learned and adapted the application ownership step as part of the onboarding process if necessary.

The next step in the onboarding process was to identify application authentication mechanisms and determine if they met the minimum thresholds established by our IAM program. In our case, that minimum threshold was based on the OpenID Connect standard (built on top of the OAuth 2.0 protocol). This proved to be a challenge for some application owners, for two reasons:

1. The longer an application has been in existence, the higher the probability that its organizational ownership has changed, possibly multiple times; and
2. When these transitions occur, knowledge transfer for the application’s business logic often overshadows IAM knowledge transfer, due to time constraints and other factors.

The result is that the basic application security knowledge is sometimes lost. In such cases, the ZTA team works collaboratively with application owners to understand their authentication frameworks.

And finally, it goes without saying that implementing a comprehensive ZTA often requires application owners to undertake multiple security workstreams, sometimes in parallel. For example, an application team may be asked to simultaneously:

1. Modernize their authentication mechanisms
2. Migrate to a microsegmented network architecture
3. Onboard their application for Access Proxy integration

This can be a tall order for even the most well-resourced application owners. Multiply this single example by the total number of candidate applications in the environment, and one can see that this can seem like an overwhelming set of asks for an organization's application community.

It is imperative that the ZTA team convey and communicate to application owners "security is a team sport" and that everyone is in the ZTA journey together. Defining ZTA priorities at a program level will help build trust between the security and application teams and prevent the perception that ZTA is a series of individual asks, but rather a set of unified strategies to improve the overall security posture of the enterprise.

7.3. Business Partner Collaboration

We spent considerable time defining the quantitative goals for ZTA and the impacts across our internal business units. We had big goals (e.g., "harden X thousand Resources by Q4") and we had to define the intermediate milestones to get to those goals. The challenge was that we didn't always know exactly how we were going to get there. We persistently took stock of what we needed to address and defined risk-based priorities (e.g., "harden Internet-exposed Resources first"). Our business partners appreciated that we put so much thought into the strategy.

Still, this required a high degree of collaboration between Cybersecurity and our business and technical partners. Successful ZTA requires partnership across an organization, not competition. It was important to us that the collaboration meant that the outcomes were better because of teams working together. There were a lot of synergies across the Pillars and meeting security needs also addressed other needs for our business.

At the core of a team is solving those unsolved problems and making the most of our ecosystem. We had to become a provider not just of standards but of solutions for our enterprise. Our business units expect us to look out for them and to anticipate the needs to secure the business.

8. Future Work

In this section, we briefly examine opportunities for enhancement and improvement to our ZTA.

All estimates of this nature are subject to change due to the ever-dynamic needs of our business, as well as unforeseeable events and trends in the information security landscape. As the entire Internet community learned with the Log4j vulnerability of 2021, even the best plans can be upended by a single ill-timed Zero-Day.

8.1. The Multi-Year Journey

In Section 2.5, we noted that for all but the smallest organizations, ZTA is a multi-year endeavor. Overall, Comcast is making good progress on its ZTA journey.

We expect much (if not all) work in the Security Hygiene domain to be complete by the end of 2022. We have also been microsegmenting in cloud/virtual environments for the past three years, and our Access Proxy infrastructure entered production in 2021. These were all big mountains to climb. We are lucky to be a part of an organization where our employees appreciate leaning into a challenge.

Continued expansion of microsegmentation and Access Proxy infrastructure will be primary focus areas of 2023 and beyond. Most net-new deployments already align with our ZTA; it is the uplift of existing and/or legacy Resources that will command most of the effort.

8.2. Centralized and Continuous Risk Evaluation

One of the key tenets of the 2022 OMB memorandum (Section 2.6) is the need for a Continuous Diagnostic and Mitigation (CDM) process. A CDM process ensures that continuous policy evaluation and risk assessment can be conducted on Users and Devices during their active sessions. The memorandum further states that:

A necessary foundation for any enterprise-wide zero trust architecture is a complete understanding of the devices, users, and systems interacting within an organization.

The generally accepted practice in introducing and incorporating an enterprise wide CDM system into a company's security ecosystem is to:

1. Enable risk signal collection for all Users, Devices and Resources
2. Evaluate the signals in a policy driven manner
3. Gradually and incrementally enable and enforce the policies to improve the overall enterprise security posture.

Comcast has developed its ERE component, currently part of the Access Proxy infrastructure. This essentially serves as a CDM engine. Our future plans are to enhance the ERE with more risk signaling from more risk sources, ultimately including Resources and other end-user applications. The latter possibility is especially intriguing, as it would allow for suspicious application-level activity (e.g., Mallory transferring unusually large amounts of money in a finance application) to serve as a source of risk signals.

8.3. Microsegmentation

Our initial forays into microsegmentation have largely centered around cloud deployments (both on-premise cloud environments as well as hosted/public cloud environments). It is generally easier to microsegment in a virtual environment, as individual application workloads can be allocated their own virtual microperimeter environments (virtual switches, virtual networks, virtual firewalls, etc.)

In 2022 and beyond, we seek to retrofit the microsegmentation concept into “bare metal” and other non-virtualized environments. Several strategies have been considered (see Section 5.2 for examples) and our ultimate strategy may yet employ multiple approaches.

8.4. Converged Application Access

As an organization with a significant mobile/remote workforce (made even further mobile/remote due to COVID-19) we deployed our Access Proxy model initially to support remote access use cases, as an alternative to VPN. Over the long term, however, we wish to route all application access through Access Proxies, both on-premises as well as off-premises.




Figure 10 – Converged Access Proxy Architecture

Depicted in this figure are two network paths – one from an organization premises network, the other from an untrusted network such as the Internet. Ultimately, both paths converge on an Access Proxy – one network path is not favored or preferred over another. This is the ultimate end-state architecture for application access in ZTA.

9. Service Provider Considerations

In many ways, ZTA for Multi-System Operators (MSOs) and Internet Service Providers (ISPs) is not fundamentally different than other organizations.

9.1. Infrastructure Hardening

For many years, it has been common practice for operators to harden their access networks to Zero Trust security standards, long before the term “Zero Trust” even entered commonplace use. The Data-Over-Cable Service Interface Specifications (DOCSIS) 3.0 Security Specification, which dates to 2006, states its goals as follows:

1. To provide cable modem (CM) users with data privacy across the cable network;
2. To prevent unauthorized users from gaining access to the network’s RF MAC services.

These goals align neatly with ZTA, in that:

1. Assume that the network provides no assurances of privacy. If privacy is desired, it must be established by higher-layer communications protocols.
2. No assumptions are made regarding the identity or authenticity of devices (CMs) on the network. Devices must establish identity through strong authentication before services are provisioned.

DOCSIS then specifies solutions that achieve these goals:

1. For privacy, the Baseline Privacy Interface (BPI) protocol.
2. For authentication, X.509 digital certificates.

Similarly, providers have long hardened the network infrastructure – routers, switches, Cable Modem Termination Systems (CMTSes), etc. – that are used to provision DOCSIS services. Examples include:

- Centralized Authentication, Authorization & Accounting (AAA) frameworks, such as RADIUS, Diameter, TACACS+, LDAP, Kerberos, etc.
- Strong administrative authentication (OTP, public key, etc.)
- Role-based authorization (read only, read/write, CLI command authorization, etc.)
- Encrypted administrative interfaces (HTTPS, SSH, etc.)
- Access Control Lists (ACLs)

In general:

- Service providers should begin regarding their enterprise networks as akin to their access networks, in terms of default trust levels and security guarantees.
- Service providers can and should use their experience hardening access network infrastructure to similarly secure enterprise infrastructure to the same degree.

9.2. Network Visibility

Service providers often deploy significant network monitoring infrastructures, to as to monitor their access networks for performance, faults, and other useful diagnostic information. Simple Network Management Protocol (SNMP) – now in its third version and with significant security enhancements – is the most well-known monitoring protocol.

Existing investments in access network monitoring can and should be deployed to monitor enterprise networks as well, for the reasons discussed in Section 5.1.5.

10. Conclusions

ZTA represents a fundamental shift away from traditional approaches to information and network security. In the past, “internal” networks were safe places, walled off from external threats by “the firewall” and other perimeter controls. Now, however, we are in a state of constant vigilance, where networks do not provide the security assurance they once did.

ZTA recognizes this reality and proposes an alternative approach, one where Users and Devices must establish trust before gaining access to Resources. Trust must be established with confidence (e.g., strong authentication) and network location is not to be used as an authenticator or a credential.

Finally, ZTA is not an experiment – it is a practical, pragmatic approach to security that can be realized with modern IT support structure. If not already on their journey, organizations should strongly consider the merits of ZTA and act accordingly.

Abbreviations

AAA	Authentication, Authorization & Accounting
ACL	Access Control List
BISO	Business Information Security Officer
BPI	Baseline Privacy Interface
CARTA	Continuous Adaptive Risk and Trust Assessment
CDM	Continuous Diagnostic and Mitigation
CISA	Cybersecurity and Infrastructure Security Agency
CM	Cable Modem
CMDB	Configuration Management Database
CMTS	Cable Modem Termination System
DOCSIS	Data-Over-Cable Service Interface Specifications
DNS	Domain Name System
EDR	Endpoint Detection & Response
ERE	Enterprise Risk Engine
FIDO	Fast IDentity Online
IAM	Identity & Access Management
IdP	Identity Provider
ISP	Internet Service Provider
MDM	Mobile Device Management
MFA	Multi-Factor Authentication
MSO	Multi-System Operators
NCSC	National Cyber Security Centre
NIST	National Institute of Standards and Technology
OAuth	Open Authorization
OIDC	OpenID Connect
OMB	Office of Management and Budget
OTP	One-Time Password
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PKI	Public Key Infrastructure
RADIUS	Remote Authentication Dial-In User Service
SAML	Security Assertion Markup Language
SCTE	Society of Cable Telecommunications Engineers
SIEM	Security Information and Event Management
SNMP	Simple Network Management Protocol
SSS	Security Support Structure
TACACS	Terminal Access Controller Access-Control System
UBA	User Behavior Analytics
VPN	Virtual Private Network
ZTA	Zero Trust Architecture

Bibliography & References

“Compartment (ship).” Wikipedia. Available: [https://en.wikipedia.org/wiki/Compartment_\(ship\)](https://en.wikipedia.org/wiki/Compartment_(ship))

Executive Office of the President, Executive Order on Improving the Nation’s Zero Trust Cybersecurity Posture, May 12, 2021. Available: <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>

Cybersecurity and Infrastructure Security Agency (2021) Zero Trust Maturity Model. (Department of Homeland Security, Washington, D. C.), June 2021. Available: <https://www.cisa.gov/zero-trust-maturity-model>

“Data Breach QuickView Report – 2019 Q3 Trends,” RiskBased Security, Richmond, VA, 2019. Available: <https://pages.riskbasedsecurity.com/data-breach-quickview-report-2019-q3-trends>

Data-Over-Cable Service Interface Specifications (DOCSIS) 3.0 – Security Specification, CM-SP-SEC3.0-C01-171207, Cable Television Laboratories, Inc., 2006. Available: <https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=9b3e83b9-daf0-4434-9258-7b6b4e8f2c2e>

Executive Office of the President, Office of Management and Budget (2022, January 26). *M-22-09, Moving the U.S. Government Toward Zero Trust Cybersecurity Principles*. [Online]. Available: <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>

L. Cittadini, B. Spear, B. Beyer and M. Saltonstall, “BeyondCorp Part III: The Access Proxy,” *logins*, vol. 41, no. 4, pp. 28-33, 2016. Available: <https://research.google/pubs/pub45728/>

J. Kindervag, “No More Chewy Centers: Introducing The Zero Trust Model of Information Security,” Forrester Research, Cambridge, MA, Sep. 2010.

S. P. Marsh, “Formalising Trust as a Computational Concept,” Ph.D. dissertation, Department of Computing Science and Mathematics, University of Stirling, Scotland, UK, 1994.

National Cyber Security Centre (2021) Device Security Guidance. (Government Communications Headquarters, London), June 2021. Available: <https://www.ncsc.gov.uk/collection/device-security-guidance/infrastructure/network-architectures>

National Institute of Standards and Technology (2020) Zero Trust Architecture. (Department of Commerce, Washington, D.C.), Special Publication 800-207, August 2020. Available: <https://doi.org/10.6028/NIST.SP.800-207>

S. Potti. “BeyondCorp Enterprise: Introducing a safer era of computing.” Google. Available: <https://cloud.google.com/blog/products/identity-security/introducing-beyondcorp-enterprise>

R. Ward and B. Beyer, "BeyondCorp: A new approach to enterprise security," *login:*, vol. 39, no. 6, pp. 6-11, Dec. 2014. Available: <https://research.google/pubs/pub43231/>

"Zero trust security model." Wikipedia. Available:
https://en.wikipedia.org/wiki/Zero_trust_security_model

ISBN 0-940272-01-6; 0-940272-08-3; 0-940272-10-5; 0-940272-11-3; 0-940272-12-1; 0-940272-14-8; 0-940272-15-6; 0-940272-16-4; 0-940272-18-0; 0-940272-19-9; 0-940272-20-2; 0-940272-21-0; 0-940272-22-9; 0-940272-23-7; 0-940272-24-5; 0-940272-25-3; 0-940272-26-1; 0-940272-27-X; 0-940272-28-8; 0-940272-29-6; 0-940272-32-6; 0-940272-33-4; 0-940272-34-2; 0-940272-35-0; 0-940272-36-9; 0-940272-37-7; 0-940272-38-5; 0-940272-39-3; 0-940272-40-7; 0-940272-41-5; 0-940272-42-3; 0-940272-43-1; 0-940272-44-X; 0-940272-45-8; 0-940272-46-6; 0-940272-47-4; 0-940272-48-2; 0-940272-49-0; 0-940272-50-4; 0-940272-51-2; 0-940272-52-0; 0-940272-53-9; 0-940272-54-7

© 2015 National Cable and Telecommunications Association. All Rights Reserved.