Evolution of the Automated Pre-Vetting Process to Validate Access Network Maps

A technical paper prepared for presentation at SCTE TechExpo24

John Schlack
Distinguished Architect
Comcast Cable
John_Schlack2@comcast.com

Kelin Kruithoff
Sr. Manager, Product Management
Comcast Cable
Kelin_Kruithoff@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

1. Introduction

Accuracy of Hybrid Fiber Coaxial (HFC) plant maps is critical to the operation of the cable network. Designers reference HFC plant maps to plan node capacity upgrades and plant extensions, while the operations group uses them to pinpoint service interruptions. Two years ago, the automated pre-vetting process was introduced. This automated process analyzed the drafted HFC network to identify methods to enhance the drafted map data. This process occurs before designing plant extensions and capacity upgrades to ensure accuracy of HFC plant maps and provide insight for node splits and segmentation.

This paper examines the usage of the original automated pre-vetting process[1], feedback from users, features that provided the most impact, and how the automated pre-vetting process has evolved and added more capabilities. The evolution takes an in-depth look at differences in the automated pre-vetting process for the integrated cable modem termination system (iCMTS) and virtualized cable modem termination system (vCMTS) and the benefits of including telemetry data. Details of the output of the enhanced automated pre-vetting process are included such as layered plant maps, device plots, and utilization data.

2. Review of Original Pre-Vet Automation Capabilities

When planning a plant extension or capacity upgrade for an existing HFC node, a pre-vetting step was manually performed to characterize the currently drafted plant maps and provide the designer with valuable information regarding node splits and node segmentation. This manual process requires aggregation and analysis of data from many sources. Automation of the pre-vet process reduces the time allocated to the pre-design analysis and provides consistent data output for the HFC nodes.

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;
- United States Postal Service (USPS) address database;
- Subscriber account database; and
- Device telemetry.

Automated pre-vet data collection, correlation, and analysis, generates a set of data for use by the designers including:

- Graph representation of all elements in the radio frequency (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 and amplifier
- 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

Visualization of the pre-vet analysis results is overlaid onto the drafted access network map. The pre-vet application presents the analysis 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 for analysis.

Homes passed address locations are color coded based according to the associated bus leg. Service location exceptions are highlighted with an icon. 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. The figure below shows a mock representation of the map view.

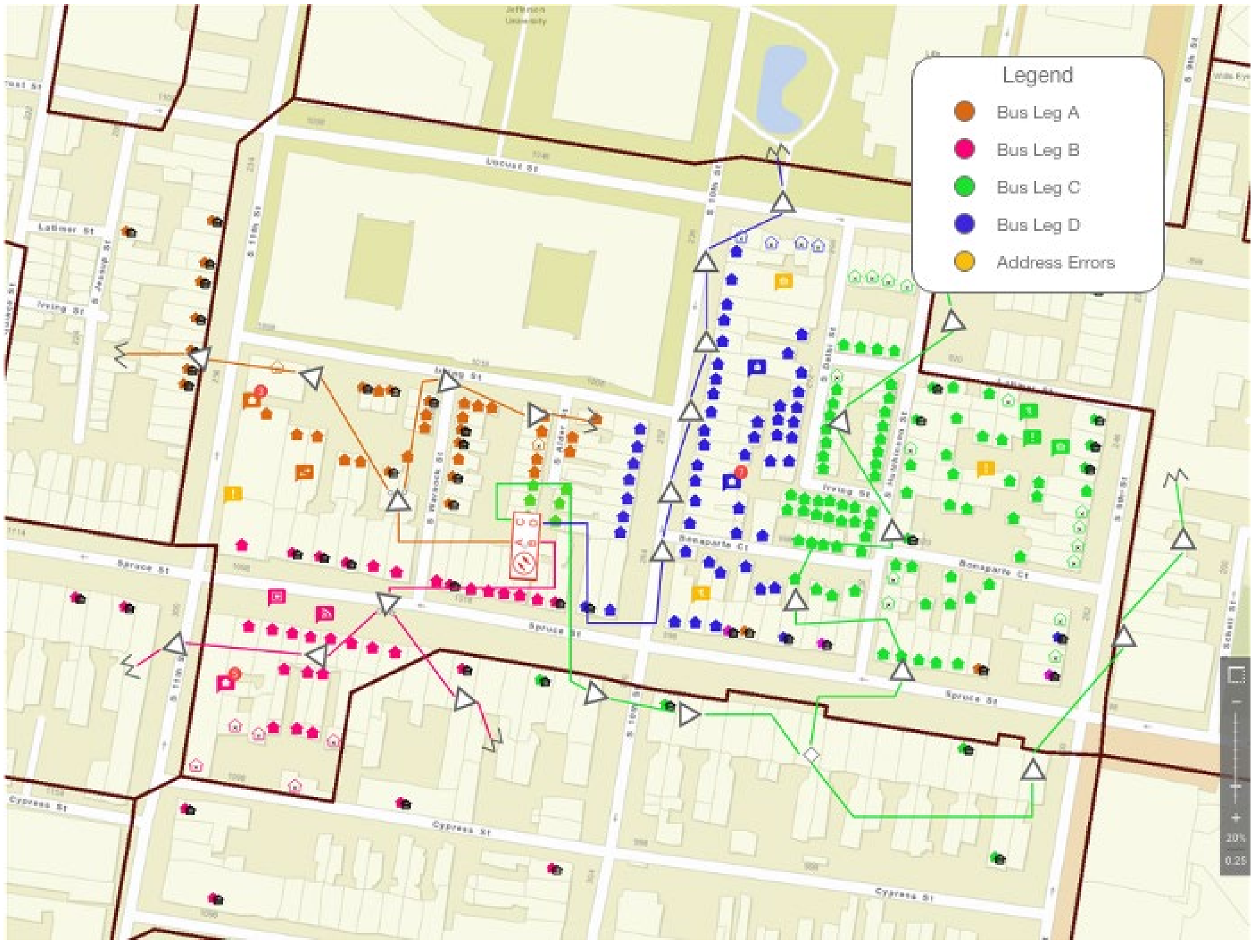

Figure 1: Map of Plant with Service Locations Color Coded by Bus Leg

The system aggregates counts of both homes passed and of high-speed data (HSD) capable devices, such as cable modems (CM), set-top boxes (STB), and media terminal adapters (MTA) at each amplifier. These counts indicate values between amplifiers and from each amplifier to the end of line. This information aids the designer in determining where there is pressure in the HFC plant and an opportunity for capacity upgrade.

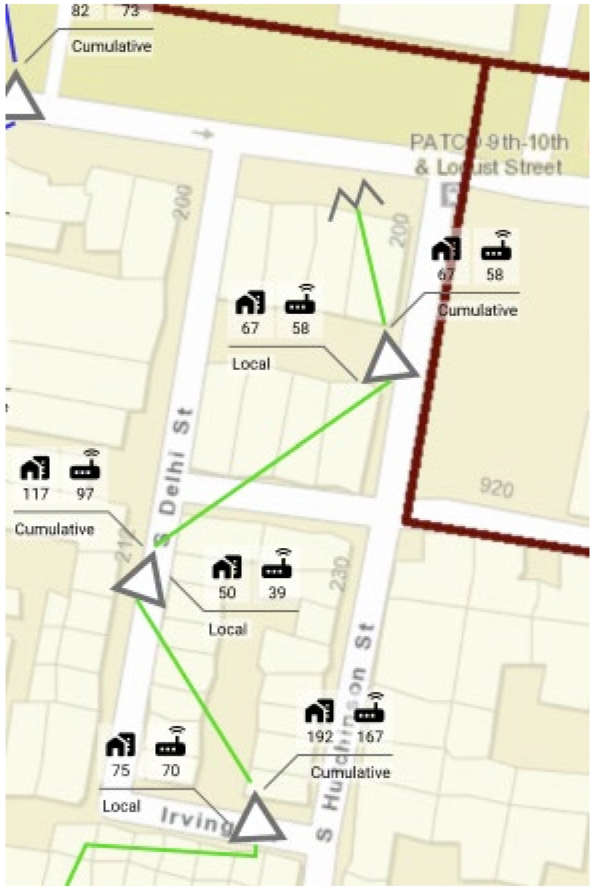

Figure 2: Device and Homes Passed Count at Amplifier

In addition to the map view, the user interface provides a summary table for each bus leg of the node housing.

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 3: Bus Leg Summary Table

3. Feedback and Enhancements

3.1. Opportunities

The original automated pre-vet process was well received but did trigger several enhancements to further drive efficiencies and value for the designer.

3.1.1. Device Plots

The original automated pre-vet access network maps displayed the homes passed locations in relation to the RF network and it displayed the device counts rolled up to the amplifiers. This proved more challenging to interpret as compared with plotting the device locations on the map. The device locations make it easy to see subscriber density, which contributes to balancing the load during a node split.

The device plot also provides a method to identify where network was drafted but not constructed. Drafters sometimes publish the network upgrade before construction since existing design tools did not provide visibility of "in flight" designs to all designers. If that construction job is canceled, the drafted network needs to be reverted to its original design. Failure to revert the network when a job cancelation occurs means that the drafted network includes plant extensions that were never constructed. A device plot that shows there are no subscribers for a given plant extension is an indicator that the plant may not have been constructed.

3.1.2. Utilization and Consumption

The original automated pre-vet solution provided aggregated counts of homes passed and high-speed data (HSD) devices by bus leg and at each amplifier. The purpose was to use these counts to plan capacity upgrades. Although these values provide a rough framework for the capacity upgrade, actual upstream and downstream peak utilization details were required for determining exactly where and how to best perform the capacity upgrade.

The MSO typically has thresholds of peak utilization that will require a node segmentation or node split. Current and historic peak utilization data indicates which upstream and downstream paths are in an alarm state, which ones are close and may require some pro-active upgrades, and which ones have sufficient capacity for the foreseeable future.

In some cases, the utilization data is not sufficient to determine how to segment or split an existing node. The upstream and downstream communications channels may be connected in such a way that multiple channels are aggregated together in the peak utilization data. Consumption data provides a method to estimate each of those channels' contribution to the aggregate bandwidth. The consumption data indicates the amount of data consumed by each of the HSD devices, which is then aggregated by node segment.

3.1.3. Layered DWG Maps

The original automated pre-vet application provided a method to view a layered network map through the application itself. However, the original pre-vetting application exported portable document format (PDF) versions of the map and analysis. The PDF file only included a single layer, where individual classes of devices could not be hidden. Due to the density of equipment on the maps, a layer-based AutoCAD drawing file (DWG) was requested. The DWG file format supports the toggling of layers to show or hide sets of information.

Each class of equipment is assigned its own layer. For example, aerial support, underground support, coax cables, fiber cables, amplifiers, RF Taps, etc. all can be individually toggled.

3.1.4. Address List

The original automated pre-vet process identified several address-related challenges including incorrect node segment, no connection to the RF network, address not serviceable, and service location outside the node boundary. These issues were displayed on the access network map. The user community requested a list-based format so that this list could be distributed to multiple users to correct issues in parallel.

A new address list export solves the issue by outputting the addresses in a comma separated value (CSV) file. The address list identifies issues related to each address. The automated pre-vet process will further enhance the address information by providing a geographic distance of the address to the node housing. The distance is calculated based on the position output by a geocoder, which provides a method to verify the accuracy of the address information.

3.2. Feedback

Below is some sample feedback provided to indicate the value of the capabilities introduced with the enhanced automated pre-vets. This feedback was provided after delivery of the enhancements discussed in this paper.

“The composition of the results grid, the automated DWG export, and the pre-vet package PDF is [equivalent to] 2 hours in manual labor.” - Aaron Reck (Engineer, Planning and Design)

“Found a 2X2 node that was being represented as a 4X4 [node]. We were able to leverage the [pre-vet] information to better learn where the discrepancy is.” - Don Simmons (Manager, Capacity Operations)

3.3. Usage

When the automated pre-vet capability was launched, the initial usage was limited. Reasons for the limited usage include an initial trial and feedback period, training time, and introduction of requested capabilities. As the additional capabilities listed above were delivered, acceptance and usage greatly increased. This usage and timeline are shown in the figure below.

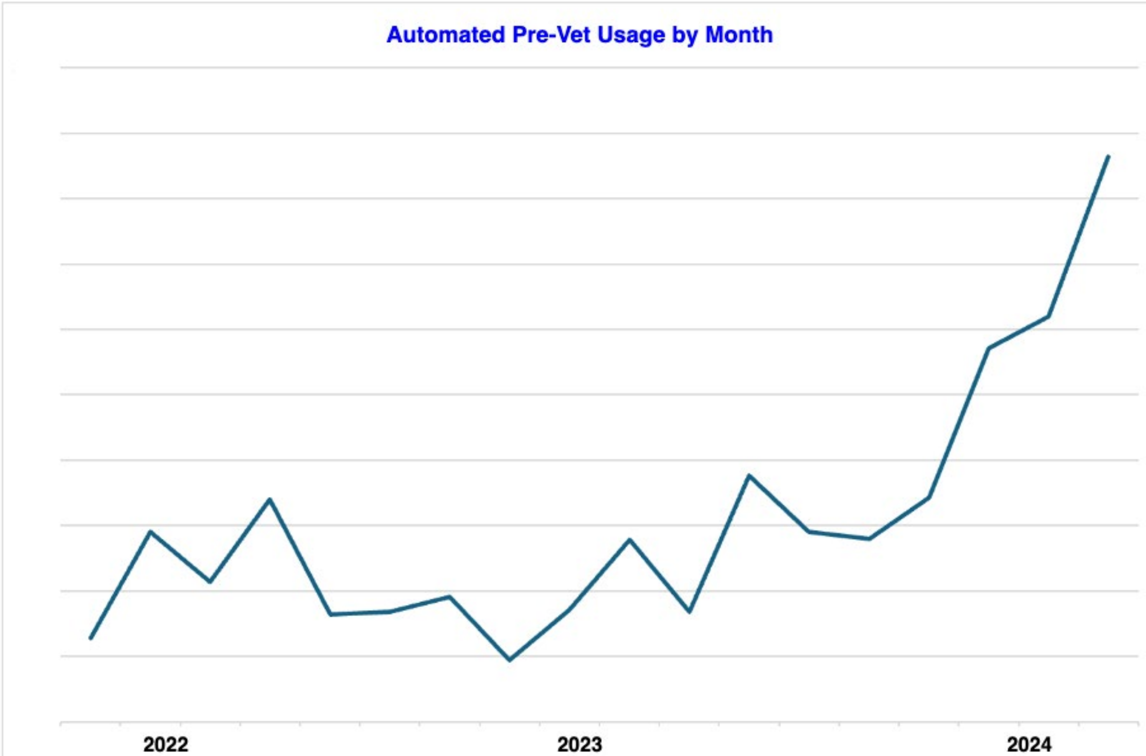

Figure 4: Automated Pre-Vet Usage

The initial usage was limited since it represents the launch feature set. Testing, trials, and feedback proceeded into 2023. Enhancements were iteratively delivered, which greatly improved capabilities and thus usage rates.

4. Evolution Details

The sections below provide more details about the evolution of the automated pre-vetting application. These sections detail the data collection, processing, and output for the new capabilities.

4.1. Device Plot

The device plot diagram provides a mechanism to visualize the location of the HSD devices for a given node housing. Its purpose is to identify missing network elements that may not have been drafted either with the original design, with the "as built", or due to an eventual plant extension. It also provides a method to identify plant extensions that were drafted but never constructed. For example, all service locations must be within 300 feet of an RF tap. If the device plot indicates that there are devices more than 300 feet from an RF tap, then drafted network maps may be missing details of a plant extension.

The automated pre-vet process already collects device information as part of its plot of aggregated device counts by amplifier. This existing device information is used to create the device plot. The device plot filters the service locations to remove addresses without an HSD device. The device plot places an icon at each address location with an HSD device and color codes the icon by upstream node segment so a designer can easily see the distribution of devices by node segment. Further, if a device's upstream node segment differs from the upstream node segment of billing or telemetry, the device icon has a red outer ring indicating that there is a node segment mismatch (also discussed in section **4.5.1**). Drafters can use this information to make corrections in the design and/or billing data sources.

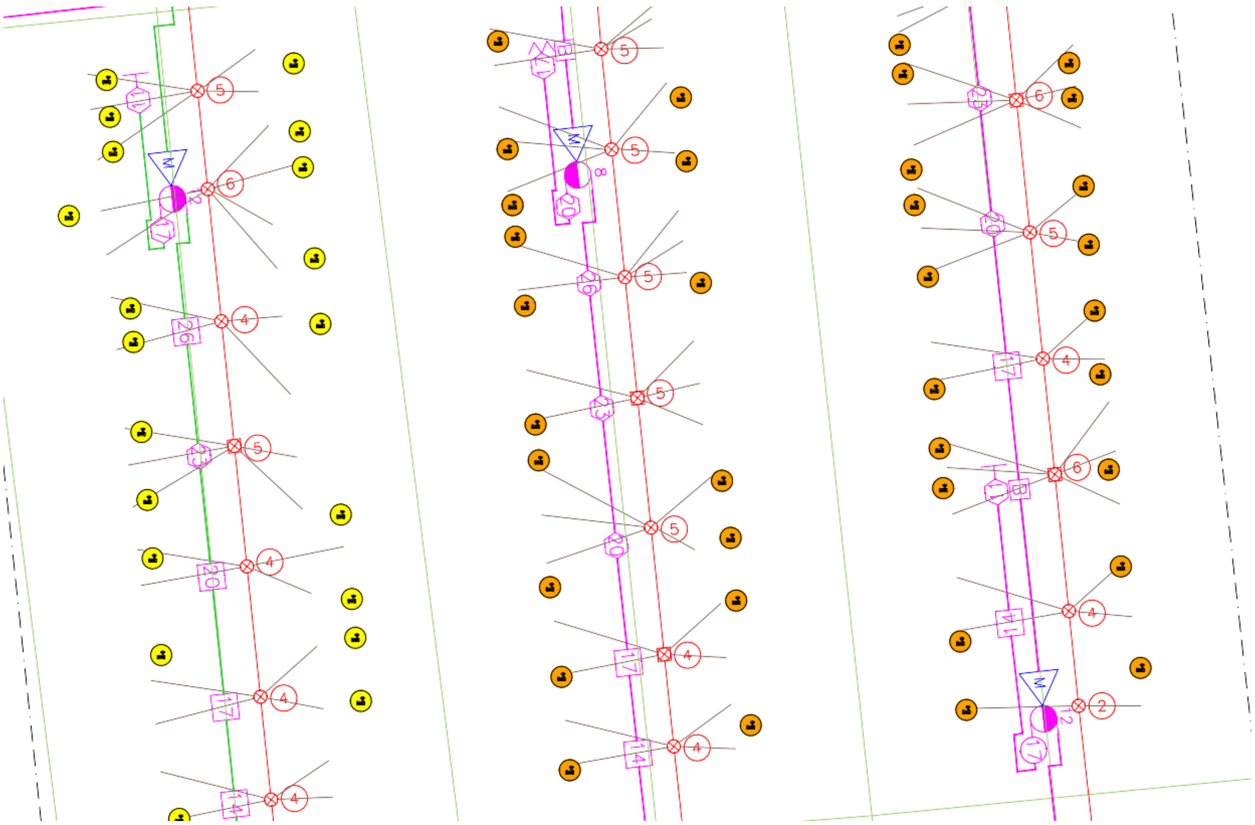

Figure 5: Device Plot with Device Locations Color Coded by Bus Leg

4.2. Peak Utilization History

The peak bandwidth utilization feature provides a historic view of upstream and downstream HSD bandwidth utilization. Utilization data collection and display differs based on whether the CMTS is an integrated CMTS (iCMTS) or a virtualized CMTS (vCMTS).

4.2.1. iCMTS Utilization

Utilization data retrieved from the iCMTS platform aggregates at the medium access control (MAC) domain level. The MAC domain may be a combination of multiple downstream and upstream channels. For this discussion, there will be a single downstream service group and one or more upstream bonding groups per MAC domain. The service group may have multiple channels, such as data over cable service interface specifications (DOCSIS®) 3.0 and DOCSIS 3.1® channels. Note that this section focuses on how to extract and examine the peak utilization details for a particular node segment. The task of extracting utilization data from the iCMTS is outside the scope of this document.

There are challenges related to processing the iCMTS utilization data including:

- Multiple node segments may correspond to the same bonding group
- Node segments from different node housings may appear in the same bonding group
- MAC domains may contain one or more bonding groups

- Node segment to bonding group associations may change over time due to combining changes in the headend

For example, a 2x2 analog node housing may have 2 downstream service groups and 2 upstream node segments. The node segments are named NH01A and NH01C. There will be 2 MAC domains, one per downstream service group. The initial configuration for week ending April 6 has two bonding groups in MAC domain 7/0/7. The assignments may change over time. For week ending April 13, the node segments were moved to a different MAC domain and node segment NH01A is no longer combined.

Table 1: Example iCMTS MAC Domain to Bonding Group Association

Week Ending	iCMTS	MAC Domain	Service Group	Bonding Group(s)	Node Segment(s)
2024/04/06	CMTS_1	7/0/7	40	14 15	NH01A NH02A
2024/04/06	CMTS_1	7/0/5	38	11	NH01C
2024/04/13	CMTS_1	8/0/3	40	14	NH01A
2024/04/13	CMTS_1	8/0/1	38	11	NH01C

Peak utilization data is tracked at the MAC domain level. Utilization data is available for the downstream DOCSIS 3.0 and 3.1 channels and each of the individual bonding group(s). Example data is shown below.

Table 2: Example iCMTS Peak Utilization Data

Week Ending	iCMTS	MAC Domain	Downstream DOCSIS 3.0	Downstream DOCSIS 3.1	Upstream
2024/04/06	CMTS_1	7/0/7	29.3%	48.3%	73.6% (BG 14) 59.2% (BG 15)
2024/04/06	CMTS_1	7/0/5	36.1%	52.6%	60.9% (BG 11)

The automated pre-vet process has a list of all upstream node segments for the node housing. As seen above, the association to MAC domains may change over time. When building the chart of peak utilization data, the node segment to MAC domain association for target date range is joined with the peak utilization data for the MAC domain over the same date range, to produce the set of node segment to MAC domain to peak utilization values for each date in the date range.

Since the utilization data is aggregated by MAC domain, there could be bonding groups in a MAC domain for node segments from other node housings. The query results will include the utilization data for those additional bonding groups. For this case, the downstream service group is servicing devices from multiple node segments. Determining the contributions of devices from different node segments to the overall downstream utilization is explained further in the "Consumption Data" section below.

The figure below shows an example peak utilization for a MAC domain associated with a specific node segment. The historic peak utilization data shows both upstream utilization for the node segment along with downstream utilization for the DOCSIS 3.0 and DOCSIS 3.1 channels.

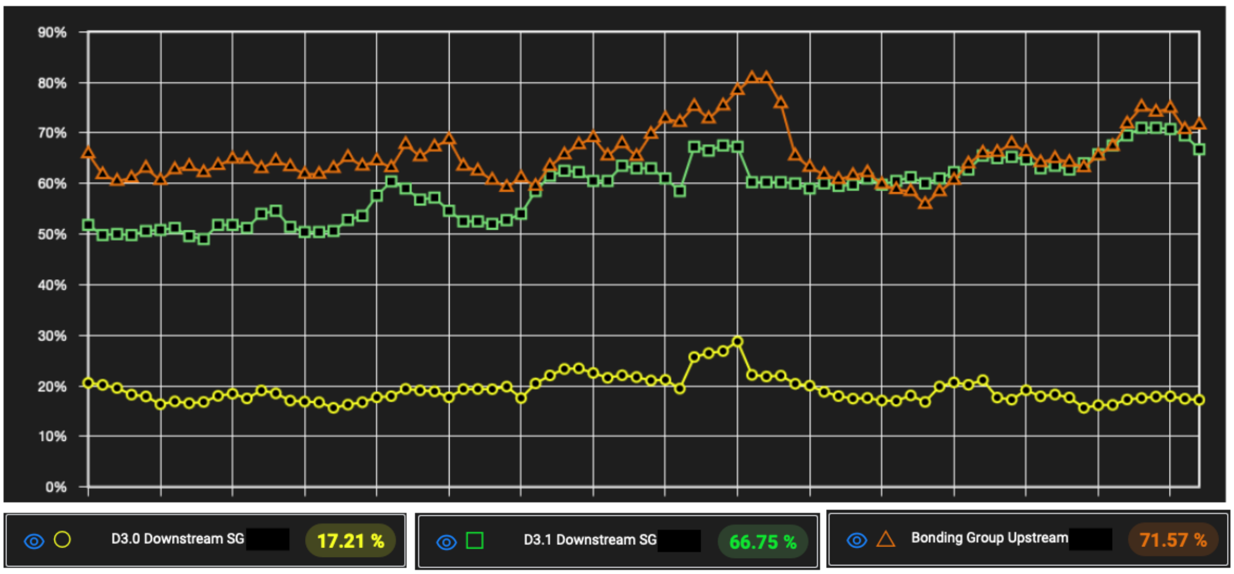

Figure 6: Peak Utilization History

Summary data is provided in tabular format. The example below shows multiple bonding groups per MAC domain. This data has been anonymized with the iCMTS, MAC domain, service group, and bonding group replaced with example values.

MAC Domain Info													
FQDN	Mac Domain	Service Group(s)	Bonding Group(s)	Node Segment(s)	Consumption Week Endings	DS (GB) last month	DS (%) of SG	US (GB) last month	US (%) of BG	Utilization Week Ending	D3.0 DS Utilization	D3.0 US Utilization	D3.1 DS Utilization
xxx.xxx.xxx.xxx.comcast.net	MD01	SG1	14 15	NS01A (BG 15)	07/06/2024 07/13/2024 07/20/2024 07/27/2024	21,792	64.25%	1,670	100%	07/27/2024	42.55	42.19 70.59	67.93
xxx.xxx.xxx.xxx.comcast.net	MD02	SG2	11	NS01B	07/06/2024 07/13/2024 07/20/2024 07/27/2024	18,527	100%	1,350	100%	07/27/2024	30	68.43	55.67
xxx.xxx.xxx.xxx.comcast.net	MD03	SG3	6 7	NS01C (BG 7)	07/06/2024 07/13/2024 07/20/2024 07/27/2024	20,403	76.42%	1,190	100%	07/27/2024	33.4	24.38 61.82	77.01

Figure 7: iCMTS Utilization Summary Chart

4.2.2. vCMTS Utilization

The virtualized CMTS platform moves the PHY device, which handles RF waveform generation, out of the core CMTS. This allows the software functions of the CMTS to be virtualized. The remote PHY device (RPD) is typically placed in the node housing, thus moving RF waveform generation toward the edge of the network.

Each digital node housing can contain up to 2 RPDs. Each RPD supports a single downstream channel (service group) and two upstream channels (node segments). Thus, a node housing with a single RPD is typically 1x2 segmentation while a node housing with 2 RPDs is typically 2x4 segmentation. Other variations can occur, but there cannot be more than 2 downstream and 4 upstream channels in a digital node housing.

A deterministic association exists between the node housing, RPDs, service groups, and node segment assignment. The figure below shows two example node housings, one with a single RPD and 1x2 segmentation and one with two RPDs and a 2x4 segmentation. These associations are important for the peak utilization reporting.

Fibers	Direction	Optical Node Ports	Direction	Optical Node Ports	Direction	RF Node ports
Fiber 1	=>	SFP1-DS	=>	DS_SG_1 Forward Service Group 1	=>	Out A
					=>	Out B
					=>	Out C
					=>	Out D
Fiber 2	<=	SFP1-US	<=	US_NS_A Node Segment A	<=	Out A
					<=	Out B
				US_NS_C Node Segment C	<=	Out C
					<=	Out D

Fibers	Direction	Optical Node Ports	Direction	Optical Node Ports	Direction	RF Node ports
Fiber 1	=>	SFP1-DS	=>	DS_SG_1 Forward Service Group 1	=>	Out A
					=>	Out B
Fiber 2	<=	SFP1-US	<=	US_NS_A Node Segment A	<=	Out A
				US_NS_B Node Segment B	<=	Out B
Fiber 3	=>	SPF2-DS	<=	DS_SG_2 Forward Service Group 1	=>	Out C
					=>	Out D
Fiber 4	<=	SFP2-US	<=	US_NS_C Node Segment C	<=	Out C
				US_NS_D Node Segment D	<=	Out D

Figure 8: vCMTS Port Mapping from Fiber to Housing to RPD to RF Ports

The vCMTS reports peak utilization data by RPD and RPD port. Since there is a deterministic association between RPD, service group, and node segments, the peak utilization graphs can be directly created from the peak utilization data.

The graph of peak utilization history looks exactly the same as the iCMTS peak utilization history. Refer to the "iCMTS Utilization" section above.

The figure below shows the vCMTS peak utilization summary data in tabular format. This is very similar to the iCMTS summary data, except that instead of MAC domains and bonding groups, data is displayed by RPD and node segment. The "VCVS DS Port 0" value is the downstream utilization. The "VCVS US Port 0" and "VCVS US Port 1" values are the upstream utilization values for the two node segments associated with this RPD.

MAC Domain Info							
FQDN	VCVS	RPD	Node Segment(s)/ US VCVS Port	Utilization Week Ending	VCVS DS Port 0 Utilization	VCVS US Port 0 Utilization	VCVS US Port 1 Utilization
xxx.xxx.xxx.xxx.comcast.net	99:1	ABCDD00103	ABCDD0010C/0	07/27/2024	14.65	35.67	8.33
	99:1		ABCDD0010D/1				
xxx.xxx.xxx.xxx.comcast.net	99:2	ABCDD00101	ABCDD0010A/0	07/27/2024	34.33	75.94	33.99
	99:2		ABCDD0010B/1				

Figure 9: vCMTS Utilization Summary Chart

4.3. Consumption Data

The peak utilization data provides accurate details for the upstream bandwidth usage by node segment. However, for downstream utilization, there are many cases where the contribution by node segment cannot be directly determined. For the iCMTS case, if there are multiple bonding groups per MAC domain, one cannot determine the contribution of the devices on each bonding group to the downstream utilization. Similarly, for vCMTS, when there are multiple node segments per RPD, one cannot determine the contribution of the devices on each node segment to the downstream utilization.

Consumption data provides a method to estimate the contribution of the devices from each node segment to the downstream bandwidth utilization. Consumption tracks the upstream and downstream bandwidth usage on a per device basis. This data is likely already collected for data cap calculations. However, unlike data caps, the consumption data is only collected at the time of peak utilization.

The consumption total by node segment is calculated by summing the individual upstream and downstream consumption data from each HSD device on that node segment as reported by the telemetry data. The result will be a total data measured in gigabytes of upstream and downstream data usage.

For a given service group, the contribution from an individual node segment is determined by dividing the downstream consumption data of that node segment by the sum of all downstream consumption from the node segments in that MAC domain or RPD.

An example of the contribution calculation is shown in the table below. Two bonding groups (14 and 15) are in the same MAC domain with a single downstream service group. The table shows a sample set of devices split across two node segments in the MAC domain along with their respective consumption data values.

Table 3: Example Consumption Data Calculation

Device ID	Upstream Usage	Downstream Usage	Node Segment	Downstream Total	Downstream % Contribution
00C001010101	0.9 GB	1.3 GB	NS01A (BG 15)	8.6 GB	42.6% (8.6 / 20.2)
00C001010102	0.1 GB	2.1 GB			
00C001010103	0.2 GB	0.7 GB			
00C001010104	1.4 GB	3.2 GB			
00D001010105	0.2 GB	0.4 GB	NS01C (BG 14)	11.6 GB	57.4% (11.6 / 20.2)
00D001010106	0.8 GB	7.0 GB			
00D001010107	1.1 GB	4.2 GB			

The contribution value appears in the summary table as shown in the "iCMTS Utilization" section.

4.4. DWG Layered Map

The original automated pre-vet provided a PDF based output mechanism for the drafted maps. The PDF map contained a single layer. However, the design community preferred a layered AutoCAD based DWG file format. Each class of equipment on the map appears in its own layer, allowing designers to toggle the layers independently.

From an implementation standpoint, the main challenge about moving from web-based map generation or PDF based map generation to a DWG format is the format of the images and icons for the map. Images and icons used for web and PDF maps use the scalable vector graphics (SVG) file format. However, the SVG format was not supported by the library used for DWG creation. Instead, computer-aided design (CAD) blocks were used for DWG creation. In most cases, the CAD blocks already existed since these were originally used to create the SVG images. However, new SVG images that never existed in a CAD application had to be re-drafted in a CAD application and saved as a CAD block.

The image below shows an example DWG output file along with its layers menu.

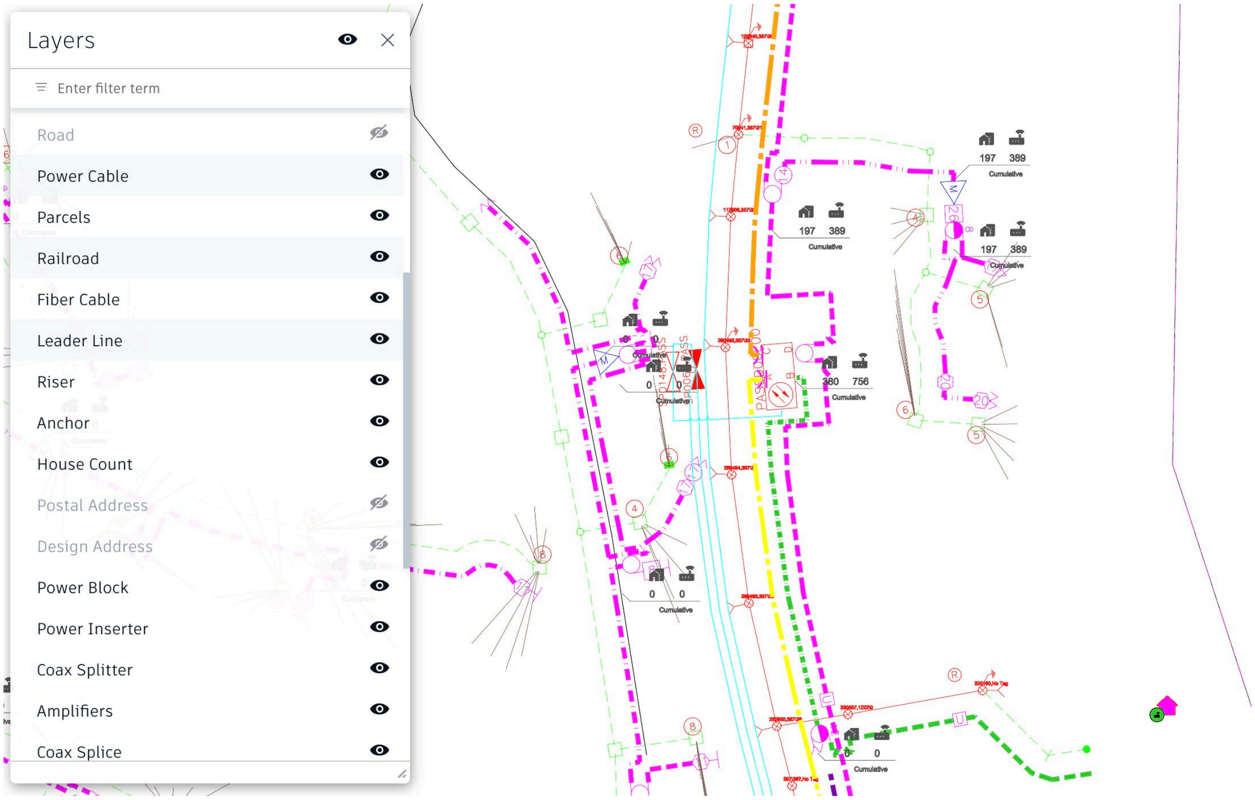

Figure 10: Layered DWG Export File

4.5. Address List

The original automated pre-vet included overlaying service location data of homes passed on the access network map. Automated pre-vet now includes an export of the service location addresses to a comma separated value (CSV) file. This format simplifies the address correction process since portions of the address list can be easily distributed to multiple teams.

In addition to specifying the USPS addresses of the service locations, the export provides some enhancements to isolate mismatches in node segment details and potentially identify data entry errors on the service address.

4.5.1. Node Segment Mismatches

Maintaining consistent information on the association of service locations to node segments is critical when performing a capacity upgrade. The node segment can be used to determine the bandwidth utilization and consumption data, which are used in node segmentation and node split analysis.

When the designer drafts the access network for an RF node, the designer specifies the node segmentation that provides sufficient bandwidth within specific cost constraints. The node segment assigned to the service location from the designer is called the design node segment.

Between the design and construction process, the service locations are entered into the billing system along with the node segment. The node segment in the billing system is called the billing node segment.

Note that entry of the service location into the billing system requires validation and serviceability checks. The USPS address value may be updated or changed during this process. Therefore, the actual USPS addresses in the drafted network and billing system may differ. The automated pre-vet imports the USPS addresses from the billing system and links them to RF taps by analyzing the RF network graph.[1]

After construction is completed, the construction team provides the "as built" to indicate any construction changes to the design. These changes may impact the service locations assigned to each node segment. Design and billing need to update their systems appropriately, which could lead to mismatches.

Finally, when subscriber devices communicate with the CMTS, the network determines the actual node segment from the communications. This is called the telemetry node segment.

Ideally, the design node segment, billing node segment, and telemetry node segment match. However, mismatches may occur due to data entry error or construction changes that were not captured. Reporting all three node segment values for each service location in the address list helps identify mismatched values.

Note that small numbers of devices may indicate a different telemetry-based node segment than other devices on the same bus leg. This is an indication that the device was connected to the network at a different location than the service location of the subscriber.

4.5.2. Service Address Validation

Data entry errors may occur during manual entry of service location addresses in the design and drafting system or into the billing system. In cases of new construction, a USPS address may not be available when creating the design in the drafting platform. The result is that an incorrect USPS address may be entered as a design address or billing address.

To isolate errors in the captured addresses, the address text is passed through a geocoder service to obtain the latitude and longitude of the USPS address. The distance between the geo-located address and the drafted node housing location is calculated using the Haversine formula.[2] If the calculated distance between the geolocation and node housing location exceeds a configurable threshold, then it is likely that the address entered for the service location is incorrect.

The Haversine formula is shown below:

```
dlon = lon2 - lon1
dlat = lat2 - lat1
a = (sin(dlat/2))^2 + cos(lat1) * cos(lat2) * (sin(dlon/2))^2
c = 2 * arcsin(min(1,sqrt(a)))
dist = R * c
where R is the radius of the Earth in feet
Latitude and longitude are in radians. Multiply decimal degrees by pi/180.
```

Since the radius of the Earth is supplied in feet, the output distance of the function is measured in feet.

Here is an example calculation:

Table 4: Example Haversine Calculation

	Latitude (Deg)	Longitude (Deg)	Latitude (Rad)	Longitude (Rad)
Point 1	41.65121	-86.17239	0.726950752	-1.50399193
Point 2	41.65671	-86.17221	0.727046745	-1.503988788
Earth Radius	3958.8	miles	20902464	feet
Dlon	3.14159E-06			
Dlat	9.59931E-05			
A	2.30483E-09			
C	9.60174E-05			
Dist	2007.00	feet		

4.5.3. Example CSV File

The table below shows an example of the address list CSV file. Note that this data is not actual homes passed or subscriber data.

Table 5: Example Address List CSV Export

Street Address	City	State	Zip Code	Distance to Node (ft)	Design Node Segment	Billing Node Segment	Telemetry Node Segment	Notes
101 Main Street	Philadelphia	PA	19101	460	PAFMD0110A	PAFMD0110A	PAFMD0110A	
103 Main Street	Philadelphia	PA	19101	485	PAFMD0110A	APFMD0110A	PAFMD0110A	
105 Main Street	Philadelphia	PA	19101	515	PAFMD0110A	PAFMD0110A	PAFMD0110A	
1107 Main Street	Philadelphia	PA	19101	12,250	PAFMD0110A	PAFMD0110A	PAFMD0110A	
122 Second Street, Unit 1	Philadelphia	PA	19101	720	PAFMD0110A	PAFMD0110A	PAFMD0110A	Outside node boundary
122 Second Street, Unit 2	Philadelphia	PA	19101	720	PAFMD0110A	PAFMD0110A	PAFMD0110A	Outside node boundary
122 Second Street, Unit 3	Philadelphia	PA	19101	720	PAFMD0110A	PAFMD0110A	PAFMD0110A	Outside node boundary

Highlighted items illustrate elements that require investigation. In one case, the distance of the service location to the node is over 12,000 feet, which may be unexpected. In another case, the billing node segment does not match the design node segment and telemetry node segment counterparts. Finally, some of the service locations were drafted outside the node boundary.

5. Conclusion

The manual pre-vetting process, which characterizes the currently drafted RF plant maps for a node housing and incorporates inputs from billing and telemetry data, required 2-8 hours to complete based on the complexity and size of the RF plant. The original automated pre-vet process succeeded in characterizing the network and providing some efficiencies, which reduced a portion of the manual labor.

Based on testing and feedback, subsequent additions were performed to enhance the automated pre-vetting process. These enhancements included:

- providing device plots to identify missing RF equipment and RF equipment that was not constructed,
- displaying a history of peak bandwidth utilization and integrating consumption data to provide a detailed view of pressure points in the upstream and downstream network,
- exporting a layered AutoCAD DWG file for offline viewing of the RF plant,
- and exporting the address list to a CSV file and indicating potential node segment mismatches and address data entry inconsistencies.

The automated pre-vetting application has evolved to a point where all major features have been developed and integrated into the tool. The user community continues to highlight opportunities for enhancing existing features and adding new features that expand the coverage of automated pre-vet to further increase time savings. The overall sentiment from the user community is that when pre-vet automation is run on a node boundary that has no drafting errors, the dual output of the pre-vet package and layered DWG map export result in considerable time savings for the pre-vetting teams as well as a standardized data gathering process.

Note that the automation in the pre-vetting process is the foundation for automation of the design and drafting process. The ingestion of RF plant information, integration of the billing and telemetry data sources, characterization of the network, and representation of the RF network using graphs, are all critical elements for design and drafting automation.

Abbreviations

CAD	computer-aided design
CM	cable modem
CMTS	cable modem termination system
CSV	comma separated value
DOCSIS	data over cable service interface specifications
GB	gigabyte
HFC	hybrid fiber coaxial
HSD	high-speed data
iCMTS	integrated CMTS
MAC	medium access control
PDF	portable document format
RF	radio frequency
RPD	remote PHY device
SCTE	Society of Cable Telecommunications Engineers
STB	set-top box
SVG	scalable vector graphics
USPS	United States Postal Service
vCMTS	virtualized CMTS

Definitions

bonding group	For the iCMTS platform, this is the set upstream channels grouped together so that communication packets can be distributed across any of them for communications.
DWG file	A layer-based AutoCAD drawing file
MAC domain	A MAC domain is a grouping of downstream and upstream ports in the iCMTS platform.
node segment	A node segment is the upstream communication for a set of subscribers. This value is typically a specific CMTS port for an iCMTS or RPD port in the case of a vCMTS. A node segment may be comprised of one or more bus legs from the node housing.
service group	This is the set of downstream channels in a MAC domain for the iCMTS platform.

Bibliography & References

[1] Schlack, J. & Hayashi, R. (2022) Validating Access Network Maps for Plant Extensions and Capacity Upgrades. *Society of Cable Telecommunications Engineers (SCTE) Expo 22*. https://account.scte.org/documents/5776/FTF22_TOOLS01_Schlack_3712.pdf

[2] Sinnott, R.W. (1984) Virtues of the Haversine. *Sky and Telescope*, vol. 68, no. 2, p. 159.