Bandwidth Planning During the Age of CoVID

A Technical Paper prepared for SCTE by

Keith Alan Rothschild, Ph.D.
Senior Principal
Cox Communications
6305-B Peachtree Dunwoody Rd, Atlanta GA 30328
+1 404 269 8122
kar@cox.com

Andrew Robinson, Network Planning Engineer III, Cox Communications

Derek Bantug, Network Engineer II, Cox Communications

Kris McNally, Senior Technical Project Manager III, Cox Communications

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

1. Introduction

What initially started in early March 2020 as a multi-day trial to determine if the company could rapidly move to support a remote workforce, as part of emergency preparedness, quickly pivoted into a remote workforce for roughly 18-months in response to the pandemic. This was not only a radical change in the way we work, but also in the way our customers use our products. Disruptions in supply-chain, construction, and customer interactions added challenges to what would have otherwise required a tremendous-level of effort to ensure the customer experience remained as consistent and positive as possible.

Fortunately, we have a standing process for managing our bandwidth on an ongoing basis. This process has four components: (1) strategize, (2) model, (3) plan, and (4) deploy. This paper reviews each of these four components in general as well as how they enabled us to rapidly respond to the increased demand driven by the pandemic.

2. Strategize

The last decade has seen an evolution from planning for the co-existence of multiple products with discrete bandwidth requirements on the HFC network to all products utilizing IP transport on multiple access network technologies. Telephony evolved from circuit-switched telephony to IP/packet-switched. Video evolved first from analog to digital, then from MPEG-2 to MPEG-4, and dedicated QAM-based to shared IP-based delivery. This paper will focus on meeting the bandwidth needs resulting from the converged delivery of products over IP.

The general formula we use to determine the amount of bandwidth required (C, capacity) as being at least the sum of the Peak Traffic (we use P_{95}) and Max Tier (T_{Max}):

$$C \geq P_{95} + T_{Max}$$

The Peak Traffic is sometimes looked at as the product of the number of subscribers and the per-subscriber contribution to peak traffic, however, we use the measured value that the node is at or below for 95% of the time (P_{95}). When additional capacity is available, there are peaks in traffic that can be handled by the network without loss of data and which typically should not impact the customer experience, so we utilize this measure to prevent those peaks from driving unneeded capacity.

The maximum tier is typically the maximum advertised speed plus some overhead to ensure a positive quality of experience. Ulm & Cloonan (2017) list these as separate variables, and many argue that the overprovisioned we include in the tier is not the same as the quality of experience variable they intend. Cloonan (2014) refers to the QOE experience as K. Rather than treating these separately, we typically overprovision the customer by ~10% and use the overprovisioned value as T_{Max} in our calculations.

Capacity calculations for upstream and downstream are treated separately. For example, if the maximum tier on a node was 300/30 (300 Mbps downstream, 30 Mbps upstream), and these were overprovisioned by 10%, then we would use T_{Max-D} of 330 and T_{Max-U} of 33. In a node with 32 SC-QAM (37.5 Mbps) channels on the downstream and 4x6.4 MHz QAM (26.88 Mbps) channels on the upstream, then we can figure out the maximum amount of traffic that the nodes can bear before a node action is required as follows:

$$C \geq P_{95} + T_{Max} \xrightarrow{rearranged} P_{95} \leq C - T_{Max}$$

Downstream

$$P_{95} \leq C - T_{Max}$$

$$P_{95} \leq (32 \times 37.5) - 330$$

$$P_{95} \leq 1200 - 330 \leq 870$$

Upstream

$$P_{95} \leq C - T_{Max}$$

$$P_{95} \leq (4 \times 26.88) - 33$$

$$P_{95} \leq 107.52 - 33 \leq 74.52$$

For QAM, we use 7/8 of the raw throughput to estimate the available payload capacity. Although the traffic limits are based on these formulas, operations typically use a "70%" rule, which has worked as these numbers are close to 70% of the capacity, though in the future, this will need to be addressed as it is likely that this will not remain close to 70%. Today, nodes are considered targets for node-actions at 70%, congested at 80%, and heavily congested at 90%, with the latter two corresponding to increases in trouble-calls.

Approaching the target maximum P_{95} value in either the upstream or downstream will typically drive a node-action. In practice, we use a model to project when a node will approach the capacity limits to plan node actions much further in advance, and closely monitor actual data to make modifications to the plan based on changes that could not be predicted.

Downstream traffic is estimated to grow at a compound annual growth rate (CAGR) of 35% while upstream traffic grows at a 26% CAGR. It is not feasible for us to accommodate this type of growth in every part of the plant each year, instead we plan to address 20% of our footprint each year in anticipation of completing an upgrade cycle each 5-years. This means that our target design needs to result in an increased peak traffic capacity of about 4.5x on the downstream and 3.2x on the upstream. These targets do not include increases otherwise needed to support higher maximum advertised tiers.

Increases in product offerings don't follow the same growth patterns as usage. The maximum offered downstream speeds tend to increase at about 3.2x over a 5-year period. There appears to be pressure on the upstream to increase to symmetric speeds, but it is unclear if that will continue over the long term. It is probable that the primary driver is more likely to be around latency rather than sustained throughput, but marketing from competitors will likely keep the need for increased speed/bandwidth.

There are two types of bandwidth levers: optimize existing spectrum and increase available spectrum. Optimizing existing spectrum could be from optimizing the underlying technology, increasing the efficiency with regards to the payload, and/or by reallocating spectrum between services. Increasing available spectrum can come from repurposing spectrum, reducing the number of subscribers sharing the spectrum (segmentation/node-split), or increasing the frequency range (spectrum) that the plant is able to support.

If we assume that the downstream traffic was 870 Mbps and needs to accommodate growth of 4.5x over the next 5 years, and the maximum tier was 330 Mbps and needs to accommodate a growth of 3.2x over the next 5 years, we can apply our formulas as follows:

$$C \geq P_{95} + T_{Max}$$

$$C \geq (870 \times 4.5) + (330 \times 3.2) \geq 3915 + 1056 \geq 4971$$

This means that we would either need to increase the amount of downstream capacity to nearly 5 Gbps, segment the node, or do some combination of the two. Having the benefit of hindsight for this example, we knew that over that timeframe we were looking to expand the downstream to 48-SC-QAMs and 192 MHz of OFDM, which would give us a downstream capacity of 3.6 Gbps, meaning that some degree of segmentation would be required. We also know that the T_{Max} was actually a 1 Gbps service with a 10% overhead, or T_{Max}=1100 Mbps. We assume that a simple node-split typically results in a 60/40 split in the traffic, a three-way split results in a 40/36/24 split, and a double split results in a 36/24/24/16 split in the traffic. T_{Max} remains unchanged.

$$P_{95} \leq C - T_{Max}$$

$$P_{95} \leq 3600 - 1100 \leq 2500$$

In order to achieve this the higher traffic leg of the segmentation would be targeted to grow to less than 2500 Mbps, which assuming a growth factor of 4.5x means it should be ≤ 555.5 Mbps, or 63.8% of current peak capacity, meaning that most simple node-splits should suffice.

2.1. CoVID Impacts

One of the major impacts of CoVID was to drive people to shelter at home, meaning all activities, including work and school, became dependent upon the residential network. Nearly overnight we noticed an increase in network utilization equivalent to what we would expect after a full year's growth, and at the 1-year mark we measured the expected 35% increase on the downstream but an unprecedented 55% increase on the upstream. The graph below shows a typical weekday before CoVID in orange and after CoVID in blue. It is also worth noting that while the downstream peak remains in the evening (and in-line with typical annual increases), the upstream peak moved to start in the late morning and usage remains significantly higher for roughly 10-hours.

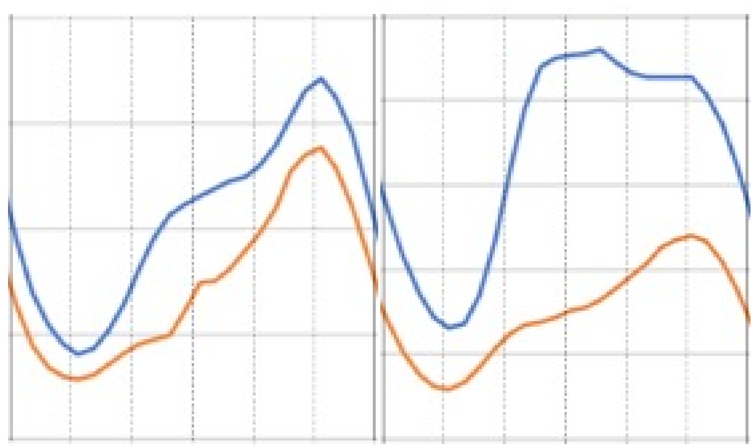

Figure 1 – Typical Weekday Bandwidth Impact, Downstream (left) and Upstream (right)

Our plans to accommodate downstream growth had already included reclaiming spectrum from other services and growing the amount of spectrum allocated to DOCSIS, especially focusing on DOCSIS 3.1 enabled devices that support OFDM capacity. Our plans to accommodate upstream growth were based on DOCSIS 3.1 and a migration to Mid-Split, including the use of OFDMA in the mid-split spectrum. We did not feel that accelerating the move to Mid-Split would enable us to respond as quickly as required to customer demand.

The first option that was proposed was to reallocate some of the ATDMA carriers to OFDMA to increase capacity. This is known internally as Sub-Split OFDMA (SOFA). The challenges with this approach were multi-fold, including requiring a higher penetration of DOCSIS 3.1 devices than was present at the time and a lack of maturity of OFDMA itself. While not implemented for this immediate circumstance, we kicked off a program that will allow us to implement SOFA should the upstream impact continue to last for many years.

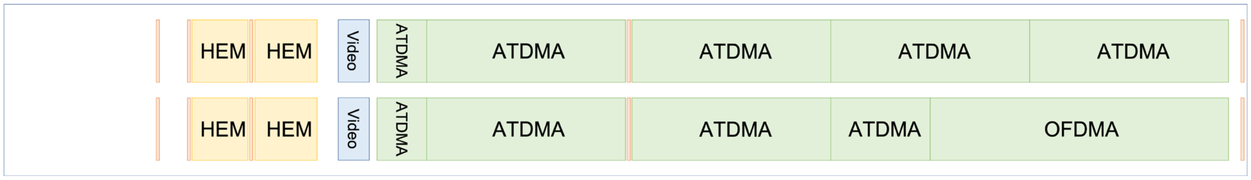

Figure 2 – Pre-CoVID Spectrum vs. Initially Proposed SOFA

We reached out to customers in heavily utilized nodes who were heavy users either intentionally (for example, to incentivize them to move to one of our Fiber products if that option was available) or unintentionally (for example, increasing outreach to customers whose devices may be impacted by malware or may be transmitting higher volumes of traffic due to configuration issues).

The option we ultimately implemented is known internally as fifth-carrier (5C) as it is the fifth carrier in the DOCSIS sub-split upstream that is intended to be used to deliver internet service (there is also a 1.6 MHz ATDMA carrier designated for DSG, which is why the initial state shows five items labeled ATDMA). This required us to work closely with all our markets to make alterations to how the upstream spectrum was utilized by non-Data products. Thanks to significant ongoing efforts resulting in moves from circuit-switched telephony to IP-Telephony, legacy video return to DOCISIS-based return, and dedicated telemetry return paths to using embedded DOCSIS modems in our plant gear, we had the opportunity to retire and/or relocate many legacy carriers, giving us enough spectrum to launch an additional 3.2 MHz upstream channel.

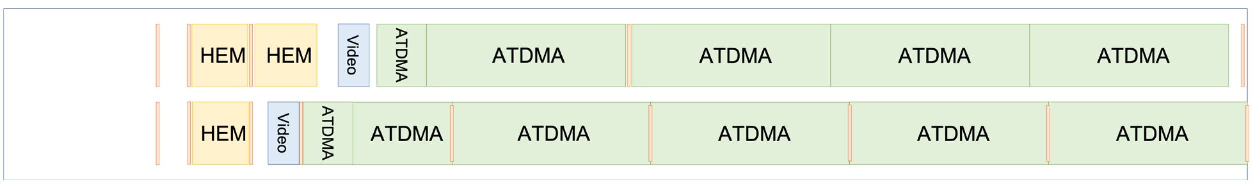

Figure 3 – Pre-CoVID Spectrum vs. Fifth Carrier

This brought along many challenges which will be discussed later in this paper, but this addition of 12.5% more raw spectrum gave us 18% additional capacity towards the P_{95} (as T_{Max} remained the same). This served to reduce the number of congested nodes back down to within 0.5% of the pre-CoVID levels, with an expectation that as this did not completely offset the 26% increase, the number of nodes approaching congestion could be expected to reach a level between 3% and 6% for the remainder of the 5-year upgrade cycle unless another activity brought additional relief. That other program is the previously mentioned SOFA program.

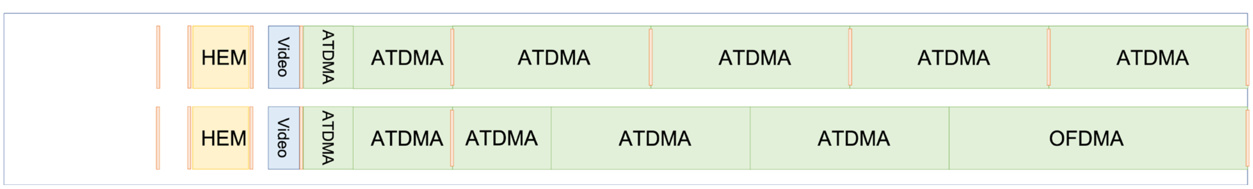

Figure 4 – Fifth Carrier vs. Fifth Carrier with SOFA

When coupled with the 5th Carrier, the newly proposed SOFA option has 2x3.2 MHz and 2x6.4 MHz ATDMA carriers in addition to the OFDMA in the subsplit spectrum region.

3. Model

Due to the sudden change in demand as customer behavior changed due to CoVID, it was crucial for our teams to quickly analyze the growing bandwidth concerns, predict the range of future growth patterns, and determine the impact of different strategies in handling each scenario. The automation of our planning tools allowed such a wide variance to be analyzed within days, and for solutions to be proposed and adjusted with similar speed.

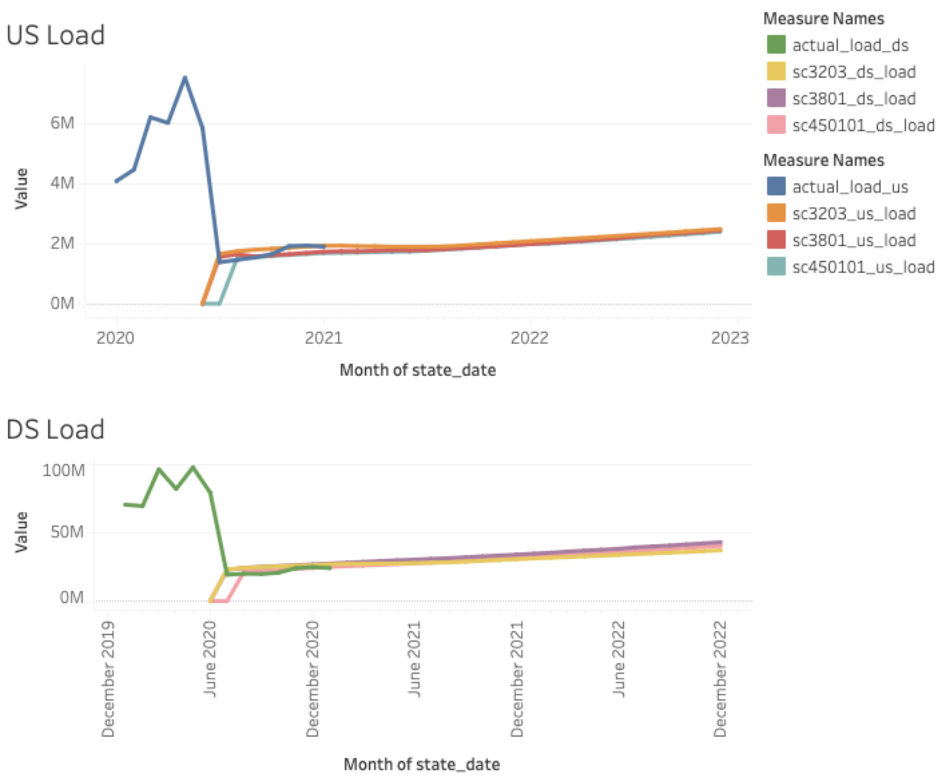

Figure 5 – Enterprise node load compared across forecast scenarios

Our forecasting model predicts network load based on historical node-level traffic data, weighted by seasonality and outlying factors, like the recent node action spike due to COVID. This predictive modeling is constantly updated to reflect new information, or to confirm existing patterns and narrow down the range of possible outcomes. Even with unexpected shifts to otherwise predictable data, the forecast gives a statistically sound expectation.

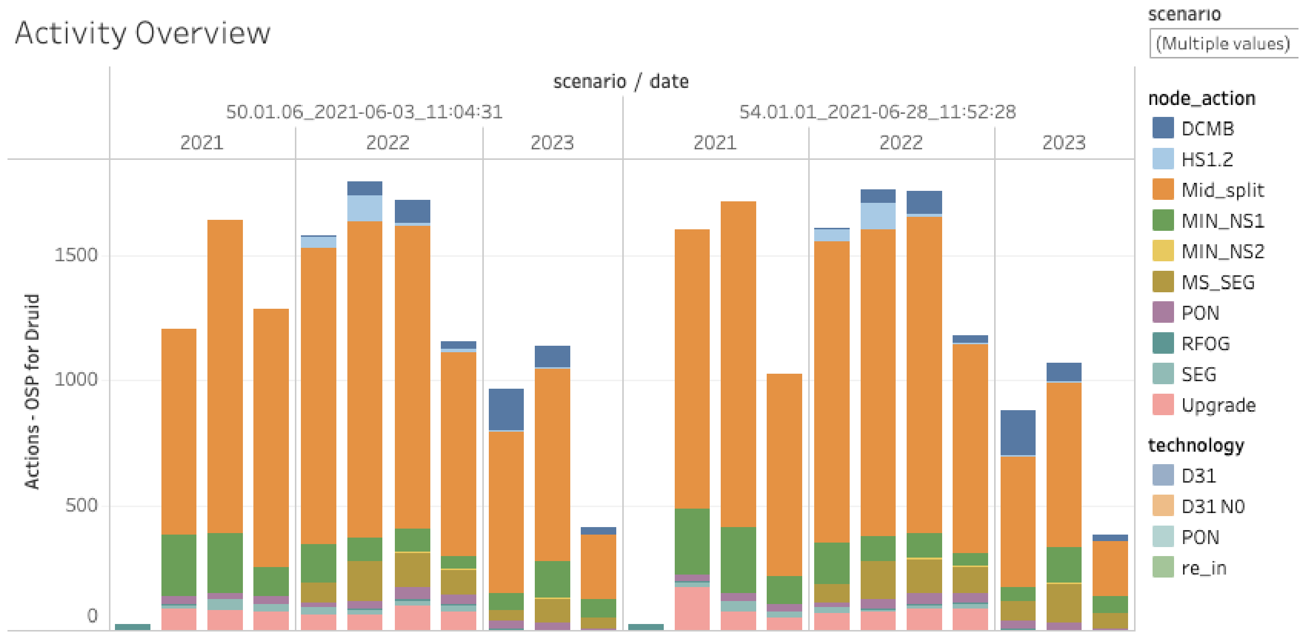

Figure 6 – Automated node action comparison example

Given any one set of expected network utilization, the Capacity Response module compiles a logical action priority based on a determined set of business requirements. As the strategy is adjusted by the planning team, the model us updated to reflect the appropriate action thresholds, technology definitions and restrictions, and deployment rates. The output of any given scenario provides an overview of the deployment schedule required at minimum to manage utilization, which is then processed through the Volume Model for the respectively required hardware and licensing per facility.

With these modeling capabilities, proposals for new technology considerations underwent rapid approval processes as we were immediately able to give relevant estimations of impact for each case. Ultimately, the abstraction of the underlying logistics to each solution allowed the planning teams to focus on solving these urgent problems at a much larger scale.

4. Plan

On Tuesday afternoon of March 10, 2020, we were told to pack up and head home early. The explanation was that the company wanted to conduct an emergency preparedness test with us working from home for a few days. By the end of the week, it was obvious that this was due to a virus with a peculiar name that was starting to get more headlines and an increasing sense of concern. By that weekend, it was clear that we would not be back in the office for at least a few weeks with few, if any of us, expecting that it would be nearly 18 months (or longer) before any of us were back at our desks.

By the following week, the node-action escalation team was seeing an uptick in node-action tickets as people started working from home and students began taking classes online and driving up upstream traffic utilization. Almost overnight, Zoom went from being a relatively obscure app to part of our lexicon. Daytime commercial usage shifted from high-capacity business infrastructure to residential nodes that until then were used primarily for daytime TV watching, personal email, and gaming. At nearly the same time, an internal node-action ticketing automation tool went live, and the node-action tickets went from a typical one or two a week to nine or ten per day, eventually averaging 24 per day in April. The small team responsible for processing the node-action tickets was quickly overwhelmed, leading to recruiting other employees to sort and triage the incoming flood of tickets.

The months of March, April, and May of 2020 saw over 1,200 node action tickets processed in those three months alone, versus a total of 77 by the same team for all of 2019. The previously discussed traffic growth predictions had gone completely out the window as we experienced over a year's worth of demand increase in less than a month.

After it was decided to use a 5th upstream ATDMA carrier to provide more capacity, the initial thought was that we would add this to every node in all markets. This had the advantage of being a global response, however over two-thirds of our nodes were still below the node-action thresholds, and the 5th carrier would have simply been excess capacity. It also would have meant many nodes consuming an unneeded upstream license at a cost in millions for the unnecessary licenses. Thus, the decision was made to add the 5th carrier on a node-by-node basis.

Once that was decided, members of the bandwidth management team were tasked with reaching out to the markets, the video teams, and the regional DOCSIS teams to prepare them for the effort, while the EMO team began pulling the node utilization data from the ACOE group each week. As that was happening, trial nodes were being identified in multiple markets. Where needed, HEM channels to support the circuit-switched telephony were being moved or removed, and the non-DOCSIS based video-OOB signals were relocated (SCTE 55-1 and SCTE 55-2). This was across the enterprise and was a coordinated effort between multiple corporate and field teams and their boundary partners.

Figure 7 depicts a typical transition as it was implemented.

1. The top row represents a typical starting configuration prior to pandemic
2. Where applicable, plant sweep channels and telephony HEM channels were moved or removed
3. The legacy video upstream OOB was relocated
4. The DOCSIS carrier used for DSG was moved downward by the regional DOCSIS team to create a gap
5. New plant sweep channels were placed
6. Applicable nodes had the 5th DOCSIS ATDMA carrier added for use by cable modems and EMTAs

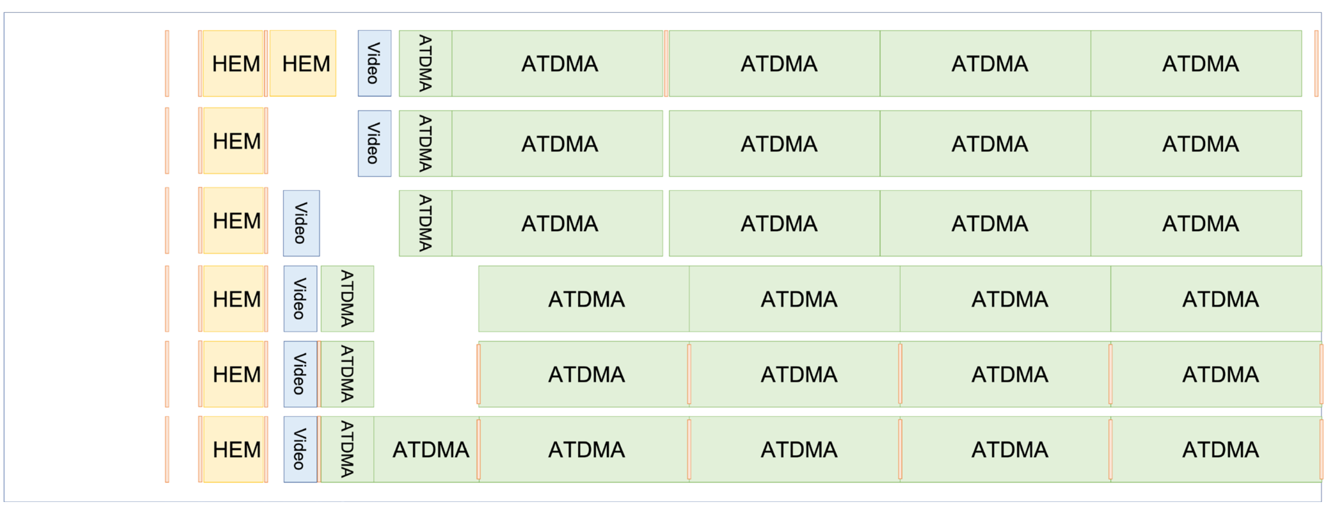

Figure 7 – Preparing to Implement 5th Carrier

The spectrum preparation (through Step 5) was implemented across all nodes in each market, leaving the space for the 5th Carrier to be inserted as required. As nodes hit the Upstream 70% utilization threshold, they were added on a weekly basis to a list to have the 5th carrier turned up. This response not only saved money but allowed for an agile response to prevent congestion where utilization approached node-action levels during the lockdowns. Such a measured response also kept the number of nodes per week manageable and prevented overwhelming the resources needed for deployment of the solution. For nodes not receiving a 5th Carrier, they will remain at the 5th row of configuration until such time as they are migrated to a new spectrum, such as mid-split or high-split. Current estimates show approximately half the enterprise nodes will eventually receive a 5th carrier.

Nodes continued to be added for the 5th Carrier throughout the year and the effort continues to this current day. At any point in time, the total number of 5th Carrier nodes is less than the total number of deployments as previously planned node actions have resulted in many 5th Carrier nodes being addressed with Mid-split.

The 5th Carrier solution is expected to run for the next several years but will gradually taper off as we execute on our plans to implement mid-split or higher capacity upstream solutions. Thus the 5th Carrier may have a fairly short duration within Cox network planning but is notable for the much-needed relief it provided at a critical time due to extraordinary and unforeseen circumstances.

5. Deploy

Typical pre-COVID deployments unsurprisingly involve high-level controls through program management and Engineering governance to standardize customer and employee impacting projects. The phase gates ensure cross-functional collaboration and awareness but also establishes metric baselines, success criteria, trial performance reviews, and early-life metrics to gauge deployment success. Our deployment process framework is internally referred to as Solution Certification. This framework integrates these cross-functional workstreams, fulfills governance requirements, and readies a service for production availability.

The Solution Certification process is used as an integration management tool to organize the Technology Readiness and Operational Readiness efforts through trials. The process' output is a scalable and sustainable deployment model to minimize operational exposures to critical activities. It combines the strategy, modeling, and spectrum planning outputs with the technology & operational readiness workstreams to culminate in production readiness of the deployment fulfillment and service guides. This process is used for large or small spectrum initiatives such as the plant and service enablement of Mid-Split with four ATDMA carriers or expanding the downstream OFDM from 96-192MHz. Overall, this process may take anywhere from 6 to 18 months depending on complexity and impact to field operations.

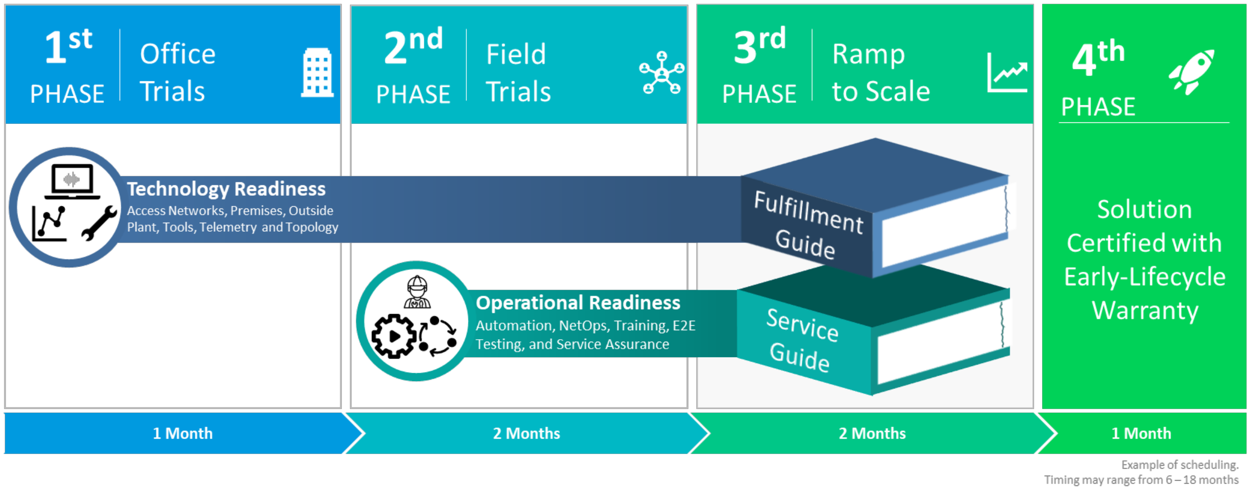

Figure 8 – Solution Certification Process

Office Trials: Technology Readiness' first major milestone that takes the lab-tested deliveries and converges multiple technical workstreams into an office test bed. Trial activities include end-to-end testing, speed tests, node validations, utilization rate checks, and validations for video, voice, and data devices. Pre-COVID, the office trial proved to be an ideal initial gate as employees quickly identify defects that could not be lab tested.

Field Trials: Ops Readiness converges on the certification path to include direct communication with local field teams to understand common issues, validate training, and identify defects that otherwise cannot be captured in analytics. Major activities include the second round of end-to-end testing, on-site node tests, residential & business test account validations, and automation initial testing iterations. Field trials may extend to soaks across multiple markets to widen the trial sampling as a final verification measure before scale.

Ramp to Scale: Handoff to Field Engineering & Operations begins when technology and ops readiness meet trial success criteria. This phase is focused to validate fulfillment and service guides operability in scaled deployment through iterative learning to refine operational processes and automation solutions. Additionally, it is used to test and improve reporting, management, and monitoring tool functionality before the solution is certified.

Solution Certified: Transition to the field is complete and the launch is functioning as designed while operating at full scale. The solution is warrantied for a minimum of 30 days by the program development team to monitor performance and address defects, however, a continuous feedback loop in place to ensure performance and enhance as the solution matures in lifecycle.

Dedicated analytics support is critical to the Solution Certification process for both technology and operational efforts. These reports are essential to trials in not only tracking utilization rates and customer transaction rates but also MER values, D3.1 penetration rate, ticketing, and even track devices in partial service. These trial-focused dashboards are continually modified to evolve it into an operational dashboard as the certification progresses in scalability. For demand-based programs, these dashboards are used as the input to identify node candidates, validate service changes, and also monitor performance to identify anomalies that require escalation to DOCSIS operations and Access engineers. This enables rapid response for our teams to review, troubleshoot, and action on nodes which may necessitate reverting nodes back to original configuration to ensure no harm is done. Ultimately, this dashboard is the key management and reporting tool through a project's lifecycle.

Automation is also a key operational readiness workstream in the effort to streamline changes and processes to meet the high operational tempo that are demanded by field engineering teams. Even small automation updates better enable enterprises to reach scale, improve efficiency, and stay ahead of the curve meet future bandwidth demand and ensure the quality of service expected by our customers. Early collaboration to standardize processes during Solution Certification with operations and engineering teams better enables automation effectiveness and increases agility of future deployments and automation updates.

5.1. Certifying 5th Upstream

COVID's demand for immediate bandwidth increase necessitated the prioritization of a technological solution. The executive leaders' *all hands on deck* call delivered the strategy, modeling, and spectrum plan but now required the development and deployment of more bandwidth. This top-down approach supported the teams by prioritizing this workstream over all others but also limited formalized project and program management as daily direct communication kept leadership abreast of 5th Upstream Carrier status.

Access engineers immediately configured 5th US Carrier in the lab, drafted a MOP in a half-day and worked alongside CPE engineers to form a test plan while outside-plant engineers tested thresholds to determine how to adjust for the low-end 3.2MHz ATDMA. Projects normally estimate 60 days for a single engineer dedicating 25% of their time for lab work, the COVID-team deployed 5th US at Cox Communications Atlanta office in three days. The greatest contributing factors for 5th US Carrier's development expediency were: 1) Leader-driven project crashing with multiple dedicated engineers, 2) standardized CMTS configurations enterprise-wide greatly limiting the volume of testing permutations, 3) and unoccupied offices permitting more aggressive entry to office trials. Labs and 5US carrier was ready for field trial in less than 10 days.

Field trial planning calls involved over a hundred engineers across Advanced Access, Outside Plant, and regional field DOCSIS teams. Each region volunteered to trial 5US, however, several markets required final spectrum alignment moves before it was 5US ready. Five markets across three regions were selected for field trials with Southern California heavily targeted due to exceptionally high utilization rates and heavy concentration of high-profile customers which would be a true test to 5US effectiveness. Overall, 5US was trialed on 15 nodes for a week before ramping to scale.

Concurrent with 5US Carrier's technology readiness workstream our analytics team, the ACoE, developed an MVP dashboard to assist teams to identify and prioritize 5US node candidates based off utilization rates and total customers in the early phases of COVID utilization. Progressive updates were added to develop a Node Prioritization dashboard to support the operational readiness to decisively target the most troublesome nodes with metrics that include high-speed subscription counts, average utilization rates ranging between 1-week to 3-months, and scheduled node actions.

By April 6th, 5US carrier graduated to the Ramp-phase and was soaking on 59 nodes. Engineering teams continued to improve 5US performance by adjusting thresholds but also addressed other issues like CSR VOIP performance, however, 5US full deployment management transitioned to Sustaining Engineering and the field's Bandwidth Management team.

By mid-April, 5US was at full-launch in spectrum-ready markets while the Bandwidth Management team coordinated with local teams to move HEMs, consolidate SCTE 55-1, and move the OOB/Upstream. They established a weekly cadence to review the Node Prioritization reports and queue nodes for the regional DOCSIS managers to review, schedule, and service enable with 5US. On April 20th, 5US had reached scale enterprise-wide with 2,291 nodes enabled by the end of the month.

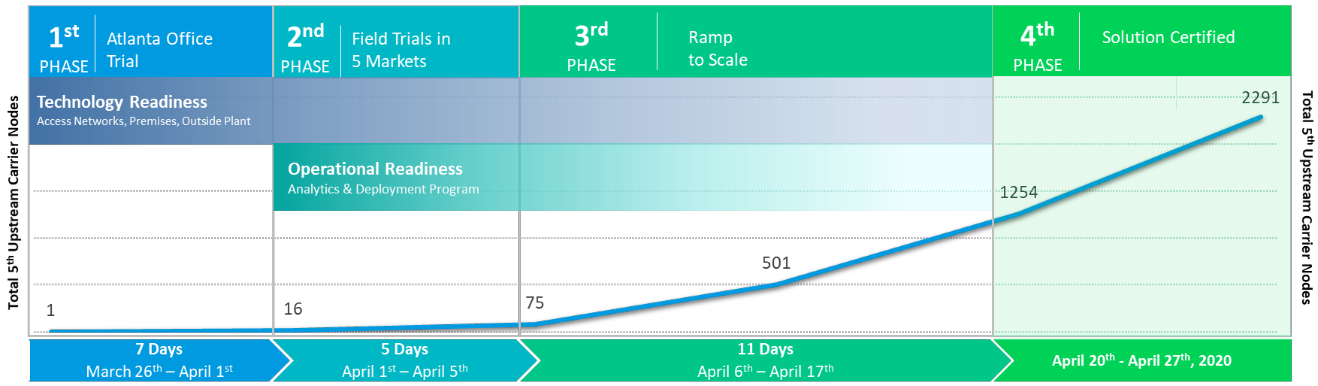

Figure 9 - 5th Upstream Carrier Solution Certification

5.2. 5th Upstream Carrier Deployed

By mid-May, the node volume utilization rates decreased to more manageable levels with teams gradually shifting focus back to the pre-COVID strategic roadmap with small contingents still focused on pandemic related demand. A total of 4,539 nodes were 5US enabled by end of May reducing the total volume of congested nodes ($P_{95} \geq 80\%$) from the 7.1% peak down to 2.6%.

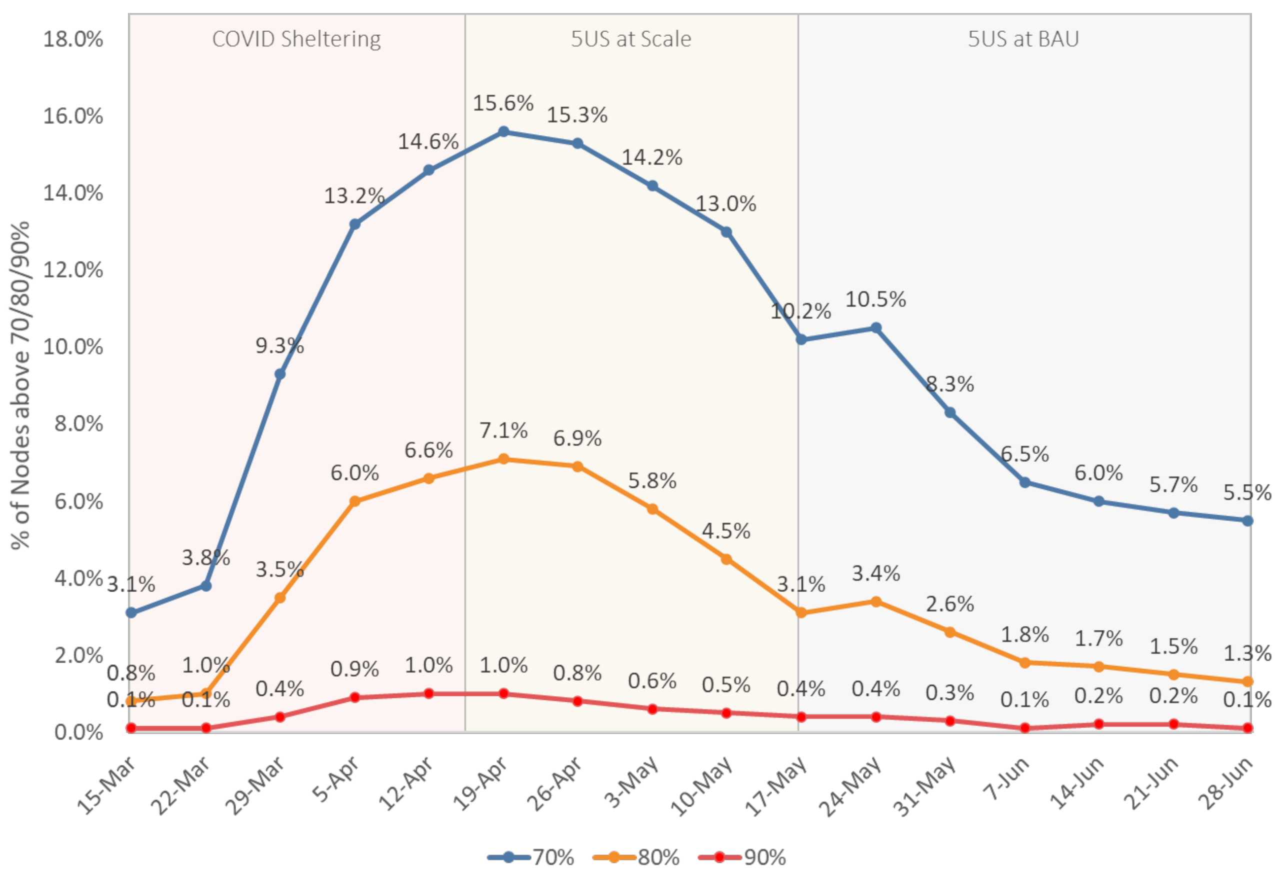

Figure 10- COVID-Driven Utilization and 5US Carrier

By Summer 2020, most nodes that could be addressed by 5US were already enabled and the team standardized the 5US enablement threshold to target nodes with a 4-week ≥70% P_{95}. Deployment rates sharply decreased to approximately 300 enablements per month but spiked to 645 once the remote-learning school year commenced. At the one-year Work from Home anniversary, 5th Upstream Carrier had been deployed to 9,958 nodes with over 42% of deployments occurring within the first 60 days of COVID sheltering.

As of Summer of 2021, the rate of 5th Upstream actions tapered off with some regions going weeks without enabling a fifth ATDMA channel. The total number of active 5US carrier nodes have progressively decreased over recent months as they continue to be consumed by mid-split migrations.

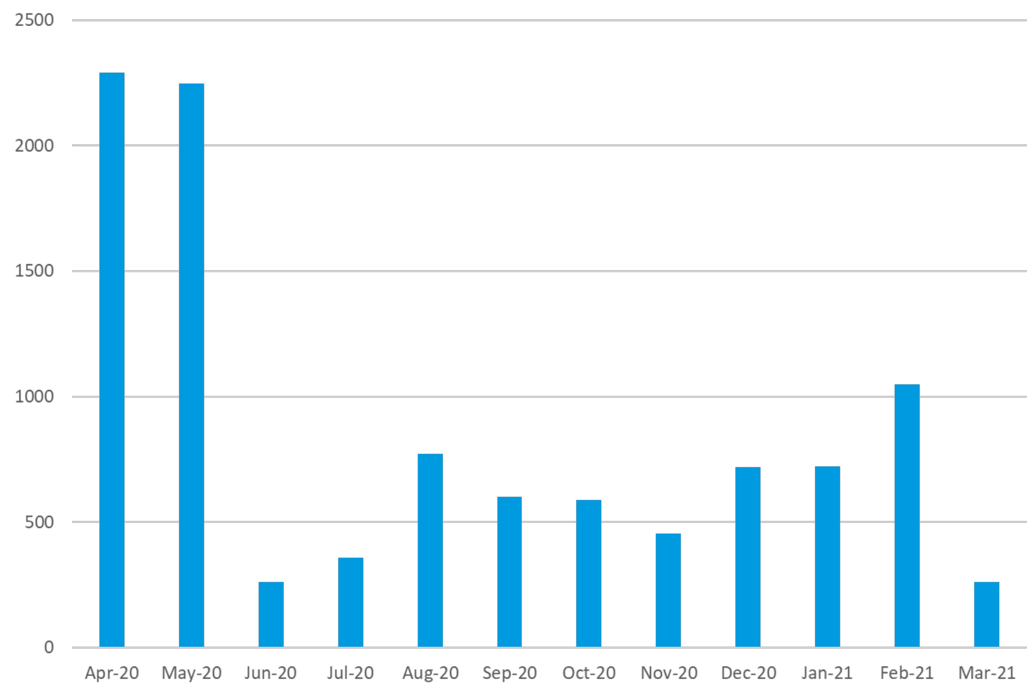

Figure 11- 5th Upstream Carrier Enablements

Operational readiness updates progressively rolled out to support the enablement process like Network Automation's service change update for 5US on RPD to enable scheduling and automated MOP execution. Additionally, the ACoE updated the Node Prioritization report to validate 5US enablements via utilization monitoring of the extra ATDMA channel. This removed the manual spreadsheet tracking of 5US nodes by both the deployment PM and the DOCSIS Operations managers. The Node Prioritization dashboard has since been expanded to also monitor downstream utilization rates and identify OFDM carrier widths to prioritize nodes that require expanded D3.1 capacity.

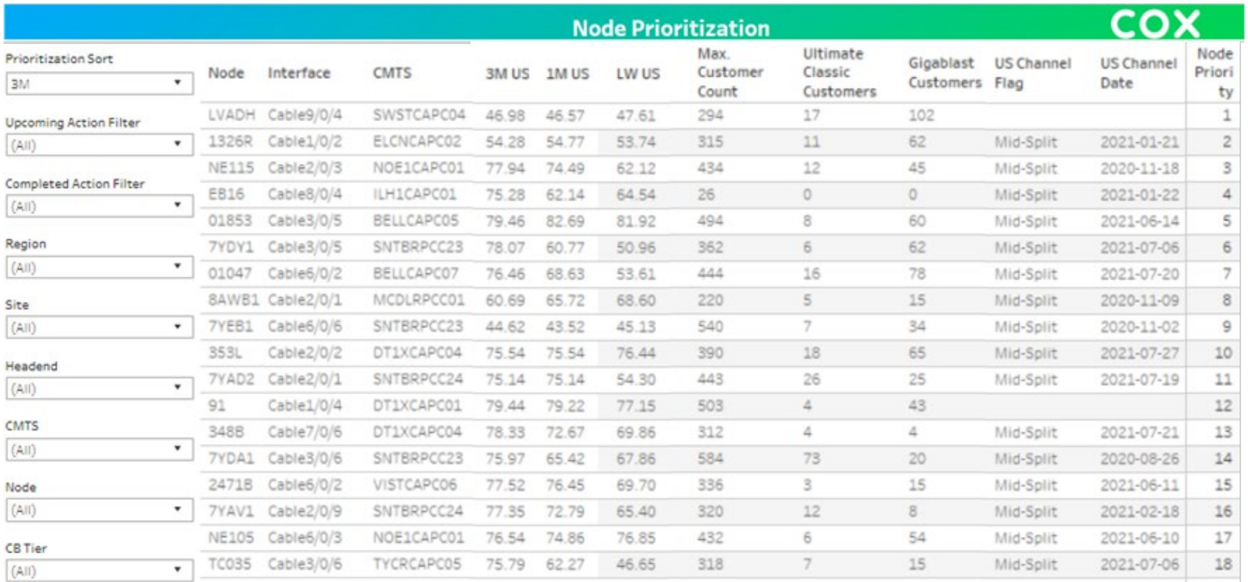

Node	Interface	CMTS	3M US	1M US	LW US	Max. Customer Count	Ultimate Classic Customers	Gigablast Customers	US Channel Flag	US Channel Date	Node Priority
LVADH	Cable9/0/4	SWSTCAPC04	46.98	46.57	47.61	294	17	102			1
1326R	Cable1/0/2	ELCNCAPC02	54.28	54.77	53.74	315	11	62	Mid-Split	2021-01-21	2
NE115	Cable2/0/3	NOE1CAPC01	77.94	74.49	62.12	434	12	45	Mid-Split	2020-11-18	3
EB16	Cable8/0/4	ILH1CAPC01	75.28	62.14	64.54	26	0	0	Mid-Split	2021-01-22	4
01853	Cable3/0/5	BELLCAPC05	79.46	82.69	81.92	494	8	60	Mid-Split	2021-06-14	5
7YDY1	Cable3/0/5	SNTBRPCC23	78.07	60.77	50.96	362	6	62	Mid-Split	2021-07-06	6
01047	Cable6/0/2	BELLCAPC07	76.46	68.63	53.61	444	16	78	Mid-Split	2021-07-20	7
8AWB1	Cable2/0/1	MCDLRPCC01	60.69	65.72	68.60	220	5	15	Mid-Split	2020-11-09	8
7YEB1	Cable6/0/6	SNTBRPCC23	44.62	43.52	45.13	540	7	34	Mid-Split	2020-11-02	9
353L	Cable2/0/2	DT1XCAPC04	75.54	75.54	76.44	390	18	65	Mid-Split	2021-07-27	10
7YAD2	Cable2/0/1	SNTBRPCC24	75.14	75.14	54.30	443	26	25	Mid-Split	2021-07-19	11
91	Cable1/0/4	DT1XCAPC01	79.44	79.22	77.15	503	4	43			12
348B	Cable7/0/6	DT1XCAPC04	78.33	72.67	69.86	312	4	4	Mid-Split	2021-07-21	13
7YDA1	Cable3/0/6	SNTBRPCC23	75.97	65.42	67.86	584	73	20	Mid-Split	2020-08-26	14
24718	Cable6/0/2	VISTCAPC06	77.52	76.45	69.70	336	3	15	Mid-Split	2021-06-11	15
7YAV1	Cable2/0/9	SNTBRPCC24	77.35	72.79	65.40	320	12	8	Mid-Split	2021-02-18	16
NE105	Cable6/0/3	NOE1CAPC01	76.54	74.86	76.85	432	6	54	Mid-Split	2021-06-10	17
TC035	Cable3/0/6	TYCRCAPC05	75.79	62.27	46.65	318	7	15	Mid-Split	2021-07-06	18

Figure 12– ACoE Node Prioritization Dashboard Example

In contrast to our typical Solution Certification which normally takes 6-18 months, 5US for the integrated chassis fast tracked the process in 23 days after entering the lab with 5US for RPD following three weeks later. This acceleration was due to a multitude of factors highlighted above, however, the greatest factor to the lightning-fast deployment was due to the strategic decision to use ATDMA. This proven and mature technology was developmentally and operationally less burdensome than OFDMA and could be deployed with minimal risk. Of particular note is that the 5th Carrier was implemented without any measurable impact to call volumes or truck rolls. To implement such a project without the customers noticing was an achievement in itself, and a testament to how the compressed schedule did not sacrifice attention to customer experience.

6. Conclusion

For the near future, Cox network planning will continue per capacity calculations as described in the early sections of this paper. As part of our efforts to accommodate constantly rising demands on utilization, we will be applying capacity levers such as 5th Upstream, SOFA, and Mid-split. In parallel will be the continued march toward the phase-out of legacy video to enable downstream OFDM expansions and transition to all-IP. The pandemic created an unforeseen spike in demand that likely will never be repeated, if only because post-pandemic trends indicate that many people will continue working from home, which will be factored into future growth models. However, for that once in a lifetime event, Cox had the resources, leadership, and corporate agility to able to react to the unprecedented growth and increase our network's resiliency to take care of our customers' needs.

Abbreviations

5US	Fifth Upstream ATDMA Carrier
ACoE	Analytics Center of Excellence (Cox Specific)
ATDMA	Asynchronous Time-Division with Multiple-Access
BAU	Business as Usual
DOCSIS	Data Over Cable Service Interface Specification
HEM	Head-End Modem (Circuit-Switched Telephony)
MER	Modulation Error Ratio
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division with Multiple-Access
OOB	Out of Band
SCTE	Society of Cable Telecommunications Engineers
SOFA	Sub-Split OFDMA
TNPM	(IBM) Tivoli Netcool Performance Manager

Bibliography & References

Ulm, John, and Cloonan, Tom. (2017) *Traffic Engineering in a Fiber Deep Gigabit World.* A Technical Paper prepared for the 2017 Fall Technical Forum of SCTE-ISBE/NCTA/CableLABS.

Cloonan, T., Emmendorfer, M., Ulm, J., Al-Banna, A., & Chari, S. (2014). Predictions on the evolution of access networks to the year 2030 and beyond. *The Cable Show NCTA/SCTE Technical Sessions, Spring*, *38*.