Repair The Ides Of March:

COVID-19 Induced Adaption of Access Network Strategies

A Technical Paper prepared for SCTE•ISBE by

Dr. Robert Howald
Fellow
Comcast
1800 Arch Street, Phila, PA 19103
215 286 8037
robert_howald@comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1.0 Introduction

The instant network stress test that resulted from COVID-19 is something the most precognizant of network engineers could not have foreseen. Even if they had, the most convincing of them would have little chance of swaying financial teams to increase investment for a traffic spike to come, and a ramp to all-time high sustained usage plateaus that would come in a matter of days.

Most importantly, the performance of the network during this most demanding and critical period shone a light on the industry in a way no press release, marketing campaign, or new service feature could. Nothing was more essential to our customers than delivering their known services, on steroids, robustly and 24/7.

This paper will review the COVID-19 period in five parts. First, we look at the raw numbers as work-from-home, school-from-home, zoom-from-home, and virtual happy hour became everyday behaviors. We will review the data and applications, downstream and upstream, against network capabilities and explain why things went well and where there were challenges.

We will then look at the response from three perspectives. First, we will assess utilization metrics in the most challenged areas, evaluate vulnerabilities, and describe quick response actions that maintained network health.

Secondly, we discuss acceleration of near term initiatives to further enhance capacity in these areas, but also more broadly to address the national rising traffic tide.

Lastly, we will zoom out and discuss how the dynamics of 2020 effect the strategic network plan.

In the final part of the paper, we will discuss the "new normal" of usage, applications, and behaviors, and describe how network planners should consider them in strategy development. Among the silver linings of the pandemic is confirmation that the industry has been focusing on, and executing on, all of the right things over the years leading up to it, and the network is indeed prepared for scenarios that are, literally, beyond imagination.

2.0 By the Numbers…..

The two months after the first reported case of COVID-19 in the United States (January 21, 2020) dispelled the notion that the country might escape catastrophic ravages of a pandemic that began overseas, the dreaded "community spread" of the virus kicked into gear. In a period of time on the order of weeks, a population that had been largely not focused on COVID-19 was experiencing some of the most extreme changes to life-as-we-know-it in a generation.

Just as incredibly, an initially bewildered population adapted to this new normal, following rules and recommendations from government leaders for what would otherwise be considered draconian civil liberties restrictions. Furthermore, the rules and recommendations changed frequently and varied from state to state. Nonetheless, the US population, at least initially, rallied behind "we're all in this together" to achieve the objective of flattening the curve. Indeed, all "you-know-what" broke loose later in the Spring, but that's left for another genre of publication to discuss.

Figure 1 captures this COVID-19 case-tracking period, showing the rapid rise that began in the second half of March.

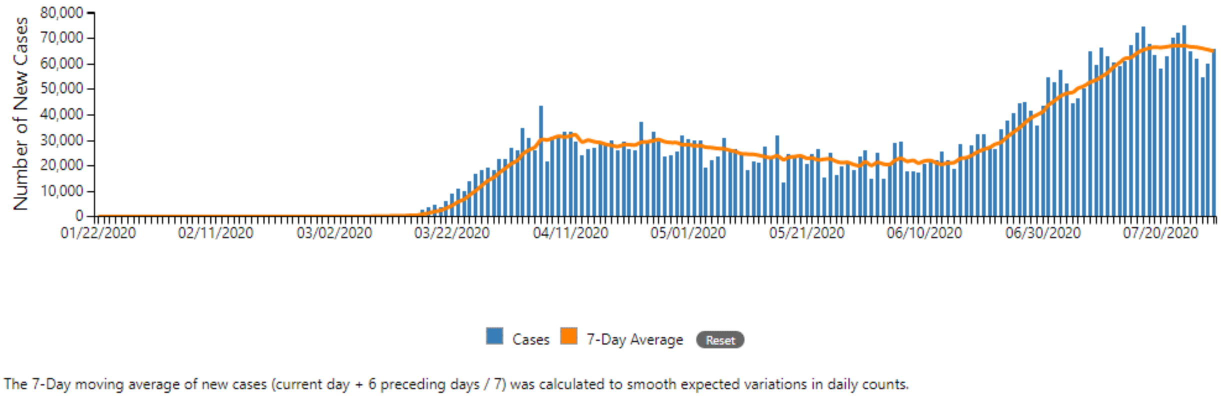

Figure 1 - COVID-19 Cases vs Time in the United States [15]

As admirable as the American people may be, one of the major contributors to the adaptability and relative smoothness of the transition into this unusual new routine was the capability and performance of the broadband network. It allowed for the virtual connectivity essential for the transition, and enabled high volume use of bandwidth-heavy, video-rich applications in homes to be exercised at unprecedent scale, often to substitute for activities no longer available to individuals under "shelter in place" orders. In addition, while massive furloughs and increased unemployment are tragedies caused by the pandemic, the broadband network provided a means for a large percentage of the population to work from home and do so effectively. It is estimated that about 29% of the US workforce is able to effectively work from home [11]. The network also enabled schools to continue to operate, albeit at a reduced level of intensity and effectiveness.

Figure 2 and Figure 3 show the Downstream (DS) and Upstream (US) traffic growth, respectively, over time, as shelter-in-place struck. Most honest broadband technologists would tell you that if you had predicted on March 1 that their networks would see traffic spikes over a period of two weeks like those shown in Figure 2, they would have expected some crash-and-burns in the field, and accompanying negative media coverage.

Rather than a spike of Internet outages, however, what was observed was that while there was an occasional connectivity "housefire" (which was able to be extinguished relatively quickly), in fact there were simply not huge numbers to address. Generally speaking, the impact was absorbed well due to years of experience of managing network traffic, understanding peak and average utilization thresholds that effect the customer experience, and continuously preparing and optimizing the network for traffic growth [1]. Internet traffic growth has not abated since the launch of cable broadband in the late 1990s; only the pace of growth changes year to year – the so-called Compounded Annual Growth Rate, or "CAGR." Network engineers responsible for traffic management of their systems speak and think CAGR in their sleep, with an occasional nightmare about 8k streaming holograms.

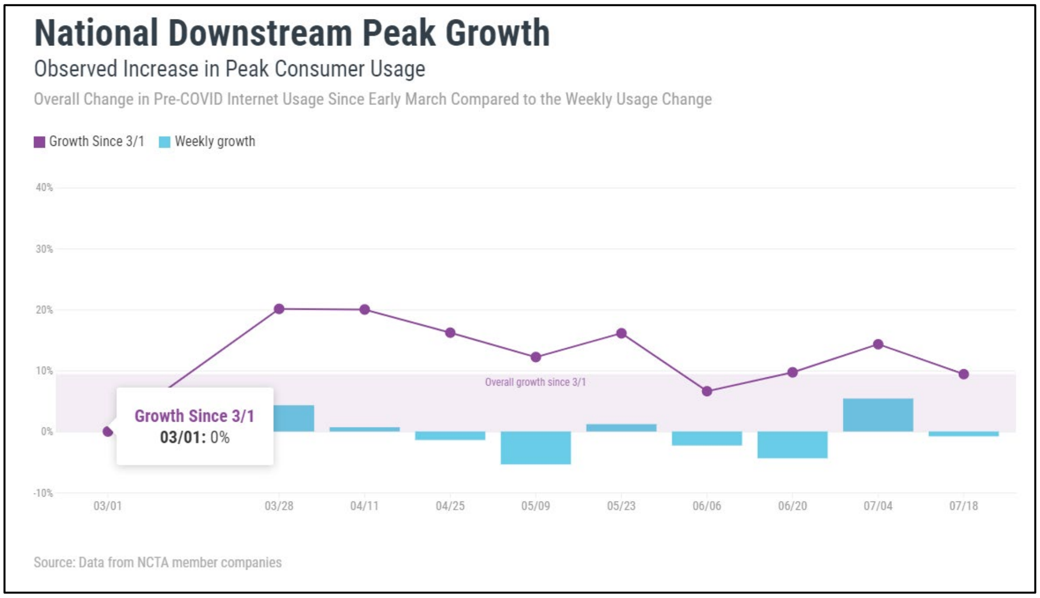

Figure 2 - COVID-19 Induced Downstream Traffic Growth [14]

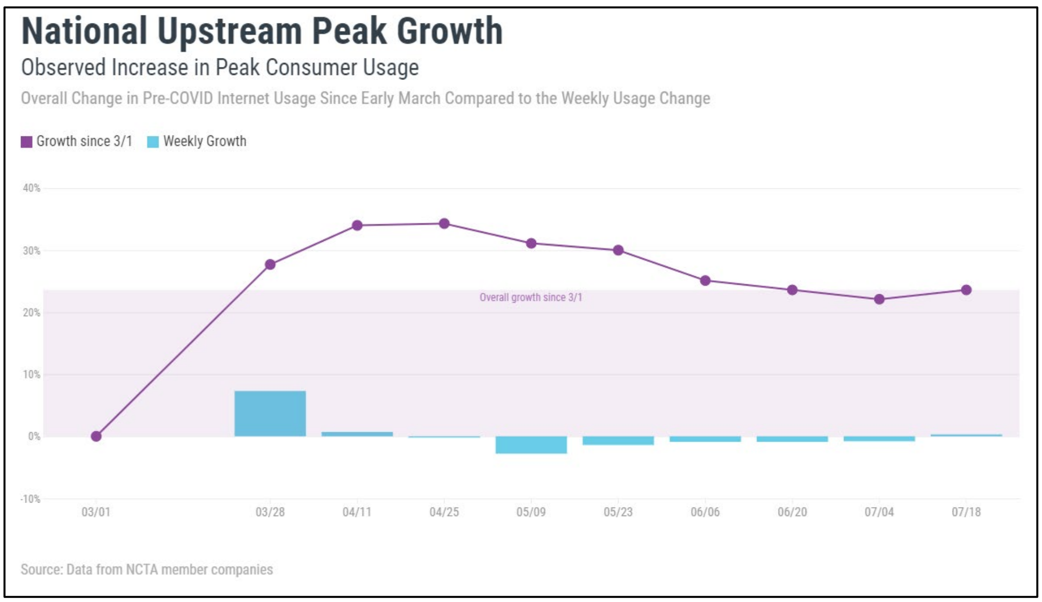

Figure 3 - COVID-19 Induced Upstream Traffic Growth [14]

Let's look at the COVID-19 spike in the context of CAGR. DS and US CAGR vary, of course, and historically the DS has grown 40-50% per year, while the US has been closer to 20-25% per year and has been more dynamic year-to-year. Upstream years of 15% have been recorded as well as very aggressive years, such as when peer-to-peer file sharing took off in the early 2000s. More recently, home security cameras streaming video are becoming noticeable contributors as their popularity increases, and this increase is tied to smart home automation trends in general.

In recent years for the Downstream, many over-the-top (OTT) video service options have fully matured into the marketplace. And yet, people have not grown more than two eyeballs to consume the voluminous programming. It is a bit early to make the call of a downward trending DS CAGR, but it has settled into the 30% range.

Below are some very useful "napkin math" rules of thumb (pre-COVID speak alert – oh, how great will it be to be able to grab a bar napkin to do math, diagram architectures, and sign contracts!). We can relate CAGR to the so-called Traffic Doubling Period (TDP):

TDP = 2 years when CAGR is ~40%

TDP = 3 years when CAGR is ~ 25%

TDP = 4 years when CAGR is ~ 20%

Now, consider the COVID-19 traffic spikes in Figure 2 and Figure 3. Downstream growth peaks at about 20% and has since receded to 10%. This tapering to 10% coincides with the loosening of restrictions, but also with the onset of spring and summer weather. This year more so than any other people were anxious at the first opportunity to get outdoors. Furthermore, the summer broadband season is typically slower. So, it may be premature to assume that the 10% is a settling point for DS CAGR after the pandemic. If we use the more conservative guideline (from an architect's perspective) of 20%, at least until we know more about what a post-COVID new normal looks like, then the virus has accelerated Downstream traffic growth by 7-8 months, using 35% as a reference DS CAGR.   The Compounded Monthly Growth Rate (CMGR – unfortunately not as easily expressible as "CAGR" in conversation!) for 35% works out to 2.5%. It seems clear that some behavior and societal impacts will remain after the virus is no longer a threat large enough to warrant restrictions.

For the Upstream, an initial spike of 35% has receded to about 25% as of July of 2020. Using an US CAGR of 25% means that traffic has accelerated by just about one year. The 35% spike, however, if sustained, is an acceleration of about 16 months (1.9% per CMGR). Generally, cable operators are driven to action in the network by congested upstream. The primary reason for this is simply that there is much less upstream, home-to-headend spectrum available, and it is more difficult to use. The network was built as a one-way video broadcast network originally, so Downstream spectrum was emphasized. Fortunately, analog video is so sensitive to noise and intermodulation distortion that the network was built with very high fidelity, even after being carried through dozens of amplifiers, in the earliest days of cable, enabling today the opportunity with fiber and DAA to have tremendous digital communications efficiency. So, the limited upstream has been able to be made quite powerful, given the available spectrum, which is typically 5-42 MHz in North America – or about 5% of total available capacity, depending on where you draw the upper spectral boundary for the downstream signal path. Cable operators have dealt with this limitation of available US spectrum since the launch of High Speed Data services (HSD), in the late '90s, and of course benefit from the fact that the residential asymmetry ratio of DS to US traffic is about 12:1 today, and has never ventured below about 7:1.

The bottom line is that while the US has been exploited very effectively, it more often than not drives plant upgrade strategies. Several access network initiatives in the industry are about adding US spectrum via 85 MHz Mid-Split, 204 MHz High-Split, and DOCSIS 4.0 Full Duplex and Extended Spectrum. The COVID-19 acceleration of 16 months is especially meaningful in terms of its impact to US ports close to the congestion threshold. A traffic spike such as this would normally occur over many months, and under normal circumstances that yields sufficient time to plan the upgrades necessary to avoid oversubscribing the link.

On the flip side, operator awareness of US bottlenecks had already been quite high, so the tools and strategies to deal with it were well known and in place to manage it. The difference with COVID-19 is that the multi-year plans lost the luxury of much of that prep time and budgets, when it accelerated forward. It is in scenarios where fiber nodes were on or approaching utilization alarms within a year that network congestion was triggered – the aforementioned "house fires" – where work to resolve it needed to begin immediately.

2.1 By the Geography….

Averages provide a general understanding of network behaviors and trends. However, they can be dangerous as well, as regional traffic variations and network performance are tied to the issue that matters most – the customer experience. Where technology-centric businesses, professionals, and universities are key tenants of a region's environment, populace and demographic, average and peak traffic metrics and growth will be correspondingly higher than in locations without these characteristics. So, as valuable as National averages are, operators have to be much more surgical in their knowledge, and in their response. Regional variations are too significant to have a one-size-fits-all network strategy, under normal circumstances, much less considering something as dynamic and unprecedented as a pandemic spike.

Figure 4 is a snapshot the network numbers for several representative states, all of which are served by Comcast as well as other cable operators.

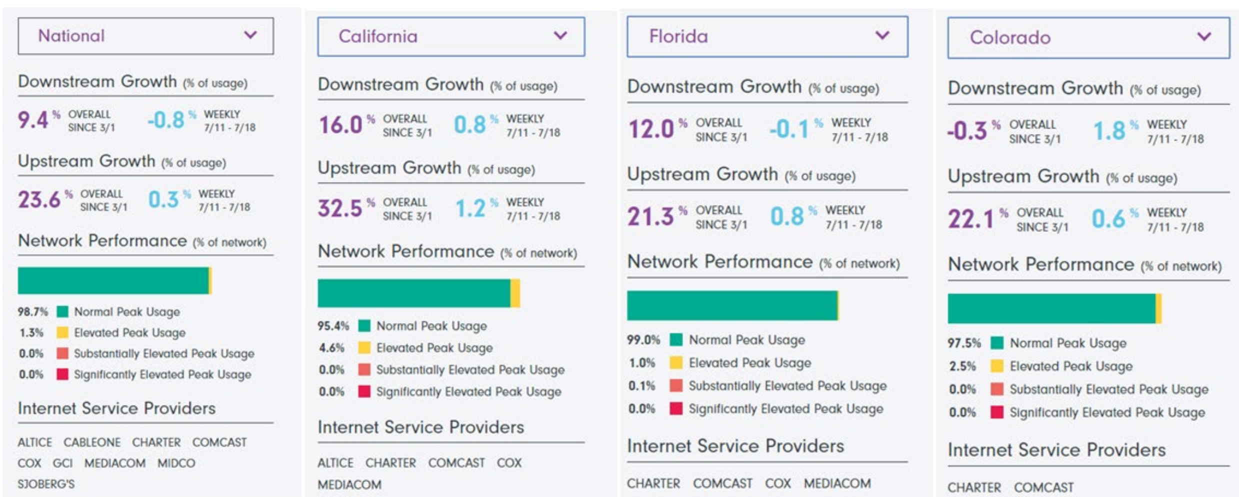

Figure 4 - Regional Traffic Variation are Significant [14]

What is clear from Figure 4 is that the National average would be insufficient as a guide for a COVID-19 response plan for California but might hit the mark in Florida. Alternatively, Colorado appears to be an area where prioritizing US should definitely take precedence over the DS. It is often the case that when addressing congestion in one direction, such as upstream, it is cost effective to address the other direction at the same time.

A conclusion that can be drawn, and a reality of the experience in the early weeks of the pandemic, is that reactive operations to put out the "housefires" were not uniformly distributed. High tech corridors, such as exist in California and other parts of the country, accounted for a disproportionate share of rapid response efforts. Pandemic or no pandemic, non-uniform implementation or staggered timing of network strategies is a common component of access evolution.

3.0 Online During a Pandemic: What Are We Doing?

It is clear and by now intuitive what applications fueled the surge in traffic, including why the US has seen the larger increase. Figure 5 shows the rapid rise of video conference and video chat applications in response to the shelter-in-place lockdowns. Such a rapid spike over a short period of time is exactly the kind of thing that a network engineer loses sleep over. The increasingly popular business conferencing tool Microsoft Teams is overshadowed by the Zoom application, which exploded as the pandemic hit, as a physically-distanced, family-friendly chat experience. However, in relative growth terms, Teams also exploded exponentially, going from 560 Million minutes on March 12 to 2.7 Billion minutes by March 31, or about 70% growth per week during that period of time [4].

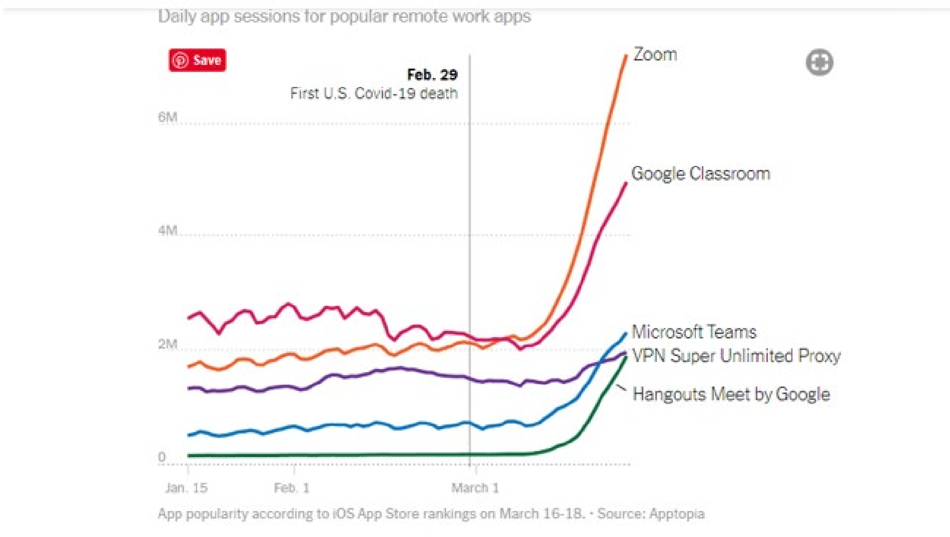

Figure 5 - COVID-19 Impact on Video Conference and Chat Applications [9]

Figure 6 shows the traffic on a weekly timeline, highlighting how much can be attributed to work-from-home (WFH) employees, as evidenced by the dip in traffic over the weekend.

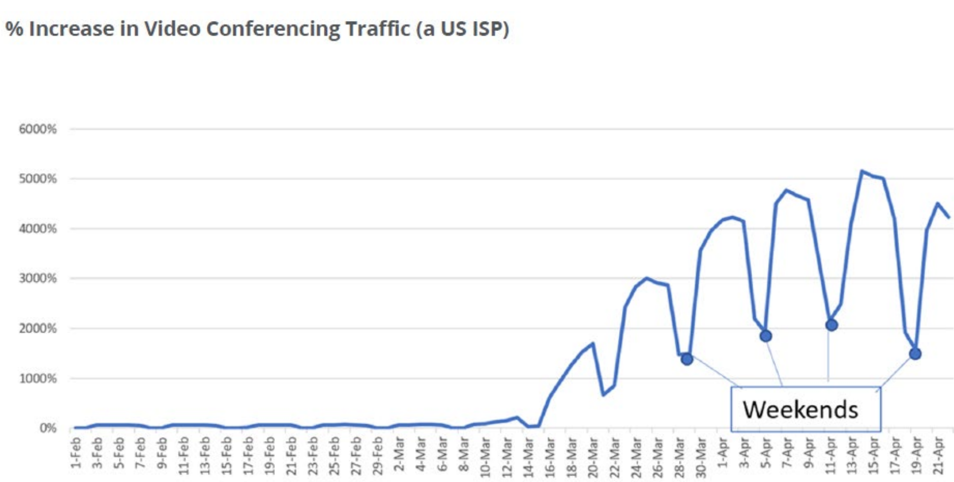

Figure 6 - COVID-19 Impact on Weekly Video Conference Traffic [2]

To punctuate this conclusion, Figure 7 shows how the upstream has experienced a shift in its Peak Busy Hour (pbh) usage, from about 9 pm local time, to 2 pm local time. The biggest percentage increase over pre-COVID utilization is actually at 1 pm, but because the characteristic of the US is a gradual rise throughout the day, the absolute peak has moved to about 2 pm [1].

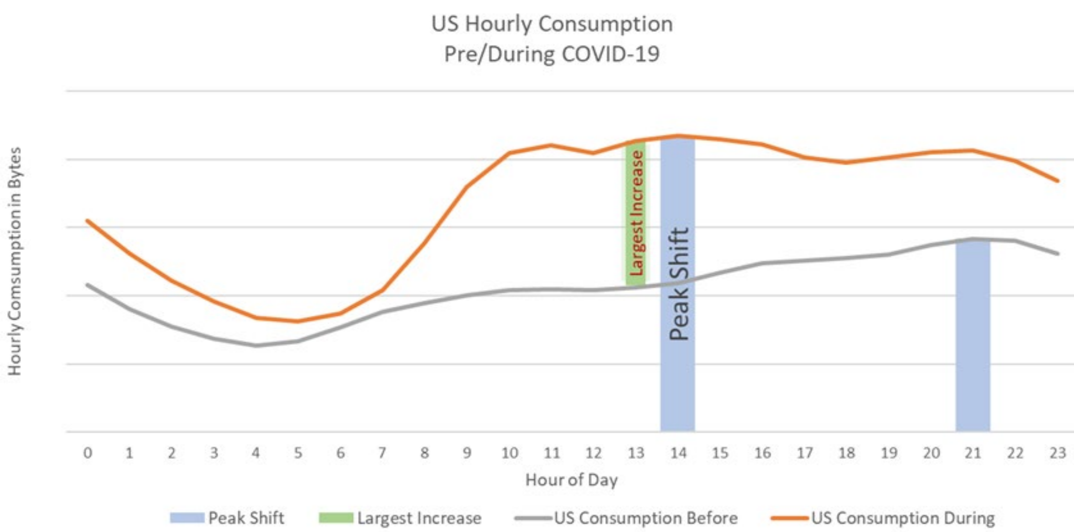

Figure 7 - COVID-19 Impact on Peak Busy Hour Time Shift of Upstream Traffic [1]

Finally, Figure 8 shows the predictable spike of virtual private network (VPN) traffic at the onset of the lockdown, such as would be expected when a massive shift to WFH has taken place.

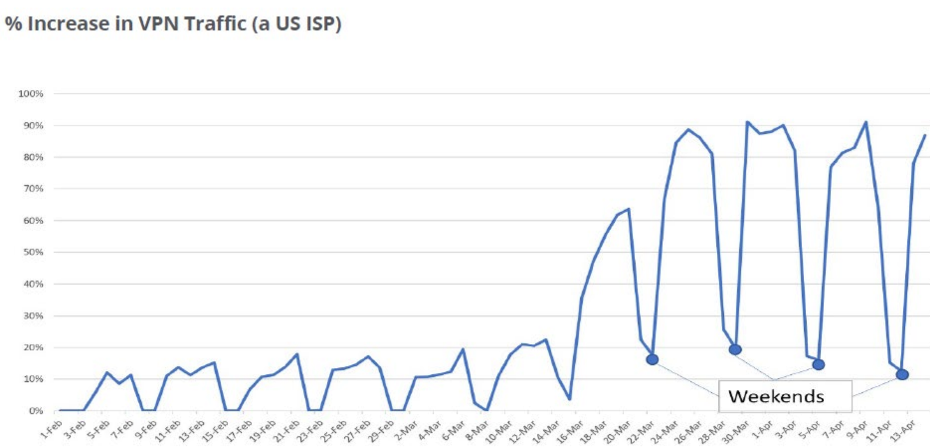

Figure 8 - COVID-19 Impact on Weekly Virtual Private Network Traffic [2]

Moving onto other expected online implications of shelter-in-place, Figure 9 shows the increase in use of streaming video usage, as alternative entertainment options – movies, concerts, etc. – were eliminated. The taper not long after the upward trend is attributed, again, to a combination of spring weather providing opportunities to relieve cabin-fever, and possibly, “binge watching fatigue,” as reality set in following initial estimates of short term lockdowns. Note that the call-out labeling on this figure appears to be in error – the highlighted dates appear to be weekends.

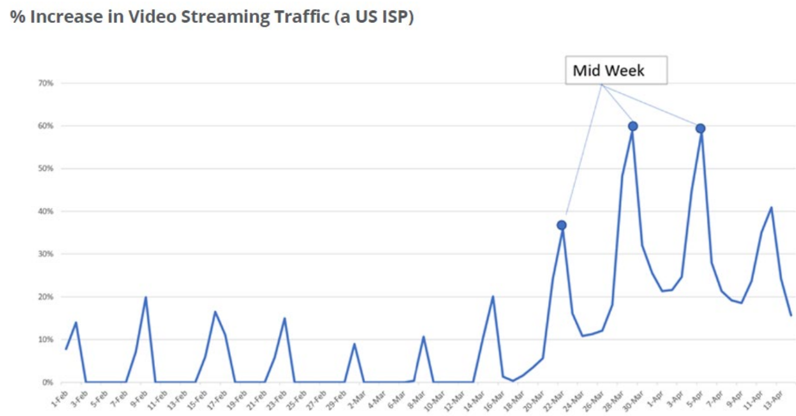

Figure 9 - COVID-19 Impact on Streaming Video Traffic [2]

Not captured in any of the charts but also intuitive is that the applications identified have seen growth in being accessed through web URLs instead of mobile applications. Figure 10 shows how trends for popular applications representing streaming video shifted rapidly away from apps and towards websites. The relevance to this paper is somewhat secondary. The large increase in streaming video traffic, which makes up approximately 70% of all web traffic, is still the main story. However, the secondary effects are twofold:

1) Applications accessed on small screen devices will on average translate to lower bit rates of video compared to their desktop or larger screen/tablet counterparts
2) Web access as well as the mobile devices themselves are more often accessing content via Wi-Fi instead of mobile networks, driving more traffic to the wireline broadband network.

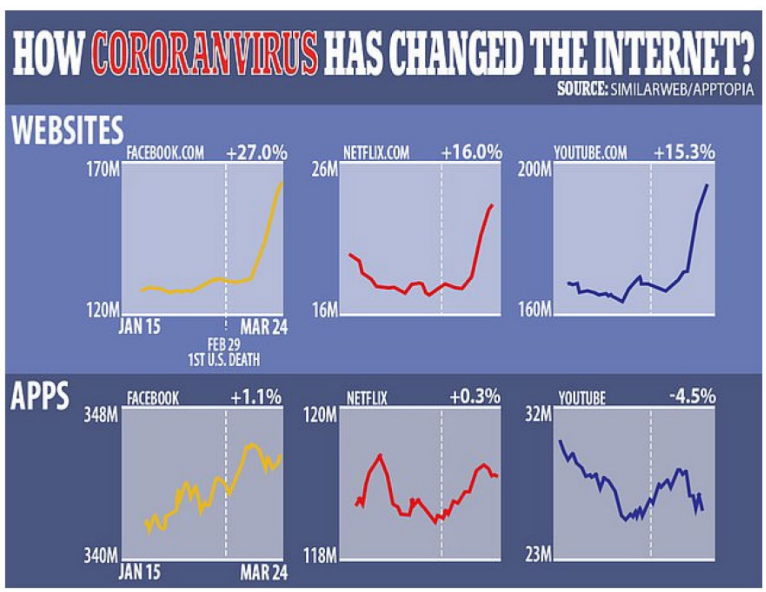

Figure 10 - COVID-19 Impact: Web Access vs Mobile Access [10]

A final interesting traffic dynamic is shown in Figure 11, which captures gaming application growth. Again, with entertainment options limited, and gaming already a popular and growing online activity, regular gamers had less reason to back away from the console, and casual gamers had good reason to increase usage. An interesting observation is that, while Figure 9 has the familiar peak and lull pattern common to weekly residential traffic behavior, the pattern itself was altered in the gaming space, at least in the early weeks. Weekdays and weekends blended into one continuous gaming binge! It will be interesting to see over the long term if gaming was a temporary diversion or if it also has achieved a new baseline level of significance and permanence due to once-casual gamers becoming more hardcore participants.

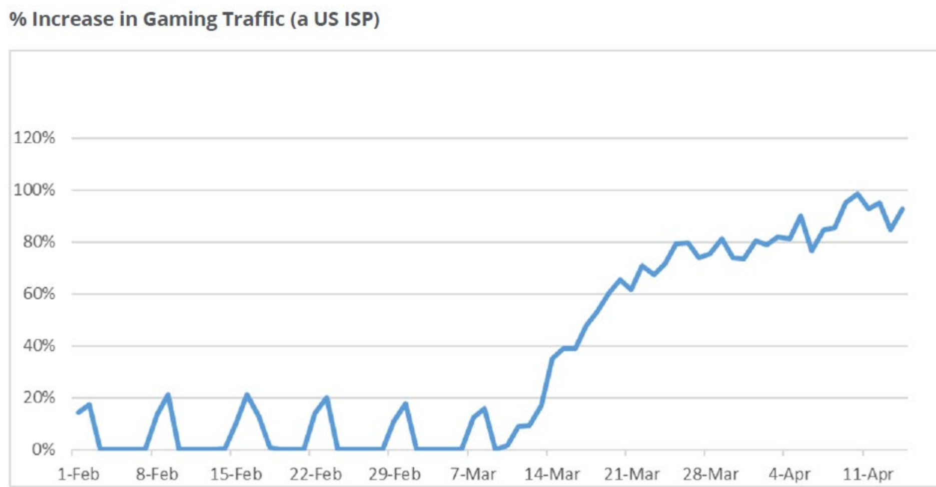

Figure 11 - COVID-19 Impact on Gaming Traffic [2]

4.0 Post-Pandemic: So….What Might Normal Look Like?

As mentioned at the outset of this section, in retrospect we have a largely intuitive outcome from the pandemic restrictions. While it is difficult for anyone to "crystal ball" in such an unprecedented situation, there are some consensus beliefs emerging about the Post-COVID-19 "new normal" with respect to traffic and the network. Permanent changes of the traffic baseline above the pre-pandemic steady state is an expected outcome. Summarizing some of this emerging conventional wisdom:

- It will be truly awesome when we will be in a position to compare projections with the real Post-COVID-19 traffic data!

- *Working from Home (WFH)* – WFH will play a larger role for office workers. The comparable productivity observed, for companies where this was the case, will have companies reconsidering the potential savings associated with the operating expenses of commercial office complexes. People density in an office building may be changed forever, and trends like open floor plans re-thought. Video conferencing had been adapted at Comcast by the end of 2019, and in many other companies as well. With an expected permanence to more WFH, applications like Microsoft Teams, WebEx, Zoom, etc., which have been indispensable from the outset of the pandemic, will be more essential. Also, their capabilities are likely to improve with a higher focus on scale, emulating face-to-face features, and blending in attributes that f2f meetings cannot obtain, such as the virtual backgrounds (I'm in Hawaii! I'm on a ski lift!) and associated computer-aided graphics seen in some of these applications already today.

- *Normalization of Online Education* – Universities have been offering online classes for many years, and some universities are entirely online. However, for most full-time universities and certainly primary and secondary schools, face-to-face instruction is still considered as the most effective means of instruction. Online courses offer much more convenience, if a somewhat less effective alternatives. However, tools focused on a classroom experience will now get significantly more attention in a bid to emulate true classroom environments, while offering capabilities only a digital environment can easily do, such as shared whiteboarding, offline non-disruptive chat, and an audio and video record of the experience.

- *Normalization of Virtual Social Gatherings* – The "virtual happy hour" has become a ritual among my spouse and her friends and with family. Despite no ritual clinking of glasses to "cheers", shelter-in-place orders caused a much wider interest in applications like Zoom. A positive side effect of requiring such applications for any social contact is that through this new familiarity, families who were physically separated by virtue of living in different parts of the country for years began to use it to virtually visit each other far more often than they otherwise would have. In this way, the pandemic has brought families which are naturally distanced from one another into more frequent contact. It remains to be seen if that continues post-COVID, but it seems likely that the impact can only be a net increase in these applications.

- *Preparation for the Next Apocalypse* – Many businesses reliant on personnel in factories, warehouses, logistics centers, etc. suffered substantially with the absence of these employees, especially in those industries not considered "essential." Production, logistics, supply chains, employment itself, etc. were disrupted, magnifying economic woes. Companies with bread and butter in goods and services tied to manufacturing, shipping, and logistics are expected to invest in increased digitization and automation of operations, including use of more machine-to-machine communications (M2M) as part of robust contingency planning, to limit the financial damage that future pandemics or unforeseen interruptions of operations can cause.

An interesting pandemic perspective, from a sales point of view is that, as a salesperson, face-to-face meetings, lunches, dinners, etc. are a key element of the business relationship. Sales teams are generally in a similar boat regarding access to face-to-face meetings. While established accounts are generally as vulnerable (or not) as always, given this equal playing field, new business accounts are increasingly difficult to land, as the reliance on relationship building and face-to-face discussions are magnified.

If this is so, the pandemic effect may be detrimental to disruptors, as well as to new business areas in the incubation phases for established OEMs. Beyond the survival mode drag placed on technology development already, this could result in a stall to innovation in the industry as a whole. As technologists and representatives of the industry with a mutual desire to keep our ecosystems healthy, this is something to keep an eye on. There has been an appropriate focus on "keeping the lights on" activities and minimizing risk to keep people at least virtually connected. However, given traffic acceleration impacts of COVID-19, it is even more important that innovation does not suffer a setback as technologies to deal with it have been essential.

5.0 Next Generation Nimble

The above data tells a clear story and suggests an expected post-pandemic permanence to some of these trends. That said, the job of any access network team is to determine what revisions may be necessary for the path forward. We will note that the broadband network is but one of many potential traffic bottlenecks and addressing it in isolation will not necessarily eliminate all customer issues. Many of the mature applications identified within, that are shifting with the trends of the pandemic were, like the network, operating in a "business as usual" (BAU) pace of upgrades of features, capabilities, and scalability. The application servers themselves, and interfaces to and from them, also experienced an unforeseen step function increase in demand.

VPN servers, for example, perhaps accessed routinely by 15% of a company's remote workforce --with compute and I/O sized accordingly -- suddenly saw perhaps 4-5x the number of remote employees vying for access, slowing everyone down. A video conferencing application, for example, may be fully suitable for cross-country meetings from several remote offices, but ill-suited for 30 individual streams. Lastly, we have seen in the news the challenges the Zoom application has had with security of sessions, a feature important to all, but with very significant implications to business data compared to the virtual family video chat. So, as mentioned, while network engineers revisit their strategies and plan for evolving the network, there are many parallel efforts across the products and applications such as these, and the companies who own or operate them. The end performance for customers involves successful migration and execution on both fronts.

Lastly, we note that the focus will be on the edge and access network, which is generally the cable system bottleneck. But we should note that the in-home broadband network, in particular the multiple and simultaneous user Wi-Fi experience, may also come under duress as a bottleneck. The new generation of Wi-Fi 802.11ax (Wi-Fi 6), and the new 6 GHz band are arriving at just the right time to improve the in-home experience. In-home Wi-Fi mesh extenders have also matured to provide a significantly improved coverage experience. Depending on their current Internet service and gateways, customers may need to upgrade their in-home network.

5.1 The Near Future Has Arrived Ahead of Schedule

We have seen the drivers of new behaviors, and, to the extent that we can use the word "predictable" in a discussion based on an unfinished global pandemic, the network traffic aspects and key applications of pandemic life are, in fact, predictable – except perhaps that gamers evidently never EVER take breaks. Furthermore, we can recognize that the pandemic has not changed very much materially about what the important initiatives for advancing the access network are. All of them are things that we should have been thinking about already.

What has changed, however, is the execution timeline of these strategies, including the relative priorities and the proper sequence of initiatives. And, importantly, it has increased greatly the level of awareness of our customers, and the nation as a whole, to the criticalness of the network and availability. This plays a role in the approach to phasing, execution, and risk assessment, as initiatives are brought to trials and rolled out. In access network engineering, it has normally been the case that the job is well done when major network-impacting technology upgrades take place and the customer does not notice you were there, unless it's having a new product offering to consider, or because they cannot remember the last time they had to call to talk to a care agent on the phone. As the shelter-in-place orders came out and companies started dispatching their employees to work from home, ALL eyes were on the network. And broadband engineers everywhere should be very proud of what they built and how this situation has played out so far.

Armed with the trend data shown earlier in this paper, the recognition that the trends are likely not temporary, that precise accuracy of end-state traffic behavior will be uncertain for some time – yet the ability to surgically address rapidly changing dynamics has never been more essential – a recalibration of the network evolution makes sense. This should be done with the following considerations in mind:

1. *Enable Agile Coaxial Capacity* – Quickly add capacity, above and beyond "BAU" increases, necessary to recover the pre-COVID bandwidth-versus-time runway, and in anticipation of increasing US CAGR and that the observed trends so far get rooted beyond the first wave of adopters.
2. *Hardening and Resiliency* – Enhance the reliability and availability to enable a seamless WFH and online education experience for all types of remote workers and students. Increase visibility and proactivity to the customer experience. Create a resilient and self-healing network.
3. *Responsivity* – Drive down network-induced latency and jitter to support the increase in real-time applications, such as video conferencing, and to ensure the increased in gaming traffic does not impact the gamer experience.
4. *Enabling of Enterprise Services* – Perhaps this seems out-of-touch with the obvious softness in office-based business services. However, economic losses suffered due to shutdowns of industries dependent on "non-essential" workers are expected to lead a drive to digitization of processes and industry. This includes M2M communications, coupled with data analytics, in order to be less vulnerable in the future.

5.1.1 It's the Bandwidth, Stupid

Bandwidth, bandwidth, bandwidth. When is it ever NOT about the bandwidth when it comes to access network evolution? Well, this is NOT the exception.

One of the valuable lessons confirmed by the pandemic is the value of building ahead of the traffic curve. How far ahead is an ongoing debate and a target that shifts around many technical, business, competitive, market, financial dependencies. Generally it is good to have several years of network runway ahead at all times, simply because we KNOW traffic will continue to grow and because it is only every so often that the opportunity to touch the network takes place, and certainly not continuously. It is this built-ahead-of-the-curve discipline that allowed network operators to not just survive the pandemic storm but excel in it. The spike in traffic was unprecedented but did not exceed the amount of time "runway" operators routinely plan for.

Despite pre-planning, the magnitude of the spike did create some of the aforementioned housefires where, for example, a portion of the network may have had an upgrade cycle a number of years ago, and at the time of the onset of the pandemic was not far from its next upgrade. In a "BAU" cycle there was more than sufficient capacity to make it to the scheduled upgrade – typically a node split. Capacity management teams carefully project "by when" an upgrade must take place or else risk a utilization/congestion alarm, by monitoring the network traffic at all sites and for all service groups, and building dashboards to make ongoing decisions, as shown in Table 1.

Table 1 - Capacity Management Dashboards Guide Targeted Network Upgrade Investment [1]

	Month 1		Month 2		Month 3	
	% >80% Utilization	Count >80% Utilization	% >80% Utilization	Count >80% Utilization	% >80% Utilization	Count >80% Utilization
Site A	0.02%	2	0.04%	5	0.17%	22
Site B	0.05%	9	0.13%	22	0.49%	80
Site C	0.21%	33	0.43%	67	1.65%	251
Site D	0.09%	13	0.06%	9	0.33%	49
Site E	0.00%	0	0.02%	2	0.10%	11
Site F	0.44%	20	0.93%	42	1.47%	66

Legend - % of CMTS SGs
0.00%
<.05%
<.1%
<.5%
<1%
>1%

5.1.1.1 Rich Fiber Diet

Like all MSOs, Comcast continues to split nodes as traffic grows – it is a great advantage of the HFC architecture, that it is so incrementally scalable in a success-based way. However, several years ago, Comcast embarked on a network upgrade path that complements the typical node splits with a more aggressive approach that effectively acts as multiple node splits in a single upgrade. By design, this strategic plan will last for many years, in order to get ahead of the cycle of repeatedly going into the field in a matter of a couple of years. This is often referred to as a "Fiber Deep" or node plus zero ("N+0") architecture. Millions of households passed have been built in this manner since introduction of this adjustment to the evolution plan was introduced. Much has been written about this architecture [5,6], which is shown in Figure 12.

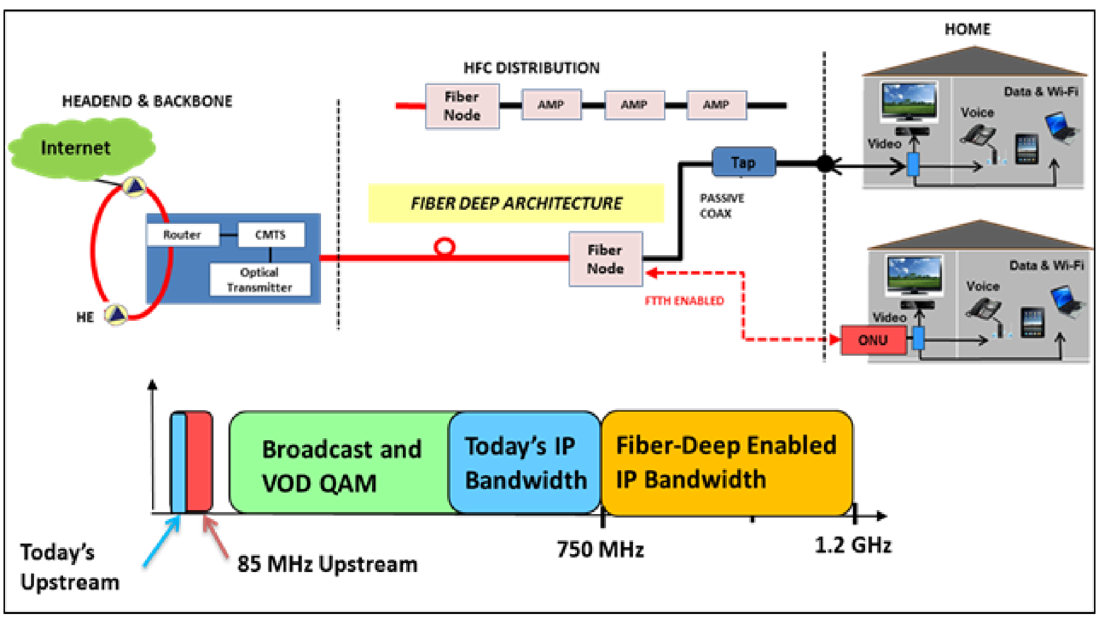

Figure 12 - N+0 Architecture Basics: No RF Amplifiers and Expanded DS and US BW

A fundamental premise of Fiber Deep that led to its adoption is that it could be shown in certain areas to provide not just a simpler, higher performing, more consistent, and generally better last mile architecture of increased bandwidth, but it was also more cost effective over time than the series of node splits that would otherwise be needed to manage the traffic, as shown in Figure 13. It is not shown in Figure 13 but the "breakeven" time frame between those two options was about 7 years.

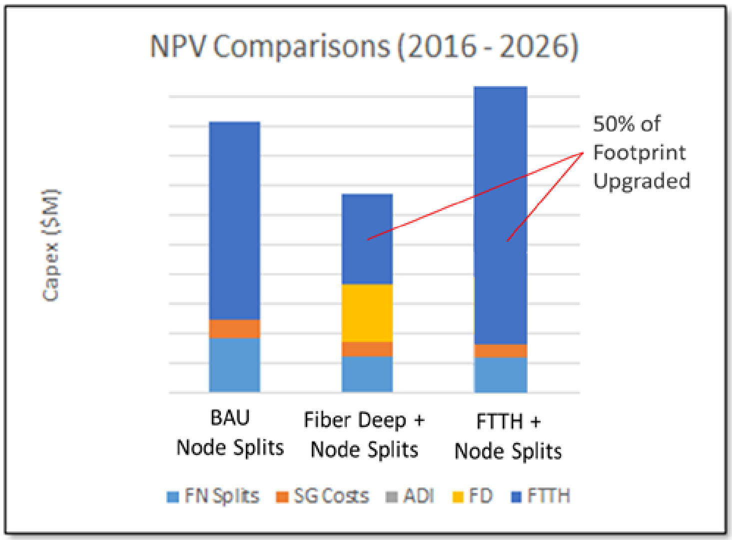

Figure 13 - Fiber Deep Cost Effectiveness Compared to Continued Node Splitting Only [6]

In summary, Fiber Deep is essentially the equivalent of executing multiple node splits at once, while expanding the spectrum in both directions. The architecture decreases the size of the node service group by typically 4-5x, while nearly tripling the US capacity and adding 60% more DS bandwidth.

Most notably, as the pandemic struck, utilization issues above thresholds marking them for augmentation were nearly non-existent in the footprint where this architecture was in place, owing to its proactively taken large bandwidth step. These networks will not require follow-up upgrades for many years, and thus were generally not under duress, despite the traffic spikes caused by COVID-19. From the CAGR and timeline perspectives, an architecture installed with a lifespan of 7-10 years holds up pretty well, even if a year, or 10-14% of it, is eliminated.

5.1.1.2 Widen the Lanes

The Fiber Deep architecture identified increases the available spectrum in both DS and US. The US expansion takes the available bandwidth from 37 MHz up to a limit of 80 MHz, also known as the Mid-Split architecture. 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 DS video out-of-band signals used by legacy QAM set-top boxes (STBs) widely deployed today. Per the earlier discussion, it is typically the US that drives network upgrade activity.

While Mid-Split is deployed in current N+0 networks at Comcast, it can also be deployed in N+X HFC networks, by installing different diplex filters in the RF amplifier cascade as well as the node. Because of the average per-user peak-busy-hour US is only in the hundreds of kbps range, US generally grows slower than DS, and the new US spectrum is much cleaner, the impact on network lifespan of mid-split is very powerful.

Figure 14 shows the time runway generated by the three options – node split, node split plus upgrade to Mid-Split, and finally N+0. While N+0 offers the longest runway of the three, it is clear that an N+X migration tied to a node split is also a very effective way to extend HFC lifespan to nearly 7 years. Even with a COVID-19 traffic hit, it is easily absorbed while maintaining a multi-year runway.

Another benefit of N+X with spectrum migration is that it addresses the COVID-19 inspired objective to “add capacity quickly”, when compared to N+0. With the COVID-19 spike eliminating months of lifespan margin, N+X upgrades bring more US bandwidth to the network quickly to reset the lifespan timeline, and also support the growth of applications and new service speeds. The naturally slower pace of deep fiber construction will leave areas unaided by this architecture for a period of time, and with the “time” erased due to the pandemic, alternatives such as drop-in HFC upgrades that are both fast and effective make a sensible COVID-19 remediation step. Having a diverse strategy, not one-size-fits-all, adds important flexibility to deal effectively with adjustments to COVID-19.

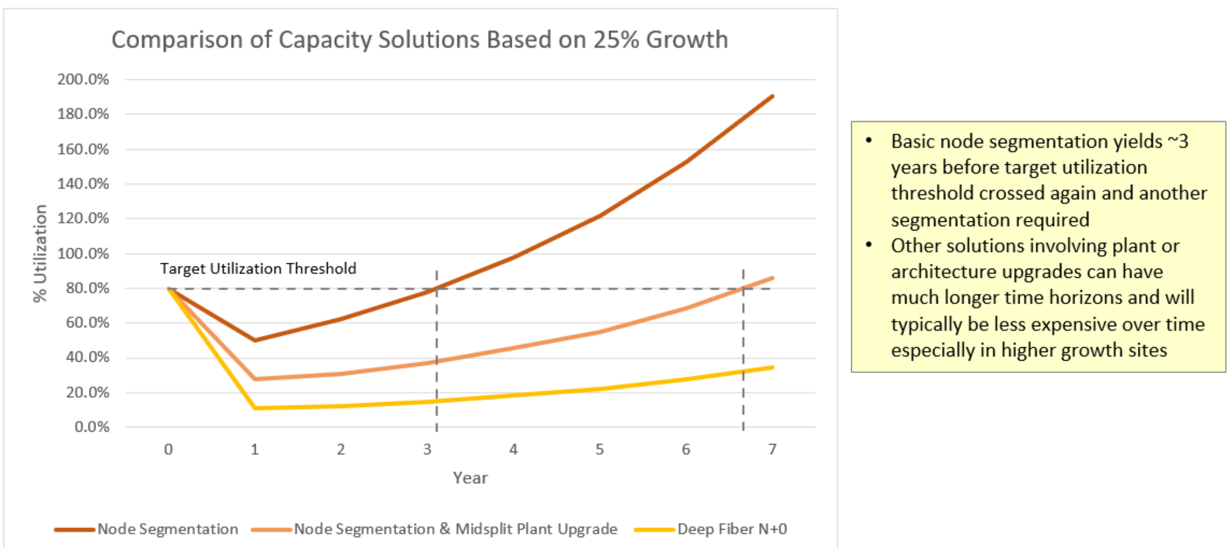

Figure 14 - Upstream Lifespan Expansion Options [1]

Lastly, with speed to increased bandwidth support in mind, coupled with the desire to push fiber deeper into the network whenever possible, adjustments to Fiber Deep architectures and practices are worth considering. For example, adding fiber in the underground network is an inherently slow process. However, given the freedom to allow a strategically placed amplifier (e.g. to allow an N+1 network), then lessening fiber construction, increasing the node size, and decreasing cost all speed the pace of the network upgrade and deliver the added bandwidth to more households passed (HHP)/year, if the long-term capacity benchmarks can still be obtained.

5.1.1.3 Every Bit Counts

Both Fiber Deep and upgrades to new frequency splits are plant operations, and as such require planning and coordination. This is because they carry regional dependencies associated with construction processes, permits, municipal intervention, etc. In addition, the ability to access some portions of the network during a pandemic itself is limited. Fortunately, there is another important tool in the DOCSIS toolbox that can address bandwidth directly and quickly and can provide a significant measure of relief. The tool is the DOCSIS 3.1 Profile Management Application, or PMA. It is closely related to the DOCSIS 3.1 feature known as Multiple Modulation Profiles, or MMPs. It fits the bill well as a quick relief valve to "recover" from COVID-19 traffic effects, while architectural plans of larger scope take longer to come together.

First, recall two of the major new changes between DOCSIS 3.0 and DOCSIS 3.1:

1) Single Carrier QAM (SC-QAM) was replaced with Orthogonal Frequency Division Multiplexing (OFDM) and Orthogonal Frequency Division Multiple Access (OFDMA). DOCSIS 3.1 recognized that it was the right time to break the mold and adopt more modern and mature OFDM. OFDM has the unique ability to optimize capacity across an uncertain channel in a way that SC-QAM cannot, at least not as easily. Instead of a 6 MHz chunk of spectrum to manage the signal over, the DOCSIS 3.1 flavor of OFDM uses granular subcarriers of 50 kHz or even 25 kHz in making up the composite signal. OFDM allows for narrow slices of spectrum to be optimized, and the overall channel to be used at its most efficient.
2) More QAM profiles and more efficient QAM modulation options were introduced, enabled by very powerful new Forward Error Correction (FEC). The base DOCSIS 3.0 of 256-QAM was increased to up to 4096-QAM, accounting for 50% more efficient use of bandwidth. All half-steps between 256-QAM and 4096-QAM were also enabled (512-QAM, 1024-QAM, 2048-QAM). QAM orders higher than 4096-QAM were added, but not as MUST requirements, and the efficiency bang for the dB-buck for these profiles is not very attractive.

Why is this so important after many successful years of 256-QAM? Figure 15 shows the recorded range of signal-to-noise ratio for millions of DOCSIS 3.0 cable modems, taken as part of the early work to develop the DOCSIS 3.1 standard. What it clearly shows is that, with the new FEC of DOCSIS 3.1 applied, many modems could achieve better than 256-QAM performance, and therefore many bits-per-second could be left on the table. How many bits are thrown away would vary depending on the particular modem.

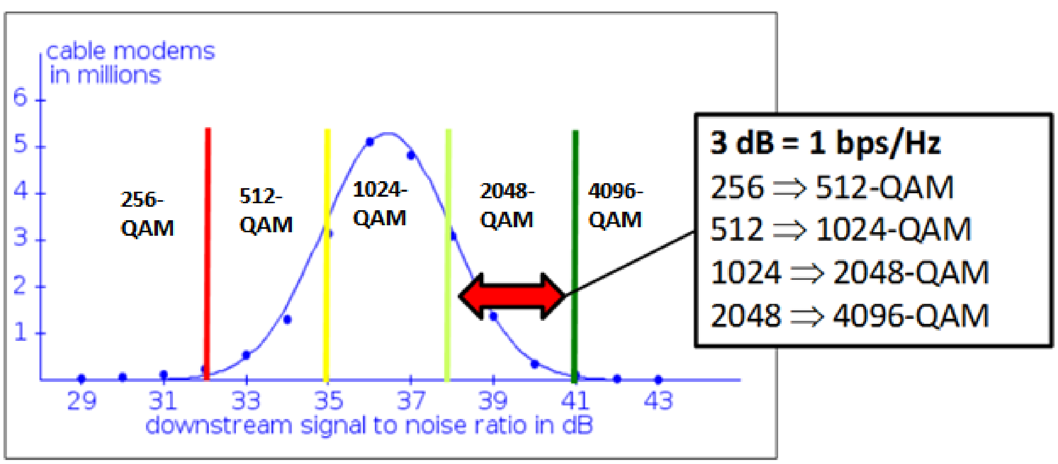

Figure 15 - Multiple Modulation Profile Potential of D3.1

Because of the characteristics observed in Figure 15, multiple modulation profiles (MMPs) were defined in the DOCSIS 3.1 spec, allowing clusters of modems seeing similar fidelity across the band to be formed into groups that operate with common QAM modulation orders on their subcarriers. In the DS, each set of modems clustered by a common profile configuration would have their time on the wire per profile, as shown in Figure 16.

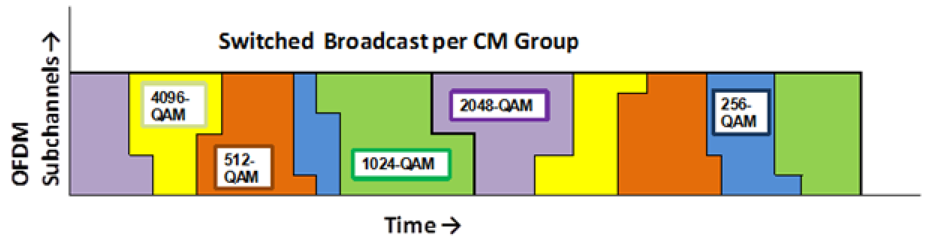

Figure 16 - DOCSIS 3.1 Multiple Modulation Profiles vs Time and Freq [8]

Figure 17 represents the workflow of fully capable, PMA-based capacity relief. It is coming to market in a very timely fashion to recover capacity taken away by COVID-19, and ensure that going forward the most efficient use of the channel is always available for the expected growth ahead, and furthermore executed as efficiently as possible with minimal human intervention. With MMP rather than a "flat" profile, for example, of all modems set at 1024-QAM, MMP fueled by PMA allows different modems to have a range of QAM profiles based on MER value. PMA brings an automated implementation, rather than thinking of MMP in a static sense or as a configuration set by the Headend Technician. PMA automates the process and allows for periodic checks and, if warranted, pushes updates to the modulation profiles without technician intervention. PMA will continually optimize algorithms as DOCSIS 3.1 continues to scale and big data metrics reveal more insights into temporal fidelity behavior on the coaxial plant.

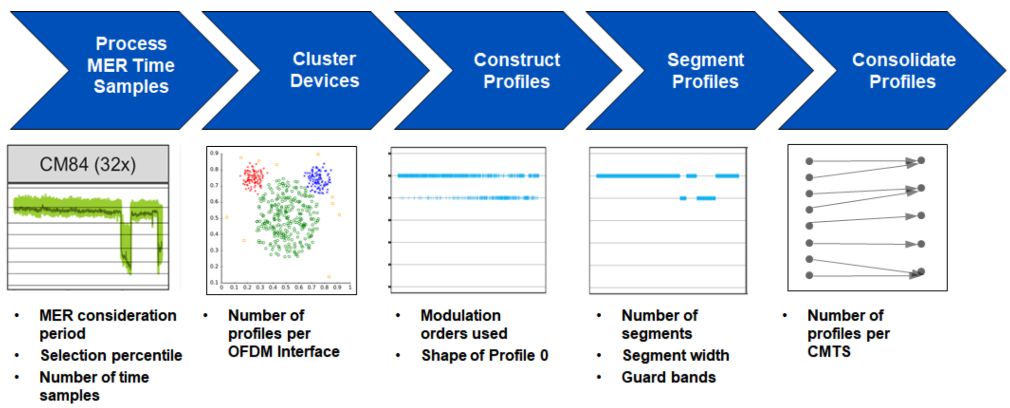

Figure 17 - Programmable Modulation Application Workflow [12]

While the example above was based on the DS, the DOCSIS 3.1 US uses OFDMA, which has the same fundamental signal structure as OFDM. Like the DOSCSIS 3.0 US, however, it must schedule and grant time-sliced access to the spectrum. Nearly all of the same logic of the DS applies with respect to optimizing the US channel in DOCSIS 3.1 However, most operators still have the majority of their network deployed with DOCSIS 3.0 modems, with DOCSIS 3.1 devices in the field supporting the higher DS speed tiers. Because of this larger percentage of D3.0 modems in the field, adding DOCSIS 3.1 in the US does not move the needle much to relieve capacity, since most devices cannot use it. Fortunately, with the US feeling more pressure from the COVID-19 traffic spike, and generally being more vulnerable, there are still DOCSIS 3.0 options available for capacity increases.

A common US configuration for operators is four bonded US 64-QAM carriers, each 6.4 MHz wide, which fills most of the high quality spectrum of a 5-42MHz US. However, there is additional spectrum, perhaps lesser in quality, that can be activated for more bps. Example use of the "5th and 6th" upstream in what had generally been empty spectrum is shown in Figure 18.

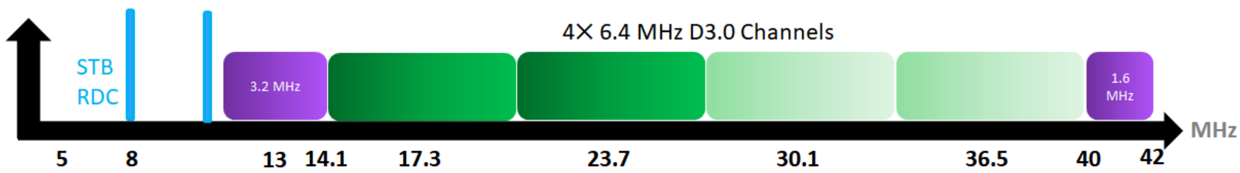

Figure 18 - Exploiting More DOCSIS 3.0 Upstream in 5-42 MHz Systems

Figure 18 suggests the presence of about 19% more useable spectrum, so in a 20% CAGR example would buy almost a year of time to work through a COVID-19-like spike. However, the spectrum available is likely to be impaired more than the existing 6.4MHz carriers, either due to filter roll-off and group delay variation at the high end, or low MER (noise aggregation) at the low end. So, the net capacity gain may be less than the spectrum gained, and the time bought by it shorter – but still meaningful for addressing a short-term crunch like COVID-19 while planning larger scale remediation steps.

Also, the DOCSIS 3.0 US offers several (four) settings for FEC overhead, where higher overhead means stronger FEC, and better protection for the packets on the DOCSIS 3.0 channels. DOCSIS 3.0 introduced 64-QAM about ten years ago, and, thinking in terms of in HFC network quality vs time, it was naturally set to maximum protection at launch, and rarely looked at since. However, since that time, node splits have shrunk service group sizes, and amplifier cascades have shortened. Initiatives that significantly enhance fidelity such as Fiber Deep and Distributed Access Architecture (DAA) are underway and in production at scale. As a result, some of those US channels, in particular the ones in the 25-40 MHz range, may no longer need maximum FEC protection. As such, there is a percentage of extra payload bits that can be gained in exchange for FEC overhead, depending on the quality of the highest MER US channels, from about 6% to, realistically, about 17% on a particular channel.

We will describe more about it in the following section, but modulation profile efficiency benefits considerably from the migration to DAA, which is being built in Fiber Deep markets and moving into N+X configurations also. The 10GbE digital fiber connectivity used in DAA significantly enhanced DS fidelity, by removing the analog optical links that typically had set the MER limit. It also enables higher-fidelity US signals to be present at the digital receiver. These higher MERs translate to more bandwidth efficient QAM signals.

In a situation like COVID-19, where a rapid response is essential, prioritizing software-based network enhancements, like PMA for capacity relief, can be executed and deployed generally more quickly than construction projects. They become exceptionally valuable tools for low-touch capacity augmentations, well suited to the kind of immediate response needs produced by the pandemic.

5.1.2 Be There for Me Always

One of the key principles of the Fiber Deep architecture described is the simplification of the network. It reduces the number of active devices in the plant, removes the oldest actives in the plant, and provides a repeatable template for a last mile architecture that breeds consistency of process, equipment, RF level and fidelity performance, and aligns with expectations of future network applications and services.

Figure 19 shows a large sample of CMs taken across N+0 deployments. When compared to Figure 15, it is clear that the average MER has increased by 3-4 dB. Note that MER is slightly different than SNR, in that MER considers not just the noise degradation added to the signal, which SNR technically represents, but all of the other impairments as well. However, they have, unfortunately, come to be used interchangeably in the industry. Also, the bell-shaped curve seen previously is now favorably skewed towards higher MERs. In fact, the red arrow identifies the approximate average MER in the Figure 15 BAU N+X case when compared to the N+0 network. With the combination of significantly higher fidelity and the PMA tool, the network is well positioned to provide a higher quality, more consistent customer experience and effectively serve the increased volume of traffic the pandemic has wrought.

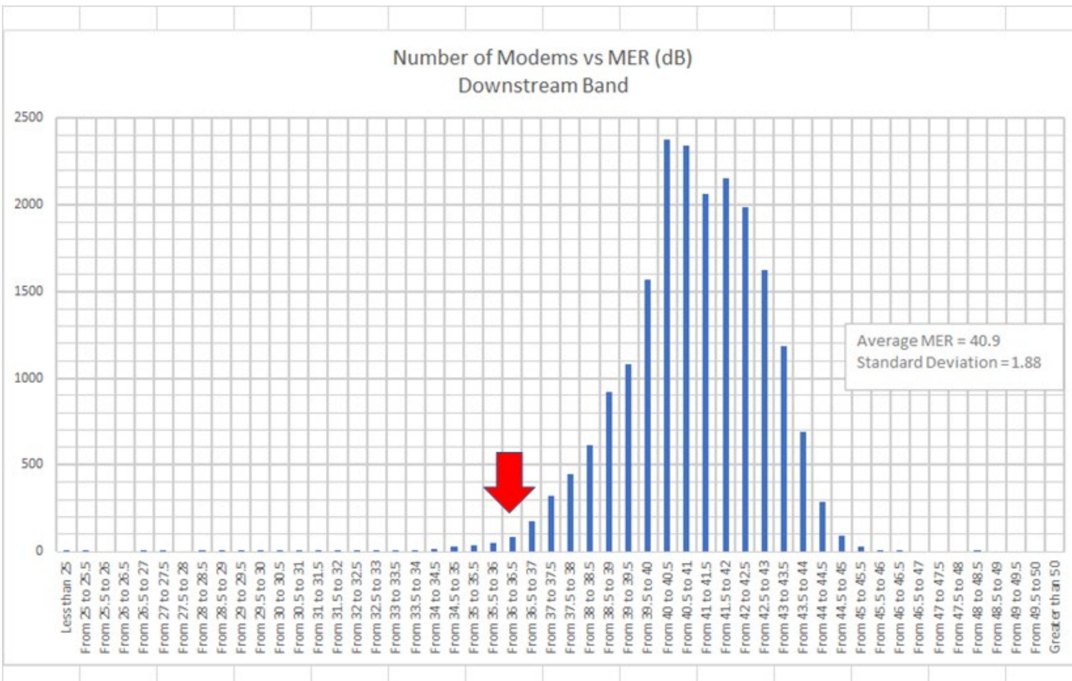

Figure 19 - MER Distribution for a Large Sample of CMs on a N+0 System

Indeed, trouble calls run consistently lower on the Fiber Deep parts of the network – ranging from 10-15% for the calendar year 2019, confirming the expectation that it is not only a superior network, it is also more cost effective to operate, and causes less disruption for busy WFH employees and online students.

Figure 20 shows a characteristic view of trouble call (TC) tracking in one Comcast Division that was very actively building N+0 in 2019. Note that the onset of COVID-19 during the initial "shock and awe" of shelter-in-place restrictions froze everything in place for a period of time (3/26/20) as companies and customers focused on getting a wide range of new priorities and changes in their day-to-day.

Not very long after – broadband being among the highest priorities – a return to normal TC interaction was taking place – but note this return to "normal" is with significantly increased COVID-19 traffic, now more essential traffic, and high awareness to network behavior among customers in the early days of the pandemic response.

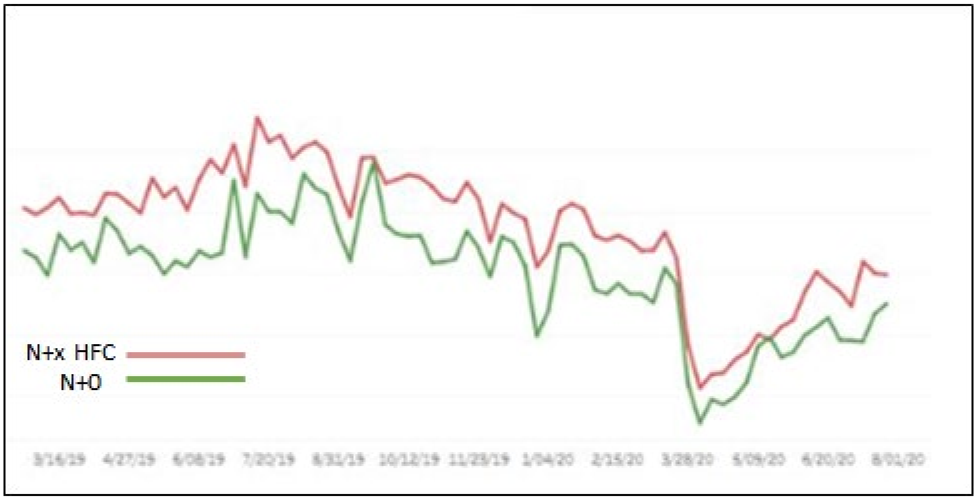

Figure 20 - Reduction of Trouble Calls in Fiber Deep Deployments

In addition to prioritizing dense fiber access as an architecture, Comcast has pivoted to deploying Distributed Access Architecture (DAA) at scale, based on Remote PHY technology, using Remote PHY Device (RPD) Nodes from multiple established vendors. In conjunction with the transition to DAA, we launched a Virtual CMTS (vCMTS) core, leveraging 10 GbE connectivity from facilities to the RPD nodes in the field. The vCMTS is built as a containerized set of micro-CMTS cores, such that each RPD node has its own packet processing compute and is arrayed in a robust redundancy scheme with a very small blast radius, increasing availability and redundancy to guarantee the always-on WFH experience. The 10GbE connectivity sets the network up well to be last-mile access agnostic by implementing a range of last mile access technologies fed by simple Ethernet links deep into the plant.

The introduction of DAA and vCMTS delivered orders of magnitude more real time visibility into the network via streaming telemetry, and with it the ability to process advanced metrics and use machine learning in operations. The increased visibility is essential to deliver the proactivity critical to network availability, enabling tools and views to quickly assess and make determinations about anomalous network behavior, raise alarms, trigger self-healing and re-routing as necessary, and capture detailed logs to continue to learn about and optimize network robustness and flexibility in modern ways. Proactive visibility and availability are, of course, among those high priority areas of focus in support of the WFH and online education trends now and ahead.

A sample dashboard summarizing the network status, with visibility on the order of seconds, is shown in Figure 21, encompassing information on CMs, RPDs, switches and routers. Virtualization in operator access networks leverages years of development and maturation in large scale data centers, driven by the likes of Amazon, Google, Facebook, etc. These providers, of course, have very stringent requirements for availability and reliability of their services, much like network operators do. While there are some important differences between operating a large data center that is handling "horizontal" data flows, compared to operating a network directly serving customers, the technological roots and the ecosystem scale created in establishing these data centers provides much of the foundation for virtualizing the access network edge.

Figure 21 - Real-time Dashboard of Network and Component Availability and Uptime

Another benefit of the virtualized core is its ability to throw packet processing compute in the direction where it is needed most, as shown in Figure 22. The pre-pandemic common example was daytime business services traffic in one part of the network geography, shifting to distributed residential web traffic during peak busy hour of the evening at home. In other areas of the footprint, such as universities and in apartment complexes tuned to a demographic, studying, gaming, chatting, or streaming late into the night is more the norm. The nimbleness of a virtual cloud could fall under the category "agile capacity" noted in the prior section. It landed here to separate this technology directional shift towards DAA and virtualization from the RF capacity focus of the prior section.

Enterprise services for office complexes are obviously not viewed as a segment that is expected to see a return to growth imminently. However, a network looking ahead should be considering the changes to business services and what might be their renewed directions of focus and investment. As implied in the example of Figure 22, business services traffic is a growing and increasingly important part of a cable operator's services and revenue. Business for mid-size and large commercial enterprises often have specific, quantified, requirements for throughput, latency, jitter, packet loss, and availability. While the coaxial network is optimized for residential use, it does support a wide range of small-to-medium businesses (SMBs) with DOCSIS-based cable modems, typically with specific software loads that include certain features important for businesses, but with more limited Service Level Agreement (SLAs) than enterprise class services. And, while the coaxial system is undergoing development to improve availability and resiliency, as well as latency and jitter (to be discussed) – all key aspects of post-COVID applications – larger scale enterprise services are still better served by fiber-based Ethernet architectures.

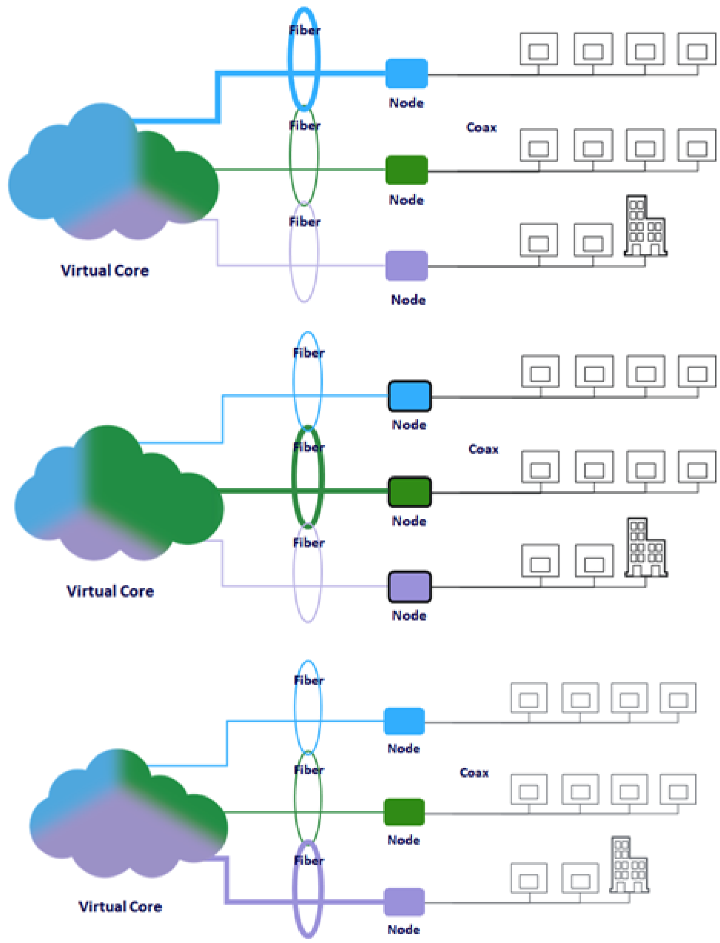

Figure 22 - The Virtual Cloud Efficiently Shifts Capacity Where it is Needed

A further key principle behind the Fiber Deep architecture is that it is an investment in a platform that supports the growth in business services, many of which are fiber-based. With fiber reaching to less than 1000 feet of every coaxial connection, there will be increasing access to fiber for nearby business services. Furthermore, the migration to a DAA architecture means that this fiber infrastructure is becoming digital fiber based on 10 GbE. As the residential network evolution takes place, a more scalable, resilient, high visibility network is being put in place. On it, Layer 2 Ethernet services can share the wavelength space, the latter of which is multiplied substantially when going from traditional analog fiber optics, and 16 wavelength limitations, to the DWDM grid supporting up to 80 wavelengths. This 5x bandwidth multiple into these dense pockets has the network well-prepared for post-pandemic applications and traffic growth, such as commercial opportunities in industrial automation, M2M, and digitization applications tied to the COVID consequences on the manufacturing economy.

5.1.3 Be Snappy

Latency and jitter have historically applied primarily to voice services, before becoming dominated by the increasingly popular gaming segment. Gamers are no longer kids hiding in their parent's basements. They are professionals who still love their games, and the games have become ever more elaborate and globally interactive. There have been many studies highlighting the sensitivity of games of certain types to network performance. Games generally do not generate high traffic volumes, but the customer experience can be significantly impacted if the packets are held up excessively between the gamer and the game server (or host). The bandwidth statement above does not necessarily apply to the category of cloud gaming or Augmented Reality/Virtual Reality (AR/VR) scenarios.

Of course, yesterday's voice sensitivities to real-time attributes is today's video conference. A major difference is that while there remains a sensitivity to latency and jitter, the bandwidth of these applications is much higher because they are video streams, and in a "Teams" meeting, for example, there can be a streaming video for each attendee using a camera. So, while capacity, bandwidth, and speed always dominate the discussion of access network evolution, latency is as much a factor if not more to the customer experience, and certainly becomes magnified in the pandemic-driven video conference world. Speed represents the amount of data that can be sent across a connection over a fixed period of time, whereas latency is the time it takes for a packet to cross the network and get a response. It is the "ping" time.

Seeing latency as an emerging priority for real time services such as gaming, video conferencing, IoT M2M, etc., working to address the components of the latency budget attributed to the DOCSIS network is an industry priority. A program was initiated at CableLabs to standardize the approach across the technology ecosystem. Obviously, the added impetus given by WFH has only increased the interest in developing Low Latency DOCSIS (LLD). In the DOCSIS 3.1 network, the dominant causes of latency are queuing and media acquisition [13], as shown in Figure 23.

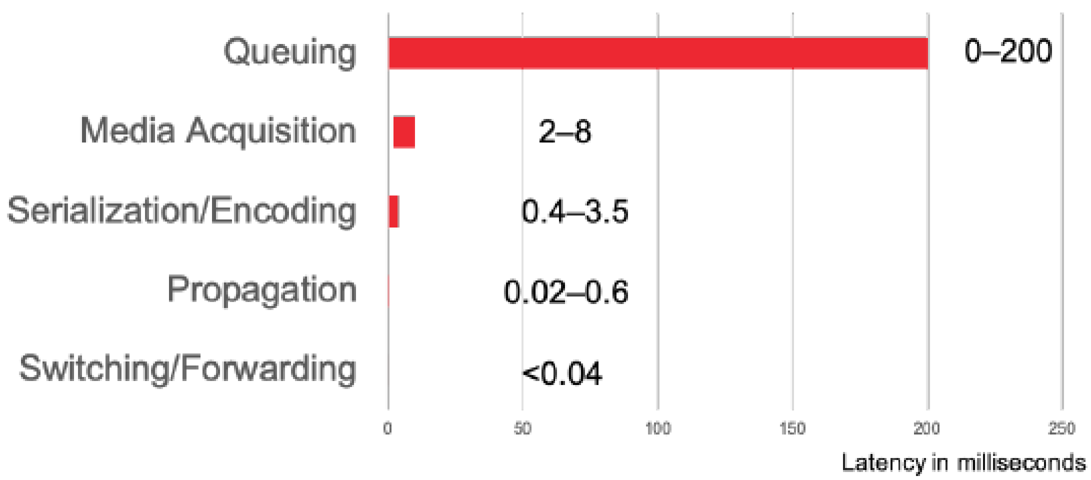

Figure 23 - Sources of Latency in DOCSIS 3.1 Networks [13]

The lowest hanging fruit is in the area of queuing. Once controlled, there is a potential to drive latency into single-digit milliseconds in the "media acquisition" area. The term "media acquisition" in this case refers to the request-grant mechanism by which a CMTS makes determinations about access to the upstream channel time slices and provides that information to the modems. Low millisecond values are instantaneous in "human" time and open the door to future IoT and machine-to-machine communications where millisecond and sub-millisecond delays can have meaningful impact, for example, in stock trading or industrial machinery.

The potential for major improvements by managing queueing more effectively is clear from Figure 23. These queueing drivers are primarily the nature of the Transmission Control Protocol (TCP), the Layer 4 mechanism used throughout the Internet. High bandwidth applications ramp to a rate faster than the network bottleneck can support, and rely on TCP congestion control, which inform the sender to back off when buffers holding the traffic bursts fill, at which point it begins to drop packets to manage the queue. Because of this, latency and packet loss are inherent in TCP links. LLD aims to improve upon this behavior.

LLD does not create a fast path and a slow path or prioritize traffic one way or another. The approach is to identify traffic characteristics in applications that contribute to building of the queues which create latency, and steer traffic to a second queue if it is a so-called Non-Queue-building (NQB) traffic. This "dual queue" approach, shown in Figure 24, shares the bandwidth allocation to the user amongst all of their active applications, and also means less packets dropped out of the queue-building buffer. It is easily enabled using DOCSIS Service Flow (SF) definitions, available today, that shape the aggregate traffic, and the individual flows defined by the dual queues have shaping turned off, avoiding prioritization. This approach improves the customer experience for applications that are latency sensitive without affecting the experience of the other services.

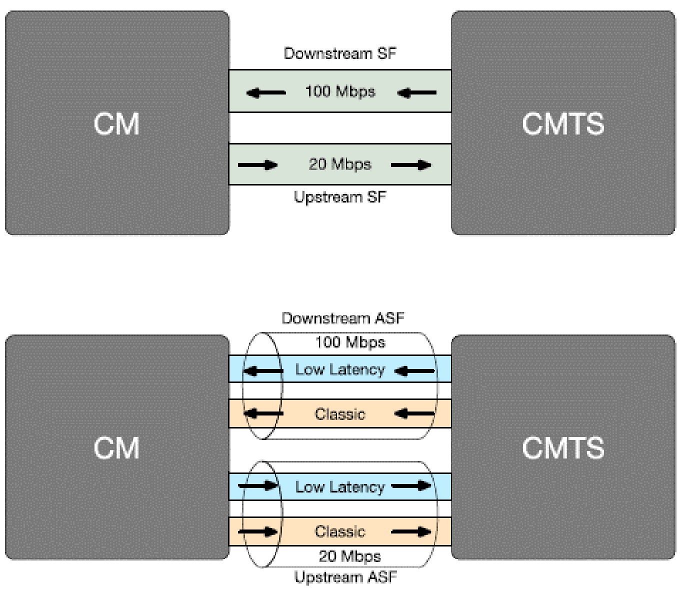

Figure 24 - Traditional and LLD Service Flow Architecture [13]

In addition, statistics are generated on the queueing delay by the CM and CMTS for use by MSOs, enabling them to fully understand the contribution of the DOCSIS network and help itemize their end-to-end latency budget. The customer experience can be closely correlated to these latency metrics. We already have tools in place that observe network latency further northbound in the network, and use this data, filtered and correlated across devices and bonding group assignments, to make proactive determinations to drive optimizations and customer QoE.

CableLabs is engaged with the IETF and the broader industry to identify the proper queue-using classifiers within the fields of IPv4 and IPv6 traffic. Most of today's traffic volume is queue-building

(QB), in particular streaming video. These services transmit packets faster than the network bottleneck can support, relying on the TCP algorithms to buffer and alert the source to adjust accordingly. By contrast, NQB traffic sources generally do not send data faster than the network can directly support. Skype (and similar) and online games generally fall into this category, as does basic web browsing. Despite their inherent "good behavior" relative to utilization of network resources, their packets are unnecessarily mingled into the queue created because of the QB applications today.

The LLD effort includes additional options to further reduce latency, in particular to speed up the DOCSIS request and grant cycle of scheduling packet transmission US. These include shortening the window between request options and proactively granting access based on traffic trends and utilization, which can be learned over time, such that inefficiency is minimized. Figure 25 shows the potential for reducing the DOCSIS network latency as these features become available over time.

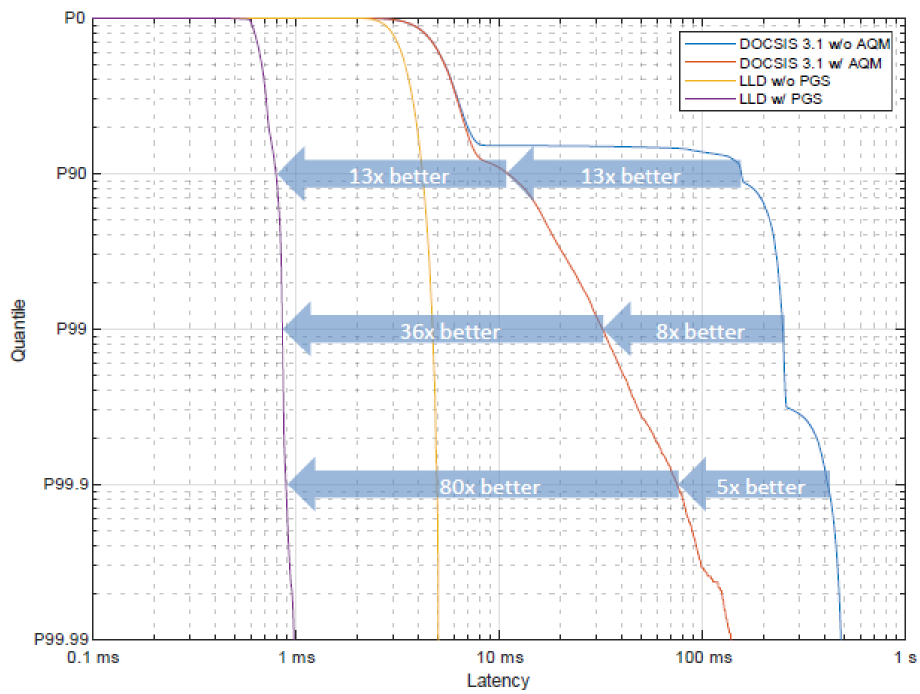

Figure 25 - Industry LLD Efforts are Promising for Supporting Real Time Services Such as the Recent Explosion of Video Conferencing Due to WFH [13]

As with other access network initiatives, existing technology development underway is aligned very well with addressing the impacts of the pandemic. The immediate WFH thrust and online schooling is certainly a reason MSOs have cause to consider adjusting priorities to move LLD higher on the list of To-Do's. These performance improvements also align well to broadband requirements that are expected to be part of industrial digitization.

6.0 Conclusion

One way to look at the COVID-19 situation as it applies to the network is to acknowledge two things that are somewhat in opposition to one another:

1) We have never before seen a traffic dynamic like this – the magnitude of the increase plus the abbreviated period over which it occurred – in our broadband cable lifetime
2) Simultaneously, there is little new that needs to be learned to figure out how to handle it.

We identified four areas of priority from which to re-calibrate access evolution plans, given the observed trends and anticipated permanent impacts on the network attributable to COVID-19:

Enable Agile Coaxial Capacity – Take advantage of existing DOCSIS 3.0 platforms to add new channels, optimize DOCSIS 3.0 US profiles, and DOCSIS 3.1 PMA to add capacity to recover the "time" lost by the COVID-19 traffic wave. These techniques complement more time-consuming operator network upgrade plans, such as changing the frequency split in the network by swapping actives, splitting nodes, and pulling fiber deeper.

Hardening and Resiliency – Synergistic with pulling fiber deeper, adding more nodes, and removing the oldest actives in the plant – to reduce trouble calls and increase availability – is a transition to DAA and virtualization of the CMTS core. Granular streaming telemetry keeps an eye on the network more closely and in real time than ever before, data analytics trigger proactive alerts and notifications, and the micro-CMTS architecture further maximizes availability and adds resiliency.

Responsivity – Drive down network-induced latency and jitter to support the increase in real-time applications, such as video conferencing and online classrooms, ensure the increased in gaming traffic does not impact the gamer experience, and to support new opportunities that arise in business services that have real-time and near-real time broadband implications.

Enabling of Enterprise Services – Change often leads to opportunity, and major change, on the COVID-size scale, will likely lead to a range of new industrial opportunities aimed and helping companies be prepared to "weather the storm" in the future. What may at first be a softness in business services opportunities is likely to take a favorable turn in the future, with perhaps a change in the types of business markets served. The fiber-rich, Ethernet-based DAA system has the flexibility to be well-suited to deliver a complete range of last mile solutions and capabilities that may be uncertain today.

If there was ever a scenario to be caught by surprise and in a desperate scramble to recover, this would have been that time. While there was indeed some scrambling, the broadband cable network gets an "A+" for its performance in the face of this unprecedented traffic wave. The reasons for this are many, and boil down to a few key points

- Planning several years ahead is the nature of the job for access network architects and capacity managers, and it paid off in this crisis
- Very familiar muscles can be exercised efficiently to deliver additional capacity, built upon many years of experience honing these processes
- Additional network technology was already in the works that was aimed at addressing the needs observed during the pandemic and what is expected to be ongoing beyond it

- Efficient-to-deploy, SW-based new technologies, such as PMA and LLD, are emerging at a very timely moment, because the benefits they bring had already been deemed important for future applications that happened to arrive at scale a bit earlier
- For specific applications of increased priority, the scale and penetration pace of them changed quickly, but there are not additional technology barrier dependencies
- Cable operators have been highly focused on the upstream very recently, because bandwidth from the home is close to filling the availability capacity
- Cable operators have been highly focused on network monitoring and availability as the cable product offering has shifted towards always-on Internet access as a top priority

Let’s hope this is the last pandemic any of us witness in our lifetimes. However, having passed the network stress test of a generation, we can all be comforted by the recognition that the power of the broadband network is there to ensure that some semblance of “life goes on” is still achievable. And undoubtedly, there will be lessons learned that can make us even better prepared should such a situation strike again.

Abbreviations

AR/VR	Augmented Reality/Virtual Reality
BAU	Business as Usual
CAGR	Compounded Annual Growth Rate
CM	Cable Modem
CMGR	Compounded Monthly Growth Rate
CMTS	Cable Modem Termination System
DAA	Distributed Access Architecture
DS	Downstream
ISBE	International Society of Broadband Experts
SCTE	Society of Cable Telecommunications Engineers
FEC	Forward Error Correction
HHP	Households Passed
HSD	High Speed Data
IETF	Internet Engineering Task Force
loT	Internet of Things
LLD	Low Latency DOCSIS
MER	Modulation Error Ratio
M2M	Machine to Machine
MMP	Multiple Modulation Profiles
MSO	Multiple Systems Operator
N+0	Node Plus Zero Amplifiers
N+X	Node Plus "X" Amplifiers
NQB	Non-Queue-Building
PBH	Peak Busy Hour
PMA	Profile Management Application
QB	Queue Building
OEMs	Original Equipment Manufacturer
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
QoE	Quality of Experience
RPD	Remote Phy Device
SC-QAM	Single Carrier QAM
SF	Service Flow
SLA	Service Level Agreement
SMB	Small-to-medium business
TC	Trouble Call
TCP	Transmission Control Protocol
TPD	Traffic Doubling Period
US	Upstream
vCMTS	Virtual CMTS
VPN	Virtual Private Network
WFH	Work from Home

Bibliography & References

[1] Barker, Bruce E, and Claude Bou Abboud, Erik Neeld, "Access Capacity Planning: Staying Well Ahead of Customer Demand Helped Ensure Stability During COVID-19," SCTE Cable-Tec Expo, Oct 13-16.

[2] Bienkowski, Tom, "Pandemic Life Online: What is Everybody Doing," www.netscout.com

[3] Bienkowski, Tom, "COVID-19 Network Traffic Patterns: A Worldwide Perspective from Our Customers," www.netscout.com

[4] Branscombe, Mary, "The Network Impact of the Global COVID-19 Pandemic," https://thenewstack.io

[5] Howald, Dr. Robert L, Comcast, "The Fiber Frontier," INTX Spring Technical Forum, May 16-18, 2016, Boston, MA.

[6] Howald, Dr. Robert L, Comcast, "Aboard the Technology Wave: Surf Report," SCTE Cable-Tec Expo, Sept 26-29, 2016 Philadelphia, PA.

[7] Howald, Robert L, and Dr. Sebnem Ozer, Robert Thompson, Saif Rahman, Dr. Richard Prodan, Jorge Salinger, "What is 10G – The Technology Foundation," SCTE Cable-Tec Expo, Sept 30-Oct 3, 2019, New Orleans, LA.

[8] Howald, Robert L, , "Network Preparation: Maximizing Capacity ROI," SCTE Cable-Tec Expo, Oct 21-24, 2013, Atlanta, GA.

[9] Koeze, Ella and Nathaniel Popper, "The Virus Changed the Way We Internet," www.nytimes.com

[10] Liberatore, Stacy, "Coronavirus has changed the way American use the Internet," Dailymail.com.uk

[11] Moritz, Scott, "Empty Offices, Full Homes: COVID-19 Might Strain the Internet," www.bloomberg.com

[12] Rice, Daniel, "DOCSIS 3.1 Operational Hardening," CableLabs Summer Conference, August 4, 2019, Keystone, CO.

[13] White, Greg, and Karthik Sundaresan, Bob Briscoe, "Low Latency DOCSIS: Technology Overview," https://cablela.bs/low-latency-docsis-technology-overview-february-2019.

[14] COVID-19: How Cable's Internet Networks Are Performing: Metrics, Trends & Observations, https://www.ncta.com/COVIDdashboard

[15] (https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/cases-in-us.html)

Acknowledgments

The author would like to thank Brian O'Neill of CommScope for some enlightening discussion and insights regarding the view on COVID-19 impacts from the OEM perspective.

The author would also like to thank his indispensable administration professional, Karen Murray, for taking the lead on researching publications regarding the pandemic's effect on networks.