THE FIBER FRONTIER

Dr. Robert L. Howald
Vice President, Network Architecture
Comcast

Abstract

There is a natural conflict between the pace of technology change and the practical pace of change that can occur across the network in a large scale operation, such as that of an MSO. Because of this, it is always the case that there will be significant uncertainties to weigh when making network investment decisions.

A commonly encountered uncertainty is assessing the merits and timing of an investment path associated with development and mass deployment of a new technology. In particular, the question that must rightfully be asked at each technology turn is whether one investment path that appears sensible is instead more wisely passed over in favor of a succeeding emerging technology path. This reflection and assessment must continually be made even should the former path be "plan of record," and even if it is already underway. With the pace of technology change accelerating and in many directions, the timing of new technology introduction is never perfect. It is rarely even good.

There is no simple solution to this puzzle, only a logical "best effort" path that requires operators to increasingly understand that business plans are living documents, and sometimes short-lived at that. Investment decisions should place a premium on flexibility and open standards, and be based on a vision of a long term end state to serve as a guidepost.

There are many example of the above dynamic for MSOs. One of the more timely examples this paper will focus on is fiber evolution for Cable Operators. Operators have been putting more and more fiber into their networks for years. This will surely continue. The question that the industry is grappling with more so than ever is the nature, scale and appropriate path forward for fiber usage beyond its long-standing HFC and business services roles, and the implications this has for operators given the aforementioned imbalance of rate of change between technology and network evolution.

This paper will take a look at how the various stakeholders of one major US operator defined the network evolution problem statement with a long term quantified perspective, determined a vision forward, and developed a plan to deliver on this vision. We will discuss and enumerate some of the assumptions, levers, compromises, platforms, and customer objectives driving the recommendations. We will consider alternative scenarios and reflect on the deviations this may cause in a long-term plan. Finally, we will postulate long term states of services and architecture, and discuss heavy-hitting variables and their potential impact. In short, we will provide an example template for the Fiber Frontier, with an eye towards investment criteria, practical change, and customer experiences.

PROBLEM STATEMENT: PERSISTENTLY AGGRESSIVE CAGR

Part of being a network operator in the age of the "lifeline" service known as broadband Internet access is having an obsessive attention to capacity management. As is commonly referenced in nearly every network architecture paper, the growth of consumer bandwidth has been persistent and aggressive since the introduction of the Internet, and particularly so in the downstream-to-the-home direction. This all makes logical sense as the technology quickly advanced to support all of our historical communication and media types – text characters to web pages to pictures to video and HD video, massive storage and file transfer, and now towards emerging Internet of Things (IoT) applications. While predicting the next "big" thing to drive Compound Annual Growth Rate (CAGR), something continually emerges to keep the CAGR train moving forward. A simple and sound rule of thumb for planning purposes is to assume this persistently aggressive CAGR charges forth at 50% per year. This means a doubling of traffic roughly every 21 months.

Figure 1 is a cable-specific look at how a 50% growth rate and a starting point of an existing capacity deployment reflect on cable system thresholds of capacity over time.

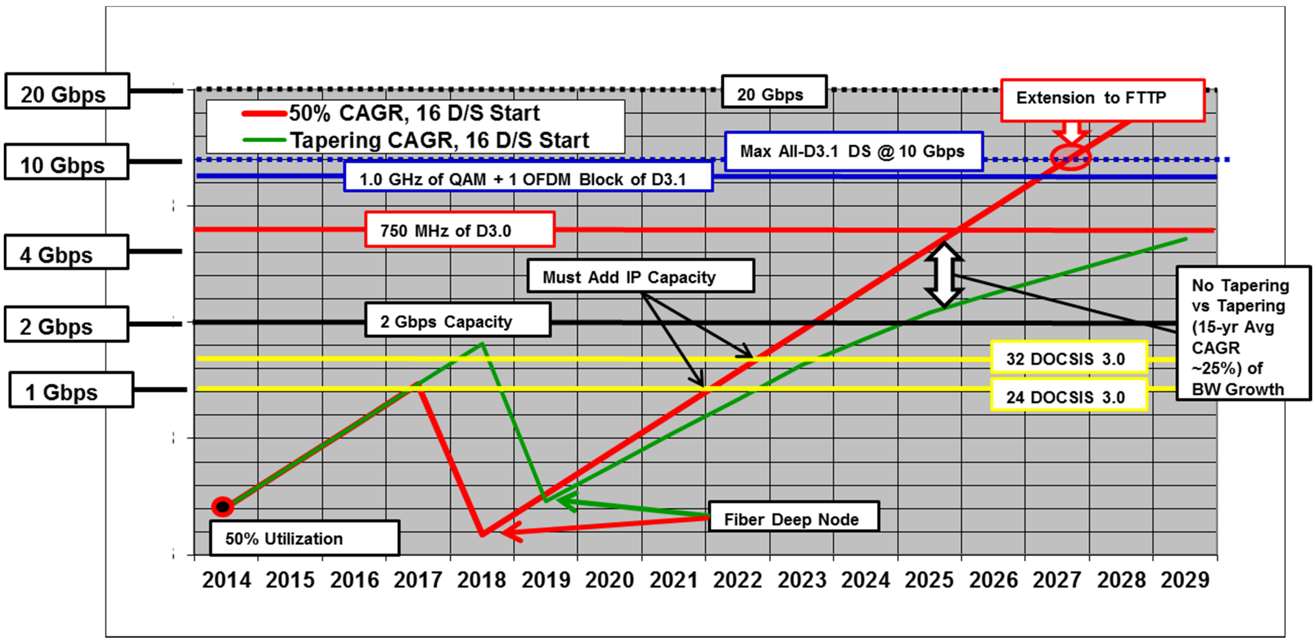

Figure 1 – Capacity Management Timeline view of HFC Long-Term Capacity [6]

Cable operators utilize this form of Capacity Management Timeline to understand the implications of growth on their network and the effects that various knobs and levers at their disposal have in keeping a healthy network capacity margin. This ensure a high quality of experience for customers.

Figure 1 is complex-looking but is in fact relatively self-explanatory when broken down into its descriptive parts. It can be considered in five understandable components that taken together tell an illustrative story:

1) The first year depicted on the chart – 2014 – is when Comcast was widely deployed with 16 downstream DOCSIS 3.0, which is approximately 600 Mbps of capacity for the IP network across a downstream service group. The utilization assumed of this allocation is 50%. This point is the black dot circled in red at the bottom left corner of the curve, and is expressed in expressed in decibels (approximately 25 dB).

2) Capacity is expressed logarithmically (dB) because traffic growth is exponential. The result is that the 50% CAGR then displays as a linear slope upwards. This makes long-term analysis very straightforward to visualize and threshold comparisons – the most important component of this type of tool – much clearer.

3) It was also in 2014 that Comcast began migrating the HFC network purposefully to a "Fiber Deep" architecture [6]. This will be detailed later in this paper, but, by definition, this means that no amplifiers exist after the HFC node. This was foreseen and proposed as the optimal MSO path for managing long term growth, for the reasons to be identified in 2009 [5].

A Fiber Deep migration will take place over time across the national footprint. Fundamentally, it is a more aggressive variant of the natural node splitting cycle that takes place today, and is based on priorities associated with relative network congestion, nature of the competition, status of existing infrastructure, nature of the regulatory environment, and several other criteria.

In Figure 1, this Fiber Deep (FD) phase happens in 2018 (red) or 2019 (green). This is the explanation for the step downward – the service group size is reduced considerably by Fiber Deep, increasing the average bandwidth available per subscriber and "resetting" the growth trajectory to this new reality. In practice, not all of this serving group reduction may take place at once (there may be HE aggregation at first), but a fundamental principle of FD is a last touch of the HFC network for as long as it is reasonable to consider "last," with any additional disaggregation – including all the way to 1:1 service group-to-node – easily applied at the HE when needed.

4) After this step, the traffic growth continues at the same slope. In the case of the green curve, some tapering of the CAGR is shown to understand the very powerful impacts if this were shown to be a trend in the future. The working assumption at Comcast is not to bet on the Internet slowing down, especially given the historical trend which has seen the aforementioned growth for roughly two decades.

5) Lastly, there are horizontal lines on Figure 2 that represent various thresholds of capacity the network can support depending on the total spectrum available, the percentage of it allocated to IP data, and the technology utilized (such as DOCSIS 3.1). When a red or green curve crosses a threshold, that architecture is no longer sufficient to support the capacity demand and new steps must be taken.

Note that capacity and data speed are not synonymous. Operators must overprovision their network capacity relative to the highest speed tier offered because the network resources are shared. After years of providing HSD service, operators understand the relationships between capacity and speed very well, at least under the existing deployment scenarios. Speed growth is predictable in the sense of expectation, but not necessarily with respect to timing, as the phenomenon is based on more competitive market dynamics rather than customer demand.

Growth Management and Today's Toolkit

Through nearly 20 years of HSD service and the associated bandwidth growth, cable operators have managed growth effectively and efficiently. There is a simple explanation why. When HSD service was launched, capacity and speeds to homes were set by the capability of the telephone wires, which was, and still is, very limited. The cable medium to the home provided orders of magnitude more IP capacity that MSOs were able to take advantage of, which has led to cable broadband dominating the broadband access business to this day.

The tools most frequently deployed to manage the growth were then and still are the addition of new DOCSIS channels to the CATV spectrum, and the adding of HSD ports at the Headend to reduce the sharing of resources into smaller pockets of users. Ports being added in the context of the HFC plant correlates with node splitting – a fiber node that once served, say, 1000 homes, would be segmented, or the network augmented with a new node, such that roughly 500 homes were served from each, and two 500-home groups each terminated at a CMTS port. This "BAU" process has been very effective, is mature and well-understood.

Fast forward 20 years, and as MSOs do their obsessive and long-term capacity accounting, it has increasingly become clear through analysis such as Figure 1 that the once-tiny network burden that the HSD service entailed, because of this persistent, aggressive CAGR, has become a roaring beast not as routinely accommodated. The relatively comfortable 20-yr phase is evolving to a need to respond more aggressively in the network itself as the heyday of excess resources on the cable are over.

Another mathematically descriptive way to state the problem is to recognize that, while traffic has continued to grow exponentially, operators cannot realistically upgrade the network at an exponential pace. Furthermore, they would not want to do so even if it were possible because of the disruptiveness this would mean for customers. Simply put, the pace of technology change at the root of traffic growth exceeds the pace of network infrastructure change – at least physical infrastructure – every time.

Instead of incremental repetitive network adaption that is difficult to sustain, new and more impactful bets must be placed that reset the game with a new set of thresholds of capability and technology. This is precisely where strategy, vision, timing, and uncertainty collide. Network infrastructure investment can be capital intensive. Because of this, investment that has directional alignment to a broader vision is essential. And, as new technology is considered and with an objective of longer term vision alignment, the visionary promise also comes with increased uncertainty.

We elaborate on the known and deployed tools, discuss and then evaluate their role in strategy development, and identify potential next-generation disruptors.

Node Splitting

All MSOs know how to exercise this muscle. When the growth of narrowcast (VOD, HSD) traffic of a node exceeds some threshold of available capacity, beyond which experience says customer experience could be effected, the node is split. Theoretically, the node service group size is halved. In practice the more nodes are split, the more unbalanced they are likely to become. In some cases, additional steps must be taken to balance them or risk not achieving enough capacity benefit for customers.

In any case, node splitting divides up a serving area into smaller pockets as shown in Figure 2, and each pocket afforded its own narrowcast bandwidth. In so doing, the average capacity of each user is increased – the same Mbps of capacity divided over fewer subscribers.

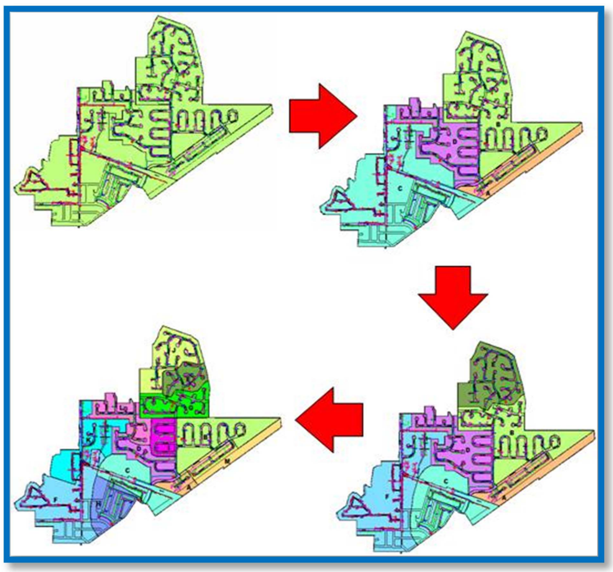

Figure 2 – Node Splitting is a Common Tool for Managing HFC Capacity [8]

Node splitting is efficient in the sense that it puts new capacity exactly where it is needed, and also because MSOs are very well-versed in this practice and can execute it quickly and with minimal customer interruption.

Fiber Deep

Fiber Deep is often referred to as "Node Splits on Steroids." The term "Fiber Deep" has different meanings for different operators. For Comcast, this is the term used for "Node + 0," whereby there are, by definition, zero RF amplifiers after the fiber optic node. Fiber Deep is largely similar to network activity MSOs have been engaged with for years – namely RF upgrades and node splits. The primary difference is that Fiber Deep is a more methodically implemented holistic architecture with specific objectives to maximize coverage and network capability.

Specifically, the many compelling advantages of Fiber Deep are:

1) Familiar HFC tools
2) Much smaller serving group size
3) More freedom of spectral allocation
4) Better end-of-line SNR and MER
5) Lower operational costs
6) Close access to direct FTTH connection for homes
7) Long-term HFC capacity runway under aggressive CAGR assumptions

Figure 3 describes the key principles associated with Fiber Deep.

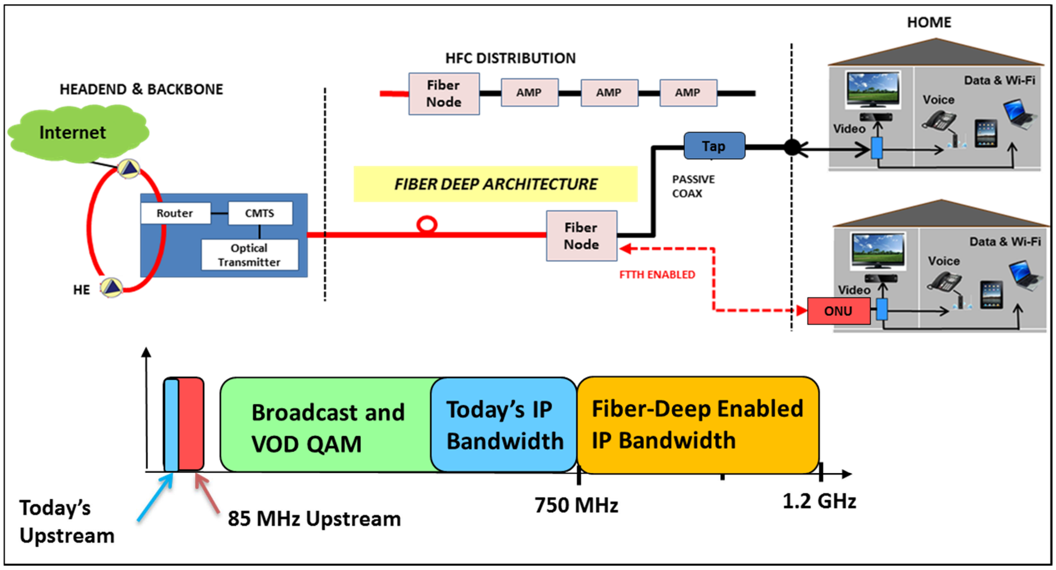

Figure 3 – Fiber Deep Architecture and Key Principles [6]

Spectrum Allocation

Operators have historically added downstream spectrum as the appetite for more video options increased. For Comcast, the transition to completely digital video (QAM) delivered 10:1 efficiencies of spectrum for SD carriage, and 2:1 or 3:1 for HD, essentially staving off the need for new spectrum since that transition took place. Those efficiencies have improved over time with performance advances and optimizations in video encoding. However, the margin in spectrum resources put in place by the use of digital only video (QAM) has nonetheless now been mostly consumed by the explosion of demand for IP data services, so the objective to add new spectrum re-emerges – both downstream and upstream.

Spectrum is a cable operator's most precious resource, so a reallocation must be very carefully thought through and justified. Tools for long term capacity must consider the ramifications to spectrum and the benefits of having it more flexibly allocated in the face of the uncertainty of new demand trends.

Also, while the allocation of spectrum must be done with support for projected CAGR for the long term in mind, it also recognize the market reality of the data speed wars. While CAGR is tied to average bandwidth per user sharing the resources, speed requires MSOs to allocate a dedicated amount of spectrum available over and above the maximum speed offering to ensure that customers will consistently experience the service speed they have signed up for.

For Fiber Deep, a proper balance of the aforementioned objectives is achieved with an 85 MHz upstream and 1.2 GHz downstream.

DOCSIS 3.1

The latest DOCSIS standard offers both the opportunity to expand spectrum beyond bandwidths that had previously been unavailable, both downstream and upstream, and to use the spectrum itself more efficiently. The latter is accomplished primarily by updating the physical layer (PHY) tools.

The combination of technologies in DOCSIS 3.1 provide the means to make the HFC network capable of approximately 10 Gbps downstream and 2 Gbps of upstream capacity, thresholds which promise runway for long term bandwidth growth when used synergistically with others tools, in particular Fiber Deep. Besides the extension of spectrum, the key new technologies of DOCSIS 3.1 are

1) Full spectrum of multiple simultaneous blocks of up to 192 MHz wide Orthogonal Frequency Division Multiplexing (OFDM)-based QAM. DOCSIS 3.0 is single carrier QAM (SC-QAM) and has a maximum of 32 bonded 6 MHz wide channels (192 MHz), which makes 1 Gbps service speeds impractical in scale. OFDM is the go-to signal structure for virtually all modern communications systems, leveraging the adaptability and flexibility of many narrowband carriers, as shown in Figure 4, to important system advantages and robustness to impairments that can be very deleterious to SC-QAM.

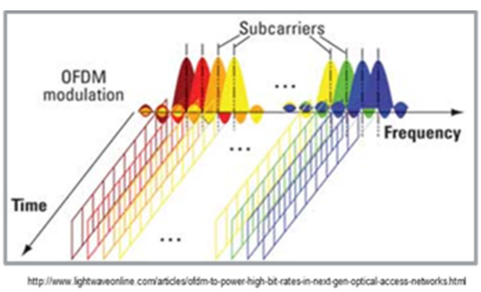

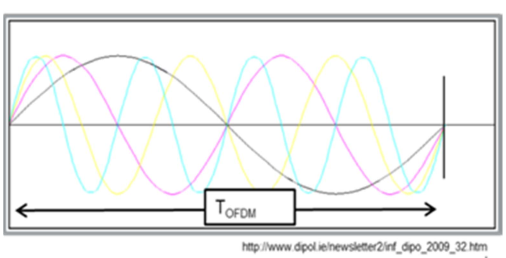

Figure 4 – OFDM as the New Signal Structure in DOCSIS 3.1 (note the underlying modulation is still QAM)

2) Updated Forward Error Correction (FEC) from Reed-Solomon based to Low Density Parity Check codes (LDPC). LDPC enables existing QAM at lower SNR than DOCSIS 3.0 by about 4 dB, or higher order QAM at the same SNR.

3) Availability of higher order QAM formats up to 4096-QAM downstream (from 256-QAM, or 50% more efficient) and 1024-QAM Upstream, (from 64-QAM, or 67% more efficient). Usability of these advanced formats takes advantage of both the FEC upgrade as well as the gradually improving network performance over time that has occurred as fiber has been pulled deeper and HFC optical technology has improved. Figure 5 compares the QAM constellation of DOCSIS 3.0 based 256-QAM and DOCSIS 3.1 based 4096-QAM.

Fiber-to-the-Premises (FTTP)

Many cable operators have been deploying FTTP for years, largely for commercial customers that rely on symmetrical data traffic needs. Residential data traffic is, and has always been, significantly asymmetric in nature in favor of the downstream. The Downstream-to-Upstream ratio has increased over the past few years due primarily to over-the-top (OTT) IP Video (i.e. Netflix, Hulu, etc). Nonetheless, as shown in Figure 6, the data speed wars continue, and FTTP options offer operators immediate paths to delivering these speeds if demanded by customers. The vast majority of customers will benefit from the simplicity of mass deployment of DOCSIS 3.1 with Gigabit speeds based on the coaxial connection to the home that already exists.

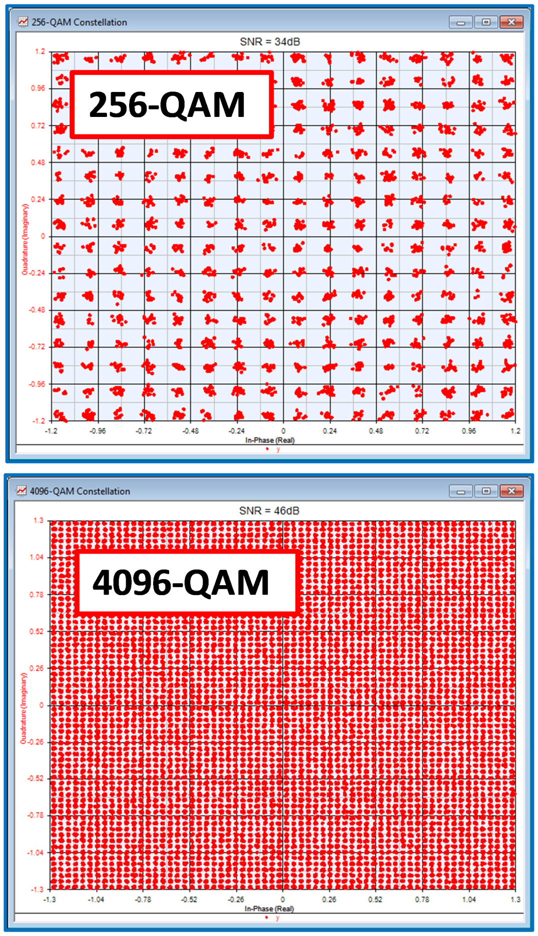

Figure 5 – DOCSIS 3.1 Bandwidth Efficiency Increase due to M-QAM Format

Figure 6 – Market Speeds Continue to Increase

Many operators have already deployed FTTP in residential footprint, and this is becoming more commonly the case for new construction as the costs for coax and fiber builds are nearly at parity. The evolution of HFC that takes advantage of the coaxial network without disturbing neighborhoods and customers will leverage many of the aforementioned tools of the previous section. For fiber systems, there are multiple options to choose from also – SCTE 174 [10], EPON, Ethernet in particular. The move to all-IP will further simplify and unify service delivery across different physical architectures. Indeed, this has always been one of the key value propositions of all-IP – any IP pipe can deliver the services, regardless of how the bits are framed, how the waveforms are generated, or what medium is used.

Lastly, as was shown in Figure 3, it is a fundamental premise of the Fiber Deep architecture that it is deployed with a "2nd putt" strategy to be FTTP capable in the future if needed. Fiber Deep (FD) provides physical access to fiber within approximately 1000 feet of all subscribers. By proactively designing a Master Architecture with a possible FTTP end-state, the optical distribution network (ODN) is deployed in a way that prepares the network for FTTP during the FD cycle. Since the deployment of new fiber is the most intensive and costly phase of new construction, it is therefore utilized to maximize long-term evolution efficiency via consideration of this end-state architecture expectation.

In general, it is a powerful strength of cable operators that they can deliver their full service suite of broadband services over coaxial or fiber last miles.

Solutions Options to the Problem Statement

Figure 1 demonstrates the fundamental problem statement of cable operators. Under an assumption of a persistently aggressive CAGR, such as the verifiable 50% per year, the point at which capacity is exhausted on the network as implemented today is moving into the long-range planning window – in the calculation of Figure 1, this is 2027.

The key specifics of this projection are a transition to a fully deployed, all-IP DOCSIS 3.1 network achieving 10 Gbps, a Fiber Deep (N+0) serving group size. Under these steps, the growth curve crosses the threshold of available capacity in 2027. While this is a long runway, the significance of this for operators and potential shifts in network investment type requires beginning the planning steps for “what’s next.”

Note that Figure 1 captures only the Downstream growth, as it is this aggressive CAGR dynamic that determines the ability of the network to support future capacity. There is a parallel analysis for upstream that is not shown because the driving dynamic for HFC evolution turns out to be the persistently aggressive DS growth, once the upstream relief of the 85 MHz mid-split is applied. The upstream does indeed grow – less consistently but at a measureable average CAGR – it simply turns out that without a major change in the traffic dynamics – to be wary of – it is secondary to the downstream impacts.

Now consider Figure 7. Figure 7 shows one scenario of many that were run to evaluate investment levels (in Net Present Value terms) for various migration approaches to the HFC Brownfield, including the BAU option. The scenario shown is a culmination of the optimization of a 6-month engagement process with key business, financial, and technical stakeholders, among others, at the table with their various perspectives:

1) Capacity Engineering – How much will we need and when?
2) HSD and Video Product Management – What product targets must be met for our Customers?
3) Architecture and Technology – How do we meet capacity and product needs, when, and at what cost?
4) Construction Engineering – How and how fast do we build it and at what cost?
5) Finance – What does the investment profile vs time look like for the various options?
6) Human Resources – How to achieve the necessary balance of skill sets?

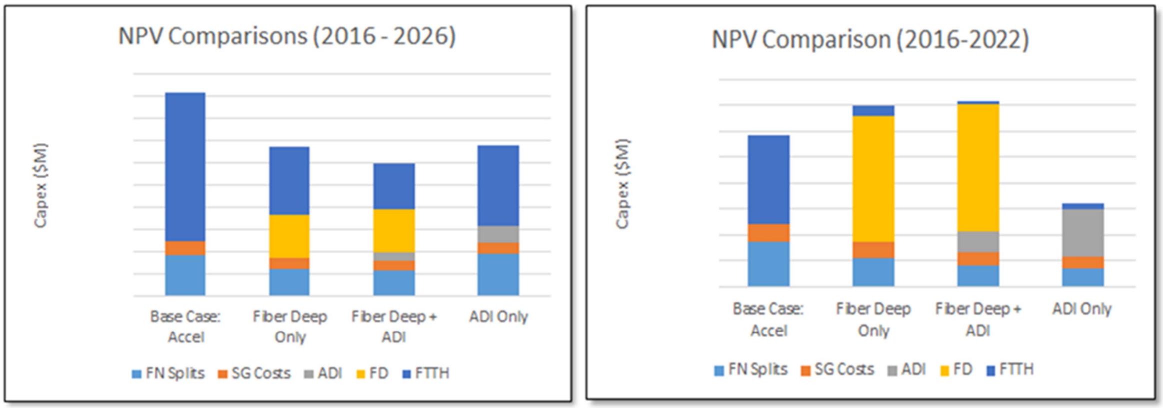

Figure 7 – Network Evolution NPV Investment Comparisons

The evolution scenarios shown in Figure 7 are:

<u>Base Case</u>

This represents the existing baseline level of network investment projected over 6 and 10 years. The "BAU" or Baseline foundation is node splitting on-demand at a standard threshold of capacity utilization and adding DOCSIS spectrum. Both of these are methodically planned, budgeted, and implemented as standard practice today.

Included in the DOCSIS spectrum are existing assumptions around the introduction and growth in allocation for DOCSIS 3.1 "channels," as a path to an all-IP network. All remaining channel slots (we use 6 MHz equivalent slots even though this channel basis is eliminated in D3.1) are for 6 MHz QAM channels for broadcast or narrowcast video, and a small number of slots carved out for overhead functions, OOB, etc. These QAM channel slots are gradually reduced over time with technology, shrinking serving groups, and transition to IP Video delivery.

Notably for "Base Case," while the service group size is reduced because of the node split, the basic HFC architecture is unchanged. Fiber is not necessarily any deeper, the same cascade of amplifiers generally exists, and the network spectrum allocation has not changed either downstream or upstream. Using 750 MHz as the available downstream bandwidth, approximately 108 total channels would be available in an all-IP end state (2025). This is less than would be available using a more aggressive drop-in upgrade (Active Drop-in or "ADI", to be described) or as compared to FD which extends the downstream to 1.2 GHz.

Finally, "base case" assumes that there is some threshold of HHP (several scenarios run to optimize) below which it does not make sense to split further because of the inefficiency. Instead, a proposed next step is to offload a portion of the "top talker" traffic to FTTP over nearby existing deployed fiber and/or via WDM and reduce the load on the HFC network for the majority of customers.

<u>Fiber Deep Only</u>

This category is as previously detailed – the network is taken to N+0, also known as passive coax. Fiber Deep by definition also includes an 85 MHz Upstream band edge and 1.2 GHz Downstream band edge.

An example scenario for DOCSIS channel growth transition (and thus QAM reclamation) is shown in Table 1 for FD, which has the most available spectrum of the HFC-centric options under consideration.

A key dynamic in the transition to DOCSIS spectrum is the practical pace at which QAM can be reclaimed as it relates to existing deployed QAM CPE. This requires extensive logistical planning and coordination, very similar to what was executed in the transition from analog video to all-QAM using the Digital Terminal Adaptors (DTAs).

ADI Only

The acronym "ADI" stands for Active Drop-In, which is also known as "plant upgrade." Nodes (if necessary) and amplifiers are upgraded to 1 GHz, thereby adding significant new bandwidth, in particular for 750 MHz systems. ADI also includes 85 MHz upstream, as expansion of the return is an important capacity strategy to achieve the long-term health of the network. Note that the upstream is not changed in the BAU base case – a significant disadvantage to that approach. The effect of this is to drive more splitting to handle congested upstream even at the moderated CAGR compared to Downstream.

	DOCSIS 3.0	DOCSIS 3.1
2016	24	16
2017	24	16
2018	24	40
2019	24	40
2020	16	64
2021	16	80
2022	8	104
2023	8	182
2024	8	182
2025	0	182

Table 1 – Example DOCSIS Growth vs Time for All-IP (Fiber Deep case only)

Because of the expansion of the downstream spectrum, about 140 channels would be available by installing actives that are 1 GHz – about 30% more spectrum. Note that some of the spectrum added is converted to upstream, where spectrum increases by over 200%. Also, it is anticipated that the network may not quite reach 1 GHz everywhere, at least not with a flat frequency response, because spacings of amplifiers may not support it without some roll-off attenuation.

ADI is expected to provide a fast, cost effective capacity relief valve to complement a Fiber Deep transition strategy. The nature of the Fiber Deep migration is that it takes time to complete in large scale as it system-wide (vs on-demand by node), and involves design optimization, the addition of fiber, and facilities work. There are also practical resource and logistical limitations to simultaneous construction projects

The strategy limitations of ADI are:

1) Node splits are still required, but are deferred due to the spectrum added both Downstream and Upstream
2) Loss of spectrum as compared to Fiber Deep
3) Lower RF performance than Fiber Deep (but the same as today), leading to lower D3.1 capacity
4) Not positioning the fiber any deeper to enable more potential FTTP customers now or in possible future transition
5) As we shall see – and perhaps falling into the category we have alluded to of a plan put in place that becomes superseded by the "next" plan as technology advances – not consistent

with the fundamentals of a full-duplex DOCSIS architecture

The category "*Fiber Deep + ADI*" is self-explanatory, meaning deploy some of both. The scenario depicted in Figure 7 assumes roughly a mix of both of the above, in approximately equal amounts.

<u>As of May 2016, a Winning Architecture Evolution Strategy is……</u>

What is immediately interesting in the scenario comparison depicted in Figure 7 is that in the very long-term view – in this case the 10-year view – the most costly path is the BAU path of attempting to node split your way through the growth. This is due the nature of the traffic growth (exponential), and that this approach does not address the spectral limitations of BAU, nor simple access to nearby fiber for FTTP offload. Note that there is in fact a large portion of expenditure allocated to FTTP. In the analysis, this takes place when s node sizes get below a certain (adjustable in the model) homes-passed threshold after which it becomes impractical to efficiently split further. Note also that the size of the bar for "FTTP" in Figure 7 is not a proxy for households passed (HHP), since FTTP in the HFC brownfield is costlier on a per-HHP basis compared to other tools shown and different depending on whether it is a Fiber Deep architecture or not, such as this case.

Just as interestingly, in the 6-yr view, this is not the case even with an FTTP transition underway. This 6-yr and 10-yr difference is insightful in what it says about the natural life of the as-is network when a "Base Case" approach is taken. Capacity constraints are within the planning window, and this is what drives the need to consider alternative strategies and assess their implications.

Note that what appears as most cost effective approach in the 6-yr view is "ADI" – suggesting a plant upgrade cycle is preferred. ADI addresses the priority capacity needs of Downstream and Upstream spectrum such that node splits and offload to FTTH are deferred. However, as the 10-yr view shows, node split re-emergence takes place and the cycle begins again. Node splits again become prominent in trying and keep up with exponential traffic growth, and FTTH offload begins as the deferment period ends and node size thresholds come back into view.

It is clear from the long-term view that the "winner" under the assumptions of the scenario in Figure 7 is a mix of Fiber Deep and ADI. Fiber Deep offers the order-of-magnitude sized capacity upgrade necessary to paint a long-term lifespan view of the HFC network in one methodical and well-understood step. In addition, it positions the fiber with close proximity to affect FTTP offload if necessary, so this "2nd putt" is then most efficient and cost effective. Perhaps most importantly, it defers FTTP offload as evidenced by the smallest of FTTP investment bar colors in the 10-yr view shown in Figure 7.

A long runway of lifespan with which to assess trends in applications, services, and technology is a comforting situation. It is not possible to predict where these trends lead ten years hence, nor would it be fruitful to try and predict these. But, we can use historical experience as a meaningful guide to what's next. So, while adjustments to the plan can be assured, the probability of an urgent, major shift in plan is lowered with such runway ahead.

<u>How Long Does a Long Term Plan Last?</u>

A dedicated team worked many scenarios across a range of assumptions, projections of demand, and product offerings to arrive at a go-forward near term and long-term strategy. There is no illusion that 10-yr plan remains unchanged for the duration of the migration. Acceleration of the pace of technology change is nothing new. New technologies will be developed that demand consideration because of the performance, savings, simplicity, or all-of-the-above that they bring.

With this in mind, unsurprisingly, just as the next generation access evolution journey begins, a wave of new technology possibilities are emerging. This is already happening a mere 18 months into the plan. These advances are certain to have an impact on the journey, suggesting a veering off of the defined path for the promise of some of those very compelling benefit mentioned above.

We discuss several of the most appealing opportunities, their effects, and implications below.

Distributed Access Architectures

A Distributed Access Architecture (DAA) breaks the CCAP function into two pieces and distributes a portion to the node. The functional portion distributed can vary. While there are a range of possibilities, the debate has mostly narrowed to two architectures for cable "DAA" – Remote PHY and Remote MACPHY. Regardless of which architecture wins the marketplace, DAA delivers some of the most powerful advantages *ever* deployed into the HFC network, primarily through the use of digital Ethernet optics. This relatively simple change creates a necessary reconsideration for the plan. Every node deployed today is implemented with AM optics. The Fiber Deep architecture defines a node with specific requirements for spectrum plan, output level, segmentability, all based on HFC optics – downstream AM optical receivers and upstream digital returns. However, just as the Fiber Deep plan is in the midst of implementation, along comes the opportunity to drastically improve upon it with DAA.

The major effects of a DAA implementation are as follows

1) Wavelength multiplier of DWDM vs the limitations of WDM for AM optics

With a migration to fiber deep, a significant number of new nodes are installed, as amplifiers are eliminated. Note that a "Fiber Deep" design is not the same as a "Fiber-to-the-Last-Active" (FTLA) approach, whereby existing amplifier stations are converted to nodes. Fiber Deep optimizes the location of the new nodes for HHP to optimize economics via reach. The addition of many new nodes demands the use of WDM to maximize the use of the existing fiber infrastructure. The use of digital optics means that there is a path to 80 DWDM wavelengths/nodes, for example, aggregated on a single fiber (or more). Standard HFC AM optics are restricted typically to 16 wavelengths to manage the nonlinear effects on MER for common HFC optical links lengths. More are possible, but there are performance and distance dependencies to manage that do not exist with digital optics. This means a more cost-effective deployment and a much lower probability of installing new trunk fiber to support the nodes.

2) Reach of digital optics vs AM optics

The longer the optical link from Headend transmitter to node, the fewer wavelengths that can be used at a fixed end-of-line (EOL) MER. Digital optics completely eliminates this dependency in practice for common HFC optical links. In fact, the CCAP core with digital optical outputs could be moved further back into the core if Headend consolidation is an objective.

The above aspects of reach and wavelength multiplication taken together massively simplify the HFC network and add never-before-seen flexibility that can pay large economic dividends as more fiber is pulled deeper into the network, which is exactly what fiber deep does.

3) SNR and MER performance improvement of eliminating the DS and US AM HFC optics

DOCSIS 3.1 enables higher order modulation profiles, increasing the bandwidth efficiency up to 50% over DOCSIS 3.0 n the Downstream and 67% in the Upstream. It also allows the QAM profiles to be varied and optimized to the network performance. In standard HFC and especially in Fiber Deep using HFC-style nodes based on AM optics, the EOL MER is dominated by the performance of the AM optics. The CCAP device transmits a very high MER (set by the Downstream RF Interface Specifications, or DRFI) that is standardized by CableLabs specifications. The DRFI requires higher fidelity transmission for DOCSIS 3.1 ports then for DOCSIS 3.0 ports.

In standard HFC, the DRFI spec sets a high level of fidelity to the optical transmitter – 43 dB minimum for DOCSIS 3.0 and 43 dB up to 48 dB for DOCSIS 3.1, depending on frequency band. This higher fidelity is set to support the higher order DOCSIS 3.1 QAM profiles.

A Headend optical transmitter then acts on the signal and degrades this, and a common requirement is to exit the node at a minimum MER of 38 dB. There is therefore significant fidelity loss in the AM optics of 5-10 dB. The SNR/MER that exists at the node degrades slightly more as it is passed through amplifiers. In Fiber Deep, this degradation does not take place because there are no amplifiers.

With DAA, the link to the node eliminates the AM optics, and eliminates the SNR/MER degradation caused by this link. Instead, the DRFI requirement is met at the node equivalent CCAP port, gaining back the lost fidelity of the AM optics nearly completely (the node has RF actives between CCAP port and output). This maximizes the capacity possible for D3.1 output signals.

The most capable HFC network that can be deployed from a network capacity, operational robustness, and consistency of performance perspective is an N+0 Fiber Deep architecture implemented with a DAA approach, and a point-of-entry home gateway architecture.

4) Space, power, cooling efficiencies in the Hubs and Headends

It seems obvious that moving some of the CCAP functionality into the plant (i.e. the node), leaves less behind in the Hub or Headend to power, cool, and consume space. This is, in fact, the case. However, the benefits are perhaps more outsized than intuitively obvious. The density of RF connectors and supporting electronics and isolation requirements on today's CCAPs tend to set the density of these chassis. With the RF port of the CCAP distributed into the plant, the density of the CCAP core can now be redefined since the density of optical connectors is much higher.

Secondly, the use of DAA drives an accelerated migration to CCAP convergence for video, eliminating Edge QAM boxes and RF combining networks.

The benefits to facilities of DAA cannot be understated in a migration to Fiber Deep, which installs many new nodes for every single existing node. A degree of aggregation occurs initially in Fiber Deep to avoid excess investment in capacity before it is needed in practice. Nonetheless there is a significant equipment burden placed on facilities to activate a Fiber Deep network.

5) Alignment to strategic NFV and SDN evolution across multiple last-mile coaxial and fiber access technologies

This will be discussed in the next section.

In summary, the significance of DOCSIS 3.1 and Fiber Deep to a long term evolution strategy are magnified by the value of bringing DAA systems to market given the wide ranging, incredibly powerful benefits identified above.

NFV / SDN

Cable operator's networks are a collection of purpose-built video, voice, and data platforms of integrated hardware and software. This has been a successful formula for over 30 years of service introduction including digital video, HD video, voice, data, and VOD.

As detailed throughout this paper, developing any long range plan must deal with the tsunami of persistently aggressive traffic growth that threatens to exhaust network capacity without some of the tools described herein. Another way to think about this is that the "linear" scale of new hardware development does not keep pace with exponential traffic increases now that traditional HFCs resources are fully deployed or nearly so. This extends beyond just network capacity, but also into service velocity enabled by the transition to all-IP, cloud, and software-based systems.

Network Function Virtualization (NFV) and Software Defined Networking (SDN) are enablers of cost effective, efficient, exponential network and service change velocity. The deployment of Fiber Deep and DOCSIS 3.1 to support the continued growth in traffic traditionally means the addition of more and more RF ports and CCAP line cards.

Bringing this lofty concept down to operational practice, in a DAA implementation, these RF ports are distributed into the field serving smaller groups of subscribers. However, supporting line cards of the CCAP core would still necessary with the disaggregation of an existing Integrated CCAP (I-CCAP). Figure 8 conceptualizes this logical architecture [9].

In a virtualized implementation, however, this purpose-built hardware core, designed to be tightly coupled to the output RF interfaces, is instead implemented in Commercial Off-the-Shelf (COTS) server hardware. This is made possible simply through the continuance of Moore's Law. The compute power and resources needed is available in standard processors today, allowing CCAP functions to be executed in such platforms and be software-based.

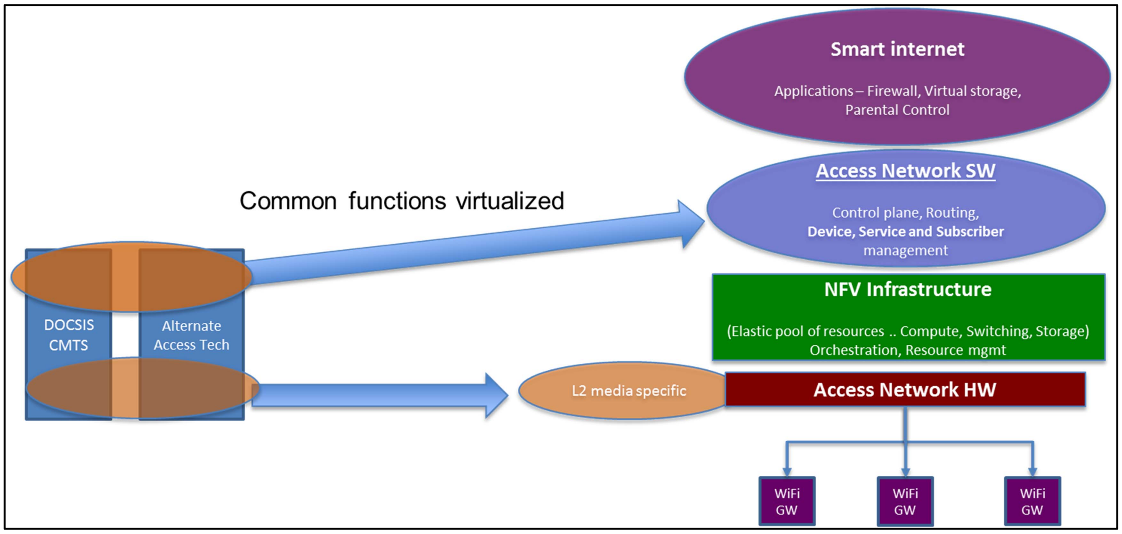

Figure 8 – Virtualization of Access Edge Platforms [9]

The potential to create SW-based CMTS on standard server platforms has enormous implications for cost, facilities, and service velocity – cheaper, faster, smaller, and more cost effective. This virtualized path applies in the case of I-CCAP and to DAA. A virtualized I-CCAP system is actually a form of DAA, since a COTS server-based CCAP core is never going to have RF outputs. And, similar to the DAA topic previously, and again representing the natural conflict of the acceleration of technology and limitations of practical infrastructure cycles, the NFV/SDN development activity will lead to a significantly superior architecture within a few years of deploying DAA systems. Initial DAA systems being invested in and deployed into the field will likely include some that are based on purpose-built hardware cores, such as today's CCAPs.

Also just as in DAA, the advantages of a migration to virtualized edge platforms are so compelling that figuring out this transition from integrated architectures to distributed architectures to virtualized *and* distributed is essential.

Full Duplex (FDX) DOCSIS

Whereas DOCSIS 3.1 pushed the network to 10 Gbps / 2 Gbps, a "Full Duplex" version of DOCSIS, currently in feasibility analysis phase, aims even higher. Taking advantage of another key tool – echo cancellation – borrowed from the telco world, FDX DOCSIS targets up to 10/10 Gbps of symmetric or near-symmetric capacity. Figure 9 demonstrates the fundamental concept behind full duplex DOCSIS.

FDX DOCSIS comes with some network expectations of various levels of dependency. Among them, in decreasing level of essentialness, are:

1) Fiber Deep, N+0
2) Distributed Access Architecture
3) Transition to all-IP

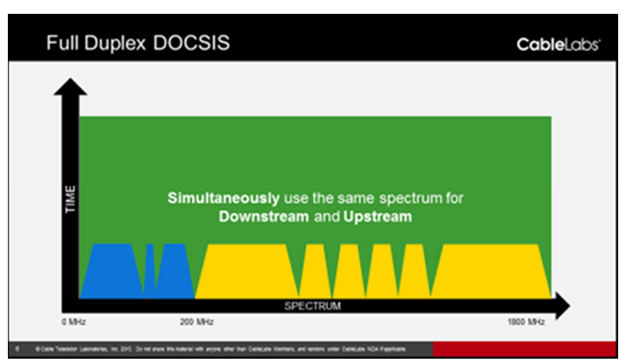

Figure 9 – Full-Duplex DOCSIS Fundamentals [2]

The concept of full duplex for DOCSIS borrows heavily from technology developed in other wireline and wireless industries, but comes with its own complications unique to the cable architecture that are still being evaluated. However, the potential benefits are so valuable – symmetrical capacity and speed regardless of whether customers are connected to a coaxial or fiber drop – this technology could also cause an adjustment to the Fiber Deep plan to not just migrate to DAA, but DAA upgradeable to full duplex.

Wireless

The vast majority of consumers receive their data services wirelessly in the home, and by far prefer they this be the case. Thus, the capabilities of WiFi and where this technology is headed are critical to customer Quality of Experience (QoE). Recognizing the value of WiFi, Comcast has deployed millions of WiFi hotspots that provide customers access to their HSD services when they are outside the home.

While WiFi is an extremely powerful and critical component of a service provider's end-to-end architecture, it is not ubiquitous and is arguably the 2nd most valuable wireless link to the end consumer – their mobile connectivity being even more essential to daily life.

Cable operators have flirted with the "quad-play" for many years. However, with the transformation from mobile phone to smart phone, the role of these devices for media consumption and broadband access has never been higher and continues to increase. As such, renewed inspiration to engage between cable and mobile operators exists.

Today, as 5G comes into view, it comes with promises of tremendous speeds and capacities, but also some critical RF, Millimeter Wave, and home architecture dependencies to achieve its loftiest objectives. One of these dependencies is accessibility to many accessible fiber connections and a widely distributed powering grid. In this way, it is quite synergistic to the infrastructure being built as markets deploy the Fiber Deep architecture. As such, we again come across a long-term plan in motion being confronted with some exciting new possibilities to consider.

The same logic applies to various "Internet of Things" (IoT) municipal and enterprise applications. The "Smart Cities" initiative is based on ubiquitous coverage of low power sensors of various types with RF connectivity to wireline (or wireless) backhaul points.

Thus, as a Fiber Deep architecture is being designed and with a consideration for possible wireless integration, it makes sense to ask: What module types should a modular Fiber Deep node support? What other access point types (in addition to WiFi) might need to be powered and enabled on the strand?

More complex still, in particular in the case of 5G applications, is to go beyond the adaption to a wireless edge on a Fiber Deep architecture, but also to think about the processing core. A value of virtualized is COTS platforms with application-specific software and common CPU resource pool. The platform itself can thus be simultaneously supporting DOCSIS platform, PON, and wireless cores as well. Each access technology having their own farm of server silos seems to be at odds with basic efficiency and flexibility arguments for NFV and SDN. This is where open interfaces and standards offer tremendous value by allowing differences in service provider architectures delivering common services across them to be integrated efficiently.

Somehow, this long-term vision must be coordinated with the already underway plan, including the nearer term path adjustments described above.

<u>Summary</u>

Service providers of all kinds are grappling with the same problem statement – a consumer bandwidth explosion that shows no signs of settling down. Capacity management using common tools of the trade that have been so successful for so long are being rapidly consumed, and the next round of network evolution is now at hand with many new options on the table.

Traffic growth acceleration is easily quantifiable, but other aspects of service needs and new technology that effect decisions on investment are not as predictable and are moving faster and faster, with shorter and shorter life cycles. It is the architect's job to develop a plan to support the fact that service and technology change outpace the ability to adapt the network at each turn. Bigger, more impactful, longer term bets must be placed using the most information that can be gathered or projected over the expected lifecycle of capital investments such as network infrastructure.

The goal of this paper was to describe some of the processes, analysis, and personnel that one operator used to develop a model and determine a possible path forward in this challenging environment.

References

[1] Chapman, John, Mike Emmendorfer, Dr. Robert Howald, Shaul Shulman, *Mission Is Possible: An Evolutionary Approach to Gigabit-Class DOCSIS*, Cable Show Spring Technical Forum, Boston, MA, May 21-23, 2012.

[2] Hamzeh, Belal, CableLabs Winter Conference Innovation Seminar, Feb 8-12, 2016, Orlando, FL

[3] Howald, Dr. Robert L, *Boundaries of Consumption for the Infinite Content World*, SCTE Cable-Tec Expo, New Orleans, LA, October 20-22, 2010.

[4] Howald, Dr. Robert L, *Breathing New Lifespan into HFC: Tools, Techniques, and Optimizations*, Cable Show Spring Technical Forum, Washington, DC, June 10-12, 2013.

[5] Howald, Dr. Robert, “Fueling the Coaxial Last Mile,” SCTE Conference on Emerging Technologies, Washington, DC, April 3, 2009.

[6] Howald, Dr. Robert, “Growth Architectures: Built to Last, Built to Launch,” Cable Show Spring Technical Forum, Los Angeles, CA, Apr 29-May 1, 2014.

[7] Howald, Dr. Robert L, *Looking to the Future: Service Growth, HFC Capacity, and Network Migration*, SCTE Cable-Tec Expo Capacity Management Seminar, sponsored by the Society for Cable Telecommunications Engineers (SCTE), Atlanta, Ga, November 14, 2011.

[8] Howald, Dr. Robert L and Robert Thompson, Tomorrowland: Customers, Servcies, Networks, SCTE Cable-Tec Expo, Denver, CO, Sepember 22-25, 2014.

[9] Nandiraju, Nagesh, NFV and SDN – A Comcast Perspective, Open Networking Summit, March 14-17, Santa Clara, CA.

[10] SCTE 174 2010, “Radio Frequency Over Glass Fiber-to-the-Home Specification.”