Growth Architectures: Built to Last, Built to Launch

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

Abstract

Since the introduction of DOCSIS in the late 90's, cable operators have embarked on an aggressive phase of enhancing existing services and adding new ones. The triple play and the tipping point of mass HD soon followed DOCSIS, and bigger and better service offerings continue today with the addition of whole-home services, multi-screen video, vast on-demand libraries, Wi-Fi access points, and cloud DVR (cDVR). Ironically, while the transparent nature of better services is testimony to the tremendous flexibility of the HFC architecture, it can be perceived by customers as a lack of attention to investment in network upgrades – i.e. out-of-sight, out-of-mind – whereby other would-be service providers announce network installations often with extensive media fanfare. Of course, this perception of cable system evolution is far from the reality.

Referencing the launch of DOCSIS means that this renewed phase of investment has been going on for about 15 years. In many cases, upgrades are required simply to keep pace with uninterrupted traffic growth. While today's incremental upgrade approach has been effective, the trend of needing to deliver even more and at an accelerated pace makes this approach less cost effective and less practical going forward. In essence, the necessary pace of service evolution exceeds that of conventional network evolution. Instead, an even more aggressive response aligned with the pace of technology change and service demand is required – and it must take place with at least the same transparency to the end user achieved today.

Operators have cost effectively evolved HFC since its inception, relying on a proven, robust, and flexible architecture able to adroitly match architecture and technology to service evolution and are evaluating the avenues and timing for the next phase of network investment. A key objective going forward is to address evolution comprehensively, synergistically, and proactively in anticipation of next-generation service and customer expectations. Alignment across all impacted areas – Headend/Hub, access network, CPE, and cloud/software – will maximize ROI, create agility and service velocity, and optimize the customer experience. There are many simultaneous and interdependent parts to assess given the pace of consumer demand and technology change. In this paper, operator guidance reflecting key areas around services, technology, architecture and system engineering will be discussed. Operators have important bets to place to be prepared for the future, and will team with key industry partners to help drive the continuous improvement of the customer experience enabled by sound investment strategies. This paper will outline operator thinking around future network evolution paths, and offer insight to solution partners in order to fulfill this mission.

INTRODUCTION

Cable operators have seen downstream IP traffic sustain a Compound Annual Growth Rate (CAGR) of 40-50% for many consecutive years since introducing data services. Upstream has grown as well, although on a more irregular trajectory. These trends are a useful foundation to consider when evaluating long term capacity, investment, and drive strategic decisions on architecture and technology. Reasonable debates can take place over the long-term growth trends of media consumption [2,8,9], although to-be-seen future applications may emerge beyond media consumption. Placing strategic bets based on anticipating a slowdown in bandwidth growth is not prudent.

Regardless of how future long-term trends actually play out, the prior decade plus represents a "long term" of periodic investment to support service growth. This has particularly been the case since high-speed data (HSD) and HDTV became service cornerstones. Operators are now evaluating potential avenues and timing for the next phase of network investment. Operators have consistently kept ahead of the "need for speed," and in providing the aggregate network capacity to meet demand. Cable's uniquely nimble and cost-effective HFC architecture has delivered more and better services over an explosive 20 year stretch of media innovation that has included the rise to ubiquity of consumer Internet. It is now poised to deliver even more with strategic evolution and the integration of key architecture and technology components.

In this paper, we will quantify the service growth challenges, define the architecture and technology steps that deliver a cost-effective sustainable evolution playbook, and describe how the right architecture investments and strategy create a touch-once network migration path. We will show how the right choices install a long term capacity runway that supports the pace of service evolution and positions the network for a continuing bright future in this era of HFC evolution and beyond.

CAPACITY MANAGEMENT TIMELINE APPROACH

To quantify the challenge of service growth and capacity management, we can turn to a Capacity Management Timeline style analysis [5,8,9]. A sample analysis is shown in Figure 1. This tool, a one-page snapshot capturing all aspects of services, traffic growth, service group splitting, available spectrum, and technology implications, is invaluable to operator planning. The logarithmic scale enables the plot to cover the wide range of Gbps inherent in compounding long-term bandwidth, which can be otherwise be difficult to capture given a foundation of exponential growth.

Figure 1 shows various horizontal capacity boundaries for four different possible thresholds. Thresholds shown include allocations of 24 and 32 DOCSIS QAM carriers of 256-QAM, as well as 2 Gbps of capacity. The latter threshold could be, for example, 32 DOCSIS 256-QAM carriers joined by approximately 96 MHz of DOCSIS 3.1.

In general, cable operators tend to use their entire spectrum to maximize services for their customers. It takes proper planning and investment to increase the spectrum allocated to DOCSIS channels, and thus the awareness of when there will be a need to act, such as Figure 1 provides, is extremely valuable.

A fourth threshold is drawn as a reference to compare what may be available today for IP data (such as 24 DOCSIS QAMs) and what the network is capable of providing in terms of total Gbps. The example shown is the threshold labeled 750 MHz of QAM (for 750 MHz systems, obviously), in this case referencing only DOCSIS 3.0 capability, which is 256-QAM.

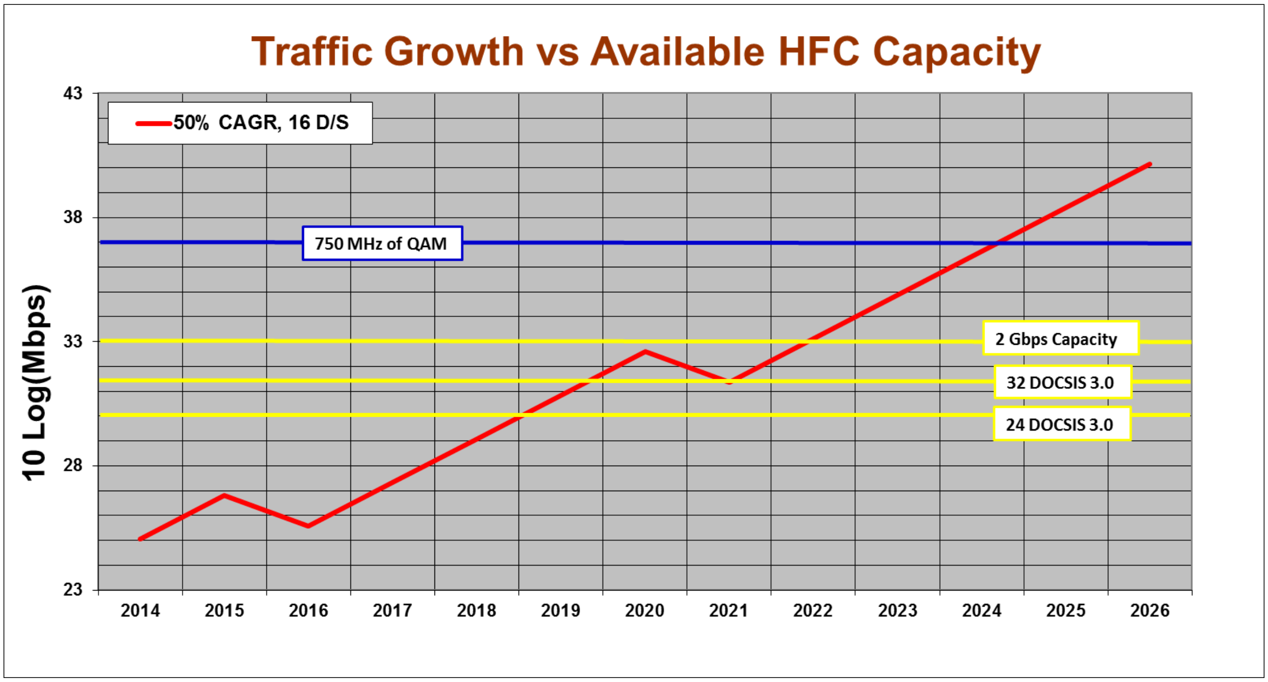

Figure 1 – A Capacity Management Timeline Guides Service and Architecture Evolution

Interpreting the rest of Figure 1, the growth of IP data (DOCSIS) is shown by the red trajectory trending upward with a slope that represents a 50% CAGR. It begins with 16 downstream channels deployed – the 2014 objective – utilized at inception at 50% at peak busy hour. The upward trajectory is broken twice along the way, representing service group splits. These splits (node segmentation or physical node splitting) represent a common, straightforward "business-as-usual" (BAU) approach to capacity management. In splitting service groups, operators manage capacity growth by sharing the existing bandwidth over a smaller number of subscribers, adding the narrowcast ports to support the new service groups, and in so doing (ideally) double the average bandwidth per subscriber.

IP video, or cloud TV (cTV), is assumed part of the engine that keeps the 50% CAGR going. There is not a separate "set-aside" of DOCSIS bandwidth for IP video. It's effect is embedded in the 50% CAGR itself the way over-the-top (OTT) video has been for many years. Various prior analysis consider an alternate approach of a steady 50% CAGR with IP video services introduced as an additive block of bandwidth [3]. In the long-term, such as over the time scale of Figure 1, the difference becomes very small as an offset of fixed channels is overwhelmed by the persistently aggressive compounding growth assumption itself. The assumption used thus only plays a role only in the details of managing near-term simulcast.

Figure 1 shows continued aggressive CAGR exceeding thresholds set by the shown DOCSIS channel allocations about 4.5-5.5 years down the road with one node split assumed. Using the two node splits shown provides 10 years of life on a 750 MHz plant of 256-QAM (DOCSIS 3.0) if all of the bandwidth were available for IP growth by that time. Not all of this bandwidth is available until the IP transition is 100% complete. In fact, this analysis approach is precisely valuable for its ability to assess the timing recommended for providing more DOCSIS channels over time, and ultimately the timing it suggests for completing the all-IP transition.

Basic conclusions from Figure 1 are that existing capacity and service growth for 750 MHz networks may present long term challenges under the assumption of persistently aggressive downstream CAGR. Also, current BAU evolution practices offer a solid runway of time to analyze trends and, as we are describing in this paper, develop a plan that overcomes the challenges and continues to deliver more and better services for the long term.

Upstream Lifespan Perspective

Similar to the calculations in the downstream, we can project upstream lifespan under various growth and service scenarios. CAGR for the upstream tends to be less predictable – spiking when Napster and YouTube were introduced for example, and lagging during other periods of time, including recent years.

With a varying CAGR range, the format shown in Figure 2 handily provides a sensitivity analysis around upstream CAGR. The chart is very straightforward to interpret – for the given scenario identified for each curve on the plot, the lifespan of the upstream is shown on the vertical axis against an average year-on-year (YoY) CAGR on the horizontal axis.

An underlying assumption for lifespan calculations is that four 64-QAM carriers (DOCSIS 3.0) of 5.12 Msps/6.4 MHz represent a full 5−42 MHz upstream. That is, there no DOCSIS carriers centered below 18 MHz and S-CDMA is not turned on. This represents a typical band usage assumption today.

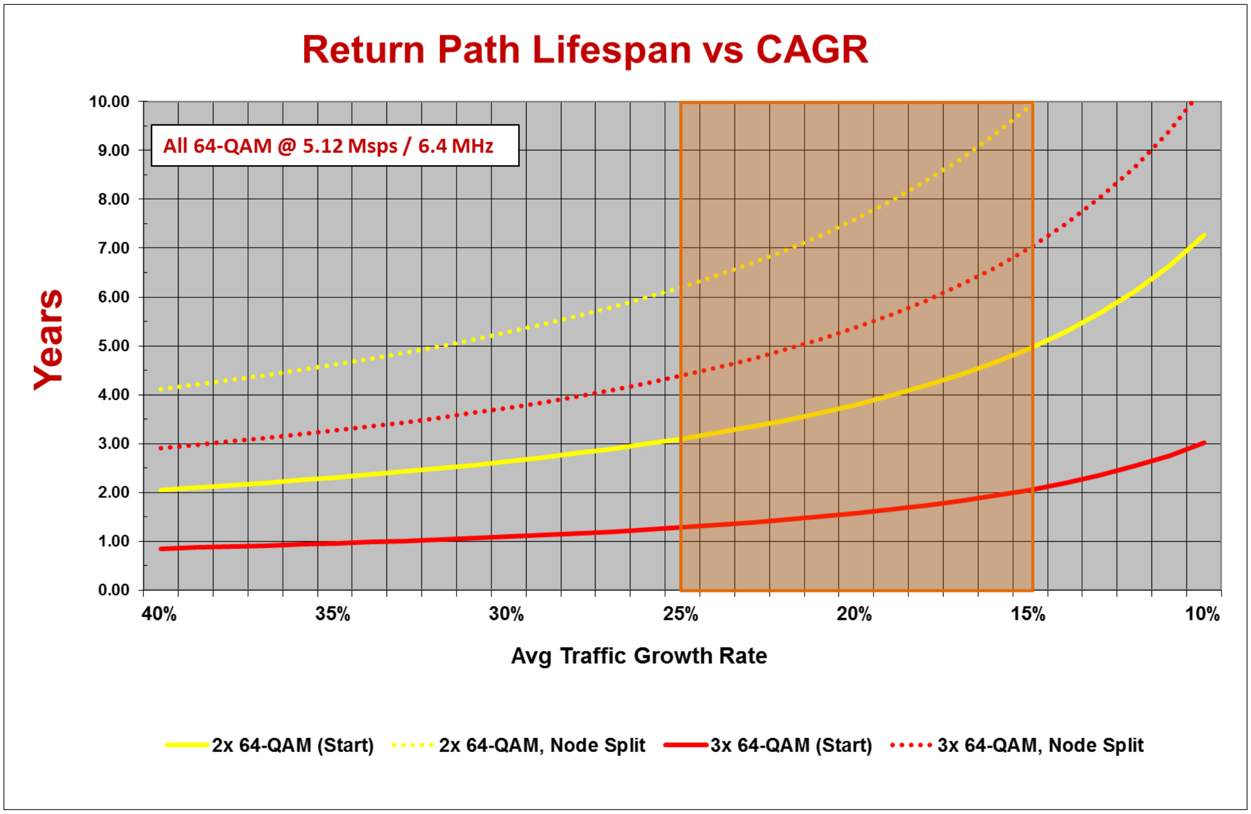

Figure 2 – Quantifying Upstream Lifespan

From Figure 2, we can conclude that where there are three 64-QAM carriers deployed today at full load, then for an upstream CAGR of 15-25%, there are 1-2 years of lifespan left before the upstream is fully utilized. As with the downstream, fresh runway is created by a BAU node split, which buys at least 3 years. Also, there is inherent margin in the years of lifespan where utilization is not 100%, which is typically the case.

For example, if the channel utilization is 20% below the threshold of congestion that would trigger a node split (i.e. it is 80% full), then a 25% CAGR means there is one extra year of margin to what the chart shows. In the 20% underutilization case with 25% CAGR and three 64-QAMs deployed, there are therefore about 2.5 years of life before another node split is required. Including this second node split there are about 5.5 years. Again, this runway provides for a comfortable time window for planning. The timing is right now to set this direction. Since areas of high traffic usage tend to be the ones with 3 or 4 upstream carriers, it is these cases that are of most interest for crafting strategy and projecting investment timing.

Note that perhaps as important as aggregate capacity of 5−42 MHz return is the limit placed on peak burst rate. In addition to delivering higher and higher upload speeds to customers, this is important for most efficiently enabling Gbps downstream speeds.

CAPACITY OPTIMIZATION LEVERS

Theoretical capacity is straightforward and based on two variables – bandwidth (B) and Signal-to-Noise Ratio (SNR). The well-known Shannon Capacity limit is the maximum error-free rate that can be obtained in an additive white Gaussian noise (AWGN) channel, and given by

$$C = [B] \mathrm{Log}_2 [1+\mathrm{SNR}] \qquad (1)$$

This is further simplified in cable, in particular for the downstream, by using high SNR assumptions typical of cable networks. In high SNR conditions, capacity is closely proportional directly to bandwidth, B, and SNR *expressed in decibels* (dB):

$$C \approx [B] [\mathrm{SNR\ (dB)}] / 3 \qquad (2)$$

The simple message of (2) is that increasing the available capacity involves increasing spectrum, increasing SNR, or both. Architecture evolution should thus aim for these goals.

Let’s first consider SNR.

QAMazing

Improving link SNR enables more bandwidth efficient modulation formats. Forward Error Correction (FEC) has an important role in new capacity in how it relates to SNR, and therefore capacity. While nothing about FEC shows up in (2), better FEC enables a given M-QAM format to operate at a lower SNR. Simply put, the best FEC makes for the most efficient use of the “SNR” in (1) and (2).

As an alternative to a given M-QAM operating at a lower SNR, we can also say that for a given SNR, better FEC enables M-QAM formats with higher bandwidth efficiency. DOCSIS 3.1 takes advantage of this, and is therefore a critical technology component to the evolution path forward. Regardless of improvements forthcoming in end-of-line (EOL) SNR, DOCSIS 3.0 (and below) limits the downstream to 256-QAM, or 8 bits/symbol. However, HFC commonly delivers performance higher than what is needed to support 256-QAM. What is wasted system margin today is taken advantage of in DOCSIS 3.1 by allowing higher order M-QAM profiles. An architecture evolution path should take advantage of these DOCSIS 3.1 innovations by delivering better SNR to most efficiently use spectrum.

Figure 3 shows the higher modulation formats associated with DOCSIS 3.1 including two of the key DOCSIS 3.1 additions – 1024-QAM and 4096-QAM. Each is shown for an *equivalent uncoded BER* of ~1e-8. The 46 dB SNR for uncoded 4096-QAM highlights both the expectation of an architecture with improved network performance, as well as the power of the modern FEC technology being adopted.

Compared to 256-QAM, the bandwidth efficiency increase of 1024-QAM, 2048-QAM, and 4096-QAM is 25%, 37.5%, and 50%, respectively.

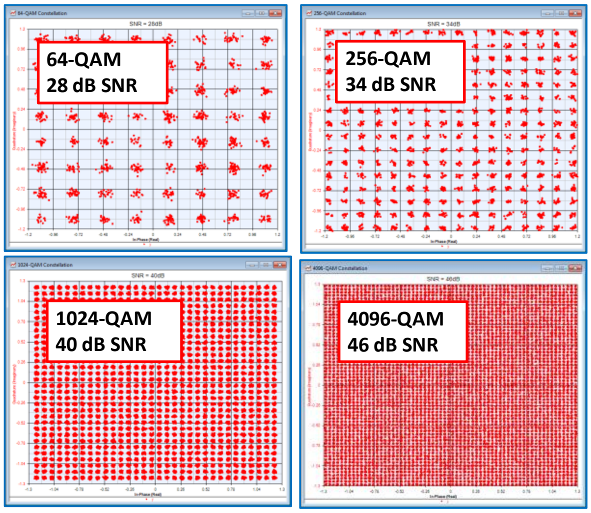

Figure 3 – DOCSIS 3.1 M-QAM Format Comparison

Shannon Nuances

Not obvious in the capacity equations in (1) and (2) is the assumption of a fixed SNR over the bandwidth, B. In real systems, variations in SNR are likely across the band and across an HFC physical footprint. DOCSIS 3.1 also aims to take advantage of this by providing modulation format flexibility, as well as having modulation formats targeted to the performance of CMs on a "binned" basis – a tool known as Multiple Modulation Formats (MMP). Figure 4 illustrates this concept [3].

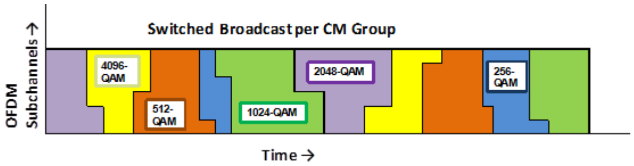

Figure 4 – DOCSIS 3.1 Includes Use of Multiple Modulation Profiles (MMP)

Fielded cable modems show a wide enough range of received SNR that to enforce the same M-QAM format on every user would be to deliver lowest common denominator performance, leaving many Mbps on the table. Figure 5 shows the reported SNR of a large sample of CMs [11]. It shows the majority of CMs are spread across three QAM profiles, while at least five profiles are observable when considering the non-zero percentage of CMs at the top and bottom end of the distribution. Figure 5 is a clear justification for the use of MMP to optimize capacity while preserving the basic simplified broadcast structure of the downstream.

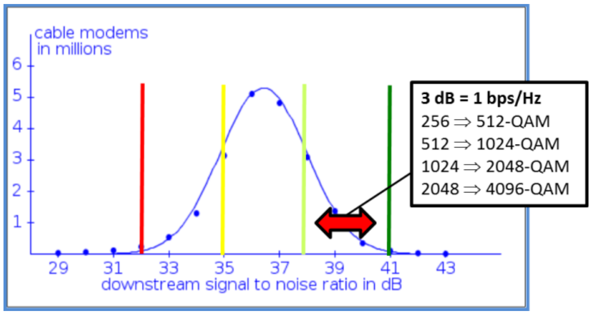

Figure 5 – SNRs Reported by Fielded CMs

Role of OFDM

To deal with non-constant SNR across the band as it relates to the "B" in the capacity equations, DOCSIS 3.1 introduces to cable the use of Orthogonal Frequency Division Multiplexing (OFDM) as the waveform on the wire. OFDM is notable for its ability to maximize capacity on a channel, and in particular on unknown or complex channels.

Why is this important to cable? An example of spectrum variability is the new bandwidth above 1 GHz, support for which is included in DOCSIS 3.1. Figure 6 [1], shows insertion loss characteristics of various Tap models above 1 GHz for "1 GHz" specified taps. A cascade of Taps will have wider variation.

Waveforms such as OFDM, with its narrow subcarriers that can be fitted to the channel characteristics, allow maximum capacity extraction from such channel performance variations. The DOCSIS 3.1 OFDM parameters (subcarrier spacing, cyclic prefix values, FFT size) were chosen based on HFC channel characteristics developed as part of a renewed channel modeling exercise [12, 13]. Both OFDM and MMP are tools selected for DOCSIS 3.1 to deal with the practical aspects of maximizing capacity in ways that are not accounted for in the simple capacity calculation of (2).

Summarizing, per capacity equation (2), higher SNR is an important component of increased HFC capacity. DOCSIS 3.1 has a primary objective of optimizing capacity by providing bandwidth efficient modulation tuned to the available SNR through:

1) Higher QAM formats
2) Better FEC
3) Optimal waveform (OFDM)

Note that while more capacity is indeed available with a higher SNR, it is with logarithmic proportionality. For example, either 50% more spectrum (B), or 50% more SNR (in dB) yield 50% more capacity.

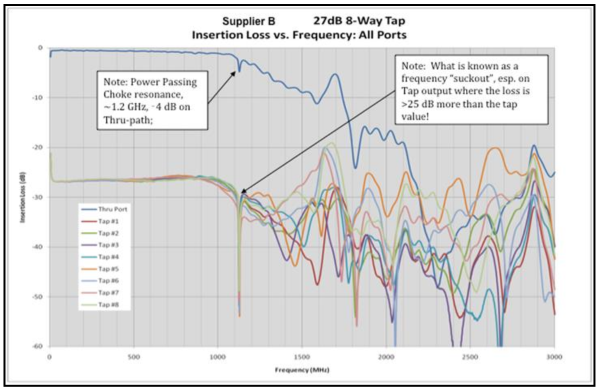

Note: Power Passing Choke resonance, ~1.2 GHz, -4 dB on Thru-path;

Note: What is known as a frequency "suckout", esp. on Tap output where the loss is >25 dB more than the tap value!

Figure 6 – Varying Tap Responses Above 1 GHz [1]

However, turning a 40 dB network SNR into a 60 dB SNR is a very tall if not impossible order, practically speaking. The implication of this is that, while deploying architectures with higher SNR is a significant objective, ensuring more spectrum offers more bang for the buck, and should be high priority.

We've identified both more SNR and more bandwidth as clear avenues that a network architecture should pursue to deliver new capacity, and DOCSIS 3.1 as the technology tool to take advantage of the higher SNR. We now discuss an evolution path that delivers on these objectives.

FIBER DEEPEST

We've emphasized the important capacity levers to target for architecture evolution. "More" capacity is of course better, but it is also critical that the investment lead to a quantifiably long-term capacity runway such as we can project via analysis such as Figures 1 and 2, as opposed to temporary gains that lead to repeated investment. As always, efficiency of capital expense is important.

For "brownfield" HFC, a logical conclusion that aligns with each of the objectives emphasized is to drive fiber directly to the coaxial last mile for carriage into the home. This is referred to as a "Fiber Deep" (FD) architecture. The foundation of the FD approach is an N+0 network architecture.

We can identify some of the fundamental objectives of the FD approach:

1) Leverage the high capacity coaxial last mile for cost effective growth
2) Provide a one-time touch sufficient for long-term bandwidth needs
3) Improve EOL performance through the elimination of the RF amplifier cascade
4) Deliver all current and emerging new services including Gbps speeds, 4kHD, Cloud TV (cTV), near-ubiquitous HSI via WiFi hotspot access, and is enabling of FTTP and DOCSIS-based business services
5) Significantly reduced operations and maintenance costs and enhanced customer QoE
6) Serve as the launch architecture for FTTP alternatives for long-term migration as needed

Executing a Fiber Deep network evolution is the sought after one final outside plant (OSP) touch operation, providing long term sustainable bandwidth, while delivering operational efficiencies and savings for as long as we can reasonably expect to project "long term" for a business built on technology.

Fiber Deep System Benefits

Figure 7 shows phased BAU service group segmentation via deeper fiber penetration as it might take place over an actual serving area footprint. It illustrates how a natural architectural end-state of HFC in a BAU sense would culminate in an N+0 system with an all-passive coaxial last mile through repeated node splits. In Figure 7, a single node serving area (top left) is ultimately broken into serving group sizes, which may be as small as (approximately) 50 hhp. Note that physical node boundaries of the map do not capture virtual segmentation of the nodes, of which configurations up to 4x4 are typical.

A fiber deep strategy instead bypasses the incremental phases and migrates directly to Fiber Deep (N+0), understanding that the persistency of CAGR inevitably drives the network to this optimal end state. Of course, a FD evolution, just like today's capacity upgrades, is incrementally introduced market-by-market based on service demand. The history of capacity upgrades suggests a rippling through the footprint over the course of many years.

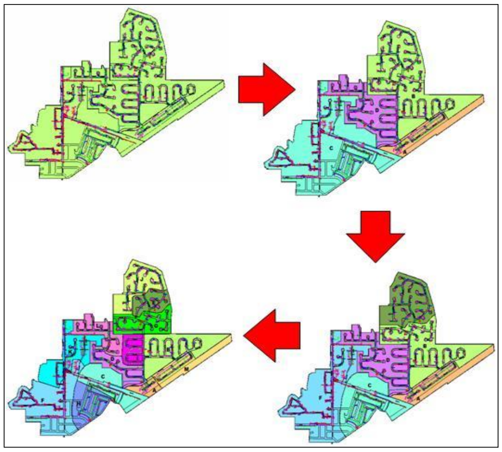

Figure 7 – Segmentation Phases of BAU Fiber Deeper Migration

The approach delivers the sought-after components of more capacity per equations (1) and (2):

1) Improved EOL SNR
2) Enabling of New Spectrum, B

What does N+0 mean to EOL performance? The performance of HFC networks in the downstream is well understood from decades of achieving fidelity acceptable for analog video. RF cascade reduction improves EOL SNR due to the reduced impact of the noise accumulation associated with active components. Physical node splitting shrinks average cascade depths. Note that service group splitting may still occur "virtually" though a segmentable node, but in this case there is no effect on the amplifier cascade depth.

On the upstream, SNR performance tends to be dominated by the optical link, with minor contributions from aggregated return path RF amplifier noise. However, the shrinking service group size has the significant effect of decreasing external interference funneling.

Quantifying the downstream performance, we can determine the effect when all RF degradations beyond the node are removed. The increased EOL SNR means an increased likelihood of supporting higher modulation formats.

Downstream performance is also largely set by the AM optical link. However, the subsequent amplifiers each noticeably nick away at SNR performance down the cascade. Using M-QAM SNR requirements, we can derive what QAM format can be supported over a range of HFC and home architecture variables. This is shown in Figure 8 [3,7].

Shown on the x-axis is expected EOL Composite Carrier-to-Noise using typical link length 1310 nm AM optics of 42 dB (N+6), 44 dB (N+3) and 47 dB (N+0), each labeled via the pink vertical lines. CCN is effectively the same as SNR, just comprised of AWGN *and* digital noise-like distortions), Note CCN is shown from *left to right* as highest to lowest.

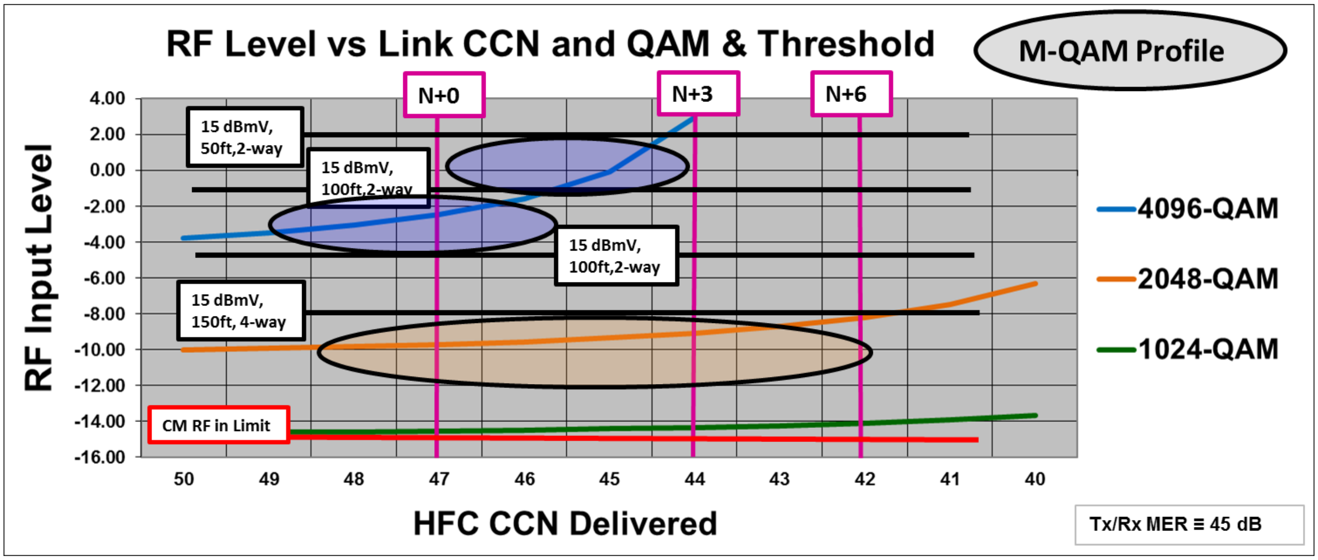

Figure 8 – HFC Cascade and Home Architectures vs. M-QAM Supported

In general, we can observe from Figure 8 that the higher the order of M-QAM format, the higher on the chart the base of its curve is and the further to the left that its upward trajectory begins. The upward trend indicates the CCN point below which it begins losing the ability to support the format. In other words, unsurprisingly, the higher CCNs of shorter cascades more readily support the most bandwidth efficient M-QAM formats, and we can predict under what conditions.

Specifically focusing on Fiber Deep, we can zero in on the "N+0" vertical line in estimating QAM expectations. We note that, for example, that a Fiber Deep (N+0) system comfortably supports 4096-QAM over a range of input levels, while for N+3, the CM must see about 2 dBmV or above at its input, and for N+6 we cannot support 4096-QAM at all for the range of scenarios shown.

Additionally, for N+6, the input level to the CM must exceed -8 dBmV in order to use 2048-QAM. Otherwise, only 1024-QAM can be supported. This chart and variants of it [3,7] quantify the N+0 advantages in EOL SNR performance in terms of bandwidth efficiency.

Use of Digital Optics

While Figure 8 is encouraging with respect to the ability to extract more bandwidth in Fiber Deep systems, another potential avenue for future Fiber Deep migrations is to employ a digital optical solution in place of classical AM optics.

The example of Figure 8 assumes classic 1310 nm point-to-point optics. Meanwhile, digital narrowcast service growth, continued node splitting, and site consolidation has driven operators to use of WDM technologies to maximize fiber usage. The trade-off parameter space for AM optics includes link length, wavelength band, number of wavelengths, and MER performance. These are carefully balanced for optimal deployment, creating architecture or performance constraints.

A digital optical link largely eliminates optical link length issues, increases the number of wavelength per fiber, and delivers a fixed (DAC limited) CCN performance at the node output. The node output SNR performance matches the performance of DRFI-compliant equipment installed in Headends, such as the output port of a CMTS or EQAM. Therefore, when digital optics is combined with the removal of RF amplifiers, nearly all the EOL performance variations are removed and link performance becomes extremely consistent and predictable. As such, it will enable consistent use of the most bandwidth-efficient DOCSIS 3.1 modulation profiles.

There are multiple solution types that offer a digital-to-the-node solution, and the industry is still vetting these options. In general, the available options are less about technology, and more about overall architecture, interfaces, and long-term network strategy. Fiber Deep merely creates the unique opportunity to define and implement a modular node solution enabling of a digital optical approach. Furthermore, the transition to digital optics is a natural and necessary step to a long-term plan that enables FTTP.

Downstream Spectrum: 1.2 GHz

In addition to the improvement in system performance as amplifiers are removed, the passive last mile provides a unique opportunity for considering changes to spectrum allocations, which are now much more efficiently available than with a cascade of RF amplifiers. In the downstream, with the defined extension in the DOCSIS 3.1 standard to 1.218 GHz and optionally to 1.7 GHz, Fiber Deep creates that unique opportunity to implement such an extension. (Note we will use “1.2 GHz” throughout the rest of this paper to represent 1.218 GHz).

One of the significant instantaneous advantages of Fiber Deep is the immediate access to 1 GHz of available bandwidth across most areas using current products. With plants more likely to be limited by RF actives than Taps, the bypassing of the amplifier cascade for an optimally placed N+0 node frees up approximately 130-250 MHz of heretofore unused spectrum in 750 MHz and 870 MHz systems, because a standard new node product today will be at least 1 GHz. As we shall see, this new bandwidth is offset somewhat by an extension to 85 MHz that will take place in the upstream, but this significant new bandwidth adds valuable shelf space for capacity growth.

While the bandwidth added by 1 GHz is quite powerful, there are nonetheless several key reasons to take advantage of the DOCSIS 3.1 extension of bandwidth beyond defined to 1.2 GHz. Among these are:

- Guarantees that bandwidth is set aside for DOCSIS 3.1 regardless of services added below 1 GHz
- Ensures a full 192 MHz block (and Gbps services) of DOCSIS 3.1 can be turned up in fully utilized 870 MHz systems
- Enables the potential for 10 Gbps of total downstream capacity via DOCSIS 3.1
- Offers more bandwidth flexibility for known service additions:
    - Time flexibility to complete all-IP transition
    - Addition of 4kHD video
    - Move to Cloud-based DVR (cDVR)
    - Implement Multi-Gbps service rates
    - Growing footprint of WiFi APs
- OSP equipment housings are already bandwidth-capable and the necessary RF actives available in CY14.
- Fiber deep (N+0) offers the uniquely efficient opportunity to make spectrum adjustments
- Taps that are in place are suitable as is in some cases, or need a faceplate change only
- 1.2 GHz keeps MoCA™ available (through channels defined from 1300-1675 MHz) as an HLAN option

MoCA™ has been widely deployed as Home LAN (HLAN) technology in recent years. In using 1.2 GHz for downstream services, most of the defined MoCA™ band is not overlapped. However, this would not be the case extending to the 1.7 GHz DOCSIS 3.1 option.

In addition to MoCA™ management, other potential issues were evaluated as part of the assessment for moving to 1.2 GHz. Two of the more important ones were powering and SNR loss.

1) Power requirements of the RF actives will increase roughly proportionally to the RF loading increase. However, in the larger picture of a power-saving Fiber Deep architecture, the effect is expected to be a wash, if not a net gain (less power) overall
2) There is slight SNR loss due to RF loading (using AM optics), but again a net gain in EOL performance from Fiber Deep is expected overall.
3) There are unknowns of optical nonlinearity performance at the extended bandwidth. However, the risk of major issues is considered low.

Note that we have seen in Figure 5 that bandwidth to 1.2 GHz in Taps can vary. Nonetheless, some families of deployed Taps are actually quite well-behaved above 1 GHz, and virtually all major remaining Tap vendors have developed technology that enables the existing fielded Taps to be expanded to up to 1.7 GHz with a simple change of the faceplate. Others have built such faceplates for a different vendor's Taps.

Example Tap response showing "excess bandwidth" on a 1 GHz Tap, and the faceplate-style bandwidth extension for it to extend to 1.7 GHz, are shown in Figures 9a and 9b.

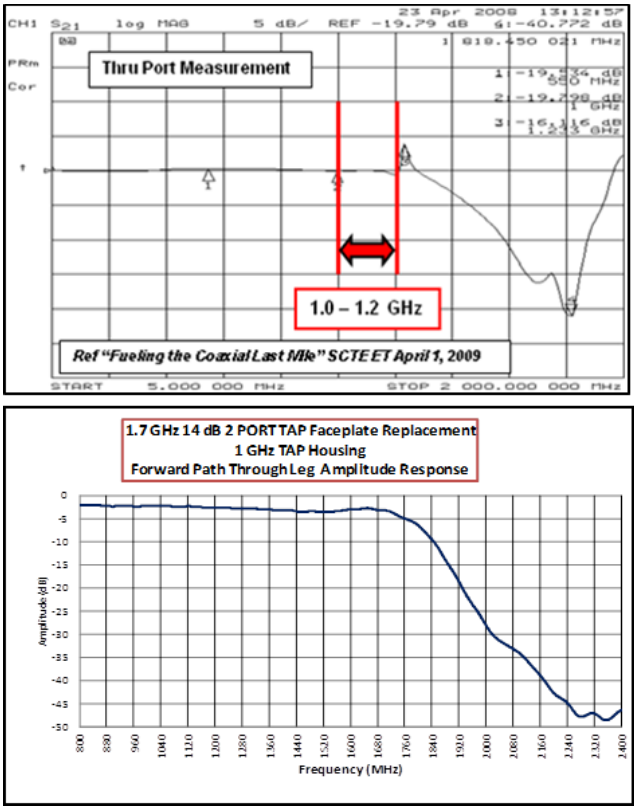

Figure 9 – a) (Top) Some 1 GHz Tap Models Have Quality "Excess Bandwidth" b) (Bottom) Tap Bandwidth Can Sometimes Be Easily Expanded with new Faceplates

No Free Launch

Now let's take a closer look at 1.2 GHz actives. For an N+0 architecture, of course, this means only the node. For the optics, because the transmitter loading is flat, the SNR cost is only about 0.6 dB relative to 1 GHz loading. Again, however, the small delta is more than offset by the SNR gains associated with the elimination of amplifiers altogether.

More significantly, however, the RF spectrum launched onto the coax is up-tilted. Thus, adding RF loading at the top of the spectrum disproportionately increases total power. Figures 10 demonstrates this effect.

In Figure 10, the extra 200 MHz to 1.2 GHz slightly more than doubles – 4 dB – the total RF power load using a particular higher N+0 tilt design. Fortunately, modern Gallium Nitride (GaN) RF technology is trending towards this higher total output power at equivalent distortion performance. It is expected that this technology will enable 1.2 GHz of bandwidth at, or very close to, equivalent 1 GHz performance and levels, including extending the tilt line all the way to 1.2 GHz. Or, conversely, these devices will deliver identical performance at or very close to the full 1.2 GHz bandwidth.

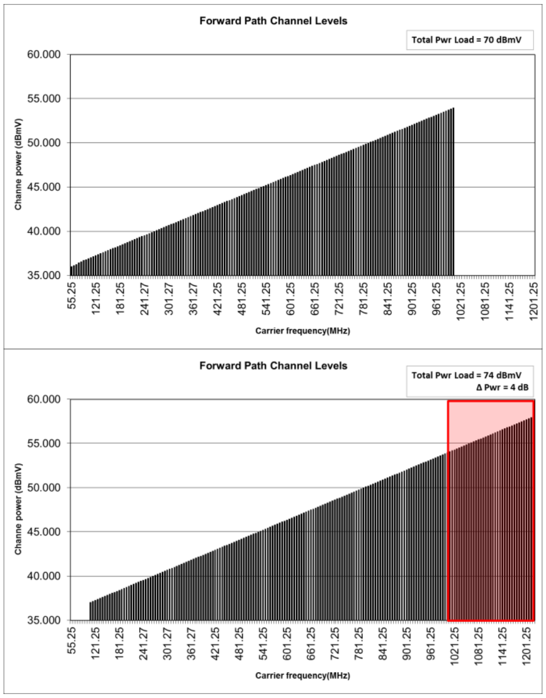

Figure 10 – Tilted RF Outputs Increase Total Power Disproportionately with Increasing Frequency

Now consider Figure 11, where the bandwidth is extended to 1.7 GHz. In this case, loading the spectrum at the identical tilt and equivalent relative PSD requires 9 dB more RF power, *nearly an 8x increase*. The impact on powering of nodes and plant powering overall to drive this additional RF power is substantial. It is not practical within the constraints of HFC as we know it, or can imagine reasonably evolving it cost-effectively in a useful timeframe.

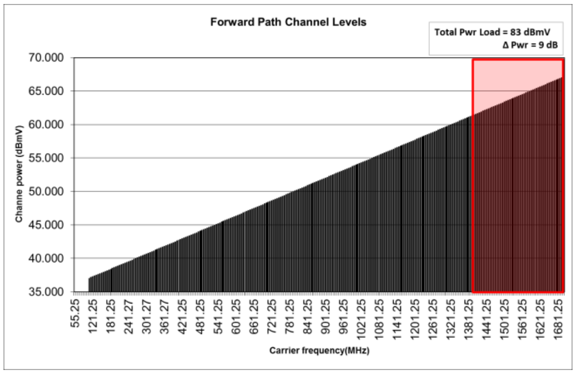

Figure 11 – Downstream Bandwidth Extension to 1.7 GHz Creates an RF (and AC) Power Dilemma

Other considerations that finalize the case against a 1.7 GHz extension are:

1) Optical loading SNR loss for the 1.7 GHz case would be significant (~2.5 dB) should it be carried over AM optics. Other optical nonlinear effects are unknown for what is now almost a full new octave. Only an overlay (parallel) architecture would make sense.
2) The OSP housings and RF interfaces themselves do not easily extend to 1.7 GHz, requiring substantial redesign of node platforms and possibly of fundamental materials used such as circuit board dielectric.
3) When extending to 1.7 GHz, the entire MoCA™ band is overlapped. A point-of-entry demarcation architecture is a requirement in this case. By restricting the downstream to 1.2 GHz, most of the MoCA™ band remains accessible without this requirement.
4) Sufficiency of 1.2 GHz plus FD serving group size reduction to support aggressive long term CAGR.

In summary, the addition of 1.2 GHz of bandwidth is a manageable extension using current techniques for optical loading, RF distribution, and powering, including tilt line extensions of tilts used today. It supports projected bandwidth needs for the long term, and technology availability is just around the corner. This is not the case for enabling up to 1.7 GHz. The recommended evolution path is therefore 1.2 GHz.

Upstream Spectrum: 85 MHz

A key component to the long-term enabling of HFC is more upstream capacity. The capacity bounds of 37 MHz, and considering that the bottom spectral portion can be less capable, drives the need to pursue a wider spectrum allocation to manage continued, albeit slower, traffic growth. Multiple upstream carriers are deployed in most markets today, and each new carrier consumes more of the total available bandwidth. Thus, looking ahead to project the timing of introducing more spectrum, while factoring in what DOCSIS 3.1 has to offer to 5-42 MHz, is important. The actual implementation of spectral re-allocation requires planning and coordination.

DOCSIS 3.0 already defines the 85 MHz mid-split for a wider upstream spectrum. Fortunately, technology supporting this band is already mature. In Figure 12, typical performances of an upstream DFB at nominal link length and Digital Return over the 85 MHz mid-split bandwidth are shown [3]. M-QAM performance thresholds, with typical margin associated with ensuring robustness in the less predictable upstream, is built-in using DOCSIS 3.1 QAM profiles.

Digital returns have a distinct advantage of not degrading NPR as a function of RF loading, and not being sensitive to the link length for performance or receiver output level setting. Performance is completely determined by the effective number of bits (ENOB) of the A/D converters.

Figure 12 shows that with new DOCSIS 3.1 FEC, the link is 1024-QAM capable using high performance DFB optical links and digital return systems available today over the full 85 MHz and with the dynamic range of NPR typically required. In fact, the NPR performance looks capable of supporting 2048-QAM in theory (3 dB higher than the 1024-QAM threshold). However, in practice, the margin that exists between the estimated 2048-QAM threshold and the NPR curve itself is less than what is typically seen as robust in the upstream.

Nonetheless, Figure 12 indicates why all of the M-QAM formats in Figure 2 are worth considering for the upstream as well as the downstream. Technology and FD architecture variables are falling into place to make these possible in the plant, shifting the performance burden to the complex task of high fidelity burst transmitters and high sensitivity burst receivers.

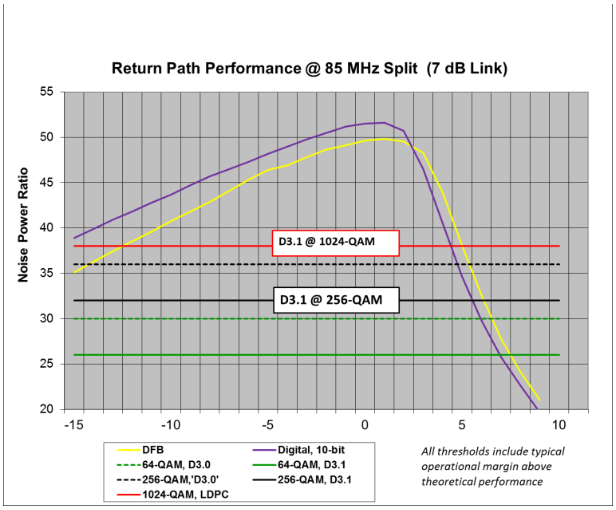

All thresholds include typical operational margin above theoretical performance

Figure 12 – Mature 85 MHz Technology Supports DOCSIS 3.1 QAM Profiles

New Capacity = New Spectrum

The highest spectral efficiency gain to expect in the downstream is 50% (256-QAM @ 8 bps/Hz to 4096-QAM @ 12 bps/Hz). In the upstream, the increase is now up to 67% using 1024-QAM. These are the "SNR" related capacity gain components in (1) and (2).

We now turn to the "Bandwidth" component of new capacity in (1) and (2). In the upstream, the relatively large (>2x) bandwidth extension to 85 MHz is particularly attractive. In the downstream, extending to 1.2 GHz is extending a bit into the unknown and a more modest addition to the current maximum forward band definition of 1 GHz. By contrast, in the upstream we are instead extending a partially troubled channel into an area where it will typically be better behaved. The upstream generally becomes increasingly cleaner with frequency above about 15 MHz in North America. And, as we have shown, the 85 MHz mid-split is available in current DOCSIS 3.0 and HFC technology.

Perhaps most importantly, as we shall quantify, an 85 MHz upstream combined with FD segmentation offers a long lifespan window. In addition, the 85 MHz mid-split offers the opportunity for Nx100 Mbps service rates, whereas 100 Mbps can be a challenge in DOCSIS 3.0 systems with 5-42 MHz of spectrum.

Similar to the downstream spectrum extension recommendation, the upstream "optional" edge defined in DOCSIS 3.1 to 212 MHz is *not* the recommended path, doe several reasons:

1) Unacceptable SNR degradation using DFB return optics, due to increased RF loading sharing a fixed total power, incurring about a 7 dB penalty over 5-42 MHz
2) Costly digital return A/D conversion and optics for high-fidelity sampling of very wideband spectrum
3) Significant implications to return path set-up, alignment, and technology implementation due to the increase in frequency dependent cable losses
4) Perhaps most importantly, removes significant downstream spectrum, where the higher traffic growth exists

The recommended return path architecture is therefore an 85 MHz mid-split.

KEY ARCHITECTURE PRINCIPLES

The major defining principles of HFC architecture migration are therefore:

- Implement Fiber Deep (N+0) for better SNR, efficient access to spectrum reallocation, reduction of service groups, and savings in maintenance, MTBF, and Opex.
- Target the N+0 design to be service group sized for long-term bandwidth runway as well as alignment with future potential architectures.
- Take advantage of cost-effectively accessible additional spectrum for the efficient enabling of new capacity.
- Use all spectrum most efficiently by taking advantage of key advances in M-QAM technology, Forward Error Correction techniques, and optimal waveform design (i.e. DOCSIS 3.1).
- Architect the FD investment – the “last” fiber node – and define the OSP equipment requirements as a template that positions the network for FTTP extensions if or when a fiber last mile is required on either a targeted or large scale basis.
- Consider the FD migration as an efficient path to transition the network to digital optics.

Applying the above tools – Fiber Deep, 85 MHz Mid-Split upstream, expansion to 1.2 GHz downstream, DOCSIS 3.1, a modular Fiber Deep node template, plus a migration plan to all-IP – creates a one-touch sustainable HFC architecture for many, many years to come. Key architecture principles above are visualized in Figure 13.

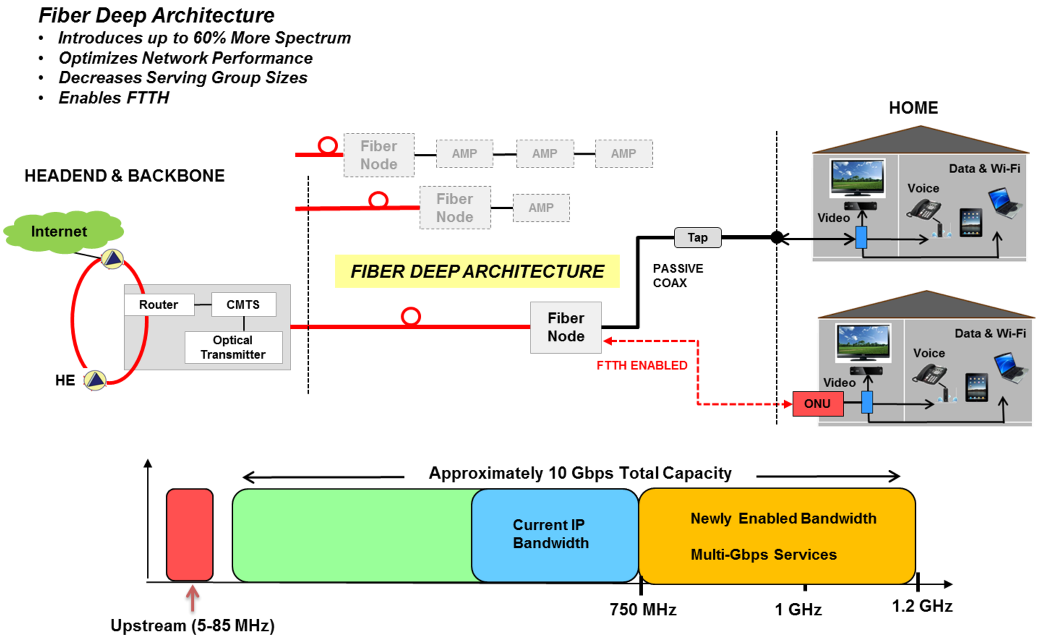

Figure 13 – Fiber Deep: Architecture Principles and Spectrum Allocations

A FINAL STEP

Using the above principles, an updated version of the downstream Capacity Management Timeline first demonstrated in Figure 1 is shown in Figure 14. It quantifies the fundamental case for one-touch long-term network lifespan based on a Fiber Deep architecture foundation. As the figure shows, the one-step FD migration path combined with the technology refresh components mentioned herein and designed to deliver long term capacity, does just that. As always, trying to predict service evolution, technology breakthroughs, and network evolution options ten years down the road, or even five, is difficult if not impossible.

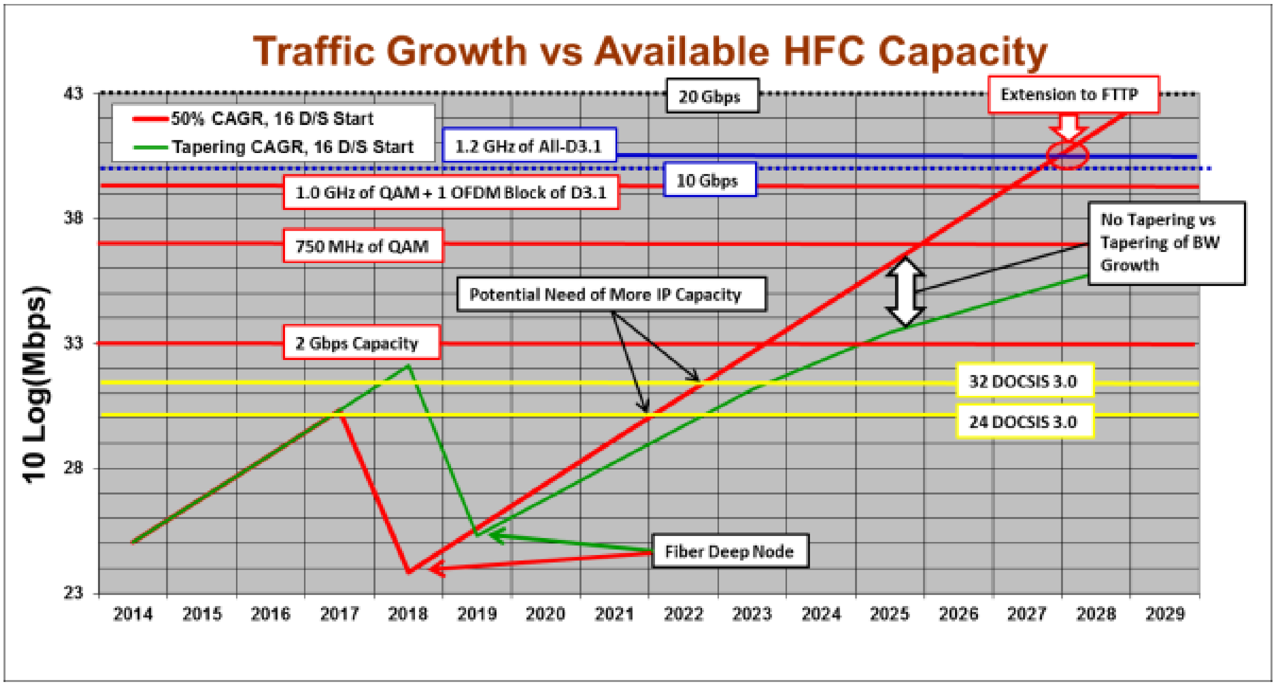

Figure 14 – Long-Term Downstream Capacity Enabled by the Fiber Deep Architecture

Looking at Figure 14, there are two trajectories – a constant 50% CAGR (red) with a Fiber Deep step taking place in 2017, and a trajectory with a built-in tapering of bandwidth growth with time (green). This is not to suggest the expectation of tapering, although this argument can be made [9]. More so, it is to show the power of compounding – or in this case, the power of *not* compounding as aggressively over time. It can result in what is essentially a “forever” network.

Note that the long runway of 15 years, even accounting for steady, aggressive CAGR, is comforting in that there is a solid window of time to observe trends and plan or re-plan accordingly. Even with this perspective of growth runway, it is important to point out that introducing Fiber Deep or FTTP throughout the footprint would inherently follow the common approach of targeting investments as driven by local service demand, as is done with any capacity upgrades, and thus the spreading out the capital investment over many years,

Examining Figure 14 further, the thresholds of 24 and 32 slots of DOCSIS 3.0 are again shown. The IP bandwidth growth trajectories reveal when these thresholds would be breached with continued bandwidth growth. As discussed, operators generally utilize their entire spectrum to provide the fullest set of services for customers. Various tools are used to manage the evolution of services, technology, and spectrum, such as for adding IP bandwidth to optimize the service mix to customers. Adding DOCSIS channels and splitting service groups are BAU operations used whenever the need arises.

The 24 and 32-channel thresholds shown in the figure show the modest impact in terms of years that incremental increases in IP bandwidth alone provide in the face of exponential growth. Without Fiber Deep or a node split, these thresholds are exceeded after 3-4 of years.

New thresholds have been added to Figure 1 that demonstrate the benefits of a Fiber Deep strategy and DOCSIS 3.1. As described previously, ensuring the full bandwidth efficiency benefits of DOCSIS 3.1 may also include a digital optical transition as part of the migration plan for Fiber Deep, and several solutions are under consideration for this path forward. Observing the persistently aggressive 50% CAGR IP growth curve and the implementation of the FD and technology strategy, it is not until 2028 (red circle on Figure 14) until there is a potential need to consider more capacity – 14 years away.

This conclusion is obviously important, but the intervening thresholds along the way inform us for setting the guidelines around the pace of the all-IP transition.

Summarizing, we note the following as key components that together achieve long term, one-touch network sustainability:

- Fiber Deep migration
- Downstream BW to 1.2 GHz
- DOCSIS 3.1
- Transition (implicit) to all-IP

Forever CAGR?

Perhaps the more intriguing long term case is if there is in fact a “Tapering CAGR.” As previously described, it would not be prudent to base a strategy on the expectation or requirement that consumer bandwidth growth will slow or end. History does not support this as a logical assumption. However, it would be negligent not to evaluate the possibility and recognize the implications.

The tapering trajectory is based on the understanding that streaming video has been the recent driver of persistently aggressive CAGR. Furthermore, a quantifiable maximum video bit rate and concurrency of subscribers at peak-busy-hour can be calculated. The tapering example here also assumes 4kHD plays a role in driving continued aggressive CAGR, but that formats beyond this are not significant contributors in scale for a number of practical reasons [2]. The conclusion of these assumptions suggest that a long term HFC capacity of 10 Gbps to a Fiber Deep sized serving group may be all that is necessary – *ever* – to satisfy consumer broadband as it relates to media consumption-centric applications.

As always, yet-to-be uncovered non-media applications may replace video as the CAGR growth engine. This is a key reason why keen attention to service growth trends over the next decade or so is critically important to determining future architectural evolution and timing required.

Every strategic plan, of course, is essentially a living document. Coarse corrections are constantly being made. However the actual growth of consumer bandwidth plays out, the lifespan enabled using the architecture principles described here provides a very comfortable window of time to assess the trends in services and technology, and develop appropriate responses. Based on our understanding of the key variables and driving forces, the strategies outlined represent the best plan of attack today.

Upstream

Figure 15 updates Figure 2 with the lifespan extension implications of the 85 MHz mid-split, Fiber Deep migration, and the eventual implementation of DOCSIS 3.1, which increase spectral efficiency up to 67% (64-QAM to 1024-QAM).

With a standard node split on a 42 MHz architecture, the 25% CAGR case managed about four years with a 3x 64-QAM (red dashed) starting point – again ignoring margin offset due to any potential peak busy hour underutilization on Day 1. With an 85 MHz mid-split, the lifespan is extended to last over 8 years (blue dashed).

Implementing Fiber Deep when combined with the 85 MHz mid-split architecture extends the lifespan to over 12 years. The 12-year range now closely synchronizes upstream lifespan to the downstream. It is this long-term capacity, the compatibility with legacy STB out-of-band signal frequency range (130 MHz maximum), the support of Nx100 Mbps services, and the opportunity that FD migration efficiently presents, that makes 85 MHz the right approach.

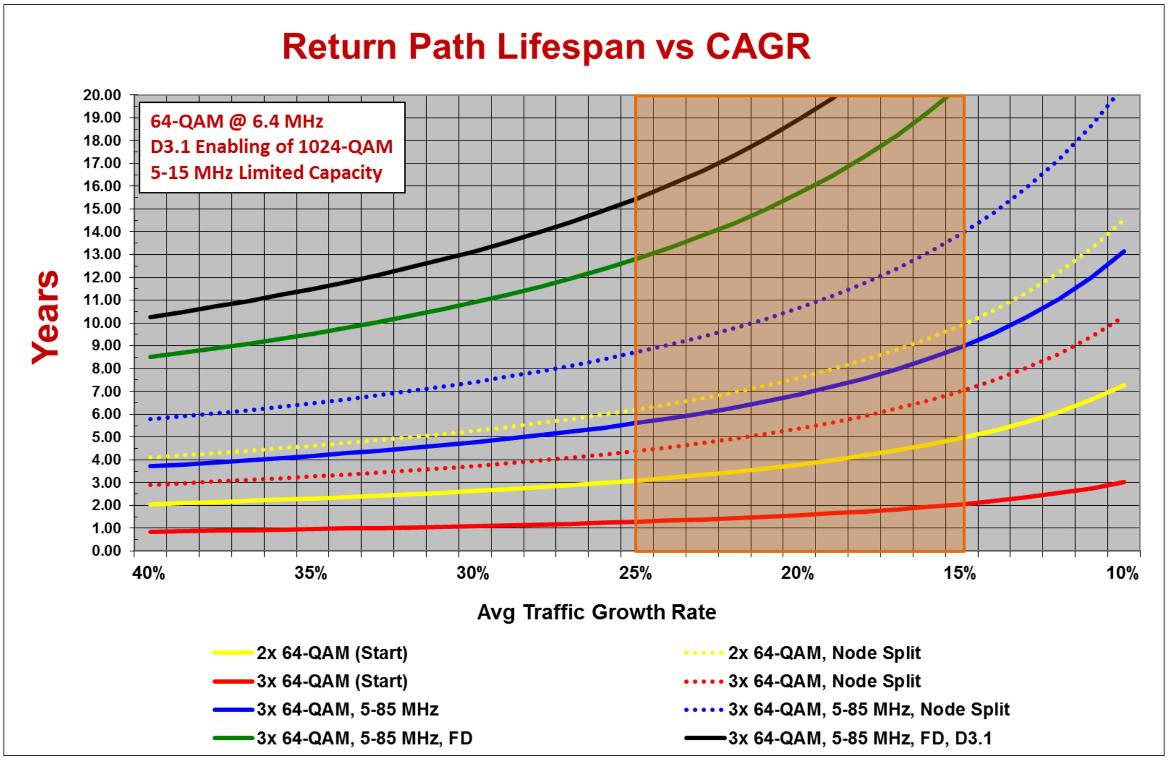

Figure 15 – Long-Term Upstream Capacity Enabled by the Fiber Deep Architecture

We now factor in that DOCSIS 3.1 is around the corner, so even more capacity potential can be baked into the calculation. This extends (black curve) the upstream expiration date out 15 years from an aggregate capacity perspective. These projections quantify and confirm why an 85 MHz upstream is the recommended spectrum option for Fiber Deep migration.

Lastly, one more aspect is worth looking at that further solidifies the upstream recommendation. As with the downstream, it will be important to keep a watchful eye on nominal CAGRs. We've discussed the upstream growth mostly around the high end of its recent CAGR range. However, with the more dynamic (relative to downstream) CAGR history of upstream in mind, note that the 85 MHz mid-split and a DOCSIS 3.1 Fiber Deep upstream still delivers *10 years* of life at a 40% CAGR – a CAGR range associated with actual *downstream* YOY growth in the past several years. In other words, if upstream CAGR does begin to resemble the downstream and see a large increase for a period of time, then the lifespan provided by implementing the above set of upstream strategies still delivers on a 10-year lifespan.

BEYOND FIBER DEEP?

In the N+0 FD architecture, the physical path from the Headend or Hub to the home is extended to now be (approximately, in physical distance) 98-99% fiber, with only the last 1-2% coaxial cable. The infrastructure itself has been heavily based on fiber optic technology and has taken advantage of the continued technology breakthroughs since all-coaxial networks first became HFC, and this will obviously continue.

Beyond maintaining the fundamental premise of a coax-to-the-home last mile, a core FD architecture strategy includes the concept of enabling the final launch of a fiber last mile if and when services drive this. In this sense – a network adaptable to the demands of customers – it is exactly how cable operators have been evolving their networks and services since they first built them. The cable industry, which pioneered AM optics, has also been deploying FTTP digital optical technologies for many years in very demanding business services and cellular backhaul applications. Thousands of FTTP Gigabit connections, readily available through years of investment in a rich fiber infrastructure and optical technology, exist in cable architectures. These FTTP systems, among other mature FTTP technology solutions, are natural candidates for use in residential services where it makes sense to do so, and enabling them is therefore considered a basic core requirement of a Fiber Deep node.

Summarizing, then, cable services can be run over traditional HFC, or Fiber Deep, or an all-fiber infrastructure. Furthermore, the shift to all-IP is erasing historical differences of service delivery and access, and IP is agnostic to whether it is an RF or fiber medium. The optimal solution will vary by situation, re-emphasizing the importance of a cost effective architecture based on flexibility and modularity, which cable operators have always embraced. Fiber Deep, and in particular the definition of the Fiber Deep node requirement itself, will continue this approach, and RF and fiber last mile access technologies will be options available.

Lastly, while architectures and engineering options dominate the thought processes in engineering departments, it is important to not lose sight of the fact that the choice of service provider by consumers is largely based on the services provided themselves and the servicing of them.

SUMMARY

Cable networks have demonstrated a remarkable adaptability to service demand, as well as the integration of the advanced technology required to support this demand. In recent years, operators have settled into a steady cadence of capacity upgrades that have delivered enhancements to best-in-class video, high-speed data, and voice, and have introduced a range of new services. These investments have been very effective, but the periodicity suggests a more comprehensive, "big picture" approach to network architecture and evolution may be more efficient.

We aimed to describe such a path in this paper, mapping out where the services and customer requirements are headed, how architecture and technology alternatives align to deliver on these requirements, the logic behind some of the key decisions, and the approach to positioning the network on a long runway of sustainability. The Fiber Deep approach infused with targeted architecture and technology refreshes produces a sound path for the next phase of network investment, and one that shall not need to be revisited for many, many years to come.

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

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

[4] Howald, Dr. Robert L, *Fueling the Coaxial Last Mile*, SCTE Conference on Emerging Technologies, Washington DC, April 2, 2009.

[5] 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.

[6] Howald, Dr. Robert, *The Grown-Up Potential of a Teenage PHY*, Cable Show Spring Technical Forum, Boston, MA, May 21-23, 2012.

[7] Howald, Dr. Robert L, *Time on Our Side: Using the Past to Pave the Way Forward*, SCTE Cable-Tec Expo, Orlando, FL, October 17-19, 2012.

[8] Howald, Dr. Robert L, Dr. Sean McCarthy, *Accounting for Techies: Taking it to the Ultra*, Cable Show Spring Technical Forum, Washington, DC, June 10-12, 2013.

[9] Howald, Dr. Robert L, Dr. Sean McCarthy, *Bits, Big Screens and Biology*, Cable Show Spring Technical Forum, Boston, MA, May 21-23, 2012.

[10] Howald, Dr. Robert L. and John Ulm, *Delivering Media Mania: HFC Evolution Planning*, SCTE Canadian Summit, Toronto, ON, Canada, March 27-28, 2012.

[11] http://www.ieee802.org/3/bn/public/adhoc_mmp/adhoc_mmp_urban_01.pdf

[12] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_01_0313.pdf

[13] http://www.ieee802.org/3/bn/public/mar13/howald_3bn_02_0313.pdf