

NETWORK MIGRATION STRATEGIES FOR THE ERA OF DAA, DOCSIS® 3.1, AND NEW KID ON THE BLOCK... FULL DUPLEX DOCSIS

AYHAM AL-BANNA
TOM CLOONAN
JEFF HOWE



TABLE OF CONTENTS

INTRODUCTION	3
DRIVERS BEHIND GIGABIT PER SECOND SERVICES.....	4
TECHNOLOGY ENABLING BANDWIDTH EXPANSION	8
Service group splits	8
Distributed access architectures.....	9
Remote PHY (R-PHY)	10
Remote PHY with virtual core (vCore)	10
Remote MAC/PHY (R-MACPHY).....	11
DOCSIS 3.1.....	12
FTTx	13
Fiber to the tap (FTTT)	13
Fiber to the home (FTTH), Fiber to the premise (FTTP).....	13
Full Duplex DOCSIS.....	14
ALTERNATIVE NETWORK MIGRATION PATHS	16
CONCLUSION	20
ABBREVIATIONS.....	21
REFERENCES	24

INTRODUCTION

The cable industry has achieved tremendous progress in offering high-speed data since the first DOCSIS specification was released in 1997. MSOs had the dual goal of meeting customers' demand for higher speeds and defending itself against competitive threats of speed wars with alternate technologies. As MSOs continue their network evolution, they are currently faced with no clear path since many options are available to augment their existing HFC networks.

For example, Figure 1 shows multiple potential evolutionary paths that the MSOs can select. The network architecture (e.g. I-CCAP/DAA/PON) is plotted against the topology which is presented here as the depth of the fiber in the network (e.g. HFC, FTTLA/FTTC, FTTT and FTTH).

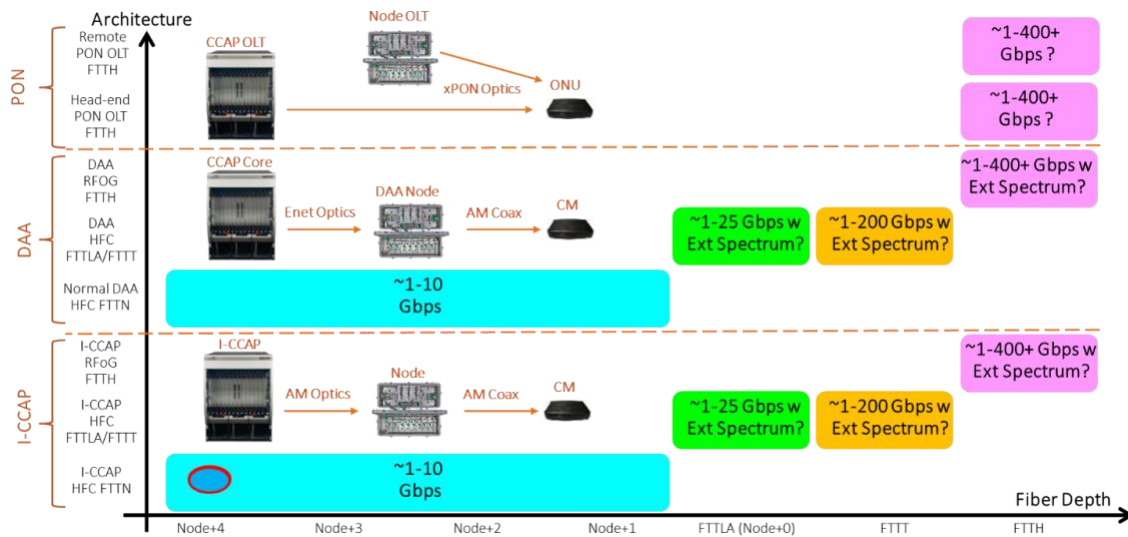


Figure 1 - Evolution of Cable Networks in the Next 2-3 Decades

The transitions between different phases of the same architecture or moving from one architecture to another will depend on the priorities and conditions of each MSO. Some MSOs may select to transition to a particular alternative at different times and different locations. For example, current HFC networks using the normal practice of node splits going to Node+0 (N+0) may be able to continue that practice until year 2025. At the same time, other MSOs may choose to move to an N+0 architecture in the immediate future. Similarly, if an extended spectrum DOCSIS technology develops, MSOs may choose to move to FTTT architecture as soon as 2025 in order to access even higher speeds. Finally, it is assumed that most MSOs may eventually choose to migrate their networks to FTTH over the next decade or two. Note that the capacities of all architectures (I-CCAP/DAA/PON) in an FTTH environment in the 2030 time frame are

assumed to be similar (~400 Gbps+) because those architectures will likely leverage similar technologies at that time.

Given the large combinations of the various network architectures shown in Figure 1 and different fiber depth topologies, selecting the appropriate architecture/topology transition path is not a trivial task. The challenge at hand is to understand the available technology enablers to assist in selecting the appropriate transition path. These technology enablers include node splitting, DAA, DOCSIS 3.1, spectrum management and reclamation, FTTx, selective subscriber migration (SSM), extended spectrum DOCSIS, Full Duplex DOCSIS (FDX) and others. This paper will examine the forces that are driving MSOs to provide symmetric multi-gigabit per second service, the technologies that will assist them in getting to those services, and the factors that will help guide them down the alternative migration paths that are available.

DRIVERS BEHIND GIGABIT PER SECOND SERVICES

For many years, studies have indicated that downstream internet traffic has been experiencing a ~50% compound annual growth rate (CAGR). For almost 35 years, this growth rate has shown itself in the maximum downstream sustained traffic rates, also known as the “billboard bandwidths” that service providers have offered to their subscribers. The 50% CAGR of maximum downstream sustained traffic rates is often reported as Nielsen’s Law and is depicted in Figure 2. The same trend, with slightly more variation, can also be seen in the average downstream bandwidth consumption rates of subscribers. Upstream billboard bandwidths and average upstream bandwidth consumption rates are more varied and typically have CAGRs with growth rates less than the 50% in the downstream at different MSOs. Looking at the upstream long term trend and projecting the curve over the next 15 years indicates a significant increase in bandwidth consumed by subscribers. Therefore, MSOs need to map out their network migration strategies to meet those future needs.




Figure 2 - Nielsen's Law of Internet Bandwidth

Although Nielsen's Law held fairly true to form over the past 35 years, there are indicators that it may deviate from its historical trend. One area is the growth of upstream bandwidth. Historically, the upstream bandwidth has lagged behind the downstream usage. Much of this trend was driven by how subscribers used the Internet. Users primarily accessed content online and downloaded it to their PCs, either doing large file transfers or viewing web pages. The upstream traffic was typically limited to protocol acknowledgements. As a result, access protocols such as DSL and DOCSIS evolved along an asymmetric path. Recent factors are compelling us to re-examine this trend.

PON technology supports symmetric bandwidth for the upstream and downstream. Although subscribers did not initially have a need for symmetry, MSOs were subjected to competitive pressure from PON providers because symmetric service was something that MSOs could not easily offer. Further, the usage of the Internet itself is experiencing a shift. Historically, usage was primarily in the downstream direction. However with new applications like YouTube that allow users to upload video, and cloud-based services like file storage and backup, there is a dramatic increasing in demand for upstream bandwidth. This shift is expected to result in a discontinuity in the upstream bandwidth which will take a step function upward to become close to the downstream curve.

Tempering this projected jump in upstream bandwidth demand is an apparent slowing in bandwidth consumption. Historically, access technology was the gating factor in usage, in that demand for bandwidth exceeded the ability of the MSOs to provide it. Any

time that the service limit was increased, the average usage went up by a corresponding amount. However, data that has been collected from several MSOs appears to indicate that although the billboard bandwidth continued to grow at 50% CAGR, the average bandwidth during the current decade did not grow nearly as fast.

During the decade of 2010, the average MSO downstream bandwidth usage grew at a 36% CAGR (Figure 3). The average MSO upstream bandwidth usage grew at a 17% CAGR (Figure 4). One way to interpret these statistics is that technology has finally allowed bandwidth supply to catch up with, and exceed, bandwidth demand. Another possible interpretation is that Internet traffic is becoming more bursty indicated by an increase in spread between peak and average utilization.

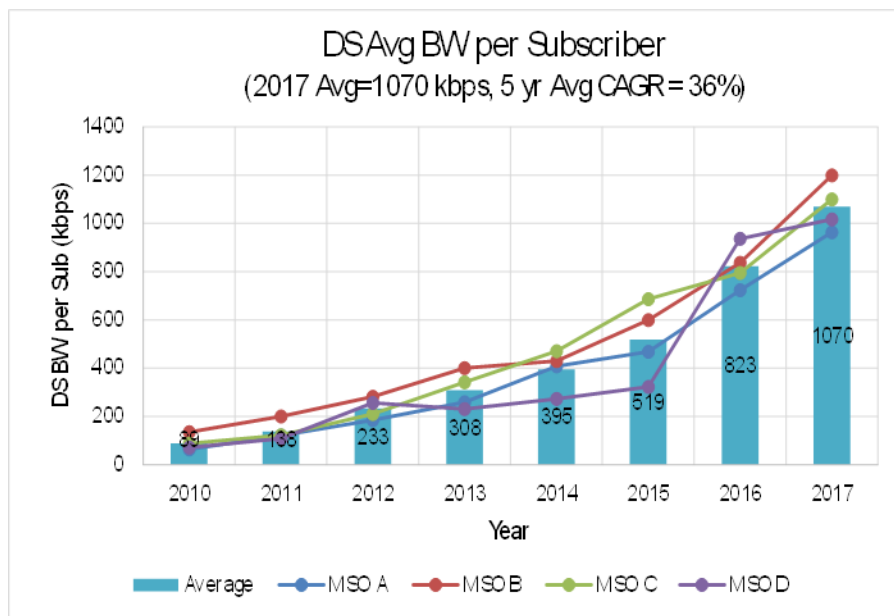


Figure 3 - Average DS Bandwidth in the 2010s




Figure 4 - Average US Bandwidth in the 2010s

Solutions such as DOCSIS 3.1 and the upcoming FDX technology which increase the available upstream bandwidth, have effectively addressed the increase in demand for upstream bandwidth. If MSOs have indeed aligned supply with demand, they may be able to slow down the growth in their advertised billboard bandwidth rates. Combining upstream bandwidth supply with a slowing of the downstream growth rate to 40%, results in a modified Nielsen's Law curve, shown in Figure 5. While network migrations will present challenges for MSOs, the modified bandwidth growth curve indicates that current HFC infrastructures will remain viable for at least another 15 years.

NIELSEN'S "SLOWED" LAW OF INTERNET BANDWIDTH (Growth Rate =40%/YEAR)




Figure 5 - Modified Nielsen's Law Curve

TECHNOLOGY ENABLING BANDWIDTH EXPANSION

Service group splits

Service group (SG) splits, sometimes referred to as node splits, have long been a trusted tool used by MSOs to reduce the bandwidth demands. The method is to divide the subscribers connected to a single SG into two smaller groups. Ideally, roughly half the group is connected to the original SG, while the other half of the subscribers is re-connected to a different SG.

There are several different ways that an MSO can perform SG splits. Traditionally, when there were multiple nodes per CMTS SG, the splitting would occur solely in the headend or hub, adding new CMTS ports and reconfiguring the RF combining network. Once there was a one to one mapping between an SG and a node, the next step was to segment the node into multiple SGs (e.g. 1x1 to 2x2 to 4x4). This was often accomplished using wavelength division multiplexing.

Today, since many nodes have already been segmented, the next step is to actually "split" the fiber node. This method pulls fiber deeper into the network and installs new nodes in the system. These new nodes may have one or more new SGs connected to

them. Thus, two separate fiber nodes and their associated feeds are required to support the new pool of subscribers, adding a cost for the MSO.

It's important to note that node splits offer no change in the service group bandwidth requirements for broadcast services. The signal is merely replicated by RF splitting in the headend and sent to the new fiber node. For example, if the MSO needed 50 quadrature amplitude modulation (QAM) channels to support broadcast video prior to the node split, the MSO will still require 50 QAMs to deliver that service after the node split.

However, the principle benefit of the SG split is associated with narrowcast services such as switched digital video (SDV), VoD, and DOCSIS. The key benefit of the split is to effectively double the capacity per subscriber for the narrowcast services. SG splitting oftentimes permits MSOs to free up some of the narrowcast video QAM spectrum. However, since the reason for doing the split is usually to increase the DOCSIS bandwidth per subscriber, the number of DOCSIS channels typically remains the same.

As node splits are performed and fiber is deeper in the network and closer to the subscribers, the network eventually reaches the point where the fiber node is the last active device in the outside plant. This plant topology is referred to as Fiber to the Last Active (FTTLA) and as Node+0 (N+0), as the fiber node has no amplifier or other active component following it.

As each split occurs and the number of subscribers in the service group is reduced, the number of amplifiers is also reduced. Therefore, the noise funneling effect from multiple subscribers is minimized. However, once the topology reaches N+0, there are diminishing returns for doing further node splits. Reaching an N+0 topology is an important milestone for an MSO, because it is also a prerequisite for migrating to FDX technology. More in-depth discussion on node splits can be found in [CLO1].

Distributed access architectures

Some MSOs will likely be able to support their video and HSD services using traditional headend-based integrated CCAPs (I-CCAP). However, other MSOs may be planning to perform node splits more frequently, needing to support more service groups than would be supported by an I-CCAP chassis. Adding additional I-CCAPs may cause issues related to the required power and/or rack-space. However, there is an alternate access architecture that helps MSOs address the required power and rack-space challenge within their headends. This technique employs distributed access architectures (DAAs) which are also a necessary component of implementing Full Duplex DOCSIS.

There are several types of DAAs being proposed for use in the future, and each proposal has its own sets of pros and cons [EMM1]. This paper provides a brief description for three of the architectures.

Remote PHY (R-PHY)

This approach separates the PHY (upstream and downstream) from the headend and places the full PHY layer, including the forward error correction (FEC), symbol generation, modulation and digital to analog converter (DAC)/analog to digital converter (ADC) processing, into the fiber node. This requires that these functions be removed from the headend equipment - CCAPs, CMTSs and EQAMs. The DOCSIS MAC processing remains in the MAC core within the headend. This approach is slightly disruptive, as it requires many pieces of headend equipment to be modified.

The R-PHY approach is an evolution of the modular headend architecture (MHA) approach with many significant requirements, such as the need to support upstream MAC/PHY separation, the need to support new timing interfaces that work over Ethernet and the need to add DOCSIS 3.1 support within downstream external PHY interface (DEPI) and upstream external PHY interface (UEPI). However, this approach offers benefits as well. Remote PHY helps with the nonlinear optical noise problem by using digital optics instead of analog, and it also addresses the headend power and rack-space challenge. Another benefit of the Remote PHY approach is that it provides investment protection because it allows MSOs to continue to re-use their headend-based CCAPs as part of the solution.




Figure 6 - Remote PHY

Remote PHY with virtual core (vCore)

An extension of the R-PHY approach is to virtualize the MAC core. Rather than using dedicated hardware in the headend to provide the CCAP core functionality, the MAC core functionality is virtualized and run on commercial off-the-shelf (COTS) compute platforms. This architecture will typically take more headend space than dedicated hardware specifically designed for the core functionality, but this alternative provides

other benefits. By decoupling the hardware from the software functionality, each can be updated independently. The virtual platform can be shared with other applications and can be easily scaled up and down to meet demand. Since the hardware can be scaled back when not needed, power savings can result. The architecture is based on software defined networks (SDN) and network function virtualization (NFV) techniques, which provide an infrastructure for rapid feature development.




Figure 7 - Remote PHY with Virtual Core

Remote MAC/PHY (R-MACPHY)

This approach places the entire upper and lower MAC (upstream and downstream) and the entire PHY layer functionality (upstream and downstream) into the fiber node. In effect, this places all of the CMTS, Edge QAM and CCAP functions into the fiber node and only requires a switch or router to remain in the headend. As a result, this approach is not as disruptive. Remote MAC/PHY also helps with the nonlinear optical noise problem, and also provides the maximum amount of power and rack-space savings within the headend (even more than the Remote PHY approach). By placing both the MAC and the PHY in the same location, it eliminates the DEPI and UEPI protocol overhead. When appropriately modified, existing headend CCAPs can be used as dense aggregation routers (or repurposed PON OLTs) feeding the remote CCAPs as well.

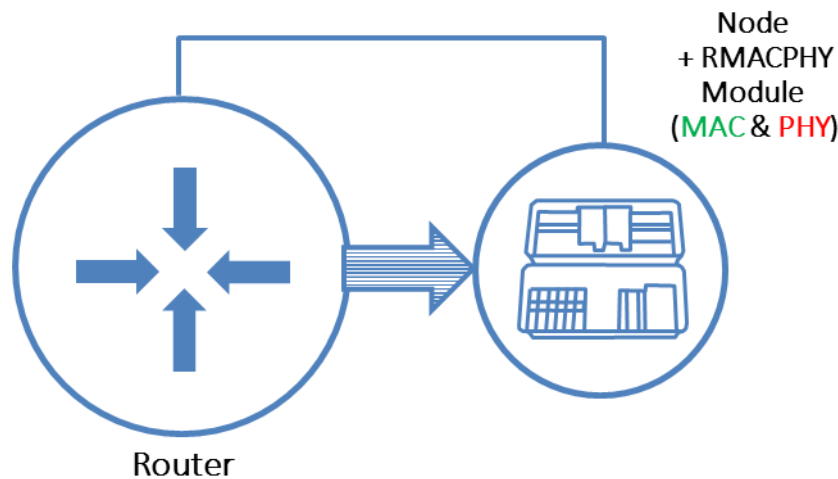


Figure 8 - Remote MAC/PHY

DOCSIS 3.1

DOCSIS 3.1 is a backwards-compatible augmentation to the DOCSIS 3.0 specification that provides better spectral efficiencies (more bps/Hz) and wider channels for both the downstream and upstream paths. The specification provides improved spectral efficiencies via many techniques, including the use of:

- Orthogonal frequency division multiplexing (OFDM) modulation
- Higher modulation orders (4096QAM and higher)
- More efficient low-density parity check (LDPC) forward error correction
- Bit-loading to custom-fit the modulation orders to the varying SNRs across the spectrum of the HFC plant
- Multiple modulation profiles to provide different modulation rates to different CMs depending on their specific noise characteristics

Backwards compatibility is guaranteed because DOCSIS 3.0 and DOCSIS 3.1 channels can co-exist on the HFC spectrum. In addition, pre-DOCSIS 3.1 CMs will work with DOCSIS 3.1 CMTSs, and pre-DOCSIS 3.1 CMTSs will work with DOCSIS 3.1 CMs.

As a result of its power, flexibility and backwards-compatibility, many MSOs are looking to DOCSIS 3.1 to extend the life of their HFC plants by at least several years. The actual HFC plant life extension that will result from the use of DOCSIS 3.1 depends on many different factors, including: the annual subscriber bandwidth growth rates; the number of node splits; the level of investment that MSOs are willing to put into their plants to extend spectral width; and the quality of the HFC plant (i.e. SNRs).

FTTx

Migrating to an N+0 architecture means pushing fiber deeper into the outside plant and closer to the subscriber. This is just one flavor of what is referred to as FTTx, where x depends on how deep into the plant the fiber goes. In the case of N+0, this is also called FTTLA (fiber to the last active) or FTTC (fiber to the cabinet or fiber to the curb). There are also other types of FTTx architectures that can support subscriber bandwidth growth.

Fiber to the tap (FTTT)

Fiber can be taken beyond the traditional node location and run all the way to the subscriber tap. From this location, the coax cable run is much shorter, resulting in less attenuation that would enable an extended spectrum DOCSIS solution. Extended-spectrum DOCSIS refers to extending the spectrum used in cable networks above and beyond of what DOCSIS 3.1 can support [CLO2]. This can be effective in network topologies where no amplifiers or diplexers are present. The coaxial cables can support very high frequencies such as 25 GHz for RG-6 drop cables. Although attenuation will cause a reduction in the modulation orders that can be used, the extremely wide spectrum will allow much higher total bandwidths.

Fiber to the home (FTTH), Fiber to the premise (FTTP)

Running fiber all the way into the premises is the next and final step in running fiber deeper into the network. Passive optical networks (PON) is a technology that provides a direct optical link between the headend and the subscriber home. The device in the headend is called an OLT, and the device in the home is called an ONU or an ONT. Many ONUs or ONTs can share a single FTTH optical feed from the OLT in the headend, so the bandwidth capacity provided by a PON is always shared by all of the ONUs or ONTs connected to the PON feed.

PON technologies today include bandwidth capacities such as 1 Gbps, 2.5 Gbps, and 10 Gbps. Ultimately, 40+ Gbps bandwidths will also likely be provided. For MSOs, this is an overlay technology to the DOCSIS HFC delivery system, since it does not offer any form of backwards-compatibility to DOCSIS. PON will likely be used initially in business services and MDU environments, but it will also be able to serve elite residential subscribers once bandwidth demand exceeds what traditional DOCSIS systems can provide.

PON may find a few competitors in the FTTH space. One FTTH competitor to PON is RF over glass (RFoG). RFoG technology enables MSOs to transmit their standard RF signals (e.g. DOCSIS, MPEG-TS video, analog) all the way to the subscriber homes over fiber. It requires a special ONU to be placed within each home, where it is responsible for

performing an optical-to-electronic conversion function, similar to the function performed by a typical fiber node. RFOG offers several benefits to MSOs. It allows MSOs to begin transitioning their HFC plants into a FTTH plant (which is likely to be the plant of the future) while maintaining backwards compatibility with their large, existing CPE investment. RFOG eliminates the coaxial portion of the HFC plant, which can lead to improved SNRs and higher modulation orders. RFOG can extend DOCSIS 3.1 transmission systems to spectral widths that exceed the 1.2-1.7 GHz spectral limits of typical coaxial distribution systems within the HFC plant. Initial RFOG systems suffered from a type of noise called optical beat interference (OBI) that is sometimes generated when multiple ONUs transmit at the same time. However, there are now forms of OBI-free RFOG systems that eliminate this type of interference.

Full Duplex DOCSIS

Full Duplex DOCSIS is an enhancement to the DOCSIS 3.1 specification to enable significant increases upstream bandwidths. The goal is to be able to provide 10 Gbps downstream bandwidth and 5 Gbps upstream bandwidth within a single service group. In order to expand the upstream bandwidth while having minimal impact on the downstream bandwidth, FDX allows certain portions of the spectrum to be used for upstream and downstream transmissions simultaneously. The spectrum from 108 MHz to 684 MHz has been designated for these bi-directional transmissions. The updated spectrum usage is depicted in Figure 9.

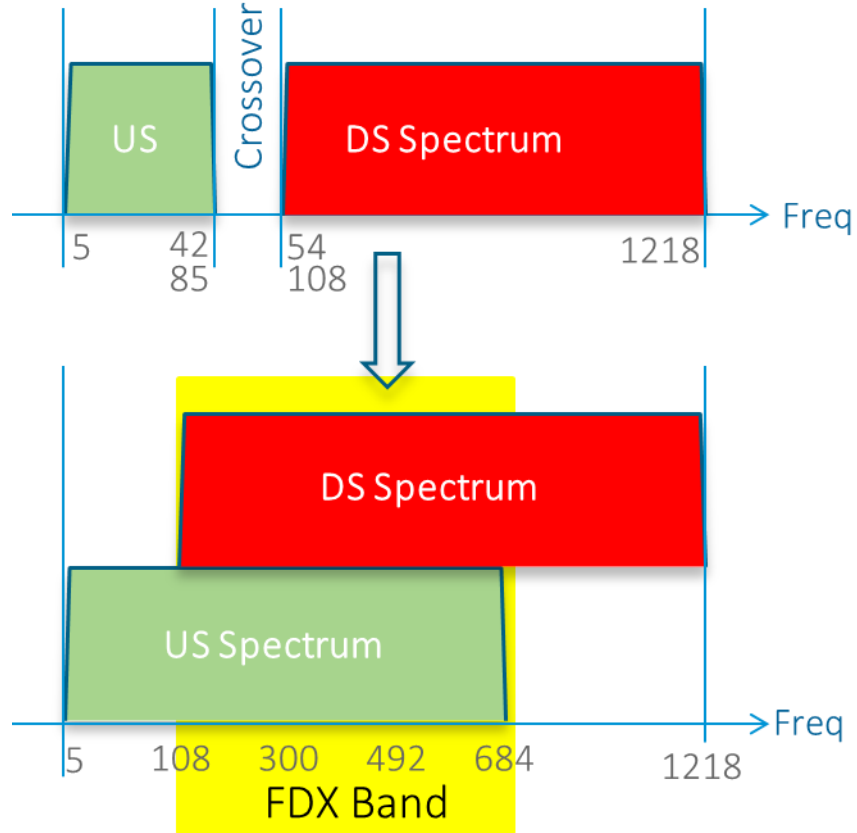


Figure 9 - FDX Spectrum Usage

It should be noted that the simultaneous transmission and reception of packets is from the fiber node point of view. Each individual CM will still be operating in a frequency division multiplexing (FDD) mode. CMs will be grouped together into transmission groups (TG). Each TG will use some channels in the FDX band as upstream channels and the other channels as downstream channels. However, one TG may be using one part of the spectrum as an upstream channel while another TG may use that same part of the spectrum as a downstream channel. In addition, usage of the spectrum for upstream and downstream within a TG can be changed over time. From a CM point of view, the FDX band operates as a dynamic FDD system, as illustrated in Figure 10.




Figure 10 - Dynamic FDD Operation from Cable Modem Perspective

ALTERNATIVE NETWORK MIGRATION PATHS

As described in the previous section, there are many tools for MSOs to choose from to support their network migration plans. These tools include node splitting and segmentation, DAA vs. centralized architectures, DOCSIS 3.1, HFC vs. FTTH, RFoG vs. PON for FTTH and selective subscriber migration. The optimal choice depends on the network parameters, demand and statistical distribution of subscribers among services, and the specific situation for each MSO with regard to logistics, operational and resource constraints, current infrastructure and budget. Additional factors to consider include current service group size and target service group size, the timing for transition from QAM video to IP video and time frame for needing to deliver symmetric services. Therefore, the ideal solution for one MSO may not be optimal for another.

As previously discussed, the rate of downstream bandwidth growth appears to be slowing to 40% per year. In the absence of any other network changes, this implies that nodes will need to be split or segmented approximately every 2.1 years in order to keep up with bandwidth demands. As the analysis in [CLO1] shows, the effectiveness of node splits is reduced each time it is split into a smaller service group. This happens because the peak bandwidth of a single subscriber impacts quality of experience more than the average bandwidth of all the subscribers in the service group. It is estimated that the point of diminishing returns is reached when the service group reaches approximately 50 subscribers. Depending on the current average service group size, node splits can

provide an effective migration strategy for many years to come. Table 1 summarizes how many years it takes to reach an average service group size of 50 subscribers.

Table 1 - Estimate HFC Plant Life Using Node Splits

Current Average Service Group Size	Years
100	2.1
200	4.1
300	5.3
400	6.2
500	6.8
600	7.4
700	7.8
800	8.2
900	8.6
1000	8.9

The above table assumes no other changes are made to the network. An additional migration strategy involves spectrum allocation. Although high-speed data (HSD) is the fastest growing service within an MSO's HFC spectrum, MSO managed video services still consume the largest percentage of the spectrum today. To accommodate the growing HSD bandwidth, MSOs may look to various technology paths that offer to squeeze the bandwidth of MSO managed video service into a smaller portion of the HFC spectrum. The future will likely see different MSOs using different mixes of SD broadcast video, HD broadcast digital video, SDV, VoD, IP video and analog video.

Over time, the analog video spectrum will be heavily reclaimed with many MSOs having entirely reclaimed it already. DTAs offer a good, low-cost technique for accomplishing that goal. Future media gateways with low-cost IP-STBs may also provide similar low-cost alternatives. SDV is another technique that can help to reclaim spectrum from the broadcast digital video tier, whereby video streams are only transmitted over a service group if a subscriber is viewing that stream. As SG sizes become smaller, SDV becomes more effective and can reclaim more legacy video spectrum.

In addition to transitioning away from analog video to digital video, as well as transitioning away from broadcast video to SDV, many MSOs are also looking to transition away from MPEG-TS-based QAM digital video delivery to IP-based video delivery over DOCSIS. There are several reasons for this trend, one being that DOCSIS provides better spectral efficiency over QAM digital video delivery [CLO1]. In addition, over-the-top (OTT) video delivery is becoming popular with subscribers and is delivered over IP. Over time, MSOs may migrate away from their managed QAM digital video delivery to their own OTT video delivery. Figure 11 depicts how downstream spectrum

may migrate over time, increasing the amount of spectrum allocated to DOCSIS, which increases the amount of available bandwidth.




Figure 11- Downstream Spectrum Migration

MSOs also have options when migrating their upstream to meet subscriber demands. The migration is dependent on transitioning CPEs to DOCSIS 3.1 CMs and eventually FDX CMs. The migration also depends on how quickly MSOs are required to increase upstream to provide symmetric services, typically to meet competitive pressures from providers like Google. Some MSOs may be able to meet the short term upstream bandwidth demands by migrating to an 85 MHz mid-split. Other MSOs may need to migrate quickly to provide large upstream bandwidths and migrate to a 204 MHz high-split.

The starting point of a split may dictate which path is taken to achieve the end goal of using the FDX band for upstream. When FDX becomes available, the manner in which the usage is shared between legacy D3.1 CMs and the new FDX CMs may depend on the diplexer in the legacy CMs. If the legacy D3.1 CMs are on an 85 MHz plant, they would not be able to participate in the whole upstream of the FDX band of 108 to 684 MHz (although with a software upgrade they could share the downstream FDX spectrum with the FDX CMs). The FDX CMs would be able to use the FDX band for upstream or downstream transmissions. On the other hand, if the legacy D3.1 CMs are currently

configured for a 204 MHz split, they will be able to share the spectrum from 108 to 204 MHz in the upstream direction with the FDX CMs, while the FDX CMs will be able to additionally use the spectrum from 204 to 684 for upstream bandwidth.

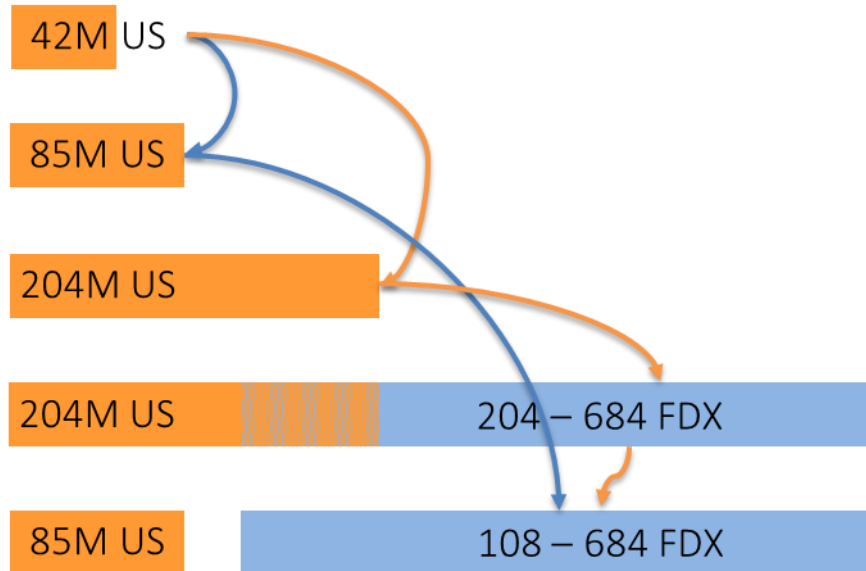


Figure 12 - Upstream Spectrum Migration

Although the FDX band is 576 MHz wide, MSOs may not want to use all the spectrum for FDX initially, or may limit how much spectrum can be freed from other services. The specification for FDX allows FDX CMs to use only a portion of the FDX band for FDX channels. However, the portion of the FDX band that is not being used for FDX channels can only be filled with legacy video QAM channels. This is because FDX CMs can only transmit and receive FDX channels in the portion of the spectrum reserved for the FDX band. The possible FDX band migration steps are shown in Figure 13. Depending on other factors such as spectrum availability, upstream bandwidth demand and FDX CMs penetrations, MSOs can choose how quickly or slowly to move through these migration steps.

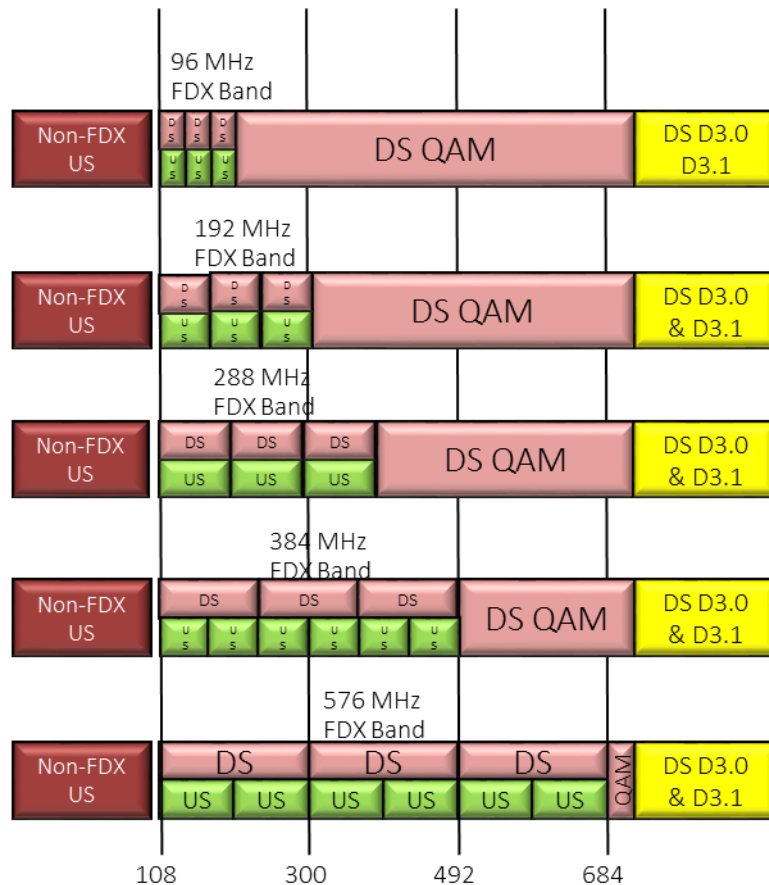


Figure 13 - FDX Spectrum Migration

CONCLUSION

With new demands driving bandwidth growth such as the competitive pressure to provide symmetrical upstream and downstream bandwidth services, some see FDX DOCSIS as the answer. However, not all MSOs are the same, so one technology is not going to be the right answer in every situation. MSOs will want to have a whole toolkit of technologies and procedures to address their network migration needs. Those tools include node splitting and segmentation, DAA vs. centralized architectures, DOCSIS 3.1, HFC vs. FTTH, RFoG vs. PON for FTTH and selective subscriber migration. Utilizing those tools offers network migration options as shown in Figure 14. Each MSO has a unique set of circumstances, so they must apply the set of tools in a combination that meets their specific goals and objectives. Aside from devising the right mix of tools, the MSO also has the flexibility to apply different tools at different times for different sites.

Figure 14 - Network Migration Options to Meet Bandwidth Demands

ABBREVIATIONS

ADC	Analog to Digital Converter
Bps	bits per second
CAGR	Compounded Annual Growth Rate
CAPEX	Capital Expense
CCAP	Converged Cable Access Platform
CM	Cable Modem
CMTS	Cable Modem Termination System
COTS	Commercial Off-The-Shelf
CPE	Consumer Premise Equipment
D3.1	Data Over Cable Service Interface Specification version 3.1
DAA	Distributed Access Architecture
DAC	Digital to Analog Converter
DCA	Distributed CCAP Architecture
DEPI	Downstream External PHY Interface
DOCSIS	Data Over Cable Service Interface Specification
DS	Downstream
DSL	Digital Subscriber Line
DTA	Digital Television Adapter
EQAM	Edge Quadrature Amplitude Modulator
FDD	Frequency Division Multiplexing
FDX	Full Duplex DOCSIS
FDX CM	Full Duplex Cable Modem
FEC	Forward Error Correction
FTTC	Fiber to the Cabinet or Curb
FTTH	Fiber to the Home
FTTLA	Fiber to the Last Active
FTTT	Fiber to the Tap
FTTx	Fiber to the 'x' where 'x' can be any of the above
Gbps	Gigabits Per Second
GHz	Gigahertz
HD	High Definition
HFC	Hybrid Fiber Coax
HSD	High- Speed Data
Hz	hertz
I-CCAP	Integrated Converged Cable Access Platform
ISBE	International Society of Broadband Experts
LDPC	Low-Density Parity Check
MAC	Media Access Control interface

MACPHY	DCA instantiation that places both MAC & PHY in the Node
MDU	Multiple Dwelling Unit
MHA	Modular Headend Architecture
MHz	Megahertz
MSO	Multiple System Operator
N+0	Node+0 actives
NFV	Network Function Virtualization
OBI	Optical Beat Interference
OFDM	Orthogonal Frequency Division Multiplexing
OLT	Optical Line Termination
ONU	Optical Network Unit
OTT	Over-The-Top
PHY	Physical interface
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
RF	Radio frequency
R-MACPHY	Remote MAC-PHY
R-PHY	Remote PHY
RFoG	RF over Glass
SCTE	Society of Cable Telecommunications Engineers
SDN	Software Defined Networks
SDV	Switched Digital Video
SSM	Selective Subscriber Migration
TG	Transmission Group
UEPI	Upstream External PHY Interface
US	Upstream
vCore	Virtual Core
VoD	Video on Demand

MEET OUR EXPERT: Jeff Howe

Jeff Howe is VP, Systems Engineering at ARRIS, responsible for directing architectural work and future product planning. Prior to that, Jeff served as Director, System Architecture at the start-up company Cadant, where he led architecture and integration work for the development of the next-gen C4 Cable Modem Termination System. Jeff has also served in managerial, architectural, hardware and software development, and test roles at Lucent Bell Laboratories, working on various data communications platforms including the Local Area Data Transport product, the 1PSS X.25 Switch, a frame relay prototype platform, the Globeview-2000 ATM Switch, and the PacketStar 64000 Advanced IP Router. Jeff has a BSEE from the University of Wisconsin, an MSEE from Stanford University, and an MBA from Capella University. His research interests include high-speed switching and QoS for next-generation converged services platforms.

REFERENCES

- (1) [CLO1] T. Cloonan, M. Emmendorfer, J. Ulm, A. Al-Banna, and S. Chari, "Predictions on the Evolution of Access Networks to the Year 2030 & Beyond," The NCTA Cable Show Spring Technical Forum 2014
- (2) [CLO2] Tom Cloonan et. al., "Lessons from Telco & Wireless Providers: Extending the Life of the HFC Plant with New Technologies," Spring Technical Forum conference INTX 2015
- (3) [EMM1] M. Emmendorfer, T. Cloonan, and J. Ulm, "A Side-by-side Comparison of Centralized vs. Distributed Access Architectures," The NCTA Cable Show Spring Technical Forum 2014