



Managing DOCSIS Capacity over HFC Networks

DOCSIS 2.0, 3.0 and Future 3.1

A Technical Paper Prepared for Cisco Live 2015

By

John J. Downey

CMTS Technical Leader

Cisco Systems

RTP, NC

919-392-9150

jdowney@cisco.com

Martin Mattingly

Manager, Business Development

Cisco Systems

Lawrenceville, GA

770-236-1338

mattinm@cisco.com

Table of Contents

Title	Page Number
Introduction _____	3
Number One Capacity Question _____	4
DS Capacity Overhead _____	5
Powerboost™ _____	5
Baseline Capacity _____	6
Speed Variables _____	7
D3.1 Capacity Planning Concerns _____	7
Spectrum Planning _____	8
D3.1 Spectrum Thoughts _____	10
US Level Swing over Temperature _____	12
Five Potential Migration Steps and Ideas _____	13
DOCSIS 3.0/3.1 DS Channel and Spectrum Allocation _____	13
DOCSIS 3.0/3.1 US Channel and Spectrum Allocation _____	14
1. 12 DSs - 5-42 MHz US & 750 MHz DS _____	16
25x Over Subscription Estimate _____	17
2. 24 DSs - Extended DS to 862 MHz _____	18
3. 32 DSs & D3.1 OFDM Block - Advanced DS to 1 GHz _____	19
b. 32 DSs & 32 DSs - Stick with D3.0? _____	20
4. 32 DSs & D3.1 OFDM Blocks - Goal of 85 MHz US & 1.2 GHz DS _____	21
5. 32 DSs & Multiple D3.1 Blocks - Remote-PHY _____	22
What Does This Bandwidth Graph Represent? _____	23
JD's 10 Rules for Capacity Planning _____	24
10 Points to Ponder _____	25
Checklist _____	26
Closing Points _____	26
Abbreviations _____	27
White Papers, References and Resources _____	31

List of Figures

Title	Page Number
FIGURE 1 - DOWNSTREAM SPECTRUM PLANNING	8
FIGURE 2 - UPSTREAM SPECTRUM PLANNING	9
FIGURE 3 - D3.1 DOWNSTREAM SPECTRUM THOUGHTS	10
FIGURE 4 - D3.1 UPSTREAM SPECTRUM THOUGHTS	10
FIGURE 5 - US LEVEL SWING OVER TEMPERATURE	12
FIGURE 6 - DS CHANNEL ALLOCATION	13
FIGURE 7 - DS SPRECTUM ALLOCATION	14
FIGURE 8 - US CHANNEL ALLOCATION	14
FIGURE 9 - US SPECTRUM ALLOCATION	15
FIGURE 10 - STEP 1 MODEM DISTRIBUTION AND OVER SUBSCRIPTION	16
FIGURE 11 - 25X OVER SUBSCRIPTION CALCULATION EXAMPLE	17
FIGURE 12 - STEP 2 MODEM DISTRIBUTION AND OVER SUBSCRIPTION	18
FIGURE 13 - STEP 3 MODEM DISTRIBUTION AND OVER SUBSCRIPTION	19
FIGURE 14 - STEP 3B MODEM DISTRIBUTION AND OVER SUBSCRIPTION	20
FIGURE 15 - STEP 4 MODEM DISTRIBUTION AND OVER SUBSCRIPTION	21
FIGURE 16 - STEP 5 MODEM DISTRIBUTION AND OVER SUBSCRIPTION	22
FIGURE 17 - DS BANDWIDTH GRAPH	23

Introduction

There are 7 layers to the OSI model. Layer 1 (physical layer), is a good place to start when talking about capacity planning. The basics of DOCSIS must be understood to plan for upstream (US) and downstream (DS) performance and capacity for any high speed data (HSD) network. Moving past layer 1, the transport layer of the OSI model (layer 4) will be reviewed to understand TCP vs UDP and its effect on capacity planning.

Capacity improvements, configurations, and best practices will be explained such as: DOCSIS 3.0 (D3.0) US and DS bonding, resiliency, "Powerboost™", ack suppression, load balancing (LB) and CM steering. The emphasis of this session will be on D3.0 operation with bonded channels and the role that D3.1 will play in future capacity allocation, offerings and delivery. Capacity exercises will be demonstrated utilizing D3.0 and D3.1 scenarios with existing cable plant deployments along with options for plant expansions and/or new architectures such as Remote-PHY.

Issues to be addressed will be: cable attenuation, diplex filter options, existing customer premise equipment (CPE) issues such as adjacent device interference (ADI), temperature variations, level fluctuations, etc. The DS spectrum to be accessed will be 750 MHz, 862 MHz, 1.002 GHz, and 1.218 GHz. The US spectrum to be accessed will be 42 MHz, 65 MHz, 85 MHz, and 204 MHz.

Historically, cable systems would monitor utilization over time to decide when a node-split was necessary or increased LB was needed. If spectrum was scarce, then no more DOCSIS channels could be added, so no LB. Now with the advent of Powerboost, different applications, and changing traffic patterns; it's difficult, if not impossible, to rely on the old monitoring techniques to prove congestion is really happening. On top of this, what exactly is "utilization"? How do we calculate a percentage accurately and over what rolling, average, time window? Some municipalities are requiring "proof" of advertised speed, so this must be a requirement of monitoring.

Setting expectations is the first step in understanding capacity constraints. Understanding what causes those constraints and ways to work-around them are the next steps. Creating a perception of higher throughput and capacity is one tool to pacify bandwidth (BW) intensive subscribers, but providing higher capacity and controlling fair use is also critical. Most calculations are based on historic data and estimates of traffic patterns and human nature. Over subscription (OS) can be a moving target and allowing a subset of users to abuse their fair-use affects everyone.

Number One Capacity Question

The number one capacity question to ask is, "Which type of over subscription should a cable company or MSO consider?" Concurrency is a measure of simultaneous usage in a pool of subscribers. It can be expressed as a percentage of users simultaneously using the service or it can be expressed as an average throughput per subscriber by monitoring actual traffic to determine average bits per second (bps). Determining concurrency for capacity planning is an on-going process.

Another option used is to track actual utilization and decide on a "watermark" that makes sense. Something like; if utilization reports > 75% for > 40% of the day, then it's time for a node split or add more capacity. A third, simpler, option is to come up with a number of devices per service group (SG) or per US channel. The running joke for this question, "What is the recommended number of devices per US channel" is: "anywhere between 1 and 1000". There are way too many variables to nail this down.

What about future requirements? More IP video streaming with Hulu, NetFlix, (a result of cord-cutters), 4K TV, DVR of multiple linear channels, live channels broadcast over the Internet, Slingbox, Periscope, home security, nanny cams, etc. The last few examples could drive much more US traffic.

One major thing to consider would be what devices must be accommodated? Assuming we need to support DOCSIS 2.0, 3.0 and eventually D3.1; what about D1.x and DOCSIS Settop Gateway (DSG)?

Over subscription factors of 25X to as high as 100X have been used historically by operators for capacity planning and continue to be used today.

The next major question to ask is; do I do node splits or allocate more spectrum? Competitive pressure for peak speeds is driving more spectrum allocation which could negate the need for node splits. More node splits equals more optics. More optics in the headend (HE) equals more US ports, which may require more US RF combining. Making nodes smaller to turn-around and combine RF again is not something we should promote. A symmetrical design of 1:1 (DS to US) is desirable, but 1:2 is ok as well. Architectures of 1:4 can cause some concerns with US port connectivity on the CMTS and US channel allocation per MAC domain.

Capacity is typically being added these days for existing subscribers' higher rates, not usually to add significantly more customers or make service groups bigger. This is the basis for our examples.

Downstream Capacity Overhead

Each US in the MAC domain can decrease DS "usable" rate by .25 Mbps. This is estimated from 500 MAPs/s * 64 B * 8 b/B = .256 Mbps, A typical MAC domain contains 4 USs, but could be as high as 8 USs, leading to potentially 2 Mbps of overhead to support MAP messages. This is replicated on every "primary" DS in that MAC domain.

"Usable" rate refers to the OSI layer 2 speed reporting by the CMTS and CM, which means 18 bytes (B) of Ethernet overhead are included in the calculations. OSI layer 3 reporting by a PC could be 5 - 10% less if average Ethernet frames are less than 1518 bytes. For example, 18/64 = 28%! 18/1518 = 1.2%. 18/229 = 7.8%. The last example is the typical size of a DS VoIP packet of 229 bytes.

There is even more overhead that is typically forgotten or overlooked such as: DOCSIS (6 B), BPI+ (5 B), Bonding (6 B), and possibly VLAN (4 B). This overhead will affect usable rate as well. IPv6 also adds 20 more bytes vs IPv4, but that is at layer 3.

Note: If using DSG boxes, multicast could create another 1 Mbps sent down each "primary" DS. This can be removed with; `(config)# cable downstream dsg disable` on the integrated & modular interfaces where not deemed necessary or needed.

"Powerboost™", a name trademarked by Comcast, is a simple manipulation of when to rate-limit a CM to allow a burst of speed until a set amount of "DS max burst" in bytes is exceeded. Other cable companies may call it by a different name. This feature can affect the "perceived" speed seen by customers.

Powerboost™

To enable this feature we need a very large DS max burst setting in the 1.1 CM file. For example; 5 MB vs the default of 3044 bytes. Depending on the CMTS vendor or device, a global command may need to be invoked, such as: `(config)# cable ds-max-burst burst-threshold {threshold in KB} peak-rate {rate in kbps}`. The burst-threshold for this command defaults to 1000, which is 1 MB.

Note: The CM needs to reboot for these settings to take affect and the throughput trace could appear oscillating slightly up and down because of the nature of rate limiting and queue build-up.

D3.0 CMs support a type length value (TLV) for per-CM peak rates, which adds more flexibility and control for Powerboost. A command exists for Cisco CMTSs to withhold those TLVs if problems persist for older CMs;

```
(config)#cable service attribute withhold-tlvs ?
      peak-rate  Peak Traffic Rate TLV 24/25.27
```

An US Powerboost can also be implemented by setting a very large US traffic burst (i.e. 2 MB) & 8 kB max concat burst in the 1.1 CM file. The max concat burst setting has no effect on D3.0 US bonded CMs, but if that CM registers in D2.0 US mode, it would still be able to get > 5 Mbps US speed.

Baseline Capacity

Downstream

Annex B DOCSIS using 256-QAM has a “raw” rate of 42.88 Mbps. “Usable” rate is normally listed within a CMTS as 37.5 Mbps, but the actual usable rate is closer to ~36 Mbps. This assumes the DS is primary-capable. If allocating bandwidth/utilization for minimum guaranteed services or Committed Information Rate (CIR), then the admission control for this could be .8 of “usable”. When the DS is “secondary-only”, the usable rate is ~37.5 Mbps because there is no signaling or MAP overhead.

Annex A DOCSIS uses 8 MHz wide channels with a 15% filter alpha regardless if 64 or 256-QAM. This leads to 6.952 Msym/s and 55.6 Mbps raw rate when using 256-QAM. The usable rate is closer to ~ 47 Mbps and ~ 50 Mbps for secondary-only.

DOCSIS 3.1 using a 192 MHz OFDM block at 1024-QAM equals ~ 1.5 Gbps (~ 47 Mbps / 6 MHz equivalent). When using 4096-QAM, this usable rate is ~1.8 Gbps (~ 56.25 Mbps / 6 MHz equivalent). These modulation schemes may seem far-fetched, but with Remote-PHY architectures and limited amplifier cascades, this may be possible. A DS MER of 50 dB and US MER of 40 dB at the node are within reason to support 4096-QAM and 1024-QAM on the DS and US, respectively.

Upstream

DOCSIS 2.0 using 64-QAM at 6.4 MHz ch width has a “raw” rate of 30.72 Mbps. Actual usable rate is closer to ~27 Mbps assuming 1518 B frames and ~7% overhead for forward error correction (FEC).

Note: IPv6 may necessitate modulation profile changes, i.e. 20 more bytes could be added to each US VoIP packet leading to inefficient minislot utilization. It may be wise to change the a-ugs burst from this:
cab modu 226 atdma a-ugs 9 232 0 40 64qam scram 152 no-diff 64 short qpsk1 1 2048. To:
cab modu 226 atdma a-ugs 7 128 0 40 64qam scram 152 no-diff 64 short qpsk1 1 2048. A typical VoIP packet is 232 bytes, but could be 256 when using IPv6 and accounting for a potential 4 byte VLAN tag. Allocating 2 FEC codewords (CW) of 128 is more efficient than one CW at 232 and a second CW that is shortened, but with more FEC overhead.

D3.0 allows for 8-ch US bonding and an 85 MHz upper bandedge, but limited spectrum in North America (NA) and lack of CMs could limit US bonding to 4. NA systems are specified for 5-42 MHz, but most stop at 40 MHz. To achieve 8-ch US bonding, we would need $8 * 6.4 = 51.2$ MHz of spectrum. D3.0 CMs are specified out to 85 MHz, but many deployed D3.0 CMs have 42/54 MHz diplex filters along with most cable plants. Euro-DOCSIS US uses 5-65 MHz and will have more spectrum to work with.

Note: 8-ch US bonding could be limited to 48 dBmV max output/ch unless the Cablelabs ECN for extended US power in the CM is implemented!

Note: D3.1 US is thought of in blocks of 96 MHz, but the industry will probably settle on licensing in blocks of 6 MHz for better granularity and a closer match to D2.0.

When using D3.1 US at 256-QAM and a 6 MHz ch width, a usable rate of ~ 36 Mbps can be achieved. Remote-PHY may allow higher D3.1 modulation schemes of 1024-QAM to potentially achieve ~ 45 Mbps with that same ch width. A case could be made for an US of 5 - 85 MHz leading to 78 MHz ($6 * 13$) of D3.1 at 256-QAM to achieve ~ 470 Mbps.

Speed Variables

Downstream

There are many factors that affect “real” data throughput. These range from: MAP & DOCSIS overhead, Ethernet frame size or MTU, VPN and/or wireless encapsulation, etc. The CM itself could be an issue with its config file, CPU which translates to packets per second (PPS), ack suppression, and Ethernet connectivity.

Layer 4 of the OSI model, transport layer, could be TCP or UDP. US speeds can affect DS TCP flows & TCP windowing affects the amount of DS frames per US ack.

The max DS burst in the CM file can give the perception of faster speed and even the computer operating systems and Windows® stack have an effect.

Upstream

There is a plethora of limiting factors that affect “real” US data throughput. These range from the “request/grant” cycle to aggregate “pipe” size. Understanding the limitations will aid in expectations and optimization. For detailed information about DOCSIS throughput factors, refer to “Understanding Data Throughput in a DOCSIS World” also listed in the reference section at the end of this paper.

Modem rate limiting and/or congestion along with physical layer issues like, dropped frames and/or older CMs, are a good starting point. These lead to a higher bit error ratio (BER). DOCSIS protocol for US traffic utilizing the request/grant cycle is one of the leading misunderstood per-CM US speed variables. Many factors affect this such as: concatenation, fragmentation, and DS MAP advance. Map advance can be affected by DS modulation & interleaving, MAP advance safety, CM time offsets (distance), US bonding, SCDMA, M-CMTS (CIN), etc.

US concatenation and CM file settings of max concat & traffic burst along with the amount of fragmentation contribute to per-CM US speed. While higher modulation schemes and larger ch widths create a bigger “pipe”, for more aggregate speed, per-CM speed could benefit as well. Because of shorter packet serialization time and extracting piggybacked requests before processing the entire segment, a quicker request/grant cycle can occur for faster per-CM US speed. D3.0 US bonding using continuous concatenation and fragmentation (CCF) avoids this inherent limitation found in D2.0 and lower CMs.

DOCSIS 3.1 Capacity Planning Concerns

How will D3.1 DS capacity planning work if some subcarriers are running different modulation and/or some subcarriers muted? Also, what is the actual block size? CMs in the same SG may have different profiles. One potential answer to this is to use the highest profile calculation, which could be misleading.

How will D3.1 US capacity planning work when sharing time with D2.0 & D3.0 CMs? There could also be different modulation schemes on many subcarriers with a varying block size. CMs in the same SG could have different profiles as well. US utilization is easier to assess by relying on minislot allocation also known as ch utilization vs BW utilization, which is bps divided by some estimate of usable rate.

Spectrum Planning

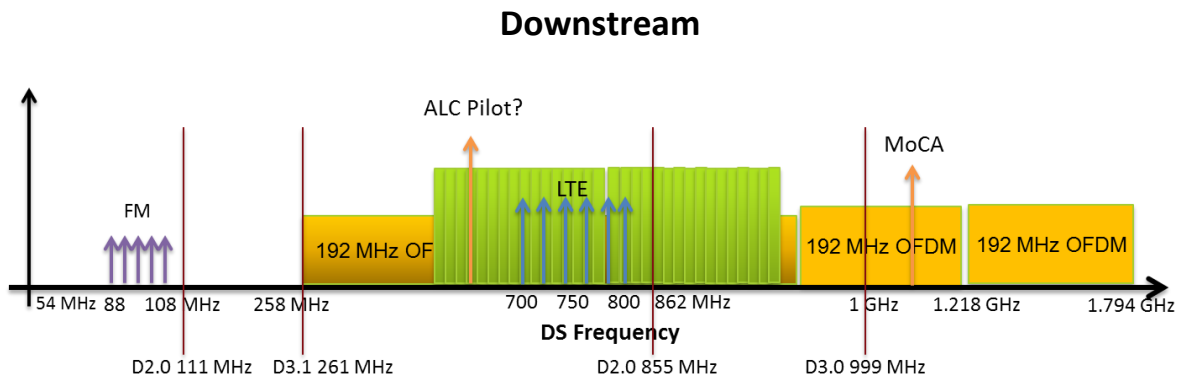


Figure 1 - Downstream Spectrum Planning

Many North American cable systems utilize DS spectrum to 750 MHz, 862 MHz and even some out to 1 GHz as shown above in Figure 1. The move to 1.218 GHz or 1.794 GHz is a very big, costly decision.

D2.0 allows a DS range of 111 MHz up to 855 MHz center frequency. D3.0 allows up to 999 MHz center frequency for the last ch. Currently, D3.1 specifies 1.218 GHz as the upper band edge with an option for 1.794 GHz. The lowest frequency of 111 MHz is optional for D3.1, while 261 MHz is mandatory.

Since D2.0 CMs stop scanning after 855 MHz, spectrum above could be utilized by D3.0 and D3.1 CMs. A decision could be made to place D3.0 from 861 to 999 and D3.1, 192 MHz OFDM block(s) above 860 MHz or even above 1 GHz if the plant allowed.

Note: There is no DS spectrum sharing between D3.0 and D3.1 CMs. Spectrum must be allocated for each individually. Even though the specification lists OFDM blocks in 192 MHz, the industry will more than likely settle on 24 MHz for licensing granularity. This is also a nice fit for Annex A (8 MHz) and Annex B (6 MHz).

Long term evolution (LTE), also known as 4th generation cellular, is located between ~700 to 800 MHz and may create potential issues, not just from the cell tower but also from the phones themselves. We could avoid this spectrum, but we would lose capacity. Another option would be to not make DSs "primary" in this region of spectrum and use resilient bonding groups (RBGs) in case of interference. If we make DSs in this spectrum secondary-only, we lose D2.0 capacity. On the plus side, we gain more speed since "secondary-only" chs have no DOCSIS overhead. By upgrading from D2.0, operators may utilize the CM full bandwidth capture (FBC) feature in newer CM/gateways to "see" ingress remotely.

Upstream

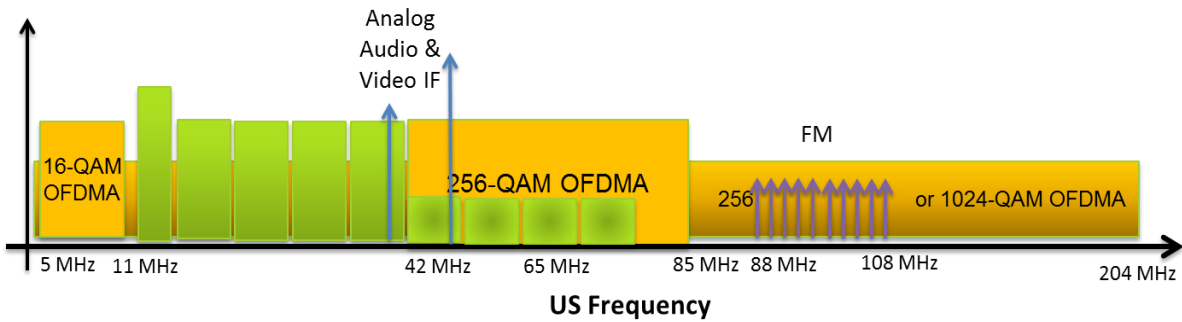


Figure 2 - Upstream Spectrum Planning

Because US traffic in a DOCSIS network is “bursty” and time-shared, we can share spectrum between D3.0 and D3.1 CMs. The CMTS will be responsible for synchronizing D3.0 and D3.1 traffic to avoid scheduling on top of each other. One big question will be; “How to allocate licensing if time-sharing the spectrum?” Why pay for D3.0 and D3.1 licensing if the CMs are using the same spectrum and is there a big enough difference in speed to justify it?

Note: SCDMA was dropped from the D3.1 specification and also made a “may” instead of a “must” for D3.0. Also, the lowest channel width is now 1.6 MHz for D3.0.

US spectrum allocation for D3.0 could go to 85 MHz while the D3.1 specification lists 204 MHz or potentially higher as shown in Figure 2 above. Most cable systems utilize D2.0 & D3.0 up to 42 or 65 MHz with a few systems utilizing D3.0, 8-ch US bonding to 85 MHz.

As mentioned earlier, D3.1 OFDMA block size is listed as 96 MHz, but will more than likely have a granularity of 6 MHz. Because of the robustness of D3.1 OFDMA, a case could be made for D3.1 below 18 MHz & above 42 MHz.

Caution: Don’t forget about adjacent device interference (ADI). Many of the CPE in North America like STBs, VCRs, and older analog TVs could have issues when a CM potentially transmits at 50 dBmV above 42 MHz and “bleeds” over into the device. Some filters may be justified in the house, but tap isolation and attenuation should be adequate for house-to-house issues.

DOCSIS 3.1 Spectrum Thoughts

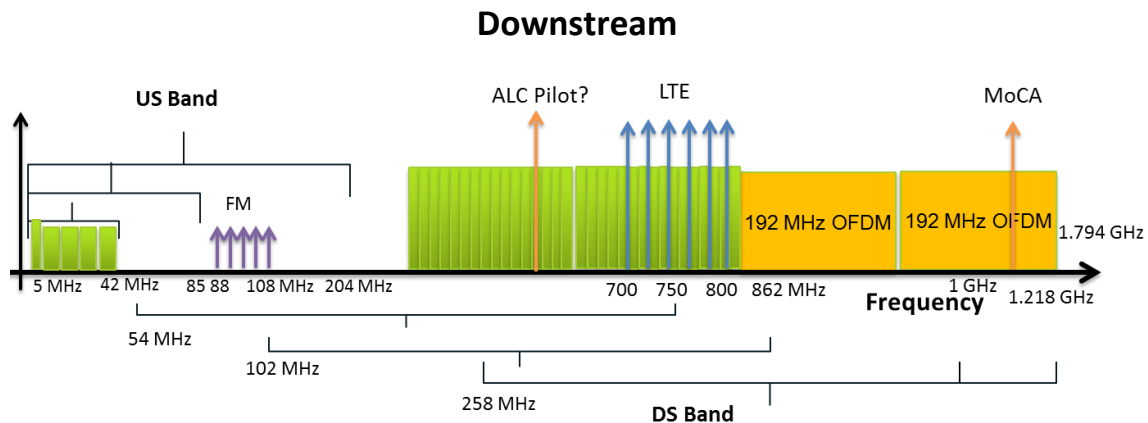


Figure 3 – D3.1 Downstream Spectrum Thoughts

DOCSIS 3.1 pushes the DS from 1 GHz in D3.0 to 1.218 GHz with a higher option to 1.794 GHz as shown in Figure 3 above. The idea of upgrading the entire HFC plant for actives and passives could limit the decision to 1.218 GHz.

A 1.218 GHz upgrade is probably the least expensive upgrade because newer amplifiers use Gallium Nitride (GaN) instead of Gallium Arsenide (GaAs) allowing module upgrades without changing housings or re-spacing. Taps and other passives would probably have to be replaced, but most coaxial cable supports well past 1 GHz. MoCA around 1.150 GHz and higher may still be an issue to deal with.

Higher frequencies past 1.218 GHz could create more problems with actives, passives and especially coax attenuation. The only way to avoid the coax issue would be to limit the amount of coax with fiber-deep designs such as: node + 0 amplifiers, RFoG, or Remote-PHY.

Keep in mind that US spectrum decisions also affect the low-end of the DS spectrum. This means lower DS analog chs must be moved and the amount of DS spectrum could be affected. Also, what about DS automatic gain control (AGC) pilots? Will the amplifiers support a QAM pilot or even a D3.1 OFDM signal pilot or will they require a dedicated continuous wave (CW) carrier for a pilot?

Upstream

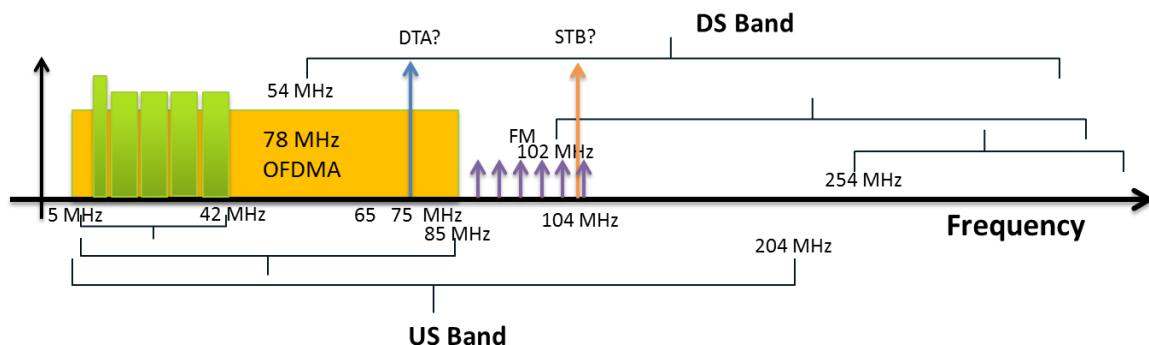


Figure 4 – D3.1 Upstream Spectrum Thoughts

D3.1 allows US spectrum out to 204 MHz as shown in Figure 4 above with higher options, but then it “eats” into the DS spectrum. Upgrading the diplex filter split is a big decision and cannot be taken lightly. Having the option of pluggable diplex filters for changes later may be advantageous.

Some systems may not upgrade at all because of millions of digital terminal adapters (DTAs) hard-set at 75 MHz for DS addressability.

Many cable systems may settle on 85 MHz for the following reasons:

1. FM-band potential interference.
2. Some legacy STBs have a DS addressability ch at 104 MHz, so an 85/102 split makes sense.
 - a. A 204/254 split forces certain STB upgrades.
3. Increasing US spectrum creates much more coax attenuation between the CM and CMTS.
 - a. This could force US AGC and what about pilots?
 - b. A novel idea of using the DS AGC circuitry to drive an US AGC circuitry is being discussed.
4. US Tx level issues;
 - a. D2.0 = 54 dBmV; D3.0, 4-ch = 51 dBmV, 8-ch 48 dBmV.
 - b. Typically a 6 dB swing annually at 42 MHz due to thermal fluctuations.
 - c. No US ALC/AGC, the operator must rely on the CM/CMTS long-loop-level control.
 - d. Potential fixes
 - i. US thermal EQs to help stabilize negative fluctuations.
 - ii. D3.0 with Extended Power ECN provides increased US transmit by 3 - 6 dB.
 - iii. D3.1 is ~ 2 dB more power per equivalent D3.0.

Note: For Euro-DOCSIS systems, an upgrade from 65 to 85 MHz may not make sense and the STB issue may not be an issue at all leading to a 204 MHz deployment/upgrade. FM carriage must go away though.

Reasons to go to 204 MHz:

1. Some STBs are frequency agile and not relegated to 104 MHz.
2. Upgrading to newer IP STBs could generate more features and services.
3. 5 – 204 MHz allows enough space to do two complete 96 MHz OFDMA blocks.
4. If Remote-PHY and Node + 0, then higher modulation schemes and a higher split are easier to achieve.
 - a. If 1024-QAM modulation is feasible, then ~ 1.44 Gbps aggregate is possible.

US Level Swing over Temperature

In regards to cable attenuation and temperature swings, Figure 5 below depicts a potential +5 to -5 dB swing at 85 MHz. Although this is slightly exaggerated because of -40 C to +60 C and it was with a 5 amplifier cascade with ~ 7000 total feet of .5" coax, it shows the extreme swing in levels. Going to 204 MHz would be even worse!




Figure 5 - US Level Swing over Temperature

The CM US transmit reserve (hopefully 3 dB) and CMTS range configuration (maybe 6 dB) will help on the positive side, but the minus side creates a higher noise floor. When it is cold out, there is less coax attenuation so the CMs lower their transmission level. This is not a problem, but the noise from lower value taps sees less attenuation so the noise floor rises for all subscribers. This could necessitate US thermal equalizers. The 204/254 MHz split may necessitate actual US AGC or much less coax!

Five Potential Migration Steps and Ideas

There are five potential scenarios and ideas we will cover for US and DS expansion and the type of services that could be offered. Migration step 3 includes a second option that will be reviewed as well. We will start out with some assumptions for the SG such as: 600 total HSD subs and 1000 to 1500 DSG STBs. The five migration steps are:

1. 12 DSs: all primary-capable
 - a. Accommodate DSG, D1.x, D2.0, & D3.0 CMs
2. 24 DSs: 12 primary-capable & all D1.x removed
3. 32 DSs: 16 primary-capable & D3.1 OFDM block
 - a. Step 3b: 32 + 32 DSs
4. 85 MHz US x 1.2 GHz DS with more D3.1
5. Remote-PHY: 204 MHz US x 1.2 GHz DS

DOCSIS 3.0/3.1 DS Channel Allocation

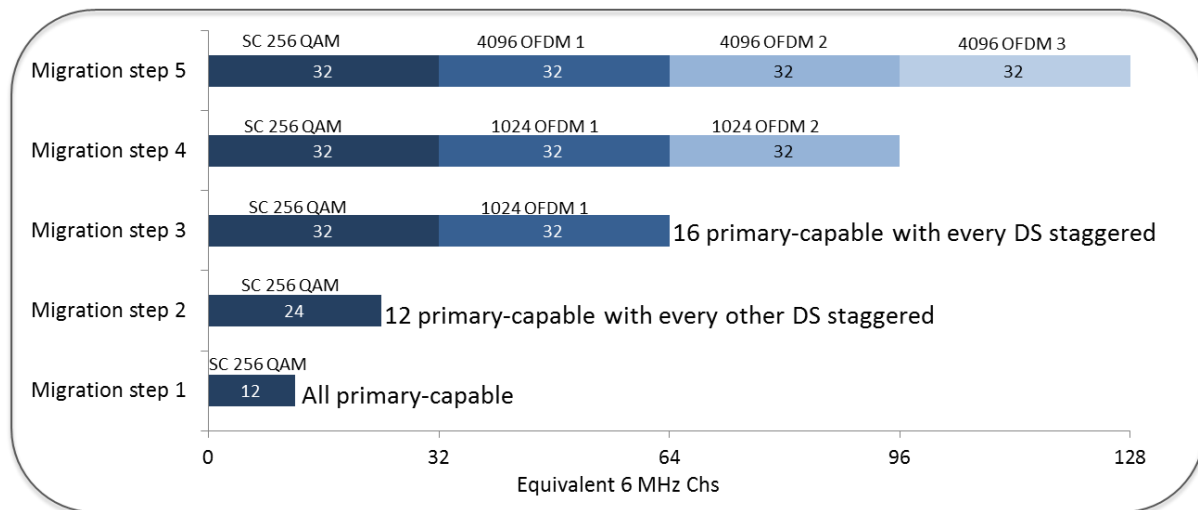


Figure 6 – DS Channel Allocation

Figure 6 above lists DS channel allocation for all 5 migration steps and examples. D3.1 is listed in 6 MHz ch equivalent. The first three migration steps can work in 750 MHz, 860 MHz, or 1 GHz plant. A 1.2 GHz upgrade can be done when desired, but is shown in steps 4 and 5. Remote-PHY with higher modulation is displayed in step 5.

Note: The primary DSs were staggered for better distribution of 4-ch BG and 8-ch BGs and also for RBG selection. Not all SC-QAMs are made primary since D2.0 CMs will be going through attrition and secondary-only DSs provide more throughput.

DOCSIS 3.0/3.1 DS Spectrum Allocation




Figure 7 – DS Spectrum Allocation

Figure 7 above lists US and DS spectrum allocation for all 5 migrations steps. Keep in mind that different modulation can be utilized in some spectrum, in different D3.1 sub-carriers, and even for specific D3.1 CMs.

DOCSIS 3.0/3.1 US Channel Allocation

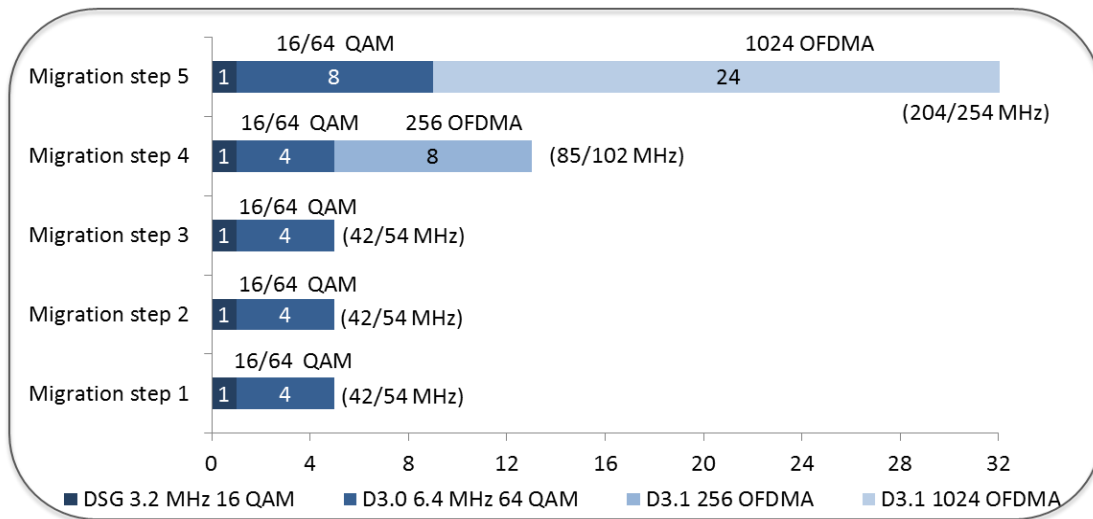


Figure 8 – US Channel Allocation

Figure 8 above lists US channel allocation for all 5 migration steps and examples. One quarantine channel is allocated for DSG boxes and four to eight ATDMA chs for D2.0 and D3.0 CMs. D3.1 is listed in 6 MHz equivalent chs. A diplex filter upgrade was done at step 4 to allow more US speed. Remote-PHY with a 204 MHz US was used for step 5 and allowed higher modulation schemes to be used.

DOCSIS 3.0/3.1 US Spectrum Allocation

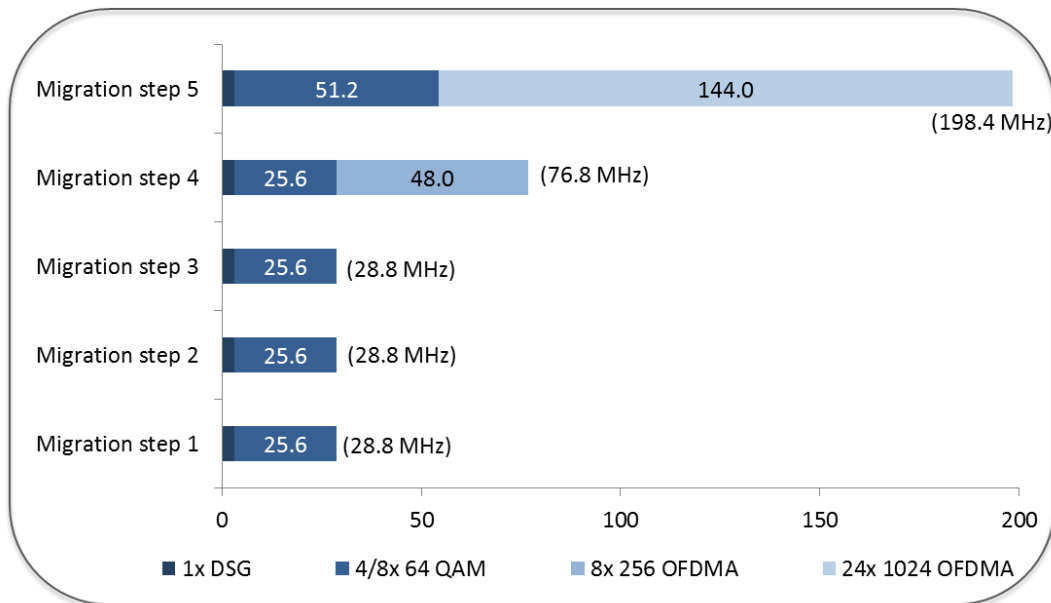


Figure 9 – US Spectrum Allocation

Figure 9 above lists US spectrum allocation for all 5 migrations steps. All migrations steps start DSG boxes at 11 MHz band edge. D3.1 can utilize all US spectrum, but will be relegated to unused spectrum. Keep in mind that different modulation can be utilized in some spectrum, in different D3.1 sub-carriers, and even for specific D3.1 CMs.

Step #1: 12 DSs 5-42 MHz US & 750 MHz DS

Trying to accommodate legacy CMs along with DSG STBs affects LB decisions, spectrum allocation, and overall capacity planning. This may require an operator to perform analog reclamation. Starting with 12 single channel QAMs (SC-QAMs) and assuming Annex B, 256-QAM, we could place all 12 from 681 to 747 MHz.

US would assign one quarantine ch at 1.6 or 3.2 MHz wide using 16-QAM or QPSK and three to four ATDMA chs at 6.4 MHz using 64-QAM. The quarantine ch would be more robust and used for D1.x CMs and DSG to allow higher transmit levels, which may be required if STBs are located deep in the house.

For this offering we start with 1000 STBs and could offer 5 x 1 Mbps to 102 D1.x CMs. 246, D2.0 CMs could get 15 x 5 Mbps while D3.0 CMs would be split based on their physical capabilities. We start with 4 and 8-ch capable bonding and could offer 50 x 5 and 150 x 15 to ~ 222 and 30 CMs, respectively.

Note: The highest offering is half of CM capability; i.e. highest offering for D2.0 CM is 15 x 5 Mbps.

As expected and shown in Figure 10 below, CM distribution forms a bell curve and the middle tiers make up the bulk of the subscribers. DS is the focus here since US is not a limiting factor at this point. Over subscription (OS) hovers near the 25X factor for the median tier. D1.x OS could easily be greater, but the extreme low and high ends of the bell curve are typical of subscriber take-rates.




Figure 10 – Step 1 Modem Distribution and Over Subscription

25X Over Subscription Estimate

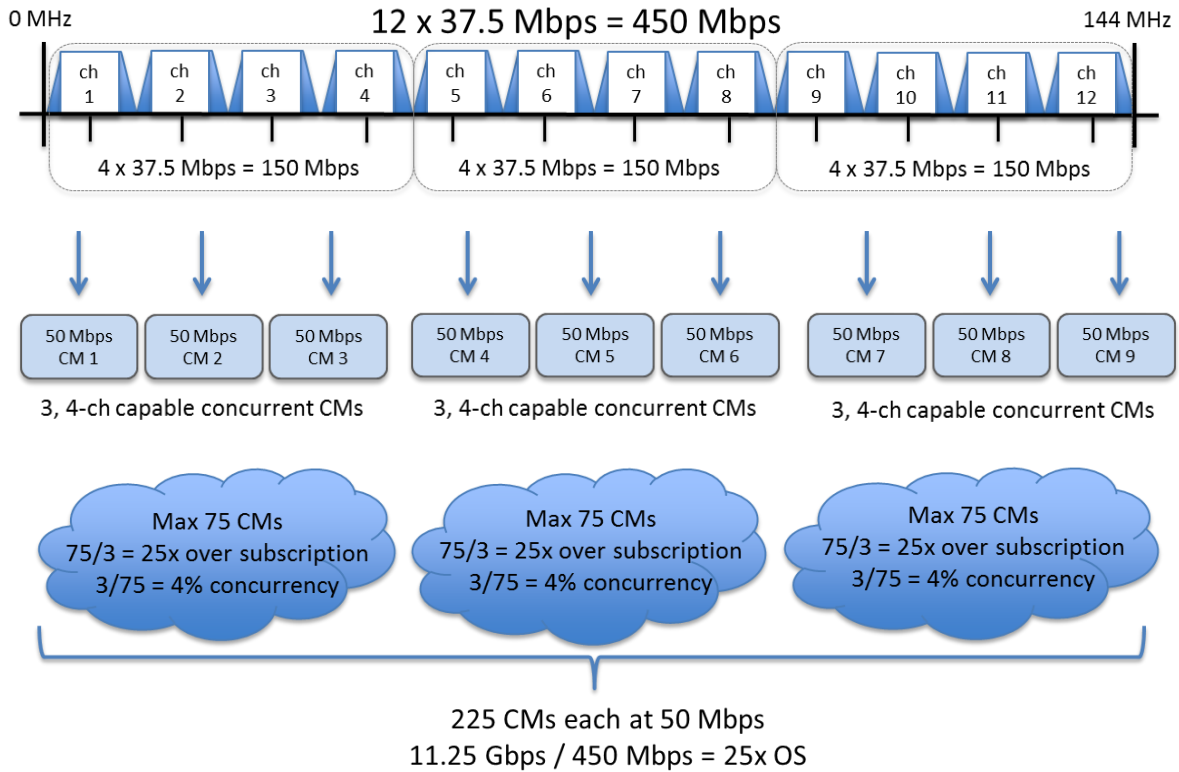


Figure 11 – 25X Over Subscription Calculation Example

Figure 11 above attempts to explain how this OS is calculated. In the above example, a 50 Mbps DS service is offered using a 4-ch capable D3.0 CM. The usable rate is based on 2 primary DSs at 36 Mbps each and 2 secondary at 37.5 Mbps each, giving a total of 147 Mbps. For simplicity, the example assumes 150 Mbps as usable. $150/50 \text{ Mbps} = 3$ CMs could be supported simultaneously. There are 225 total 4x4 D3.0 CMs distributed across 3, 4-ch BGs leading to: $225/3 = 75$ CMs per BG. $75/3 = 25X$ OS factor.

Step #2: 24 DSs Extended DS to 862 MHz

We remove all D1.x CMs and subscribers upgrade to higher speeds while we migrate to 24 SC-QAMs from 717 to 855 MHz. This may require more spectrum allocation and/or analog reclamation.

The US has not changed and same as previous scenario, but more DSG STBs have been deployed (1250) and now at a lower speed of 1 x 1 Mbps.

Note: D2.0 CMs are limited to 15 x 5, D3.0, 4x4 = 50 x 5, D3.0, 8x4 = 150 x 15, D3.0, 16x4 = 300 x 30, and D3.0, 24x4 = 500 x 40 Mbps.

Note: Although 32x8 D3.0 CMs are available, the aggregate of 1200 x 100 Mbps could limit a realistic offering to 600 x 60 Mbps. The US is limited to four chs because of spectrum. Also keep in mind that gigabit Ethernet (GigE) ports on the CM or PC are limited to ~ 960 Mbps.

There are now 132, D2.0 CMs at 15 x 5 Mbps and D3.0 CMs look like: 210 CMs at 50 x 5, 216 CMs at 150 x 15, 30 CMs at 300 x 30, and 12 CMs at 500 x 40 Mbps.




Figure 12 – Step 2 Modem Distribution and Over Subscription

Step #3: 32 DSs & D3.1 OFDM Block Advanced DS to 1 GHz

This scenario adds D3.1 DS OFDM to offer a 1 Gbps tier. The assumption in this step will be that D3.1 CMs will bond D3.0 and D3.1 chs together. Expanding the plant to 1 GHz or 1.2 GHz also makes it easier to “squeeze” in more DOCSIS without analog reclamation or even digital reclamation. Digital reclamation would entail converting MPEG-2 to MPEG-4 or even higher orders of compression with H.265 using IP STBs. We migrate to 32 SC-QAMs from 621 to 807 MHz and can now easily allocate 1, 192 MHz OFDM block from 810 to 1002 MHz.

Note: We are confident that where we run 256-QAM in the DS today, D3.1 with its more robust features can run 1024-QAM reliably.

Note: D3.1 US OFDMA is planned for much later than D3.1 DS and assumed to be less of a requirement than DS in the near future. With this said, we as an industry are placing ourselves in a “pigeon hole” by sticking with the “rule-of-thumb” of 10 x 1 DS to US ratio offerings. Before we get too far, we should start “reeling” this back slightly. Maybe 20:1 or even 30:1 for really high DS offerings would be wise.

The US has not changed and same as previous scenarios. The same amount of DSG boxes have been deployed at 1250 STBs using 1 x 1 Mbps service. The D2.0 CMs at 15 x 5 Mbps service have gone through some attrition and dropped to 102 total CMs. D3.0 CMs look like: 192 at 50 x 5, 240 CMs at 150 x 15, 36 CMs at 300 x 30, 18 CMs at 500 x 40 Mbps and 12 D3.1 CMs at 1000 x 50 Mbps.




Figure 13 – Step 3 Modem Distribution and Over Subscription

The D3.1 OS is based on an aggregate “pipe” of 1.5 Gbps for the OFDM block plus 1.2 Gbps for the 32 SC-QAMs. $(1.5 + 1.2) / 1 = 2.7$. $12 \text{ D3.1 CMs} / 2.7 = 4.4\text{X OS}$.

Note: All these speeds could be offered in previous steps, but more spectrum required may equate to more analog reclamation or a pure-digital DS design.

Step #3b: 32 DSs & 32 DSs Stick with D3.0?

The big question in the industry lately is whether to wait for D3.1 or just deploy more D3.0 DS chs and 32x8 capable CMs. Although we suggest an aggregate “pipe” ~ 2 times the amount of the offering, 32 Annex B DS chs can provide ~ 1.2 Gbps. Trying to offer 1 Gbps from this aggregate pipe is not recommended and severely hampers the statistic multiplexing calculations. But, it could be done if it’s mainly for a peak-speed offering and possibly marketing purposes.

Assuming we start with 32 DS chs bonded, we could offer D2.0 and D3.0 CMs with 300 Mbps and lower tiers. An extra 32 DS chs narrowcast or overlaid for .5 to 1 Gbps customers-only could be designed and it would be the same spectrum allocation as D3.1.




Figure 14 – Step 3b Modem Distribution and Over Subscription

Some Cons to this sub-step idea may be:

1. Will it be a software-based change or physical layer changes (splitting/combining) are required?
2. Does the amount of SGs per CMTS shrink? This may require more hardware.

Some Pros to this idea:

1. No D3.1 CPE to wait on or pay for. US silicon for the D3.1 CPE could be awhile, although it may be acceptable to deploy D3.1 DS-only CMs in the beginning since very low numbers.
2. ~ 1.2 Gbps pipe for Annex B vs 1.5 Gbps for D3.1 is not a huge difference. Annex A could offer ~ 1.6 Gbps but more spectrum would be needed.

Step #4: 32 DSs & D3.1 OFDM Blocks Goal of 85 MHz US & 1.2 GHz DS

The ultimate goal is higher speeds and more US allocation. We start with the same 32 SC-QAMs, but now from 453 to 639 MHz and 3, 192 MHz OFDM blocks from 642 to 1218 MHz. More OFDM blocks affords us the capability to not have to bond between D3.0 and D3.1 chs alleviating congestion on the existing D2.0 and D3.0 CMs.




Figure 15 – Step 4 Modem Distribution and Over Subscription

On the US we allocate multiple 6 MHz blocks of D3.1 OFDMA using 256-QAM and still provide one quarantine channel for DSG boxes and now four to eight ATDMA chs. The D3.1 US OFDMA can either burst for the entire spectrum or above and below the existing TDMA/ATDMA allocated spectrum. For this step, we did not have to share time on the US spectrum between D3.1 and D3.0. D3.0 total allocated spectrum was 28.8 MHz leaving enough spectrum for 8, 6 MHz OFDMA chs (~266 Mbps aggregate).

Note: Cable systems currently running 64-QAM in the US should be able to run D3.1, 256-QAM OFDMA confidently.

DSG boxes at 1 x 1 Mbps have gone up slightly to 1500 STBs. The D2.0 CMs at 15 x 5 Mbps service have gone through more attrition and dropped to 60 total CMs. D3.0 CMs look like: 150 at 50 x 5, 168 CMs at 150 x 15, 150 CMs at 300 x 30, 36 CMs at 500 x 40 Mbps, 24 D3.1 CMs at 1000 x 50 Mbps and 12 D3.1 CMs at 2000 x 100 Mbps. Three OFDM blocks equals ~ 4.5 Gbps aggregate to be able to offer these types of speeds and quantities and keep the OS around 5X.

Step #5: 32 DSs & Multiple D3.1 Blocks Remote-PHY

The long term solution for broadband providers is getting fiber and/or the CMTS functionality closer to the subscribers. RF over Glass (RfOG) may have its place and GPON/EPON sounds good, till you run all the numbers for cost. Placing the CMTS physical layer functionality in the node itself has some great advantages. This is known as Remote-PHY. One of those advantages is longer distance between the CMTS and CM. Some preliminary testing has shown > 2000 km! Another advantage of this architecture is potential MERs of close to 40 dB for the US and 50 dB for the DS. This will lead to realizing the higher modulation schemes provided in D3.1.




Figure 16 – Step 5 Modem Distribution and Over Subscription

We start with the same 32 SC-QAMs from 453 to 639 MHz for all D2.0 and 3.0 CMs. We can now provide 3, 192 MHz OFDM blocks from 642 to 1218 MHz using 4096-QAM. On the US we allocate multiple 6 MHz blocks of D3.1 OFDMA using 1024-QAM now and still provide one quarantine channel for DSG and four to eight ATDMA channels.

DSG boxes at 1 x 1 Mbps have stayed at 1500 STBs. The D2.0 CMs at 15 x 5 Mbps service have gone through more attrition and dropped to 48 total CMs. D3.0 CMs look like: 138 at 50 x 5, 162 CMs at 150 x 15, 132 CMs at 300 x 30, 54 CMs at 500 x 40 Mbps, 42 D3.1 CMs at 1000 x 50 Mbps and 24 D3.1 CMs at 2000 x 100 Mbps. Three OFDM blocks using 4096-QAM equals ~ 5.4 Gbps aggregate to be able to offer these types of speeds and quantities and keep the OS at a reasonable multiplication factor <10X.

Note: It would not be unreasonable to allow double the amount of subs in a SG to 1200 and just cut the middle tiers of service to .5 to .3 of the current rates. The upper and lower ends of the CM distribution bell curve have low numbers and low OS and shouldn't need to be adjusted. The middle tiers make up the "lion's share" of usage. Here's an example: DSG boxes at 1 x 1 Mbps are close to 3000 STBs. The D2.0 CMs at 15 x 5 Mbps service are set at 96 total CMs. D3.0 CMs look like: 276 at 25 x 3, 324 CMs at 50 x 5, 264 CMs at 150 x 15, 108 CMs at 300 x 30 Mbps, 84 D3.1 CMs at 1000 x 50 Mbps and 24 D3.1 CMs at 2000 x 100 Mbps. With these types of speeds and quantities, the OS for the higher tiers is still at a reasonable multiplication factor <17X.

What Does This Bandwidth Graph Represent?

Running 256-QAM with an annex B, 6 MHz ch width, will yield a maximum of ~ 36 to 37 Mbps usable rate. This is at layer 2 (Ethernet) and above. When a user reaches the maximum rate provisioned for their service level, a flat line in traffic vs. time will be seen. But notice in the traffic output shown in Figure 17 below that there are multiple flatlines seen in the traffic rate vs. time.




Figure 17 – DS Bandwidth Graph

The first flatline is the usual 100% capacity of ~ 37 Mbps. This would be closer to 47 Mbps for annex A (Euro-DOCSIS). The beginning of this graph represents a single modem and displaying what happens with Powerboost where the speed will peak up to maximum (if available) for a few seconds. The amount of Powerboost time is based on the DS Max Burst setting in the D1.1 cm file. When the max burst is done, rate limiting will kick in and the CM will be limited to the Max DS Rate configured in the cm file.

The second flatline is the max info rate (MIR) of 20 Mbps as configured for this customer. Powerboost will not happen again for this customer until it goes below the MIR and “tokens” are built back up.

The rest of the graph is used to represent the entire DS “pipe”, not just one customer. The third flatline indicates 100% capacity, but much lower than what was previously believed to be the max of 37 Mbps. This usually indicates many small frames. The 37 Mbps assumes traffic with big frames such as: 1518 bytes. Ethernet is 64-1518 B. One application that utilizes small frames would be VoIP. The DS VoIP frame is typically 229 B. Small frames like this have the same overhead per frame as the 1518 B frame leading to a higher percentage of overall overhead. The DS “pipe” can easily reach 100% capacity

around 35 Mbps under these circumstances. The fourth flatline appears to be near 27 Mbps. This could indicate the DS is only running 64-QAM instead of 256-QAM.

JD's 10 Rules for Capacity Planning

The following "loose" rules are explained in more detail in the proceeding section titled "10 Points to Ponder". These 10 rules are more like guidelines and industry best practices.

1. Aggregate speed = 2 x highest offering.
2. Configured speed = 1.1 x "Marketed" speed.
 - Layer 3 vs layer 2 reporting. Minimum 1.05 multiplication factor.
3. DS:US = 10:1. Account for US acks from DS TCP (ex: OTT video). 20:1 minimum initially. Possibly as low as 30:1.
4. Powerboost for median tiers, 5, 10, 15 Mbps. US Powerboost possible as well.
5. Powerboost for high tiers to bypass rule 2. Be sure a peak rate TLV is used.
6. Concurrency = over subscription.
 - Aggregate = 2x = 10-20:1
 - Aggregate = 4x = 25-50:1
 - Aggregate = 10x = 50-100:1
7. SG size = 1:1 or 1:2. This means 1 DS x 2 US connectors. HHP per SG depends on DSs/SG. i.e. 4 DS = 500 avg for 3 tiers, 8 DS = 1250 avg for 4 tiers, 16 DS = 1800 avg for 5 tiers.
8. Base math on 768B frame size because most traffic is 64 and 1518 B frames.
9. Monitor actual utilization and keep in mind that Powerboost can skew decisions for node splits. Quality of Experience (QoE) is required and not just speed. Maybe utilize D3.0 CM FTP feature to monitor.
10. Control "abusers" and prioritize some flows with deep packet inspection (DPI) and byte counting with STM/IPDR. Keep track of Denial of Service (DoS) from attacks and cloning. Offer higher priority for gaming, video, call signaling, CM registration, etc.

10 Points to Ponder

- Many speed test sites report at layer 3 of the OSI model, but the CMTS and CM files are at layer 2, which include 18 bytes of Ethernet headers. It's usually best to configure the CM config file for 5 to 10% higher speeds than marketed to account for these headers.
- Another potential problem with speed test sites is the lack of control over the actual frame size (64-1518 B) used for testing. Smaller frames have more overhead than larger frames; 18/64 (28%) vs 18/1518 (1.2%). The Maximum transmit unit (MTU) is further reduced by overhead from VPN encapsulation, wireless, etc.
- Small frames make DOCSIS "pipes" worth less than originally assumed. A DS of 37 Mbps may hit 100% utilization at 35 Mbps when all the frames are small as could be the case with DS VoIP of 229 B.
- DS TCP traffic requires US acknowledgments (acks). As should be stressed to all DOCSIS engineers, the US pipe can slow down DS speed tests. To make things worse, small US acks make the US "pipe" worth less. This is because each frame will have DOCSIS overhead, which is usually 11 bytes per frame. 10.24 Mbps raw speed equals approximately 9 Mbps usable, but only 7.5 Mbps when only TCP acks are transmitted!
- Powerboost™ is a trademarked name from Comcast, but used throughout the cable industry. It is a feature that manipulates how and when DOCSIS enforces rate limiting. This feature can give the perception of greater speed, but only for short amount of time. This is great for typical webpage downloads and specifically speed test sites. The traffic pattern created by this feature could cause issues when deciding to do node splits. Now the DS could appear >85% utilized, but no real signs of actual congestion. The other question is, "how to control peak rate when Powerboost™ is being used?" Having a different peak rate for different tiers of service would appear fair in the customers' eyes.
- More frames equals more packets per second processing, which means more central processing unit (CPU) usage. At some point the CPU in a cable modem could (and will) be the bottleneck. TCP is typically 2 DS frames per 1 US ack, so 3 PPS for each transaction. If the modem were doing 100 Mbps DS TCP, this could be, $100 \text{ Mbps} / (1518 * 8) = 8 \text{ kPPS}$ for DS and 4 kPPS for US acks = 12 kPPS total.
- Load balancing of single channel modems is good, but a decision has to be made when a certain speed tier pushes a customer to a D3.0 CM with bonding. Somewhere close to 50% of line rate is suggested. Having many CMs with high traffic load balancing between DS channels will cause "flip-flopping"/thrashing and instability.
- DS IP video is usually UDP-based, but over-the-top (OTT) apps like Netflix/Hulu TV are using adaptive bit rate (ABR), which is TCP-based. This will cause significantly more US traffic in the form of acks. New CMs may have ack suppression, but it may not be as efficient as assumed. The author has seen a typical US to DS TCP ratio of ~2%. With ack suppression, this can drop below 1%. Ack suppression alleviates acks on the wire and US traffic, but doesn't alleviate CM CPU load. Engineers are still testing with actual DS IP video since most will be 3-7 Mbps and ack suppression may not be that efficient anyway.

- Many “tweaks” are needed to get per-CM US speeds greater than 3 Mbps. Some tweaks require lots of concatenation, which then leads to fragmentation, and thereby stresses the CMTS scheduler. Fragmentation also adds headers, preamble & guardtime to each fragment, thereby reducing again the maximum throughput. On the other hand, D3.0 US bonding can do continuous concatenation and fragmentation (CCF) and possibly keep packets less than 2000 B and not require fragmentation, thereby leading to less overhead.
- During congestion, you will want to provide priority for VoIP signaling, maybe video acks, and CM registration. Many of the service flows for call setup and CM registration will default to a priority of 0. Under heavy US congestion, it may be difficult for these service flows to be provided with adequate time on the wire.

Checklist

- ✓ Highest offering < ½ of aggregate.
- ✓ Use Powerboost™ judiciously.
 - Use peak rate TLV.
- ✓ Utilize CMTS features for robustness & “self-healing”.
 - Load balance (2.0 & 3.0), RBGs, dynamic modulation, etc.
- ✓ Manage fair use of network (DPI, shaping, STM/IPDR?).
- ✓ Monitor actual traffic load and usage.
 - Large rolling time average can mask results.
- ✓ Monitor modem states and service flows are correct.
 - Partial mode, registered-traditional-docsis, etc.
- ✓ Use periodic speed test to prove no congestion.
 - Newer CMs have built-in FTP functionality.

Closing Points

Understand theoretical rates before testing. Math and new designs are beginning steps for traffic engineering and capacity planning. Monitoring actual traffic is next and then deciding what to do to alleviate congestion or perceived congestion is paramount. Keep in mind that bigger “pipes” are more efficient, but that doesn’t necessarily mean less latency. One may need to “steer” latency/jitter-sensitive services to smaller, subset bonding groups and determine when additional capacity is necessary and how long it could take to implement. Different service tiers on different channel-capable CMs could necessitate new load balance rules. Something like: don’t move if traffic < 20% or > 60% of aggregate. Also, don’t assume DOCSIS is your only bottleneck; it could be WAN, CPU, memory, TCP, etc.

Abbreviations

ABR	adaptive bit rate
ADI	adjacent device interference
ack	acknowledge
AGC	automatic gain control
ALC	automatic level control
ATDMA	advanced time division multiple access
B	byte
BER	bit error ratio or bit error rate
bps	bits per second
BPI+	baseline privacy interface +
BW	bandwidth
CCF	continuous concatenation and fragmentation
Ch	channel
CIN	converged interconnect network
CIR	committed information rate
CM	cable modem
CMTS	cable modem termination system
CNR	carrier-to-noise ratio
CPE	customer premise equipment
CPU	central processing unit
CW	continuous wave or codeword
dB	decibel
dBmV	decibel referenced to a millivolt
DOCSIS	Data Over Cable Service Interface Specification
DoS	denial of service
DPI	deep packet inspection
DS	downstream

DSG	DOCSIS set-top gateway
DTA	digital terminal adapter
ECN	engineering change notice
EPON	Ethernet passive optical network
EQ	equalizer
FBC	full-bandwidth capture
FEC	forward error correction
FTP	file transfer protocol
G	giga or generation
GaN	gallium nitride
GaAs	gallium arsenide
GPON	gigabit Ethernet passive optical network
HE	headend
HFC	hybrid fiber-coax
HEVC	high efficiency video coding
HSD	high speed data
Hz	Hertz
I	integrated
IP	Internet protocol
IPDR	Internet protocol detail records
IPv4	Internet protocol version 4
IPv6	Internet protocol version 6
k	kilo
LB	load balance
LTE	long term evolution
M	mega or modular
MAC	media access control
MER	modulation error ratio or rate

MoCA	multimedia over coax alliance
MPEG	moving pictures experts group
MSO	multiple system operator
MTU	maximum transmit unit
OBI	optical beat interference
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OS	over subscription or operating system
OSI	Open Systems Interconnection
OTT	over the top
PC	personal computer
PHY	physical layer
PPS	packets per second
QAM	quadrature amplitude modulation
QPSK	quadrature phase shift keying
RBG	resilient bonding groups
RF	radio frequency
RFoG	RF over glass
SC-QAM	single channel QAM
SCDMA	synchronous code division multiple access
SCTE	Society of Cable Telecommunications Engineers
SG	service group
SNR	signal-to-noise ratio
STM	subscriber traffic management
sym/s	symbols per second
TCP	transmission control protocol
TDMA	time division multiple access
TLV	type length value

UDP	user datagram protocol
US	upstream
VLAN	virtual local area network
VoIP	voice over IP
VPN	virtual private network

White Papers, References and Resources

Understanding Data Throughput in a DOCSIS World

http://www.cisco.com/en/US/tech/tk86/tk168/technologies_tech_note09186a0080094545.shtml

Troubleshooting Slow Performance in Cable Modem Networks

http://www.cisco.com/en/US/tech/tk86/tk89/technologies_tech_note09186a00800b123c.shtml

Upstream FEC Errors and SNR as Ways to Ensure Data Quality and Throughput

http://www.cisco.com/en/US/tech/tk86/tk319/technologies_white_paper09186a0080231a71.shtml

How to Increase Return Path Availability and Throughput

http://www.cisco.com/en/US/products/hw/cable/ps2209/products_white_paper09186a00801be4ad.shtml

Understanding MAP Advance

http://www.cisco.com/en/US/tech/tk86/tk89/technologies_tech_note09186a00800b48ba.shtml

Carrier-to-Noise Ratio in Cable Networks

http://www.cisco.com/en/US/products/hw/cable/ps2209/products_white_paper0900aecd800fc94c.shtml

Determining RF or Configuration Issues on the CMTS

<http://www.cisco.com/c/en/us/support/docs/broadband-cable/radio-frequency-rf-hybrid-fiber-coaxial-hfc/7137-rf.html>

Cable Modem Provisioning Scenarios Doc

http://www.cisco.com/en/US/tech/tk86/tk89/technologies_white_paper09186a008025c169.shtml