# CAPACITY PLANNING AND DOCSIS TRAFFIC ENGINEERING

John J. Downey, Broadband Network Engineer – CCCS CCNA
Cisco Systems, USA

## ABSTRACT

As MSOs are looking to increase capacity, cable systems are now offering higher bit rate service level agreements to meet market needs and customer demand. This paper will discuss strategies for approaching, implementing and managing capacity planning and DOCSIS traffic engineering across commercial and residential cable customer bases.

This paper will outline when and how to expand capacity using singular techniques or a combination of technologies and the impact of DOCSIS 3.0 on cable's bandwidth capacity. It will demonstrate how to calculate data requirements and throughput. It will then outline the importance of ongoing monitoring actual traffic loading in determining the fair use of the network and determine when additional capacity is necessary.

The session will discuss near term approaches such as the use of multiple downstream frequencies to one service area and the use of intelligent load balancing.

## INTRODUCTION

Before attempting to measure the cable network performance, there are some limiting factors that one should take into consideration. In order to design and deploy a highly available and reliable network, an understanding of basic principles and measurement parameters of cable network performance must be established.

Historically, most High Speed Data (HSD) content was downstream (DS) to the end users and required an asymmetrical design of DS bandwidth vs upstream (US) bandwidth. This traffic, consisting primarily of web-based downloads, e-mail and chat rooms, was very bursty in nature. Downstream throughput is once again becoming the "bottleneck" with new applications like, Slingbox and Youtube. If multiple system operators (MSOs) want to begin providing higher bandwidth to users over the DS DOCSIS channel, they need to decrease the number of users sharing that channel. This can be done with fiber node splits or adding more channels into each node at the expense of frequency bandwidth and more cable modem termination systems (CMTSs).

The issue with this becomes what to do with all the upstreams that are associated with those downstreams. A 1x4 mac domain, is one DS and 4 associated US ports and may not be needed. US ports are unused or extremely under utilized. Also, cable modems (CM) that move to another DS, must move to a new US associated with that DS. This is where DOCSIS 3.0 comes to the rescue by breaking the 1xN mac domain constraint and allowing mac domains of MxN.

In Data over Cable Service Interface Specification (DOCSIS) 1.0 systems, CMs contend with each other to make transmissions and compete for data capacity. This mode of operation—known as “best effort” service—is suitable for e-mail and web browsing applications. Voice over Internet Protocol (VoIP) and Moving Picture Experts Group (MPEG) video traffic, however, require an assured rate of throughput and have strict requirements on latency and jitter that are not supported in a “best effort” (BE) environment.

One of the foremost issues in traffic engineering is to know where the different bottlenecks are in a network. If we open one bottleneck, we need to know the next one that will present itself.

**DOCSIS THROUGHPUT PART 1**

Bit rate, or throughput, is measured in bits per second (bps) and is associated with the speed of the data through a given medium. For example, this signal could be a baseband digital signal or perhaps a modulated analog signal conditioned to represent a digital signal. One type of modulated analog signal is Quadrature Phase Shift Keying (QPSK).

This is a modulation technique, which manipulates the phase of the signal by 90 degrees to create four different signatures as shown in Figure 1. We call these signatures "symbols", and their rate is referred to as baud. Baud equates to symbols per second.

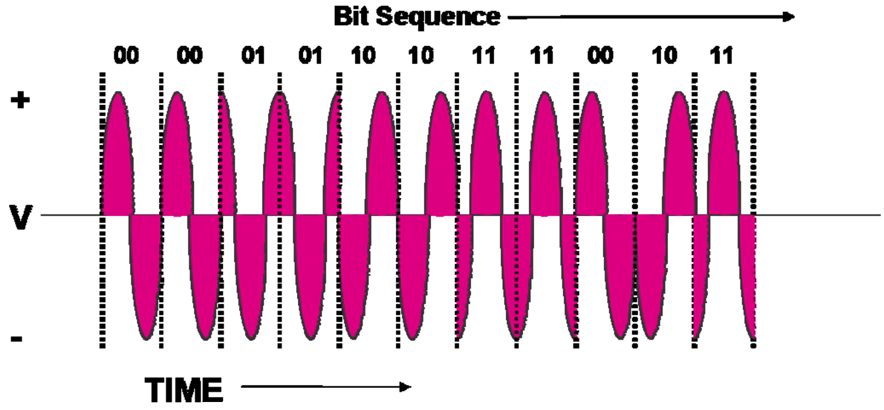

Figure 1 - QPSK Diagram

You may also be familiar with the term PPS, which stands for packets-per-second. This is a way to qualify the throughput of a device based on packets regardless of whether the packet contains a 64-byte or a 1518-byte Ethernet frame. Sometimes the “bottleneck” of the network is the power of the CPU to process a certain amount of PPS and not necessarily the total bps.

**What is Throughput?**

Data throughput begins with a calculation of a *theoretical* maximum throughput then concludes with *effective* throughput. Effective throughput available to subscribers of a service will always be less than the theoretical maximum, and that's what we'll try to calculate.

The age-old question of how many subscribers can one put on the US or DS is still being asked today and the answer is still the same, it depends. Some systems may get away with 1000 devices on an US port with only 50% utilization during peak periods, while another system is relegated to 75 subscribers per US and their utilization is “pegged” at 80% for most of the day. Another system could be constrained to 100 subscribers per US because the mac domain of 1x4 would put the DS at 400 subscribers and that was their “bottleneck”. Demographics can have a drastic affect on bandwidth utilization.

DOCSIS modems rely on a reservation scheme where the CMs request a time to transmit and the CMTS grants time slots based on availability. CMs are assigned a service ID (SID) that's mapped to class of service (CoS)/quality of service (QoS) parameters. In a bursty, Time Division Multiple Access (TDMA) network, we must limit the number of total CMs that can simultaneously transmit if we want to guarantee a certain amount of access speed to all requesting users.

Traffic engineering, as a statistic used in telephony-based networks, signifies about 10% peak usage. This calculation is beyond the scope of this paper. Data traffic, on the other hand, is different than voice traffic, and will change when users become more computer savvy or when VoIP and VoD services are more available. For simplicity, let's assume 50% peak users * 20% of those users actually downloading at the same time. This would equal 10% peak usage.

All simultaneous users will contend for the US and DS access. Many CMs could be active for the initial polling, but only one CM will be active in the US at any given instant in time. This is good in terms of noise contribution because only one CM is adding its noise complement to the overall affect.

Some inherent limitations with the current standard are that when many CMs are tied to a single CMTS, some throughput is necessary just for maintenance and provisioning. This is taken away from the actual payload for active customers. Also, per-modem US speeds can be limited because of the request and grant mechanisms as explained later in this document.

**Throughput Calculations**

Assume we are using a CMTS card that has one DS and six US ports. The one DS port is split to feed about 12 nodes.

The 500 homes/node multiplied by an 80 percent cable take-rate and multiplied by a 20 percent CM take-rate equals 80 modems per node. The 12 nodes multiplied by the 80 CMs per node equals 960 CMs per DS port.

The US signal from each of those nodes will be combined on a 2:1 ratio so that two nodes feed 1 US port. 6 US ports * 2 nodes/US = 12 nodes. 80 CMs/node * 2 nodes/US = 160 CMs/US port.

**Downstream**

DS symbol rate = 5.36 Msymbols/s or Mbaud. A filter roll-off (alpha) of 12% gives 5.36 * (1+0.12) = ~6 MHz wide "haystack" as shown in Figure 3.

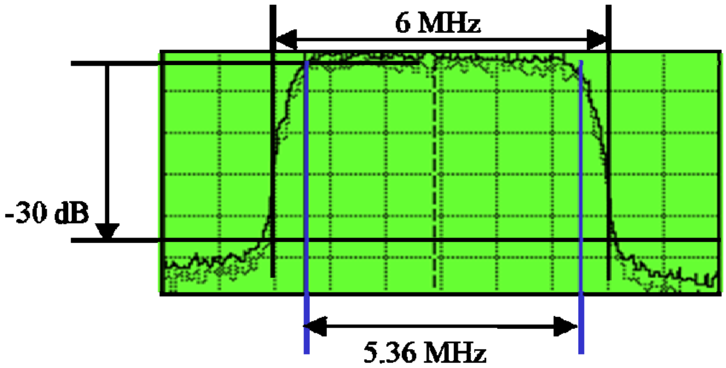

Figure 3 - Digital "Haystack"

Assuming 256-QAM, 256 = 2 to the $8^{th}$ power.  Using the exponent of 8 means eight bits per symbol for 256-QAM and would give 5.36 * 8 = 42.88 Mbps. After the entire FEC and MPEG overhead is calculated, this leaves about 38 Mbps for payload.  This payload is further reduced because it's also shared with DOCSIS signaling.

**Upstream**

The DOCSIS US modulation of 16-QAM at 4 bits/symbol would give about 10.24 Mbps. This is calculated from the symbol rate of 2.56 Msymbols/s * 4 bits/symbol. The filter alpha is 25% giving a bandwidth of 2.56 * (1+0.25) = 3.2 MHz wide.  We would subtract about 8% for FEC.  There's also approximately 5-10% overhead for maintenance, reserved time slots for contention, and "acks".  We're now down to about 9 Mbps, which is shared amongst 160 potential customers per US port.

**DOCSIS THROUGHPUT PART 2**

One of the main reasons for this increased urgency toward smaller node sizes or less customers sharing the "pipe" is the increasing competitive threat.  Cable systems need to provide faster network service to commercial and residential customers to compete against new fiber-to-the-home (FTTH) offerings from Verizon, such as; 10x2 Mbps for residential & 20x5 or 30x5 Mbps for commercial.  This service is being labeled as FiOS, which stands for fiber optic systems.  This new service promises to offer faster speeds and dedicated lines, but the Internet is shared somewhere!

Upgrading to DOCSIS 2.0 increases US throughput to satisfy those new services provided by DOCSIS 1.1.  Also, using the existing HFC network is much more economical than running fiber to the house.

Part of network management is not just adding more speed, but controlling abusers.  If 10% of the users are "abusing" the service and consuming 80% of the bandwidth, then utilizing equipment to shape that traffic or redirect it may keep things manageable.  Some ideas for deep packet inspection (DPI) include P-Cube, Sandvine, and Allacoya. Since these devices can sit in series or parallel with the WAN link, different results can be achieved.  The question that looms: is traffic that remains within the HFC network (on-net) shaped and/or controlled or just traffic that goes outbound?  Some traffic control may need to be done within the CMTS such as subscriber traffic management (STM).  STM allows "byte counting" to control the total amount of bytes subscribers use in a given time frame. Subscribers abusing the system can be dynamically given a lower quality of service (QoS) with a new CM configuration file. Any traffic that is created on the “pipe” steals time off the wire for “real” traffic, but we don’t want to restrict customers for Address Resolution Protocol (ARP) traffic. ARP filtering can be used to control non-compliant end-devices that respond to many ARPs and cause an ARP storm to the CMTS. Other “denial of service” (DoS) attacks occur from “hackers” stealing service, cloning mac address, depleting IP address space, and/or creating a ping of death. These “denial of service” attacks need to be controlled.

**Downstream Bottlenecks**

The first bottleneck to keep in mind is the total usable rate of the DS “pipe”. It may not be as big as first assumed. The usable rate is very dependent on the actual frame size. A 256-QAM DS carrier may be quoted as 37 or 38 Mbps, but it may be closer to 34 Mbps if all the frames are small, 229-B VoIP frames. Another bottleneck is the modem’s DS max rate setting in its configuration file or its Ethernet port. Don’t forget about the PC’s network interface card (NIC). Many people don’t realize that DS speed can be affected by US speed. This becomes apparent when the transport protocol is TCP, which require US acknowledgements before more DS TCP frames can be sent.

Another consideration is the modem’s CPU. For example, if the total CPU packet per second (PPS) limit for the CM is 13 kPPS, then DS TCP could be limited to 8600 DS frames for 4300 US acks. 8600 PPS for DS * 1518 bytes * 8 bits/byte = 104 Mbps. This will be much worse if the frame size is any smaller than 1518 bytes! Concatenation is great for DS TCP speed to increase the rate of US acknowledgment codes (ACK). ACK suppression may be good for CPU load on the CM, but the actual implementation may not save any cycles on the CPU. At least it should save some US bandwidth.

One novel idea is the use of a very large Max DS Burst to give the perception of faster service. This essentially would allow a customer to burst with the entire DS “pipe” for a few seconds before the rate limiting feature would activate. For example; A 3 MB Max Burst would allow a 3 MB file to download in 3*8*1024*1024/36000000 = ~ .7 seconds. After the first second, the rate would be limited to whatever was configured in the modem’s Max DS Rate field. The biggest issue with this idea is the possibility of DS VoIP jitter. If a typical VoIP packet is sent every 20 millisecond (msec), having to wait 700 msec is way too long! That would be 35 VoIP frames lost. The CMTS implementation should provide priority to DS VoIP packets regardless of the best effort (BE) data max burst settings.

Another thing to keep in mind is per-modem speeds. DS speed is affected by frame size and the transport layer 4, which is TCP or UDP. TCP is affected by US speeds because TCP requires US acknowledgements and this affects TCP windowing. Upstream speed is affected by the DOCSIS protocol, map advance, DS interleaving, concatenation, max concatenation and traffic burst, CM and CMTS fragmentation, and modulation profiles. It may be prudent to upgrade the routing engine in the CMTS to meet performance and new feature requirements.

**Upstream Bottlenecks**

Simple math would say that 5 Mbps US for acks is more than enough to achieve super fast DS speed. Typically, TCP windowing is 2 for 1 ack. Assuming the CM is concatenating acks, I would suspect 250 PPS*64*8*x = 5 Mbps, therefore x = 39 concatenated acks possible. Let's suppose only 20 acks are concatenated with an US rate of 250 PPS giving 5000 PPS for acks leading to a DS TCP rate of 5000*2*1518*8 = 121 Mbps. If the DS frame size is only 512, you'll only get 40 Mbps.

Another US rate limiter is map advance delay, which is affected by DS interleave settings and DS modulation, which in turn affects total interleave delay.

The idea of multiple grants per request or outstanding requests with DOCSIS 3.0 is also a good idea, but still doesn't fix the CPU issue on the CM.

Since DOCSIS 1.0 CMs do not support fragmentation, their speed could be very much slower when UGS grants are present than DOCSIS 1.1/2.0 CMs, which support fragmentation. The CMTS scheduler should allocate enough contiguous minislots for a 1518-B frame every so often.

### VoIP Bottlenecks

There are many bottlenecks that must be considered and one the first that comes to mind is the US capacity. If the plant is not very clean, then it may be prudent to use quadrature phase shift keying (QPSK). This, however, will decrease the amount of throughput and subscriber calls that can be supported.

Another bottleneck not usually understood is US scheduling. The scheduler may only be 80% capable, but efficiency really depends on the "pipe" size. This is why a bigger "pipe" is better because of greater statistical multiplexing.

### Future Evolution Options

Cable companies have some choices in regards to throughput offerings that consist of:

- Do nothing and watch the competition erode the subscribers base.
- Segment the fiber nodes, which could cost approximately $10,000 per node. Physical segmentation is the best in regards to frequency re-use and less noise funneling, but very costly.
- FTTC (fiber to the curb), FTTH (fiber to the home), FTTP (fiber to the premise), FTTWAP (fiber to the wireless access point). This is an expensive proposition and wireless may have its own interference unless it's a licensed band.
- Combining of multiple US &/or DS DOCSIS channels. This is the option being promoted, but carries its own pros and cons.

## DOCSIS THROUGHPUT PART 3

Frequency division multiplexing (FDM) allows multiple carriers to be on the same plant. Using 3.2 MHz channel width at 16-QAM on both upstreams can give approximately 9 Mbps per US frequency, but depends on the frame size. Using 256-QAM on both downstreams can provide ~36 Mbps per DS frequency, but also depends on the frame size. Usable rate is very dependent on the actual Ethernet frame size. For example, using 256-QAM on the DS can provide ~ 36 to 37 Mbps of usable rate, but if all the packets on the DS are small voice frames, then the actual usable rate can be diminished by approximately 4 to 5% because of DOCSIS headers on every small frame. This would put the usable rate at 36*.95 = 34.2 Mbps.

One scenario with FDM of multiple carriers is the utilization of load balancing. The caveat to DS load balance is that when a CM moves to another DS frequency, it must move to a different mac domain, which means a new US associated with that DS. The load balance technique may also force the CM to re-register if it doesn't support the proper version of dynamic channel change (DCC).

Mapping two DS frequencies and two US ports on a line card into the same node allows frequency A to serve residential subscribers and frequency B to serve new commercial subscribers. Using client-class processing steers residential subscribers to A and new commercial subscribers to B. This entails setting the DS frequency and/or US channel ID in the cable modem's configuration file. This requires "high-touch" back-office provisioning.

**Pros**

This provides a simple "get up and run" approach. The only real modifications are combining to map USs and DSs. Two DSs at 256-QAM give approximately 72 Mbps of usable rate and 2 USs at 16-QAM /3.2 MHz give about 18 Mbps of total usable rate. More involved modifications require the provisioning system to steer CMs to the proper US and/or DS channel or load balancing can be used for a simpler approach.

The use of US and/or DS load balance allows the subscribers to share the combined USs and/or DSs without requiring CM configuration file changes. The use of US and DS load balancing could be considered a "poor man's" redundancy scheme. If one DS upconverter dies, then all the CMs would register on the other DS upconverter. This assumes the CM has physical connectivity to the appropriate US and doesn't use provisioning to redirect the CM to a specific US channel ID or a specific DS frequency. We do not support dynamic load balance between line cards because each card with an on-board processor and timing source aren't synched between each other.

If the US port is configured as a DOCSIS mixed-mode, then DOCSIS 2.0 CMs could burst with 64-QAM and send their packets faster leading to more open time for other traffic from 1.x CMs. Using 64-QAM on the US at 3.2 MHz channel width can achieve approximately 13 Mbps of usable rate for the 2.0 CMs.

Another option is to configure the second US frequency in the node to ATDMA-only. Refer to figure 4 for a graphical display. Utilizing ATDMA-only on US ports for commercial CMs allows only the 2.0 CMs to "see" that US port and register on it. Existing residential subscribers are blind to an ATDMA port because they don't understand mac message 29, which is included in the upstream channel descriptor (UCD). UCDs are sent for each US that belongs to the specific DS every 2 seconds in a DS map.

**DOCSIS 2.0 Benefits**

DOCSIS 2.0 hasn't added any changes to the DS, but many to the US. The advanced physical layer specification. in DOCSIS 2.0 has added 8-, 32-, and 64-QAM modulation schemes; 6.4 MHz channel width; and up to 16 T bytes of Forward Error Correction (FEC). It allows 24 taps of pre-equalization and upstream interleaving. This adds robustness to reflections, in-channel tilt, group delay, and upstream burst noise. Also, 24-tap equalization in the CMTS will help older, DOCSIS 1.0 modems. DOCSIS 2.0 also adds the use of S-CDMA in addition to Advanced Time-Division Multiple Access (A-TDMA).

DOCSIS 2.0 increases US capacity to 30.72 Mbps. The modulation profiles add interval usage codes (IUCs) for atdma mode. This allows 2.0 modems to burst with different modulation schemes vs 1.x CMs in a mixed environment. DOCSIS 1.x modems use IUCs 5 & 6 for short and long grants while 2.0 CMs use IUC 9 for a-short, 10 for a-long, and IUC 11 for a-ugs.

DOCSIS 2.0 introduces a “DOCSIS-mode” concept, “cable upstream x DOCSIS-mode {}” to configure the US channel to a desired mode. These modes are; TDMA (traditional) mode, ATDMA-TDMA (mixed 1.x and 2.0) mode, and ATDMA (2.0-only) mode.

The advantage of DOCSIS 2.0 and ATDMA usage is wider channels, which results in better statistical multiplexing (a 6.4 MHz channel is better than 2, 3.2 channels), greater spectral efficiency, better use of existing channels and basically more capacity. This provides higher throughput in the US direction and greater per-modem speed with better packets-per-second (PPS) rates. Refer to the references section for more details about DOCSIS Throughput considerations.

An added feature with DOCSIS 2.0 is advanced physical features, which include ingress cancellation, 24-tap pre-equalization (Pre-EQ), analog-to-digital (A/D) conversion, and modulation profile advances. Although not part of the DOCSIS 2.0 spec, ingress cancellation allows higher orders of modulation. Ingress cancellation makes the US port robust against certain plant impairments, opens unused portions of spectrum, and is insurance for life-line services.

### Cons

The cons to this approach include:

- Requires new “high speed” users to have 2.0 CMs.
- Requires provisioning work to "block" 2.0 CMs from registering on the Res freq.
- If Residential subscribers buy their own 2.0 CMs, they could lock to the commercial upstream w/o provisioning interdiction.
- Load balancing can’t be utilized unless DOCSIS mixed-mode is configured with utilization-based load balancing.
- May require dynamic frequency hopping or modulation changes to provide insurance against plant impairments that cause high packet loss for the higher channel width and/or modulation schemes.

## PERFORMANCE IMPROVEMENT FACTORS

### Throughput Determination

There are many factors that can affect data throughput, such as:

- Total number of users
- “Bottleneck” speed
- Type of services being accessed
- Cache and proxy server usage

- Media access control (MAC) layer efficiency
- Noise and errors on the cable plant
- Many other factors such as limitations inside Windows TCP/IP driver

More users sharing the “pipe” will slow down the service and the bottleneck may be the web site being accessed, not your network. When you take into consideration the service being utilized, regular e-mail and web surfing is very inefficient as far as time goes. If video streaming is used, many more time slots are needed for this type of service.

**Increasing Access Speed**

Many systems are decreasing the homes/node ratio from 1000 to 500 to 250 to PON or FTTH. PON stands for passive optical network and, if designed correctly, could pass up to 60 people per node with no actives attached. Fiber-to-the-home (FTTH) is being tested in some regions, but it’s still very cost prohibitive.

The most obvious segmentation technique is to add more fiber optic equipment. Some newer designs are decreasing the number of homes per node down to 50 to 150 households passed (HHP). It does no good to decrease the homes per node if you’re just combining them back again in the headend anyway. If two optical links of 500 homes/node are combined in the headend (HE) and share the same cable modem termination system (CMTS) upstream port, this could realistically be worse than if one optical link of 1000 homes/node were used.

As technology becomes more advanced, new ways will emerge to compress more efficiently or send information with a more advanced protocol that is either more robust or less bandwidth intensive. This could entail using DOCSIS 1.1 Quality of Service (QoS) provisioning, payload header suppression (PHS), or DOCSIS 2.0 features.

There is always a “give-and-take” relationship between robustness and throughput. More speed out of a network is directly related to the bandwidth used, resources allocated, the robustness to interference, and/or cost.

## OTHER FACTORS

There are other factors that can directly affect performance of your cable network such as the QoS Profile, noise, rate-limiting, node combining, over-utilization, etc.

## NEW TECHNOLOGY CORNERSTONES

New technologies are being pursued to address the DS bottleneck conundrum. DOCSIS 3.0 uses a channel bonding technique to achieve higher capacity links, enable faster high speed data (HSD) service, and provide M x N mac domains to enable Video over IP solutions. Using a CM that can tune 3 DS channels gets us to the illustrious 100 Mbps with a less expensive CM. Most customer PCs are limited to 100 Mbps anyway.

The modular CMTS (M-CMTS) architecture is promoted to achieve better DOCSIS economics, lower cost DS PHY, and de-couple DS and US ports. One day we may see fiber optic nodes with DOCSIS physical layer chips embedded so we can use ingress cancellation at the node, digital links from the node back to the headend without the need to amplify, and no more laser clipping. Of course, this means all traffic needs to be DOCSIS-based on the US!

Part of DOCSIS 3.0 is the support of a modular CMTS architecture to allow the physical-layer technology such as Broadcom QAM chips to be physically separated from the mac-layer chips. This could allow the demodulation of the US signal from many devices in nodes or hubs sites to then be processed by one packet-shelf which would do all the DOCSIS processing.

Providing more services and speed per customer, and more VoIP with very small packets can cause CPU overload. More packets being processed translates to more CPU utilization.

**SUMMARY**

The previous paragraphs highlight the shortcomings of taking performance numbers out of context without understanding the impact on other functions. While you can fine-tune a system to achieve a specific performance metric or overcome a network problem, it will be at the expense of another variable. There is always a trade-off and compromise relationship between throughput, complexity, robustness, and/or cost.

**CONCLUSION**

Knowing what throughput to expect is the first step in determining subscriber data speed and performance. Once it is determined what is theoretically possible, a network can then be designed and managed to meet the dynamically changing requirements of a cable system.

The next step is to monitor the actual traffic loading to determine what's being transported and determine when additional capacity is necessary to alleviate any bottlenecks.

Service and the perception of availability can be key differentiating opportunities for the cable industry if networks are deployed and managed properly. As cable companies make the transition to multiple services, subscribers' expectations for service integrity move closer to the model established by legacy voice services. With this change, cable companies need to adopt new approaches and strategies that ensure networks align with this new paradigm. There are higher expectations and requirements now that we are a telecommunications industry and not just entertainment providers.

Whatever the future has in store, networks will get more complex and the technical challenges will increase. The cable industry will only be able to meet these challenges if it adopts architectures and support programs that can deliver the highest-level of service integrity in a timely manner.

**References**

1. Advanced PHY - http://www.cisco.com/en/US/products/hw/cable/ps2217/products_white_paper09186a008017914d.shtml
2. Understanding DOCSIS Throughput Issues http://www.cisco.com/en/US/tech/tk86/tk168/technologies_tech_note09186a0080094545.shtml
3. DOCSIS 2.0 ATDMA Configuration on MC5x20S and MC28U Linecards http://www.cisco.com/en/US/tech/tk86/tk168/technologies_white_paper09186a0080231fc3.shtml
4. Virtual Interfaces and Frequency Stacking Configuration on MC5x20S and MC28U Linecards http://www.cisco.com/en/US/tech/tk86/tk804/technologies_white_paper09186a0080232b49.shtml
5. How to Increase Return Path Availability and Throughput http://www.cisco.com/offer/sp/pdfs/cable/cable_land/returnpath_wp54.pdf
6. 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
7. Understanding Map Advance http://www.cisco.com/en/US/tech/tk86/tk89/technologies_tech_note09186a00800b48ba.shtml
8. Carrier-to-Noise Ratio in Cable Networks http://www.cisco.com/en/US/products/hw/cable/ps2209/products_white_paper0900aecd800fc94c.shtml
9. Determining RF or Configuration Issues on the CMTS http://www.cisco.com/en/US/tech/tk86/tk319/technologies_tech_note09186a008012fb28.shtml
10. Cable Modem Provisioning Scenarios Doc http://www.cisco.com/en/US/tech/tk86/tk89/technologies_white_paper09186a008025c169.shtml
11. Chapter 4: Connecting the Cisco uBR7200 Series Router to the Cable Headend http://www.cisco.com/en/US/products/hw/cable/ps2217/products_installation_guide_chapter09186a008022610b.html
12. Implementation of 16-QAM to Increase the Value of Cable Modems http://cisco.com/en/US/tech/tk86/tk319/technologies_white_paper09186a0080237c17.shtml
13. Upstream Modulation Profiles for Cable Linecards http://www.cisco.com/en/US/tech/tk86/tk319/technologies_tech_note09186a008020560f.shtml