



AN OVERVIEW OF THE DOCSIS (CABLE INTERNET) PLATFORM

January 14, 2010

Matt Tooley, Vice-president, Consulting Solutions

Don Bowman, Chief Technology Officer

Introduction

Sandvine was established in 2001 and employs over 400 people in Canada, the United States, Israel, and in remote offices globally. Sandvine was recently named to the Deloitte Technology Fast 500 list of fastest growing technology companies in North America.

Sandvine is the global leader in network policy control solutions. Sandvine's solutions make the Internet better by protecting and improving the Internet experience for subscribers. The solutions comprise network equipment and software that help cable, DSL, FTTx, fixed wireless and mobile operators understand network traffic and trends, mitigate network congestion, protect the quality of experience for sensitive applications, offer subscribers new services, mitigate malicious traffic, and improve customer service.

Sandvine's technology is used by more than 180 Internet service provider customers in over 70 countries, including over 60 service providers in the United States alone, almost half of which are cable broadband operators. Together, Sandvine's customers serve over 80 million fixed line broadband subscribers and more than 200 million mobile subscribers.

Matt Tooley is a veteran of the broadband networking equipment and cable industries and is currently Sandvine's Vice-president of Consulting Solutions. Matt has been Chief Technology Officer for two network solutions providers to the cable industry. In a senior engineering role at 3Com, Matt was instrumental in developing the industry's first Data Over Cable Service Interface Specification (DOCSIS) Cable Modem Termination System (CMTS). Matt has been and continues to be a key contributor and author to the cable industry's PacketCable Multimedia (PCMM) Specification¹ and PacketCable 2.0 QoS Specification², both critical tools for the delivery of broadband service over cable networks.

At Sandvine, Matt's team has global responsibility to provide technical consultation for pre and post sales support for all of Sandvine's solutions. Prior to joining Sandvine, Matt was Chief Technology Officer of CableMatrix, a developer of policy management platforms for broadband service providers, most notably in the cable market. CableMatrix was purchased by Sandvine in 2007. Matt was also a co-founder and Chief Technology Officer for Xinnia Technology, a startup developing Operations Support Solutions software for dynamic Quality-of-Service (QoS) management for the broadband industry, including cable operators, which was ultimately acquired by CableMatrix.

Among other roles, Matt has also held senior engineering positions with networking equipment companies Tellabs and 3Com. At Teradyne, Matt was a product manager for a DOCSIS system testing

¹ Cable Television Laboratories, *PacketCable™ Specification, Multimedia Specification, PKT-SP-MM-I05-091029*. See <http://www.cablelabs.com/specifications/PKT-SP-MM-I05-091029.pdf>

² Cable Television Laboratories, *PacketCable™ 2.0, Quality of Service Specification, PKT-SP-QOS-I02-080425*. See <http://www.cablelabs.com/specifications/PKT-SP-QOS-I02-080425.pdf>

solution. Matt started his career at NASA, where he designed spacecraft communication systems for three different satellites.

Matt holds a Bachelors of Science degree in Computer Engineering and a Masters of Business Administration from the University of Chicago.

Don Bowman is Sandvine’s Chief Technology Officer and a co-founder of the Company. Don is a globally recognized expert in network policy control in cable and other access networks. He is an inventor on two United States patents³ related to network policy control, and has other patents pending. Don is a regular contributor to the International Engineering Task Force (IETF) efforts in the development of TCP/IP standards. Don has spoken at notable events such as Broadband World Forum Asia, SCTE Cable-Tec Expo, ISS and the Communications Futures Program at MIT.

At Sandvine, Don is currently responsible for leading Sandvine’s technical vision, including the strategic development, direction and future growth of Sandvine’s products and solutions. Previously at Sandvine, Don was Vice-president, Consulting Systems Engineering and led Sandvine's development engineers and technical service consultants.

Prior to Sandvine, Don led the firmware and software engineering efforts for PixStream, a manufacturer of networking equipment and software to help network service providers and enterprises reliably distribute and manage digital video. In 2001, PixStream was purchased by Cisco Systems, where Don managed software engineering efforts for the resulting Video Networking Business Unit.

In 2007, Don and his fellow Sandvine co-founders were honored with the Ernst and Young Entrepreneur of the Year award in the technology category. Originally from Ottawa, Don attended the University of Waterloo’s bachelor of applied science program for systems design engineering.

Executive Summary

Investments in cable networks to date have delivered excellent results. Today, cable networks offer very high bandwidth broadband networks in the United States – up to 105 Mbps. Cable is also the first network to offer compatibility with IPv6, the next-generation of IP that will allow for the ongoing proliferation of IP-enabled devices. These innovations have resulted from extensive investments by cable network operators, such as:

- Moving networking equipment further to the edge of the network, in order to better use available RF spectrum;

³ United States patent 7,571,251 “*Path optimizer for peer to peer networks.*” See <http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetacgi%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=7571251.PN.&OS=PN/7571251&RS=PN/7571251>” and United States patent 7,376,749 “*Heuristics-based peer to peer message routing.*” See <http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=%2Fnetacgi%2FPTO%2Fsrchnum.htm&r=1&f=G&l=50&s1=7376749.PN.&OS=PN/7376749&RS=PN/7376749>

- Extending the deployment of the fiber portion of the network closer to the edge of the network;
- Advancing the DOCSIS platform, through innovations such as PacketCable Multimedia for delivering application-specific QoS, improved RF modulation, and RF channel-bonding.

However, theoretical and practical limits for the network are approaching, which demand that the network be managed more intelligently. There has always been intelligence throughout the network, and current and future traffic trends are likely to make that more and more necessary. Several properties of a cable network, all of which will be expanded on in this document, need to be understood before determining how regulation of broadband services could impact the network.

1. Bandwidth growth has limits.
 - The physical properties of the coaxial cable component of cable's hybrid fiber-coaxial network and the lack of available RF spectrum (most cable networks have 750 MHz of spectrum, which is all allocated to the network's various services) limit theoretical and practical growth in bandwidth on the network.
 - To take advantage of additional bandwidth offered by latest version of DOCSIS, subscribers have to upgrade their modems, and adoption can be slow.
2. The network is shared at the edge.
 - At the access edge of the network, anywhere between 25 and 2000 homes may share the network connection; typically, about 500 homes share a single "node" in this fashion. This architecture raises user-to-user fairness issues, as certain users may disproportionately consume network resources.
 - As the number of homes increases, the amount of bandwidth consumed by the data needed to manage these modems increases quickly, resulting in an efficiency loss and a practical maximum number of modems per node.
 - Emergence of a new application, or a network attack, can damage the experience for an entire cable node (typically 500 users) or even an entire Cable Modem Termination System (CMTS) (typically 10,000 users).
3. Upstream capacity is limited.
 - Cable broadband networks are built asymmetrically, with more RF (radio frequency) spectrum (and therefore bandwidth) in the downstream direction than in the upstream. This design results from the network's roots in transporting television signals downstream.
 - Scheduling of the upstream capacity between many users can result in variable availability of the upstream link for a given subscriber, which can in turn increase latency and jitter – a problem for time-sensitive interactive applications.
4. Network noise consumes bandwidth.

- Because the coaxial portion of cable networks is analogous to a giant antenna, it is prone to “noise” on the line, particularly in the upstream direction on the network. Certain lawful devices when connected to the network can exacerbate this noise. Noise increases packet loss, which in TCP/IP forces bandwidth-consuming retransmission of packets, which in turn increases the need for network management.
5. Multiple services delivered over the same network.
 - Cable networks transport multiple services including broadband Internet, analog and digital television, and managed services such as video-on-demand and voice-over IP. Each service comes with different end user expectations, and in the case of managed services with specific service level agreements.
 6. A wide variety of QoS techniques are needed due to network limitations
 - IETF standard techniques, such as DiffServ⁴ (DSCP) marking, are not effective in cable due to inherent technical limitations in the networking equipment. Also, the marking of traffic by application vendors can no longer be trusted as some applications mask their identity in order to receive higher priority in the network. Consequently deep packet inspection (DPI)-enabled devices are needed to accurately identify traffic.
 - Cable-specific QoS tools like PacketCable Multimedia are effective but offer a limited number of “service flows” per subscriber and have a limited ability to identify traffic – again requiring DPI devices.
 - IP Multicasting is only available starting in DOCSIS 3.0, which is not universally deployed, but even then cable equipment vendors are still implementing aspects of the solution to enable multicasting.

Background on Cable Television Networks

Cable Television was invented in 1948. The first cable television networks were built purely with coaxial (copper) cable. The original “all coaxial” networks experienced high signal loss and were prone to noise at the connections. To address the noise and maintenance issues, cable operators upgraded their networks to use an architecture called the hybrid fiber-coaxial (HFC). HFC networks generally use fiber-optic facilities to carry signals to a centralized “node,” and then rely on coaxial cables to carry traffic from the node to individual homes. With HFC networks, noise could be isolated to nodes and had much lower signal loss, allowing operators to build larger networks. In addition to addressing the noise and maintenance issues, moving to the HFC architecture gave operators the option of offering additional services, such as interactive television. In general HFC networks provide cable operators upstream bandwidth, lower noise or better signal noise ratios (SNR) and therefore faster transmission rates both upstream and downstream. During the mid 1990s, many cable television operators started exercising

⁴ RFC 2474 Definition of the Differentiated Services Field (DS) in the IPv4 and IPv6 Headers, <http://tools.ietf.org/html/rfc2474>

this option by upgrading their networks to become bidirectional, paving the way for Internet-access over this coaxial network.

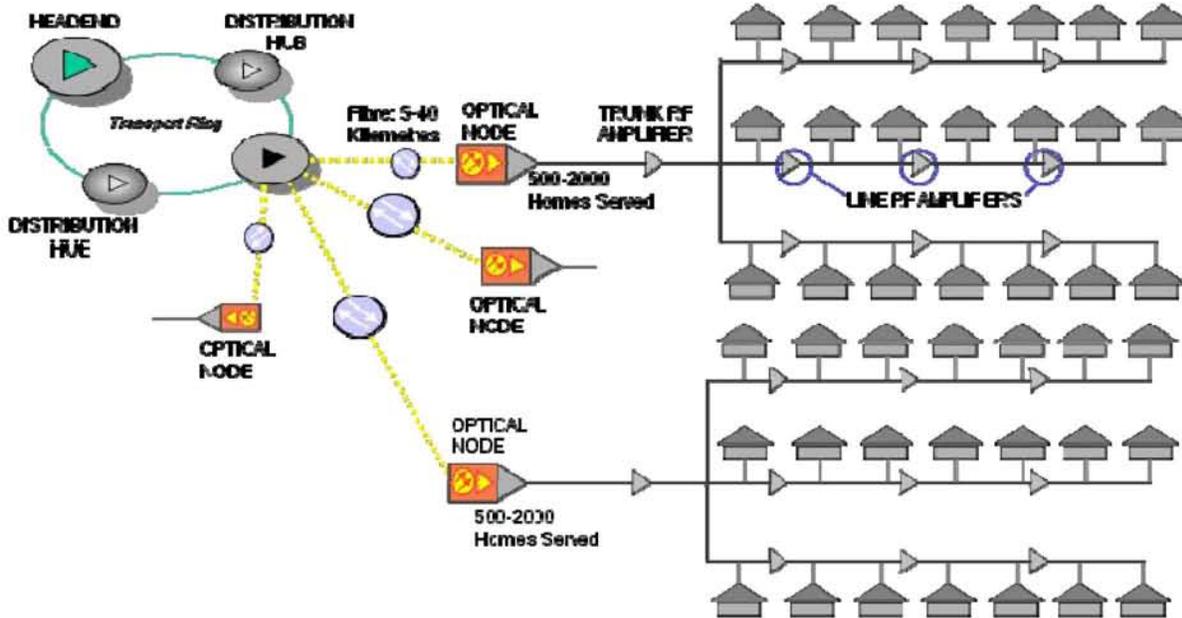


Figure 1 Hybrid Fiber-Coaxial Network

This development is shown graphically in Figure 1 Hybrid Fiber-Coaxial Network. The headend, at the top left, is analogous to the central office in a telephone network. The headend is the master facility for processing all the services that need to be distributed over the HFC network. These services include television signals, video-on-demand, Internet, and voice. The various services delivered on the network are encoded, modulated and unconverted onto RF (radio frequency) carriers, combined onto a single electrical signal and inserted into a broadband optical transmitter. This optical transmitter converts the electrical signal to a downstream optically modulated signal. The cable headend distributes these services to a large serving area via analog fiber transport rings. Distribution hubs that sit on these rings pick up the analog fiber and re-transmit this RF spectrum into a set of optical transfer nodes. Fiber optic cables connect the headend or hub to optical nodes in a point-to-point or star topology, or in some cases, in a protected ring topology. An optical transfer node then converts the fiber back into coaxial copper, and it is connected to approximately 25 to 2000 (500 is typical) homes in a tree-and-branch configuration off of the node. A cable head-end might exist in each city, and distribution hubs are located in neighborhood serving areas. The optical transfer nodes are located outside, typically slung below wires on telephone poles.

HFC networks typically support downstream channels from 50 MHz through 750 MHz, though some newer networks support as much as 1 GHz. Signals traveling from the homes to the headend (the upstream) are transmitted at frequencies between 5-42 MHz. The spectrum is divided into 6 MHz channels to mirror the over-the-air analog TV channel spectrum. A 750MHz HFC network has about 5

Gbps of downstream digital bandwidth available for all the services that are transported (analog and digital TV, broadcast and on-demand, and data services such as VoIP and Internet) using the HFC network and about 200 Mbps of upstream bandwidth, representing the bandwidth that is available in the coaxial portion of the network. The bandwidth is primarily limited by the line amplifiers. The electrical signals in the coaxial network are transmitted over great distances and experience signal loss. Therefore, the network includes a series of line amplifiers to re-amplify the signal. Line amplifiers are the shoe-box size silver boxes that are in-line with the cable strung along power lines in neighborhoods. Line amplifiers are designed to amplify a limited band of spectrum, i.e. 50-750 MHz in most cases. To increase bandwidth beyond the upper limits of the line amplifiers would require upgrading all the line amplifiers in the network. As one can imagine, such a project would be costly in terms of both labor and equipment. Ultimately the bandwidth of the network is limited by the bandwidth (quality) of the coaxial cable. A service provider can increase the available bandwidth of a HFC network by adding more optical transfer nodes, which in turn reduces the number of homes served per node. This is a costly activity as it occurs outside, in residential areas, and entails rerouting power and obtaining right-of-way locations. A service provider can also make a trade-off in number of TV channels offered versus Internet bandwidth offered.

HFC networks in general are still prone to noise particularly in the upstream direction as the coaxial network operates like a giant antenna and all the noise in the upstream gets collected and amplified at the fiber nodes. Noise has an impact on how much data can be transported in the network as more robust modulation and coding for error correction must be used to ensure the data is transmitted free of errors. Radio noise is inserted into the network by the wiring inside the homes of the consumers, coming from CB radios, microwave ovens, cellular phones, vacuum cleaners, etc.

Upstream bandwidth is limited by the aggregate effect of this noise from all households. Upstream bandwidth is also limited by the initial spectrum allocations from the early analog TV days and decisions by the designers of DOCSIS that consumers would require more bandwidth downstream than upstream (reflecting the behavior of Internet consumers during dial-up days). These limits on upstream bandwidth require that a cable operator must provision the network so that the number of upstream homes is smaller than the downstream. Regardless of capital investment, radio noise and early spectrum allocations will act as limiting factors for upstream bandwidth.

HFC networks are unique in that they:

1. Support two-way operation over the same wire – Other transport technologies may use a separate physical media (i.e. wire or fiber) for downstream and upstream respectively. Using the same wire requires that frequency be allocated for downstream and upstream respectively and this partitioning of frequency cannot be easily changed. As a result there is finite amount of frequency set aside for upstream transmission (5-42MHz) and a soft limit of spectrum for the downstream before needing to replace all the line amplifiers.
2. Are asymmetrical – More downstream spectrum than upstream spectrum as a result of having its roots in originally transporting broadcast television downstream.

3. Use common transport for multiple services - All the services use a common transport protocol, MPEG2, but MPEG2 may carry anything from a data packet to a digital television signal. HFC networks were some of the first multi-service networks to be widely deployed.

Introduction to Data Over Cable Service Interface Specification (DOCSIS)

With the demand for high-speed Internet the cable companies enlisted CableLabs (a [not-for-profit research and development](#) consortium of cable operators) to develop a standard method for transporting Internet data over the same HFC network that they were using to transport broadcast television – that standard is known as DOCSIS (Data Over Cable Service Interface Specification).

DOCSIS was developed by CableLabs and contributing cable market vendors including Arris, BigBand, Broadcom, Cisco, Harmonic, Intel, Motorola and Texas Instruments. It marked the entry of the multiple service operators (MSOs, or cable companies) into the broadband market. The first version of the specification (DOCSIS 1.0) was issued in March 1997 (although proprietary implementations had been functioning since 1994), with the first major revision (DOCSIS 1.1) following in April 1999. DOCSIS 1.1 primarily added support for quality of service (QoS) to enable multiple services to be delivered over the Hybrid Fiber Coax (HFC) architecture. DOCSIS 2.0 was introduced in December 2001, primarily to increase upstream speeds. DOCSIS 3.0 (August 2006) was introduced to provide further speed increases in both the upstream, up to approximately 100 Mbps, and downstream, in excess of 100 Mbps, as well as introducing support for IPv6 (next-generation IP)⁵. Limitations due to the original analog TV spectrum selection remain inherent in the data services that are today operated over the same infrastructure.

The strong industry support for DOCSIS, coupled with the widespread deployment of HFC, led to DOCSIS becoming an early leader in broadband in the US market, where it currently holds about 50% market share.

⁵ IPv4 (the current generation IP) uses a 32-bit address system, which is proving insufficient for the growing demand for IP addresses. IPv6 overcomes this limitation through a 128-bit address system.

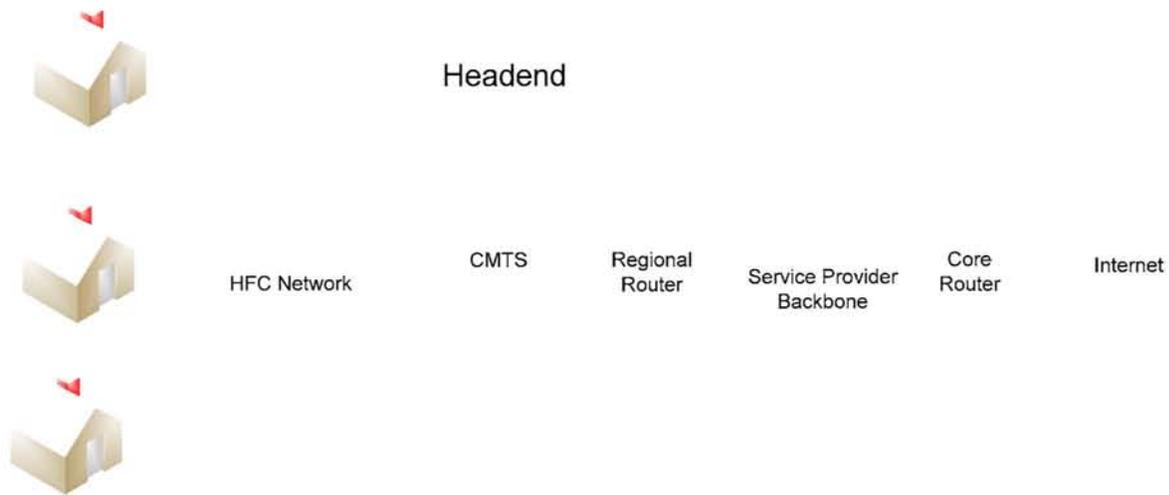


Figure 2 HFC Network with DOCSIS for IP Transport

DOCSIS networks in simple terms are IP over HFC networks with the IP packets being encoded and transported as a digital television signal intermixed among the other services being transported across the HFC network. As shown in Figure 2 the Cable Modem Termination System (CMTS) connects the HFC network to the IP network. Within the homes that connect to the HFC network are devices referred to as cable modems, which convert (encodes/decodes) the packets to/from the HFC network. To keep costs to the consumer down, economics require that the cable modem is a simple device.

Between the CMTS and the Internet is usually a backbone that networks all the service provider's headends. Depending upon the service provider, they may or may not have their own backbone. Most service providers have a 10Gbps network connecting their HFC network to the Internet. Originally most DOCSIS service providers used asynchronous transfer mode (ATM)⁶ transmission technology as their core network, but today ATM is being phased out in favor of Ethernet⁷.

Technical discussion of DOCSIS

Bandwidth in US cable is allocated into 6MHz channels of spectrum. This number comes from the broadcast TV analog spectrum, carried forward into analog cable, and from there into data services over cable. This core decision made more than 60 years ago continues to drive limitations and decisions in state of the art cable communications today. In general, the spectrum is split so that 0 to 45MHz is the upstream (data sent from the home PC to the Internet), and 50MHz to the maximum plant quality are available for the downstream (data received at the home PC from the Internet). The downstream is shared between analog television broadcast (1 TV channel/ 6MHz), digital television broadcast (approximately 8-10 channels/ 6MHz), video on demand, and DOCSIS data. DOCSIS data in the cable

⁶ ATM is a circuit-switched networking technology, and was the primary competitor to IP. ATM networks are being phased out in favor of IP networking.

⁷ Ethernet (IEEE 802.3) was originally a local-area-networking standard but has become common in the telecommunications field as a low-cost means of connecting devices via IP.

network acts like a digital cable TV channel in the downstream. The evolution of DOCSIS has been to move the switching of 6MHz increments closer to the consumer edge so as to re-use as much of the spectrum as possible. The net effect is that the number of households per node has been going down over time and this starts to become analogous to other networks like cellular phone, Wi-Fi, pico-cells, femto-cells, and WiMAX where consumers use a shared media to connect to the first network hop. The re-use of spectrum is also analogous to allowing lower-power transmission in smaller serving areas with radio or TV.

In DOCSIS 1.0 and 1.1, the theoretical upstream bandwidth limit is 10 Mbps per allocated RF channel and the theoretical downstream bandwidth limit is up to 38Mbps per channel. The bandwidth available per consumer becomes a function of how many homes are sharing the same channel. In general, a consumer shares their upstream with a different group of users than they share their downstream. In the downstream direction, approximately 500 homes share the same radio spectrum. This means that, if on average 1 home in 4 subscribed to DOCSIS Internet service, there would be 125 homes that would be actively sharing the 38Mbps. If we assumed that at peak time there were 60% of these users online, this would imply approximately 300kbps of bandwidth per user is available. Each user can burst up to the speed governed by their data plan with their provider, typically speeds such as 10Mbps. This 'high-peak low-average' property is not unique to DOCSIS, but rather is a function of almost all network technologies ranging from corporate LANs to DSL to mobile, and dates back to the inception of the Internet. If an operator were to try and guarantee the 10Mbps plan in the access network, it would mean that only 3 users could share the downstream RF, an increase in cost of 42 times. Statistical oversubscription on cable, like on other networks such as roads, water, electricity, provides an economical means of providing a high quality service.

DOCSIS 2.0 increased the upstream bandwidth limit per allocated RF channel from 10 Mbps to 30 Mbps by both doubling the channel bandwidth from 3.2MHz to 6.4MHz and adding support for two additional modulation techniques that encode more bits per Hz.

DOCSIS 3.0 introduced the concept of 'channel-bonding', which allows multiple channels to be joined together to linearly increase the per-user bandwidth and with it the current capability to offer bandwidth up to 160 Mbps in the downstream direction and 100 Mbps in the upstream direction. Note that DOCSIS 3.0 does not change the amount of RF spectrum that is available, it merely allows more of it to be used by a single user. The economics of a cable service provider do not improve with DOCSIS 3.0. The current DOCSIS 3.0 specification specifies the minimum bonding group size that must be supported. Currently this is set at four channels in each direction which provides 4x38Mbps or approximately 160 Mbps for the downstream and 4x30Mbps or approximately 120 Mbps for the upstream. The DOCSIS 3.0 specifications do not limit the channel bonding. In theory, much larger bonded channel groups can be defined. The limiting factors are: 1) free 6 MHz RF spectrum, and 2) equipment from vendors that support larger groups. Currently most cable plants have 750 MHz of spectrum available, which is partitioned into 6 MHz channels. Typically almost all the 6 MHz channels are in use so for a cable operator to increase the DOCSIS 3.0 bandwidth they must re-allocate one of the 6 MHz channels for data. In other words, the bandwidth doesn't come for free, and comes at the expense of television or video-on-demand services. The second limiting factor is equipment availability. At this time, very few

vendors can economically offer equipment that can bond more than 4 channels due to high marginal costs related to a relatively low level of demand.

In general, DOCSIS 1.0, 1.1, 2.0, and 3.0 users may share some channels as part of the interoperation. This means that bandwidth is shared across a group of four channels for a DOCSIS 3.0 consumer, while DOCSIS 2.0 customers might be on one of the 4 channels alone. This can lead to load-balancing and fairness issues since the higher bandwidth consumers can ‘starve’ the lower bandwidth consumers. Management of the network is necessary to avoid this result. This is shown graphically in Figure 3: Mixed DOCSIS 2.0 and DOCSIS 3.0. In this example network, there are three different, overlapping, sharing groups in the downstream direction. In this example the users in DBG1 are all assumed to be DOCSIS 2.0 users as they are receiving only one downstream channel and the users in DBG2 and DBG3 are DOCSIS 3.0 users.

This means that users of type ‘DBG1’ all share with, and interfere with, all users in type ‘DBG1’. The actions of any one or more of these users will have an impact on all of the remainder. Since the total theoretical bandwidth available to a single user of type ‘DBG1’ is 38Mbps, and since there are typically 125 active subscribers per cable channel, this means that about 250 users could in the worst case be competing for the same bandwidth (assuming the load – balancing was worst-case). The theoretical bandwidth of type ‘DBG2’ is 4x38Mbps (152Mbps). Users in type ‘DBG2’ might have a service plan of ‘up to 20Mbps’, and in type ‘DBG1’ they might have ‘up to 10Mbps’. Without proper network management it is possible for just two users of type ‘DBG2’ to consume all of the bandwidth available to all users of type ‘DBG1’. An active network management approach is required to enforce fairness in this shared network.

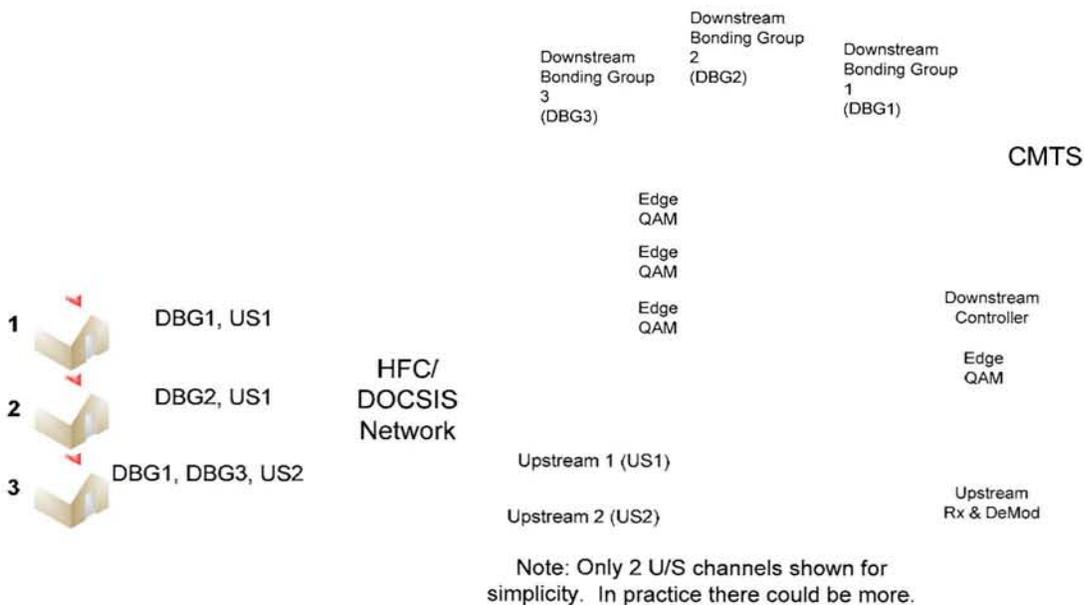


Figure 3: Mixed DOCSIS 2.0 and DOCSIS 3.0

Although technically DOCSIS is backwards compatible (older modems continue to work), generally the end user must upgrade his or her device to get the new features and bandwidth of later DOCSIS versions.

DOCSIS defines a data-link protocol for the transmission of Internet Protocol (IP) packets across the cable TV distribution plant. DOCSIS uses a shared RF channel in each direction, and implements an RF scheduler to control the arbitration to the network.

In the upstream DOCSIS uses a time-division multiplexing scheme to manage when endpoints can use the shared media. The time is arbitrated into 'mini-slots' of 6.25us (microseconds), and all allocation is centrally managed by the cable modem termination system (CMTS), networking equipment that enables communications with subscribers' modems. These 'mini-slots' are sent in a Media Access Protocol (MAP) message, which gives the transmit opportunities for approximately the next 15ms at a time. These mini-slots give a coarse quanta of the bandwidth that a given user can consume in the upstream. It is impossible to consume non-multiples of 6.25us. Long packets consume multiple mini-slots, and make the upstream network temporarily unavailable for all other users sharing the same channel. A single user transmitting at high rate with large packets can create significant latency issues for other users using multiple small packets, such as VoIP. A simple analogy is the check-outs at the grocery store where each checkout represents mini-slot and consumer's grocery carts represent packets. Some grocery carts hold more items and take longer to process than others. So even though all the checkout lines may have the same number of people waiting, some lines go faster and the time it takes to process each grocery cart varies.

As a result of the upstream scheduling, there can be variance between the opportunities when an endpoint can transmit. This variance is a function of how many other users are contending to use the shared media as well as the size of their transmissions. Big packets take longer to transmit than short packets. This variance causes packet jitter and latency which can impact time-sensitive applications.

In the downstream, DOCSIS effectively use a first-come first-served approach to packet forwarding where packets from the Internet side are forwarded as DOCSIS-encoded packets and modulated on to a downstream RF channel for transmission. DOCSIS 1.x and DOCSIS 2.0 use a single downstream RF channel at a time. With the channel bonding introduced in DOCSIS 3.0 there is also a downstream manager to manage the multiplexing (load-balancing) of packets on to the bonded channel groups. The downstream manager routes the IP packets to the channel with the best matching transmission characteristics – speed, latency, and jitter. This improves the overall delivered quality for a given amount of RF spectrum, in the typical case, but may have worst case behavior that is no better.

To mitigate the effects of first-come first-served packet forwarding, DOCSIS includes a logical concept called a "service flow". A service flow is a uni-cast virtual circuit with specific bandwidth, jitter, and latency characteristics. DOCSIS has two default service flows for data - one for the upstream and one for the downstream. This data pair is defined to provide best-effort service where best-effort means no QoS guarantees. Other services flows that require QoS may be provisioned for managed services like voice, where the QoS is managed by assigning scheduling attributes appropriate to the service flow.

Historically, DOCSIS networks use a minimum of two service flow pairs, one pair for VoIP, and one pair for all other data.

Non-managed (or ‘over-the-top’) services like third party voice and video providers all use the shared best-effort data service flows since there is no signalling control mechanism to allow them to request a guaranteed QoS. Services that require more bandwidth than provisioned for the user’s default services flows may experience quality issues (delays or retransmitted packets). To ensure that these services achieve sufficient quality third-party content providers must design their application to work with the available bandwidth provisioned for the subscriber, and with the typical observed jitter and loss. This means that an application provider must be aware of the empirical limitations of a DOCSIS network to achieve optimal quality. As DOCSIS providers continue to innovate and improve, this goal moves and expands.

DOCSIS limitations

DOCSIS networks use simple IP forwarding rules and route IP packets very close to the consumer (in the CMTS). This can lead to very specific challenges for identifying traffic, and for acting on the traffic.

DOCSIS has a number of limitations.

1. Like all networks, DOCSIS is shared amongst multiple users. In common with wireless, this sharing occurs on the last mile of the network. DOCSIS networks suffer from RF “noise” introduced in the upstream channel from each home. One user with badly shielded wires, or with a CB antenna, or running HomePNA⁸ networking can create a significant noise problem affecting 100 or more people. These noise problems can then impact the shared throughput for all the users on a node. A node in a cable network is the junction where the optical signal carried by fiber is converted to an electrical signal that fans out and is transmitted the last mile using coaxial cable to the homes. A single node typically supports 25 to 2000 homes. Noise can also be introduced by other acts of nature such as cable breaks from squirrels chewing through cables, home owners cutting cables with shovels, and rain leaking into equipment. All of these coupled with age and wear on the network can cause the signal to deteriorate and impact the overall throughput in the various nodes in the network. When impacts from the noise become severe enough DOCSIS operators are forced to find and fix the source of the noise and/or take some other kind of action to match bandwidth supply with bandwidth demand, such as network management.
2. DOCSIS 1.0 and DOCSIS 1.1 modems have a significantly reduced upstream capacity compared to DOCSIS 2.0 and 3.0. Because the user needs to upgrade their modem to reduce this limitation, adoption rates tend to be slow.

⁸ HomePNA is a home-networking technology that re-uses the existing coaxial wire in your home to form a local area wired network. It is fundamentally incompatible with DOCSIS since it uses the same upstream spectrum. A single user running HomePNA without having disconnected from the public cable plant can affect the upstream bandwidth of all shared users.

3. As the number of cable modems per channel increases, the amount of bandwidth consumed by the data needed to manage these modems increases quickly. This results in an efficiency loss and a practical maximum number of modems per node.
4. DOCSIS is primarily a link-layer (local-connection) specification and therefore DOCSIS networks do not, in general, interoperate with Internet Engineering Task Force (IETF)-standard QoS approaches such as Differentiated Services Code Point (DSCP)⁹ or Explicit Congestion Notification (ECN)¹⁰. These IETF standards, designed for application-specific QoS and congestion management, are not translated into the RF-domain scheduler of the DOCSIS network, and are treated strictly on the Ethernet side of the network, giving no advantage to cable operators. The IETF, amongst other standard engineering bodies, has spent considerable time focusing on application-specific quality of service and optimizing application performance on networks. Unfortunately this is still an area of active research for shared RF networks such as DOCSIS.
5. Instead of replicating and sending the same packet to multiple recipients, there is a technique in IP called multicast. Multicast is a technique used for sending the same IP packet to multiple users when they are trying to see the same content, for example broadcast video over IP, which can save significant bandwidth. Multicast is particularly useful for broadcast-like services. Generally, multicast is not currently solved in DOCSIS 1.x and 2.0. DOCSIS 3.0 solved the multicast problems, but multicast also requires a new version of the PacketCable Multimedia (PCMM) specifications (see Platform-specific concerns below for a discussion of PCMM) to be completed and implemented by the DOCSIS vendor community. This process is still underway.
6. IPv6 is not supported by DOCSIS 1.0 or DOCSIS 1.1. IPv6 is technically supported on DOCSIS 2.0, but in practical terms no implementations support it. DOCSIS 3.0 is the first version with true IPv6 support, with more QoS and better support for service flows that map end-to-end. DOCSIS 3.0 is also the *first broadband access technology* platform to support IPv6, an indication of how the adaptability and innovation underlying this time-tested network technology.

Platform Summary

The DOCSIS platform has evolved over its 12-year lifespan from a simple to complex broadband platform. DOCSIS is unique from other broadband platforms in that it:

- Runs over a combination of an electrical (coax) and optical (fiber) network that is commonly referred to as a Hybrid Fiber Coax (HFC) network. The coax network is susceptible to environmental impacts that can degrade its performance;
- Supports bandwidth in excess of 100 Mbps per customer;

⁹ RFC 2474

¹⁰ RFC 3168

- Permit shared media by 100-1000 subscribers using time division multiplexing in the upstream and first come-first served in the downstream;
- Is asymmetrical – downstream or download speeds are much faster than the upstream or upload speeds;
- The end points are simple – Cable modems are simple devices that for all intents and purposes are analogous to simple media converters to convert from Ethernet media to cable TV coaxial media. Therefore, all intelligent routing and packet filtering can only be done at the access router or further towards the core of the network.
- Service Flows for QoS enforcement allows quality of service to be enforced through the entire access layer

DOCSIS was originally designed to match the prevailing asymmetric usage patterns and expected speed demands of subscribers at the time. As Internet applications and content evolved so too did the user's behavior in both asymmetry and bandwidth requirements. The ratio of downstream to upstream changes over time, from as much as 1000:1 in the early days of dialup, to as little as 2:1 in the peak days of P2P file sharing. Today, according to Sandvine's 2009 Global Broadband Phenomena study¹¹, it is approximately 2.7:1. As with all network technologies, there is continually a lag to update the technology to match the market requirements. As a result of the lag and in the interest of ongoing network optimization, DOCSIS operators need to manage their networks to continue to deliver the expected level of service for current and emerging applications.

Platform specific concerns

The DOCSIS platform provides richer QoS tools than other access networks via its PacketCable MultiMedia (PCMM) interface. PCMM defines an IP-based platform for delivering QoS-enhanced multimedia services over DOCSIS® 1.1 (or greater) access networks. However, these tools are not unlimited. There is a limited number of service flows per subscriber (typically 4-8). The classifiers used to create these service flows have limited capacity to match traffic, operating on IP and TCP-specific constructs that are unreliable.

The shared RF nature, in both the upstream and downstream, provides additional concerns. The emergence of a new application, or a network attack, can damage the experience for an entire cable node (typically 100-500 users), or even an entire CMTS (typically 10,000 users).

Reasonable network management

The practice of managing congestion in telecommunications networks has been around for a long time. The traditional PSTN phone network used call admission control to manage congestion – a call was not

¹¹ See <http://www.sandvine.com/downloads/documents/2009%20Global%20Broadband%20Phenomena%20-%20Full%20Report.pdf>

admitted to the network unless end-to-end capacity existed to handle it. Traditional IETF standardisation for Quality of Service (QoS) calls for the communication endpoints to “mark” traffic with the desired drop/latency/quality characteristics so that routers at each hop can weight their decisions. In IP networks, packets have a field in the packet header that can be used to indicate the relative level of priority for the packet. However, this marking is not universally obeyed by current generation access devices in either DSL networks (ATM L2 backbone), or cable networks (DOCSIS layer), where congestion is highest, due to technical limitations in the devices. Also, the devices can’t trust the marks: over time applications have cheated the system by mischaracterizing their traffic in order to achieve higher priority. Accordingly, this approach no longer works.

A new class of intelligent networking equipment products emerged that could accurately classify network traffic, overcoming the classification shortcomings of DOCSIS. . These devices allow service providers to accurately identify and characterize network traffic and activity and allow for enhanced management of the network.

Application Types

Central to the notion of reasonable network management is the observation that applications differ with respect to the minimum amount of bandwidth, or maximum latency, jitter, and packet loss that they can tolerate in order to be delivered at an expected quality of service level.

- **Bandwidth:** traffic volume over time. It is usually measured over a short time, such as bits/second or megabits/second (Mbps), which is 1,000,000 bits/second.
- **Latency:** the delay for a message to get from one communications end point to the other. For example, the time it takes for a VoIP data packet to leave the speaker’s mouth and arrive at the listener’s ear. It is typically measured in milliseconds.
- **Jitter:** the variation in the latency of one packet to another, typically measured in milliseconds. For example, if the first message takes 1 ms and the second message takes 10 ms, then there is 9 ms of jitter.
- **Packet Loss:** occurs when one or more packets of data traveling across a computer network fail to reach their destination. Packet loss is typically measured in percentage terms and can be caused by a number of factors, including signal degradation over the network medium, oversaturated network links, corrupted packets rejected in-transit, faulty networking hardware, faulty network drivers or normal routing routines.

While bandwidth gets most of the attention, adding bandwidth is not always (or even mostly) the answer to improving the user’s quality of experience for an application. The other factors can play a critical role. An application can be classified into one of three categories (bulk, interactive, paced/burst-paced), based on its requirements of a network across these four characteristics:

Bulk. These applications include P2P filesharing (e.g., BitTorrent, FastTrack, etc), web surfing, usenet news (NNTP), and file transfers over FTP or HTTP, for example, and will go as fast as the network will permit. TCP is designed to achieve the maximum communication rate possible. In practice bulk applications will go as fast as the thinnest part of the network between the client and server. In the case of the server collocated within the ISP network (e.g. a content-delivery network, a cache), this will be bound by the access equipment speed. In the case of a server which is located farther away, this may be bound by transit (connection to all worldwide public networks) or peering (connection to other nearby private networks) performance. Typically servers of bulk applications (e.g. Speedtest.net, Rapidshare.com, Megaupload.com) will saturate the download speed of the consumer's modem, as they typically download-only. In the case of P2P, it is bi-directional so it can also have the same effect in the upstream direction.

Most bulk applications can run unattended by the user. File transfers are initiated by the user who may then walk away – often for hours or even overnight – while the process completes. Bandwidth is the primary determinant of transfer speed and performance will generally improve linearly with increases in bandwidth. As a result, latency, jitter and packet loss matter much less – users likely would not even notice their effect.

Web surfing represents an exception in the Bulk category. “Web 2.0” sites have introduced interactive components to web surfing – the user typically attends the activity and data is traveling bi-directionally as users have started to be content providers in their own right. Increases in bandwidth do not translate linearly to increased performance, because it takes several “round trips” between a personal computer and the related web servers to load a website – typically at least four: the Domain Name Server (DNS) lookup¹² and the three-way handshake established by Transmission Control Protocol, or TCP, one of the core protocols of the Internet Protocol Suite¹³. Each of the four round trips is subject to the latency in the network, and when added together this delaying effect becomes the limiting factor in the transmission such that additional bandwidth does not dramatically improve loading times for a website.

Interactive. These applications are paced by the consumer. In the case of VoIP, bandwidth largely depends on silence suppression and the codec bandwidth chosen, but it is typically 8-30kbps. The bandwidth requirements of interactive applications are often modest (though in the case of video conferencing the rates are significantly higher: 200-500kbps is common), but they typically require very low latency, jitter and packet loss to achieve a satisfactory quality of experience. For example, a VoIP user can perceive latency of 150 milliseconds on a call, and delays greater than 300 milliseconds will prevent voice communication¹⁴. As with web surfing, adding bandwidth will not necessarily address quality of service issues. In general, because of the sensitivity of Interactive applications to latency, jitter and packet loss it is particularly important to protect the quality of service for these applications.

¹² See IETF RFC 1035 at <http://www.ietf.org/rfc/rfc1035.txt>

¹³ See IETF RFC 793 at <http://www.ietf.org/rfc/rfc0793.txt>

¹⁴ See <http://voip.about.com/od/glossary/g/latency.htm>, or T. Blajic, D. Nogulic, M, Druzijanic, *Latency Improvements in 3G Long Term Evolution*, p. 1-2, available at http://www.ericsson.com/hr/about/events/mipro_2007/mipro_1137.pdf, or http://www.telephonyworld.com/training/brooktrout/iptel_latency_wp.html.

Paced/Burst-paced. Streaming applications such as YouTube and SHOUTcast fall into this category. The media involved has a natural bit rate, and the connection tries to achieve this rate on average over its lifetime, though for short durations the media will 'burst' to provide buffering on the client to allow for packet loss on the network (YouTube, because it uses TCP, will attempt to transmit at line rate when possible). So, these applications can be modeled by the media they carry. For typical Internet streaming today, rates of approximately 300-400kbps are common. Hulu, YouTube, and others are starting to shift to higher-definition video, for which the rate can increase to 1-6Mbps of bandwidth.

With paced/burst-paced applications it is important that a network sustain the minimum bandwidth requirements, but because of the buffering involved additional bandwidth only marginally improves performance, by making the applications less sensitive to latency, jitter and loss in the network.

The following table provides some representative benchmarks to achieve a minimum quality of service for certain popular applications. Such figures require significant assumptions, which Sandvine has included as Appendix 1:

Application Category	Application Class	Minimum Bandwidth	Maximum Latency	Maximum Jitter	Maximum Loss
Bulk	P2P	19Kbps	n/a		
	Web surfing	1Mbps (Web 2.0)	166ms (latency + jitter)		n/a
	Email	60Kbps	n/a		
	Usenet news	195Kbps	n/a		
	FTP file transfers	195Kbps	n/a		
Interactive	VoIP	16Kbps	300ms (latency + jitter)		< 0.5%
	Video gaming	50Kbps	75ms (latency + jitter)		< 0.5%
	Video Conferencing	250Kbps	300ms (latency + jitter)		< 0.05%
Paced (and burst-paced)	Video streaming	300Kbps, to not have much of a wait time	< 1s for "channel change"	<50ms	<0.05%
	High def video	1-3Mbps depending on quality of HD.	< 1s for "channel change"	<50ms	<0.05%
	Audio streaming	Audio: 128Kbps for CD quality. 56Kbps for radio	< 1s for "channel change"	<50ms	<0.05%

Network Management Approaches

There are two commonly used technological approaches to managing data packets. The first is to limit the rate at which the packets of a specified class enter the network (through Traffic Policing), which can delay session transmission. The second is to affect the priority of access to the network. These approaches are outlined below.

- Traffic Policing. A method of ensuring that packets of a specified class do not exceed a specified bit rate. For example, a 4Mbps service offered by a service provider is policed to 4Mbps per subscriber. Traffic policing can be applied per subscriber or per application.
- Traffic Prioritization. A method for selecting which packet is transmitted next as the network has capacity. For example, by default routers give all packets equal priority, and when a packet drop occurs it may be on a packet that is more sensitive to loss or is more valuable to the subscriber at that moment. Traffic policing can be enhanced through prioritization techniques, which apply different classes of services to packets, giving each class a different priority. Prioritization can be strict. For example, if ever there is a packet of the top priority class queued for delivery, it will go first. Alternatively, to avoid the risk of starvation, priority can be weighted so that packets of lower priority classes are assigned a probability of delivery to the proportion represented by their relative weighting (e.g., every fifth packet if in a class with 20% weighting). Priority assignments and their weighting are entirely configurable. The classification of packets can be done on a per-flow, device, application and/or subscriber basis. When the network does not have sufficient capacity, the lowest-priority packets are usually delayed and then dropped. Prioritization is a very efficient method of network management since it achieves the highest possible utilization of a network, and very good overall quality. Without active network management, prioritization is random, and drops occur without regard to effect on applications.

Application-centric and Subscriber-centric Policies. The ability to enact network management policies on a per-application and per-subscriber basis greatly enhances the quality of experience to the end user and the “fairness” of the network. Such policies can allocate scarce network resources efficiently by taking into account the different characteristics of bulk, interactive and paced traffic, as well as different subscriber usage patterns for these applications.

- Application-centric. Prioritize interactive, real-time network applications that are sensitive to latency and jitter (e.g., VoIP, online gaming) and that most affect the quality of the users’ immediate Internet experience. Bulk protocols that are only bandwidth-intensive and that are typically unattended by the user (such as FTP or P2P file transfers) can accept lower priority without any meaningful impact to the user’s quality of experience.
- Subscriber-centric. Prioritize the traffic of subscribers who are not contributing disproportionately to congestion over a given time period, so that others are free to consume their “fair share”. A “fair share” based policy is designed to ensure fairness across users. A

recent study conducted by Sandvine showed that over a month the top 1% of subscribers are responsible for 25% of total bytes on the network and the top 20% of subscribers account for fully 80% of total Internet traffic¹⁵. Sandvine's customers' data also shows that the demand of disproportionate users is largely inelastic during peak network loads. Based on 15-minute sample periods, 85% of the time the same users consuming a disproportionate share of the upstream bandwidth in one sample period will also be doing so in the next sample period. For the downstream direction, 60% of the time a user consuming disproportionate bandwidth in one sample period will also be doing so in the next sample period. One of the goals of a "fair share" policy can be to encourage disproportionate users to shift their usage to off-peak times so that bandwidth can be divided fairly among users throughout the day.

- Subscriber- and Application-centric. Apply priority in a subscriber-centric manner and also prioritize latency- and jitter-sensitive interactive applications even for disproportionate users. This highly targeted approach preserves the quality of experience of affected subscribers to the maximum extent possible. In the future, Sandvine believes that cable service providers will be able to provide tools to their subscribers to let them directly select which applications should be prioritised in the event that they are identified as users contributing disproportionately to network congestion.

Upstream versus Downstream Policies. Network management policies can also be applied specifically to upstream and/or downstream traffic. As cable networks have historically been designed with less upstream bandwidth than downstream bandwidth, this is a particularly important capability for MSOs. As more applications and websites have encouraged the transmission of data upstream, the upstream path has frequently been the first point in the network where congestion is experienced. Consequently, separate network management policies for upstream and downstream traffic may need to be considered.

DOCSIS-specific Network Management

Owing to the shared nature of the network and the relatively coarse (15ms) time scheduling in the upstream, latencies per subscriber can increase or become highly variable without network management. As DOCSIS has matured from 1.0 to 3.0, however, the requirements in the upstream have become simpler as spectrum has become less scarce.

Network management is required in DOCSIS networks to ensure optimal operation of the network including:

- Enforce user to user fairness.
- Ensure QoS of time-sensitive Internet applications
- Mitigate abusive traffic from SPAM, malware, mis-behaving applications

¹⁵See <http://www.sandvine.com/downloads/documents/2009%20Global%20Broadband%20Phenomena%20-%20Full%20Report.pdf>

- Enforcement of service level agreements (SLAs)
- Managed Services work as intended

While all the network management approaches identified earlier, (e.g., traffic policing, traffic prioritization by application or subscriber etc.) are possible in a cable network through the deployment of DPI-enabled policy control solutions, DOCSIS itself only offers a limited set of tools to manage their network, including:

- Static provisioning of bandwidth caps (i.e. speed tiers) to police the subscriber’s maximum bit rate.
- Service Flows can be used to create virtual circuits across the DOCSIS network to make paths with specific QoS (jitter, latency, and bandwidth) characteristics. The service flows may use a form of traffic management to achieve the desired QoS. The service flows can be either dynamic or static in nature. The service flows are limited to classifying packets based upon what is commonly referred to as a 5-tuple classifier (source IP address, destination IP address, source port, and destination port, and IP protocol). A single end-point can typically support no more than eight simultaneous service flows.
- Access Control List (ACLs) – ACLs can be used at the first-hop router to filter using a 5-tuple classifier. ACLs are useful for downstream traffic management, but are limited in use in the upstream due to the fact that the traffic has already been transported over an expensive leg of the network before it can be filtered.

DOCSIS itself can only provide management of traffic using a 5-tuple classifier. The 5-tuple classifier may not be enough to correctly identify and manage the traffic. In addition, the number of simultaneous 5-tuple classifiers that can be in use by an end-point or are router is limited. In these situation DOCSIS operators have to use non-DOCSIS specific technology, such as DPI-enabled network policy control solutions, to enable them to manage traffic using alternative identification techniques at layers 3-7.

The DOCSIS Network and the “Open Internet” Rules

The Federal Communications Commission (the Commission) has proposed six rules in connection with its October 22, 2009 Notice of Proposed Rulemaking (NPRM), “In the Matter of Preserving the Open Internet, Broadband Industry Practices” (FCC 09-93), GN Docket No. 09-191, WN Docket 07-52. These rules could potentially have significant impact on a cable network, as addressed below.

Additionally, the phrasing of all rules as obligations of service providers poses a problem, as all participants in the Internet – subscribers, application providers, content providers and network providers – can affect one another. In such a symbiotic environment, it is troublesome to assign responsibilities to one and only one network participant, as that participant may not even be at “fault” for a network problem. For example, one application (or subscriber) can consume sufficient bandwidth to preclude the use of other applications by subscribers sharing the same cable network connection to

the Internet. In this case, how would the cable operator fulfill its obligation to ensure that subscribers can run all lawful applications (i.e. Rule #2)?

Rule #1 – Sending and receiving lawful content and Rule #2 – Running lawful applications and services

The bandwidth of a cable network has natural limits. The lack of available RF spectrum (most cable networks have 750 MHz of spectrum, which is all allocated to the network's various services) limit theoretical and practical growth in bandwidth on the network. Also, radio noise and the historical allocation of radio spectrum between the upstream and downstream parts of the network limit upstream capacity in particular. Typically, upstream bandwidth is shared at the edge of the network among 500 homes.

All these limitations can make it difficult (and will make it more difficult in the future) to deliver all applications and content satisfactorily at all times. Upstream bandwidth will need to be managed both on a per-subscriber and per-application basis to ensure that the maximum number of users receive the experience they expect from each application (and its related content). In particular, latency-and jitter-sensitive interactive applications will need to be protected against bulk applications that can consume a disproportionate amount of bandwidth yet aren't typically time-sensitive. Reasonable network management will need to be defined with the user's needs in mind with respect to these applications. Not all network management techniques available to other access technologies, such as such as DiffServ¹⁶ (DSCP) marking, are available in cable networks.

Additionally, certain applications and content, which may be lawful, may also be malicious or objectionable to users, such as Spam, spyware, worms, child pornography, etc. Removing such applications and content from the network is beneficial to all Internet participants, and network providers need to retain the flexibility to aggressively manage or block such traffic as appropriate.

Also, certain content providers, such as ESPN360, charge network providers (not end users) for the right to have their subscribers access the content. To be clear, that means that subscribers can only access ESPN360 if their network provider has an agreement in place with ESPN360. A strict interpretation of a rule stating that network providers are obliged to guarantee that their users can send and receive all lawful content could be construed as an obligation for all network provider to carry ESPN360 since, it is lawful content.

Finally, subscribers' access to lawful content can be blocked or impaired, at least temporarily, as a result of surges in subscriber activity, not through any actions of the network provider. Recent events such as Michael Jackson's death resulted in widespread reports of subscribers' inability to access related content as web servers and networks were overwhelmed.¹⁷ The event didn't take place during typical "peak" times for broadband networks, so capacity would likely have otherwise been available. The

¹⁶ RFC 2474 Definition of the Differentiated Services Field (DS) in the IPv4 and IPv6 Headers, <http://tools.ietf.org/html/rfc2474>

¹⁷ MSNBC.Com, *Texts and tweets spread news about Jackson; Twitter, Facebook, cell phone companies, Web sites report surge in traffic*. See http://www.msnbc.msn.com/id/31566668/ns/technology_and_science-tech_and_gadgets

proposed rule could require cable network providers to build their networks to guarantee that subscribers have unrestricted access to content at all times, even if doing so means networks must be built un-economically and inefficiently.

Rule #3 – Connecting and using lawful devices

There are instances in which lawful devices attached to a cable broadband network can affect the performance of multiple users on the network, but it is unclear whether such performance degradation would qualify as “doing harm” to the network as anticipated by the Rule #3.

Certain lawful devices when connected to the network can exacerbate radio noise in the upstream direction. For example, one user running HomePNA networking can create a significant noise problem affecting 100 or more people. HomePNA is a home-networking technology that re-uses the existing coaxial wire in your home to form a local area wired network. It is fundamentally incompatible with DOCSIS, yet lawful. The noise problems caused by HomePNA can then impact the shared throughput for all the users on a node, which can support 25 to 2000 homes. Such noise also increases packet loss, which in TCP/IP forces bandwidth-consuming retransmission of packets, which in turn increases network congestion that affects an even broader group of network users.

Additionally, cable broadband subscribers’ home PCs may become infected by viruses, botnets and other malicious traffic that, if left unmanaged, can affect multiple users in a network.

In these cases, cable operators require substantial flexibility to manage the network for the benefit of most users. Additionally, cable devices may require certification from Cablelabs to prevent network harm.

Rule #4 – Competition among network providers

Any implementation of Rule #4 should recognize that network management practices themselves create new opportunities for competition among cable broadband operators. Such operators are just beginning to explore the use of network management practices to help them create service offerings that are more attractive to consumers in an increasingly competitive Internet access market. In the United States, high-speed Internet services are largely offered in the form of flat-rate, monthly plans. Consumers may be interested in other types of service plans that better reflect the unique ways that they use their Internet connections. Such plans would necessitate the ability to differentiate between the traffic of individual subscribers, and between applications.

For example, “light” Internet users may be interested in a service package that ties their fees to the bytes they consume on the network. But these consumers would not want to pay for malicious traffic that affected their usage in a month, or visits to the service provider's web service portal to address service issues. Thus, a user- and application-specific policy would be required to manage the plan.

Other consumers may value their Internet connection by the quality of experience they receive for their favorite applications, like latency-sensitive Internet video gaming or VoIP. Network providers could offer a Premium Video Gaming or Premium VoIP service plan that delivers exactly the type of Internet

experience these consumers want. Such plans would need to be supported by application-specific and user-specific policies.

New service plans like these would offer consumers new choices and in so doing create new grounds for competition among network providers.

Rule #5 – Non-discrimination

It is important to distinguish between the very different meanings of the term “discrimination.” Discrimination was originally a neutral term. One of the definitions given by Merriam-Webster is “the process by which two stimuli differing in some aspect are responded to differently” – in other words synonymous with “differentiation.” However, in the context of network management discussions “discrimination” has evolved to take on the meaning of “anti-competitive” behavior when the two terms are not one and the same.

From Sandvine’s experience with many of the leading cable broadband network providers in the United States, network management solutions are deployed in order to “differentiate” between network traffic to enhance the performance of the maximum number of subscribers for the maximum period of time, *recognizing that the performance needs of applications and their content vary*. Such variance in application needs has been described in this document and has been a long-held view of technical bodies, such as the IETF (with DiffServ marking¹⁸) and the European Telecommunications Standards Institute (ETSI) through its work on the IMS and TISPAN standards, including 3GPP TS 29.211.19. Any overly broad interpretation of Rule #5 could discourage such beneficial differentiation, when only anti-competitive behavior is of concern.

Additionally, certain services and applications that consumers would value receiving over their Internet connection are currently not feasible absent differentiation that enables a minimum quality of experience. An example of such services would be telepresence, which is beyond the delivery capabilities of current HSIA networks but could be feasible with appropriate traffic differentiation. Rule #5 could discourage the development of these services, stifling the innovation and the competition that flows from it.

Rule # 6 – Transparency

Cable broadband networks, like other access technologies, are susceptible to malicious attack, such as denial of service attacks, phishing, Spam, and other constantly-evolving threats. For network and subscriber security, it is important that any requirement for disclosure of network management practices be balanced by the need to defend against such attacks. If the rule were written too broadly, providers may be required to reveal information that would leave their networks vulnerable to such attacks.

¹⁸ IETF RFC 2474. *Definition of the Differentiated Services Field (DS) in the IPv4 and IPv6 Headers*. See <http://tools.ietf.org/html/rfc2474>

¹⁹ See <http://www.3gpp.org/ftp/Specs/html-info/29211.htm>

Managed services

DOCSIS is a platform for delivering a range of consumer services over an all-IP network. One of the services is a high-speed Internet service. Other “managed services” may include voice and video. The managed service may have a specific service level agreement to ensure its operation. Managed services are created using DOCSIS service flows. Each service flow creates a virtual private channel with segregated QoS. Service flows are used to create guaranteed QoS, private addressing and security. It is not feasible to create static service flows for all possible managed services (since that would segregate the bandwidth and be inefficient), so CableLabs created PCMM for dynamic signalling. PCMM’s dynamic signalling is used by operators’ applications to signal to the network when the managed service requires network resources to deliver the service to the subscriber. Likewise, the application signals to the network when the service terminates. Dynamic resource allocation is much more efficient than using static provisioning as it allows service providers to effectively support more users with fixed network resources, which is critical in an oversubscribed network. The approach is analogous to how circuits in the over-subscribed PSTN are provisioned (where there are many more subscribers than there are circuits available). The PSTN affects reasonable network management through connection admission-control (e.g. fast-busy), as previously described.

Cable operators historically have provided managed services such as voice, television, and video-on-demand by dedicating RF channels for the delivery of those services. In order to increase the utilisation of the network and reduce costs, network architectures are converging towards an all-IP network and using the DOCSIS network to deliver managed services beyond Internet access. PCMM is used to manage the QoS for the managed services delivered using the DOCSIS network.

Law-enforcement, public-safety, national security

Cable networks support law-enforcement, public-safety and national security issues in a variety of ways.

- CableLabs created the Cable Broadband Intercept Specification (CBIS) for lawful data intercept, allowing an operator to easily implement a CALEA system. This is a cable-specific means of capturing all of the data of a single subscriber as part of a warrant.
- DOCSIS service flows for VoIP create guaranteed bandwidth and QoS and can meet the standards of E911 services. DOCSIS has no support for location-based services, meaning the E911 caller’s physical address needs to be known and provisioned into the system by the DOCSIS operator when the subscriber signs up for the voice service.
- Owing to the dynamic nature of creating DOCSIS service flows using PCMM, the cable network is very well suited to creating an emergency bandwidth segregation as needed for national security reasons.

Technology evolution and innovation

DOCSIS is currently the leader in high-speed bandwidth to the home in the US. In April 2009, Cablevision launched a 101Mbps residential service. In December 2009, Mediacom launched a 105Mbps residential service.

DOCSIS was also the leader in application-enabled QoS with PacketCable and PCMM initiatives, starting in 1997. PacketCable, in particular Network-Call-Signalling (NCS) provided the first wide-scale deployments of VoIP in the US, starting in 2002, and was rolled out widely from 2003 onwards. This marked the first wide-scale consumer application offering requiring network-based QoS.

Today innovation in DOCSIS is largely driven by lowering the cost of bandwidth through simplifying and commoditizing the RF switching equipment, making it more cost effective to move this equipment closer to the edge of the network and keeping the more complex, costly IP equipment in the core. This lowered cost of bandwidth in turn is giving consumers access to very high speed services. The industry is moving to increase managed services over this efficient, QoS-enabled infrastructure for services such as telepresence, cellular backhaul, video on demand, remote medical telemetry, and so on.

With the ever-increasing number of network appliances and Internet-based services for the home, residential subscribers will start to experience self-imposed network problems as these devices and services share the one Internet connection. To overcome this, DOCSIS operators may start to explore how to offer their customers a service to manage these devices and services over the common connection. Such a service would allow their customers to assign priorities to each device/service when contention for shared network resources demands that choices be made.

Even with DOCSIS 3.0, operators do not have infinite bandwidth as they are constrained by the physical bandwidth of the network with the chokepoint being the coaxial portion of the HFC network. Most HFC networks are designed to provide between 750 MHz and 1 GHz of spectrum (i.e. 5-7 Gbps of bandwidth), this spectrum is shared by all the services (analog TV, digital TV, video on demand, DOCSIS data, and voice). To overcome this limitation, DOCSIS operators are forced to continue to extend at a cost the fiber portion of the network further and further into the network. They cannot do this indefinitely and many cable networks are already near the practical limit. As a result, DOCSIS operators have a significant financial incentive to operate their networks in the most efficient manner. And to do this, operators use reasonable network management techniques.

In addition, operators continue to investigate how to move more services to their all-IP delivery platform (DOCSIS), including legacy services, such as broadcast television and video on demand. Moving to all-IP delivery of content, including broadcast television, VoD, voice, means more efficient use of the network, which in turn means there can be more choices available for the consumer for a given capital investment by the operator. For example, moving to all IP allows unlimited channel choices, versus the fixed channel choices of analog or digital television over fixed 6MHz chunks. It allows greater efficiency of advertising insertion, which in turn means more tailored ads, which in turn means fewer ads for the same revenue. It allows for greater interaction by the consumer (e.g. pause/rewind, save, and also 'buy-now'). It allows integration with other services such as caller-id so consumers can see the identity of

voice calls from their TV screen. It allows more customised parental controls and content filters, and so on. Of course, while rich media and QoS-sensitive applications all merging into the Internet pipe gives better user flexibility, it also requires a smarter network. Flexibility will be required to experiment with the most appropriate network management techniques.

DOCSIS has proven to date to be a resilient technology that has adapted to user behavior far beyond how the original inventors of the technology ever imagined. Even so, increasing network demands from subscribers and applications (current and emerging) can cause unforeseen consequences that can only be remedied through the combination of network management techniques and the ongoing evolution of DOCSIS technologies.

Appendix 1 – Assumptions for Application Requirements

P2P, Usenet, FTP: Bandwidth is the most important network characteristic as it affects the time required for these applications to transfer the data. A typical movie is about 700 MB. If a typical user expects, at a minimum, to download a movie overnight (i.e., 8 hours), the minimum bandwidth required would be 195 Kbps. Latency, jitter and loss do not have strict thresholds with these applications, however, the effective TCP throughput is correlated to the loss, latency and jitter for short download sessions due to TCP's slow-start²⁰ algorithm.

Web Surfing: A typical "Web 2.0" website requires approximately 10 to 20 connections to download the approximately 0.5MB to 2MB of data needed to display the page. Studies have shown that to maintain a good user experience this must be done within 2-4 seconds²¹.

To reach a website, the name must first be translated into its numeric IP address via the DNS. This happens for each server that the webpage references. Many webpages have images, videos and advertisements on different servers and thus the Internet browser must resolve each DNS name. Each time DNS is used, 2x the latency (for the round trip) is added to the total time to load the page. In most PC environments, to compensate for jitter, the PC buffers the data, so the latency time is actually (latency + jitter).

For each connection or file that needs to be downloaded, the (latency + jitter) is added multiple times as the browser initiates a TCP connection to the server to retrieve the file. This multiple is usually three, in accordance with TCP's three-way handshake for initiating connections.

Once the connection is established, the time to download the file is a function of the bandwidth available. However, given that websites often have many small files (images, text), TCP is not always able to achieve the full throughput rate due to its "slow start algorithm".

The above argument does not take into account many of the complex algorithms or tools in place such as parallel connections and HTTP pipelining, but does show that bandwidth is not the only determining factor for measuring HTTP quality of experience. In fact, latency and jitter will likely be the gating factors on the user's quality of experience with Web 2.0 websites.

For a 1.5MB webpage with 20 connections to load with a satisfactory user experience, available bandwidth must be at least 1 Mbps and latency+jitter must not exceed 166ms.

Email: A normal text email is between a few kilobytes and a few hundred kilobytes. Email is not instantaneous, however, there is a perception that it is near real-time. To send or receive an email with a large attachment in under a minute, the bandwidth required is approximately 60Kbps.

VoIP: The most basic audio codecs require bandwidth of approximately 16 Kbps (allowing for overhead of the Internet). VoIP is a real-time application that is very sensitive to latency and jitter. ITU-T G.114

²⁰ IETF RFC 2001. See <http://tools.ietf.org/html/rfc2001#ref-2>

²¹ See http://www.akamai.com/html/about/press/releases/2009/press_091409.html

suggests that the maximum one-way latency+ jitter be 150 ms²² (or round-trip 300 ms), above which it becomes noticeable to the end user. Most VoIP protocols use stateless connections (UDP) and have no built in retransmit. Loss must not be over 0.5% for calls to be audible.

Video Conferencing: Similar to VoIP, the application is bi-directional and is highly susceptible to latency, jitter and loss, however the bandwidth requirements are higher due to the addition of the video.

Video & Audio Streaming. These applications are primarily uni-directional. The average normal-definition video on YouTube requires approximately 300Kbps. High-definition videos (depending on the quality, i.e., different encodings) require bandwidth between 1-3Mbps. Because most streaming done on websites like YouTube use HTTP, latency, jitter and loss are not a major concern. Traditional streaming video (RTSP, RTP, etc) are done over UDP and are affected more by loss. Streaming of compressed CD quality audio requires approximately 160Kbps of bandwidth.

²² International Telecommunication Union. ITU-T Recommendation. G.114. See <http://www1.cs.columbia.edu/~andrea/new/documents/other/T-REC-G.114-200305.pdf>