

WIRELINE PLATFORM DECLARATION TIA COMMENTS – OPEN INTERNET DOCKET

1 Introduction

The authors of this declaration are Mr. Kenneth D. Ko and Dr. Kevin W. Schneider. The business address for both authors is 901 Explorer Blvd., Huntsville, AL 35806.

Kenneth D. Ko is currently a Senior Staff Scientist within the Chief Technology Officer (“CTO”) organization of ADTRAN, Inc. (“ADTRAN”). In this position, he performs communications research in support of ADTRAN’s Research and Development organization and represents ADTRAN in the development of communications standards. Prior to his work at ADTRAN, he has held engineering and management positions within the field of communications research and development at Paradyne Corporation and at Rockwell Semiconductor Systems. Mr. Ko earned a Bachelor of Electrical Engineering from the Georgia Institute of Technology in 1980 and a Master of Science in Electrical Engineering from the University of South Florida in 1987.

Kevin W. Schneider, Ph.D., is currently Chief Technology Officer for ADTRAN. In this position, he leads the ADTRAN corporate technical staff, which is responsible for ADTRAN’s research activities, the creation and analysis of new technologies, and participation in industry-wide standards development organizations. He currently serves on the board of the Alliance for Telecommunications Industry Solution (ATIS) where he has chaired the ATIS TOPS Council Optical Access Networks and IPTV Focus Groups, and established ATIS’ IPTV Interoperability Forum. Dr. Schneider holds a Ph.D. in Electrical Engineering from the Missouri University of Science and Technology.

This declaration is prepared in support of comments filed by the Telecommunications Industry Association (TIA) with regard to the Notice of Proposed Rulemaking addressing broadband industry practices (“Proposed Rulemaking”).¹ We make the statements in this declaration based on our personal knowledge and experience as engineers developing broadband access equipment and technologies.

ADTRAN, headquartered in Huntsville, Alabama, is a global manufacturer of networking and communications equipment, with a portfolio of more than 1,700 solutions for use in telecommunications access networks. ADTRAN’s equipment is deployed by some of the world’s largest service providers, as well as distributed enterprises and small and medium businesses. Importantly for purposes of this Declaration, ADTRAN solutions enable voice, data, video and Internet communications across copper, fiber and wireless network infrastructures. Because of the breadth of its product lines, ADTRAN is not wedded to any one last-mile technology. Rather, ADTRAN believes that copper, fiber and wireless will all continue to be used in the deployment of robust, ubiquitous and affordable broadband.

¹ Federal Communications Commission, GN Docket No. 09-191, WC Docket No. 07-052, “In the Matter of Preserving the Open Internet Broadband Industry Practices,” released 22 October, 2009.

2 Wireline Access Network Basics

The basic wireline access network architecture generally comprises three separate sections between the customer premises and one or more Internet Gateways or peering exchanges (the interface between the broadband access provider's network and the Internet or other network providers). These sections include:

- The “middle mile,” or the network between the Internet Gateways, peering exchanges, and the Central Offices (COs) operated by the access network provider. This is usually a high speed fiber network running (today) at Gbps rates or higher. The middle mile network can use different transmission technologies such as Ethernet, ATM or SONET.
- The “second mile,” or the network between the CO and remote terminals or other equipment (if any) in the outside plant (OSP). This network may run over fiber, copper, or a combination of the two, depending on multiple factors, including the number and type of customers being served and the services being offered. As with the middle mile, the second mile may be implemented with a range of transmission technologies. Currently, throughput on the second mile can range from a few Mbps (via multiple T1 links) to over a Gbps (via fiber).
- The “last mile,” or the connection to the customer premises. Depending on the service being offered, the technology used and the design of the network, the last mile may connect the customer premises either to a terminal in the outside loop plant or directly to the CO (eliminating the second mile connection). Technologies used in the last mile include:
 - Asymmetric Digital Subscriber Line 2+ (ADSL2+) over twisted pair copper subscriber loops. ADSL2+ provides throughput of more than 15 Mbps downstream and 1 Mbps upstream over 4000 ft of 26 American Wire Gauge (AWG) copper loop.² Over a Carrier Serving Area (CSA)³ with 26 AWG loops of up to 9000 ft, ADSL2+ provides up to 6 Mbps downstream and 1 Mbps upstream. ADSL2+ can be served either directly from a CO (Figure 1) or from an OSP node (Figure 2).

² Of the wire gauges that are widely deployed in the United States, 26 AWG supports the lowest rates for a given length over DSL, so the figures for rate and reach in this declaration can be considered conservative. In the heavier gauge loops which make up nearly all of the rural and much of the suburban loop plant, DSL performance is better than that indicated here.

³ Carrier Serving Area (CSA) is a telephony loop plant design method that was often used by the Bell Operating Companies when designing networks employing digital loop carrier (DLC) equipment. CSA loops were limited by attenuation (rather than by resistance, as was the case with the earlier developed resistance design methods. The attenuation was related to loop length of 24 and 26 AWG loops via a formula, which reduced to 9000 feet when the loop was all 26 AWG and 12,000 feet when the loop was all 24 AWG cable.

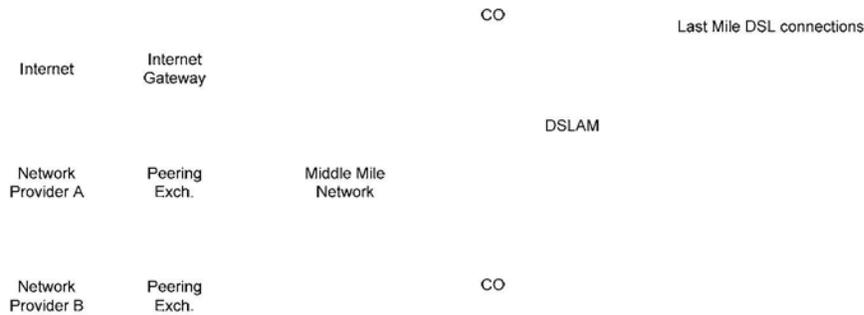


Figure 1 – DSL served from the CO

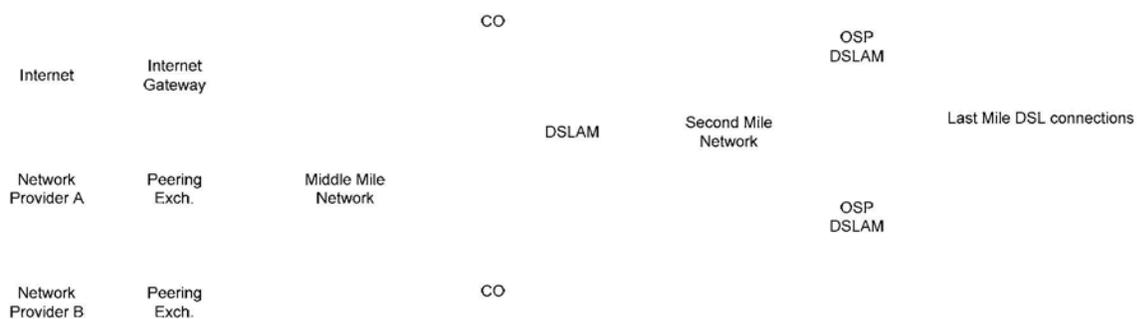


Figure 2 – DSL served from the OSP

- Very High Speed Digital Subscriber Line 2 (VDSL2) over twisted pair copper subscriber loops. As deployed in the United States⁴, VDSL2 can provide over 80 Mbps downstream and 30 Mbps upstream over short loops (1500 ft at 26 AWG) when using 17 Mhz of wireline spectrum. Over longer loops (3000 ft at 26 AWG), VDSL2 provides at least 25 Mbps downstream and 4 Mbps upstream. While VDSL can be served from the CO (Figure 1), it is usually served from a node in the outside plant (Figure 2) to shorten the effective loop length and raise the data rate.
- Symmetric High Speed Digital Subscriber Line (SHDSL) over twisted pair copper subscriber loops. SHDSL provides symmetric rates of up to 5.7 Mbps in each direction on short loops, with lower rates over longer loops. Up to eight SHDSL loops can be bonded together to provide “Ethernet over the First Mile” (EFM) services at up to 45 Mbps. SHDSL can be served either from the CO (Figure 1) or from an OSP node (Figure 2).
- Gigabit Passive Optical Network (GPON) over fiber. A single Optical Line Terminal (OLT) within the CO or in the OSP serves multiple Optical Network Units (ONUs) at customer premises through a Passive Optical network, in which

⁴ VDSL2 is standardized world-wide as ITU-T Recommendation G.993.2. This recommendation provides several regional specific option regarding the allocation of wireline spectrum for use in downstream and upstream transmission. The numbers here reflect the use of bandplan 998, which is used in many regions around the world.

the transmitted signal is optically split onto multiple fibers in the downstream direction and merged onto a single fiber in the upstream direction (Figure 3). GPON currently provides shared rates of 2.5 Gbps downstream and 1.25 Gbps upstream. The split ratio (the number of ONUs served by a single OLT) is typically 1:32, but higher split ratios of 1:64 or even 1:128 are possible. The maximum reach of a 1:32 network is approximately 20 km, with shorter reaches for higher split ratios.

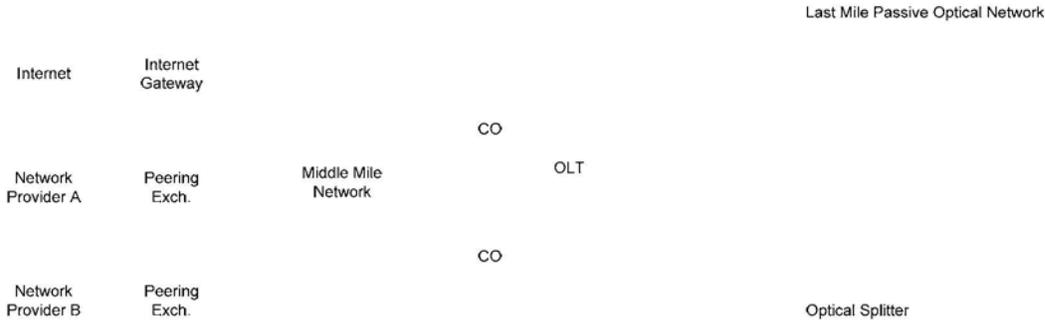


Figure 3 – GPON or EPON served from the CO

- o Ethernet Passive Optical Network (EPON) over fiber (Figure 3). As with GPON, a single OLT serves multiple ONUs with split ratios from 1:32 to as high as 1:128. EPON currently provides symmetric shared rates of 1 Mbps in each direction.
- o Active Ethernet over fiber. Active Ethernet provides symmetric service (usually at 100 Mbps or 1 Gbps) over point-to-point fiber links from Ethernet switches which are usually located in the outside plant (Figure 4).

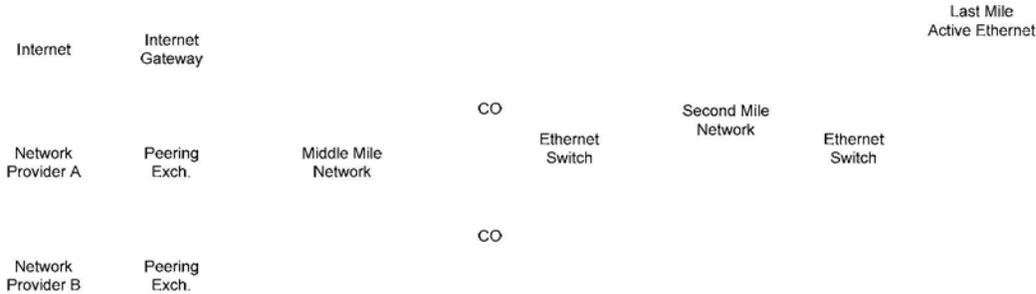


Figure 4 – Active Ethernet

All of the network architectures described above use shared facilities in the middle mile and second mile networks. GPON and EPON also use shared facilities in the last mile, while DSL and Active Ethernet use last mile circuits dedicated to a single customer.

2.1 Sources of Delay and Packet Loss

Traffic is carried across modern networks in the form of “packets,” or small chunks of data that are packaged together with “headers” that contain information regarding the data’s source, intended destination, and other information designed to facilitate its delivery. The primary

impairments that affect packets in wireline access networks are delay and packet loss. Delay has two components – a fixed component that combines the time taken to traverse network connections at the speed of light with the minimum time required for the nodes (routers, switches, DSLAMS, etc.) along the path to process the packets, and a variable component (usually called “jitter”) that results from accumulated time spent waiting within one or more nodes.

The primary cause of both delay and packet loss in wireline access networks is congestion. Congestion can occur at any point in the network where the sum total rate for the traffic entering a node can exceed the rate at which that same traffic can exit the node. This usually occurs under one of two circumstances. In the first circumstance, packets arriving at multiple inputs are forwarded to a common output. In the second circumstance, packets arriving at a fast input are forwarded to a slower output.

When a node takes packets from multiple inputs and forwards them to a common output, that node is referred to as an aggregation point. Many aggregation points exist in a typical wireline access network. In the upstream direction (from the customer into the network), packets from thousands of customers are aggregated at DSLAMS, OLTs, and other nodes in their trips from customers to the Internet or to other networks. In the downstream direction, packets from multiple networks are aggregated together in the middle mile network before being forwarded to their respective destinations.

2.1.1 Queuing Delay and Packet Loss

The bandwidth at the output of an aggregation point is frequently less than the sum of the input bandwidths. That is, the sum of the total potential bandwidths of the inputs being aggregated will often exceed the maximum capacity of the output. This approach is practical because even during times of heavy usage, individuals tend to access the network intermittently, with short bursts of traffic interspersed with longer periods of inactivity. As a result, it is rare for all users to be generating or receiving traffic at the same time even in a small access network. It is common for the sum of the individual peak rates to be many times the capacity of the network’s shared segments, yet for the average traffic (even during periods of heavy usage) to remain well within the network’s capacity. In short, even when the maximum potential bandwidth of the inputs exceeds the output capacity, the sum of the *actual* inputs only exceeds the capacity of the output during periods of particularly high usage.

When the instantaneous load from multiple users does exceed the capacity of the network, excess packets are stored temporarily in a memory buffer, or “queue,” until they can be transmitted. An example of queuing delay at an aggregation node is shown in Figure 5.

If the rate at the inputs exceeds the output rate by too high a margin or for too long a time, the cumulative size of the incoming packets may exceed the amount of memory allocated to the queue. In this case, one or more packets will be discarded. An example of packet loss is shown along with queuing delay in Figure 5.



Figure 5 – Congestion causing jitter and packet loss

A detailed look at Figure 5 explains the mechanism behind queuing delay and packet loss. The switch in the figure is receiving input packets from four ports on the left, with each packet bound for the port on the right. All five ports operate at the same speed. The relative position of each packet on the left shows when it was received at the switch. The first packet to arrive is labeled 101, followed by 301 and then 201, and so on.

Since there are four bursts of packets, all arriving at the switch at virtually the same time on different ports, they cannot all be sent on the same output port as soon as they are forwarded to that port. For instance, when packet 301 arrives at the output port, the port is busy sending packet 101. It is still busy sending packet 101 when packets 201 and 401 arrive. This is a classic example of congestion caused by more traffic received at an aggregation node than can be transmitted by that same node at a given moment in time.

The switch handles this momentary congestion by storing packets in a queue until it can send them. It does so using a simple algorithm called First In, First Out (FIFO), in which packets are read from the queue and transmitted in the same order in which they are received. The time spent by each packet waiting in the queue is the queuing delay, which is shown visually in Figure 5 by the length of each arrow in the upper right section of the figure. As shown, the queuing delay experienced by each packet can be different. In this example, packet 301 must be queued for a longer period than packet 101, and packet 201 must be queued even longer than packet 301, even though all three were received by the network at nearly the same moment. Jitter is the result of this variation in delay as it accumulates between the network source and destination.

Figure 5 also shows an example of packet loss. The first eight packets to arrive at the switch are stored in the queue, but when the ninth packet in the burst (packet 303) arrives, the queue is full. Since there is nowhere to store it and no way to send it immediately, the packet is discarded (or lost).

2.1.2 Serialization Delay

An additional source of both delay and jitter is packet serialization. A packet can contain 1500 or more bytes of data. That data cannot be transmitted across a communications link as a single entity – it must be “serialized,” or sent one bit (or more accurately, a small number of bits) at a time. The serialization delay is the time starting with transmission of the first bit of a packet and ending with transmission of the last bit. This delay increases proportionally with the size of the packet, and decreases in inverse proportion to the speed of the link. For low speed links this

delay can be significant – the serialization delay for a 1500 byte packet on a 1 Mbps link is 12 ms.

An example of serialization delay is shown in Figure 6. In the figure, two packets – a large (1500 byte) packet, immediately followed by a smaller (256 byte) packet – arrive at a high speed (100 Mbps) input to a switch, and are forwarded to a low speed (1 Mbps) output port. The sequence of events is:

1. The first bit of Packet 1 arrives at the switch. We mark the time of this event as $T = 0$.
2. The last bit of Packet 1 arrives at the switch, at $T = 0.12$ ms. The packet is now forwarded to its output port, which is idle. At the input port, the last bit of Packet 1 is immediately followed by the first bit of Packet 2.
3. The first bit of Packet 1 is transmitted from the output port. Ignoring any processing delays inside the switch, this occurs at $T = 0.12$ ms.
4. The last bit of Packet 2 arrives at the switch at $T = 0.14$ ms. The packet is forwarded to the same output port as Packet 1. Since the port is still busy sending Packet 1, Packet 2 is stored in the port's queue.
5. The last bit of Packet 1 is transmitted at $T = 12.1$ ms. Once Packet 1 has been sent, the first bit of Packet 2 follows immediately.
6. The last bit of Packet 2 is transmitted at $T = 14.2$ ms.

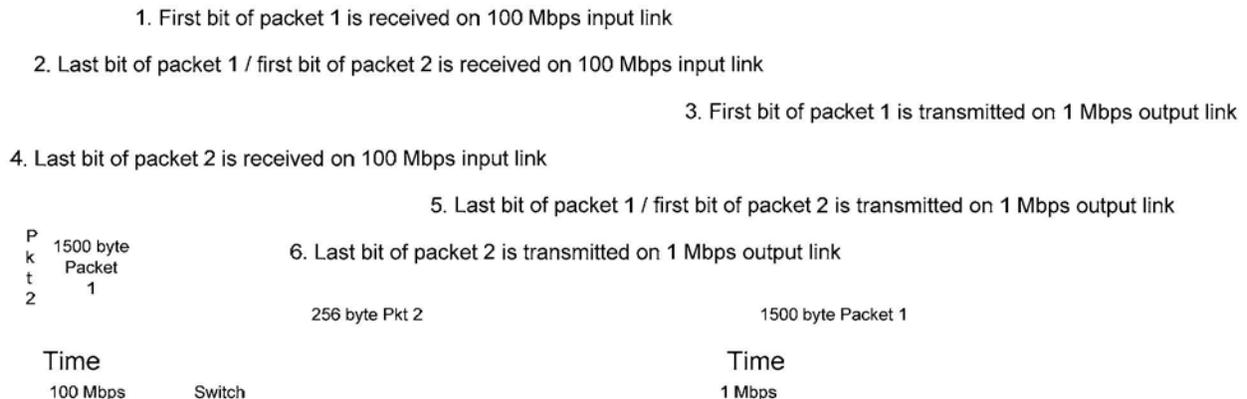


Figure 6 – Serialization delay

In the example shown, the serialization delays for Packets 1 and 2 are 12 ms and 2.05 ms, respectively. However, Packet 2 experiences queuing delay of 12 ms in addition to its own serialization delay, for a total of over 14 ms delay going through the switch. This shows that serialization of large packets on low speed links can add significant amounts of delay and jitter, not just to the large packets themselves, but to the packets that are queued behind them.

Note that since serialization delay is proportional to packet size, a small packet delays the traffic queued behind it less than a large packet does. This characteristic becomes important when we consider network management practices that may give priority to voice packets (which are both small and highly sensitive to delay) over data packets (which can be both large and more tolerant of delay).

3 Wireline Network Architecture

A more detailed and inclusive wireline access network architecture is shown in Figure 7.⁵ The figure shows a converged “multi-service network” architecture, in which multiple types of broadband access customers may be served by different types of services through a common access network and connected to multiple types of service providers.

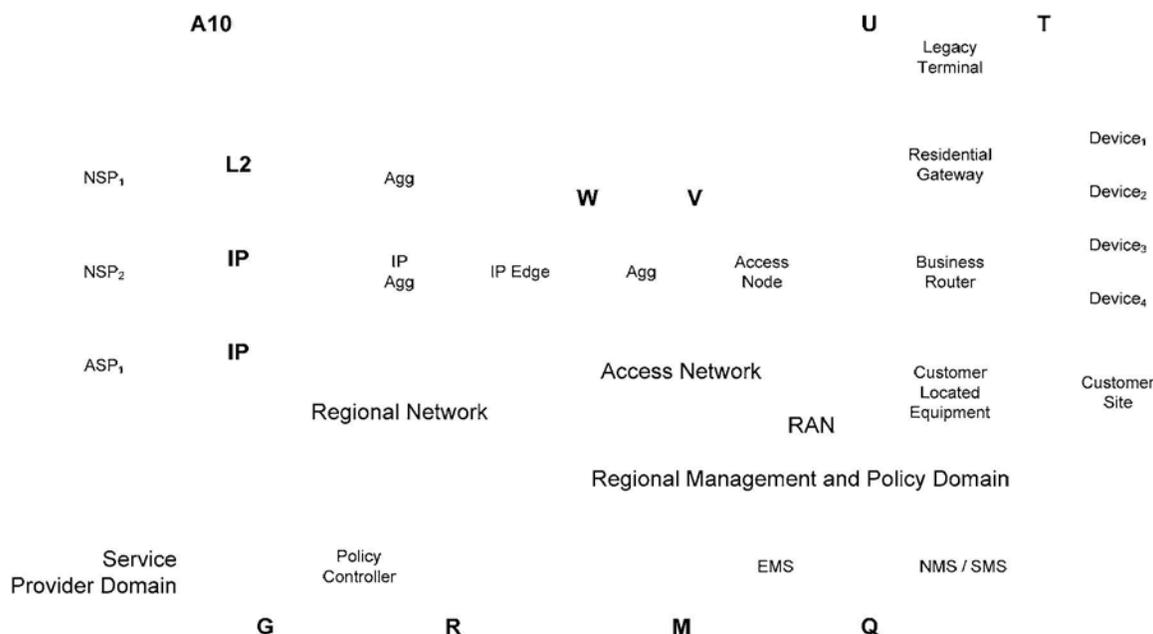


Figure 7 – Broadband Multi-Service Reference Model (from TR-144)

Starting at the right side (the customer premises) in Figure 7, residential and business customers can access multiple types of services over the same access network. Legacy Terminals may include interfaces for analog telephones, facsimile machines, dial modems and other types of equipment designed to interwork with the Public Switched Telephone Network (PSTN). Residential Gateways connect multiple devices within a home to the access network, usually using any of a number of options for home networking (two examples are WiFi and Ethernet). Business Routers provide outside network access to Local Area Networks (LANs) running within business environments. Finally, Customer Located Equipment can be used to provide services based on Time Domain Multiplexed (TDM) or Asynchronous Transport Mode (ATM) transport to the customer site, for example to provide services for legacy Private Branch Exchange (PBX) equipment.

All of the terminal variations noted above connect to their respective Access Nodes using one of the last mile transport technologies described in Section 2. The remainder of the Access Network, between the Access Node and the IP Edge, comprises the second mile and middle mile

⁵ This figure is identical to Figure 2 in TR-144 from the Broadband Forum, which develops multi-service broadband packet networking specifications addressing interoperability, architecture and management. Broadband Forum, “TR-144: Broadband Multi-Service Architecture & Framework Requirements,” August 2007. See generally <http://www.broadband-forum.org/about/mission.php>.

networks over which traffic from service providers is forwarded to the appropriate Access Node and traffic from customers is aggregated for delivery to the Regional Network.

The Regional Network provides the connection between multiple Network Service Providers (NSPs) and Application Service Providers (ASPs) and the Access Network. NSPs include network operators in adjoining regions with whom the Access Network provider may have a peering agreement, as well as backbone network operators transporting Internet traffic and other data. ASPs provide application services such as IPTV to the access network. In the Regional Network, traffic from multiple Service Providers is aggregated for delivery to the Access Network and traffic from the Access Network is forwarded to the destination Service Provider.⁶

At each of the aggregation points in Figure 7, there is the potential for congestion resulting in queuing delay and/or packet loss. This delay can occur in either direction. For instance, in the upstream direction (from the customer into the access network), momentary congestion can occur where traffic from multiple customers is aggregated into a shared link with bandwidth lower than the sum total of each of the customer links. In the downstream direction, congestion can occur at the same node if the traffic for a single customer (which may have been transported across high speed middle mile and second mile links) momentarily exceeds the bandwidth of a lower speed last mile connection.

Many types of network services can coexist on the Regional and Access Networks shown in Figure 7. These include, but are not limited to:

- High Speed Internet Access (HSIA) providing service between customers and Internet Gateways.
- IPTV services providing multiple streams of broadcast quality video at both High Definition (HD) and Standard Definition (SD).
- Voice services ranging from the equivalent of legacy telephone service to services incorporating advanced VoIP features.
- Carrier-grade services for business customers providing connections between customer sites and between customers and other networks.
- Interconnection of legacy TDM circuits which may be carried over the packet network using “circuit emulation” technology.
- Second mile and/or middle mile transport of data and voice traffic for wireless network providers (also known as “mobile backhaul”).

Each of the services listed above provides value distinct from the other services, and each service is (or at least can be) offered separately from the others.⁷ Each service also requires appropriate network management in order to ensure the performance attributes necessary for successful delivery of that service. For example:

- Voice services require low latency, jitter, and packet loss.

⁶ In the simplified network diagrams shown in Section 2, the Regional Network is contained within the middle mile network.

⁷ Some of the services are also frequently bundled together in packaged offerings.

- Broadcast quality video requires low latency and jitter and very low packet loss, as well as guaranteed bandwidth at rates that may exceed 10 Mbps for each high definition video stream.
- Voice or data TDM services may require extremely low latency, jitter, and packet loss.
- Carrier-grade business services require traffic delivery compliant with specifications for delay, jitter, packet loss, and other parameters contained in a Service Level Agreement (SLA).
- Mobile backhaul services for wireless providers may require extremely low levels of latency and jitter, as well as timing synchronization between the endpoints connected by the service. With the introduction of femtocells designed to provide cellular wireless access to an area as small as a single household, versions of this service may be carried across residential as well as business access connections.

Of the services listed above, only HSIA provides connections to the public Internet. Each of the other services is generally carried over a managed network designed to provide the performance required for that service. In converged multi-service networks, the managed network is either wholly or partially “virtual,” in that it is carried over the same physical facilities as other services. For instance, residential voice and IPTV services are carried over the same last mile facilities as HSIA, yet they must be managed separately to provide the performance characteristics that each service requires. For example:

- If voice traffic is not given priority over HSIA traffic, short bursts of HSIA traffic can cause momentary “dropouts,” or gaps in place of speech, in telephone calls. Longer bursts of HSIA traffic can cause extended dropouts, unintelligible voice quality, or even dropped calls. Conversely, voice traffic – which is transmitted at regular, predictable intervals in small packets – can be given priority over HSIA traffic with little or no observable effect on HSIA performance.
- When IPTV and HSIA services are provided over the same last mile connection, the connection is provisioned with enough total bandwidth to provide both services concurrently. However, without network management that differentiates between the two services and limits each service to its respective bandwidth allocation, large bursts of HSIA traffic can use bandwidth intended for IPTV, causing degradation of picture quality and/or loss of video at all televisions in use in the household. By the same token, network management that limits the number of concurrent IPTV streams prevents both loss of video quality (due to total video bandwidth exceeding the provisioned bandwidth) and degradation of HSIA performance.

While the above examples are specific to residential access, the concepts apply to all services carried over common facilities. Appropriate network management is necessary to allow coexistence of multiple services while preserving the performance characteristics necessary for each. For example, network services provided to business are often measured against Service Level Agreements (SLAs) that place strict limits on the amounts of delay, jitter and loss allowed. The businesses purchasing those services may rely on adherence to the SLAs to support critical applications, such as the timely and reliable replication of data across multiple business sites to prevent data loss in the event of a disaster. Without network management, network providers have no way to satisfy the SLAs required by their customers.

4 Throughput Issues

Different types of last mile wireline connections have different throughput and performance characteristics. For example, DSL connections can be slower than either second mile connections or home (or business) Local Area Network (LAN) connections, so packets forwarded to those connections can be subjected to queuing delay or loss as well as relatively high serialization delay. In contrast, GPON and EPON networks have very low serialization delay because they run at Gigabit per second speeds. However, GPON and EPON can add jitter in the upstream direction due to the Time Division Multiple Access (TDMA) protocol used to prevent data from different sources from colliding. Also, because the bandwidth is shared, the effective data rate available to a single subscriber on GPON and EPON connections can be lower than the respective rates in either the middle mile or the customer premises, causing possible queuing delay or packet loss in either direction. Currently, this is more of an issue with business customers who tend to run LANs at higher speeds than residential customers, but as Gigabit Ethernet connections become common in home networks it will be seen in residential connections as well.

In the second and middle miles, the range of bandwidths needed varies widely, particularly for second mile facilities where the number and types of customers served spans a wide range. A large Access Node serving a combination of business and residential customers with multiple services, including IPTV, could easily require 1 Gbps or more in the second mile. In contrast, a small Access Node serving a few dozen residential customers with basic HSIA service may require less than 10 Mbps total bandwidth to serve their current needs. This bandwidth may be provided by bonded DSL or even Inverse Multiplexed ATM (IMA) T1 facilities. In addition to possible congestion at the aggregation nodes, as documented above, these lower speed facilities are subject to the same issues associated with serialization delay that affect last mile DSL facilities.

As the trend continues towards increasing broadband speeds and higher performance requirements for a broad array of applications, advances in wireline technologies and architectures which enable higher speeds will see increased deployment. These advances include the following:

- **Bonding.** The bandwidth available on a single DSL loop is limited and dependent on the loop's length. However, many residences and businesses have access to multiple subscriber loops, in part due to the decrease in landline telephone usage. By providing DSL transport on each of those loops and bonding the traffic at each end, the bandwidth delivered to a subscriber can be effectively multiplied by the number of loops available.
- **Advanced spectrum management techniques.** A common limiting factor in DSL performance is crosstalk from other DSL transmitters in the same cable. By coordinating transmission power and spectrum usage between multiple transmitters, the performance of all of the DSL connections in the cable can be optimized.
- **Vectoring.** Traditionally, the crosstalk noted above as a limiting factor in DSL performance is treated as unwanted noise. In fact, the "noise" contains information from other transmitters. By applying advanced Multiple Input Multiple Output (MIMO) signal processing techniques, the loop-to-loop transmission paths inherent in crosstalk can be

exploited for their information value and the effective bandwidth on the cable as a whole can be increased significantly.

- ***Fiber to the Node (FTTN) and Fiber to the Home (FTTH).*** Wireline access networks continue to push fiber closer to the customer premises, shortening the copper loop and increasing the speed at which service can be delivered. FTTN pushes fiber to access nodes within a few thousand feet of the customer and uses VDSL2 to bridge the remaining copper loop. FTTH pushes fiber all the way to the customer premises (usually using GPON or EPON), eliminating the copper loop altogether.
- ***Next Generation GPON and EPON.*** Next generation standards for GPON and EPON are being formalized in ITU and IEEE which will increase downstream rates for PON networks to 10 Gbps in the near future.

5 Network Management

In the context of the proposed rulemaking, Network Management refers to a set of policies, techniques, and tools used to manage the different types of traffic flowing across an access network. We discuss those tools below.

5.1 Network Management Tools

The basic tools used to manage network traffic fall into four broad categories. The first category includes tools for classifying traffic. The second includes methods to ensure that the amount of traffic entering and exiting the network is consistent with contracted levels. The third set of tools prioritizes and schedules traffic within the network, based on classifications identified using the first category of tools. The fourth category includes traffic filtering and other techniques related to network security and blocking of illegal, harmful or objectionable content.

In most cases, before other network management tools are applied to traffic, it must be classified. “Classification” involves identifying packets by one or more relevant characteristics. For example:

- Specific fields in the packet’s headers may indicate a virtual network (and by association, the customer, service or application) with which a packet is associated. The optional VLAN tag in an Ethernet header is frequently used for this purpose.
- Other header fields are used to identify the “Class of Service” (CoS) associated with a packet. The Priority Code Point in Ethernet VLAN tags and the DiffServ Code Point in IP headers are examples of these fields. CoS labels are frequently used in managed networks to identify different types of traffic (for example, VoIP traffic and “Best Effort” traffic) that require different performance attributes.
- Packets can also be classified by the source address header field, which can identify traffic by customer.
- The above examples all use information contained within either “Link Layer” or “Network Layer” packet headers to classify traffic. One or both of these headers are always present and network switches and routers use them for classification and forwarding of packets. Other, more advanced methods for classifying packets are based on information contained in higher layer protocol headers (which are embedded deeper

within the packet structure) or even within the packet payload, and are known collectively as “Deep Packet Inspection” (DPI) techniques.

Network service providers create contracts (with customers and with each other) which, among other things, specify the amount of traffic they will carry. Depending on the contract, this specification may take the form of a guaranteed bit rate, a maximum rate that is not guaranteed, a volume of traffic measured over some time period, or a more complex specification that includes any or all of the above elements. This specification provides network engineers the information they need to design and configure their networks to support the contracted services. The network design will only support those services, however, if all traffic entering the network is consistent with the contracted values. Otherwise, unexpectedly high traffic volume from just one customer can create unacceptable levels of congestion that affect performance for all customers. To prevent this from happening, network engineers use the following tools:

- Traffic “policers” generally operate at network ingress points to limit the rate at which traffic enters the network. Policers can operate on all of the traffic entering the network through a specific point, or they can focus on specific packets based on their classification. For example, multiple policers can be configured at the same physical point to police traffic classified by service or by virtual network. If packets arrive at the policer at too high a rate, the policer can discard some of them so that the remaining packets that get forwarded conform to the contracted terms. A policer can also be configured to allow limited excursions above the contracted rate, or even to mark excess packets so that they can enter the network, but are the first to be discarded if congestion is encountered later on.
- “Shapers” also limit the rate at which traffic is transmitted, but instead of discarding packets (as policers do), they store them in a queue until it is time to send them. This has the effect of “smoothing out” bursts of traffic that might otherwise momentarily exceed a given rate. Shapers are frequently used on traffic just before it is delivered to another network, to make sure that the traffic being delivered to the other network conforms to the contracted rate or other terms of the contract.
- Some services make use of higher layer admission control techniques to prevent a new traffic stream from being initiated on a network that is already operating at capacity. For example, a Voice over Internet Protocol (VoIP) service may use Resource Reservation Protocol (RSVP) or a similar method to confirm that the necessary resources are available before allowing a new call to be set up across the network.

Once packets have been admitted into a network, they encounter many potential points of congestion. In a multi-service network of the type described in Section 3, traffic associated with different services, which may have very different performance requirements, can arrive at the same aggregation node at the same time. Some of that traffic may be highly sensitive to delay and jitter, such as VoIP. Other traffic may be associated with a business service that is subject to stringent contractual delivery requirements. Yet other traffic may require a guaranteed high bandwidth, such as IPTV video. Still other traffic may be tolerant of both delay and jitter, such as web browsing traffic on its way to or from the Internet. How can packets with such different requirements be forwarded through the same network connection while preserving the performance required by each type?

One answer is illustrated in Figure 8. Packets arriving at a point of congestion are classified and placed in separate, parallel queues based on their classification. In the example, there are separate queues for VoIP, IPTV, Defined SLA, and Best Effort traffic. Each queue operates on a First In, First Out basis, but each time a new packet can be transmitted, a scheduler determines the queue from which the packet is selected. Schedulers can use different types of algorithms in combination with each other:

- **Strict Priority**, in which traffic from a queue with higher priority will always be transmitted before traffic from a queue with lower priority. In the example, Strict Priority might be used to transmit VoIP packets whenever they arrive, regardless of the packets waiting for transmission in other queues. VoIP can be a good candidate for Strict Priority scheduling in some cases because a) VoIP traffic is highly sensitive to delay and jitter, b) VoIP packets are small, and c) the packets from any one VoIP flow are sent at regular intervals, with time for other traffic in between them, so allowing them to “jump to the front of the line” during momentary periods of congestion has little impact on other traffic.⁸
- **Round Robin**, in which each queue is served in its turn. This can limit the delay caused to other classes by a large burst of packets in one class by interleaving the transmitted traffic with packets from all classes.
- **Weighted algorithms**, such as Deficit Round Robin. These algorithms select queues in proportion to the weight assigned to each queue. This allows bandwidth to be assigned proportionally as needed to different classes, which can, for example, guarantee the IPTV class the bandwidth it needs.

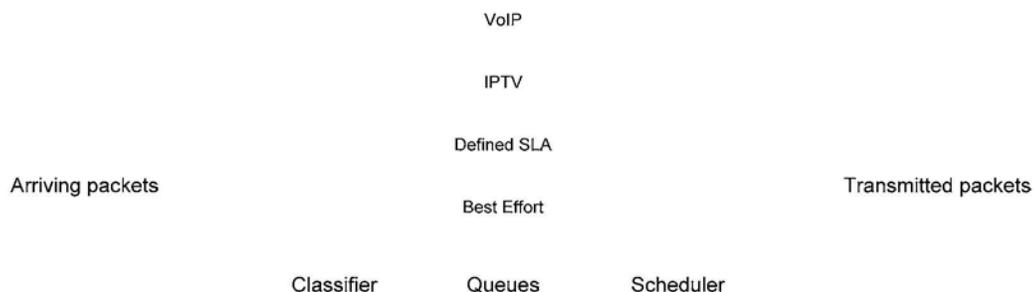


Figure 8 – Classification and Scheduling

The final category of network management tools includes access control lists and filters, used to maintain separation between different parts of the network, and to block attacks and unwanted content:

- **Access control lists**, which can be configured throughout the network, can block traffic from being forwarded to unintended or unauthorized destinations in the network. This is a critical feature in multi-service networks, as well as any access network where traffic from many different customers is carried on a common infrastructure.
- Other, more specific functions are implemented at key points in the network. For example, email may be filtered for spam, malware, and viruses before delivery. DPI

⁸ Whenever Strict Priority scheduling is used, the network management design must ensure that the bandwidth used by higher priority traffic is not so large that it “starves” lower priority traffic.

capabilities operating at wire speed enable filtering for malware, viruses and many types of malicious network attacks. As these capabilities become more prevalent, they can be distributed further throughout the network, providing much greater degrees of protection to customers against infection and network outages.

5.2 Network Management for Managed Services

In Section 3, we noted some of the different types of managed services provided by broadband access providers in addition to High Speed Internet Access over a converged multi-service network. We also noted that these services frequently have specific performance requirements relative to the applications being delivered over the services. To reiterate:

- Voice services require low latency, jitter, and packet loss.
- Broadcast quality video requires low latency and jitter and very low packet loss, as well as guaranteed bandwidth at rates that may exceed 10 Mbps for each high definition video stream.
- TDM services may require extremely low latency, jitter, and packet loss.
- Carrier-grade business services require traffic delivery compliant with specifications for delay, jitter, packet loss, and other parameters contained in a Service Level Agreement (SLA).
- Mobile backhaul services for wireless providers may require extremely low levels of latency and jitter, as well as timing synchronization between the endpoints connected by the service.

Note that in many cases, the applications supported by the services listed above extend beyond the scope of the broadband access provider's network. Voice services are a classic example of this, as voice users are conditioned by long experience with the Public Switched Telephone Network to expect that they can place calls to virtually anywhere in the world. The performance requirements for these applications generally extend from end to end across the full network path, rather than being defined as within the narrow scope of a broadband access provider's network.

In order to meet end-to-end performance requirements that exceed the scope of their own networks, NSPs contract with each other as necessary to carry traffic such that it meets defined SLA requirements. Such traffic is generally maintained separate from the Internet, which (for the present) is limited to Best Effort delivery of traffic.

5.3 Network Management for Internet Applications

As noted above, at the present time traffic on the Internet is limited to only one Class of Service, known as Best Effort delivery. As the name implies, Best Effort delivery provides no guarantees with regard to delay, jitter, or packet loss. This can constrain the performance of applications that rely on transport across the Internet.

The Internet is constantly evolving, however, with many thousands of people involved in continuously innovating and improving the protocols, equipment, and standards used. What is true today regarding network management practices on the Internet may not be true tomorrow. A combination of existing tools and new techniques, adopted by a critical mass of equipment

vendors and NSPs, could enable tremendous improvement in the performance of applications such as high quality video conferencing and telepresence. New features and applications could be enabled that are not currently under consideration or even possible. Innovative network management techniques will be especially important given the rapid exponential growth of traffic across the Internet,⁹ which will only exacerbate congestion on backbone networks over time.

Even without introducing new protocols or capabilities on the Internet itself, customers can realize significant value from network management practices offered by their broadband access network providers on those providers' own networks. Some practices are in common usage today, and others are still in development. Still others have yet to be envisioned. An incomplete list includes the following:

- Preventing spam, spyware and other malware from being downloaded to the customer's home.
- Allowing parents to control their children's access to objectionable content.
- Allowing customers to choose how to prioritize traffic to and from their home. For instance, allowing customers to specify that traffic to and from an interactive gaming site be given priority over email traffic.
- Automatically managing HSIA traffic by application class, to align the performance of each class with its requirements. For example, a customer's video conferencing traffic, which is sensitive to delay and jitter, could be given priority over that same customer's web browsing or peer to peer traffic, which is more tolerant of such impairments.

6 Potential Implications of Proposed Rulemaking

The four existing Principles and the two proposed Principles regarding broadband industry practices provide valuable guidance to the broadband community. To the degree that these Principles are codified as rules with the power of enforcement, they will be subjected to considerable interpretation during (and after) the rulemaking process. As we have explained in the preceding material, there is considerable complexity involved in managing a broadband access network and there are many potential opportunities for the rules to result in unintended consequences that would be in conflict with the general principles they were trying to ensure. Some examples of such consequences are provided below.

Principle 1: Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from *sending or receiving the lawful content of the user's choice* over the Internet.

Service providers can add value and protect consumer interests by providing filters that can block certain types of content. An overly broad codification of Principle 1 could prevent some or all of these practices. Examples include:

⁹ See Cisco, "Cisco Visual Networking Index – Forecast and Methodology, 2008-2013," available at http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html

- Spam, adware and spyware, some of which may be arguably legal but objectionable to the vast majority of recipients. A network should not be prohibited, for example, from blocking spam sent willfully by a user.
- Other objectionable content that parents may wish to block from entering the household.
- Files that are infected by viruses or worms. These files may be otherwise desirable and may be infected without the sender's or receiver's knowledge. An access provider should not be prohibited from blocking a file, even one specifically requested by a user, if that file contains an infection that could cause harm to the user's equipment and spread to other users in the network.

Principle 2: Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from ***running the lawful applications or using the lawful services of the user's choice***.

The codification of Principle 2 could be understood to require that all applications must be guaranteed optimal or even basic performance, regardless of how stringent their requirements or the level of broadband service purchased by the user. Consider as one example very high quality interactive video conferencing, or telepresence, established over an HSIA connection. This application may require bit rates that cannot be sustained on some HSIA services, especially in the absence of application-oriented network management. The resulting performance may be unacceptable or even render the application completely unusable. This is not, however, due to any blocking, or other action or negligence, on the part of the broadband access provider. Instead, it is the result of a basic mismatch between the requirements of the application and the capabilities of the access service. The situation is analogous to a software application which will not run on computer systems that do not meet the minimum system requirements for the application.

Principle 3: Subject to reasonable network management, a provider of broadband Internet access service may not prevent any of its users from ***connecting to and using on its network the user's choice of lawful devices that do not harm the network***.

“Harm to the network,” as first defined in FCC Part 68, focused on requirements such as transmit power, transmitted spectrum, return loss, and other physical layer attributes associated with the local connection. Some of these requirements, such as limitations on transmitted spectrum, are intended to prevent one device from unfairly limiting the performance of other devices on the same access network. Advances in transmission technology require a broader definition of “harm” to fulfill the same intent. In addition, the ongoing epidemic of malware highlights the fact that “harm to the network” can occur well outside the narrowly defined parameters of the original definition.

Given the dynamic nature of both technology advances and malicious activity, some latitude is required to determine what constitutes “harm to the network” in order to be able to maximize performance and security for all users. For example:

- Some DSL modems located at the customer premises may comply with basic network requirements, but may not support advanced features required to maximize performance for all users on the network. In many cases, these features must be

implemented by all devices on a network in order for the performance advantage to be realized. For instance, vectoring (discussed above in Part 4) requires that the modem support an Embedded Operations Channel (EOC) which sends channel information back to the DSL Access Multiplexer (DSLAM). A modem that doesn't support this feature would limit the performance of all other modems served by subscriber loops in the same physical cable, thus causing "harm" to the network.

- A device can be several devices removed from the access provider's network and still cause harm. For instance, a host computer within a subscriber's network may stream traffic out into the network due to a hardware or software malfunction. Or, the host may be infected by a virus and attempt to infect other computers. Or, due to a malware infection, the host may be an inadvertent participant in a "botnet" initiating a Denial of Service or other attack on the network. In any of these cases, the access provider must be able to remove the harmful device's access to the network until the issue is resolved. The only feasible way to do so may be to disconnect the subscriber's service at the customer's point of interconnection to the access network.

Principle 4: Subject to reasonable network management, a provider of broadband Internet access service may not deprive any of its users of the user's entitlement to *competition among network providers, application providers, service providers, and content providers*.

As noted below in the discussion of Principle 5, there are network management techniques (many as yet unforeseen) that can add considerable value for the customers being served by the access network. As noted in Section 5.3, access network providers may allow customers to choose how to prioritize traffic to and from their home. Access network providers may also optimize performance for traffic flows based on the requirements associated with their respective application classes. While these features make use of prioritization and other network management techniques, they are not anti-competitive – in fact, by enabling enhanced performance, they promote innovation among all application and content providers.

These features and the benefit derived from them may go unrealized if an overly broad interpretation of Principle 4 has a chilling effect on innovation. As just one example, there will always be a steady stream of new applications and services being introduced into the market. It will be impossible for access providers to be aware of every new product in real time as it is introduced, much less create the support necessary to classify flows associated with that product or make it available for customer prioritization. If an access provider is subject to claims of discrimination and legal action for not providing enhanced support for each new application – no matter how new, no matter how small the application's customer base, no matter how little benefit may be derived from such support – then the safest course of action will be to treat all services and applications equally, meaning that none will be supported, innovation will be stifled, and no benefits will accrue to the customers.

The same argument applies to many services that access providers currently offer over managed networks. Significant system design (as well as investment) may go into providing the level of performance required by those services. Some similar services (such as VoIP and video services) are also offered by third party Internet sources, delivered via the customer's HSIA service. While such "over the top" services using

standard Best Effort HSIA may compete on other factors, they generally cannot guarantee the same performance characteristics provided via a managed network.

An overly broad interpretation of Principle 4 could create confusion as to whether facilities-based providers must optimize their networks for third-party applications, even when such optimization might degrade the services they provide themselves. Such interpretation could also deter investment in system design and managed services by over-the-top providers, who might expect the *network* providers to compensate for weaknesses in the applications providers' own offerings. This result would remove incentive for access providers to invest in their infrastructure. The customer would ultimately lose as the performance level for all services would be reduced to the lowest common denominator.

Principle 5: Subject to reasonable network management, a provider of broadband Internet access service must *treat lawful content, applications, and services in a nondiscriminatory manner*.

It is important the network service providers not be required to guarantee equivalent performance for all applications. Different applications have different characteristics in terms of traffic volume, throughput requirements, and latency and jitter requirements. For instance, Peer-to-Peer (P2P) applications are frequently used to transfer extremely high amounts of data. This one application may account for the majority of consumer traffic in the upstream direction, and that volume may be concentrated within a small percentage of the subscriber population. If fair network management practices limit the volume of traffic on a per-subscriber basis to ensure that all subscribers have equal access to the Internet, P2P traffic may be affected on average more than traffic from other applications. This example is not “discriminatory” from an equity or anticompetitive perspective, however. Instead, bandwidth is being allocated fairly to subscribers based on the services to which they have subscribed. For instance, subscribers who are generating high traffic volumes by uploading files to a web server will be affected equally to those who are using P2P applications – however, both of the above subscribers are likely to be affected more than a subscriber checking email or browsing the web, because they are both attempting to use a disproportionately large share of the network bandwidth and the light user is not.

Likewise, network providers must not be required to guarantee all applications *optimal* performance. Consider interactive video conferencing over HSIA. It may be possible for the user to configure this application for rates that cannot be sustained reliably on the network, especially in the absence of application-oriented network management. The resulting performance may be sub-optimal or even unusable – however, this is not a result of “discrimination.”

Conversely, an overly broad interpretation of Principle 5 could prevent a service provider from offering value-added features that customers would desire and would freely choose to accept. Customers may want traffic to and from their interactive game server – or, to revisit the previous example, their video conferencing provider – to be given priority over their other traffic. Fulfilling that desire adds value for the customer and enriches the broadband experience, but it could be prohibited by a non-discrimination rule.

Such value added features may not require proactive selection by the customer. The data streams for many application classes can be identified in real time, for example by the

protocols they use or by information in the associated packet headers. It is then possible for a service provider to optimize the performance of these streams relative to the specific parameters that are most critical for that application class. Service providers could then improve the application-level performance of traffic on their networks, creating value for their customers and differentiating their HSIA service offerings. However, this too could be prohibited by codification of a nondiscrimination rule.

The proposed nondiscrimination rule could also endanger value added services offered over managed networks, such as voice, IPTV, or carrier class business services. As noted above, different services require significantly different performance attributes. Providers must be allowed to use network management tools to meet the performance requirements (both implicit and explicit) of the services they offer.

Principle 6: Subject to reasonable network management, a provider of broadband Internet access service must *disclose such information concerning network management and other practices as is reasonably required for users and content, application, and service providers to enjoy the protections specified in this part.*

Transparency in the service contracts between users and service providers will bring genuine benefits to the broadband access community. In particular, we note that the actual speeds that users are able to sustain are currently less than half of advertised “up to” speeds on average.¹⁰ Disclosure of sustainable speeds and other contractual details such as limits on daily or monthly usage will help consumers make informed choices and will promote the adoption of broadband across the country.

At the same time, it is critical for the security of all network users that any requirement to disclose information concerning network management practices be tempered by the need to maintain security against the malicious attacks, malware, phishing, spam, and other threats that network providers must fight on a continuous basis. Attackers can and will use any detail provided under regulatory disclosure in their attempts to breach network security measures.

The above issue is exacerbated if codification of Principle 6 requires network providers to provide details of network management practices prior to implementation. Network security is a constantly evolving struggle, as new malware and viruses are introduced at an alarming rate. As an example, any user of modern antivirus software has probably observed that, while in years past antivirus tools downloaded new virus definitions once a week, the cycle time for updates is now measured in *minutes*. Under these rapidly changing conditions, hampering network providers with requirements for pre-notification of detailed changes – or, even worse, requiring pre-*approval* of those changes – would create a situation ripe for disaster.

Even disclosure of more general, supposedly benign details of network management techniques may lead to unintended consequences. Users and applications providers may be able to use network management disclosures to “game the system” for their own benefit, generating performance gains in unintended ways that may be unfair to other stakeholders. Such a result would run counter to the core intent of the Proposed Rulemaking.

¹⁰ FCC, “Broadband Gaps,” update at the November 18, 2009 open meeting.

7 Special Requirements

In addition to the more general issues associated with provisioning and management of services on the access network, access providers need to be able to support special requirements related to law enforcement, public safety, national and homeland security. Two requirements are of particular importance:

- The normal call quality performance requirements associated with interactive voice communications take on new urgency for 911 calls. Dropouts, unintelligible speech, and lost calls which are a nuisance under normal circumstances can contribute to a tragedy when emergency services are involved. Access providers must be able to use the appropriate network management tools to provide the QoS required for E911 services.
- Communications Assistance for Law Enforcement Act (CALEA). Communications providers, including certain Internet access providers, are required to cooperate in the interception of communications for law enforcement purposes. Compliance with these requirements requires that a particular subscriber's voice and data traffic that travels over a public network be identified for interception. Sophisticated DPI capabilities may be beneficial in this process.

8 Investment and Innovation

This section addresses ongoing investment in the wireline network and the ways in which the proposed rules might impact that investment. Wireline access networks are being transformed in a number of ways as customers' requirements and expectations continue to evolve. Some of the ways include:

- Fiber is being pushed further into the network. Many second mile copper facilities are being upgraded to fiber, enhancing the capacity in that portion of the network by orders of magnitude.
- At the same time, the outside plant nodes terminating copper subscriber loops are being pushed closer to the customer premises. In many cases, higher speed VDSL is deployed on the shorter loops, replacing existing ADSL connections. In other cases, fiber is being pushed all the way to the customer premises in the form of GPON or other fiber-based last mile technologies.
- Bonding of multiple subscriber loops is enabling 2x or higher speeds in last mile copper facilities.
- Dynamic Spectrum Management is being deployed in many networks to increase performance on DSL connections.
- Vectoring, which uses advanced signal processing to exploit the information content in crosstalk paths to significantly increase performance, is moving from the lab to practical implementations.

The above last mile and second mile upgrades also enable the deployment of managed services such as IPTV and voice, which require specific performance attributes including (in some cases) high bandwidth, over a rapidly increasing number of wireline networks. Network providers are

also integrating consumer and business services on common platforms to an increasing degree, while broadening the range of services available to customers.

In addition to the above transformation in wireline networks, a transformation in wireless devices and wireless access methods is blurring the lines between wireline and wireless networks. Many wireless devices are multi-mode, meaning that they can communicate over a variety of wireless standards. When mobile, they use traditional cellular wireless access networks for voice and data communications. When they can use a WiFi or other wireless LAN for access, however – for instance, at home or in the office – they use that network as the preferred transmission path. Additionally, an increasing number of homes and other buildings are making use of femtocells – small cellular base stations, provided by the wireless access provider, which operate within a single building and which provide wireless service using the same radio interface as for mobile communications. The data and voice traffic is then “backhauled” from the femtocell’s base station to the wireless access provider’s network over the customer’s wireline access network.

Increasing capabilities for flow classification and network management create the opportunity for wireline access providers to offer new features and added value to customers, both for managed services and for HSIA. Advanced network management tools can incorporate awareness of the performance attributes of different applications classes to optimize the Quality of Experience (QoE) for customers. Customers may be able to go a step further, customizing the QoE for their applications and services to their individual needs.

Potential advances in the capabilities described above for flow classification and network management are not limited to the access network or to managed services. There is enormous potential for innovation in the management and transmission of traffic across the Internet, which is presently limited to Best Effort traffic. Improvements in Internet traffic management can enable reliable end-to-end performance for demanding applications such as high quality consumer video conferencing, and can open the door to entire classes of applications and uses which are simply not feasible on a “one size fits all” Internet.

Network and user security enhancements are also developing at a rapid pace. Continuing advances in filtering and Deep Packet Inspection are required to fight the constant escalation of spam, viruses, malware, and other malicious attacks that threaten both individual users and the network as a whole.

The evolution described above affects all aspects of wireline access networks, from architecture and physical layer technology on up to applications. The investment required to drive this evolution at the current rapid pace of change is considerable. If the proposed rulemaking results in overly burdensome or ambiguous regulations, they could have a fundamental impact on that investment, and on the continued evolution of the access network. If access providers are constrained in their use of “reasonable network management” methods – or if a lack of clarity in the definition of such methods causes uncertainty and opens the door to potential legal action – there will be a disincentive to use such methods. The unintended consequences of such a disincentive could include: lack of advanced QoS and security features; limits on network integration, leading to limits on efficiency and increased costs as traffic demands continue to climb; and stagnation rather than innovation on the Internet itself. At the least, the Commission should consider limiting the potential for such disincentives by identifying a large and non-inclusive list of practices that would always fall under the umbrella of “reasonable network management.” In addition to the practices listed in the current proposed definition of

“reasonable network management,” such practices could include, but not be limited to, the following:

- Management of different services over common facilities, including the policing, scheduling, and prioritization of traffic as appropriate to each service.
- Network management to fairly allocate resources between users during periods of network congestion.
- Quality of Service related network management, including flow classification by application class and network management aligned with the performance attributes of known application classes.
- Blocking of spam, malware, viruses, and other malicious traffic.

Of particular concern is any result of the Proposed Rulemaking that would extend beyond the Internet and affect the managed network facilities and services that are part of the converged multi-service network. Those services exist to deliver performance that cannot be guaranteed over the Internet. In many cases, such as IPTV, broadband access providers can deliver that performance only by creating a design that integrates content, application, network management, and network infrastructure as a complete system. Trying to apply requirements such as those contemplated in this proceeding to such services could stifle investment and innovation in these types of solutions.

Finally, investment and innovation in wireline broadband access infrastructure and technology can be sustained only if any rules adopted are applied fairly to all stakeholders. Neither wireless operators, wireline operators, application providers, service providers, nor any other stakeholders should obtain an artificial advantage from disparate regulatory treatment. All access networks, for example, must deal with the issue of finite capacity, and need to have access to the tools required to meet that challenge.

In addition, because of the ongoing convergence of wireless and wireline networks (via multi-mode devices and femtocells), increasing amounts of wireless access network traffic will actually be carried over wireline access networks. Because of this convergence, any resulting rules that treat different technologies differently could result in the same traffic being subjected to one set of rules in the last mile and a different set of rules in the second mile. Moreover, there may be points in the network where it simply isn't clear which rules apply.

Finally, applications and service providers can sometimes have much more capability than access providers to treat content, applications, and services in a discriminatory manner, and to do so in ways that may not be apparent to the users of those services. Web portals, search engines, browsers, and network caching services are just some examples of applications and services capable of influencing users' access to content, applications and services, and in some instances an access provider, using network management techniques, could be competing against these alternatives. Applying the Proposed Rulemaking only to access providers would fail to provide the protections for users that are the intent of the proposed rulemaking, since a significant source of potential discrimination would be unaffected by the resulting rules.