Comcast Corporation
2001 Pennsylvania Ave., NW
Suite 500
Washington, DC 20006
202.379.7100 Tel
202.466.7718 Fax
www.comcast.com

September 19, 2008

VIA ECFS AND HAND DELIVERY

Ms. Marlene H. Dortch
Secretary
Federal Communications Commission
445 12th Street, S.W.
Washington, D.C. 20554

Re: **In the Matter of Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, File No. EB-08-IH-1518**

In the Matter of Broadband Industry Practices; Petition of Free Press et al. for Declaratory Ruling that Degrading an Internet Application Violates the FCC's Internet Policy Statement and Does Not Meet an Exception for "Reasonable Network Management," WC Docket No. 07-52

Dear Ms. Dortch:

In accordance with the Commission's August 20, 2008 *Memorandum Opinion and Order* regarding Comcast's network management practices for our High-Speed Internet ("HSI") service,[1] Comcast hereby complies with the three filing requirements set forth therein. Specifically, consistent with Paragraphs 54 and 59 of the Commission's *Order*, we submit the following:

(1) a description of our current approach to managing network congestion (Attachment A);

(2) a description of the new protocol-agnostic congestion management practices to which we are transitioning no later than year-end 2008 (Attachment B); and

(3) a compliance plan setting forth the benchmarks that we will meet as part of this transition (Attachment C). We have also included in this document our plans for direct communication with our customers during this transition.

[1] *In re Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications; Broadband Industry Practices; Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC's Internet Policy Statement & Does Not Meet an Exception for "Reasonable Network Management*," Mem. Op. and Order, FCC 08-183 (Aug. 20, 2008) ("*Order*").

These filings are consistent with our previously announced commitment to transition away from the congestion management practices we currently use to prevent peer-to-peer ("P2P") traffic from degrading our customers' use and enjoyment of our HSI service to a new set of protocol-agnostic congestion management practices, and to do so across our network by December 31, 2008. Over the last several months, we have conducted technical trials to determine how best to implement a new protocol-agnostic approach to congestion management. We are making excellent progress and are on track to complete the transition as scheduled. As in everything we do, our goal is to ensure continued delivery of a world-class service to all of our subscribers, while minimizing the impact on any individual users whose traffic must be managed as part of this process.

We continue to refine the details of our new practices, so we commit to make supplementary filings in this docket as necessary to keep the Commission (and the public) informed of any material changes in our plans before we complete the transition to protocol-agnostic congestion management by year-end. Separate and apart from the requirements of the *Order*, we have an ongoing commitment to our customers to provide a world-class Internet experience. To do so, we must always preserve the flexibility to manage our network in lawful and appropriate ways. Moreover, we know that clear communication with our customers is essential to a successful long-term relationship. So we are committed to ensuring that our customers receive clear, concise, and useful information about the services that we provide.

Even as we adopt the new network management practices described in Attachment B, we continue to make the investments in network upgrades that will permit us to better prevent congestion and meet our customers' ever-increasing demands for bandwidth. For example, earlier this year we doubled, and in many cases tripled, the upload speeds for almost all of our existing HSI customers. In addition, since our initial rollout of DOCSIS 3.0 (which currently offers consumers wideband download speeds of up to 50 Mbps and upload speeds of up to 5 Mbps) in the Twin Cities Region in April, we have continued preparations to deploy DOCSIS 3.0 to up to 20 percent of our footprint by the end of this year, and in many more markets in 2009.

As all of the Commissioners recognize, the Internet is an engine for innovation and economic growth. We are proud to be a leader in bringing broadband Internet to consumers all over the country, adding fuel to that engine. We will continue to work hard to deliver a world-class service that gives all of our subscribers access to the content, applications, and services that they demand.

Please contact me should you have any questions regarding this submission.

Sincerely,

/s/ Kathryn A. Zachem
Kathryn A. Zachem
Vice President, Regulatory Affairs
Comcast Corporation

cc: Chairman Kevin J. Martin
Commissioner Michael J. Copps
Commissioner Jonathan S. Adelstein
Commissioner Deborah T. Tate
Commissioner Robert M. McDowell
Daniel Gonzalez
Dana Shaffer
Scott Bergmann
Scott Deutchman
Kris Monteith
Amy Bender
Greg Orlando
Nick Alexander

ATTACHMENT A:

COMCAST CORPORATION
DESCRIPTION OF CURRENT NETWORK MANAGEMENT PRACTICES

COMCAST CORPORATION
DESCRIPTION OF CURRENT NETWORK MANAGEMENT PRACTICES

Pursuant to Paragraphs 54 and 59 of the Commission's *Memorandum Opinion & Order* regarding how Comcast manages congestion on its High-Speed Internet ("HSI") network, Comcast hereby "disclose[s] to the Commission the precise contours of the network management practices at issue here, including what equipment has been utilized, when it began to be employed, when and under what circumstances it has been used, how it has been configured, what protocols have been affected, and where it has been deployed."[1]

I. INTRODUCTION

Comcast's HSI network is a shared network. This means that our HSI customers share upstream and downstream bandwidth with their neighbors. Although the available bandwidth is substantial, so, too, is the demand. Thus, when a relatively small number of customers in a neighborhood place disproportionate demands on network resources, this can cause congestion that degrades their neighbors' Internet experience. In our experience, over the past several years, the primary cause of congestion (particularly in the upstream portion of our network) has been the high-volume consumption of bandwidth associated with use of certain peer-to-peer ("P2P") protocols. In order to tailor our network management efforts to this reality, Comcast's current congestion management practices were designed to address this primary contributor to congestion. Our objective in doing so was to provide all our customers with the best possible broadband Internet experience in the marketplace.

As described in Attachment B, in response to significant stated concerns of the Internet community, Comcast had already announced plans to transition away from its P2P-specific

---

[1] *In re Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications; Broadband Industry Practices; Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC's Internet Policy Statement & Does Not Meet an Exception for "Reasonable Network Management*," Mem. Op. and Order, FCC 08-183 ¶¶ 54, 59 (Aug. 20, 2008) ("*Order*").

congestion management practices and terminate them entirely by December 31, 2008. Paragraph 54 of the *Order* directs Comcast to describe these current practices, and we do so here.[2]

At the outset, we provide some background on how these practices came into being and how they work in a general sense. We then provide the greater detail required by the *Order*.

II. BACKGROUND

To understand exactly how Comcast currently manages congestion on its network, it is helpful to have a general understanding of how Comcast's HSI network is designed.[3] Comcast's HSI network is what is commonly referred to as a hybrid fiber-coax network, with coaxial cable connecting each subscriber's cable modem to an Optical Node, and fiber optic cables connecting the Optical Node, through distribution hubs, to the Cable Modem Termination System ("CMTS"), which is also known as the "data node." The CMTSes are then connected to higher-level routers, which in turn are connected to Comcast's Internet backbone facilities. Today, Comcast has approximately 3300 CMTSes deployed throughout our network, serving our 14.4 million HSI subscribers.

Each CMTS has multiple "ports" that handle traffic coming into and leaving the CMTS. In particular, each cable modem deployed on the Comcast HSI network is connected to the CMTS through the "ports" on the CMTS. These ports can be either "downstream" ports or "upstream" ports, depending on whether they send information to cable modems (downstream) or receive information from cable modems (upstream) attached to the port. Today, on average,

---

[2] Although the *Order* focuses entirely on Comcast's current practices with respect to controlling network congestion, Comcast's efforts to deliver a superior Internet experience involve a wide variety of other network management efforts beyond congestion control. As Comcast has previously explained, we actively manage our HSI network in order to enhance our customers' Internet experience by, among other things, blocking spam, preventing viruses from harming the network and our subscribers, thwarting denial-of-service attacks, and empowering our customers' ability to control the content that enters their homes.

[3] The reader may find it useful to refer to the attached glossary for additional explanation of unfamiliar terms.

about 275 cable modems share the same downstream port and about 100 cable modems share the same upstream port. As will be described later in this document, Comcast's current congestion management practices focus solely on a subset of *upstream* traffic.

Internet usage patterns are dynamic and change constantly over time. As broadband networks deliver higher speeds, this enables the deployment of new content, applications, and services, which in turn leads more and more households to discover the benefits of broadband Internet services. Several years ago, Comcast became aware of a growing problem of congestion on its HSI network, as traffic volumes, particularly for upstream bandwidth (which is provisioned in lesser quantities than downstream bandwidth[4]), were growing rapidly and affecting the use of various applications and services that are particularly sensitive to latency (i.e., packets arriving slowly) or jitter (i.e., packets arriving with variable delay).

In order to diagnose the cause of the congestion and explore means to alleviate it, in May 2005, Comcast began trialing network management technology developed by Sandvine, Inc. The Sandvine technology identified which protocols were generating the most traffic and where in the network the congestion was occurring. After jointly reviewing significant amounts of usage data, Comcast and Sandvine determined that the use of several P2P protocols was regularly generating disproportionate burdens on the network, primarily on the upstream portion of the network, causing congestion that was affecting other users on the network.

As previously explained on the record and described in greater detail below, in order to mitigate congestion, Comcast determined that it should manage *only* those protocols that placed

[4] This asymmetric provisioning of bandwidth is based on how the vast majority of consumers have historically used the Internet, i.e., most consumers have been far more interested in how fast they could surf the web, how fast they could download files, and whether they could watch streaming video than in uploading large files. Even today, with the widespread proliferation of services that place greater demand on upstream resources, most consumers still download much more than they upload, and so we continue to architect our network to optimize the experience of the vast majority of our users. As usage patterns change over time, so, too, will our provisioning practices.

excessive burdens on the network, and that it should manage those protocols in a minimally intrusive way utilizing the technology available at the time. More specifically, in an effort to avoid upstream congestion, Comcast established thresholds for the number of simultaneous unidirectional uploads that can be initiated for each of the managed protocols in any given geographic area; when the number of simultaneous sessions remains below those thresholds, uploads are not managed. The thresholds for each protocol vary depending upon a number of factors discussed in detail below, including how the particular protocol operates and the burden that the particular protocol was determined to place on our upstream bandwidth. These management practices were not based on the type (video, music, data, etc.) or content of traffic being uploaded.

The Sandvine equipment has been used (1) to determine when the number of simultaneous unidirectional upload sessions for a particular P2P protocol in a particular geographic area reaches its pre-determined threshold, and (2) when a threshold is reached, to temporarily delay the initiation of any new unidirectional upload sessions for that protocol until the number of simultaneous unidirectional upload sessions drops below that threshold.

III. WHAT EQUIPMENT IS UTILIZED?

The specific equipment Comcast uses to effectuate its network management practices is a device known as the Sandvine Policy Traffic Switch 8210 ("Sandvine PTS 8210"). Literature describing this product is attached. The following sections explain where and how Comcast uses the Sandvine PTS 8210.

IV. WHERE HAS THE EQUIPMENT BEEN DEPLOYED AND WHEN AND UNDER WHAT CIRCUMSTANCES HAS IT BEEN USED?

Comcast initially began technical trials with the Sandvine PTS 8210s starting in May 2005. Commercial (i.e., not trial) deployment of this equipment took place over an extended period of time, beginning in 2006. We achieved wide-scale deployment in 2007.[5]

On Comcast's network, the Sandvine PTS 8210 is deployed "out-of-line" (that is, out of the regular traffic flow)[6] and is located adjacent to the CMTS. Upstream traffic from cable modems will pass through the CMTS on its way to upstream routers, and then, depending on the traffic's ultimate destination, onto Comcast's Internet backbone. A "mirror" replicates the traffic flow that is heading upstream from the CMTS without otherwise delaying it and sends it to the Sandvine PTS 8210, where the protocols in the traffic flow are identified and the congestion management policy is applied in the manner described in greater detail below. In some circumstances, two small CMTSes located near each other may be managed by a single Sandvine PTS 8210.[7] The following graphics provide a simplified illustration of these two configurations:

---

[5] Some locations currently have a network design that is different from the standard Comcast network design because we are trialing new protocol-agnostic congestion management practices in those locations, we are preparing those locations for evolution to DOCSIS 3.0 (which has already been launched in one market), or we acquired those systems from other operators and are in the process of standardizing them. The congestion management practices described herein are not used in those systems. The locations of our trials have been widely publicized, but disclosure of proprietary plans regarding the order and timing for network investments and service upgrades would cause substantial competitive harm.

[6] Comcast deploys the Sandvine PTS 8210 "out-of-line" so as to not create an additional potential "point-of-failure" (i.e., a point in the network where the failure of a piece of equipment would cause the network to cease operating properly). The Sandvine equipment can also be deployed "in-line," which can make the management effectuated by the equipment nearly undetectable, but Comcast does not employ this configuration.

[7] Although the PTS generally monitors traffic and effectuates policy at the CMTS level, the session management interface is administered at the Upstream Router, one layer higher in the overall architecture.

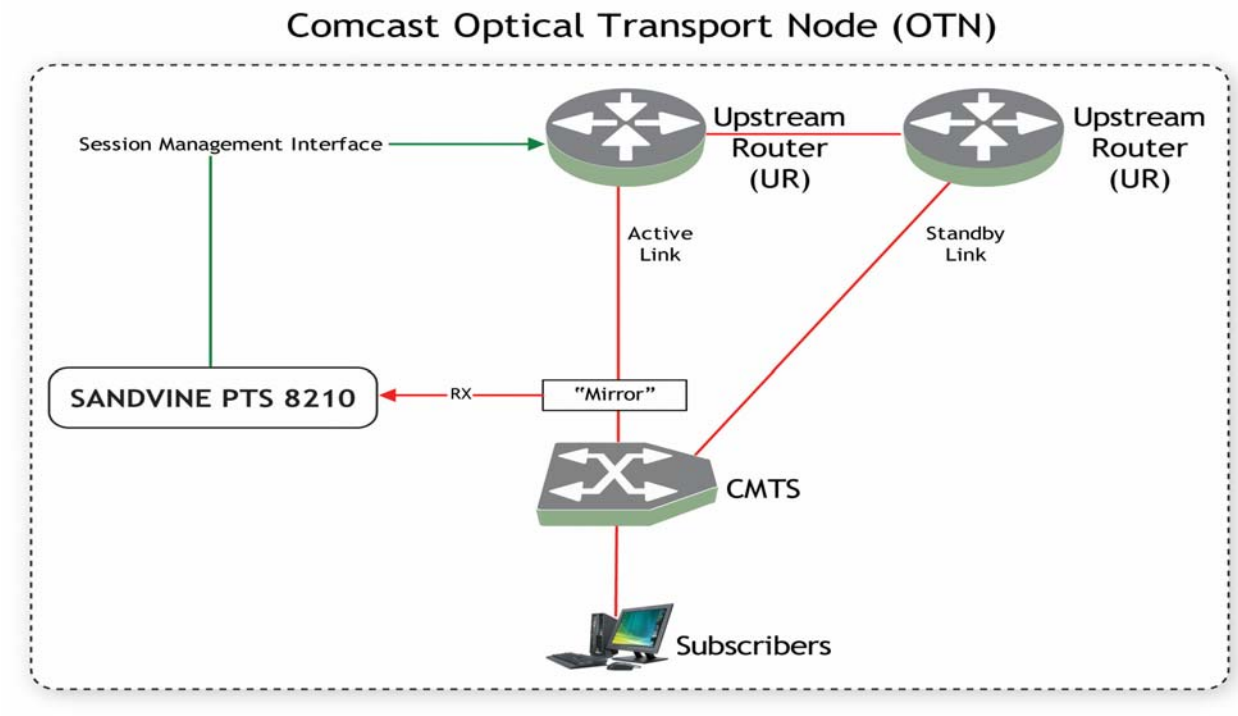

Diagram 1: Sandvine PTS Serving One CMTS.

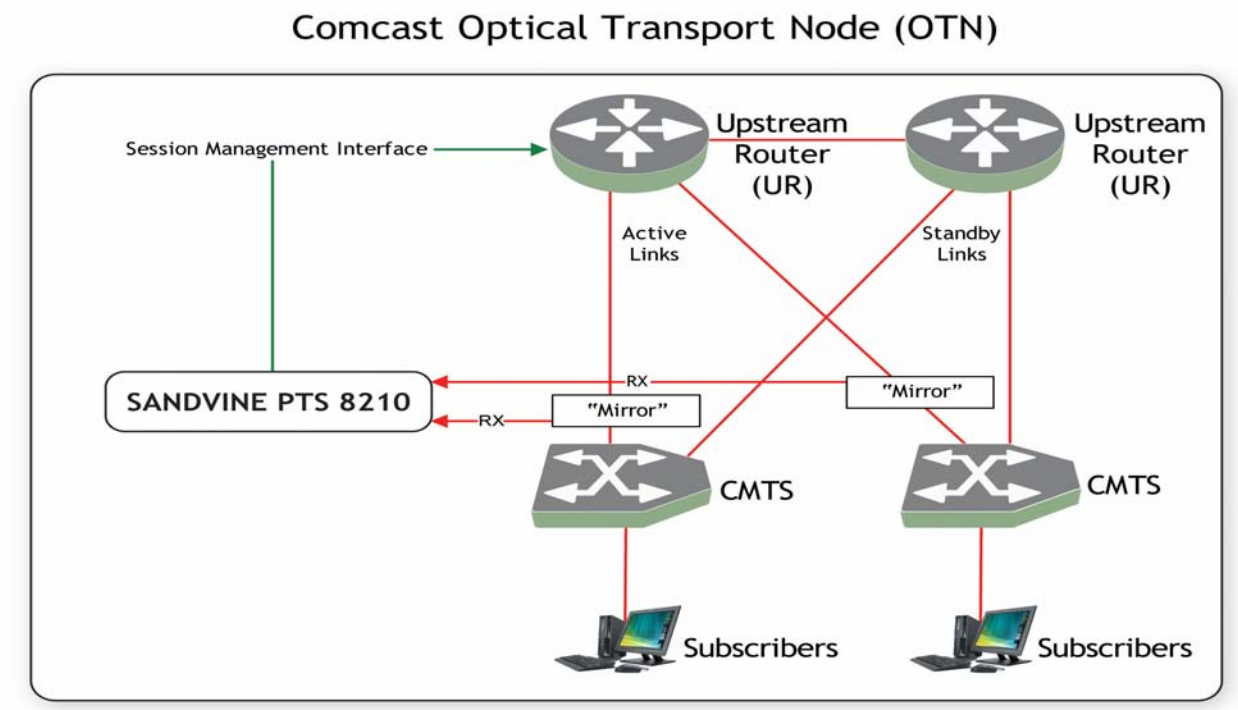

Diagram 2: Sandvine PTS Serving Two CMTSes.

V. HOW HAS THE EQUIPMENT BEEN CONFIGURED AND WHAT PROTOCOLS HAVE BEEN AFFECTED?

For purposes of managing network congestion,[8] the Sandvine PTS 8210 has been configured to identify unidirectional P2P uploads for the protocols -- identified below -- that were determined to be the primary causes of upstream congestion.[9] To do this, the Sandvine PTS uses technology that processes the addressing, protocol, and header information of a particular packet to determine the session type. The Sandvine PTSes, as deployed on Comcast's network, *do not inspect the content*. These devices only examine the relevant header information in the packet that indicates what type of protocol is being used (i.e., P2P, VoIP, e-mail, etc.). The equipment used does *not* read the contents of the message in order to determine whether the P2P packet is text, music, or video; listen to what is said in a VoIP packet; read the text of an e-mail packet; identify whether any packet contains political speech, commercial speech, or entertainment; or try to discern whether packets are personal or business, legal or illicit, etc.

The following diagram graphically depicts the session identification technique undertaken by the Sandvine PTS 8210 as deployed on Comcast's network. The first layers include addressing, protocol, and other "header" information that tells the network equipment what kind of packet it is. The "content" layer is the actual web page, music file, picture, video, etc., and is not examined by the Sandvine equipment.

---

[8] The Sandvine PTS 8210 has not been used solely to manage congestion. It also performs numerous functions related to network management and security, including traffic analysis, anti-spam measures, denial-of-service attack prevention, and other similar functions.

[9] A "unidirectional upload" session is different from an upload associated with a "bidirectional upload" session. A session is considered bidirectional when the user is simultaneously uploading to *and* downloading from another individual using a single TCP flow. Two of the protocols that are managed, BitTorrent and eDonkey, use bidirectional sessions; the other protocols only use unidirectional sessions. A large percentage of P2P traffic is bidirectional and is not managed by these techniques.

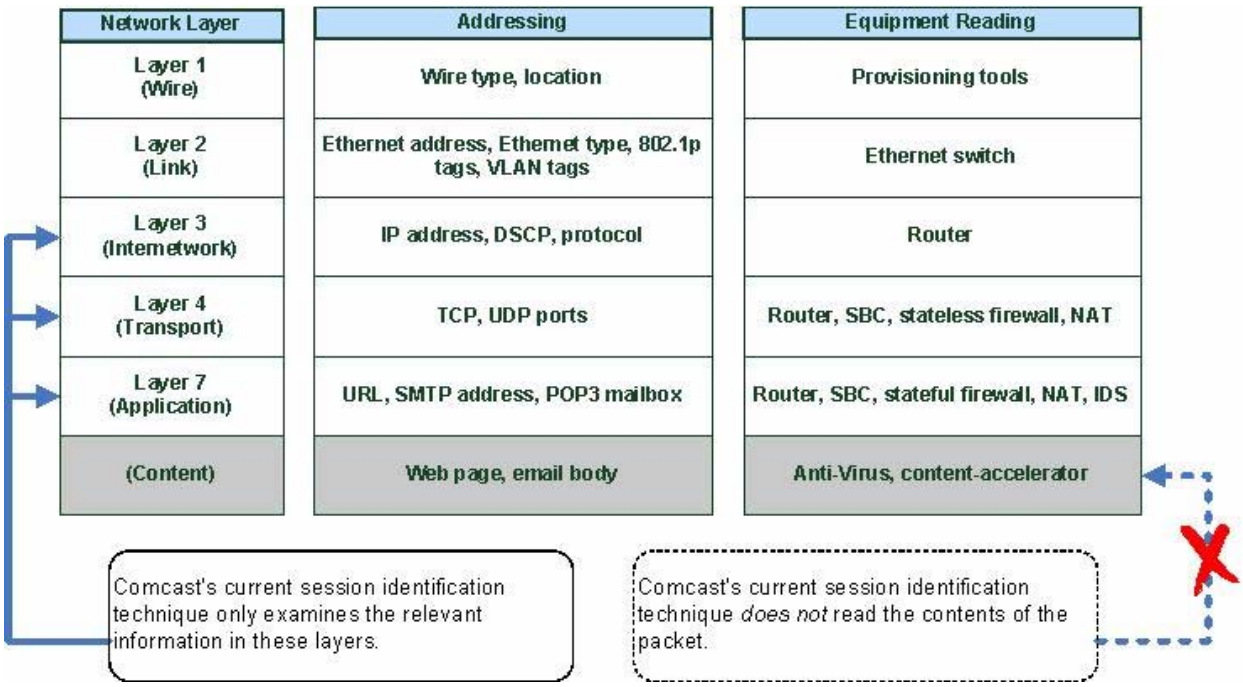

Diagram 3: Session Identification Technique.

In selecting which P2P protocol uploads to manage, network data were analyzed that identified the particular protocols that were generating disproportionate amounts of traffic. Based on that analysis, five P2P protocols were identified to be managed: Ares, BitTorrent, eDonkey, FastTrack, and Gnutella. Four of those protocols have been subject to Comcast's management practices since Comcast first implemented these practices. Ares was added in November 2007 after traffic analysis showed that it, too, was generating disproportionate demands on network resources.

For each of the managed P2P protocols, the PTS monitors and identifies the number of simultaneous unidirectional uploads that are passed from the CMTS to the upstream router. Because of the prevalence of P2P traffic on the upstream portion of our network, the number of simultaneous unidirectional upload sessions of any particular P2P protocol at any given time serves as a useful proxy for determining the level of overall network congestion. For each of the protocols, a session threshold is in place that is intended to provide for equivalently fair access

between the protocols, but still mitigate the likelihood of congestion that could cause service degradation for our customers.

Developing session thresholds for each P2P protocol must take into account the unique characteristics and behavior of each particular protocol. For example, BitTorrent and eDonkey use both bidirectional and unidirectional upload sessions, whereas Ares, FastTrack, and Gnutella only use unidirectional upload sessions.[10] And even between BitTorrent and eDonkey, there are significant differences. The BitTorrent protocol more heavily promotes bidirectional uploads as compared to eDonkey, so, while they both may have the same total number of sessions, BitTorrent would have a much higher percentage of bidirectional sessions than eDonkey. Differences also arise between Ares, FastTrack, and Gnutella. For example, each protocol consumes different amounts of bandwidth per session (e.g., a high percentage of Ares unidirectional uploads consume negligible bandwidth).

The following table lays out by protocol the simultaneous unidirectional upload session thresholds for each protocol as well as the typical ratio of bidirectional to unidirectional traffic observed on our HSI network for those P2P protocols that use both, and other factors that contribute to the overall bandwidth consumption by protocol.

[10] Session thresholds are not applied to bidirectional uploads so as to not interfere with the corresponding download.

Protocol	Ratio Bi:Uni	Session Equivalence[11]	Uni Threshold	Notes
Ares	(N/A)	150	150	Many overhead flows exist for signaling, using little or no bandwidth. The session limit is set higher to account for this. Ares is typically used for small files.
BitTorrent	~20:1	~160	8	High ratio of bidirectional to unidirectional flows. The bidirectional to unidirectional ratio varies. Typically used for large files.
eDonkey	~.3:1	~42	32	Low ratio of bidirectional to unidirectional flows. Used for large files.
FastTrack	(N/A)	24	24	Typically used for large files.
Gnutella	(N/A)	80	80	Typically used for small files.

Table 1: Managed Protocols, Relevant Thresholds, and Other Notes

When the number of unidirectional upload sessions for any of the managed P2P protocols for a particular Sandvine PTS reaches the pre-determined session threshold, the Sandvine PTS issues instructions called "reset packets" that delay unidirectional uploads for that particular P2P protocol in the geographic area managed by that Sandvine PTS. The "reset" is a flag in the packet header used to communicate an error condition in communication between two computers on the Internet. As used in our current congestion management practices, the reset packet is used to convey that the system cannot, at that moment, process additional high-resource demands without creating risk of congestion. Once the number of simultaneous unidirectional uploads falls below the pre-determined session limit threshold for a particular protocol, new uploads using that protocol are allowed to proceed. Some significant percentage of P2P sessions last

[11] This number reflects the total number of sessions that we estimate are on-going at any moment in time when the number of simultaneous upload sessions has met the threshold that has been established for that protocol.

only a few seconds, so, even when the thresholds are met, new opportunities for unidirectional uploads generally occur quite frequently.

VI. CONCLUSION

Data collected from our HSI network demonstrate that, even with these current management practices in place, P2P traffic continues to comprise approximately half of all upstream traffic transmitted on our HSI network -- and, in some locations, P2P traffic is as much as two-thirds of total upstream traffic. The data also show that, even for the most heavily used P2P protocols, more than 90 percent of these flows are unaffected by the congestion management. Data recently collected from our network show that, when a P2P upload from a particular computer was delayed by a reset packet, that same computer successfully initiated a P2P upload within one minute in 80 percent of the cases. In fact, most of our customers using P2P protocols to upload on any given day never experienced any delay at all.

Nonetheless, as Comcast previously stated and as the Order now requires, *Comcast will end these protocol-specific congestion management practices throughout its network by the end of 2008.*

Basic Glossary

Cable Modem:

A device located at the customer premise used to access the Comcast High Speed Internet (HSI) network. In some cases, the cable modem is owned by the customer, and in other cases it is owned by the cable operator. This device has an interface (i.e., someplace to plug in a cable) for connecting the coaxial cable provided by the cable company to the modem, as well as one or more interfaces for connecting the modem to a customer's PC or home gateway device (e.g., router, firewall, access point, etc.). In some cases, the cable modem function, i.e., the ability to access the Internet, is integrated into a home gateway device or embedded multimedia terminal adapter (eMTA). Once connected, the cable modem links the customer to the HSI network and ultimately the broader Internet.

Cable Modem Termination System (CMTS):

A piece of hardware located in a cable operator's local network (generally in a "headend") that acts as the gateway to the Internet for cable modems in a particular geographic area. A simple way to think of the CMTS is as a router with interfaces on one side leading to the Internet and interfaces on the other connecting to Optical Nodes and then customers.

Cable Modem Termination System Port:

A CMTS has both upstream and downstream network interfaces to serve the local access network, which we refer to as upstream or downstream ports. A port generally serves a neighborhood of hundreds of homes.

Channel Bonding:

A technique for combining multiple downstream and/or upstream channels to increase customers' download and/or upload speeds, respectively. Multiple channels from the HFC network can be bonded into a single virtual port (called a bonded group), which acts as a large single channel or port to provide increased speeds for customers. Channel bonding is a feature of Data Over Cable Service Interface Specification (DOCSIS) version 3.

Coaxial Cable (Coax):

A type of cable used by a cable operator to connect customer premise equipment (CPE) -- such as TVs, cable modems (including embedded multimedia terminal adapters), and Set Top Boxes -- to the Hybrid Fiber Coax (HFC) network. There are many grades of coaxial cable that are used for different purposes. Different types of coaxial cable are used for different purposes on the network.

Comcast High Speed Internet (HSI):

A service/product offered by Comcast for delivering Internet service over a broadband connection.

Customer Premise Equipment (CPE):

Any device that resides at the customer's residence.

Data Over Cable Service Interface Specification (DOCSIS):

A reference standard that specifies how components on cable networks need to be built to enable HSI service over an HFC network. These standards define the specifications for the cable modem and the CMTS such that any DOCSIS certified cable modem will work on any DOCSIS certified CMTS independent of the selected vendor. The interoperability of cable modems and cable modem termination systems allows customers to purchase a DOCSIS certified modem from a retail outlet and use it on their cable-networked home. These standards are available to the public at the CableLabs website, at http://www.cablelabs.com.

Downstream:

Description of the direction in which a signal travels. Downstream traffic occurs when users are downloading something from the Internet, such as watching a YouTube video, reading web pages, or downloading software updates.

Headend:

A cable facility responsible for receiving TV signals for distribution over the HFC network to the end customers. This facility typically also houses the cable modem termination systems. This is sometimes also called a "hub."

Hybrid Fiber Coax (HFC):

Network architecture used primarily by cable companies, comprising of fiber optic and coaxial cables that deliver Voice, Video, and Internet services to customers.

Internet Protocol (IP):

Set of standards for sending data across a packet switched network like the Internet. In the Open System Interconnection Basic Reference Model (OSI) model, IP operates in the "Network Layer" or "Layer 3." The HSI product utilizes IP to provide Internet access to customers.

Internet Protocol Detail Record (IPDR):

Standardized technology for monitoring subscribers' upstream and downstream Internet usage data based on their cable modem. The data is collected from the CMTS and sent to a server for further processing. Additional information is available at: http://www.ipdr.org.

Optical Node:

A component of the HFC network generally located in customers' local neighborhoods that is used to convert the optical signals sent over fiber-optic cables to electrical signals that can be sent over coaxial cable to customers' cable modems, or vice versa. A fiber optic cable connects the Optical Node, through distribution hubs, to the CMTS and coaxial cable connects the Optical Node to customers' cable modems.

Open System Interconnection Basic Reference Model (OSI Model):

A framework for defining various aspects of a communications network in a layered approach. Each layer is a collection of conceptually similar functions that provide services to the layer above it, and receive services from the layer below it. The seven layers of the OSI model are listed below:

Layer 7 – Application
Layer 6 – Presentation
Layer 5 – Session
Layer 4 – Transport
Layer 3 – Network
Layer 2 – Data Link
Layer 1 – Physical

Port:

A port is a physical interface on a device used to connect cables in order to connect with other devices for transferring information/data. An example of a physical port is a CMTS port. Prior to DOCSIS version 3, a single CMTS physical port was used for either transmitting or receiving data downstream or upstream to a given neighborhood. With DOCSIS version 3, and the channel bonding feature, multiple CMTS physical ports can be combined to create a virtual port.

Provisioned Bandwidth:

Comcast-specific definition The peak speed associated with a tier of service purchased by a customer. For example, a customer with a 16 Mbps/2 Mbps (Down/Up) speed tier would be said to be provisioned with 16 Mbps of downstream bandwidth and 2 Mbps of upstream bandwidth.

Quality of Service (QoS):

Set of techniques to manage network resources to ensure a level of performance to specific data flows. One method for providing QoS to a network is by differentiating the type of traffic by class or flow and assigning priorities to each type. When the network becomes congested, the data packets that are marked as having higher priority will have higher likelihood of getting serviced.

Transmission Control Protocol (TCP):

Set of standard rules for reliably communicating data between programs operating on computers. TCP operates in the "Transport Layer" or "Layer 4" of the OSI model and deals with the ordered delivery of data to specific programs. If we compare the data communication network to the US Postal Service mail with delivery confirmation, the Network Layer would be analogous to the Postal Address of the recipient where the TCP Layer would be the ATTN field or the person that is to receive the mail. Once the receiving program receives the data, an acknowledgement is returned to the sending program.

Upstream:

Description of the direction in which a signal travels. Upstream traffic occurs when users are uploading something to the network, such as sending email, sharing P2P files, or uploading photos to a digital photo website.

Sandvine Policy Traffic Switch Portfolio

Real Performance. Real Networks. Real Intelligence

Sandvine Policy Traffic Switch (PTS) portfolio helps service providers to better profit from application traffic. Our policy management solutions address key challenges such as managing bandwidth-intensive traffic, controlling malicious threats, enabling new services and identifying application quality trends. These policy management solutions are deployed on a single intelligent platform to simplify the network architecture and ensure a fast return on investment.

Service providers today are focusing on policy management solutions to differentiate and protect their IP service offering. This critical technology must be easily deployed throughout the network and delivered reliably for today's demanding subscribers. Sandvine has recognized these important requirements with the introduction of 10 GE interfaces, scalable performance with the largest network deployment base of 80Gbps and an innovative approach to transparently handle asymmetry in large tier-one networks.

With extensive deployments worldwide, Sandvine has proven experience in large service provider networks. Let us show you how we can help your business.

Policy Management Solutions

Sandvine's portfolio of policy management solutions addresses immediate challenges throughout the service provider: Engineering must manage growing bandwidth pressures; Marketing must launch new services; Security must control network threats; and Network Operations has to ensure service quality. Sandvine provides visibility into application traffic along with a powerful set of policy solutions.

Traffic Optimization

File-sharing traffic still dominates networks with even more subscribers actively using P2P today. The implication is clear – Network Engineering must preserve the subscriber experience while managing traffic intelligently. Sandvine's Intelligent Traffic Management solution offers the widest range of subscriber fairness options including P2P session management and adaptive traffic shaping. And our powerful network-class capability enables service providers to easily tailor traffic policies for each service category.

Network Integrity

A service provider's own subscribers have become a major source of malicious traffic leading to inconsistent service quality. Sandvine's Network Integrity solution protects the network against unpredictable attacks such as worms, DoS attacks, spam trojans and DNS server attacks. Our unique combination of DPI signature recognition plus behavioral analysis means effective zero-day protection and fewer customer support calls.

Service Creation

Today's best-effort networks are poorly suited to deal with jitter and delay sensitive multimedia applications. This means service providers have an immediate opportunity to introduce new prioritized services for gaming, VoIP and other multimedia applications. The latest capabilities of the PTS enable service providers to generate an additional revenue stream from targeted advertising.

Operations Management

Residential subscribers are demanding higher service quality as Internet applications, particularly VoIP, become a key part of our lifestyles. The result for service providers? Escalating costs as Customer Support deal with difficult complaints about service degradation. Sandvine's solution uses comprehensive VoIP QoE metrics to alert Network Operations about VoIP quality trends so they can resolve problems before subscribers are affected.

What are your subscribers doing online?

Subscriber behavior has always been difficult to characterize. What applications are popular? How does usage vary by service? Are third-party services increasing? Sandvine Network Demographics reporting, without impinging on subscriber privacy, provides valuable insights into application behavior and trends.

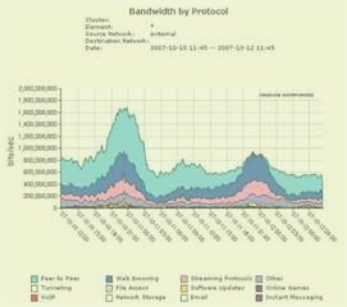

And with an understanding of subscriber behavior comes the ability to implement policies that improve the subscriber experience.

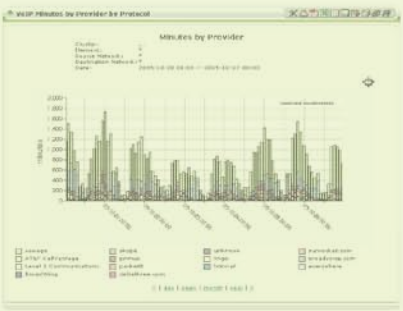

PTS Platform Advantages

Gain Control & Visibility

Major service provider networks are prone to asymmetric routes where inbound and outbound traffic for a subscriber follow different routes - making stateful identification impossible. Sandvine has overcome this challenge with unique cluster technology that ensures accurate protocol identification without any changes to the network routing. Today, most service providers are deploying PTS platforms towards the services edge of the network for maximum visibility of subscriber traffic – this ensures application traffic is properly prioritized as it first enters the network and provides complete network-wide reporting.

True Scalability

Our PTS portfolio provides true scalability for extremely cost-effective deployments from small remote hubs to large centralized sites. The PTS 14000 and PTS 8210 platforms offer a broad selection of performance, interface types and port densities that are easily tailored to the wide-ranging sites across large service provider networks. Both platforms deliver the same proven feature set to ensure that subscriber services and traffic management policies are universal across the entire network.

The PTS 14000 architecture delivers up to a remarkable 80 Gbps easily handling today's growing multimedia traffic. In addition, Sandvine's Virtual Switch Clusters combine the inspection performance and data ports from multiple PTS 14000 units to deliver the performance necessary at major aggregation points.

Leading Redundancy Options

High availability is critical throughout the services architecture so Sandvine provides a complete range of redundancy features and options. For many service providers, the preferred approach is to introduce n:n+1 redundancy, that is not only extremely cost-effective, but offers more comprehensive protection than simple hardware replication alone. Our integrated interface-bypass capability ensures that subscriber traffic continues uninterrupted in the event of a power failure or major system event.

Centralized Policy Management

Our Service Definition Manager grants service providers complete control over their network. The intuitive and user-friendly GUI enables them to configure and deploy network-wide policies quickly and easily from a central management system. The Service Definition Manager also provides centralized physical element administration, deployments management and network association control.

Let Sandvine recommend the best deployment options for your unique service provider network.

With over 150 fully-customizable reports, Sandvine Network Demographics provides decision-critical information to major organizational departments - including marketing, operations, security, and support.

Characterize and apply policy to:

- Peer-to-Peer
- Web Browsing
- News Groups
- Instant Messaging
- Email
- Database
- Streaming Protocols
- Tunneling (including VPN)
- VoIP
- Remote Connectivity
- Network Administration
- File Access Protocols
- Network Storage protocols
- Gaming protocols
- ...and many more.

Technical Summary

General

- PTS platforms are highly scalable network elements enabling Sandvine's full suite of policy management solutions
- Leading stateful traffic identification using advanced signature recognition and behavioral analysis
- Broadest range of policy management options including unique subscriber fairness approaches to preserve your subscribers' Internet experience
- Cluster technology delivers unmatched scalability and easily solves network asymmetric route challenges

Platform and Interface Options

Platforms can be deployed independently or as clusters with PTS 14000 Virtual Switch Clusters delivering up to 80 Gbps performance.

Platform	1452X	1451X	1421X	8210
Performance	10 Gbps	10 Gbps	4 Gbps	1 Gbps
10 GE	2 port	-	-	-
GE MM or SM	-	12 port	12 port	4 port
10/100/1000T	-	12 port	12 port	4 port
1GE Integration	4	8	8	1 port
10 GE Integration	6 port	2 port	2 port	-

* Performance shown as aggregate throughput
* Maximum number of interface ports is shown
* Bypass interfaces can be internal and/or external
* 10GE Integration ports are normally used for virtual switch cluster connectivity

High Availability

- n:n+1 for simple, cost-effective redundancy
- Interface bypass option for power failures and major system events
- Optional features such as hardware replication and link-pair redundancy

Protocol Support

- Peer-to-peer, web browsing, news groups, instant messaging, email, database, streaming protocols, tunneling including VPN, VoIP, remote connectivity, network administration, file access protocols, network storage protocols, game protocols

Policy Management

- Powerful network-class capability to apply policy management by service type
- Session Management – optimizes upstream bandwidth - intensive application sessions
- FairShare™ - allocates equitable network resouces during periods of congestion
- Traffic Shaping - manages aggregate or per subscriber application traffic
- Integration Interfaces - targeted advertising, caching, content filtering and traffic mirroring
- Applicationn Marking - sets application IP TOS/DSCP field bits to identify prioritized traffic
- Captive portal communications - redirects user HTTP sessions to target URL for better service management
- Attack Traffic Mitigation detects and mitigates network threats before services are affected

Subscriber Policy Broker

- Subscriber Policy Broker for simple, centralized policy definition & management
- Service Definition Manager - user friendly and intuitive GUI for quick and easy centralized policy configuration and deployment
- Industry-standard API to easily integrate with external systems
- Subscriber identification option via DHCP, RADIUS, Diameter or OSS integration

Network Demographics

- Advanced application reporting with full network-wide, region and subscriber views
- Unique network-class reporting to better understand subscriber behavior
- Value-added analysis including VoIP call details and QoE reporting

Physical Specifications *(PTS 14000 / 8210)*

- Dimensions: W 17" x H 7" x D 23" (432mm x 177.8mm x 584.2mm) / W 16.75" x H 1.75" x D 22.5" (419mm x 44.4mm x 571.5mm)
- Mounting: 19" rack - 4 RU / 19" rack - 1 RU
- Weight: 75 lbs (34 kg) / 24 lbs (11 kg)
- AC 100 – 240 V or DC 42 – 60 V input
- Operating temperature 0°C to +40°C
- Humidity 5% to 85% non-condensing

Approvals

- NEBS Level 3 compliance
- Product Safety and EMC approvals for Argentina, Australia, Canada, Europe, Mexico, New Zealand, Russia, South Africa, South Korea, Taiwan and USA

sandvine
Intelligent Broadband Networks

Headquarters
Sandvine Incorporated
Waterloo, Canada (GMT-5)
Phone: +1 519.880.2600
Email: sales@sandvine.com

European Offices
Sandvine Limited
Basingstoke, UK (GMT)
Phone: +44 (0) 1256 698021
Email: sales@sandvine.co.uk

Visit www.sandvine.com

Last Revised:2008-05-14

ATTACHMENT B:

COMCAST CORPORATION
DESCRIPTION OF PLANNED NETWORK MANAGEMENT PRACTICES TO BE DEPLOYED FOLLOWING THE TERMINATION OF CURRENT PRACTICES

COMCAST CORPORATION
DESCRIPTION OF PLANNED NETWORK MANAGEMENT PRACTICES TO BE DEPLOYED FOLLOWING THE TERMINATION OF CURRENT PRACTICES

Pursuant to Paragraphs 54 and 59 of the Commission's *Memorandum Opinion & Order* regarding how Comcast manages congestion on its High-Speed Internet ("HSI") network, Comcast hereby "disclose[s] to the Commission and the public the details of the network management practices that it intends to deploy following the termination of its current practices, including the thresholds that will trigger any limits on customers' access to bandwidth."[1]

I. INTRODUCTION & SUMMARY

Comcast's HSI network is a shared network. This means that our HSI customers share upstream and downstream bandwidth with their neighbors. Although the available bandwidth is substantial, so, too, is the demand. Thus, when a relatively small number of customers in a neighborhood place disproportionate demands on network resources, this can cause congestion that degrades their neighbors' Internet experience.[2] The goal of Comcast's new congestion management practices will be to enable all users of our network resources to access a "fair share" of that bandwidth, in the interest of ensuring a high-quality online experience for all of Comcast's HSI customers.[3]

---

[1] *In re Formal Complaint of Free Press & Pub. Knowledge Against Comcast Corp. for Secretly Degrading Peer-to-Peer Applications; Broadband Industry Practices; Petition of Free Press et al. for Declaratory Ruling That Degrading an Internet Application Violates the FCC's Internet Policy Statement & Does Not Meet an Exception for "Reasonable Network Management*," Mem. Op. and Order, FCC 08-183 ¶¶ 54, 59 (Aug. 20, 2008) ("*Order*").

[2] Although the *Order* focuses entirely on Comcast's current practices with respect to controlling network congestion, Comcast's efforts to deliver a superior Internet experience involve a wide variety of other network management efforts beyond congestion control. As Comcast has previously explained, we actively manage our HSI network in order to enhance our customers' Internet experience by, among other things, blocking spam, preventing viruses from harming the network and our subscribers, thwarting denial-of-service attacks, and empowering our customers' ability to control the content that enters their homes.

[3] These congestion management practices are independent of, and should not be confused with, our recent announcement that we will amend the "excessive use" portion of our Acceptable Use Policy, effective October 1, 2008, to establish a specific monthly data usage threshold of 250 GB per account for all residential HSI customers. This excessive use threshold is designed to prevent any one residential account from consuming excessive amounts

Importantly, the new approach will be protocol-agnostic; that is, it *will not* manage congestion by focusing on the use of the specific protocols that place a disproportionate burden on network resources, or any other protocols. Rather, the new approach will focus on managing the traffic of those individuals who are using the most bandwidth at times when network congestion threatens to degrade subscribers' broadband experience and who are contributing disproportionately to such congestion at those points in time.

Specific details about these practices, including relevant threshold information, the type of equipment used, and other particulars, are discussed at some length later in this document. At the outset, however, we present a very high-level, simplified overview of how these practices will work once they are deployed. Despite all the detail provided further below, the fundamentals of this approach can be summarized succinctly:

1. Software installed in the Comcast network continuously examines aggregate traffic usage data for individual segments of Comcast's HSI network. If overall upstream or downstream usage on a particular segment of Comcast's HSI network reaches a pre-determined level, the software moves on to step two.

2. At step two, the software examines bandwidth usage data for subscribers in the affected network segment to determine which subscribers are using a disproportionate share of the bandwidth. If the software determines that a particular subscriber or subscribers have been the source of high volumes of network traffic during a recent period of minutes, traffic originating from that subscriber or those subscribers temporarily will be assigned a lower priority status.

3. During the time that a subscriber's traffic is assigned the lower priority status, such traffic will not be delayed so long as the network segment is not actually congested. If, however, the network segment becomes congested, such traffic could be delayed.

4. The subscriber's traffic returns to normal priority status once his or her bandwidth usage drops below a set threshold over a particular time interval.

---

of network resources as measured over the course of a month. That cap does not address the issue of network congestion, which results from traffic levels that vary from minute to minute. We have long had an "excessive use" limit in our Acceptable Use Policy but have been criticized for failing to specify what is considered to be "excessive." The new cap provides clarity to customers regarding the specific monthly consumption limit per account. As with the existing policy, a user who violates the excessive use policy twice within six months is subject to having his or her Internet service account terminated for one year.

We have made considerable progress in recent months in formulating our plans for this new approach, adjusting them, and subjecting them to real-world trials. Market trials in Chambersburg, PA; Warrenton, VA; Lake City, FL; East Orange, FL; and Colorado Springs, CO have enabled us to validate the utility of the general approach and collect substantial trial data to test multiple variations and alternative formulations.

Comcast appreciates the *Order*'s recognition that Comcast "may not have finalized the details of the network management practices that it intends to deploy following termination of its current practices" by the date of this report,[4] but our progress to date is sufficient that we do not need to make the certification contemplated by the *Order* or postpone disclosing the details of our current plans. Certainly some additional adjustments -- and possibly material changes -- will be made as we continue our trials and move forward with implementation. Thus, consistent with the spirit of the language quoted above, Comcast commits that, until we have completed our transition to the protocol-agnostic congestion management practices described below, we will inform the Commission and the public of any material changes to the practices and plans detailed here, at least two weeks prior to implementation of any such changes.[5]

II. IMPLEMENTATION AND CONFIGURATION

To understand exactly how these new congestion management practices will work, it will be helpful to have a general understanding of how Comcast's HSI network is designed. Comcast's HSI network is what is commonly referred to as a hybrid fiber-coax network, with coaxial cable connecting each subscriber's cable modem to an Optical Node, and fiber optic cables connecting the Optical Node, through distribution hubs, to the Cable Modem Termination

[4] *Order* ¶ 55 n.246.

[5] We recognize that clear communication with our customers is an important part of a successful long-term relationship. On an ongoing basis, we will provide our customers with clear, concise, and useful information about the services that we provide.

System ("CMTS"), which is also known as a "data node."[6] The CMTSes are then connected to higher-level routers, which in turn are connected to Comcast's Internet backbone facilities. Today, Comcast has approximately 3300 CMTSes deployed throughout our network, serving our 14.4 million HSI subscribers.

Each CMTS has multiple "ports" that handle traffic coming into and leaving the CMTS. In particular, each cable modem deployed on the Comcast HSI network is connected to the CMTS through the ports on the CMTS. These ports can be either "downstream" ports or "upstream" ports, depending on whether they send information to cable modems (downstream) or receive information from cable modems (upstream) attached to the port.[7] Today, on average, about 275 cable modems share the same downstream port and about 100 cable modems share the same upstream port. Both types of ports can experience congestion that could degrade the broadband experience of our subscribers and, unlike with the previous congestion management practices, both upstream and downstream traffic will be subject to management under these new practices.

To implement Comcast's new protocol-agnostic congestion management practices, Comcast will purchase new hardware and software that will be deployed near the Regional Network Routers ("RNRs") that are further upstream in Comcast's network. This new hardware will consist of Internet Protocol Detail Record ("IPDR") servers, Congestion Management servers, and PacketCable Multimedia ("PCMM") servers. Further details about each of these pieces of equipment can be found below, in Section III. It is important to note here, however,

---

[6] The reader may find it useful to refer to the attached glossary for additional explanation of unfamiliar terms.

[7] The term "port" as used here generally contemplates single channels on a CMTS, but these statements will apply to virtual channels, also known as "bonded groups," in a DOCSIS 3.0 environment.

that, even though the physical location of these servers is at the RNR, the servers will communicate with -- and manage individually -- multiple ports on multiple CMTSes to effectuate the practices described in this document. That is to say, bandwidth usage on one CMTS port will have no effect on whether the congestion management practices described herein are applied to a subscriber on a different CMTS port.

The following diagram provides a simplified graphical depiction of the network architecture just described:

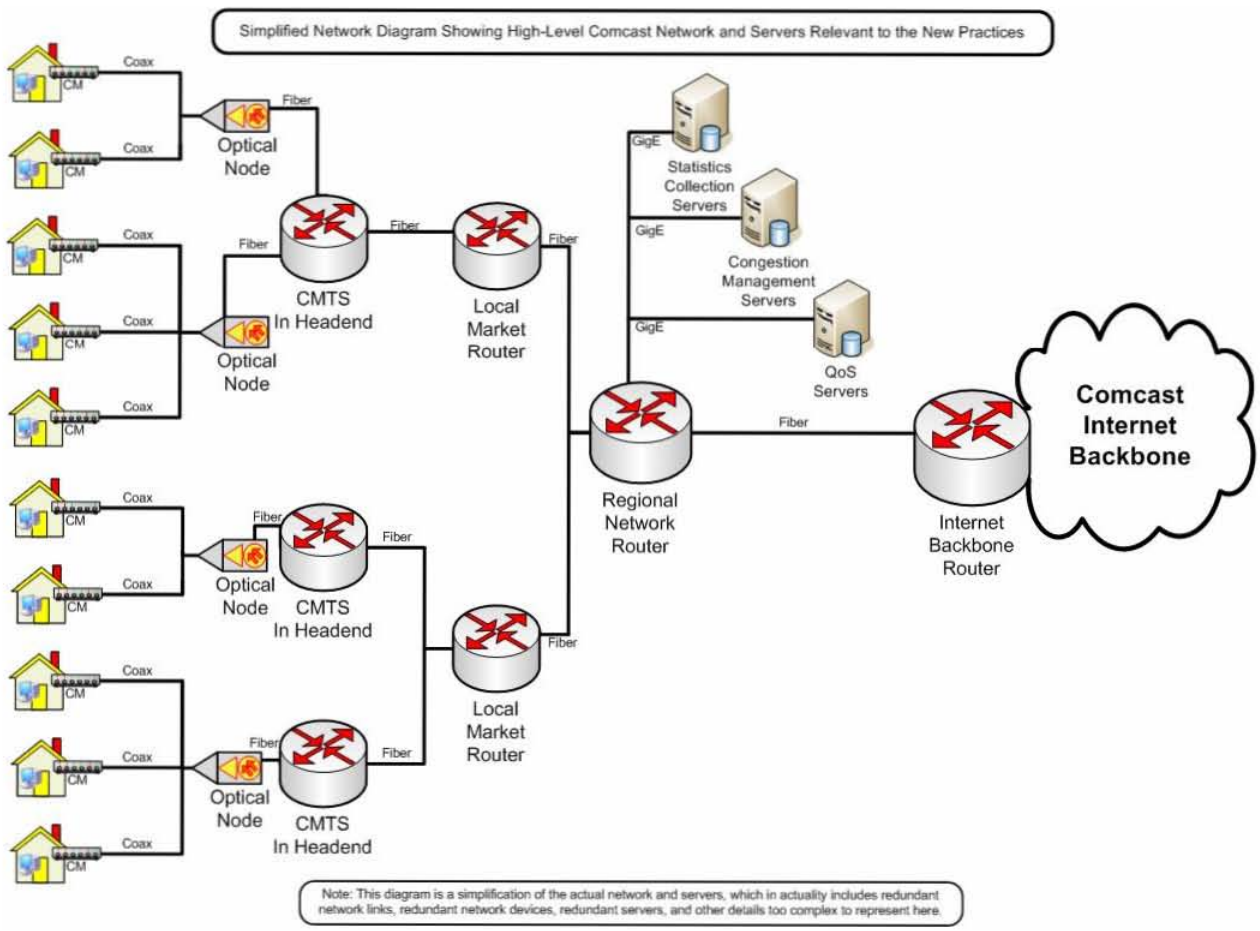

Diagram 1: Comcast Network Design

Each Comcast HSI subscriber's cable modem has a "bootfile" that contains certain pieces of information about the subscriber's service to ensure that the service functions properly.[8] For example, the bootfile contains information about the maximum speed (what we refer to in this document as the "provisioned bandwidth") that a particular modem can achieve based on the tier (personal, commercial, etc.) the customer has purchased. Bootfiles are generally reset from time to time to account for changes in the network and other updates, and this is usually done through a command sent from the network and without any effect on the subscriber. In preparation for the transition to the new practices, Comcast will send new bootfiles to our HSI customers' cable modems that will create two Quality of Service ("QoS") levels for Internet traffic going to and from the cable modem: (1) "Priority Best-Effort" traffic ("PBE"); and (2) "Best-Effort" traffic ("BE"). As with previous changes to cable modem bootfiles, the replacement of the old bootfile with the new bootfile requires no active participation by Comcast customers.[9]

Thereafter, all traffic going to or coming from cable modems on the Comcast HSI network will be designated as either PBE or BE. PBE will be the default status for all Internet traffic coming from or going to a particular cable modem. Traffic will be designated BE for a particular cable modem only when both of two conditions are met:

- First, the usage level of a particular upstream or downstream port of a CMTS, as measured over a particular period of time, must be nearing the point where congestion could degrade users' experience. We refer to this as the "Near Congestion State" and, based on the technical trials we have conducted, we have established a threshold, described in more detail below, for when a particular CMTS port enters that state.

---

[8] No personal information is included in the bootfile; it only includes information about the service that the subscriber has purchased.

[9] A very small percentage of Comcast's HSI customers use first-generation cable modems that cannot support the new congestion management practices. These cable modems will not receive the new bootfiles and, after December 31, 2008, those cable modems will not be subject to congestion management and all their traffic effectively will be designated PBE. These older cable modems have less capability to utilize significant amounts of bandwidth and will, in any event, be replaced over time.

- Second, a particular subscriber must be making a significant contribution to the bandwidth usage on the particular port, as measured over a particular period of time. We refer to this as the "Extended High Consumption State" and, based on the technical trials we have conducted, we have established a threshold, described in more detail below, for when a particular user enters that state.

When, *and only when*, both conditions are met, a user's upstream or downstream traffic (depending on which type of port is in the Near Congestion State) will be designated as BE. Then, to the extent that actual congestion occurs, any delay resulting from the congestion will affect BE traffic before it affects PBE traffic.

We now explain the foregoing in greater detail.

A. Thresholds For Determining When a CMTS Port Is in a Near Congestion State

For a CMTS port to enter the Near Congestion State, traffic flowing to or from that CMTS port must exceed a specified level (the "Port Utilization Threshold") for a specific period of time (the "Port Utilization Duration"). The Port Utilization Threshold on a CMTS port is measured as a percentage of the total aggregate upstream or downstream bandwidth for the particular port during the relevant timeframe. The Port Utilization Duration on the CMTS is measured in minutes.

Values for each of the thresholds to be used as part of this new management technique have been tentatively established after an extensive process of lab tests, simulations, technical trials, vendor evaluations, customer feedback, and a third-party consulting analysis. In the same way that specific anti-spam or other network management practices are adjusted to address new issues that arise, it is a near certainty that these values will change in both the short-term and the long-term, as Comcast gathers more data and performs additional analysis resulting from wide-scale deployment of the new technique. Moreover, as with any large network or software system, software bugs and/or unexpected errors may arise, requiring software patches or other

corrective actions. As always, our decisions on these matters will be driven by the marketplace imperative that we deliver the best possible experience to our HSI subscribers.

Given our experience so far, we have determined that a starting point for the upstream Port Utilization Threshold should be 70 percent and the downstream Port Utilization Threshold should be 80 percent. For the Port Utilization Duration, we have determined that the starting point should be approximately 15 minutes (although some technical limitations in some newer CMTSes deployed on Comcast's network may make this time period vary slightly). Thus, over any 15-minute period, if an average of more than 70 percent of a port's upstream bandwidth capacity or more than 80 percent of a port's downstream bandwidth capacity is utilized, that port will be determined to be in a Near Congestion State.

Based on the trials to date, we expect that a typical CMTS port on our HSI network will be in a Near Congestion State only for relatively small portions of the day, if at all, though there is no way to forecast what will be the busiest time on a particular port on a particular day. Moreover, the trial data indicate that, even when a particular port is in a Near Congestion State, the instances where the network *actually* becomes congested during the Port Utilization Duration are few, and managed users whose traffic is delayed during those congested periods perceive little, if any, effect, as discussed below.

B. Thresholds For Determining When a User Is in an Extended High Consumption State and for Release from that Classification

Once a particular CMTS port is in a Near Congestion State, the software examines whether any cable modems are consuming bandwidth disproportionately.[10] For a user to enter an

[10] Although each cable modem is typically assigned to a particular household, the software does not (and cannot) actually identify individual users or analyze particular users' traffic. For purposes of this report, we use "cable modem," "user," and "subscriber" interchangeably to mean a subscriber account or user account and not an individual person.

Extended High Consumption State, he or she must consume greater than a certain percentage of his or her provisioned upstream or downstream bandwidth (the "User Consumption Threshold") for a specific length of time (the "User Consumption Duration"). The User Consumption Threshold is measured as a user's consumption of a particular percentage of his or her total provisioned upstream *or* downstream bandwidth (the maximum speed that a particular modem can achieve based on the tier (personal, commercial, etc.) the customer has purchased, e.g., if a user buys a service with speeds of 8 Mbps downstream and 1 Mbps upstream, then his or her provisioned downstream speed is 8 Mbps and provisioned upstream speed is 1 Mbps).[11] The User Consumption Duration is measured in minutes.

Following lab tests, simulations, technical trials, customer feedback, vendor evaluations, and a third-party consulting analysis, we have determined that the appropriate starting point for the User Consumption Threshold is 70 percent of a subscriber's provisioned upstream or downstream bandwidth, and that the appropriate starting point for the User Consumption Duration is 15 minutes. That is, when a subscriber uses an average of 70 percent or more of his or her provisioned upstream or downstream bandwidth over a particular 15-minute period, that user will be in an Extended High Consumption State.[12] As noted above, these values are subject to change as necessary in the same way that specific anti-spam or other network management practices are adjusted to address new issues that arise, or should unexpected software bugs or other problems arise.

---

[11] Because the User Consumption Threshold is a percentage of provisioned bandwidth for a particular user account, and not a static value, users of higher speed tiers will have correspondingly higher User Consumption Thresholds.

[12] The User Consumption Thresholds have been set sufficiently high that using the HSI connection for VoIP or most streaming video cannot alone cause subscribers to our standard-level HSI service to exceed the User Consumption Threshold. For example, while Comcast's standard-level HSI service provisions downstream bandwidth at 6 Mbps, today, streaming video (even some HD video) from Hulu uses less than 2.5 Mbps, a Vonage or Skype VoIP call uses less than 131 Kbps, and streaming music uses less than 128 Kbps.

Based on data collected from the trial markets where the new management practices are being tested, on average less than one-third of one percent of subscribers have had their traffic priority status changed to the BE state on any given day. For example, in Colorado Springs, CO, the largest test market, on any given day in August 2008, an average of 22 users out of 6,016 total subscribers in the trial had their traffic priority status changed to BE at some point during the day.

A user's traffic is released from a BE state when the user's bandwidth consumption drops below 50 percent of his or her provisioned upstream or downstream bandwidth for a period of approximately 15 minutes. These release criteria are intended to minimize (and hopefully prevent) user QoS oscillation, i.e., a situation in which a particular user could cycle repeatedly between BE and PBE. NetForecast, Inc., an independent consultant retained to provide analysis and recommendations regarding Comcast's trials and related congestion management work, suggested this approach, which has worked well in our ongoing trials and lab testing.[13] In trials, we have observed that user traffic rarely remains in a managed state longer than the initial 15-minute period.

Simply put, there are four steps to determining whether the traffic associated with a particular cable modem is designated as PBE or BE:

1. Determine if the CMTS port is in a Near Congestion State.
2. If yes, determine whether any users are in an Extended High Consumption State.
3. If yes, change those users' traffic to BE from PBE. If the answer at either step one or step two is no, no action is taken.

[13] NetForecast, Inc. is an internationally recognized engineering consulting company that, among other things, advises network operators and technology vendors about technology issues and how to improve the performance of a network.

4. If a user's traffic has been designated BE, check user consumption at next interval. If user consumption has declined below predetermined threshold, reassign the user's traffic as PBE. If not, recheck at next interval.

The following diagram graphically depicts how this management process would work in the case of a situation where upstream port utilization may be reaching a Near Congestion State (the same diagram, with different values in the appropriate places, could be used to depict the management process for downstream ports, as well):

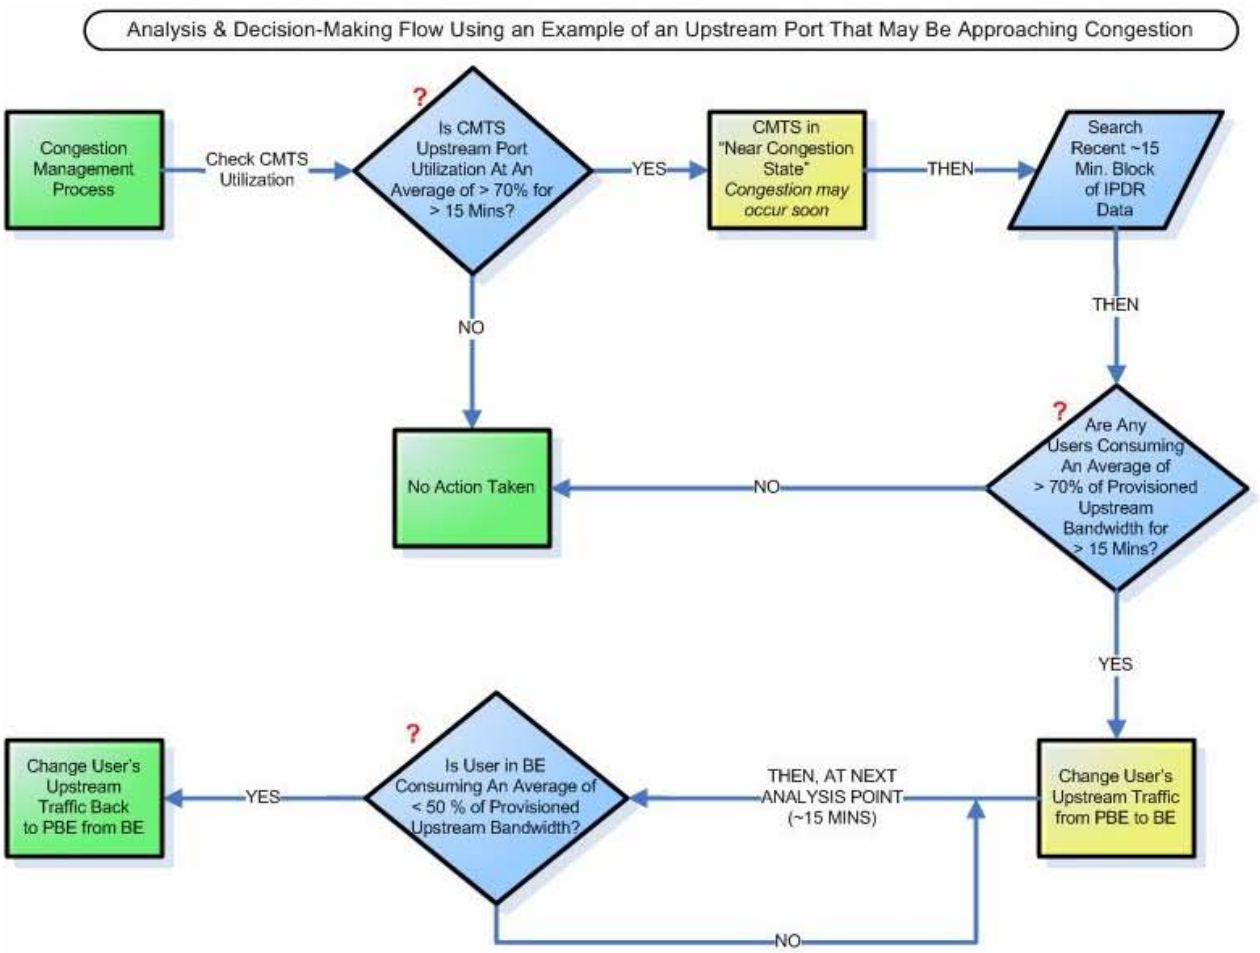

Diagram 2: Upstream Congestion Management Decision Flowchart

C. Effect of BE Quality of Service on Users' Broadband Experience

When a CMTS port is in a Near Congested State and a cable modem connected to that port is in an Extended High Consumption State, that cable modem's traffic will be designated as

BE. Depending upon the level of congestion in the CMTS port, this designation may or may not result in the user's traffic being delayed or, in extreme cases, dropped before PBE traffic is dropped.[14] This is because of the way that the CMTS handles traffic. Specifically, CMTS ports have what is commonly called a "scheduler" that puts all the packets coming from or going to cable modems on that particular port in a queue and then handles them in turn. A certain number of packets can be processed by the scheduler in any given moment; for each time slot, PBE traffic will be given priority access to the available capacity, and BE traffic will be processed on a space-available basis.

A rough analogy would be to busses that empty and fill up at incredibly fast speeds. As empty busses arrive at the figurative "bus stop" -- every two milliseconds in this case -- they fill up with as many packets as are waiting for "seats" on the bus, to the limits of the bus' capacity. During non-congested periods, the bus will usually have several empty seats, but, during congested periods, the bus will fill up and packets will have to wait for the next bus. It is in the congested periods that BE packets will be affected. If there is no congestion, packets from a user in a BE state should have little trouble getting on the bus when they arrive at the bus stop. If, on the other hand, there is congestion in a particular instance, the bus may become filled by packets in a PBE state before any BE packets can get on. In that situation, the BE packets would have to wait for the next bus that is not filled by PBE packets. In reality, this all takes place in two-millisecond increments, so even if the packets miss 50 "busses," the delay only will be about *one-tenth of a second*.

[14] Congestion can occur in any IP network, and, when it does, packets can be delayed or dropped. As a result, applications and protocols have been designed to deal with this reality. Our new congestion management practices will ensure that, in those rare cases where packets may be dropped, BE packets will be dropped before PBE packets are dropped.

During times of actual network congestion, when BE traffic might be delayed, there are a variety of effects that could be experienced by a user whose traffic is delayed, depending upon what applications he or she is using. Typically, a user whose traffic is in a BE state during actual congestion may find that a webpage loads sluggishly, a peer-to-peer upload takes somewhat longer to complete, or a VoIP call sounds choppy. Of course, the same thing could happen to the customers on a port that is congested *in the absence of any congestion management*; the difference here is that the effects of any such delays are shifted toward those who have been placing the greatest burden on the network, instead of being distributed randomly among the users of that port without regard to their consumption levels.

NetForecast, Inc. explored the potential risk of a worst-case scenario for users whose traffic is in a BE state: the possibility of "bandwidth starvation" in the theoretical case where 100 percent of the CMTS bandwidth is taken up by PBE traffic for an extended period of time. In theory, such a condition could mean that a given user whose traffic is designated BE would be unable to effectuate an upload or download (as noted above, both are managed separately) for some period of time. However, when these management techniques were tested, first in company testbeds and then in our real-world trials conducted in the five markets, such a theoretical condition did not occur. In addition, trial results demonstrated that these management practices have very modest real-world impacts. To date, *Comcast has yet to receive a single customer complaint in any of the trial markets that can be traced to the new congestion management practices*, despite having broadly publicized its trials.

Comcast will continue to monitor how user traffic is affected by these new congestion management techniques and will make the adjustments necessary to ensure that all Comcast HSI customers have a high-quality Internet experience.

III. EQUIPMENT/SOFTWARE USED AND LOCATION

The above-mentioned functions will be carried out using three different types of application servers, supplied by three different vendors. As mentioned above, these servers will be installed near Comcast's regional network routers. The *exact* locations of various servers have not been finalized, but this will not change the fact that they will manage individual CMTS ports.

The first application server will be an IPDR server, which will collect relevant cable modem volume usage information from the CMTS, such as how many aggregate upstream or downstream bytes a subscriber uses over a particular period of time.[15] Comcast has not yet chosen a vendor for the IPDR servers, but is in active negotiations with several vendors.

The second application server is the Sandvine Congestion Management Fairshare ("CMF") server, which will use Simple Network Management Protocol ("SNMP") to measure CMTS port utilization and detect when a port is in a Near Congestion State. When this happens, the CMF server will then query the relevant IPDR data for a list of cable modems meeting the criteria set forth above for being in an Extended High Consumption State.

If one or more users meet the criteria to be managed, then the CMF server will notify a third application server, the PCMM application server developed by Camiant Technologies, as to which users have been in an Extended High Consumption State and whose traffic should be treated as BE. The PCMM servers are responsible for signaling a given CMTS to set the traffic for specific cable modems with a BE QoS, and for tracking and managing the state of such CMTS actions. *If no users meet the criteria to be managed, no users will have their traffic managed.*

---

[15] IPDR has been adopted as a standard by many industry organizations and initiatives, such as CableLabs, ATIS, ITU, and 3GPP, among others.

The following diagram graphically depicts the high-level management flows among the congestion management components on Comcast's network, as described above:

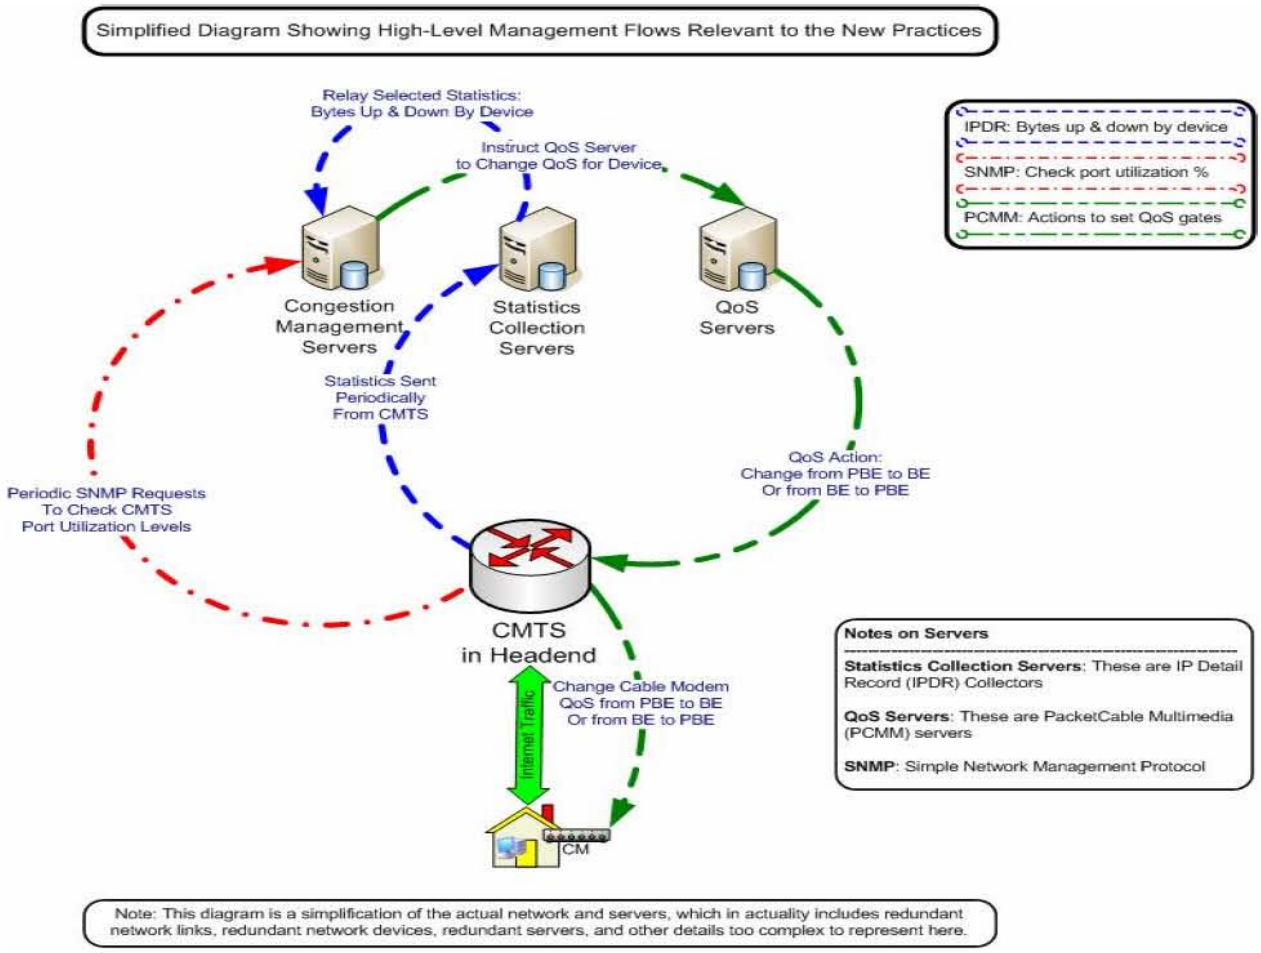

Diagram 3: High Level Management Flows

IV. CONCLUSION

Comcast's transition to protocol-agnostic congestion management is already underway, and Comcast is on schedule to meet the benchmarks set forth in Attachment C in order to complete the transition by December 31, 2008. As described above, the new approach will not manage congestion by focusing on managing the use of specific protocols. Nor will this approach use "reset packets." Rather, the new approach will (1) during periods when a CMTS port is in a Near Congestion State, (2) identify the subscribers on that port who have consumed a

disproportionate amount of bandwidth over the preceding 15 minutes, (3) lower the priority status of those subscribers' traffic to BE status until those subscribers meet the release criteria, and (4) during periods of congestion, delay BE traffic before PBE traffic is delayed. Our trials indicate that these new practices will ensure a quality online experience for all of our HSI customers.

Basic Glossary

Cable Modem:

A device located at the customer premise used to access the Comcast High Speed Internet (HSI) network. In some cases, the cable modem is owned by the customer, and in other cases it is owned by the cable operator. This device has an interface (i.e., someplace to plug in a cable) for connecting the coaxial cable provided by the cable company to the modem, as well as one or more interfaces for connecting the modem to a customer's PC or home gateway device (e.g., router, firewall, access point, etc.). In some cases, the cable modem function, i.e., the ability to access the Internet, is integrated into a home gateway device or embedded multimedia terminal adapter (eMTA). Once connected, the cable modem links the customer to the HSI network and ultimately the broader Internet.

Cable Modem Termination System (CMTS):

A piece of hardware located in a cable operator's local network (generally in a "headend") that acts as the gateway to the Internet for cable modems in a particular geographic area. A simple way to think of the CMTS is as a router with interfaces on one side leading to the Internet and interfaces on the other connecting to Optical Nodes and then customers.

Cable Modem Termination System Port:

A CMTS has both upstream and downstream network interfaces to serve the local access network, which we refer to as upstream or downstream ports. A port generally serves a neighborhood of hundreds of homes.

Channel Bonding:

A technique for combining multiple downstream and/or upstream channels to increase customers' download and/or upload speeds, respectively. Multiple channels from the HFC network can be bonded into a single virtual port (called a bonded group), which acts as a large single channel or port to provide increased speeds for customers. Channel bonding is a feature of Data Over Cable Service Interface Specification (DOCSIS) version 3.

Coaxial Cable (Coax):

A type of cable used by a cable operator to connect customer premise equipment (CPE) -- such as TVs, cable modems (including embedded multimedia terminal adapters), and Set Top Boxes -- to the Hybrid Fiber Coax (HFC) network. There are many grades of coaxial cable that are used for different purposes. Different types of coaxial cable are used for different purposes on the network.

Comcast High Speed Internet (HSI):

A service/product offered by Comcast for delivering Internet service over a broadband connection.

Customer Premise Equipment (CPE):

Any device that resides at the customer's residence.

Data Over Cable Service Interface Specification (DOCSIS):

A reference standard that specifies how components on cable networks need to be built to enable HSI service over an HFC network. These standards define the specifications for the cable modem and the CMTS such that any DOCSIS certified cable modem will work on any DOCSIS certified CMTS independent of the selected vendor. The interoperability of cable modems and cable modem termination systems allows customers to purchase a DOCSIS certified modem from a retail outlet and use it on their cable-networked home. These standards are available to the public at the CableLabs website, at http://www.cablelabs.com.

Downstream:

Description of the direction in which a signal travels. Downstream traffic occurs when users are downloading something from the Internet, such as watching a YouTube video, reading web pages, or downloading software updates.

Headend:

A cable facility responsible for receiving TV signals for distribution over the HFC network to the end customers. This facility typically also houses the cable modem termination systems. This is sometimes also called a "hub."

Hybrid Fiber Coax (HFC):

Network architecture used primarily by cable companies, comprising of fiber optic and coaxial cables that deliver Voice, Video, and Internet services to customers.

Internet Protocol (IP):

Set of standards for sending data across a packet switched network like the Internet. In the Open System Interconnection Basic Reference Model (OSI) model, IP operates in the "Network Layer" or "Layer 3." The HSI product utilizes IP to provide Internet access to customers.

Internet Protocol Detail Record (IPDR):

Standardized technology for monitoring subscribers' upstream and downstream Internet usage data based on their cable modem. The data is collected from the CMTS and sent to a server for further processing. Additional information is available at: http://www.ipdr.org.

Optical Node:

A component of the HFC network generally located in customers' local neighborhoods that is used to convert the optical signals sent over fiber-optic cables to electrical signals that can be sent over coaxial cable to customers' cable modems, or vice versa. A fiber optic cable connects the Optical Node, through distribution hubs, to the CMTS and coaxial cable connects the Optical Node to customers' cable modems.

Open System Interconnection Basic Reference Model (OSI Model):

A framework for defining various aspects of a communications network in a layered approach. Each layer is a collection of conceptually similar functions that provide services to the layer above it, and receive services from the layer below it. The seven layers of the OSI model are listed below:

Layer 7 – Application
Layer 6 – Presentation
Layer 5 – Session
Layer 4 – Transport
Layer 3 – Network
Layer 2 – Data Link
Layer 1 – Physical

Port:

A port is a physical interface on a device used to connect cables in order to connect with other devices for transferring information/data. An example of a physical port is a CMTS port. Prior to DOCSIS version 3, a single CMTS physical port was used for either transmitting or receiving data downstream or upstream to a given neighborhood. With DOCSIS version 3, and the channel bonding feature, multiple CMTS physical ports can be combined to create a virtual port.

Provisioned Bandwidth:

Comcast-specific definition The peak speed associated with a tier of service purchased by a customer. For example, a customer with a 16 Mbps/2 Mbps (Down/Up) speed tier would be said to be provisioned with 16 Mbps of downstream bandwidth and 2 Mbps of upstream bandwidth.

Quality of Service (QoS):

Set of techniques to manage network resources to ensure a level of performance to specific data flows. One method for providing QoS to a network is by differentiating the type of traffic by class or flow and assigning priorities to each type. When the network becomes congested, the data packets that are marked as having higher priority will have higher likelihood of getting serviced.

Transmission Control Protocol (TCP):

Set of standard rules for reliably communicating data between programs operating on computers. TCP operates in the "Transport Layer" or "Layer 4" of the OSI model and deals with the ordered delivery of data to specific programs. If we compare the data communication network to the US Postal Service mail with delivery confirmation, the Network Layer would be analogous to the Postal Address of the recipient where the TCP Layer would be the ATTN field or the person that is to receive the mail. Once the receiving program receives the data, an acknowledgement is returned to the sending program.

Upstream:

Description of the direction in which a signal travels. Upstream traffic occurs when users are uploading something to the network, such as sending email, sharing P2P files, or uploading photos to a digital photo website.

ATTACHMENT C:

COMCAST CORPORATION
NETWORK MANAGEMENT TRANSITION COMPLIANCE PLAN

COMCAST CORPORATION
NETWORK MANAGEMENT TRANSITION COMPLIANCE PLAN

1. **New Network Management Practices.** Comcast is preparing to transition to new, protocol-agnostic practices for managing congestion on our High-Speed Internet ("HSI") network ("congestion management"). We will complete that transition across our HSI network by December 31, 2008. We provide more details about these new practices, and detailed information about some of the hardware and software referenced in this document, in Attachment B.

2. **Trials.** Comcast is currently performing technical trials of the new congestion management practices in the following communities: Chambersburg, PA; Warrenton, VA; Lake City, FL; East Orange, FL; and Colorado Springs, CO. If Comcast management deems it necessary to conduct additional trials, they will be announced on Comcast's Network Management Policy page, located at http://www.comcast.net/networkmanagement/.

3. **Benchmarks.** Comcast expects to meet the following benchmarks in our transition to the new protocol-agnostic congestion management practices:

    a. **October 15, 2008.** Comcast will have completed installation of the PacketCable Multimedia and Internet Protocol Detail Record servers, and will have begun installation of the Congestion Management Fairshare servers. These servers, and other hardware used for the new congestion management practices, are described in detail in Attachment B.

    b. **November 15, 2008.** Comcast will have begun commercial (i.e., not trial) "cut-overs" to the new congestion management practices on a market-by-market basis. Once the equipment is in place in a particular area, this involves Comcast installing a software update to our customers' cable modems in that area, launching the software for the new protocol-agnostic congestion management practices in that area, and disabling the current congestion management techniques in that area.

    c. **December 31, 2008.** Comcast will have completed the deployment of all hardware and software needed to implement our new congestion management practices, and will have completed the "cut-overs" to the new, protocol-agnostic congestion management practices. We will also have discontinued the protocol-specific congestion management practices throughout our network.

    d. **January 5, 2009.** Comcast will report to the FCC that we have discontinued our protocol-specific congestion management practices throughout our network, and that we have completed transitioning to the new congestion management practices.

4. **Information Sharing.** Comcast will take the following steps to provide timely information to our customers about the transition to our new congestion management practices. We intend for our disclosures to be clear, concise, and useful to the average consumer.

a. **Congestion Management Trials.** Comcast already provides information about the trials of our new congestion management practices on our Network Management Policy page. Information about any additional trials will be posted there.

b. **Revision of Acceptable Use Policy.** Comcast will take the following two steps with regard to revising our Acceptable Use Policy ("AUP").

   i. Comcast will revise our AUP to explain that our network congestion management practices may include temporarily lowering the priority of traffic for users who are the top contributors to current network congestion. This new AUP will be published on October 1, 2008.

   ii. By January 1, 2009, Comcast will publish an amended AUP to reflect the discontinuation of the current protocol-specific congestion management practices, as well as any other necessary and appropriate updates.

c. **Customer Disclosures.** Comcast will take the following steps to inform our customers of the new congestion management practices.

   i. Attachment B, detailing Comcast's planned network management practices, as filed with the Commission on September 19, 2008, will be posted by midnight on that date to Comcast's Network Management Policy web page.

   ii. Comcast will, by midnight on September 19, 2008, provide new Frequently Asked Questions that explain these developments clearly, and will continue to post on our Network Management Policy web page updated information about the new congestion management practices.

   iii. At least two weeks prior to the first commercial (i.e., not trial) deployment of the new congestion management practices, Comcast will send e-mail notifications to the primary Comcast.net e-mail address associated with each customer regarding the new congestion management practices, informing them of the AUP revisions, and directing them to Comcast's Network Management Policy page for FAQs and other information. These developments will be further publicized through announcements at http://www.comcast.net.

d. **Customer Support.** Comcast will also answer customer questions on our Customer Support Forums page, located at http://forums.comcast.net/, which is available to all Comcast HSI customers. A link from the Network Management Policy page to the Customer Support Forums will also be provided.

5. **Management Responsibility.** The transition to these new practices and the discontinuation of the old practices is a high-priority effort. The project is being led and overseen at a senior executive level. The actual engineering and operations work is a joint project of the Office of the Chief Technology Officer and National Engineering & Technical Operations. In addition, regular customer communications and messaging are overseen by the company's Online Services business unit representatives.

6. **Employee Training.** Educational materials about the new protocol-agnostic practices are being developed for broad distribution throughout the relevant business units in Comcast. All affected employees in those business units will receive appropriate training about Comcast's transition to the new protocol-agnostic congestion management practices. Detailed technical customer inquiries about the new practices will be directed to the representatives in the Online Services business unit who will be trained to deal with such questions.

7. **FCC Notification of Material Changes.** Comcast will make supplementary filings with the Commission as necessary to keep the FCC (and the public) informed of any material changes in our plans before the transition to protocol-agnostic congestion management is completed at year-end.