

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):

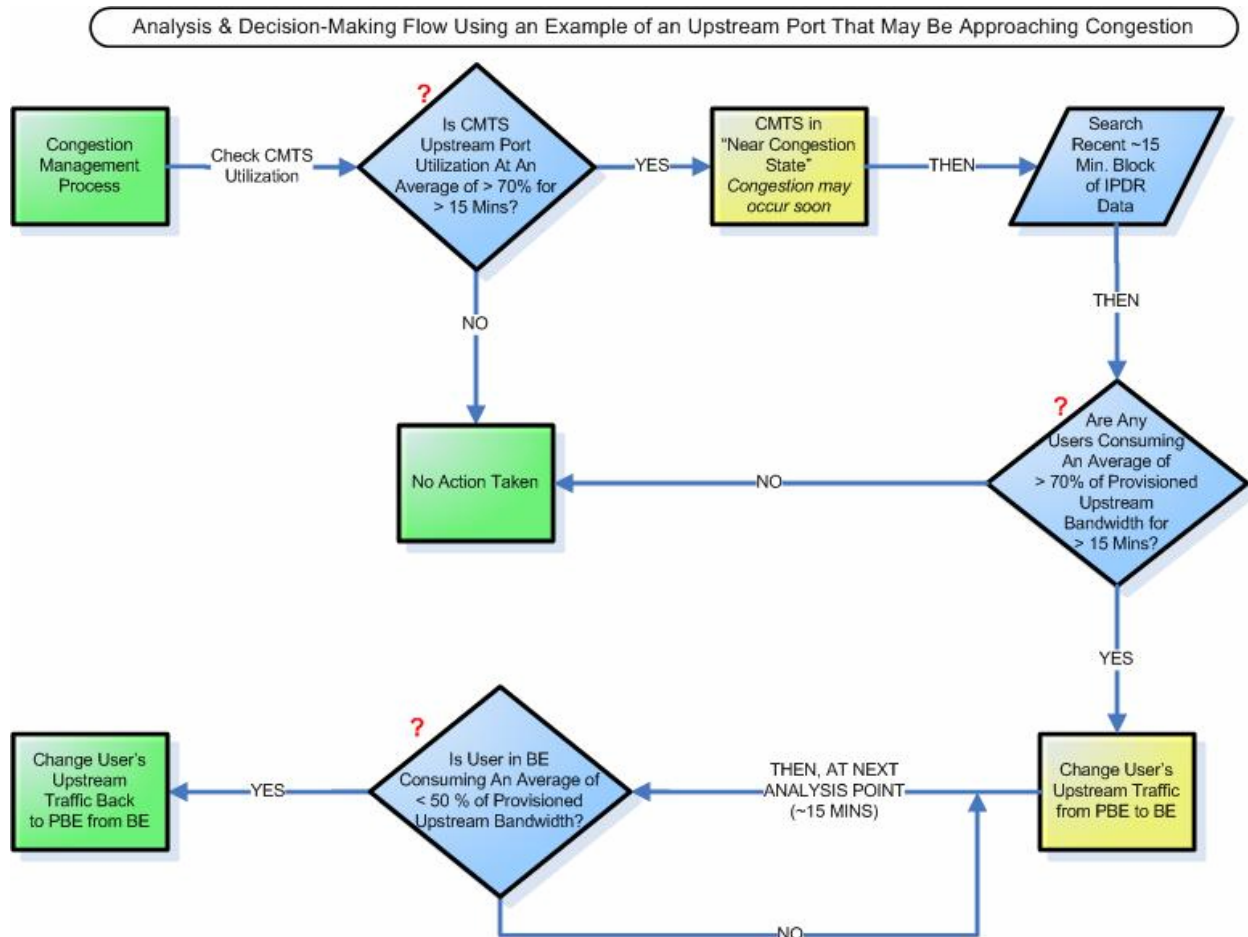


Diagram 2: Upstream Congestion Management Decision Flowchart

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:

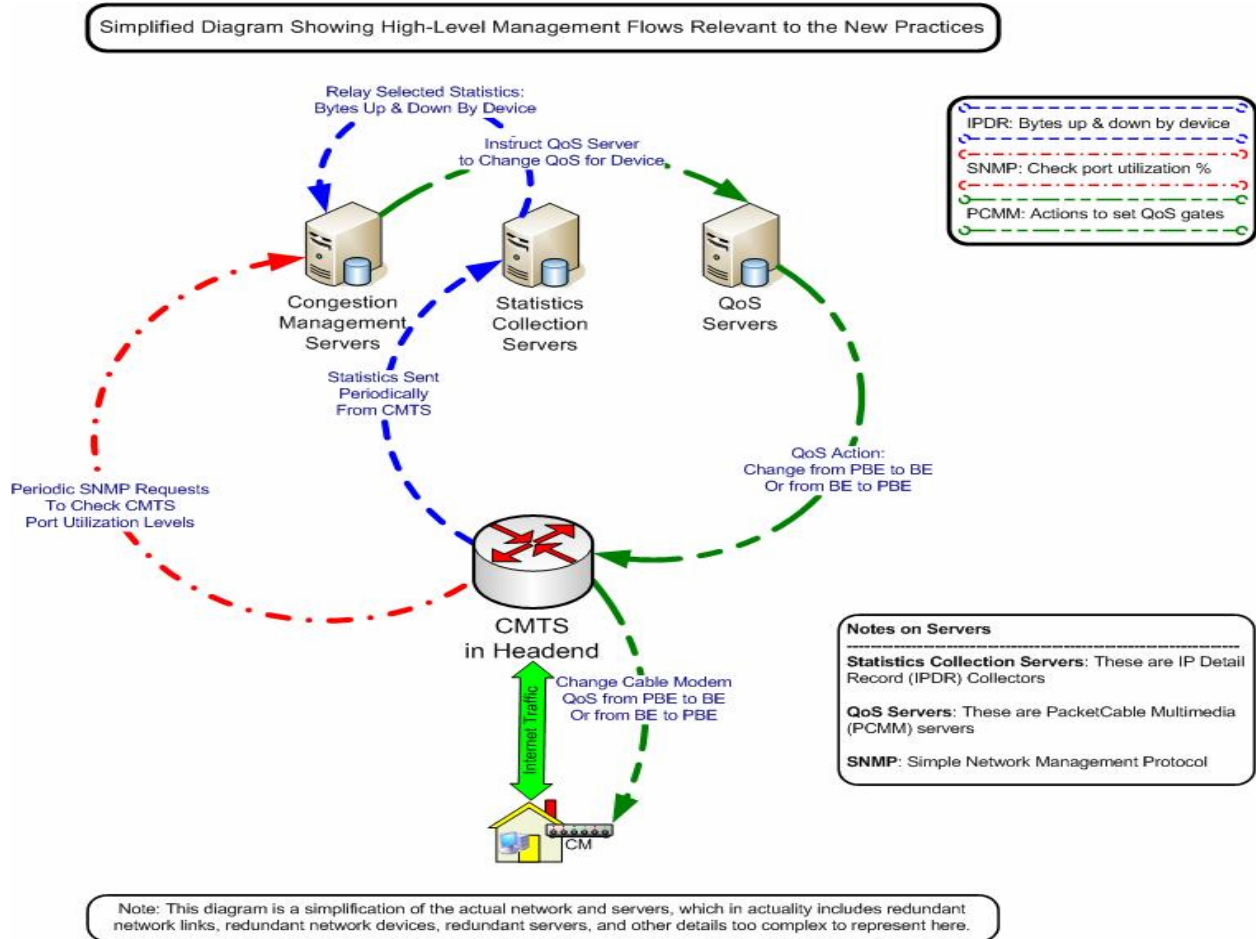


Diagram 3: High Level Management Flows

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.