BEFORE THE

Federal Communications Commission

WASHINGTON, D.C.

In the Matter of	)	
	)	
Preserving the Open Internet	)	GN Dkt. No. 09-191
	)	
Broadband Industry Practices	)	WC Dkt. No. 07-52
	)	
	)	
	)	

REPLY COMMENTS OF TW TELECOM

Willkie Farr & Gallagher LLP
1875 K Street, N.W.
Washington, D.C. 20006
(202) 303-1000

ATTORNEYS FOR TW TELECOM INC.

April 26, 2010

TABLE OF CONTENTS

Page

I. INTRODUCTION AND SUMMARY 1

II. MANAGED NETWORKS ENABLE INNOVATIVE APPLICATIONS WHICH WOULD OTHERWISE NOT BE POSSIBLE 3

III. MANAGED NETWORK SERVICES REDUCE COSTS AND PRICES FOR ALL USERS OF THE NETWORK 11

IV. THE FCC SHOULD ESTABLISH AS A PUBLIC POLICY GOAL THE PROVISION OF MANAGED SERVICES ACROSS MULTIPLE NETWORKS 15

V. CONCLUSION 19

BEFORE THE

Federal Communications Commission

WASHINGTON, D.C.

In the Matter of	)	
	)	
Preserving the Open Internet	)	GN Dkt. No. 09-191
	)	
Broadband Industry Practices	)	WC Dkt. No. 07-52
	)	
	)	
	)	
	)	

COMMENTS OF TW TELECOM

tw telecom inc. ("TWTC"), by its attorneys, hereby files these reply comments in response to the notice of proposed rulemaking in the above-referenced proceedings.[1]

I. INTRODUCTION AND SUMMARY

TWTC applauds the Commission's decision to commence this proceeding, and TWTC supports the Commission's overarching objective of ensuring that the Internet remains an open platform for innovation and civic engagement. TWTC is concerned, however, that the rules adopted in this proceeding could have an unintended negative effect on enterprise customers. In particular, TWTC is concerned that the rules adopted in this proceeding could inadvertently harm the development of managed services that are so critical to enterprise customers and their ability to compete in the global marketplace. This is not an issue that has received adequate attention in the comments. Accordingly, in these reply comments, TWTC describes the manner in which managed services, and in particular, those used by enterprise customers, are provided today, the efficiencies

[1] *See Preserving the Open Internet, Broadband Industry Practices*, Notice of Proposed Rulemaking, 24 FCC Rcd 13064 (2009) ("*NPRM*").

yielded by the provision of managed services, and emerging developments in the enterprise managed services marketplace. TWTC's hope is that this information will enable the Commission to adopt rules in this proceeding that do not harm the present or future provision of managed services for enterprise customers.

There are two basic factors that the Commission should account for in addressing managed services utilized by enterprise customers. First, firms that provide broadband Internet access service ("BIAS") in conjunction with applications, services and content for which enterprise customers are sensitive to latency and delay must rely extensively on network management practices. For example, in its network "core" that runs between and among its multiplexers in incumbent LEC central offices and its routers and soft switches, TWTC carries both "best-effort" BIAS traffic and voice, video and other traffic for which enterprise customers are highly sensitive to latency and packet delay. In order to efficiently manage the capacity of its core network and meet its customers' needs, TWTC must rely on network management techniques. This means, for instance, that TWTC must give higher priority to video and VoIP than e-mail and web access packets. This is because customers perceive VoIP and video service degradation if the packets are not transmitted on time, but they are less sensitive to delay in the delivery of e-mail and web access. Indeed, if the network is engineered correctly, granting priority treatment to video and voice packets should not cause a customer to notice any perceptible degradation in e-mail or web access service quality.

If TWTC and other firms that serve enterprise customers were unable to rely on this kind of traffic management, they would be forced to deliver all traffic by using the less efficient "best efforts" approach. In order to provide services for which customers

are sensitive to latency and delay in a "best-effort" environment, service providers would be required to add substantial capacity to their core and access networks to eliminate all bottlenecks. Such an approach is unlikely to be economically possible because it is very expensive for service providers to establish sufficient capacity to accommodate peak load capacity requirements without utilizing any traffic management. In any event, the increased expense required for this approach would increase the costs of providing service, and therefore the prices charged to all network users. Less innovation and higher prices would result.

Second, while the Commission's discussion of managed services in the NPRM seems to assume that managed services are provided solely over a single service provider's network, this is unlikely to continue to be true for long. Some network owners are agreeing to exchange traffic at specified levels of prioritization, and industry groups are beginning to standardize the way in which such traffic exchanges occur, facilitating further agreements. In any event, the public policy goal should be to encourage the development of industry standards for exchanging managed services between service providers because, as with other forms of traffic exchange, such standards will lower entry barriers, increase competition and enhance consumer welfare. The FCC should ensure that any net neutrality rules adopted in this proceeding do not undermine the development of such arrangements.

II. MANAGED NETWORKS ENABLE INNOVATIVE APPLICATIONS WHICH WOULD OTHERWISE NOT BE POSSIBLE

Managed services can be understood to fall into two broad categories: (1) applications, either offered by the BIAS provider or its customer, that cannot function at an optimum level in a "best-effort" environment and without appropriate network

prioritization; and (2) transmission services offered by BIAS providers that permit their customers to choose the level of prioritization that they believe is necessary for their own applications.[2] Firms that serve enterprise customers provide managed services to enterprise customers pursuant to service level agreements ("SLAs"). Under an SLA, the service provider generally agrees to deliver the customer's traffic according to quality of service ("QoS") parameters that support the applications demanded by or provided by the customer.[3] These QoS parameters address performance categories such as delay in the delivery of packets,[4] variation in such delay (i.e., jitter), packet loss, throughput and service availability.[5] BIAS providers often "package" these QoS parameters into different Classes of Service ("CoS").[6] Each "class" is assigned the QoS attributes

---

[2] For purposes of these comments, TWTC assumes that firms that provide managed services also provide best efforts BIAS over the same network, since it is the provision of those two categories of service over the same network that implicates the discrimination concerns addressed in the NPRM.

[3] Under many SLAs, a BIAS provider must provide service credits to customers for whom the BIAS provider fails to meet its QoS commitments.

[4] Delay, like many QoS parameters described here, consists of several subcategories, and SLAs sometimes establish separate requirements for one or more of these subcategories. Delay subcategories include propagation delay (the time taken for a single bit to travel from the output port on a route across a link to another router), switching delay (the time difference between receiving the packet on the incoming router interface and the queuing of the packet in the outbound interface), scheduling delay (the time difference between the queuing of the packet on the outbound scheduler and the start of clocking the packet into the outbound link) and serialization delay (the time taken to clock a packet into the link). *See* John Evans & Clarence Filsfils, *Deploying IP And MPLS QoS For Multiservice Networks: Theory and Practice*, at 4-6 (Morgan Kaufmann 2007) ("*Deploying QoS*").

[5] *See Deploying QoS, supra* note 4, at 2.

[6] Standards-setting organizations have attempted to standardize the different QoS parameters of various classes. *See, e.g., id*. at 241 ("[ITU Standard] G.1010 identifies four different categories for use traffic: interactive, responsive, timely and non-critical....[RFC4594] identifies twelve service class categories.").

necessary to support a particular application or set of applications.[7] Each traffic class is prioritized within the carrier's network so that traffic with higher CoS "labels" is given precedence over other, less time-sensitive classes. CoS provides the customer the type of guaranteed delivery of certain types of traffic (e.g., voice and video) that was possible in an environment where each service was delivered over a dedicated connection, but it achieves this goal by utilizing a network that is shared between CoS traffic and non-CoS traffic (i.e., "best-effort" traffic).

Network operators provide managed services to enterprise customers by relying extensively on QoS and CoS arrangements. This is true for both managed service applications offered by the BIAS provider itself as well as for managed service transmission offered to customers that provide their own applications. The most recognizable example of managed services applications offered via the BIAS provider's own network is video service. Video requires the transmission of data with very little delay (to send the signal from the customer remote control to the head-end and to begin streaming), high guaranteed bandwidth and low packet loss (to ensure picture quality), while "contending" with (but not degrading, at least from the customer's subjective point of view) other lower-priority internet access traffic delivered over the same connection.[8]

---

[7] *See id*. at 3 ("SLAs may be defined in absolute terms, e.g., a worst-case one-way delay of 100 ms, or may be defined statistically, e.g., a loss rate of 0.01%....How the loss rate is measured and calculated needs to be defined in order to understand what impact the 0.01% loss rate will have on the end applications; 1 lost packet in every one hundred packets may not have a significant impact on a VoIP call, but 10 consecutive packets dropped out of 1000 will cause a glitch in the call that is audible to the end-user.").

[8] *See Deploying QoS, supra* note 4, at 24 ("Excessive packet loss or delay may make it difficult to support real-time applications although the precise threshold of 'excessive' depends on the particular application.").

A commercial video service with such tight QoS requirements would simply not work over a "best-effort" connection.[9]

While it is understandable that the NPRM focuses on mass market managed applications provided over a BIAS provider's own networks[10], it is important that the Commission also focus on the video and other data services provided to enterprise customers in this manner. For example, as explained in more detail below, BIAS providers serving the business market increasingly provide video and sophisticated data services to enterprise customers.

In addition, TWTC and other firms that serve enterprise customers provide "business-class" VoIP offerings over their own networks. VoIP services provided over the "best-effort" Internet are generally not reliable enough to meet the needs of enterprise customers.[11] Therefore, network operators provide QoS to VoIP packets sent within the service provider's network to achieve the required level of service.[12]

---

[9] *See id*. at 56 ("[N]etworks supporting streaming video should typically be designed for very close to zero percent packet loss, and hence are designed to be congestionless from the perspective of the video traffic. This means that the available capacity for video traffic must be able to cope with the peak of the offered video traffic load. Further, in order to ensure that the service availability is maintained, this peak load must be able to be supported without loss while maintaining the required delay and jitter bounds for the video traffic."); SureWest Comments at 47 ("Multichannel video IPTV services require guarantees levels of latency and quality over loss of packets. Placing services such as IPTV on the managed portion of the network can help insure a high-quality HDTV service with the ability for near instantaneous channel change….*[M]anaging these services separately do[es] not detract from the expected quality of service or experience on the internet access portion of the network*.") (emphasis added). Unless otherwise noted, all comments referenced herein were filed in GN Dkt. No. 09-191 & WC Dkt. No. 07-52 on or around January 14, 2010.

[10] *See NPRM* ¶ 150 (describing AT&T's U-verse service where the bandwidth is dynamically allocated between AT&T's video offering and broadband internet access).

[11] *See Deploying QoS*, *supra* note 4, at 34 ("If congestion occurs, it is not practically possible to engineer [a best effort] network to ensure that consecutive packets from a

Enterprise voice service bundled with Internet access is also offered as a "managed service." For example, TWTC, like many other firms that serve business customers, offers a managed channelized voice/internet access product (*e.g.*, TWTC's VersiPak) in which the network dynamically allocates bandwidth between Internet access traffic and voice traffic as the need arises.[13] As the Fiber-to-the-Home ("FTTH") Council explains, in these types of channelized services "voice traffic is usually given priority over broadband Internet access traffic to avoid problems for voice traffic caused by bursts in Internet transmission. If voice services were susceptible to congestion caused by such bursts, calls could be rendered unintelligible or be dropped."[14] Unlike traditional channelized circuits, in which a particular number of 64 kbps channels are dedicated to circuit-switched voice service and cannot be reallocated to internet access, managed

---

single VoIP call are not dropped….For this reason, networks supporting VoIP are designed to be congestionless from the perspective of the VoIP traffic; that is. the available capacity for VoIP traffic is able to cope with the peak of the offered VoIP traffic load. QOS mechanisms, admission control techniques and appropriate capacity planning techniques are deployed to ensure that no packets are lost due to congestion.").

[12] *See* XO Comments at 16 ("[W]hile Internet-based VoIP services may offer many consumers adequate service quality for their typical personal use, enterprise and other institutional customers often demand a much higher level of reliability, preferring a dedicated connection with service level assurances. While consumers may tolerate occasional delays and diminished voice quality in exchange for VoIP services that are often free or very low cost, enterprise and government customers typically will not.").

[13] *See* tw telecom, VersiPak, at 2 (2008), *available at* http://www.twtelecom.com/Documents/Resources/PDF/MarketingCollateral/4101VersiPakVRS.pdf ("Dynamic Bandwidth Allocation: when voice channels aren't being used, that bandwidth is dynamically allocated to Internet access.").

[14] FTTH Council Comments at 10.

channelized services allow the customer to utilize the entire capacity of the circuit for internet access when no calls are being sent or received.[15]

In addition to providing applications themselves, service providers that serve enterprise customers provide application-agnostic managed services (e.g., bare transmission with QoS parameters) to enterprise customers. Enterprise customers use these transmission services to utilize selected applications. For example, a business could purchase circuits with high QoS guarantees to transmit private video signals between and among its different office locations or to utilize cloud computing applications. In neither case does the service provider provide the content or application. Rather, the video signals would be generated by the customer and the third-party cloud computing application might be installed and configured on the customer's network by its IT personnel. Only the bare transmission with the necessary QoS to enable the video and cloud computing services would be provisioned by the service provider.

Switched Ethernet services often serve as the "bare transmission" access inputs for enterprise customers' managed networks and are sold with specified CoS and QoS attributes to meet particular customers' prioritization needs. For example, AT&T offers three classes of service for its OPT-E-MAN switched Ethernet service: "Best Effort," "Bronze" and/or "Silver." The customer pays according to the class of service

---

[15] *See* Covad Communications, *Covad Service Customer Policies*, at 53 (last visited Apr. 21, 2010), *available* at http://support.covad.net/onlinesupportcenter/resources/legal/docs/CustomerPolicies_Direct.pdf ("Integrated Access Service is delivered over Covad's Voice Optimized Access…T1 service, which dynamically allocates bandwidth for voice and ensures that voice traffic takes precedence over data traffic. The IA service supports up to a maximum of 12 simultaneous calls, which leaves 288 kpbs for Customer's data traffic. *If fewer people are using the phone service, more bandwidth automatically becomes available for data transmission*.") (emphasis added).

provided.[16] According to AT&T, "Best Effort" is "[i]ntended for non-critical data applications with more tolerance for delay and/or data applications that are lower in priority…," while "Silver" "[s]upports applications that require minimal loss and low [jitter]."[17] AT&T provides service credits if the relevant QoS parameters for each CoS are not met.[18] Similarly, Qwest offers four classes of service for its Qwest Metro Optical Ethernet service,[19] and it charges its customers according to the CoS provided.[20] In all of these situations, the purchaser of the Ethernet service can utilize the Ethernet transmission as a means of utilizing its preferred managed service applications.

Finally, service providers also often offer the two different kinds of managed services (the QoS-enabled application and the QoS-enabled bare transmission) in

---

[16] AT&T, Inc., *AT&T Interstate Guidebook***,** Part 6 - Special Access Services - Midwest, § 24 - OPT-E-MAN, 1st Revised Sheet 2 (effective Mar. 14, 2008), *available at* http://cpr.bellsouth.com/guidebook/is/0006-0024.pdf ("*OPT-E-MAN Guidebook*").

[17] *Id*. at 1st Revised Sheet 3.

[18] *Id.* at Original Sheets 6-7, Revised Sheets 8-10.

[19] Qwest Corp., *Rates and Service Schedule Interstate No. 1*, § 8.8.2, at 270 (effective Apr. 15, 2009), *available at* http://tariffs.qwest.com:8000/idc/groups/public/documents/rss/qc_rss1_s008p261.pdf#Page=1&PageMode=bookmarks ("Quality of Service for QMOE allows customers to prioritize their traffic in four classes of service. The four different classes of service or levels are Priority 1, Priority 2, Priority 3 and Priority 4….Priority 1 – This Quality of Service class of service supports VoIP (Voice over Internet Protocol) and other near real time applications. Priority 2 – This Quality of Service class of service supports interactive Video applications. Priority 3 - This Quality of Service class of service supports business data applications. Priority 4 - This Quality of Service is the default for all other traffic that is not defined in Priorities 1, 2 or 3 above. Priority 4 traffic will have the lowest priority on the network.").

[20] *Id*. § 8.8, at 292 (effective Apr. 8, 2010) and 292.1 (effective Dec. 16, 2009), *available at* http://tariffs.qwest.com:8000/idc/groups/public/documents/rss/qc_rss1_s008p281.pdf#Page=1&PageMode=bookmarks.

combination to enterprise customers. For example, virtual private networks ("VPN") are provided in this manner where the carrier provides a variety of managed services to a multi-location customer.[21] To ensure that the each type of traffic (e.g., voice, video, data, Internet access) is transmitted optimally without perceptible delay, VPN customers' traffic is typically assigned a different CoS with particular QoS attributes.[22] For example, TWTC offers five classes which can either utilize "off-the-shelf" QoS parameters or QoS parameters selected by the customer.[23] Similarly, XO's IP VPN service is offered in four

---

[21] TWTC describes many typical multi-location VPN deployments on its website. *See, e.g.*, Press Release, tw telecom, *Indiana Blood Center Taps tw telecom for Data, Internet and IP Telephony Services*, at 1 (Dec. 9, 2009) ("tw telecom...today announced that Indiana Blood Center has selected the company to provide advanced voice, data and Internet services to connect the organization's 11 locations located throughout Indiana. Indiana Blood Center will also utilize tw telecom's metro MPLS IP network for IP VPN telephony and intra-office communications."), *available at* http://www.twtelecom.com/Documents/Announcements/news/2009/Indiana_Blood_Center.pdf; Press Release, tw telecom, *Caltrol Telecom Selects tw telecom to Deliver Multi-state IP VPN Services*, at 1 (Oct. 30, 2009) ("As part of the multi-year agreement, tw telecom will connect Caltrol's headquarters in Las Vegas to tw telecom's metro fiber network and nationwide IP backbone, as well as provide managed IP VPN services to the company's eight additional locations in California, Nevada, Arizona and Hawaii."), *available at* http://www.twtelecom.com/Documents/Announcements/News/2009/CaltrolFINAL10_30_09.pdf.

[22] Yu-Wei Eric Sung *et al.*, *Modeling and Understanding End-to-End Class of Service Policies in Operations Networks*, SIGCOMM'09, 219, 219 (Barcelona, Spain, Aug. 17-21, 2009), *available at* http://ccr.sigcomm.org/online/files/p219.pdf ("CoS designs embed inherently complex policy decisions based on the described priorities of various applications, extent of bandwidth availability and cost considerations."); *id*. at 226 (describing a hypothetical VPN for four data classes where "C1 is the real-time class, designed for jitter and latency sensitive applications like voice and video. C2 is the premium class, designed for critical business applications such as database transactions. C3 is the bulk data class, designed for medium priority business applications such as e-mail and file transfers. C4 is the best-effort class, designed for the remaining background traffic.").

[23] *See* tw telecom, *IP VPN*, at 2 (Feb. 2009) ("With tw telecom's IP VPN Premium, you can manage multiple applications that are competing for the same network resources by prioritizing traffic with 5 service classes."), *available at*

different classes of service: "standard," "priority," "critical" and "real-time," each with its own QoS parameters.[24] These classes of services are offered with both BIAS provider-provisioned services and customer-provisioned services in mind. For example, XO's highest-level VPN CoS, "Real Time," is specifically advertised as having the appropriate QoS to enable VoIP (likely provided by XO) and Citrix (a cloud computing/desktop virtualization application offered by a third party).[25]

III. MANAGED NETWORK SERVICES REDUCE COSTS AND PRICES FOR ALL USERS OF THE NETWORK

Prioritizing and classifying traffic yields substantial efficiencies for enterprise customers because these techniques enable service providers to smooth out spikes in network traffic and decrease the peak load capacity requirements of the network. In sum, these techniques enable service networks to deliver more, and more sophisticated, services with less capacity than would be the case in a "best-effort" environment.[26]

---

http://www.twtelecom.com/Documents/Resources/PDF/MarketingCollateral/twtc_IP_VPN_2250.pdf; tw telecom, *Managed IP VPN Services*, at 1 (Jan. 21, 2009) ("Managed IP VPN provides customers with five optional service classes. This Class of Service (CoS) feature provides a mechanism for managing traffic in a network by logically grouping similar traffic profiles, or applications. Each group is defined within its own 'Service Class' and Service Level Agreement (SLA), related to end-to-end tolerance specifications based on packet delay, packet loss and jitter. *You may choose to utilize tw telecom's CoS Policy template or customize your CoS policies based on your own unique applications and requirements—you decide with tw telecom's Managed IP VPN service*.") (emphasis added), *available at* http://www.twtelecom.com/Documents/Resources/PDF/MarketingCollateral/TW_22501bManagedIPVPN.pdf.

[24] *See* XO Communications, *The XO Evolution*, at 5 (2007), *available at* http://www.comtechinc.biz/Brochures/XO%20IP-VPN.pdf.

[25] *See id*.

[26] Jon M. Peha, *The Benefits and Risks of Mandating Network Neutrality and the Quest for a Balanced Policy*, 1 INT'L J. OF COMM. 644, 651 (2007), *available at*

Perhaps the best way to describe the efficiencies of traffic management for enterprise customers is by illustration. Video conferencing services is an especially telling example. Until recently, standard-definition video conferencing service was generally provided over private, single-purpose network connections (e.g., dedicated ISDN circuits running between two business locations).[27] A service provider did not utilize QoS and CoS for such dedicated connections because the video service traffic offered over dedicated facilities does not have to "contend" with other traffic.[28] Instead, the customer would be required to purchase a separate transmission service to transmit its non-video conferencing traffic. But providing video conferencing via dedicated ISDN connections is extremely inefficient because the capacity of such circuits is unused when a customer is not holding video conferences, usually a significant amount of time. Moreover, connecting multiple locations with dedicated connections (e.g., providing the capability for all 10 of a company's locations to video-conference with each other) might become economically infeasible.

---

http://ijoc.org/ojs/index.php/ijoc/article/viewFile/154/90 ("Discrimination with respect to QoS is also important because different applications have different QOS needs. In a VoIP application…packets must be received within 50 ms because any arriving after 50 ms are useless….If sophisticated traffic control algorithms take these QOS requirements into consideration, it is possible to give packets high priority when and only when they need high priority to meet QOS requirements, thereby meeting QOS requirements for many more users on a given network. Alternatively, it is possible to serve the same number of customers at the same QOS with less network capacity, making the network less costly. This benefits Internet users and network operators.").

[27] *See generally* Ira M. Weinstein, *The ISDN to IP Migration for Videoconferencing: Real World Options That Make Both Dollars and Sense,* Wainhouse Research (Jan. 2006) *available* at http://www.tandberg.com/collateral/white_papers/whitepaper_ISDN_to_IP_migration_for_videoconferencing.pdf.

[28] *See id.* at 2 ("In an ISDN videoconferencing environment the video traffic does not touch the organization's data network.").

In contrast, when a service provider enables a customer to conduct video conferencing on a QoS-enabled network, the capacity of the network is used far more efficiently. In this scenario, the video conferencing traffic is transmitted on the same network connections as other traffic (e.g., intracompany e-mail) for which customers are less sensitive to packet delivery delay. The video traffic is given priority, but other services are nevertheless normally delivered in a satisfactory manner.[29] Thus, when a two-way video conference is in session, e-mail transmission times for users on the network might slow down, but, if the network is engineered properly, reduced e-mail performance will not be perceptible to the user. When the video session is over, the bandwidth can be reallocated to lower priority internet and data traffic. Because the video conferencing traffic is routed through the carrier's core network, not via individual ISDN connections between offices, only 10 access connections into the network core are necessary to enable all 10 offices to video conference with each other.

As the video conferencing example shows, QoS lowers costs and prices because (1) the customer no longer needs to purchase two separate networks--one for video traffic and one for other traffic, and (2) traffic prioritization lowers the peak network load of the unified network, limiting the total access capacity the customer must purchase to meet its overall network needs.[30] In addition, because the network resources dedicated to each customer are more efficiently utilized, the carrier needs to provision less capacity in its

---

[29] *See Deploying QoS*, *supra* note 4, at 114 ("Typically, most basic IP packet schedulers available today support a queue services with a priority scheduler for delay intolerant traffic such as voice and video.").

[30] *See* Ira M. Weinstein, *supra* note 27, at 5 ("IP-only endpoints are less expensive to purchase…, cheaper to keep under a service plan (fees are based on purchase price), and do not require dedicated ISDN lines, resulting in a lower total cost of ownership.").

core network, creating additional savings for the service provider that are passed on to all of its customers.

These efficiencies are not possible if a service provider must rely solely or predominantly on adding additional capacity to its "best-effort" network. Such a "brute force" approach will increase the network operator's costs and the rates of all users of the network because the operator will need to increase the network's overall peak load capacity. Indeed, for a BIAS provider to ensure that voice, video and data applications function adequately on a "best-effort" basis, the service provider would have to eliminate the possibility of bottlenecks at *all* points in the network at *all* times. This would be unrealistically costly to implement.[31] As the Telecommunications Industry Association ("TIA") explains, "[t]he goal of a Managed Service offering is... to guarantee the level of service at all times, whereas the Best Effort paradigm may deliver comparable performance, but not at all times, i.e., not under all traffic load conditions or not for all applications simultaneously."[32]

---

[31] *See Deploying QoS*, *supra* note 4, at 87-88 ("If it is possible to ensure that there is always significantly more capacity available than there is traffic load then delay, jitter and loss will be minimized, throughput will be maximized, and the service requirements will be easy to meet. In practice however, ensuring that the network is always overprovisioned relative to the actual traffic load is not always cost-effective.").

[32] *See* TIA Comments, Managed Services Declaration of Marcus Weldon, at 7; *Deploying QoS*, *supra* note 4, at 211-212 ("Without such QoS mechanisms…, in order to support services and applications with tightly bound SLA requirements, the available capacity needs to be over-provisioned relative to the peak of the aggregate offered load….[Assuming that overprovisioning is not always going to be possible…, supporting real-time services with requirements for tight bounds on delay, jitter, and loss [ ] will be precluded in cases where congestion occurs.").

IV. THE FCC SHOULD ESTABLISH AS A PUBLIC POLICY GOAL THE PROVISION OF MANAGED SERVICES ACROSS MULTIPLE NETWORKS

In the NPRM, the FCC appears to assume that managed services are provided over individual providers' private networks whereas non-managed services are provided over the best-effort public Internet.[33] This public/private distinction still makes sense because most BIAS providers are only able to guarantee QoS parameters for packets which remain on the service provider's own network. BIAS providers generally cannot dictate how packets sent over other service providers' networks or over the Internet will be treated.[34] But the distinction between wholly private managed networks and the public Internet is becoming increasingly blurred, and the FCC should work to promote this development.

As many commenters and experts have noted, BIAS providers are beginning to agree to exchange QoS and CoS parameters *across* networks through private peering

---

[33] *See NPRM* ¶ 108 (noting that managed services "may be provided to end users over the same facilities as broadband internet access service, but may not themselves be an Internet access service….These services may require enhanced quality of service to work well. As these may not be 'broadband Internet access services,' none of the principles we propose would necessarily or automatically apply to these services.").

[34] For example, XO defines its "network" for the purpose of SLAs for its IP VPN service as "all XO-owned and controlled routers and circuits used to transport IP VPN traffic." XO Communications, *Service Level Agreements (SLA's) and Associated Credits; IP Virtual Private Network (IP VPN)*, at 1 (rev. Oct. 5, 2009), *at* http://www.xo.com/SiteCollectionDocuments/information/TOS_SLA_Rates/sla/IP_VPN_SLA.pdf. Similarly, XO offers service guarantees for its private line services only if "both ends of the Private Line are Connected to the XO Network." XO Communications, *Private Line Service Level Agreement (SLA)*, at 1 (rev. Oct. 13, 2006), *at* http://www.xo.com/SiteCollectionDocuments/information/TOS_SLA_Rates/sla/PrivateLineSLA.pdf.

arrangements.[35] Under these peering agreements, when one network delivers a packet with a particular level of prioritization to a second network, the second network can recognize and maintain the same level of prioritization until the packet reaches its ultimate destination. If these QoS peering agreements become more prevalent, and more and more networks agree to exchange QoS, the entire Internet may eventually become capable of supporting QoS sensitive managed services. Such an outcome would be pro-competitive and pro-consumer.[36] Ubiquitous QoS peering would allow smaller network

---

[35] *See* Comments of Verizon and Verizon Wireless, Attach. D: Declaration of Michael F. Ritter ¶ 14 ("Over time, additional integration of managed and 'Internet' content will continue, thus increasing the choices available to consumers, and further blurring the lines between private network services and stand-alone Internet access."); Cisco Comments at 16 ("Today, managed services are mostly offered over facilities entirely segregated from the 'public Internet.' Over time, however, it is likely that managed services…are likely to rely on customers' own Internet access links in the last mile, *to the extent those links can be provisions to ensure sufficient quality of service*.") (emphasis added); Hunter Newby, *(Internet Protocol + Ethernet Transport) x QoS - (The public Internet / Net Neutrality) = Enterprise Peering*, TMC.net, Dec. 6, 2007 ("No network goes everywhere all of the time and therefore no carrier can guarantee [QoS] in a dynamic off-net scenario, especially over the 'best efforts' Internet….Enterprise peering [ ] simply means there is a better way for Enterprise network operators to interconnect between their own sites, other enterprises, services providers and [access providers] directly, thus avoiding the Internet, but still using IP. This improves quality and reliability."), *at* http://www.tmcnet.com/viewette.aspx?u=http%3a%2f%2fwww.tmcnet.com%2fnews%2f2007%2f12%2f06%2f3146015.htm&kw=1*;* Shane Amante *et al.*, *Interprovider Quality of Service*, White Paper Draft 1.1, MIT Communications Futures Program, at 4 (Nov. 17, 2006) ("*MIT White Paper*") ("Some providers are now beginning to interconnect with each other via 'QoS-enabled peering' in an attempt to offer QoS that spans the networks of multiple providers."), *available at* http://cfp.mit.edu/docs/interprovider-qos-nov2006.pdf.

[36] There are two scenarios in which QoS and CoS traffic might traverse multiple networks: (1) where the traffic originates with a customer with a retail relationship with one service provider and is delivered to a customer with a relationship with a second service provider (*e.g.*, where a TWTC subscriber originates a VoIP call to an AT&T subscriber); and (2) where the traffic traverses a service provider's own network facilities as well as network facilities the service provider leases from another network owner as an input into a retail service offering (*e.g.*, where TWTC provides service to a multi-location business customer via loops TWTC owns and loops it leases from AT&T). Although

providers to offer innovative, QoS-sensitive managed services to locations not connected to their networks. As with other traffic exchange arrangements, the result would be that smaller competitors could benefit from the network effects of the combination of networks, thereby increasing competition and choices for consumers. For example, QoS peering would allow two locations to receive two-way HD video services even if each location is served by a different network operator. Indeed, TWTC has entered into an agreement with another carrier to exchange two-way HD video traffic by ensuring that the carriers' QoS parameters are maintained as the traffic traverses both networks. As a result, customers served by TWTC's network can now transmit and receive HD video to and from customers reached by TWTC's partner's network. Similar agreements are also being reached through membership in Ethernet "exchanges" such as CENX.[37]

These agreements should be encouraged, not discouraged, by FCC regulation.[38] Indeed, other regulatory bodies and industry groups have already begun to move forward to establish QoS and CoS standards and to develop the mechanics of transmitting these

---

standards for maintaining QoS and CoS across multiple networks are necessary in both scenarios, these comments address only scenario (1).

[37] *See* CENX, *Interconnect* ("With a physical connection in place, CENX members can establish virtual service-level interconnections with all other CENX service provider members connected to CENX System. The CENX System automatically aligns and manages interconnected CENX members' Carrier Ethernet services for smooth and seamless operations across all CENX's service provider members' networks."), *at* http://cenx.com/services/interconnect (last visited Apr. 25, 2010).

[38] If anything, the FCC should be on guard to intervene against Tier 1 network operators who unreasonably deny the exchange of QoS parameters in order to restrict the addressable market of smaller rivals. Indeed, the FCC should monitor the market for CoS and QoS interconnection closely to ensure that market forces prevent larger networks from discriminating against or refusing to interconnect and exchange CoS and QoS with smaller networks.

standards across all networks.[39] For example, Level 3, Verizon, Cisco, Nortel and MIT, among others, co-authored a recent white paper which "presents a proposal to enable the deployment of Inter-provider Quality of Service" with a particular focus on VPN interconnection.[40] Moreover, a report compiled for the European Community in 2008 regarding IP interconnection specifically recognized the MIT white paper as an important step in implementing end-to-end QoS across multiple providers.[41] Similarly, in 2006, the UK Regulator OfCom proposed establishing a new advisory body on next generation networks. Among other things, the advisory body would focus on "a reference architecture [ ] required for IP interconnection, covering...service characteristics and required interoperability standards. This is particularly urgent in relation to IP-based voice interconnection, but any proposed architecture also needs to make a forward-

---

[39] *See generally MIT White Paper*, *supra* note 38; *see also Deploying QoS*, *supra* note 4, at 78 ("There are some current standards efforts , which are looking to formalize the definitions of inter-provider services"); NECA Comments at n.16 ("There are a number of industry standards, including those developed by the [ITU, ITF and IEEE] that define such methods…IETF RFC5160 'Provider-to-Provider Agreements for Internet-Scale Quality of Service', ITU E.800 to E.880 'Overall Network Operation -- QoS.' These international standards are continually evolving[.]").

[40] *MIT White Paper*, *supra* note 38, at 4-5.

[41] *See* J. Scott Marcus & Dieter Elixmann, *The Future of IP Interconnection: Technical, Economic, and Public Policy Aspects*, WIK-Consult, Final Report, at 87-88 (Jan. 29, 2008) ("There has been substantially more positive movement in regard to the development of practical QoS metrics for use between providers, and for implementing measurement regimes. Particularly noteworthy in this regard is the *Inter-provider Quality of Service* white paper developed by the Communication Futures Program (CFP) at the Massachusetts Institute of Technology (MIT)….By providing standard service classes and measurement methodologies, the MIT paper potentially represents a significant advance in achieving inter-provider QoS."), *available at* http://ec.europa.eu/information_society/policy/ecomm/doc/library/ext_studies/future_ip_intercon/ip_intercon_study_final.pdf.

looking view of the requirements of new multimedia services."[42] OfCom understood that "[t]here is a need to understand the types of network intelligence which need to be exchanged between operators of NGNs, the commercial basis for such exchange , and technical interoperability issues."[43]

In sum, there is a widespread recognition that establishing CoS and QoS across different networks is an important goal for the promotion of competition in the provision of managed services. It follows that the FCC should seek to advance this objective. In all events, the FCC should ensure that the rules it adopts in this proceeding in no way undermine the development of such inter-network arrangements.

V. CONCLUSION

The Commission should ensure that the rules adopted in this proceeding enable BIAS providers to provide the managed services to enterprise customers.

Respectfully submitted,

/s/Thomas Jones

[42] OfCom, *Next Generation Networks: Developing the regulatory Framework*, at 8 (Mar. 7, 2006), *available at* http://www.ofcom.org.uk/consult/condocs/nxgnfc/statement/ngnstatement.pdf.

[43] *Id.* at 9.

Thomas Jones
Jonathan Lechter
Willkie Farr & Gallagher LLP
1875 K Street, N.W.
Washington, D.C. 20006
(202) 303-1000

ATTORNEYS FOR TW TELECOM INC.

April 26, 2010