

Appendix B

DECLARATION OF VIJAY GILL

**Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, D.C. 20554**

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

DECLARATION OF VIJAY GILL

I. INTRODUCTION AND BACKGROUND

A. Qualifications

1. My name is Vijay Gill, and I am the manager of Engineering and Global Network Architecture at Google Inc. (“Google”). In this capacity, I am responsible for all network functions at Google, including Optical, Transport, Internet Protocol (“IP”), Multiprotocol Label Switching (MPLS), Datacenter connectivity and Internet related protocol design. My curriculum vitae is attached to this declaration as an Appendix.

B. Overview and Summary of Conclusions

2. The FCC has proposed draft rules in the above-referenced proceedings for ensuring the continuation of a “free and open Internet” and has requested comment on its proposals (the “open broadband” rules). I have been asked by counsel for Google to provide my technical and engineering expertise on certain issues raised by the FCC’s proposed rules and the comments filed in the record of the above-referenced FCC dockets.

3. There are four fundamental points that are critical to understanding the unique role of last-mile broadband providers, and the potential impact of the FCC's proposed broadband openness rules on broadband providers' abilities to manage their networks.
4. First, last-mile broadband access providers are uniquely positioned as a technical matter vis-à-vis all other entities connected to the Internet. Last-mile broadband access providers' networks act as the on/off ramps for Internet traffic, so that every packet of Internet traffic must traverse the networks and devices under their control. Because they own and operate physical last-mile networks, including the routers closest to the end user customer, broadband providers are able to inspect, act upon, and apportion capacity for all online traffic – including third party Internet traffic -- that traverses their networks. By contrast, other entities on the Internet, including applications and content providers, can view and interact with only their own traffic.
5. Second, using the network router to prioritize certain data packets over others inherently results in degradation of other data packets traversing that same router.
6. Third, the proposed rules are fully consistent with ensuring that broadband providers can engage in reasonable and legitimate management of broadband networks. Nothing about managing networks to alleviate legitimate congestion and malware concerns is inconsistent with broadband openness requirements. Further, network congestion and malware issues can be addressed without permitting broadband network providers to exercise unfettered discretion and control over online traffic traversing their networks.

7. Fourth, the positive and continuing evolution of broadband infrastructure and technological progress of networks is fully consistent with the proposed FCC rules.

II. LEGITIMATE TECHNICAL AND ENGINEERING CONCERNS SUPPORT ADOPTION OF THE PROPOSED OPEN BROADBAND RULES

A. The modular, end-to-end nature of the Internet has important technical implications.

8. For decades, network engineers have employed software-based protocols, or standardized rules, in order to separate out distinct functions in data networks. Because it is difficult and highly undesirable to write a single protocol to handle every operation in a network, multiple protocols are used to partition a communications problem into discrete modules that handle each sub-problem. Functions are allocated to different protocol levels, or layers, with standardized connective interfaces operating between layers.

9. Layering plays a central role in modern-day data networks, because it allows changes to implementation at one layer without affecting others. Almost by definition, utilizing a layered protocol “stack” creates a high degree of modularity, which allows for ease of maintenance within the network, and facilitates communications between disparate networks. As a result, the use of distinct layers persists to this day for sound and enduring reasons of network engineering.

10. Since the early 1970s, engineers have developed various network design models incorporating protocols in a layered manner. Each of these models shares the same overall structure and philosophy of dividing up network tasks into

functional layers, stacked one on top of the other. While there are several generally accepted ways to divide up the protocol layers, for purposes of this declaration I will utilize a model consisting of four layers: physical, logical, applications, and content.

11. The IP suite, introduced in 1974, is the most famous and universally accepted such model. IP serves as the bearer protocol of the Internet, as well as many private data networks. One of the key virtues of IP is that it facilitates multiple layers riding on top of separate physical infrastructure. IP also is an “agnostic” protocol, in that it carries all packets indiscriminately. Some have described IP as the essential “waist” of the protocol stack “hourglass,” with the Internet’s applications and content layers above, and the physical communications network layers below.
12. The process of sending a packet over a data network involves a series of orchestrated ascents and descents through different layers. Notably, the data received from the applications at the upper layers of an IP network is broken into data packets to be handed to the TCP/IP layers; conversely, a data packet received from the TCP/IP layer is assembled into a data stream to be delivered to the upper layers. Lower layers treat data passed from upper layers as structureless payload, and place headers and/or trailers around that payload. Thus, in the resulting vertical hierarchy, a piece of content begins at the top layer and works its way down to the lower physical layer for transport to the ultimate destination, where it then ascends back to the top layer again.

13. The layered approach to engineering networks has some obvious and important implications. First, there is a clear separation between the upper layers (the applications and content) and the lower layers (the physical and logical, or IP, networks). Generally speaking, network providers operate in the physical and logical layers, while end users operate in the applications and content layers. Second, each layer in a network depends on the layers below it to transport content to its proper destination. For example, a piece of content cannot be transported anywhere in a data network without access to the physical and logical layers. This is because all the physical infrastructure necessary to carry traffic from one point to another – including the transmission lines and routing equipment – resides at the physical and logical layers. However, and importantly, the reverse is not the case; the physical networks do not rely on other layers in order to function properly. Third, the physical and logical (IP) layers have a unique ability to affect the behavior of traffic flowing over them, for example by managing congestion and assigning priority to certain packets.

B. Last mile broadband providers are uniquely positioned as a technical matter to transport, inspect, manipulate, and allocate capacity for all other entities' online traffic.

14. Given the realities of network engineering, providers of last mile broadband infrastructure to end user customers occupy a unique place of control. This control can be conceptualized as “horizontal” in nature, to correlate to the horizontal layering of the protocol stacks. This control stems from the ability to own and operate the essential physical and logical inputs to and from the end user customer, which cannot be replicated by any other entity operating in the upper

applications and content layers of the Internet. In particular, broadband providers own and control the last mile transmission facilities to the home, including the “last router” between the end user and the rest of the broadband provider’s local network.

15. As a result, broadband providers are positioned uniquely to transport, intercept, inspect, and manipulate every packet between the end user and any application or service on the Internet. This is very different from entities like application providers, which can only access packets destined to or originating from their own application layer services. For example, Application Provider A cannot intercept or interfere with packets that are meant for Application Provider B, whereas the broadband provider “sees” and can act on any application providers’ online traffic that transits its broadband network.
16. As such, it is useful to conceptualize at least four unique characteristics of a last-mile broadband provider that is routing Internet traffic and data packets between its end user customers and the rest of the Internet. First, the broadband provider is able to transport all Internet traffic and data packets to and from the end user customer. For any application or piece of content on the Internet to reach an end user customer, it must traverse the broadband provider’s network. This means the broadband provider can carry over its local network not just its “own” data packets, but in theory the traffic of each and every third party entity using the Internet.
17. Second, the broadband provider is able to intercept and inspect the contents of the data packets sent from the end user, as well as all other entities on the Internet to

the end user. This function can be most readily exercised using advanced router technologies such as deep packet inspection (DPI).

18. Third, the broadband provider is able to interact with other entities' data packets, including blocking, degrading, and/or prioritizing such traffic. This ability to manipulate all packets stems from the functionalities built into network routers, which can be set to determine unilaterally whether, when, and how Internet traffic is actually delivered.
19. Fourth, the broadband provider is able to determine which data traffic receives a share of existing broadband capacity. The broadband provider alone can decide how to apportion the total capacity of its last mile network, including what bandwidth to allocate to Internet access, versus what bandwidth to allocate to other specific services provided to end user customers (such as proprietary video and voice offerings).
20. Collectively, these four horizontal elements of last-mile physical network control amount to unique ability to control all traffic between the end user and the rest of the Internet; namely, to transport, to inspect, to manipulate, and to apportion capacity for all traffic flowing over the broadband pipe. This control inherently extends to the vast bulk of data packets owned by other parties. Thus, the broadband provider always has the final say over whether, how, and in what manner data packets make it to and from end users.
21. These four elements of horizontal control differ fundamentally from those of any other entity operating on the Internet. This includes independent content and applications providers, which (aside from their own transmission networks on the

other side of the Internet “cloud”) are limited to the upper layers of the protocol stack. Such horizontal control grants broadband providers a unique ability to use their technical and engineering place in the network to directly affect all Internet traffic as it flows across their last mile broadband transmission networks.

22. One example of a network functionality that falls far short of the control provided via last-mile broadband networks is the CDN, or content delivery network. Some of the functions of a CDN include reducing local network congestion, creating “burst” or overflow capacity events, and enhancing the end user’s experience by hosting and serving content from a location more proximate to end users. By definition, CDNs do not and cannot involve or interfere with other traffic flows to end users. Simply put, storing your own packets is not the same as routing someone else’s packets.
23. CDNs are able only to control what traffic runs through those particular content servers, and cannot otherwise affect the end user’s entire Internet experience. In technical terms, CDNs at best can improve the user experience by adding “burst” or “overflow” capacity (for example, involving “flash” events like large live concerts) and by reducing the latency with “forward caching” servers closer to the end user. By contrast, a last-mile broadband provider can use its local network of routers to affect every other Internet-based entity with which the end user exchanges traffic.
24. As a result, content and applications providers cannot control traffic and data packets beyond their own. An application service provider may elect to use a CDN to improve performance for its own users, but it is not a zero sum game. If

someone elects not to use Service Provider A, the fact that Service Provider A has elected to use a CDN to distribute Service Provider A's application makes no difference for that end user. The crucial point is that electing to use a CDN by any provider is immaterial for users who choose not to use that provider's application or service.

C. Router-based prioritization amounts to a zero-sum game, where other network traffic inherently is degraded.

25. It is important to understand the effects of using a network router to prioritize certain data packets. A router is simply a piece of equipment that sends and receives traffic within a data network. The router "reads" the headers on a data packet, and determines where next to relay it within the network. Obviously not all data packets can make it through an individual router at the same time. As a result, the router relies on a variety of technical considerations and operational policy that are programmed into its routing and forwarding logic. These considerations include the destination address, packet priority level, and via policy could also include minimum route delay, minimum route distance, route congestion level, and the "least-cost" route.
26. The routing table also can be modified to include "policy" considerations, such as prioritizing packets containing headers indicating priority treatment. Many refer to this as the router's "prioritization" of that particular data packet.
27. Because the movement of packets through a router is limited by the overall network capacity, it is intuitively clear that an Internet packet moved to the front of the line pushes back every other packet in the queue. Thus, favoring one class of traffic in a router inherently disfavors other classes.

28. In a shared network environment, then, prioritizing certain classes of data traffic creates greater delay and lower throughput for less favored traffic. Eventually, different classes of prioritized service can result in infinite delay and zero throughput for everybody else. When used solely as an even-handed network management practice intended to reduce congestion and latency, prioritization could be an acceptable means of managing data traffic. However, because prioritization of some inevitably also means degradation of others, the practice becomes far more problematic when employed for non-engineering-based reasons, such as commercial gain.
29. Some parties have analogized a third party paid prioritization as the equivalent of choosing to pay more for sending a package via FedEx, or choosing to save money by sending the same package via the regular U.S Postal mail service. This is not an apt analogy. By prioritizing certain traffic as a commercial matter, all other traffic will be slowed and, as prioritization grows, all other traffic could be degraded and negatively impact the end user's experience. Further, commercializing traffic prioritization for reasons not related to network engineering is likely to incent the last mile provider to emphasize its revenue-generating prioritization service, and to degrade the transmission of non-prioritized service. Finally, the analogy is inapt because FedEx competes with a number of other suppliers (e.g., UPS, DHL, etc.), but the last mile provider would price and operate its prioritization service under no constraints of a robustly competitive marketplace. One argument that has also been made is paying FedEx for various delivery options on a package basis – sending a package next

business day or sending it 2nd business day. FedEx will not destroy the package if sent 2nd business day. It will get there on the 2nd business day, and if they are out of capacity to deliver, the customer has an option of going to another competitive provider. With last-mile networks, not only can my packets be destroyed, I have no option of going to a second or third carrier.

D. The proposed open broadband rules are consistent with reasonable and legitimate network management.

30. Some parties assert that it will not be possible for broadband providers to manage their networks if the FCC adopts the proposed open broadband rules. In my view, this argument is incorrect as a technical and engineering matter. Network management is an important and necessary engineering imperative. However, the open broadband rules as proposed still will allow legitimate and reasonable broadband network management.
31. Verizon asserts (Verizon Comments at 81) that it is critical for broadband providers to respond to an increasing number of complex issues while still maintaining a quality of service required by consumers. However, claims that the network security functions would be hamstrung and could cause broadband providers to “target their responses too narrowly, to the benefit of terrorists and hackers” fails to distinguish between short-term responses to a real-time situation, and a systemic degradation of service.
32. Further, while Verizon states (Verizon Comments at 40) that the concept of a dumb pipe is mythical and that networks are the enabling technology for the Internet, this is not inconsistent with narrowly-tailored network management subject to FCC oversight. Network management policies are not at odds with an

“intelligent network.” Significant technical innovation has been created with so-called “dumb pipes” -- including the entire Internet ecosystem and companies that rely on the Internet. The greatest technical and economic value creations seen so far in terms of the user experience -- such as Yahoo, Ebay, Google, Dell Online, Amazon.com, and Cisco’s B2B portals -- have installed intelligence at the edges and treated broadband pipes largely as transport. For practical purposes, the Internet is an end-to-end, transparent transport mechanism, much like the interstate highway system.

33. Arguments that prioritizing data to mitigate jitter and latency does not discriminate against non-prioritized data (*see, e.g., AT&T Comments Exh. 1, at 18*) are erroneous. As a technical matter, when certain packets are prioritized and given preferential treatment, other packets either are denied or experience higher drop probabilities by being subject to either reduced scheduler time or reduced buffer space or both.
34. Contrary to some commenters’ assertions (*see, e.g., Verizon Comments at 84*), managing one’s network does not mean that engineers must have broad and unchecked discretion and flexibility to ensure functioning and secure networks. Tactical traffic management to stop malware, virus and denial of service attacks is easy to differentiate from systemic degradation of traffic based on non-engineering criteria, like business arrangements. Taking a concrete example, if there is a large denial of service attack or a virus probe from an infected user, appropriate filters or blocks can be applied in a tactical fashion to contain the attack. This does not mean, however, that a particular service provider or type of

product (such as third-party VoIP or video) must be rate-limited as a deliberate policy decision.

35. Likewise, while broadband providers may have different views on what network management is most effective (*see, e.g.*, Verizon Comments at 84, AT&T Comments at 186), and broadband network management can require technical, complex, and skilled decisions (AT&T Comments, Exh. 1 at 25), there are sound and common engineering standards that can be adopted to allow legitimate network management while prohibiting discrimination and other practices that can create traffic and packet distinctions unrelated to any legitimate engineering and technical requirements. For example, P2P discrimination can harm users using legitimate P2P applications such as Linux distribution downloads and Skype for communication. The classic traffic network management can be broken down into a few broad categories:

- 1) Tactical “security” management, as discussed above.
- 2) Long-term management of chronic congestion caused by lack of capacity. This generally is alleviated in the industry by adding capacity in places where there is a problem.

36. Some parties further assert that network congestion has been a problem that must be managed by broadband providers since the beginning of the Internet. These parties cite to an incident in 1987, where the pre-World Wide Web Internet suffered a “congestion collapse” that required an immediate response by providers. (AT&T Comments, Exh. 1 at 17). The congestion collapse problem was documented previously in RFC 896 (*see* <http://tools.ietf.org/html/rfc896>).

The technical description of the congestion collapse from RFC 896 is quoted below:

In heavily loaded pure datagram networks with end to end retransmission, as switching nodes become congested, the round trip time through the net increases and the count of datagrams in transit within the net also increases. This is normal behavior under load. As long as there is only one copy of each datagram in transit, congestion is under control. Once retransmission of datagrams not yet delivered begins, there is potential for serious trouble.

Host TCP implementations are expected to retransmit packets several times at increasing time intervals until some upper limit on the retransmit interval is reached. Normally, mechanism is enough to prevent serious congestion problems. Even with the better adaptive host retransmission algorithms, though, a sudden load on the net can cause the round-trip time to rise faster than the sending hosts' measurements of round-trip time can be updated.

Such a load occurs when a new bulk transfer, such a file transfer, begins and starts filling a large window. Should the round-trip time exceed the maximum retransmission interval for any host, that host will begin to introduce more and more copies of the same datagrams into the net. The network is now in serious trouble. Eventually all available buffers in the switching nodes will be full and packets must be dropped. The round-trip time for packets that are delivered is now at its maximum. Hosts are sending each packet several times, and eventually some copy of each packet arrives at its destination. This is congestion collapse.

The response was designed and implemented shortly by Van Jacobson *et al.* (see <http://ee.lbl.gov/papers/congavoid.pdf>). The salient point is that the solution was applied in a non-discriminatory fashion and, more importantly, was implemented in the standard protocol stacks without any network intervention needed or required. All applications of the same type (Telnet in the above example), for all users and all destinations, are subject to the same rules. Telnet for Provider A is not treated any differently than Telnet for Provider B. The solution proposed in RFC 896 is quoted below:

The solution is to inhibit the sending of new TCP segments when new outgoing data arrives from the user if any previously transmitted data on the connection remains unacknowledged. This inhibition is to be unconditional; no timers, tests for size of data received, or other conditions are required. Implementation typically requires one or two lines inside a TCP program.

37. Further, AT&T argues that engineers have long recognized the Internet involves traffic prioritization and Quality of Service (QoS) standards (e.g., use of user datagram protocol (UDP)). AT&T asserts that the engineers who actually set Internet standards (the Internet Engineering Task Force, or IETF) have long understood the importance of QoS capabilities as the best means of providing differentiated services that customers need and demand. However, AT&T then erroneously states that the proposed rule mandating nondiscrimination by broadband providers would prevent QoS from happening. (AT&T Comments at 37-38 and Exh. 1, at 16-18). These arguments are not valid because, as mentioned earlier, the applications electing to use particular mechanisms for QoS all would be treated the same, per the Telnet example in paragraph 25 above. To take a concrete example, let us say that there are three application service providers providing some service using Telnet. Telnet clients for all three service providers can set whatever QoS capabilities they desire. When the packets emitted from that service enter the broadband service provider's network, they are all subject to the same policy.
38. Other parties assert that "best efforts" service is not sufficient for all data, and that congestion leads to jitter, dropped packets and latency problems (Cox Comments at 21). AT&T also claims that "best efforts" traffic in a congested environment severely limits real-time applications (AT&T Comments, Exh. 1 at 18). Neither

claim is supported. For example, voice communication is a real-time application that has shown surprising resilience and utility in today's best-effort Internet. Further, the majority of providers in the Internet marketplace do not honor QoS markings from other networks that connect with them unless there is a customer relationship. For example, most Tier 1 providers (as defined in AT&T's comments) do not honor other Tier 1 QoS settings in their public networks. Regardless, the Internet continues to operate well.

39. Similarly, some parties argue that there is exponential online video growth that will put continuing pressure on broadband networks such that congestion cannot be eliminated simply by adding further capacity (Verizon Comments at 83). These assertions ignore the fact that the continuing growth of video traffic is merely the latest in the perennial "application that is going to cause problems for the Internet" saga. (See <http://arstechnica.com/old/content/2007/12/growth-of-p2p-leads-ietf-to-debate-fair-bandwidth-use.ars>). This argument was made as well in earlier times about "new" traffic from applications like NNTP, or netnews. With forward cache deployments, a robust services market in the CDN sector, and other similar techniques, video congestion issues likely will be substantially ameliorated in last mile broadband provider's networks.
40. Moreover, allowing different treatment for video packets, or rate-limiting the latest emerging application, likely would not a problem, as long as such type-based differentiation is implemented in a reasonable, objective, and neutral manner that applies to all packets equally. Capacity increases also can be balanced by non-discriminatory management of the packets as they flow through

the routers. By contrast, the discretion to engage in discriminatory throttling of one or more specific application service providers fails to address the basis of the argument: that exponential growth will outstrip the last mile provider's ability to expand capacity. If demand is indeed exponential, the growth will just move to the application providers that are not discriminated against, and the fundamental infrastructure problem will still not have been eliminated. However, the discretion to engage in discrimination allows the broadband provider to use its unique position in the Internet infrastructure unilaterally to select its preferred providers of applications and content.

E. Network evolution and technological progress is consistent with open broadband.

41. It is unquestioned that technological advances continue to lead to rapid changes in the Internet. However, this fact does not mean that near-term adoption of the proposed rules would effectively "lock in" the current Internet and restrict service providers' abilities to respond to changes in technology, as some argue (Verizon Comments, Attach. A at 6). It is critical to recognize that there is a logical and significant difference between potential uses of the architecture, design, construction, and operation of broadband infrastructure, and the policies that infrastructure owners choose to adopt and apply to that infrastructure.
42. Network evolution including architecture and operations is orthogonal to the proposed open broadband policies that can be implemented on the physical infrastructure and connectivity model. In several cases, a multitude of policies and operational responses have been implemented on infrastructure in a reasonable, nondiscriminatory, and transparent fashion. Examples include

Comcast's "unlimited" home Internet product switching to a capped 250 Gigabyte product without any visible disruption to the end user. The amendment can be seen at <http://www.comcast.net/terms/network/amendment/>. A quote from the URL is below:

Today, we're announcing that beginning on October 1, 2008, we will amend our Acceptable Use Policy (AUP) available at <http://www.comcast.net/terms/use/> and establish a specific monthly data usage threshold of 250 GB/month per account for all residential customers.

Similar examples exist in the wireless space as well, such as when AT&T and Verizon switched from "unlimited" datacard products to capped 5 Gigabyte products.

43. Even though new services and technologies continue to emerge, the proposed rules would not "lock in" today's Internet technology, nor would they eliminate or harm technological advances. If certain types of services require prioritization similar to the Telnet example mentioned earlier, it may be acceptable to apply prioritization to broad classes of applications in a fair and equitable manner. This would be consistent with the Internet's evolution as we have seen with Telnet, SMTP *et al.* Any allowed prioritization on an application level basis is entirely different from prioritization on a business entity basis. The suite of IETF protocols stays out of the business relationship for precisely this reason: the network evolution and connectivity methods, designs, and architectures are by design orthogonal to the business relationships at the higher layers. As designed, and as we have seen over the years, this has not caused stagnation in any protocol or architecture for connectivity.

Declaration of Vijay Gill
GN Docket No. 09-191, WC Docket No. 07-52

I declare under penalty of perjury that the foregoing is true and correct.

A handwritten signature in black ink, appearing to read "Vijay Gill", is written over a horizontal line. The signature is stylized and cursive.

Vijay Gill
Manager of Engineering and Global
Network Architecture
Google Inc.

Date: April 26, 2010

APPENDIX

CURRICULUM VITAE OF VIJAY GILL

I. PERSONAL INFORMATION

Vijay Gill
Google
1600 Amphitheatre Parkway, Mountain View CA 94089
Work phone: +1 650-253-8355
Work email: vgill@google.com

II. EXPERIENCE AND EDUCATION

Professional Experience:

Manager of Engineering and Global Network Architecture, Google Inc.

I am currently responsible for all network functions at Google, including Optical, Transport, IP, MPLS, Datacenter connectivity and Internet related protocol design.

Senior Technical Manager at AOL, LLC

Manager for the group responsible for all network design and evolution at AOL.

Manager, Architecture, Abovenet.

Manager of the group responsible for all network design, and evolution at Abovenet.

Senior Network Engineering Manager, UUNET

Lead engineer for Multicast, MPLS, IP core backbone for AS 701.

Educational Background:

University of Maryland, Baltimore County, Maryland, Computer Science.

III. PROFESSIONAL AWARDS AND AFFILIATIONS

Member, Internet Architecture Board at the IETF

(<http://www.iab.org/http://www.iab.org/>).

The IAB is chartered both as a committee of the Internet Engineering Task Force (IETF) and as an advisory body of the Internet Society (ISOC). Its responsibilities include architectural oversight of IETF activities, Internet Standards Process oversight and appeal, and the appointment of the RFC Editor. The IAB is also responsible for the management of the IETF protocol parameter registries.

APPENDIX

IV. INDUSTRY CONFERENCES

I am an experienced international presenter on network design, scaling, interconnection and cost methodologies at industry leading conferences, including:

- *Network Design for Large Scale Compute Infrastructure*. Keynote at OFC/NFOEC 2010, San Diego
- *Warehouse Scale Computing*. Invited paper/talk at ECOC 2009, Vienna, Austria.
- *Perspectives on Network Routing*. IEEE-Infocom, Barcelona, Spain
- *Design Analysis of a Global 10G Backbone*. NANOG 34, Seattle, WA
- *High-Capacity Streaming and Caching*. NANOG 32, Reston, VA
- *ATDN OSPF to IS-IS Conversion*. NANOG 29, Chicago, IL
- *Lack of Priority Queuing Considered Harmful*. NANOG 27, Phoenix, AZ
- *Services, Complexity, and the Internet*. NANOG 26, Eugene, OR
- *Operational Feedback to IP Equipment Vendors*. NANOG 26, Eugene, OR
- *Large Scale IP Networks: GMPLS and MPLS explained*. AT&T Labs. NJ
- *Global Routing System Scaling Issues*. NANOG 21, Atlanta
- *Multiservice Core Design*. NANOG 21, Atlanta, GA
- *Service Provider Route Filtering*. NANOG 20, DC

V. Publications

I have been the primary or contributing author of the following ITETF RFCs and peer-reviewed papers:

- RFC 3345 BGP Persistent Route Oscillation Condition
- RFC 3346 Applicability Statement for Traffic Engineering with MPLS
- RFC 3582 Goals for IPv6 Site-Multihoming Architectures
- RFC 5082 The Generalized TTL Security Mechanism (GTSM).
- RFC 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)
- RFC 4116 IPv4 Multihoming Practices and Limitations (multi6)
- RFC 4451 BGP MULTI_EXIT_DISC (MED) Considerations
- Report from the IAB Workshop on Routing and Addressing. The report is available at <http://tools.ietf.org/html/draft-iab-raws-report-02>
- Communications of the ACM Lack of Priority Queuing Considered Harmful (ACM Queue November 2004)