

December 30, 2014

Ms. Marlene H. Dortch, Secretary
Federal Communications Commission
445 Twelfth Street, SW
Washington, DC 20054

Via Electronic Filing

Re: **GN Docket No. 14-28, Protecting and Promoting the Open Internet**
GN Docket No. 10-127, Framework for Broadband Internet Service

Dear Ms. Dortch,

I¹ offer these comments to aid the Commission in reaching the proper conclusion in the matter of the regulatory classification of Internet Service over broadband networks. President Obama reanimated this issue by essentially directing the FCC to reclassify Internet Services under Title II of the Communications Act. Moreover, some commenters have claimed “Internet Protocol packet transfer is telecommunications” because they misconstrue the nature of Internet Protocol and Internet Service.² In reality, Internet Protocol is simply a packet format, not a means of transmission; and Internet Service encompasses much more than Internet Protocol processing.

I believe Title II reclassification of Internet Service over broadband would be an error for several reasons, but I will only address the technical nature of the service in this letter.

The Communications Act defines “information service” as “the offering of a capability for generating, acquiring, storing, transforming, processing, retrieving, utilizing, or making available information via telecommunications, and includes electronic publishing, but does not include any use of any such capability for the management, control, or operation of a telecommunications system or the management of a telecommunications service.” This definition applies to Internet Service because it emphasizes information processing in the essential nature of the service and not merely as a means of managing the service. If we take information processing out of Internet Service, nothing is left but the wired or wireless medium over which information flows and the Data Link Layer (layer two of the Open Systems Interconnection Reference

¹ I am an independent network engineering consultant and policy analyst, presently working at the American Enterprise Institute as a Visiting Scholar and at High Tech Forum as editor and founder. These remarks are offered in my personal capacity and do not necessarily represent the opinions of AEI or any client or sponsor. I have previously offered comments in the “Preserving the Open Internet” and “Broadband Industry Practices” dockets, GN 09-191 and WC 07-52 respectively, and offered testimony at the [FCC En Banc Public Hearing on Broadband Network Management Practices in Cambridge on February 25, 2008](#) as an invited technical expert. My CV is available at <http://www.bennett.com/resume.pdf>.

² Barbara A. Cherry and Jon M. Peha, “The Telecom Act of 1996 Requires the FCC to Classify Commercial Internet Access as a Telecommunications Service: Comments Before the FCC in the Matter of Protecting and Promoting the Open Internet, GN Docket 14-28,” December 22, 2014.

Model) service that actually transmits and receives information frames hop-by-hop in the overall path from client to server or peer to peer through the internals of the Internet.³

The Act defines “telecommunications” as “the transmission, between or among points specified by the user, of information of the user’s choosing, without change in the form or content of the information as sent and received” and “telecommunications service” as “the offering of telecommunications for a fee directly to the public, or to such classes of users as to be effectively available directly to the public, regardless of the facilities used.”

This definition clearly calls out telephony, the only information exchange that promises to send and receive information between end points without alteration of the information’s form or content. This definition is also applicable to the Data Link Layer of the network protocol stack; but it is not applicable at the Network Layer (layer three, where Internet Protocol resides) because layer two addresses are changed hop-by-hop as Network Layer packets transit the Internet, and for other reasons given below.

While telephony is an interaction between persons using telephone handsets that are essential elements of the telephone network, information service is an interaction between computers that involves continual interaction between computers and the transmission network as well as between computers and each other. In an Information Service, the human user – if there is one, which is not the case for Internet of Things applications – interacts with the computer, and the computer mediates this interaction with the network and the paired computer or computers. This is an entirely different and more complex information system than telephony is.

The Nature of Internet Service

For purposes of this analysis, I will stipulate that an element of telecommunications service is embedded at the Data Link Layer of the Internet Service that the Internet Service Provider (ISP) offers to the public. This is the case for both facilities-based ISPs such as Comcast, AT&T, Verizon, et al., and is also the case for unbundled, Over-the-Top (OTT) ISPs such as Sonic.net, a firm that offers ISP service over AT&T’s DSL lines in California as well as a small number of its own lines in certain neighborhoods and small communities such as Sebastopol and Brentwood, California.

Even when Sonic.net provides Internet Services over AT&T facilities, it uses its own facilities, either leased or owned, to connect DSLAMs in AT&T Central Offices to the Internet Exchanges that reach the Internet as a whole.

It is also the case that there is an element of telecommunications in the provision of MVPD video services because MVPDs transmit information of the user’s choosing from one or more network broadcast centers to one or more television sets, DVRs, personal computers, or other devices of the customer’s choosing. The user makes the choice about the information he or she will receive when subscribing to a rate plan, makes it again

³ Hubert Zimmerman, “OSI Reference Model - The ISO Model of Architecture for Open Systems Interconnection,” *IEEE Transactions on Communications* Com-28, no. 4 (April 1980): 425–32.

when programming their DVR, and makes it yet again when choosing a recorded program to view.

An MVPD is an MVPD and not a telecommunications service because of the content it delivers and the menu of programs it follows, not by virtue of the fact that the content is transmitted from one place to another. Similarly, an ISP is an Information Service because of the actions it must perform in order to make computers at the customer premise a functional part of the global system we know as the Internet. The task of an Internet Service Provider is essentially the same regardless of whether it connects to users over dial-up facilities or over broadband: on the Internet-facing side of the service, it performs the same services in either case. If a dial-up or unbundled ISP is not offering a Telecommunications Service, neither is a broadband ISP.

Internet Service is an Information Service

It's common to say that ISPs provide "access to the Internet". The Communications Act does so at §230(e)(2):

INTERACTIVE COMPUTER SERVICE.--The term "interactive computer service" means any information service, system, or access software provider that provides or enables computer access by multiple users to a computer server, including specifically a service or system that provides access to the Internet and such systems operated or services offered by libraries or educational institutions.⁴

"Access to the Internet" has to be understood as shorthand for "joining a customer premise computer to the Internet" because the actual service of the ISP makes the customer's computer a part of the Internet, just as capable of providing content and services to others as to providing the customer access to content and services hosted by other, co-equal computers on the rest of the Internet. Examples of making content and services available from customer premise equipment include sending email, engaging in chat, operating a web or ftp server, and engaging in peer-to-peer file sharing with BitTorrent. These are routine activities familiar to nearly everyone who uses a connected computer, smartphone, or entertainment device today.

The Internet is effectively a federation of interconnected computers, not simply a bulletin board system like America Online's (AOL) and CompuServe's legacy services. There can be no "access to the Internet" when the Internet is fundamentally different from a bulletin board; the mode of interaction Internet users employ is one of membership and participation, not one of access.

That being said, at §230(e)(2) the Communications Act clearly states that Internet Service is an Information Service.

Internet Protocol Packet Transfer is an Information Service

Contrary to popular myth, Internet Protocol is part of the Internet's routing function

⁴ *Communications Act of 1934, as Amended by the Telecommunications Act of 1996*, vol. 47 U. S. C. S 151 et seq., 1996, <http://transition.fcc.gov/Reports/1934new.pdf>.

rather than a transmission function. Internet Protocol is a Network Layer element in terms of the OSI Reference Model, and as such its only concern is with crossing the boundaries between networks. This is not a transmission function, since the means of crossing network boundaries is the simple movement of a packet descriptor from one area of memory within a router (a reception queue attached to the ingress port) to another area of memory (a transmission queue attached to the egress queue).⁵

The actual transmission and reception functions that characterize transmission are actually accomplished by Ethernet circuitry (or the functional equivalent) because Internet Protocol implementations lack the capability to perform transmission or reception; IP code can only request these services, it cannot perform them.

The Internet Protocol packet format identifies a desired destination network, but it cannot take any independent action to cause the packet to reach that destination. This is to say that by itself, Internet Protocol is utterly incapable of transmitting information; it is a passive information format that depends on active elements such as Ethernet and the Boundary Gateway Protocol (BGP) for actual transmission.

Internet Protocol exists in two forms today, versions four and six, which use different address formats and different options for requested transmission services. While most traffic transiting the Internet from end to end is enclosed in either the IPv4 or the IPv6 envelope, Internet Protocol is not the only format recognized by ISPs.

ISPs also use Dynamic Host Configuration Protocol (DHCP) to provision IP addresses and to set customer premise configuration options such as Domain Name Server (DNS) addresses and Classless Inter-Domain Routing (CIDR) parameters. Given that DHCP makes it possible for IP to function, it's difficult to accept arguments that IP is the elemental service provided by ISPs. In fact, IP packet transfer is one in a bundle of many information services provided by ISPs.

The declaration that Internet Service is nothing more or less than packet transmission is “cherry picking” that suppresses evidence of the many related functions that ISPs provide.⁶ IP packets move from source to destination over a multiplicity of Data Link Layer paths provided by many ISPs, including the originating ISP, the destination ISP, and some number (zero or more) of intermediary ISPs (AKA “transit providers”).

But even if we accept the argument that IP packet transfer is the elementary ISP service, it still does not follow that Internet Service is telecommunication; this is because IP

⁵ A packet descriptor is an internal information element created and processed by a router to facilitate transmission and/or reception. The end user does not create the packet descriptor, and the descriptor it is not preserved after a packet has been processed. Packet descriptors are the means by which layer 2 and 3 functions communicate with each other.

⁶ A similar claim was made in the Comcast/BitTorrent case by advocates who insisted that blocking a single TCP connection was the same as blocking an entire BitTorrent transaction. These advocates glossed over the fact that BitTorrent transactions use dozens of TCP connections, no one of which is essential the entire transaction. Blocking some but not all TCP connections in a transaction simply slows the transaction down but does not cause it fail.

packet transfer depends on routing, Ethernet services, third party networks, and agreements between and among deregulated network operators for the processing of information packets according to freely negotiated Service Level Agreements. In addition to these dependencies, users and “edge services” affect the quality of end-to-end Internet Service, as we’ve seen in the slowdowns recently inflicted on the general population of Internet users by Netflix and Cogent.⁷ Singling out one link in this chain for special regulatory treatment while leaving the others deregulated would be arbitrary.

Managing Internet Bandwidth is an Information Service

As a member of the Internet, customer premise equipment is obligated to behave in a manner consistent with Internet norms, and is required to protect itself from dangerous activities performed by other members. One example of conforming to Internet norms is the TCP Congestion Control system governed by “Jacobson’s Algorithm”.⁸ The Internet lacks a built-in mechanism for protecting itself from overload.

This is not by design; Internet protocol designers included a mechanism known as “Source Quench” in the original specification of the Internet Control Message Protocol that was meant to provide overload protection, but it didn’t work (see: RFC 777).⁹ The Source Quench mechanism wasn’t tested until Ethernet bypassed ARPANet in the mid-1980s and the Internet suffered Congestion Collapse.¹⁰

Jacobson’s Algorithm requires Internet members – known as “hosts” – to reduce their rate of transmission when signaled by an Internet router that congestion is growing to dangerous levels. The router discarding a packet typically provides this signal, but altering a bit in the Internet Protocol header can also provide it; this latter method is known as “Explicit Congestion Notification” (ECN).¹¹ While hosts that do not conform to Jacobson’s Algorithm are not kicked off the Internet, its successful operation depends on broad conformance because normal Internet operation involves hosts cycling between underload and near overload.

Hosts are owned and maintained by end users, and routers are owned and maintained by ISPs and by end users; residential-focused, business-focused, and transit-focused ISPs all participate in this system, but they don’t necessarily signal congestion at by the same means or at the same time. The use of packet discard as congestion signal is obviously ambiguous, because packet discard also takes place for other reasons: when wireless packets collide, they tend to be dropped by the receiving host because their addressing

⁷ Dan Rayburn, “Cogent Now Admits They Slowed Down Netflix’s Traffic, Creating A Fast Lane & Slow Lane,” *StreamingMediaBlog.com*, November 5, 2014, <http://blog.streamingmedia.com/2014/11/cogent-now-admits-slowed-netflixs-traffic-creating-fast-lane-slow-lane.html>.

⁸ Van Jacobson, “Congestion Avoidance and Control,” *Computer Communication Review*, ACM Special Interest Group on Data Communication, 25, no. 1 (1995): 157.

⁹ J Postel, “RFC 777 - Internet Control Message Protocol” (Network Working Group, April 1981), 777, <https://tools.ietf.org/html/rfc777>.

¹⁰ W. Stevens, “TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms” (Network Working Group, January 1997), <http://tools.ietf.org/pdf/rfc2001.pdf>.

¹¹ K. Ramakrishnan, S. Floyd, and D. Black, “RFC 3168 - The Addition of Explicit Congestion Notification (ECN) to IP,” Internet RFC, (September 2001), <http://tools.ietf.org/rfc/rfc3168.txt>.

information can't be trusted.

Routers also implement a number of different sub-algorithms of the Jacobson master scheme, such as Random Early Detection, Weighted Random Early Detection, Adaptive Random Early Detection, and Robust Random Early Detection.¹² The choice and execution of these algorithms is an example of close interaction between customer premise equipment and ISPs, not only those that directly serve the end user but also those connected directly to the customer's ISP or to other ISPs connected to the end user's ISP.

Implementing Jacobson's Algorithm requires extensive information processing in the customer's hosts and routers as well as in the ISP's routers, in part to distinguish packet loss due to congestion from that caused by other factors, and in part to select and implement the most effective congestion control mechanism. We are not at the final stage of Internet congestion avoidance and management, of course. Not only is the technical literature replete with research in this area, such policy works as *A "Third Way" on Net Neutrality* address this matter as well.¹³

Bandwidth management is essential to Internet Service because packet-switching systems allocate bandwidth on demand and apportion it among concurrent users dynamically. This differs from the way telecommunication networks operate; when telephone users take part in calls, the telephone network allocates a fixed quantum of bandwidth to callers, limited to a few kilohertz, and this allocation lasts for the duration of the call whether it is used or not. But when Internet users access a web server, the ISP allocates bandwidth dynamically.

If only one user is active on a given path or segment at a time, that user is able to use the entire capacity of that path, but if many users are active, each contends with the others for capacity and the ultimate assignment is a function of patterns of user demand, location of end points, upstream congestion, server capacity, and a host of other factors. At each point in the path between the ISP's customer and the Internet-based service that user accesses, decisions are made regarding the treatment of each packet.

This isn't simply "the management of a telecommunications service" as some may argue. The management of a telecommunications service is a set of actions that ensure the service conforms to a predetermined level of quality, but dynamic bandwidth management in a packet switched network is a real-time negotiation of service parameters.

When network load is light, users enjoy a better quality experience than they enjoy when load is heavy. On a telecommunication network the user's experience is the same whether the network is lightly or heavily loaded. An overloaded packet switched network

¹² Wikipedia, "Random Early Detection," accessed December 23, 2014, http://en.wikipedia.org/wiki/Random_early_detection.

¹³ Robert D. Atkinson and Philip J. Weiser, *ITIF: A "Third Way" on Network Neutrality*, report (Washington, DC: Information Technology and Innovation Foundation, May 30, 2006), <http://www.itif.org/index.php?id=63>.

continues to function, but does so slowly, while an overloaded telecommunications network simply refuses to connect new calls.

Consequently, Internet service is a dynamic service made available on a statistical basis and provided by a system that relies on information processing, while telecommunication is static service that can be provided on an analog network where human operators plug wires into panels to make connections, as it was for many years. It would not be possible to provide packet switching in the same manner as it relies on millisecond-by-millisecond dynamic decisions made dozens of times in the transmission of each information packet.

Attack Mitigation is an Information Service

Billions of people use the Internet through their own computers worldwide, and (unsurprisingly) some are up to no good. The Commission is aware that criminals use the Internet for the theft of intellectual property, identity theft, and extortion. This is no trivial matter as the connectionless nature of Internet Protocol makes it an ideal vehicle for Denial of Service attacks, and the insecure nature of basic Domain Name Service allows criminals to use DNS and Simple Network Management Protocol (SNMP) as amplifiers for attacks. Distributed DoS attacks are made possible by viruses that enable botnet operators to invade and take over end-user systems in order to enlist them into their botnets, where they can be used to send spam and to take part in DDoS attacks.

There is no parallel to a DDoS attack using amplification to bring a web site to its knees in the realm of plain old telephone service.

Mitigating these attacks requires ISPs to engage a multi-pronged strategy, using information technology to distribute anti-virus software to end user computers, to monitor networks for suspicious traffic and attacks, to block (or redirect) attack traffic when it is found, and to notify other ISPs of infected computers on the other ISP network so that end users can take appropriate action.

Attack mitigation is not simply a management function performed to make networks operate; it's an added-value service that is necessary to reduce the incidence of unlawful activity across the Internet.

The most obvious and well-know element of attack mitigation is the anti-virus software that ISPs make available to their customers, typically free of charge. Anti-virus software is an intensive use of information processing to search for viruses in downloads and incoming email, to monitor the integrity of system files, and to distribute attack knowledge to software producers so that they can fix bugs that allow viruses entry where they are not wanted. Internet Service without security would be a meaningless and dangerous offering.

Domain Name Service is an Information Service

Internet Service always includes Domain Name Service provided over “the largest distributed database in the world”.¹⁴ DNS is an increasingly sophisticated distributed function that translates domain names into IP addresses, its best-known function, but it does much more. DNS implements the DNSSEC protocol, an authentication service that validates the correctness of the domain name to IP address mapping and protects users from man in the middle (MITM) attacks. DNS is also a traffic direction service that connects Content Delivery Network users to the nearest and/or fastest location. DNS also provides a reverse mapping from IP addresses to domain names, and distinguishes authoritative domains from other domain names that may share an IP address.

DNS manages aliased domain names – another case of multiple domain names sharing a common IP address – and provides both IPv4 and IPv6 addresses. DNS distinguishes multiple services within a domain, such as the email “Mail Exchanger” and the web service. The database managed by a DNS server is updated in real time, with updates shared across the entire Internet as needed. DNS servers protect themselves from attacks, since a simple, unprotected DNS server is an attack vector that can amplify DDoS attacks in much the same way an unprotected SNMP agent can.¹⁵

It must be acknowledged that the largest distributed database in the world is an information service, or nothing is. It must also be acknowledged that DNS is an indispensable part of the Internet Service provided by ISPs.

With the advent of DNSSEC, the DNS service provided by ISPs today is much more information-processing intensive than was the DNS service provided by ISPs when the FCC classified cable modem Internet service as an Information Service. Consequently, the logic that guided the Commission’s previous classification decision is even stronger today than it was in 1992. DNSSEC processes much more information than simple DNS did, and for a very good reason: information security. As VeriSign describes it, DNSSEC is essential for the protection of DNS information from attacks, and DNS itself is essential to the operation of the Internet:

*The Domain Name System (DNS), the Internet's addressing system, is the most critical component of the Internet infrastructure. **Without it, the Internet could not function.***

However, it was not designed with security in mind. As a result, it is vulnerable to man-in-the-middle (MITM) attacks and cache poisoning. These threats use forged data to redirect Internet traffic to fraudulent sites and unintended addresses.

¹⁴ Fred Donovan, “DNS Infrastructure Is ‘Highly Vulnerable’ to Attacks, Warns Infonetix,” news blog, *Fierce IT Security*, (November 14, 2014), <http://www.fierceitsecurity.com/story/dns-infrastructure-highly-vulnerable-attacks-warns-infonetix/2014-11-13>.

¹⁵ Broadband Internet Technical Advisory Group, “SNMP Reflected Amplification DDoS Attack Mitigation” (Broadband Internet Technical Advisory Group, August 2012), <http://www.bitag.org/documents/SNMP-Reflected-Amplification-DDoS-Attack-Mitigation.pdf>.

Once an unsuspecting user or device reaches the fraudulent site, cyber criminals can potentially extract credit card data, steal user passwords, eavesdrop on voice over IP (VoIP) communications, plant malicious software or display images and text that defame the legitimate brand or provide misleading information. Given that a single DNS name server can act as the name-to-address resolution point for thousands of users, the potential impact of a MITM attack or cache poisoning can be considerable.¹⁶

[Emphasis added.] DNSSEC is information processing, and without it all commercial transactions over the Internet are suspect.

Routing is an Information Service

Routing is an indispensable element of any packet-switched network such as the Internet. In its most elementary form, the routing function determines whether packets of information received by a router are to be dropped, forwarded, or processed.

A packet is dropped if it comes from an unauthorized source, or if the forwarding path is congested or unavailable. A packet is forwarded if its network identifier matches a known valid route and resources are available for forwarding. Packets are processed if they contain management information such as routing map updates or network management commands.

In commercial settings, all packets potentially have implications for accounting, security, and public safety, so routers also provide these functions. Network Quality of Service is provided and ensured by routers, and these functions are present in routers in a number of different forms.

While the telephone network determines a path from the calling to the called party when calls are setup, it simply records a circuit identifier at call establishment time, which is used by all subsequent elements of the call. Packet switching routers, on the other hand, recalculate the route from source to destination every time a packet is forwarded, a much more intensive information-processing task. Packet routers also react to network failures by choosing alternate routes while a packet is in flight, sometimes reacting to network failures in small fractions of a second. Consequently, packet routers perform several orders of magnitude more computation than telephones switches do.

It is useful to compare the functions performed by common web servers with those performed by Internet routers. There is no dispute that web services are Information Services, so it follows that ISP services performed over routers must be information services as well if the two are essentially similar or if routers are more information service intensive. The following chart examines the functions of web servers and routers according to the definition of Information Service in the Communications Act, the offering of a capability for “generating, acquiring, storing, transforming, processing,

¹⁶ VeriSign, “DNSSEC Test and DNSSEC Testing,” corporate blog, *VeriSign*, accessed December 29, 2014, http://www.verisigninc.com/en_US/innovation/dnssec/dnssec-test/index.xhtml.

retrieving, utilizing, or making available information via telecommunications”.¹⁷ Each term in this definition is presented as an “Information System Property”.

Information System property	Web server	Internet router
Generating information	<ul style="list-style-type: none"> • Creates dynamic web pages with personalized content and real-time content such as Twitter streams, ads, and news streams. 	<ul style="list-style-type: none"> • Creates updates to routing table; • Resolves DNS queries; • Forwards information packets; • Signals operational options such as Quality of Service parameters; • Generates error messages in ICMP terms; • Signals congestion to end point by dropping packets; • Gathers network health; and • Responds to network management queries.
Acquiring information	<ul style="list-style-type: none"> • Obtains information about the user, browser, platform, end point location and advertisements 	<ul style="list-style-type: none"> • Receives packets, routing updates, Quality of Service parameters, end-point location, and load balancing information.
Storing information	<ul style="list-style-type: none"> • Stores software and software updates, web page elements, user IDs and passwords, comments, and log file elements. 	<ul style="list-style-type: none"> • Stores packets in queues; • Stores software and software updates, next-hop routing information, user IDs and passwords, configuration parameters and log files.
Transforming information	<ul style="list-style-type: none"> • Redirects web queries to alternate URIs and URLs and modifies web query content. 	<ul style="list-style-type: none"> • Transforms MAC addresses from original from/to pair to new pair for next hop; • Decrements TTL; • Modifies Class of Service indicators.
Processing information	<ul style="list-style-type: none"> • Chooses from multiple representations of web pages, images, audio/video 	<p>For each packet transferred:</p> <ul style="list-style-type: none"> • Looks up next hop by prefix or by port;

¹⁷ *Communications Act of 1934, as Amended by the Telecommunications Act of 1996.*

	<p>streams, and page formats according to user’s browser version;</p> <ul style="list-style-type: none"> • Compresses and decompresses streams; • Examines user cache to determine when to retransmit images; and • Interacts with local caches. 	<ul style="list-style-type: none"> • Determines success or failure of transfers and adjusts best route according to network conditions; • Chooses priority queue by SLA contract and packet preferences; • Discards packets according to various congestion algorithms for long queues; • Re-orders queues as needed; • Determines packet priority by deep packet inspection if necessary; • Determines whether inbound and outbound packets are parts of attacks; • Processes Access Control Lists for forwarding, discards, and other purposes. <p>On a longer-term basis:</p> <ul style="list-style-type: none"> • Updates routing tables in response to a variety of conditions in the forwarding path, changes in contracts, and attack/malware mitigations.
<p>Retrieving information</p>	<ul style="list-style-type: none"> • Fetches web pages and page elements from local or network storage, gathers cache information from end user system and end user’s local cache; • Obtains ads for web pages; • Gathers dynamic streams from various network and local sources. 	<ul style="list-style-type: none"> • Fetches routes from neighbors and network-wide functions; • Gathers information from network interfaces; • Receives software updates; • Stays abreast of network conditions on alternate routes; • Obtains DNS updates from more authoritative sources; and • Obtains network management information

		from other routers.
Utilizing information	<ul style="list-style-type: none"> • Uses user identity, location, and preferences to authenticate access to protected content, perform financial transactions, sell advertisements, localize services, and obtain desired content as directed by links. 	<ul style="list-style-type: none"> • Utilizes destination address to determine best route; • Utilizes destination network address to determine best route; • Utilizes port-relative network prefixes to determine best route; • Aggregates packet streams in LISP; • Utilizes SLA terms to determine port-based packet queue ordering; • Utilizes stream profiles to identify attacks;
Making information available	<ul style="list-style-type: none"> • Makes wide variety of information available in a number of forms, as this is the primary function of web server. 	<ul style="list-style-type: none"> • As parts of distributed databases of domain name and routing information, makes a wide variety of location, identity, service level and routing information available across the Internet.

In terms of overall processing power, the contemporary router is several times more powerful than the typical web server. This processing power comes at a price, as the typical router is also several times more expensive than the typical web server. Router users would not be willing to pay this price if end users did not require the router’s information service elements.

Routing service is more information processing intensive today than it was in 1992: there are many more routes, there are two types of IP address formats in use, and ISPs support new protocols such as LISP that attempt to deal with the explosion in the size of the routing table.¹⁸ Routing is also becoming more secure now, thanks to pending upgrades in BGP, the Internet’s basic routing information exchange protocol in the interest of security.¹⁹ BGP also has means for exchanging Quality of Service parameters that did not exist until some ten years ago.²⁰

¹⁸ D Farinacci et al., “RFC 6830 - The Locator/ID Separation Protocol (LISP)” (RFC Editor, January 2013), <http://tools.ietf.org/html/rfc6830>.

¹⁹ S Bellovin, “Security Requirements for BGP Path Validation” (RFC Editor, August 2014), <http://www.rfc-editor.org/rfc/rfc7353.txt>.

²⁰ Cisco, “Using BGP Community Values to Control Routing Policy in Upstream Provider Network,” corporate site, (August 10, 2005), <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/28784-bgp-community.html>.

Routing is more information-rich than it has ever been.

Relationship of Telecommunications and Information Service Elements of Internet Service Provision

The OSI Reference Model and similar constructs describe a modular distribution of functions in a network.²¹ Layers in these models describe similar functions performed at different scopes. All network layers transmit and receive information, whether they be telecommunications devices, Internet routers, or web servers and browsers, but they do so over paths that differ mainly by distance. Within the Internet, the telecommunication function consists of a set of circuits connecting one unique end point to another, such as an Ethernet cable connecting a personal computer to an Ethernet switch or an (optical) Ethernet cable from a router to an Ethernet switch. Ethernet is an un-routed service, similar to Wi-Fi, DOCSIS, DSL, or Passive Optical Networking (xPON). For simplicity, I use “Ethernet” as a proxy for any and all of these systems.

Ethernet packets carry a payload consisting of IPv4 or IPv6 data elements and their respective payload (which consists of TCP, HTTP, and similar application-oriented protocols.) The Ethernet frame carries very limited information and is not routable as the addresses Ethernet processes are local to the switch.

IPv4 and IPv6 information packets are co-mingled over Ethernet circuits and contain addresses that are unique across the entire Internet. While an Ethernet switch can interconnect dozens of computers, a router connects any computer to any other computer anywhere in the world. The means of interconnecting the dozen or so devices in a home, office, or neighborhood to each other is telecommunication, and the means of interconnecting any computer to any other computer is an information service.

The telecommunications scope is perhaps as many as a few hundred machines in a closely controlled network, while the information service scope is billions of machines attached over hundreds of millions of telecommunications connections in a loosely coordinated, unpredictable mesh in of both good and bad actors, heavy and light users, sophisticated customers and naïve amateurs with wildly different needs and desires.

The manager of a telecommunications network can over-provision the network under his or her control to ensure it never stalls or congests and never provides service to bad actors, but the manager of an Internet service can never have that much control. Routing is a fundamentally different activity from switching, and maintaining the Domain Name Service is fundamentally different from creating a table of computer name to IP address mappings with a text editor (as can be done in Windows, OS X, or Linux).

The Essence of the Classification Decision

The proper classification of Internet Service under the Communications Act requires the regulator to determine whether the functions performed by the ISP are more or less

²¹ Zimmerman, “OSI Reference Model - The ISO Model of Architecture for Open Systems Interconnection.”

similar to those performed by the millions of American consumers with Ethernet switches in their homes and offices than to those performed by web site owners in terms of the extent of their information processing intensity. The law indicates that Information Service is performed over a telecommunications facility, and that it does more information processing than the telecommunications facility does (even in the interest of its internal management).

The descriptions I have provided of congestion management, attack mitigation, domain name service and IP routing show that Internet Service is more than telecommunication. Consumers are more than capable of providing their own “telecommunication” services today with inexpensive hardware that can be bought in retail stores, just as pizza deliverers can provide delivery services with common bicycles and automobiles.

Internet service is, however, a specialized information technology-enabled service that can only be provided by highly skilled operators with a deep pool of talent and a serious investment in equipment, training, and infrastructure. It makes no more sense to classify Internet Service as simple telecommunication than it would to classify integrated circuit design and production as telecommunication simply because some chips are used in telephone networks. These are two vastly different realms, and to confuse them in a way that does injustice to the nature of Internet Service would be a clear technical error.

I urge the Commission not to change the regulatory classification of Internet Service. The FCC is meant to be an independent, expert agency that interprets the law and applies it to technical realities, and those realities are abundantly clear: Internet Service is an Information Service as defined by the Communications Act.