August 29, 2017

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

Washington, DC 20054

Via Electronic Filing

In the Matter of	)	
	)	
Restoring Internet Freedom	)	WC Docket No. 17-108

Dear Ms. Dortch,

I[1] offer these reply comments to aid the Commission in reaching the proper conclusion in construing the nature of Internet Service, determining the regulatory classification of Internet Service (IS) over broadband networks, in creating the regulations that should apply to Internet Service providers (ISP), and in repealing the regulations that should not apply to such services. These comments chiefly reply to the so-called "Engineers Letter" filed by the Electronic Frontier Foundation (EFF) in the initial round of comments.[2] The factual and analytical errors in The Letter render it misleading and useless.

---

[1] I am an independent network engineering consultant and policy analyst, presently working at High Tech Forum as editor and founder and as an independent consultant. These remarks are offered in my personal capacity and do not necessarily represent the opinions of any client or sponsor. I have previously offered comments in the "Protecting and Promoting the Open Internet" docket, GN 14-28, 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] Aaron L. Jones et al., "Joint Comments of Internet Engineers, Pioneers, and Technologists on the Technical Flaws in the FCC's Notice of Proposed Rule-Making and the Need for the Light-Touch, Bright-Line Rules from the Open Internet Order" (ID 1071761547058, July 17, 2017), https://ecfsapi.fcc.gov/file/1071761547058/Dkt.%2017-108%20Joint%20Comments%20of%20Internet%20Engineers%2C%20Pioneers%2C%20and%20Technologists%202017.07.17.pdf.

1. Summary

The EFF's Engineers' Letter ("The Letter") asserts that the Internet "in many ways parallels the telephone network" with respect to routing and other features and is therefore appropriately regulated, at least in part, by Title II of the Communications Act.[3] This assertion contradicts the story that Internet engineers have told since the 1960s as well as the technical realities of the two networks.

While The Letter attempts to provide a point-by-point comparison of the plain old telephone service (POTS) to the Internet, it fails to provide a similar comparison of the Internet to Title I Information Services such as dial-up Internet Service Providers (ISP) or the platform services offered by firms such as the Internet's six major service monopolies, Google, Facebook, Amazon, Apple, Microsoft, and Netflix. A more thorough comparison reveals that broadband ISPs share more features in common with platform services.

The letter falls short of providing the Commission with useful, actionable insight by failing to grasp the following issues:

[3] Ibid., page 10.

The Letter’s argument for retaining Title II classification for ISPs rests on a fundamental misunderstanding of transmission service in the terms of the Communications Act. It falls victim to the same factual and analytical errors that plagued the Wheeler FCC’s 2015 Open Internet Order.

The Wheeler FCC didn’t actually define the broadband transmission service in a manner analogous to Title II transmission services used by dialup ISPs. Rather, it chose to apply the “transmission service” label to the IP datagrams that define the internetworking layer of the Internet. This is a curious move because it presumes that IP datagrams are self-transmitting. In fact, no IP datagram gets anywhere unless it is carried as the payload of a Link Layer information frame.

The Pai FCC, on the other hand, treats Internet Protocol datagrams as part of the Information Services ecosystem that makes the Internet something people care about. It asserts that IP provides access to Information Services and also that IP depends on broadband technologies for simple transmission between the customer premise and the ISP as well as between ISPs and unaffiliated Information Services.

One of the more telling asides in The Letter concerns the evolving telephone network (page 26):

> Thus, this interpretation [the observation that consumers are often unaware of where online content is stored] of what it means to transmit information between or among points specified by the user, i.e. that the user must explicitly tell the network what routing decisions to take, has no basis in reality. Taken to its logical conclusion, it would require the FCC to similarly decide that telephone services are also not telecommunications services—an obviously absurd conclusion.

The Letter is forced to acknowledge that today's telephone network bears little resemblance to the POTS networks of the '30s or even of the '90s. Mobile phones require massive network intelligence to support consumer traffic, whereas, the "routing" to telephone numbers that represent wires is simpler than routing to portable numbers on foreign networks. Modern directory service isn't much like picking up the phone and asking Ernestine the Operator for a number either.

It strikes me as reasonable to consider the question of how well the statutory definition of POTS in Title II of the Communications Act even applies to modern telephone networks. Many seek an update to the Act because it's so out-of-date and ambiguous.[4]

If Title II is too narrow to even describe modern telephone networks, it's obviously not a good fit for much more powerful Internet services. This is the key takeaway from the search for similarities between the Internet and POTS.

2. What Internet technology promises to provide

The Letter attacks the Commission's understanding of the Internet with a "Brief Introduction to the Internet" originally filed with the DC Circuit in the challenge to the 2015 Open Internet order. An attack on today's FCC written for an *amicus* brief isn't really on point.

---

[4] Richard Bennett, Jeffrey A. Eisenach, James K. Glassman, Bronwyn E. Howell, Justin (Gus) Hurwitz, Roslyn Layton, Bret Swanson, "Comments on Communications Act Modernization" (Comments to Congress, Washington, D.C., February 4, 2014), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2388723.

The FCC and Chairman Pai could very well "lack a fundamental understanding of what the Internet's technology promises to provide, how the Internet actually works, which entities in the Internet ecosystem provide which services, and what the similarities and differences are between the Internet and other telecommunications systems the FCC regulates as telecommunications services" as The Letter claims. But trying to prove that with a letter written for an entirely different purpose is disingenuous at best. In fact, The Letter is guilty of all the charges it raises against today's FCC.

Any claim about the "promise of the Internet" is necessarily subjective. Some people claim the Internet promises free access to information, the freedom to communicate, and democracy; but in many cases it brings about controlled speech and surveillance. While the negatives are commonly associated with authoritarian governments such as China and Saudi Arabia, this isn't the whole story.

The financial model for most of the large platform networks (AKA "edge services") is based on advertising; this involves tracking visits to websites and the creation of dossiers of the interests and activities of individual users. The Letter doesn't mention any of this, but surveillance states are jealous of the dossiers Google and Facebook compile not just on visitors to their sites, but on visitors to the websites that use their tracking code.

In reality, the Internet promises to connect every network that wants to participate on terms of each network's choosing. When we get to specific terms and conditions of the interconnection, we're in the realm of local and national policy. The design of the Internet doesn't specify what we can do with our connections.

The claim that the Internet inherently requires "openness and non-interference" is wishful thinking. We Americans want it to be open and indiscriminate, but there is nothing in the design of the Internet that forces it to behave that way. And the idea that designers of a system are free to create the public policy for the system is indeed very troubling. Should we extend that privilege to firearms designers, chemical manufacturers, or automakers?

The Great Firewall of China proves that the Internet behaves the way lawmakers want it to behave. And that's why policy discussions are more important that technical claims.

The Internet simply provides us with an opportunity to communicate; what we do with it is up to us.

3. How the Internet is organized

Any attempt to describe the Internet in ten double-spaced pages is going to offer a big target to fact-checkers. This is only enough space to list the titles of the last six months' worth of RFCs, one per line. But we do want something more than an *Internet for Dummies* introduction from the people who offer their expert assessment of the facts as justification for a particular legal policy. Unfortunately, the letter comes up short in that regard.

The first two and a half pages simply say that the Internet is composed of multiple networks: ISPs, so-called "backbones", and edge services. These networks connect according to interconnection agreements, which the letter somewhat incorrectly calls "peering arrangements".

This description leaves out the transit networks that connect small ISPs to the rest of the Internet (for a fee) in the absence of individualized peering agreements. This is an important omission because the overview sets up a complaint about the tradition of requiring symmetrical traffic loads as a condition to settlement-free peering.

When transit networks interconnect, settlement-free peering requires symmetrical traffic loads because asymmetry would imply that one network is providing transit for the other. Because transit is a for-fee service, it obviously would be bad business to give the service away for free. The letter fails to mention this, but insists on pointing out that ISPs provision services asymmetrically.

This misleading description of the organization of the Internet is later used to justify settlement-free interconnection to ISP networks by the large edge service networks operated by Google, Facebook, Netflix, et al.

4. Which players in the Internet ecosystem provide which services

The Internet is simply a communication system that interconnects networks of computers to other networks of computers. This interconnection has two parts: 1) A physical interconnection with a common technology (most commonly Ethernet in addition to TCP/IP); and 2) An agreement to interconnect, either directly or through a third party.

Third party, for fee "transit" networks, which The Letter fails to mention, are the most common mode of interconnection for the small networks, and settlement-free peering is the common interconnection for the small number of extremely large, dominant networks such as Google, Facebook, Comcast, Verizon, and AT&T.

Ordinary businesses and regional ISPs connect to the Internet through transit providers such as Level 3. Level 3 provides both physical connections to the entire Internet and agreements with all of its constituent networks, either directly or through other transit networks. So most networks connect to each other under agreements by service providers.

The Internet consists of connections, agreements, and information systems. There are no bright lines in the Internet segregating these three things.

The large edge services – the monopolies enjoyed by Google, Facebook, et al. – have networks with global reach. They have facilities all across the US and in other nations. But so do the transit networks and the "backbone networks" owned by large ISPs. So everybody is in the connection business.

In addition to the basic interconnection agreements already mentioned, networks also agree to share routing information with each other via common routing information protocols like BGP. Networks also agree to share domain information with a DNS, a non-routing system used by large networks to organize their resources.

DNS also provides lookups between names and numbers, like the telephone directory. But unlike the telephone directory, DNS is not managed solely by connection merchants.

It's a shared database updated in close to real time by domain owners. Access to an Internet-based service can require dozens of DNS database queries in sub-second intervals of time.

5. The nature of DNS: a general-purpose database

DNS is much more than a simple directory. The Letter claims DNS's "primary purpose is to enhance functionality otherwise provided by the internetwork layer". This is such a vague statement that it is meaningless. We can also claim that DNS "enhances functionality otherwise provided by the application layer", or any other part of the Internet.

The Letter also claims that "DNS allows a vital level of abstraction" because it makes it unnecessary to remember IP addresses and allows IP addresses assigned to specific computers to change over time. But the Internet had means of doing this before we had DNS: we used to download a "hostfile" every week that had computer names and associated IP addresses in it. You can use a hostfile today to bypass DNS if you want.

The benefit of DNS over the hostfile is the same as for any database over a text file equivalent: concurrent access by multiple users, resilience, and instant updates. Having a shared database that spans the entire Internet – and allows for the storage and retrieval of any information, not just domain names and IP addresses – provides a number of benefits.

Service networks can use DNS to optimize resources, balance load, and match up content formats with users. DNS supports the application layer, not just the internetwork layer. It's completely misleading to claim DNS is merely a tool to make packet routing happen.

Arguably, DNS has nothing at all to do with routing. Routing is the job of IP and BGP.

6. Congestion and packet switching

The Internet is based on a transmission technology known as packet switching, which the engineers describe by referencing an out-of-print law school textbook, *Digital Crossroads: American Telecommunications Policy in the Internet Age*, by Jon

Nuechterlein and Phil Weiser. Nuechterlein and Weiser are both very bright lawyers – Jon is a partner at Sidley Austin and Phil runs the Silicon Flatirons Center at the CU Boulder Law School – but this is probably the first time in history that a group of engineers has turned to a pair of lawyers to explain a fundamental technology for them. (This reference leads me to believe the letter was written by EFF staffers rather than by actual engineers.)

The description of packet switching omits three key facts:

- Packet routers are stateless devices that route each packet without regard for other packets;
- Packet switching is a fundamentally different transmission technology than circuit switching, the method used by the telephone network; and
- Packet switching increases the bandwidth ceiling available to applications at the expense of Quality of Service.

The omitted facts would have been helpful in explaining congestion, an issue that the letter combines with its description of packet switching in a way that makes it appear arbitrary. Packet switched networks are provisioned statistically, hence any network that is not massively over-provisioned will undergo periodic congestion. Therefore, any well-designed packet network must include the capability to manage congestion.

In the case of the Internet, congestion management is somewhat troubled topic. It was originally addressed by Vint Cerf through a mechanism called Source Quench, discarded because it didn't work.[5] Quench was replaced by the Jacobson Algorithm (a software patch consisting of two lines of code), which was at best a partial solution.[6]

---

[5] Fred Baker, "RFC 1812 - Requirements for IP Version 4 Routers" (Network Working Group, June 1995), https://tools.ietf.org/html/rfc1812.

[6] Van Jacobson, "Congestion Avoidance and Control," *Computer Communication Review*, ACM Special Interest Group on Data Communication, 25, no. 1 (1995): 157.

Jacobson's patch was supplemented by Random Early Detection, which didn't solve the problem entirely either.[7] The current status quo is Controlled Delay Active Queue Management (CoDel), a somewhat less than ideal system that seeks to manage transmission queues more accurately.[8]

The letter claims "the sole job [of routers] is to send packets one step closer to their destination." This is certainly their main job, but not their only one. Routers must implement the Internet Control Message Protocol (ICMP) in order to support advanced routing, network diagnostics, and troubleshooting. So routers also have the jobs of helping administrators locate problems and optimizing traffic streams.[9]

Internet tools such as traceroute and ping depend on routers for packet delay measurement; networks also rely on ICMP to verify routing tables with "host/network not reachable" error messages and to send ICMP "Redirect" messages advising them of better routes for packets with specific Type of Service requirements. The Redirect message tells the sending computer or router to send the packet to a different router, for example. So the actual job of routers is to implement all of the specifications for Internet Protocol routers.[10]

Two router specifications that go unmentioned are Integrated Services and Differentiated Services.[11] The letter makes no mention of Source Routing, a system that allows applications to dictate their own paths through the Internet. While Source Routing is rare, IntServ is used by LTE for voice and DiffServ is used within networks for control of

---

[7] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking* 1, no. 4 (August 1993): 397–413, doi:10.1109/90.251892.

[8] Kathleen Nichols and Van Jacobson, "Controlled Delay Active Queue Management" (TSVWG, March 10, 2014), http://www.pollere.net/Pdfdocs/draft-02.pdf.

[9] Baker, "RFC 1812."

[10] Ibid.

[11] S. Blake et al., "RFC 2475 - An Architecture for Differentiated Services" (Internet RFC, December 1998), http://tools.ietf.org/rfc/rfc2475.txt; R. Braden, D. Clark, and S. Shenker, "RFC 1633 - Integrated Services in the Internet Architecture: An Overview" (Internet RFC, June 1994), http://tools.ietf.org/rfc/rfc1633.txt.

local traffic. Both are likely to play larger roles in the future than they've played in the past.

7. The nature of "best efforts" transmission

The letter claims all Internet traffic is sent at the same baseline level of reliability and quality: "Thus the Internet is a "best-effort" service: devices make their best effort to deliver packets, but do not guarantee that they will succeed."

The *Internet Protocol* is actually more of a *minimum effort* system that lacks the ability to perform retransmission in the event of errors and congestion, but the Internet as a whole provides applications with very reliable delivery, at least a 99.999% guarantee. But it does this because all the networks cooperate with each other, and because software and services cooperate with networks.

The term "best efforts" needs a better definition because it means two things: 1) The lack of a delivery confirmation message at the IP layer; and 2) the expectation that quality will vary wildly. The first is a design feature in IP, but the latter is not. Variable quality is a choice made by network operators that is actually a bit short of universal.

8. What protocol layering really means

After discussing packet switching in a patently oversimplified way, the letter goes utterly off the rails in attempting to connect the success of the Internet to design principles that don't actually exist. It sets up this discussion by offering a common misunderstanding of network layering: "the network stack is a way of abstracting the design of software needed for Internet communication into multiple layers, where each layer is responsible for certain functions…"

Many textbooks offer this description, but network architects such as John Day dispute it.[12] In Day's analysis, each layer performs the same function as all other layers, but over

[12] Richard Bennett, "The Emerging Future Internet," *High Tech Forum*, November 17, 2015, http://hightechforum.org/emerging-future-internet/.

a different scope. That function is data transfer, and the scopes differ by distance. A datalink layer protocol (operating at layer two of the OSI Reference Model) transfers data from one point to one or more other points on a single network.[13] Wi-Fi and Ethernet networks within a home or office are layer two networks that do this job incredibly well; they're joined into a single network by a network switch with both Ethernet ports and Wi-Fi antennas.

A layer three network – using IP – transfers data over a larger scope, such as from one point on the Internet to one or more other points on the Internet.[14] This is the same job, but it goes farther and crosses more boundaries in the process. So layering is more about scope than function.

There is a second misunderstanding with respect to cross-layer interactions. The Letter says the code implementing layers needs to be: "flexible enough to allow for different implementations and widely-varying uses cases (since each layer can tell the layer below it to carry any type of data)." This description reveals some confusion about the ways that layers interact with each other in both design and practice. Standards bodies typically specify layers in terms of services offered by lower to higher layers and signals provided by lower to higher layers.

For example, a datalink layer may offer both urgent and relaxed data transfer services to the network layer. If the design of the protocol stack is uniform, this service option can percolate all the way to the application. So applications that need very low latency transmission are free to select it – probably for a higher price or a limit on data volume over a period of time – and applications that are indifferent to urgency but more interested in price are free to make an alternate choice. The actual design of the Internet makes this sort of choice possible.

---

[13] Gorry Fairhurst, "Unicast, Broadcast, and Multicast," March 10, 2009, http://www.erg.abdn.ac.uk/users/gorry/course/intro-pages/uni-b-mcast.html.

[14] Cisco Systems, "Introduction to IP Multicast," May 2005, https://www.cisco.com/c/dam/en/us/products/collateral/ios-nx-os-software/ip-multicast/prod_presentation0900aecd80310883.pdf.

9. Limits of the end-to-end argument

From the faulty description of layers, The Letter jumps right into a defective explanation of the end-to-end argument about system design, even going so far as to call it a "principle":

> In order for a network to be general purpose, the nodes that make up the interior of the network should not assume that end points will have a specific goal when using the network or that they will use specific protocols; instead, application-specific features should only reside in the devices that connect to the network at its edge.

There's a difference between "general purpose" technologies and "single purpose" ones that the letter doesn't seem to grasp. The designs of the datalink layer, network layer, and transport layer protocols don't assume that applications have the same needs from the network. Hence, the Internet design reflects a multi-purpose system designed to accommodate the widest possible set of use cases. This is why there is both a Transmission Control Protocol and a User Datagram protocol at layer four. The letter describes them accurately as a mode of transmission that values reliability and correctness (TCP) and one that values low latency (UDP).

IP Header Fields

Fields	No of bits	Remarks
VER	4	Version of the IP protocol in use (typically 4)
HLEN	4	• Length of the header, expressed as the number of 32-bit words • Minimum size is 5, and maximum 15
Total Length	16	• Length in bytes of the datagram, including headers • Maximum datagram size : 2^{16} = 65536 bytes
Service Type	8	• Allow packet to be assigned a priority • Router can use this field to route packets • Not universally used
Time to Live	8	• Prevents a packet from travelling in a loop • Senders sets a value, that is decremented at each hop. If it reaches zero, packet is discarded
Protocol	8	Identifies the higher layer protocol being used

28

Elements of Every IP Datagram

For these two transmission modes to work correctly, lower layers – IP and the datalink – need to have the ability to tailor their services to different needs. We can certainly do this at the datalink layer, which the letter fails to describe. Datalink services in Wi-Fi and Ethernet offer the options for urgent and relaxed delivery. These modes are accommodated by the Wi-Fi 802.11e Quality of Service standards and Ethernet 802.1D options.[15]

[15] IEEE Computer Society et al., *IEEE Std 802.11e™-2005: IEEE Standard for Information Technology Telecommunications and Information Exchange between Systems--Local and Metropolitan Area Networks--Specific Requirements. Part 11, Amendment 8, Part 11, Amendment 8,* (New York, NY: Institute of Electrical and Electronics Engineers, 2003), http://ieeexplore.ieee.org/servlet/opac?punumber=10328. IEEE

The IP layer's ability to request Quality of Service from the datalink layer is implemented by the Type of Service field in the IP datagram header, and also by the IntServ and DiffServ standards. ToS is also a property of routes, and is supported by ICMP, a required feature of IP routers. So the Internet is a multi-purpose rather than a single purpose network.

You can still stress your belief in the end-to-end "principle" without pretending that the Internet lacks the ability to tailor its internal service to specific classes of applications. All the end-to-end argument really says is that it's easier to develop applications that don't require new features inside a network. It doesn't say you have to pretend one size fits all.

In fact, the classic paper on end-to-end arguments (cited by The Letter) acknowledges the role of intelligence inside the network for performance reasons:

> When doing so, it becomes apparent that there is a list of functions each of which might be implemented in any of several ways: by the communication subsystem, by its client, as a joint venture, or perhaps redundantly, each doing its own version. In reasoning about this choice, the requirements of the application provide the basis for a class of arguments, which go as follows:
> *The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.)*

Computer Society; International Electrotechnical Commission; International Organization for Standardization; Institute of Electrical and Electronics Engineers; IEEE Standards Board, *IEEE Std 802.1D™- 2004* (New York, N.Y. USA: Institute of Electrical and Electronics Engineers, 2004).

> We call this line of reasoning against low-level function implementation the "end-to-end argument."[16]

So we don't need to pretend that a perfectly dumb network is either desirable or even possible in all application scenarios.

10. What makes networks "open"?

Let's not forget that the goal of the FCC's three rulemakings on the Open Internet is to provide both users and application developers/service providers with easy access to all of the Internet's capabilities. We can do that without putting a false theory of the Internet's feature set in place of a true one. The Letter offers a defective view of the Internet in terms of packet switching, congestion, router behavior, layering, and the end-to-end argument in order to support a particular legal/regulatory argument, Title II classification for Internet Service.

Accepting a defective view of Internet service makes it difficult, if not impossible, to optimize for openness, flexibility, innovation, and progress. Imposing Title II on Internet service suggests a desire to force application developers to adapt their products to the shortcomings of a rigid framework that forbids co-development. Instead of imposing a firewall between networks and applications, innovation-friendly regulators must permit network operators and application developers to create enhancements in network service that creates new opportunities for application development.

We know, from 40 years of bitter experience, that the one-size-fits-all model of Internet service fails to serve at least two classes of applications in an ideal manner: extremely latency-sensitive applications that need a more exacting service profile than standard web and streaming apps; and extremely price-sensitive applications, such as disk backup, which don't care at all about transmission quality. Service providers are not able to offer low-latency interconnections to and from their networks for a fee under the current

---

[16] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-End Arguments in System Design," *ACM Transactions on Computer Systems* 2, no. 4 (November 1984): 277–88.

regulations. This is a barrier to innovation that makes the Internet less open than it might be.

Title II may very well be a fine way to regulate the Internet in terms of the goals of keeping the regulator out of court and preventing obscure forms of abuse in non-competitive markets. But I'm suspicious of any argument for Title II that does violence to the nature of the Internet as The Letter has clearly done. There's a very straightforward argument for classifying broadband datalink (layer 2 of the OSI model) service under Title II even though Internet Service fits better in Title I. But as long as pure transmission – a wire that moves bits from the residence to the ISP edge without any cross-network routing – is not the regulated service, Title I is the better fit.

11. The differences between the Internet and POTS

The Letter insists that Internet and POTS are substantially the same because both networks connect endpoints, both networks relay through intermediate devices, and the paths through their respective intermediate devices are all somewhat dynamic.

The Letter alleges that because both networks exhibit these specific similarities, the only way to guarantee an open Internet is to apply Title II of the Communications Act (the portion created for POTS) to Internet service providers instead of Title I, the part created for information networks.

Because The Letter endorses Title II as the basis for Open Internet rules, I believe my summary of its argument is fair and not simply a strawman. This sounds ridiculous to techies, but the letter does betray a basic lack of insight about what the Internet actually is and why it's not a telephone network.

11.1. Functional and architectural differences

Does the Internet Have the same Architecture and Function as POTS? The Letter says "Yes, it does". For this argument to succeed, The Letter would need to show that the Internet has the same basic structure as POTS. More specifically it would have to show that Internet Service is unambiguously a telecommunications service that provides

nothing more than Internet Protocol transmission as a substitute for telephone transmission. In technical terms, The Letter's argument rests on its ability to show that the last 40 years of global and domestic approbations about the exceptionalism of the Internet over POTS are hogwash.

I don't believe The Letter succeeds on either count. In fact, I don't believe that any existing service definition on the books accurately captures the services the regulators say they want to control with open Internet/net neutrality regulations. But that's a blog for another day.

If I had to choose a regulatory classification for broadband information services on purely technical grounds, I'd go with the Title I Information Service classification. I select this option because I believe it was meant to apply across the entire end-to-end scope of information networks, including Internet service provider networks as well as platform networks.

If The Letter's general approach is a sound regulatory idea, it appears that UPS, United Airlines, and the city bus line should also be regulated under Title II of the Communications Act. Because, you see, their networks bear superficial similarities to POTS as well. But that's ridiculous as well.

11.2. POTS call routing vs. Internet packet routing

The Letter argues that the Internet is just like POTS because of the way the Internet handles routing (I(D)(2, "Internet Routing's Similarities to Telephone Call Routing"), page 10-11). Their proof of this conclusion is the fact that both systems currently employ dynamic routing of some sort.

The Letter bases this claim on a ten year old short journal article on the histories of POTS and Internet routing: "Routing Management in the PSTN and Internet: A Historical Perspective".[17]

The article recounts the transition from hierarchical routing to a limited form of dynamic routing known as "Dynamic Non-Hierarchical Routing":

> DNHR is the first implemented dynamic routing scheme. It was developed by AT&T and was deployed in AT&T's long-distance telephone network in 1984; it was later retired in 1991 when [Real-Time Network Routing] RTNR was deployed.

This is an interesting historical tidbit because it shows the Internet deployed dynamic routing before the telephone network did. The Internet – defined as a public or semi-public TCP/IP network – was first deployed in the 1970s over ARPANET. It was dynamically routed from the beginning. A limited form of dynamic routing was added to POTS nearly 100 years after telephone service was initially deployed.

POTS is still not fully dynamic in the sense that the Internet has been from the beginning. And when telephone network regulations were drafted, POTS was a hierarchical rather than a dynamic system.

But the static vs. dynamic routing issue is less important than much larger technical differences between POTS and the Internet. POTS routing is "stateful" while Internet routing is "stateless"; the two networks are built on entirely different technologies from top to bottom. But the term "routing" doesn't mean the same thing for the Internet as it does for POTS.

The Internet is a packet-switched network, but POTS is circuit-switched. This difference explains why the term "routing" means something very different in the Internet context

---

[17] Deep Medhi, "Routing Management in the PSTN and Internet: A Historical Perspective," *Journal of Network and Systems Management* 1, no. 15 (2007).

than in the POTS context. While the Internet routes *packets*, POTS routes *telephone calls*, as The Letter recognizes in titling their discussion with the term "telephone call routing".

Loading a web page requires the Internet to route an average of 2000 packets independently of each other. When a packet arrives at an Internet router, the router acts on it as if it were the only packet it had ever received and it processes it without knowing other packets exist. This process is known as "stateless routing" and it's key to understanding packet switching, the Internet's foundation technology:

> The Internet's basic protocol, the Internet Protocol (IP), is an example of a stateless interaction. Each packet travels entirely on its own without reference to any other packet. When you request a Web page from a Web site, the request travels in one or more packets, each independent of the other as far as the Internet Protocol program itself is concerned. (The upper layer Transmission Control Protocol – TCP – does relate packets to each other, but uses the information within the packet rather than some external information to do this.) The term connectionless is also used to describe communication in which a connection is made and terminated for each message that is sent. IP is connectionless as well as stateless.[18]

While each Internet packet is routed dynamically through a packet switched network, POTS approaches routing in a completely different way. Telephone networks are built on circuit switching, a network architecture that goes back to the days when phone calls were connected by a human operator plugging in wires on a board.

Since the dawn of the digital telephone network in the 1960s, phone calls have been connected by sending messages from the calling party's telephone office to the called party's phone office signaling the desire to connect a call. These call creation signals

---

[18] "What Is Stateless?," dictionary, *WhatIs.Com*, accessed August 29, 2017, http://whatis.techtarget.com/definition/stateless.

cause the network to find a route between the two offices and they also reserve bandwidth for the call in each telephone switch along the way.

Phone calls need consistent bandwidth at the same rate to have a decent quality signal. Once a call is created, the network guarantees that each digital sample (a bucket of bits that approximates a packet) will follow the same path from end to end as all of the other samples in the call. And it guarantees that each sample will cross the network at the same rate.

The overall technical concern of POTS is preserving call quality, while the Internet is more concerned about speed. POTS accepts very little variation from one unit of information to the next, but the Internet is happy with variation (AKA "jitter") as long as it means that some messages will be superfast.

11.3. Access to routing resources

The net neutrality debate is largely about the packet-level resource contention that takes place on the Internet. There is no meaningful analogy to this kind of struggle in POTS.

The circuit-switched design of POTS means that resource contention only occurs at the time calls are connected. Call creation reserves network resources, ensuring consistency. It also creates "state" in each POTS switch, reminders of all calls in progress.

Hence, POTS relays bits without the need to consult a network-wide routing table the way that Internet routers do. This process makes calls slow to connect, easy to maintain, and slow to disconnect.

The POTS paradigm limits resource contention to call creation, but the Internet model pushes contention down to the packet level. This means that packets contend for resources every time they're relayed by an Internet router.

The Letter claims Internet routers' "sole job is to send packets one step closer to their destination", but neglects to say that even this job requires routers to mediate conflicting demands thousands of times a second.

The Internet is more flexible than POTS and therefore demands a different and more permissive regulatory approach.

The Internet does have a feature that's something like a circuit-switched call, the so-called "virtual circuits" managed by TCP. But TCP virtual circuits are managed entirely by software under the control of Internet users and run on their equipment.

Routers control packet-level resource contention in IP, but POTS is designed in such a way that there is no resource contention after calls are created. So it's misleading to claim that Internet routing is pretty much the same as POTS routing. [Note: See comments by Tom Evslin on "call paths" in the Appendix.]

For The Letter to claim POTS routing and IP routing are pretty much the same is like claiming that there's no real difference between TCP and IP. TCP, like POTS, is about circuits; but IP is about packets, the Internet's real currency. Users control circuit creation in TCP, while the network controls it in POTS. This fact has major regulatory implications.

There is an interesting exception to stateless routing in the Internet, however. The Internet includes a protocol called Integrated Services (IntServ) that was designed to make phone calls work over the Internet.[19] (Fred Baker explains IntServ in our podcast.[20]) IntServ is used today in LTE networks for, you guessed it, phone calls. Like POTS, IntServ is stateful, reserving resources at each router and securing a fixed route through the Internet for each call. Like TCP, IntServ illustrates the fact that the POTS model can be added to the Internet but is not essential to it.

The Letter builds on faulty use of the term "routing" in another section of its comment, III(B) "The NPRM Displays a Disturbing Lack of Knowledge of How Data is Routed on

---

[19] Braden, Clark, and Shenker, "RFC 1633 - Integrated Services in the Internet Architecture: An Overview."

[20] Richard Bennett, *A Conversation with Fred Baker of Cisco*, High Tech Forum, accessed July 5, 2015, http://hightechforum.org/conversation-with-fred-baker-of-cisco/.

the Internet (and in the Telephone Network)". This part attacks the FCC for accepting the accuracy of the Stevens Report, hence it bashes two FCCs in one go.

The Stevens Report says:

> …Internet service providers do not appear to offer 'telecommunications,' i.e., '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,' to their users. For one, broadband Internet users do not typically specify the 'points' between and among which information is sent online. Instead, routing decisions are based on the architecture of the network, not on consumers' instructions, and consumers are often unaware of where online content is stored.

The distinction the Stevens Report is making is between basic mission of POTS vs. that of the Internet, with a secondary observation about the way the Internet works. POTS is about connecting people to people for the purpose of conversation.

Over the years, additional capabilities were added to POTS such as limited information retrieval, fax transmission, and communicating computer-to-computer via modem, but these other modes of interaction have had to conform to the conversation model. They also tend to be regulated as Information Services rather than as telephone services.

POTS bandwidth is limited to the 3 KHz needed for voice, and when a speaker is silent there's no saving unused bandwidth for future use; it's use it or lose it. Telephone calls are connected from one handset to another, or from one device that looks like a handset to the network as modems do.

So POTS is a network designed for conversation that has been extended to slightly different forms of communication that can attach to the POTS network as if they were telephone handsets engaged in conversation.

11.4. Connecting people to people as opposed to connecting people to information

The Internet serves a fundamentally different purpose than POTS, especially now that it's dominated by web applications. Users of the web want the network to connect them with specific items of information, such as pages, social networks, and entertainment.

Hence, the person-to-person, conversational style of interaction has given way to a person-to-web site form of interaction. We may use social networks to communicate with other people, but it's a very different style of communication than the one that typifies POTS.

With this new form of interaction come new forms of routing and transmission.

POTS "routing" is really connecting a phone call. Internet routing is something that happens tens of thousands of times for an information transaction that could only create a single "call routing" event for POTS. When the Internet routes a packet, it can hit 10 – 20 routers as the packet traverses from source to destination. But once a call is established for POTS, the path is static for the duration.

The path information takes between source and destination through the Internet is very flexible for two reasons:

1) The Internet user is interested in information, which can be copied to many places at the same time; but POTS calls connect people, each of whom is in one and only one place at a given time; and:
2) Because the Internet user is interested in information, both the location and the specific identity of the copy the user accesses are unimportant; for the POTS user, the location of the other party is unimportant, but their identity is very important.

There are many copies of the web pages, documents, and chat sessions the Internet retrieves, but there is only one human for a POTS conversation. We can accelerate the web through CDNs that place copies of information closer to consumers, but we can only

improve call quality by managing network traffic flows, reducing contention, and increasing bandwidth.

There is actually a very fundamental difference between placing a phone call and accessing a web page quite apart from the architecture of the network. The user qualifies as a single point in the legal definition of a Title II service, but there is no corresponding point in the Internet. The information we seek from the web is distinct from its mode of transfer.

A phone number is similar to an IP address in the sense that it must be globally unique. And like a phone number (something that stands for a device), an IP address stands for a point of attachment to the Internet. But how that connection is made and how information is exchanged differ greatly in each transaction.

Internet users don't connect to an IP address. An IP address is a pass-through to a final destination. We actually *connect* to services and information *behind* the IP address. And while a phone call is connected by the telephone network between two devices, a TCP virtual circuit is connected by one user's software to another user's software without ISP intervention. On the Internet, we all act as our own phone companies in the sense of making – or refusing to make – connections.

Therefore, to apply Title II to the Internet is not just to regulate ISPs, it is to regulate users and services as well. Regulating Google and Amazon as if it were communication utilities is clearly not the intent of the Communications Act.

11.5. Reiterating the Confusion between Calling and Routing

The Letter capitalizes on the confusion it created earlier about call creation and packet routing:

> Saying that Internet users do not specify the points to which information is sent online is like saying that telephone users do not specify the phone they want their call sent to when they dial a phone number. As explained in Section I.D.2, both the Internet and the telephone network make use of dynamic routing based on the

> architecture of the network. Further, in both networks the customer is often unaware of where the endpoint is actually located—particularly in mobile networks, where a phone customer may have absolutely no way of knowing, a priori, even what country a mobile phone might be located in.

Title II is about the classical POTS network, not the mobile phone network; Title III covers wireless communication. Hence, it's not legitimate to argue that the Internet belongs to Title II because it displays characteristics of Title III devices; but The Letter does it anyway.

The Letter also tries to muddy the waters with yet another conflation of packet switching with call creation:

> Thus, this interpretation of what it means to transmit information between or among points specified by the user, i.e. that the user must explicitly tell the network what routing decisions to take, has no basis in reality. Taken to its logical conclusion, it would require the FCC to similarly decide that telephone services are also not telecommunications services—an obviously absurd conclusion.

We don't have to tell POTS anything about routing, we simply provide it with a number that uniquely identifies a device. In traditional POTS terminology, call creation is called "routing". This use of the term "routing" creates confusion in understanding the role of DNS in the Internet.

One the one hand, advocates of treating ISPs as Title II carriers compare DNS to directory service, and on the other hand they treat it as a network management function that makes the network more efficient. These two characterizations can't both be correct.

11.6. Impact of the Confusion of Calling and Routing on Regulation

During the DC Circuit's oral arguments on the Title II challenge, judges took DNS toward the directory service analogy before describing it as having a network management function. Here Judge Srinivasan uses the directory service analogy (transcript page 11, line 9 – page 12, line 23):

> KEISLER: — and when you have two services, caching provided by us, caching provided by somebody else, DNS provided by us, DNS provided by someone else, if they are performing exactly the same functions in conjunction with exactly the same service is can't be that they're telecommunications management with respect to one service, and an information service with respect to the other, because —
> JUDGE SRINIVASAN: But why not, because — why not under the language of the statute? Because the telecommunications management exception speaks in terms of the use of such capability for the management of a telecommunications service, and obviously can't be for the management of the telecommunications service if it's not the person who's, the entity that's providing the telecommunication service is providing the capability –
> KEISLER: Well, I think Your Honor is right, that obviously it can't be if someone else is providing it, but what follows from that? We certainly would not let a third party manage our telecommunications network, and so our point is that if customers can choose a third party to provide exactly the same function in connection with exactly the same service, and it's not management when a third party does it, it can't be management when we do it, either.
> JUDGE SRINIVASAN: So, is there a phone analog? So, for example, I think directory assistance was part of the basic service —
> KEISLER: Yes.
> JUDGE SRINIVASAN: — and now if a third party comes along and says, you know, I've got a better directory service, better directory assistance program because I've done stuff that makes it faster, I've got a better list, what have you, and then they came up with that as a competitor and say call my line instead to get directory assistance, it seems to me what would have happened is, although maybe it's not borne out historically, that that would have been deemed an enhanced service, or information service, but the fact that the phone company had its own directory assistance wouldn't preclude there from being that dichotomy.

If in fact the function of DNS were simply to help broadband networks make basic connections to web sites and streaming services, it could easily be subsumed in the

historical notion of directory service. A telephone network without a directory that maps names to numbers isn't very useful because POTS is all about the customer at number X calling the customer at number Y.

But DNS is more than a directory service. It's a general purpose database that networks not managed by an ISP can use to make their internal operations more efficient and to enable customers to retrieve information. DNS doesn't support the transmission function as much as it supports an information storage and retrieval function. DNS does more for Netflix than it does for Comcast.

This latter role distinguishes it from directory service. The FCC's attorney was forced to misrepresent the function of DNS in order to support his client's characterization of DNS as a network management function:

> SALLET: DNS is a routing function, it helps move traffic to the destination intended by the consumer; caching is storing of preexisting content, no original creation of content, in a way that also facilitates routing, and that allows the network, for example, to allow costs that would be incurred if it had to go to a further distance to get that preexisting content, both, therefore, fall easily within the telecommunications management exception.

DNS certainly does support call creation in the same way that directory service does for POTS. But it does not move IP packets from source to destination, nor does it help move them. Packet transmission is the job of IP, and mapping the Internet for IP is the job of BGP, not DNS.

But DNS does more than support call creation, and its additional functions are not in support of ISP network management. To the extent that DNS is a tool to be used by Netflix and similar firms to connect customers to content, it's a Title I information service that helps services find information for their customers.

And insofar as DNS is part of the bundle of services offered by ISPs to their customers, that bundle is more than simple transmission between IP addresses of the customer's choosing.

11.7. The Disappearing Internet Service

The Letter claims that Internet users change DNS providers in order to make content move through the Internet faster (page 26):

However, as explained in Section II.C., ISP-provided DNS is by no means necessary, and ISPs are often not the best at providing the service. Users can and do change their DNS provider to lower round-trip latency and thereby have faster overall Internet service.[67]

The footnote goes to a blog post that compares the latency of several third party DNS services to each other; in no way does it support the claim that any third party DNS makes traffic move faster through the Internet overall.[21] Again, The Letter is attempting to capitalize on the fact that the term "routing" is used in descriptions of POTS call creation and Internet packet transmission to mean utterly different things.

The Letter expands on this ambiguity in order to reduce the Internet's essential function as a network of information to that of a network of transmission. The NPRM asks (§29):

> We believe that consumers want and pay for these functionalities that go beyond mere transmission—and that they have come to expect them as part and parcel of broadband Internet access service. We seek comment on our analysis.

And The Letter answers:

> The NPRM is correct that consumers want and pay for these functionalities, i.e. routing. But the NPRM is incorrect when it says that routing is not part of mere transmission. To say that making routing decisions is not a necessary requirement for "mere transmission" implies the FCC is still living in a world where people dial a phone number by verbally asking an operator to connect them to a specific line. No modern telecommunications network, be it the PSTN or the Internet,

[21] Young Xu, "Comparing the Performance of Popular Public DNS Providers," *ThousandEyes*, May 15, 2017, https://blog.thousandeyes.com/comparing-performance-popular-public-dns-providers-2017/.

> requires (or could conceivably require) an end-point (including a BIAS customer) to know the details of how the network is laid out.
>
> Simply put, customers expect their ISP to route their traffic as "part and parcel of broadband Internet access service" because it is a necessary functionality for the transmission of data in any modern network.

An Internet service consists of basic transmission as well as access to information. The line between these two functionalities is subtle and can easily be confused. We know that dialup ISPs have always been treated as Title I information services even though the networks to which their modems attach – POTS – have always been treated as Title II services.

So the question the FCC is asking is what happens to the Title I part of traditional ISP service when POTS gives way to broadband. One answer to this question is to delineate the role of broadband transmission to the scope of transmission between the customer premises and the ISP premises. This matches the role played by Title II in connection with dialup ISPs.

There is no routing between an Internet user's cable modem or smartphone and the ISPs property. This is nothing but a Link Layer network connection identified by Layer Two medium access control (MAC) addresses. The Internet Service *per se* is carried out in equipment residing at the ISP offices.

The ISP's equipment routes information dynamically between its network and the rest of the Internet according to BGP route exchange and network management. BGP doesn't propagate routes to our homes.

12. Conclusion

The Wheeler FCC didn't actually define the broadband transmission service in a manner analogous to Title II transmission services used by dialup ISPs. Rather, it chose to apply the "transmission service" label to the IP datagrams that define the internetworking layer of the Internet. This is a curious move because it presumes that IP datagrams are self-transmitting. In fact, no IP datagram gets anywhere unless it is carried as the payload of a

Link Layer information frame. So the 2015 Open Internet Order labels a *non-transmitting service* a *transmission service*. This is a clear error.

The NPRM, on the other hand, treats Internet Protocol datagrams as part of the Information Services ecosystem that makes the Internet something people care about. It asserts that IP provides access to Information Services and also that IP depends on broadband technologies for simple transmission between the customer premise and the ISP as well as between ISPs and unaffiliated Information Services.

One of the more telling asides in The Letter concerns the evolving telephone network (page 26):

> Thus, this interpretation [the observation that consumers are often unaware of where online content is stored] of what it means to transmit information between or among points specified by the user, i.e. that the user must explicitly tell the network what routing decisions to take, has no basis in reality. Taken to its logical conclusion, it would require the FCC to similarly decide that telephone services are also not telecommunications services—an obviously absurd conclusion.

The Letter is forced to acknowledge that today's telephone network bears little resemblance to the POTS networks of the '30s or even of the '90s. Mobile phones require massive network intelligence to support consumer traffic, whereas, the "routing" to telephone numbers that represent wires is simpler than routing to portable numbers on foreign networks. Modern directory service isn't much like picking up the phone and asking Ernestine the Operator for a number either.

It strikes me as reasonable to consider the question of how well the statutory definition of POTS in Title II of the Communications Act even applies to modern telephone networks. Many seek an update to the Act because it's so out-of-date and ambiguous.[22]

---

[22] Richard Bennett, Jeffrey A. Eisenach, James K. Glassman, Bronwyn E. Howell, Justin (Gus) Hurwitz, Roslyn Layton, Bret Swanson, "Comments on Communications Act Modernization."

If Title II is too narrow to even describe modern telephone networks, it's obviously not a good fit for much more powerful Internet services. This is the key takeaway from the search for similarities between the Internet and POTS.

Appendix: Internet Pioneers Discussion

On August 15, 2017 a group of Internet pioneers discussed the issues raised by the NPRM in general and the EFF letter in particular. The group consisted of:

- Tom Evslin, the founder of the Internet's first telephone carrier, ITXC.
- John Day, manager of the OSI Reference Model and author of Patterns in Network Architecture: A Return to Fundamentals
- Barry Shein, founder of The World, the first commercial Internet Service Provider.

The discussion was moderated by Richard Bennett, co-creator of Ethernet over Unshielded Twisted Pair and Wi-Fi. The discussion is available on YouTube and Facebook in video format and at iTunes, Soundcloud, and High Tech Forum in audio format.[23]

A rough transcript of the discussion follows.

High Tech Forum Webinar

The Architectures of the Internet & Telephone Networks — A Primer on How They're Different

Date: Tuesday, August 15, 2017 | 11:00 am- 12:00 pm EDT

Richard Bennett, High Tech Forum*:* I pulled this group of people together today because I wanted to talk about some of the issues in the FCC's restoring internet freedom docket which is essentially the proceeding that asks the question of whether Net Neutrality legislation should be reformulated under Title I away from their current Title II status. One of the arguments that has been made for that on a technical basis came from the Electronic Frontier Foundation (EFF) who assembled a letter and had it signed by a

[23] Richard Bennett, "Internet Pioneers Discuss Network Architecture and Regulation," *High Tech Forum*, August 16, 2017, http://hightechforum.org/internet-pioneers-discuss-architecture-regulation/.

pretty large group of internet technologists. Which argued as I read it, that internet service providers should be regulated under Title II because the service that they provide is very similar to the telephone network. In fact, they even went so far to say the networks that these providers operate work pretty much the same way that the telephone network does. So that is what we are going to explore. Let me give you some bios for folks who aren't familiar with the panelists. Tom Evslin, was the founder of ITXC which was a really innovative company in the late 90s early 2000s that transported telephone calls over the internet. In 2002 they were rated as the fastest growing tech company in the country by Deloitte Touche. Tom has also been a chief technology officer for the State of Vermont and currently runs a very interesting business that transports natural gas over trucks in areas that don't have pipelines. Which I find kind of oddly similar to the kind of thing that ITXC was doing, in transporting telephone calls over the internet. Tom was also responsible for the conception, launch, and operation of AT&T's first ISP Worldnet Service. You can read more about Tom on Wikipedia, he has a page. Barry, who just joined us is the founder and operator of the world's first commercial ISP, he has been named one of the 11 people who made the internet happen because of his ISP which is called The World. Hence Barry's Twitter handle is @World_Leader – maybe a little lounge in cheek. There's a page that BT put up of the 11 people that made the internet happen in which Barry is included. It shows an ad for The World Service from 2010 with Gold Accounts that had for 20 bucks a month you got email space, web space, your own website, two mailing lists, and Unix shell access – that was quite a bundle of services. And, that's still on offer today right, Barry?

Barry Shein, President and CEO of The World: It is, we still have dial-up if you want it, not sure for how much longer.

Richard Bennett: And, incidentally the other two panelists have both been Barry's costumers. So, your impact is pretty amazing. Barry was also on the board of directors of Usenix which played an interesting role in creating UUNet which was one of the first internet commercial backbones that sort of helped in the transition of the internet from a research network to a network for the general public and for commerce. John Day is a lecturer in computer science at the Boston University Metropolitan College, computer scientist, and internet pioneer and a historian and quite a historian at that. You have

probably heard people talk about the 7-Layer Model for network architecture well John's the guy who put that together.

John Day, Computer Scientist; Internet Pioneer; Historian: That is a bit of an exaggeration –

Richard Bennett: John chose the guy who put it together, he was the rapporteur. We are going to talk about protocol layers after a while and I guess John's about as much of an authority on that as much as anyone can be. He wrote the book in 2008 called *Patterns In Network Architecture: A Return to Fundamentals* which I believe is the best by far, the best overview of network architecture as well as the politics of the development at the ISO protocols in the 1980s that at the time were believed to be what the internet was going to be based on. So the TCP/IP internet that we use today there was sort of a consensus position in the computer industry that was a transitional stage that we were going to pass through and that OSI was going to be the real thing, but politics got in the way of that and a bit of overdesign. John has also published articles on the history of cartography and if you want to talk to someone about the map not being the territory, you know he's your man. And the author of four early RFCs in the 3 Digit number range 728, 731, and 732. He publishes articles on cartography still - and which auction house?

John Day: Christie's.

Richard Bennett: Christie's – yeah. On the authenticity of some of these maps, people like to buy. And John has a Wikipedia page too. Folks, thanks for agreeing to do this, and I am kind of jazzed this thing seems to be working which is amazing to me. But, tell me you have all read the EFF letter…should we sign?

Tom Evslin, Founder of ITXC: EFF is just wrong, about facts and I disagree with them on opinion. And it is hard to say that someone is wrong on an opinion. The internet and the phone network work, as everyone on the panel, knows in very different ways – in almost diametrically opposed ways. I am very aware of that because the company I built which as you mentioned was ITXC which was a wholesaler or a Voice over IP, we were in the middle of the phone call used the internet to replace the PSTN, the Public Switched Telephone Network and our success depended upon the fact that our technology was completely different, the routing was completely different. Just to give a tiny example: a

phone call travels over something called a Call Path which means every bit of information goes along the same path as every other bit or else the call breaks. Well, no internet message, including Voice over IP is restricted to a single path.

Packets fly all over the place, and then get reassembled at the destination. Part of the reason for the strength of the internet, for the capacity of the internet and one of the reasons why the internet is not nearly so much a natural monopoly – as in regulation - as public switch telephone was. Just from a technical view, I don't want to over emphasize it because any technologist understands this, the internet just doesn't look like the phone network. More import, nor should it be regulated like the phone network under Title II of the communications act of I think it's 1935 or some time ago, 34, and my answer is definitely *No*.

Look at what we got from Title II regulation, we have a phone network that since its inception was a miracle of modern design that in the ensuing 50 years had about two innovations, dial tone and 800 numbers and regulation was one of the things, monopoly was another, that almost guaranteed that public switch telephone networks would be innovation free and network was innovation free. And, if we look at legacy today if you look and you happen to have an old fashioned phone on your desk, can it send text messages? No.

Can it do any of the things a smart cell phone can do? No. That's partly a result of monopolies and of regulation. If we look at smartphones and all of the wonderful things it can do and all of things being innovated there, that's an example of unregulated technology being able to grow more quickly. You would think that a desk phone where you don't have to worry about weight where you don't have to worry about power would be able to do much more than a phone that has to be small enough to be clipped on your belt, that has to worry about battery power, that's not the case. If we want the internet to be innovation free we can regulate the internet under Title II as we did the phone network or we can do what we have done for most of the internet's life up until 2015 which is have a very, very light regulatory touch.

This argument, and I am giving you a long winded answer, but this argument is very strange. Because back at the end of the 90s because my company depended on the

internet being unregulated the telephone companies, some of the same companies that are now ISPs were eager that the internet be regulated so that innovation on the internet wouldn't interfere with their business models. So I spent a lot of time lobbying the FCC and at the time internet players like Google, and Amazon were very much on the same side that I was, saying that regulation would kill innovation on the internet and that is what the proponents of internet regulation wanted to happen. Well, first Clinton's FCC under Reed Hundt decided not to regulate the internet that got formalized later in the George W. Bush administration with so called Pulver Decision to specifically not to regulate the internet. And, so within two years the cost of phone calls to China went from two dollars a minute to two cents a minute because we can use innovation on the internet to break the monopoly of the regulated telephone networks. You notice the regulation protected the high cost, the regulation didn't prevent the high cost and the high prices.

Now, fast forward to 2015 and companies like Google and Amazon who are relatively much more powerful than the ISPs are in favor of regulation I think maybe because maybe now they are the ones who think they can capture the regulators the same way you could argue the telcos did earlier and because now they're the incumbents and I think they're the ones who are afraid of disruptive innovation. This is all done in an Orwellian sense in the name of Net Neutrality and a free internet – and I am all for a neutral internet with small ante – but to have that accomplished under Title II regulation means I am afraid that we end up with both an innovation free internet and incidentally a great danger of government intervention in content it amazes me that some of the people who are most opposed to the current administration are in favor of having Presidential appointees under Title II be the ultimate deciders of what can be done on the internet. I will stop it, and give somebody else a chance.

Richard Bennett: So what do you think Barry, is Tom right?

Barry Shein: Well, I read the letter – the EFF letter – the first thing that impressed me was that everything that was said in there could have been said in a lot less than 50 pages. Okay, it's unfortunate because it struck me as written by somebody who desperately wanted to believe that nobody, somebody at the FCC doesn't fully understand the internet which I think is a terrible approach – to lecture and give definitions for 40 pages you

don't understand this, you don't understand this. When I would find a sentence I didn't agree with – make your point –

Richard Bennett: I thought it was kind of bizarre that when they went to define packet switching the reference they used was a law school textbook – it strikes me that when you ask engineers to define something that's *that* fundamental to the internet they're not going to refer to a law school textbook– it was a good law school textbook.

Barry Shein: And who knows maybe it was quick – but I am just saying dragging through that and saying: Who are you brow beating? Who's your target audience here? This is the type of thing that some engineers will do, challenged on a policy issue they will assume not that you disagree with them but that you must not understand and we will start explaining the issue to you and that is a terrible approach to the policy environment.

Richard Bennett: Well actually that part of it, what they call the Brief Summary, the Brief Summary of the Internet, introduction to the internet or whatever it is, the internet for dummies part of the letter it was recycled from an amicus brief that EFF filed in the US Telecom challenge to Tom Wheeler's open internet order. So, they added the section at the end about DNS but the rest of it was pretty much some material they had developed in a different context but they did add the stuff about how "Chairman Pai doesn't understand the internet."

Barry Shein: Well it invites criticism of their policy position based on going through and seeing if they're correct in every detail in their explanation and do they understand any of that, they set that up as the contest. Which side understands the internet? Which left me kind of empty. Talk about the policy, I understand the idea of– innovation – and that's all good stuff right? That's something that policy makers may respond to. They can understand, the commission wants innovation. I think they got better when they started talking about bandwidth shaping and how this may impact or intrusion/interference, where packets are intercepted and changed for vendors –

Richard Bennett: Which is something that pretty much something that hasn't happened on the internet there have been some mentions of it as far as the paid fast lanes and all that the closest we have got to that was zero rating, people that are mobile customers could access certain services like Facebook without them counting against their minutes.

Barry Shein: Zero Rating as I understood it was originally introduced in Africa to support underdeveloped nations. Because these people can't afford a nickel of time and it is all charge, there's no such thing as bundled. For many of them, it is basic cell – I don't know what they get.

Richard Bennett: Its mostly something that seems to have great appeal in countries where the only communications device that is sort of widely deployed is a feature phone and you're trying to move people up from the feature phone to the smart phone and you know cost is an issue and I think it's like the old strategy of the first ones free. You got to give people some incentive, like what is the internet good for. If you don't live in a first world country and use it every day it's kind of hard to remember, what is it good for?

Barry Shein: Well that's the concern with zero rating. Even in some undeveloped countries is it gives an unfair advantage to deep pockets because they can get away with it. Because when you have no money you are going to go to whichever vendor is going to give you the most free minutes.

Richard Bennett: On the regulation point, as the earliest ISP you weren't regulated under Title II, you actually founded The World before the 96 Telecommunications Act so there really wasn't such a thing as Title I and Title II. But, the FCC had done their computer inquires and created their distinction between basic service and an enhanced service and your ISP was concerned an enhanced service and therefore deregulated service, right?

Barry Shein: Right, it was several years. Obviously it was 1989 and we were completely off the radar as far as any of that went. The only thing that even resembled regulation was the first reaction from the national science foundation, the NSF Net people that two-thirds of it blocked me. That's all – they just blocked me in their routers. They said "this guys just letting the public on he can't do that." Imagine that, being blocked for letting the public on the internet. In fact, two years later that resulted in a letter from Stephen Wolff and the Committee giving me permission to put the public on the internet. I actually have a letter giving me permission to put the general public on to the internet.

Richard Bennett: So, permissionless innovation wasn't a thing in 1989?

Barry Shein: Well, permission to get their people, you know I could get to whoever wanted/willing to which was one third of the internet and it grew, people got used to the idea. Of course, other ISPs founded it pretty quickly and they were all in favor. So, slowly but surely that fell away these were just most –

Richard Bennett: There was a bit of a backlash when AOL allowed all of their customers on the internet because they flooded into UUNet and the IQ level of the average UUNet user dropped substantially –

Barry Shein: We used to talk about the eternal September which was a term from UUNet which was all of the students returning in September and they all get UUNet access. So you all of a sudden have a couple million new users rushing a chat – so now AOL gave us the eternal September it's no longer just in September. How are we going to explain anything to these people? It worked out. I wasn't too concerned. It was a funny observation.

Richard Bennett: So John, you were chuckling when Barry was talking about the acceptable use policies on the NSF Net and you were on ARPANET, you were like the 12th node on the ARPANET from the beginning. So, on Flag Day you were there – yeah you were on well before Flag Day.

John Day: Yeah Flag Day wasn't until 83 –

Richard Bennett: Oh yeah, that was 12 nodes, but this was another 12th node.

John Day: Yes, this was 1970. Well, I would have to agree with Tom because if you really go back and look at how packet switching came about Paul Baran's whole idea was this was a *not-voice* network. You know? In fact it was designed not to do what voice did. So, that's one thing. Then of course the next progress of that was with Pouzin coming up with the datagram, and actually I think it was fairly common in those first few years the four characteristics that I have always said characterized the paradigm shift that we saw networking taking which the internet builds on was that we had layers of different scope, distributed shares of different scopes it was a peer network there was a movement toward versus nondeterministic versus deterministic technique and not just in networking with the datagrams but this was the same time when operating systems

became interim driven as opposed to polled, and Ethernet was another good example of that –

Richard Bennett: And the rate concept in storage –

John Day: Right, and then, of course, the other one was, we all saw this as distributing computing, not telecom. And, in fact it has been really interesting. We've said that networking is inter-processing communication and lately we have gotten a reaction that, that it is a radical concept when in the early 70s nearly everyone thought that networking was operating system inter-processing communication.

Richard Bennett: Yeah, I think Bob Metcalfe was one of the people that made that point, early on –

John Day: He and actually Dave Walden has an RFC of even smaller – that calls it IPC, you know? And, so yeah it was a common idea.

Richard Bennett: So the notion that protocol layers have some sort of a functional difference at each layer in the protocol stack – I mean EFF makes that argument that there's layer two which they don't talk about, layer one and layer two they don't talk about at all. I think that would basically destroy the argument. They are talking about layer three, which is IP and they might make some mentions of application layer and there's some passing stuff about TCP. But, they make this argument that these layers each do something functionally different which is sort of a lawyer's view of the layers, right? That's not the way you see them is it?

John Day: Well that was a very early idea and that Idea came from Dijkstra's paper on the THE operating system and actually when I went back and started looking at it and by the late 80s or so we had been working this problem for some time and we always had problem that there were always functions that repeated. And so, that was breaking and at one point I just realized no, no, no the layers all do the same functions but they do them for a different range of the problem and are over a different scope. Actually the CYCLADES guys by 1972 had figured out that really the primary property of a layer is scope –

Richard Bennett: And CYCLADES by the way is the French network that is immediate precursor to TCP/IPs organization of functionS – unreliable datagram layer and an end to end transport layer that –

John Day: Right, and I think actually part of the thing that the internet people missed was in CYCLADES the link layer was an HDLC like protocol because lines lot of the time were unreliable so you did point to point error control so it was really reliable point to point. The network layer did relaying and would lose to packets due to congestion and then you had the transport layer. You know, this was Louis's big innovation that he realized that the users were never going to trust the network anyway so you didn't have to be perfect you only had to give your best effort to have an end to end transport protocol that recovered from errors.

Richard Bennett: Yeah you still need to do error checking in a computer because of memory failures, there are things that can happen in the computer that corrupt the data.

John Day: But the point is, layers only have a certain scope and really to talk about stacks is a misconception because not all systems have all the layers.

Richard Bennett: It is kind of more like an onion right?

John Day: It's like an onion with multiple centers.

Barry Shein: Can I interject something? If you're going to use – this is more my issues with Net Neutrality which is okay lets hold on to that later model and one of my problems is that Net Neutrality can go on every layer, there's Net Neutrality issues going on at every layer and they simplify it mostly in layer three, the IP layer, right? And I think there missing how many ways determined people with money could interfere with transmission and that's the fatal flaw on a fluffy model. Look at a connection model, look at a bandwidth model.

John Day: Then they would have to admit that there was processing going on. They would have to admit that there was computing going on. If you'll notice in the letter they talk about BGP being an application protocol and doing routing but they don't talk about inter-domain routing that uses the link layer to transport stuff around. In fact routing is

not an application protocol its layer management and so they leave it out because it doesn't fit their model. Even DNS is not an application.

Richard Bennett: Yeah I mean the protocol that actually does the routing is IP. BGP provides the information to routers to use to figure out –

John Day: Yeah I think that is another misconception that we have we talk about IP being stateless and that its resilient to failures which Tom was pointing out and that's all true but the reason it is true because in that model everybody knows how to route everything in the phone company model only the edges know how to route everything.

Richard Bennett: Well the only thing that gets routed on the telephone network is the call creation –

John Day: Well the thing is there is a much wider distribution of state information in that packet switch network.

Richard Bennett: Yeah I mean essentially the routing table defines the state of the network at a certain point in time and you get the resilience from the internet because every time a router forms a packet it goes and checks the routing table to make sure the state is the same as it was the last time. And so, it can pick up on changes in the path architecture from point A to point B on a fairly speedy basis.

Barry Shein: But getting a back to Net Neutrality that's a great example: BGP right? How much mischief can be done at the BGP level and it brings up the real issue here the real problem here independent of the FCC's approach of this "Regulation" like they could regulate this which is that we have the same companies providing core Tier 1 BGP as they are providing the applications which is creating the interest in this mischief – the monetary interest in all of this right? They can cause as much damage with BGP as they can with bandwidth or fooling around with peering is a wonderful tool for that and I am saying all the way up the stack you can find this kind of problem. You just have to stop for a second and think. DNS, I could slow down DNS requests you know? I can do all sorts of mischief for my competitors and you know? It's not really bandwidth per say its priorities perhaps on the DNS servers they rely on.

Richard Bennett: One of the first events that actually brought Net Neutrality to the public's attention was that Comcast/BitTorrent thing, as that issue has been explained to me what was actually happening was that Comcast had just introduced their voice service and right after they introduced their voice service and they're now the biggest voice provider in the country a lot of people who were Vonage users who were complaining that their Vonage wasn't working as well as it used to be and they suspect it has – as one would, that Comcast is deliberately degrading their VOIP so they would have to buy voice service from Comcast.

Which as you know that is not the best thing for public relations and certainly if it was happening they would want to deal with it but the engineers looked into it and found well actually the problem is the people who are complaining are in neighborhoods where there is a lot of peer to peer going on and so before peer to peer there really weren't applications that stressed the uplink because web browsing is like 100X more downlink traffic then uplink traffic you just click on the link and you do the get and it's like one or two packets and now you're downloading, it's an average of 3 megabytes per web page, but BitTorrent was able to use – to saturate the uplink in a way that it had never happened before they had to figure out how to throttle that and they didn't have a good tool to do it.

Well this is the thing the telephone network is designed in such a way that callers don't have an impact on each other. I am calling my neighbors calling we have a defined segment of network resources that we can work with at any given time and we don't really have side effects on each other except in the one case where everybody wants to make a long distance call on Mother's Day that you may not get your call connected, but once your call is connected there's no side effects between users, but the internet is all about breaking down those barriers between the calls so that the whole capacity of link is available to like whoever wants it at any given sub second interval, right?

Tom Evslin: A very good example of that difference was on 9/11 when the switches that were Verizon switches that were underneath the World Trade Center which were destroyed. So the public switched telephone network pretty much went out of business in the New York area what remaining – it had some capacity to route around but what capacity there was totally saturated which meant you couldn't make a call at all. You

couldn't get a dial tone. Which is the way sources were allocated were a wholesale Voice Over IP provider we never did domestic calls because we weren't competitive but everyone's switches were attached to us and all of a sudden at ITXC we're carrying that domestic traffic that couldn't be routed through the Verizon switches we were providing lousy quality because the internet was congested but on the internet it bends but it doesn't break. On the PSTN it doesn't bend but it breaks. So I wouldn't boast about the quality of the calls we got through but the internet as whole not just Voice Over IP but email and other forms of messaging on the internet let a lot of people know that a lot of other people were okay were the PTSN because of its inflexibility, and interestingly because its insistence that you get absolute quality or you get no quality at all just suddenly broke under the load.

Richard Bennett: And a lot of those are just one bit messages right? Like "I am okay."

Tom Evslin: That's a few more than one bit but its 140 characters anyway.

Richard Bennett: Well John you've said that it's your kind of impression that the internet is actually too much like the telephone network and because, right? It's a subtle point do you want to explain it?

John Day: I am not sure what you are referring to?

Richard Bennett: The way the internet does routing, is not as dynamic as it could be.

John Day: No, no, no the whole connectionless thing is a lot more resilient. And as Tom was just saying it bends but doesn't break as opposed to the phone system which doesn't bend but breaks. I think the point that I was making was they were pointing at BGP to push the argument that they wanted to make but if you looked at were all of the routing occurs on the internet that argument wouldn't hold. Routing has always been considered layer management part of the layer.

Richard Bennett: Well, routing – that word doesn't really mean the same thing in PSTN context, and in the telephony context that it means in the internet context. Because the thing is the telephone network routes is calls. So when I dial your number my request to connect to your telephone gets routed over some path that is chosen by the telephone

network to be the best path for that call request. But once we are connected it doesn't change. It's like Tom was saying about the paths, it goes over the same path–

John Day: It is funny, and I think Tom hit on a good aspect of this which is I think the incumbents right now believe that the Internet – they don't want the internet to change they really do want to stifle innovation in the internet. They want all of the innovation to be in the applications, the trouble is we are building on sand now. When I started to digging around to see what the fundamentals were I really was fairly of shocked. Well, some of the stuff we already knew about we knew that we were missing two thirds of the addressing structure that's required and we've known that since the 70s.

I was really shocked that when I went to write a paper about the economic policies of what was going on in the 70s between the phone companies and the computer companies - and you forget after 30 or 40 years – and I realized that in the mid-1970s the idea that everyone was pushing toward was a transport layer with transport layer addresses over a network layer with network addresses of less scope. Because it was a mobile network, and a data link layer with data link addresses over even less scope. You had three layers all at different scopes all with their own addresses, we don't have that today.

Richard Bennett: Right, we have one address that basically – the IP address which is globally unique doesn't have any structure to it particularly. It's kind of a reflection of the MAC address which even MAC addresses have 48 bits, and IP addresses only have 31 that are usable, and yeah it stands for the same thing. It identifies an interface it does not identify –

John Day: That's the second problem it identifies the interface not the node –

Richard Bennett: Well the Xerox systems they got that right they address the network and they host independently and if you were a host you could have more than one connection. You could have an Ethernet you could have a couple Ethernet addresses and something else.

John Day: Everybody got that right beside the internet.

Richard Bennett: I mean yeah, DECNet and CYCLADES. Were they being too clever? This was back in the day when you had to twiddle bits. You could take an IP address and

then based on what the two bits were, it told you how many bits in it identified the network and how many identified the node. So that sort of had a network address embedded in there but then that kind of went away -

John Day: Well no, what had happened was the ARPANET your host address was your import number interface and the internet just did the same thing. And they never changed it, I mean we realized that there was a mistake in 72.

Richard Bennett: With the Tinker Airforce Base? – the Air Force base wanted redundancy they kind of bought into the idea that this packet switch network would be resilient to attacks right? So they figured if we are going to do that we are going to need more than one connection, so why have something to fall back to. Well it turns out if your connection is identified by the same thing that your computer is identified by you can only have one connection. So we have sort of taken this wonderful, resilient technology and sort of dumbed it down into something that is less than what it could be. Right?

John Day: Well first of all as I always say, the ARPANET – nobody had ever built one of these things and we had a lot of problems to solve. Like, just having to move data because none of the computers we were connecting had ever been built to talk to another computer. And really to some extent why you often see in software that the first one is a throw away we're sort of removed from the physics and what the principals would turn out to be that you really sort of have to build one to see what you should have done. In fact, the BBN guys spent a lot of time with the CYCLADES guys working with them so they wouldn't make the same mistakes they did.

Richard Bennett: Well this is a bit of the history that isn't very well understood I think that the path from Arpanet to the internet is not a single hop path. Because what had happened was the guys at BBN in Boston that created ARPANET after it had been in operation for a few years that had a good grasp on what the limitations were and Louie Pouzin and his research group in France they were ARPANET users and Louie decided he wanted to build a network that addressed those shortcomings.

John Day: The thing was the ARPANET was built to be a production network to lower the cost of doing research amongst ARPANET contractors on other things. Okay, so

BBN was limited in how much they could change Pouzin built CYCLADES to be a network to do research on networks so he had free reign –

Richard Bennett: People are always complaining about the network is down because someone is running an experiment I need to send this email, and they were like sorry dude it's a research network.

John Day: The CYCLADES network was never used in France the way the ARPANET was used in the U.S. because they were really more interested in exploring and understanding stuff. And that's how the datagram comes about because Louie says "what's the minimum thing we have to assume, and let's see what we have to add", right? And that's what led to the whole datagram and end to end transport protocols so –

Richard Bennett: That's very Cartesian – very French.

John Day: But that's how they got to that and I think the internet once they were used to that idea. Just the whole Idea that you are naming the host is still very much in their minds.

Richard Bennett: Yeah the stuff that they wrote about CYCLADES its really fascinating to read that because one of the things that he talks about is the desire to set up a routing system that would enable a host to break a message up into parts that could be sent in parallel over multiple paths through the network at the same time. Which was kind of the idea behind IP fragmentation and reassembly the kind of resequencing that TCP does but it doesn't actually happen that much on the internet because we do tend to stick to a path until it changes.

John Day: Well there's that then there's the fact –

Richard Bennett: But could you even imagine a telephone call – having parts of a telephone call being sent in parallel on multiple paths at the same time it just not conceivable.

John Day: But remember this was a network build to not be a voice network.

Tom Evslin: Those of us that were early in voice over IP lived with that of course all of the time. The network was built not to do voice almost if it was out of hostility –

John Day: Exactly.

Barry Shein: The internet was first built on top of the phone network, we had dial up, and then the phone network was reinvented on top of the internet. And now they are going through a cycle – but can I make a comment more on the policy. My problem here is it's funny that how they are using the Communications Act of 1934, which surprises me that the language is being resisted I was reading and it says "Well regulators created various telecoms" I read, point phrase. I said "no" it created one monopoly. When they granted that monopoly is what granted it the authority, besides the monopoly of the entire phone system with a few tiny exceptions. And it was by that power to ignore antitrust monopoly regulations that the FCC got that, and the limited bandwidth, and the limited spectrum arguments with radio.

Anyhow one of the things that happened with AT&T, one of the agreements in the monopoly was they couldn't go into applications. When they came out with Unix, the Unix operating system I remember I first got it in the 70s and they would tell us we are charging you 60 dollars for a tape copying fee and you are not allowed to – and we would pay 60 dollars for the mailer and the round tape. Well I think somewhere, therein lies a big part of the solutions to the Net Neutrality question, OK. I think what is needed is for the FCC and others to step in and split the wire plans so to speak from the application plans. Comcast wants to run video or a cable TV network fine, good luck to them, but they can't also run Tier 1 networking, OK? Because that creates a conflict of interest. Especially the way we have small –

John Day: You also have to allow for innovation in the network.

Barry Shein: By whom?

John Day: The network itself is the problem. There are fundamental flaws in what we've got. Two-thirds of the addressing is missing, the congestion management stuff is the worst you could have. This thing is not built to last. In fact, there's a report out already that the ROI on the internet has gone negative.

Richard Bennett: A lot of the sort of emotional drive for Net Neutrality is the idea that the internet is damn near perfect. Certainly, when the politicians talk about it the

arguments that they make, are they want to keep the internet the way it's been, they want to preserve the norms that have allowed the internet to flourish and all of this. And the assumption is it's really kind of telephone thinking that the network just is what it is and the only interesting stuff is what happens in applications. Keeping the internet the way it is foreclosures a lot of opportunities for new applications that are not really practical to implement over today's internet because of the quality of experience issues, the congestion issues, the financial issues.

John Day: The way things are set up you're almost required to never get beyond a client server model.

Barry Shein: I think the assumption of the politicians is like the phones, like the cable TV, like any of these systems is that if it breaks some bunch of propeller heads will go fix it "not our problem" and they deal with it on a level of policy. This is what they understand. Policy makers understand Game of Thrones they are moving little pieces on the chess board of who owns what and who has more power, who has more authority. They are not going to consider these issues of technological problems I am sure they screamed with the phone company they used to actually – you used to be able to get lines that were permanently flipped - you could actually – and put a current through it on both ends. They were physically connected between here and let's say Cambridge. Anyhow you could buy these for ten dollars a month. Alarm Tips, I don't know what they were called, something like that. Anyhow, the underlying technology is just not their interest.

Richard Bennett: So where to come down on the policy side Barry? One of the options is we accept the fact that neither Title I nor Title II really is an appropriate model for internet service? And ask for Congress to give the FCC a targeted authority to prevent internet service providers from misbehaving? As the LA Times so eloquently put it yesterday "monkeying with packets."

Barry Shein: Well I said what we will ultimately need do in the long haul is split the wire plans, the packet delivery service, all that layer from the application deliveries, OK. Just like they did with the AT&T phone system that you either run the wire plans or you run the applications.

Richard Bennett: Yeah but that was a monopoly scenario, right? I mean you could split that and say AT&T can't sell Unix because its got a telephone monopoly but nobody has a monopoly over internet service.

Barry Shein: No there's four of them – five of them.

Tom Evslin: Richard, I disagree with Barry even though the argument is tempting for a couple of reasons. One the wire plan isn't the only way that we deliver internet packets anymore, so you want innovation at all levels whether it's micro grids, below earth orbit satellites, balloons. So if you're saying now the non-literally wired is also a monopoly then you prevent innovation by little guys in that space. The other problem is the one you and John were getting back. There's a hypocrisy in this – not that you're being hypocritical – but a lot of the wired network belongs to Google and Amazon but is careful not in what the FCC regulates according to Title II so those companies have their own backbone, and they have data caches it's like the local warehouses, it's the cyber equivalent.

So if you're saying OK we're going to regulate the backbone if it belongs to Comcast or AT&T or Verizon, but what are we going to do about the private networks which certainly advantage the packets of Google and Amazon are those free to remain competitive? I think they should be, but then it is hard to regulate the others.

Barry Shein: I don't think that that is a problem. We have always had that. People used to run voice telephony over barbed wire. More specifically, Sprint. What does Sprint stand for: Southern Pacific Railroad Internal Network that's what Sprint is an acronym for. The railroads ran their own telephone network, and that's not a problem because the public is not affected.

Tom Evslin: But the public is affected by it if you're hosted on Amazon you get to ride on the Amazon network. The networks you were talking about like Sprint and the barbed wire – which I love – those are splinters but if you look at Google's worldwide network or Amazon's worldwide network those are hardly splinters. If you look at whose Cisco's biggest customers it's not traditional telcos, it's not Comcast it's the very internet companies that want to have some of the other internet backbone regulated. And so I would argue they already are competitive backbone providers therefore there's

competition. The regulation far from accomplishing Net Neutrality is actually putting its finger on the scales.

Richard Bennett: So what this is starting to look like is an incumbency protection racket.

Tom Evslin: That's what I am thinking. No one sits down and says I am an evil incumbent, I can do evil. It's just that as we get older we tend to get more conservative you get more dominant, you tend to be more protective and less open to innovation and thinking a little bit of control would be a good thing.

Richard Bennett: If you've got a worldwide network like Facebook and Google do and Amazon, you're in control of the quality at which you deliver your product to customers because you can decide where to connect to the ISP networks, how much work to do, and if you have that ever increasing power in your private network and are also regulating the hell out of the ISP networks so that the progress in their quality is retarded, slowed down, then you got an advantage over incumbents that might challenge your business model right?

Tom Evslin: If I want my business to take advantage of Amazon's network I can host stuff on Amazon, and by the way I think Amazon hosting is wonderful, I use it. But suppose that I didn't want to host on Amazon I wanted to buy better connectivity to various points, Title II regulations says that the carriers who could sell me that by the drink aren't allowed to.

Richard Bennett: Yeah you can buy it from Akamai, or Limelight, but you can't buy it from the people who could actually provide it for the least cost.

Tom Evslin: That's right.

Richard Bennett: That's to me the essence of the fast lane controversy. It is very definitely the case that the large ISPs would like to be in the content network delivery business. They are really well positioned to do it they already have the infrastructure in place, they've got the agreements with the other ISPs and they could offer that service at pretty substantially lower price than the current CDNs do but they are barred by law from doing it. But that actually sort of prevents them from the kind of abuses that you might expect in a vertically integrated monopoly but it also deprives the public of lower costs to

get to those services and effectively more start-ups. And the main thing is what the ISPs could do if they were allowed to be in the business of selling customized transmission services for a fee, we could get the sort of voice equivalent of a CDN.

Tom Evslin: Well also the whole innovation thing has become not in the network has just become a morass of patches. I am an unlicensed engineer but in general id like a layered architecture where one layer doesn't know anything about the other layers so a much cleaner environment to develop it. Except when you need radical innovation. My favorite example of that is Apple's development of the Macintosh, you couldn't've taken that step forward without owning every layer. Without owning the hardware layer, the software layer, developing the first applications. It's not a good model except when you have to make changes at every level to get to the next place -

John Day: Well actually what we found was by doing precisely what you just said making sure the layers completely independent, actually led to a much simpler implementation. It turns out if you do that right, you not only reduce the router table size by a factor of at least four or five but enables mobility with no new protocols at all. But the implementations are anywhere from 5,000 to 1,000X simpler.

I absolutely agree architecturally the layers ought to stay simple and I didn't say that well, I think there's a time in order for innovation to happen the same organization has to be able to innovate in all layers at the same time. That doesn't mean to break the layering but it's to say this breakthrough is possible. And we have to think about that in every layer, still keeping the layered approach.

Richard Bennett: The engineering problem in setting up a good layered architecture is that you don't just build firewalls between the layers you define communications that takes place between the layers. And I think you can have a highly functional and very efficient layered model if you can do the right kinds of communication across layers which if you look at IEEE 802 standards that's the first thing they do, they find all of the messages and all of the signals that cross a layer boundary such that you can actually layer them independently in a highly functionally system. When a company like Apple that does it all in house, they are doing everything from the applications to the operating

system through the hardware they are ensuring that the different elements of the system can communicate with each very well. All of the information is available.

John Day: One of the things that Apple did that I think has never really been very well recognized was that they had very strong APIs.

Tom Evslin: Yeah from a developer's point of view that was the whole secret you could at the power of the layers below you without knowing a damn thing about them and so you could do wonderful things.

Richard Bennett: And the way that you did that on the IBM PC was that you read the source code to the bios. You either manipulated the hardware directly or you went right into a BIOS level routine.

Barry Shein: Of course if someone has a design for an entirely new network it's completely straightforward to simulate it on the current internet.

Richard Bennett: Well John is that person. RINA is the new internet, a recursive internet architecture. They are simulating it on the existing internet, how about that. Well we should probably – I think we are about to start losing our audience here we have been on for a little over an hour so does anybody have any pithily remarks to close with that you want to make sure the world understands? 1989 was a great year Barry that was the last time the Oakland A's won the World Series. There was that little earthquake thing though. Thanks an awful lot, it's been a great discussion I would like to get you all together again sometime and do some more of this. The only thing missing was some beer. I will write this up for the High Tech Forum Podcast and I forgot to mention Tom's op-ed in Morning Consult about who has fast lanes and who doesn't I will put a link to that in the blog post.

Tom Evslin: If you put a link in to my blog not only will it help promote it there is a lot more content in there.

Richard Bennett: I will in fact. Great to see you guys and come check us out at HighTechForum.org. And if you haven't already made your voice heard at the FCC go ahead and do that too. You can join the 20 million robot comments.