

JOINT STATEMENT OF INTERNET ENGINEERS AND PIONEERS

The undersigned submit the following statement in support of the Open Internet Order. We seek to assist the Court’s review of the Order by supplying certain facts about the structure, history, and evolving nature of the Internet. As developers, engineers, and designers, we realize that without openness and neutrality the Internet as we know it will cease to exist, because it is that openness and neutrality that gives the Internet its flexibility, leads to its growth, and has made it a vital resource for all aspects of modern life. We believe the Order, if affirmed by this Court, will help preserve those characteristics. Further, in the absence of a clear but limited Open Internet Rule, service providers could *and would* engage in the practices of blocking, throttling, and interference that would upend the Internet, representing a fundamental departure from how the system has behaved for over forty years. We urge the Court to uphold the Order.

I. A Brief Introduction to the Internet

A. A Network of Networks

Fundamentally, the Internet is a collection of tens of thousands of individual networks of computers and other devices, almost all of which are owned, operated, and maintained by different entities.¹ In order to facilitate global communication, each of these independent networks interconnects to one or more of the other networks, thus leading to the term “Internet.” While each of these networks speaks the same language (or in technical parlance, protocol), and can thus be described using the same technical tools, the actual forms of the networks vary widely, both in terms of their architecture (i.e. their size and shape) as well as the underlying technology they use to connect devices. These differences depend in large part on the purpose each network serves.

For example, the type of network that is perhaps most familiar is a Local Area Network (LAN). LAN networks, such as the wired network in an office building or a Wi-Fi network in a home, connect a relatively small number of devices together. LAN networks connect to the Internet via yet another network, that of an Internet service provider, or ISP.

¹ CIDR REPORT, www.cidr-report.org/as2.0/ (last visited Sept. 14, 2015).

A typical ISP network connects anywhere from dozens to thousands of homes and businesses (or in the case of some wireless ISPs, mobile devices) to the rest of the Internet. This connection occurs in two parts. In the first part, the ISP must connect its customers (i.e. its retail subscribers) within a given geographic area to its own network facilities. This connection can be made over a variety of mediums: coaxial cables (originally used solely for cable TV transmission), copper wires (originally used solely for telephone communication), fiber optic cables, or in the case of wireless ISPs, radio waves. Regardless of the communications medium, it is typically the case that the connection is asymmetric: the ISP reserves more of the bandwidth of the connection for downloads – data traveling to the customer – than it does for uploads from the customer.

The second part of the connection involves connecting the ISP’s network to one or more of the other networks that make up the Internet. Typically, this second connection is made to either another ISP or an entity known as a “backbone provider.” Unlike a retail ISP, a backbone provider does not sell Internet access to individuals. Instead, backbone providers are “high capacity long-haul transmission facilities” which offer to connect different networks together in what are called “peering arrangements.”²

In peering arrangements, the two connecting parties formalize the role each will play in their interconnection: what levels of traffic will be allowed to and from each party, where the interconnection will be located physically, and who will pay for upgrades to the interconnection if they are necessary. Peering between large entities is often done in a settlement-free manner, meaning that no money is exchanged as part of the peering arrangement. This sort of settlement-free peering is sometimes dependent on the two networks exchanging similar levels of traffic (i.e. each network sending as much traffic to the other as it receives).³ However, an equal traffic exchange requirement frequently does not make much sense when backbone providers or edge providers connect to ISP networks, due to the inherent asymmetric nature of ISP traffic. In other words, because most ISP customers download more than they upload, any peering arrangement

² *Verizon Communications and MCI, Inc. Applications for Approval of Transfer of Control*, 20 FCC Rcd 18433, 18493 (2005).

³ *See, e.g., Time Warner Cable's IPV4 and IPV6 Settlement-Free Peering Policy*, TIME WARNER CABLE, http://help.twcable.com/twc_settlement_free_peering_policy.html (last visited Sept. 14, 2015).

between a backbone or edge provider and a retail ISP's network will result in more traffic being sent from the backbone or edge provider to the ISP than vice versa.

Finally, it should be noted that the same company will often act in different roles: a large ISP can provide backbone service to other, smaller ISPs, or a large edge provider may own similar infrastructure to a backbone provider. Thus, it is important when discussing the roles of the major players on the Internet to focus on the specific context in which they are being discussed; to do otherwise can lead to confusion and mismatched assumptions.⁴

B. Packet-Switching and Congestion

While the above gives an accurate picture of how the Internet is laid out, it does not explain how the different networks actually succeed in communicating with one another. In this section we explain how this is done, so that we can later explain the technical ramifications of the FCC's Order.

Two major technical principles underlie how the Internet functions. The first is the concept of packet switching. In a packet switched network, the data to be transmitted (be it a webpage, images, sound files, or a video) is broken down into chunks known as packets, each of which is sent off individually to its destination.⁵ An Internet packet contains several important pieces of information: the numerical address of the device which sent the packet, known as an Internet Protocol address (or IP address); the IP address of the intended recipient; (sometimes) the type of data the packet contains; and of course the actual data.⁶ In this way, a packet is similar to a postcard—anyone who is part of the delivery chain can read who it is intended for, who sent it, and what it says. (Note that this does not hold true if the contents of the packet are encrypted—then the packet is more like a sealed, opaque envelope.)

⁴ For example, an ISP may have different customers depending on its role: as a retail ISP, its customers are the retail customers who subscribe to its service for Internet access, but if it also provides transit services as a backbone provider, then in that role its customers would be other ISPs.

⁵ JONATHAN E. NEUCHTERLEIN & PHILIP J. WEISER, *DIGITAL CROSSROADS: AMERICAN TELECOMMUNICATIONS POLICY IN THE INTERNET AGE* 42-43 (2005).

⁶ INFO. SCI. INST., UNIV. S. CAL., *DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION (1981)*, *available at* <https://tools.ietf.org/html/rfc791>.

When it comes time for a computer to transmit a packet, the computer sends it to the next “hop” in the delivery chain, typically a device known as a “router.” A router is a device that bridges the connection between multiple communications links, whose sole job is to send packets one step closer to their destination. It does this via a “routing table,” which lists all the communication links the device is attached to, and the range of IP addresses that can be found on each of those links. Thus when a packet arrives, the router compares its destination address to the routing table and then sends it off on the appropriate link.

Of course, sometimes packets arrive at a router faster than the router can process them or faster than the communications link can transmit them, leading to congestion. Internet congestion is analogous to the traffic congestion that might occur when a busy four-lane interstate splits into two smaller highways: even though there is theoretically enough capacity, if all of the cars coming from the interstate want to travel along only one of the smaller highways, a backup will ensue. Similarly, if a router receives packets faster than it can transmit them along their desired links, the packets will be stored in a buffer until they can be sent. Unlike traffic congestion, however, if too many packets fill up the buffer, any new packets will simply be “dropped”, or ignored. Thus the Internet is a “best-effort” service: devices make their best effort to deliver packets, but do not guarantee that they will succeed.⁷

C. The Principles of Neutrality and Openness Are Inherent in the Internet’s Design

The Internet is more than just a way for computers across the globe to exchange packets of data; it is a platform on which people have developed a variety of amazing new technologies, from web browsing to email to social networking to online courses. The Internet’s tremendous growth and popularity as a platform have been due to two design principles, both of which ensure that the Internet is an open, neutral platform.

The first of these design principles is the idea of the layered network communications stack (often referred to as simply “the network stack”). Essentially, the network stack is a way of abstracting the design of a computer network into different layers, where each layer is

⁷ In fact dropping packets is one of the key signals routers use to communicate to devices that they are sending packets too quickly, so that the devices can reduce their transmission rate. Thus, a packet being dropped is not only a frequent event—it is critical to the proper functioning of the Internet.

responsible for certain functions, but can implement those functions in any way that meets the specifications. For example, the “physical layer” is responsible for physically transmitting and receiving bits. It can do so over fiber optic cable, copper telephone lines, radio signals, etc., as long as it provides a way for the layer above it to access the “transmit and receive bits” function. Above the physical layer is the “network layer,” which is responsible for ensuring each device on the network has a unique address, and for sending and receiving packets of data to specific addresses. It is at this layer that the famous Internet Protocol actually resides, which provides a “send data to a certain address” function to the layer above. Similarly, further up is the “transport layer,” which is responsible for ensuring that data gets to its destination reliably and intact. While most Internet traffic use a protocol known as the Transmission Control Protocol (TCP) to accomplish this, theoretically any protocol could be used as long as it provided the next layer up a “reliably send data” function.⁸

The key takeaway from the idea of the network stack is that each layer can tell the layer below it to carry any type of data. This is why the same Internet Protocol can support such varied applications as email and real-time video-conferencing. If someone wants to develop a new Internet application or protocol, all they have to do is insert their new technology at the appropriate layer; the layers below will perform their functions regardless of the type of data the developer tasks them to handle. This openness allows developers to build new and different types of applications without having to worry about the technical details of the layers below. “Consider, for instance, how these design principles collectively facilitated the rise of the World Wide Web application. Because the network is general, its founder Tim Berners-Lee could introduce it without requiring any changes to—or permission from—the underlying physical network.”⁹ All he had to do was define the protocol, and the underlying layers transported the data as desired.

The second design principle is the “end-to-end principle.” In order for the network stack to function properly, all of the “smarts” must be located in the edge devices—the computers, servers, and smartphones which connect to the Internet—and not the routers and

⁸ JAMES F. KUROSE & KEITH W. ROSS, *COMPUTER NETWORKING: A TOP-DOWN APPROACH* (4th ed. 2007).

⁹ Brief Amicus Curiae of Internet Engineers and Technologists Urging That The FCC’s Order Be Affirmed, *Verizon v. Federal Commc’ns Comm’n*, 740 F.3d 623 (D.C. Cir. 2014) (No. 11-1355).

communications channels that make up the interior of the network.¹⁰ This enables the network to be general purpose: the nodes that make up the interior of the network are as general as possible (i.e. do little more than move packets along to their final destination), and application-specific features only reside in the devices that connect to the network at its edge.¹¹

It is easy to see how the end-to-end principle applies in the case of the Internet. The interior of the network, made up of the communications links (i.e. the physical cables) and the routers that connect them, have very little “intelligence.” In fact, the Internet Protocol, which is the protocol routers use to communicate, does not even have a way for one device to make sure the next device received a packet. All a router is expected to do is read the packet it received, figure out where it should go, and send it off. The actual “intelligence” comes entirely from the computers and servers and smartphones that connect at the “edge” of the Internet. This is how the Internet can support protocols that require guaranteed delivery of data (such as file transfer protocols), as well as protocols where guaranteeing delivery is less important than ensuring that the packets that are received have low latency (such as protocols for voice or video chat).

II. How the Internet Has Changed Since 2010

While technologies like the Internet Protocol and TCP have changed little since the early nineties, part of the Internet’s resilience and value comes from the myriad ways in which those underlying protocols can be used. It should come as no surprise, then, that the Internet as a whole is not a static, monolithic creation, but a constantly evolving system. In this section, we describe the major ways the Internet as a whole, and consumer ISPs in particular, have changed since 2010.

A. New Internet Protocols and Services Continue to be Invented

¹⁰ J.H. Saltzer, D.P. Reed & D.D. Clark, *End-to-End Arguments in System Design*, 2 ACM TRANSACTIONS ON COMPUTER SYS. 277 (1984).

¹¹ We note that many network operators and equipment vendors contest the fundamental nature of the “end-to-end” principle. However, their arguments are usually made in order to claim that they (or their equipment) can “add value” to the network by adding “smarts” to the network itself—usually as a way to try to reverse the commoditization of network hardware and services. Further, as we explain in section III.C, this insertion of “smarts” into the interior of the network frequently causes problems for developers of innovative new protocols and applications designed to run on a neutral Internet.

Although it may seem obvious, it is worth noting that new services and applications that rely on the Internet are constantly being developed. This sort of innovation is not limited just to new websites, web-based services, or even mobile applications which use the Internet for communication. It goes deeper, down the network stack to new protocols and new ways of using the network. For example, consider Bitcoin, “an open-source, peer-to-peer digital currency.”¹² Bitcoin, which was first released in 2009, relies on the decentralized nature of the Internet to ensure that all individuals using Bitcoin see the same transactions (i.e. to prevent someone from spending the same Bitcoin twice) as well as to ensure that no one individual has full control of the ledger that records transactions (i.e. to prevent someone from simply assigning themselves all of the Bitcoins in the system).¹³

Similarly, the “InterPlanetary File System (IPFS) is a peer-to-peer distributed file system that seeks to connect all computing devices with the same system of files,”¹⁴ first developed just last year.¹⁵ The goal of IPFS is to create a more permanent, more distributed version of the World Wide Web, one in which the entirety of files available on the Web are distributed to millions of computers across the globe. If successful, IPFS would make censorship of individual webpages or websites technically impossible, while also ensuring that a permanent record of all the files ever posted on the Web is always available, for archival and historical purposes. Like Bitcoin, IPFS relies on the underlying decentralized, open infrastructure of the Internet, distributing data using peer-to-peer protocols that are fundamentally different from the sorts of protocols used to transmit webpages, emails, or streaming videos.

The key takeaway from both of these examples is that innovation surrounding the Internet is ongoing—but more importantly, this sort of innovation relies on the open, neutral nature of the Internet. As we explain further in section III.C, if ISPs interfered with their customers’ traffic based on the protocol or service in use, such innovation would become impossible.

¹² JERRY BRITO & ANDREA CASTILLO, MERCATUS CTR., GEORGE MASON UNIV., BITCOIN: A PRIMER FOR POLICYMAKERS (2013).

¹³ Joshua Davis, *The Crypto-Currency*, THE NEW YORKER, Oct. 10, 2011, available at <http://www.newyorker.com/magazine/2011/10/10/the-crypto-currency>.

¹⁴ THE IPFS PROJECT, <https://ipfs.io/> (last visited Sept. 14, 2015).

¹⁵ *History for IPFS*, GITHUB, <https://github.com/ipfs/ipfs/commits/master/README.md> (last visited Sept. 14, 2015).

B. ISP Caching is Becoming Less Useful

In the early days of the Internet, many ISPs set up caching servers that would sit between their customers and the rest of the Internet. These servers would record what data customers were requesting from the World Wide Web, and store copies in a local cache that the server could send when other customers made the same request. For example, if many customers were reading the same newspaper article about net neutrality, the ISP would store a copy of that article on the caching server. Then, when a new request for the article came in, the ISP would send back the copy instead of waiting for the request to go all the way to the newspaper's server and back via the Internet. This way the ISP could reduce the amount of time it took for a customer to download the article (since the ISP's caching server would be closer to the customer than the newspaper's server), and ISPs could save on bandwidth (since they would not have to re-download the article from the newspaper's server every time a new request came in).¹⁶

However, recent changes have decreased the need for ISP caching services. This is due to the widespread use of Content Delivery Networks, or CDNs. CDNs are very similar to the caching servers described above, except they are operated by edge providers themselves (or by third-party companies who sell their CDN service to edge providers). CDNs consist of Internet-connected caching servers strategically placed in different geographic regions, often very close to or on the edge of the network of one or more ISPs. Edge providers upload their content to these caching servers, so that they can have fine-grained control of what gets cached and how long it stays cached—control they do not have over ISP-controlled caches.

In addition to becoming unnecessary, ISP caching is also becoming less feasible due to the increasing proportion of Internet traffic that is encrypted. (In 2010 less than 2% of traffic on the Internet was encrypted¹⁷, but by 2016 that number is projected to reach over 64%.¹⁸) Encryption prevents ISP caching from being effective because when a user requests a webpage

¹⁶ KUROSE & ROSS, *supra* note 6.

¹⁷ SANDVINE, GLOBAL INTERNET PHENOMENA REPORT (2011), *available at* <https://www.sandvine.com/downloads/general/global-internet-phenomena/2011/1h-2011-global-internet-phenomena-report.pdf>

¹⁸ SANDVINE, GLOBAL INTERNET PHENOMENA SPOTLIGHT: ENCRYPTED INTERNET TRAFFIC (2015), *available at* <https://www.sandvine.com/downloads/general/global-internet-phenomena/2015/encrypted-internet-traffic.pdf>.

or file over an encrypted connection, the ISP cannot see the name or location of the file the user is requesting, or the contents of the file itself. As a result, the ISP has no way of knowing what files are popular enough to cache, nor any way of knowing when a user requests a popular file. Given the inevitability of ubiquitous encryption, ISP caching is destined to become an obsolete practice.¹⁹

C. DNS and Email Are No Longer the Province Solely of ISPs

Another major change has been the dramatic surge in popularity of third-party web-based email providers. Within the US, Google.com, Yahoo.com, and Live.com, the websites of the top three webmail providers, are ranked first, fifth, and twelfth in terms of their popularity.²⁰ For comparison, the websites of the top three US ISPs, Comcast.com, Att.com, and Twc.com are ranked at 1541st, 66th, and 1551st. While not everyone who visits each of these websites is using them strictly for email, the dramatic difference in popularity illustrates the decreasing relevance ISP customers put on the information services provided by their ISPs.

Similarly, fewer people are making use of their ISP's Domain Name Services (DNS).²¹ This is because of the proliferation of free, open DNS servers online. Google Public DNS, for example, is a DNS service Google offers to any Internet user, free of charge. It handles over 400 billion DNS requests per day,²² many of which come from home users who, for reasons of speed or security, choose to use Google's DNS instead of their own ISP's service.²³

D. Customers Depend on ISPs for Internet Access, Not Information Services

¹⁹ Indeed, all major browsers have announced that they will only support the next version of the famous HTTP protocol, HTTP/2, over encrypted connections. Dan Goodin, ARS TECHNICA, *New Firefox Version Says "Might as Well" to Encrypting All Web Traffic*, April 1, 2015, <http://arstechnica.com/security/2015/04/new-firefox-version-says-might-as-well-to-encrypting-all-web-traffic/>

²⁰ *Top Sites in United States*, ALEXA, <http://www.alexa.com/topsites/countries/US> (last visited Sept. 14, 2015).

²¹ DNS is the service computers rely on to look up the numerical IP address associated with a given domain name (e.g. www.eff.org).

²² Yunhong Gu, *Google Public DNS and Location-Sensitive DNS Responses*, GOOGLE WEBMASTER CENTRAL BLOG, Dec. 15, 2014, <http://googlewebmastercentral.blogspot.com/2014/12/google-public-dns-and-location.html>.

²³ *Introduction to Google Public DNS*, GOOGLE, <https://developers.google.com/speed/public-dns/docs/intro> (last visited Sept 14, 2015).

In the early days of Internet access, customers frequently chose which ISP to subscribe to based on the content and information services that ISP supplied in addition to general Internet access. ISPs like AOL, Compuserve, or Prodigy differentiated themselves based on the different information services each provided—services like chat rooms, bulletin board systems, email, and specialized content only available to an ISP’s own subscribers.²⁴

Now, however, ISPs compete primarily on the reliability and bandwidth of their Internet connections,²⁵ and customers subscribe to an ISP’s service not because of the added information services an ISP might provide, but because the subscription enables customers to transmit and receive data to and from the wider Internet. In other words, the information services ISPs provide are simply no longer connected in any meaningful way to the data routing and transmission service they offer. The two are easily separated, as evidenced by the fact that a consumer can instead choose to subscribe to any given information service from an entity other than their ISP. In fact, saying that ISPs provide an information service to their customers because they offer caching and webmail in addition to Internet connectivity is like saying that airlines are in the business of providing an entertainment service because they offer in-flight movies in addition to transportation. While these additional services might be selling points, they are not integral to the fundamental offering ISPs and airlines make: to transport things (either data or people) at the customer’s request.

III. Technical Interpretation of the FCC’s Order

In light of the foregoing, we can better anticipate the technical consequences of the Order and the risks of losing the “rules of the road” it establishes. We focus on the parts of the Order that will have the greatest technical effect: the rule preventing broadband ISPs from slowing down traffic (or blocking it altogether) based on what type of data the traffic contains (or what

²⁴ Michael Wolff (1997). Netstudy. Dell Publishing.

²⁵ See, e.g., *Sprint 4g Commercial*, YOUTUBE, <https://www.youtube.com/watch?v=NPdkvg9Kw-M> (last visited Sept. 14, 2015) (touting the bandwidth of Sprint’s 4G wireless network); *Comcast- Fast Rabbit*, YOUTUBE, https://www.youtube.com/watch?v=h16qMJ_LCyg (last accessed Sept 14, 2015) (compares Comcast’s high-speed Internet access with “a rabbit/panther with turbines backed by an unusually strong tailwind on ice...driven by an over-caffeinated fighter pilot with a lead foot all traveling down a ski jump in Switzerland under better than ideal conditions.”).

type of service that data is for); and the rule preventing broadband ISPs from prioritizing certain types of traffic in exchange for consideration (monetary or otherwise).

A. Technical Effects the Order Will *Not* Have

First, we wish to dispel the rumor that the FCC's Order will eviscerate ISPs' ability to manage their networks, resulting in massive congestion, unchecked proliferation of spam and viruses, and slow speeds for all.²⁶ This is simply not the case. First, the Order contains an exception for reasonable network management. Thus, as the FCC has explained, it does not affect ISPs' ability to filter unwanted spam, computer viruses, and other malicious content out of their customers' unencrypted traffic, if the customer requests this sort of protection.²⁷ Similarly, it does not bar ISPs from defending their networks against attacks.

Second, the Order does not affect techniques for dealing with network congestion that do not discriminate based on service or application.²⁸ Such techniques include "weighted fair queuing," in which each flow of traffic (for example, all of the traffic coming to or from each customer) is assigned a proportion of the outgoing bandwidth along the congested portion of the network. More advanced algorithms for handling congestion, such as Random Early Detection²⁹, CoDel (Controlled Delay)³⁰, and Stochastic Fair Blue³¹, are also not impacted by the prohibition on throttling. Simply put, the Order will not dramatically change how most ISPs manage their networks. ISPs will still be able to ensure each customer gets a fair allocation of the ISP's total

²⁶ See, e.g., *The Truth About Net Neutrality*, CENTER FOR INDIVIDUAL FREEDOM, <http://www.stopnetregulation.org/wp-content/uploads/2011/08/Net-Neutrality-talking-points.doc.pdf> (last accessed Sept. 14, 2015) ("Under 'Net Neutrality' regulations, every decision to block pornography, spam, or security threats will have to be approved by the government.").

²⁷ FED. COMM'NS COMM'N, REPORT AND ORDER ON REMAND, DECLARATORY RULING, AND ORDER para. 221 (March, 12, 2015), available at http://transition.fcc.gov/Daily_Releases/Daily_Business/2015/db0312/FCC-15-24A1.pdf.

²⁸ Monica Allevan, *Nokia Networks: Necessary Network Management Still Possible Under Proposed Net Neutrality Rules*, FIERCEWIRELESSTECH, Feb. 9, 2015, <http://www.fiercewireless.com/tech/story/nokia-networks-necessary-network-management-still-possible-under-proposed-n/2015-02-09>.

²⁹ Sally Floyd & Van Jacobson, *Random Early Detection Gateways for Congestion Avoidance*, IEEE/ACM TRANSACTIONS ON NETWORKING (1993), available at <http://www.icir.org/floyd/papers/early.twocolumn.pdf>.

³⁰ Kathleen Nichols & Van Jacobson, *Controlling Que Delay*, ACMQUE, May 6, 2012, <http://queue.acm.org/detail.cfm?id=2209336>.

³¹ Wu-change Feng et al., *Stochastic Fair Blue: a Queue Management Algorithm for Enforcing Fairness*, 3 TWENTIETH ANNUAL JOINT CONFERENCE OF THE IEEE COMPUTER AND COMMUNICATIONS SOCIETIES 1520, available at http://www.thefengs.com/wuchang/blue/41_2.PDF.

bandwidth. And of course, ISPs can still sell customers different levels of service, and manage their network so that higher-paying retail customers get more overall bandwidth. The only thing the Order forbids is ISPs blocking or throttling their customers' traffic based on the content, applications, or protocols their customers choose to use.

B. Scope of the FCC's Order

Some opponents of the Order suggest that it allows the FCC to regulate the entire Internet.³² This is not the case. The Order is limited in scope, targeting only retail broadband ISPs. With that said, we do not intend to minimize the effects should the Order be struck down; data destined for retail customers make up a huge percentage of U.S. Internet traffic.

Instead, we wish to highlight that unlike large businesses or data centers, which typically have multiple connections to different ISPs in order to achieve redundancy, most retail customers have only one Internet connection. As a result, retail ISPs enjoy what is known as “gatekeeper authority”—they are the sole gatekeepers of what customers can do online, since customers have no way to bypass any blocking or filtering their ISP puts in place (except changing ISPs, which is a time-consuming process that is often not even feasible). Essentially, retail ISPs are the weak link in the chain that connects users to Internet content and services. As with a physical chain, it is the weakest link that limits what tasks a chain can perform. In the next section, we explain how that weak link could break if the Open Internet rule is struck down.

C. Risks In the Absence of Open Internet Rule

In the absence of a clear and limited Open Internet rule, ISPs will be free to block, throttle, or speed up data based on its content or what service or application generated it. ISPs could degrade (or altogether block) certain protocols or websites. A frequently given example is that of an ISP degrading traffic containing streaming movies from some or all edge providers, in order to encourage its customers to instead use its own media-streaming service. But this sort of blocking and throttling would only be the tip of the iceberg. ISPs could go further, degrading traffic for any service they do not recognize or have not previously approved of.

³² See, e.g., *supra* note 27 at 321 (Commissioner Pai's dissenting statement on the order) (“[The Order] gives the FCC the power to micromanage virtually every aspect of how the Internet works.”).

That, in turn, would mean an end to the layered model upon which the Internet was built. Developers would have to ensure that their new application or protocol would work under different specifications on each of the thousands of networks that make up the Internet. Some networks might decide to handle data differently depending on whether it represented webpages or video. Others might decide that certain data needed to be prioritized.³³ Such a haphazard mishmash of different specifications and functionality would have made the growth of the Internet as we know it utterly impossible. Instead, it would have resulted in a balkanized Internet—one in which each ISP was its own private fiefdom, where edge providers had to negotiate with the gatekeeper in order to get access to the end users.

But blocking and throttling are not the only dangers. ISPs could decide to abandon the end-to-end principle, inserting nodes in their network that tried to “enhance” their customers’ experience by augmenting or transforming some content. This might seem like a reasonable design, since conceivably an ISP might have access to information that edge providers would not. (For example, an ISP might be able to provide more relevant search results or other information since it has a complete record of its customers’ browsing histories.) But this sort of interference could not only introduce bugs into services and webpages that weren’t expecting it, it could make it impossible for some applications (including applications yet to be dreamed of) to work correctly. Worse yet, it could also introduce security vulnerabilities which a malicious actor could use to harm the ISP’s customers.

IV. Conclusion

As computer scientists, networking engineers, and professionals who deal with Internet technology on a daily basis, we realize that without openness and neutrality the Internet as we know it will cease to exist, because it is that openness and neutrality that give the Internet its flexibility, lead to its growth, and have made it a vital resource for all aspects of modern life.

³³ It is worth noting that the Internet Protocol does specify a field in the header of IP packets known as the “differentiated service” field, meant to indicate some sort of priority. However, in the over thirty years since the widespread adoption of IP, no consensus has been reached about how edge devices should populate that field for use on the public Internet. As a result, traffic prioritization on the public Internet is almost nonexistent. The closest the Internet engineering community has come to a standard on prioritization is RFC 2474, which is a *proposed* standard last updated in 1998, and which is not in force. IETF NETWORK WORKING GROUP, DEFINITION OF THE DIFFERENTIATED SERVICES FIELD (DS FIELD) IN THE IPV4 AND IPV6 HEADERS (1998), *available at* <https://tools.ietf.org/html/rfc2474>.

We also realize that the threat to the Internet’s openness and neutrality is real. None of the scenarios described in the previous section is hypothetical. Comcast has interfered with legitimate traffic based solely on its type.³⁴ Both Comcast and Verizon have also admitted to modifying their customers’ traffic without their consent—Comcast by inserting ads into the webpages its customers view,³⁵ and Verizon by inserting unique tracking ID numbers into the data its customers send.³⁶ Port blocking and interference by ISPs in general has forced developers of new protocols and services to “camouflage” their new protocols as existing ones, in order to avoid discriminatory treatment. In fact, this sort of interference has become so bad that network engineers have developed a name for it: the “ossification” of the network stack.³⁷ As a result of this interference, development of innovative new protocols and services is already being hindered.³⁸

If this sort of blocking, throttling, and interference becomes more widespread, it would transform the Internet from a permission-less environment (in which anyone can develop a new app or protocol and deploy it confident that the Internet treats all traffic equally) into one in which developers would first need to seek approval from ISPs before deploying their latest groundbreaking technology. Developers and engineers would no longer be able to depend on the core assumption that the Internet would treat all data equally. The sort of rapid innovation the Internet has fueled for the past two decades would come to a sudden and disastrous halt.

Fortunately, there is a way to prevent this worst-case scenario from occurring: uphold the FCC’s Open Internet Order.

³⁴ Peter Eckersley et al., *Packet Forgery By ISPs: A Report on the Comcast Affair*, ELECTRONIC FRONTIER FOUNDATION, Nov. 28, 2007, <https://www.eff.org/wp/packet-forgery-isps-report-comcast-affair>.

³⁵ David Kravets, *Comcast Wi-Fi Serving Self-Promotional Ads Via JavaScript Injection*, ARS TECHNICA, Sept. 8, 2014, <http://arstechnica.com/tech-policy/2014/09/why-comcasts-javascript-ad-injections-threaten-security-net-neutrality/>.

³⁶ Jacob Hoffman-Andrews, *Verizon Injecting Perma-Cookies to Track Mobile Customers, Bypassing Privacy Controls*, ELECTRONIC FRONTIER FOUNDATION, Nov. 3, 2014, <https://www.eff.org/deeplinks/2014/11/verizon-x-uidh>.

³⁷ See, e.g., TRAMMELL & KUEHLEWIND, IAB WORKSHOP ON STACK EVOLUTION IN A MIDDLEBOX INTERNET (SEMI) REPORT (2015), available at <https://tools.ietf.org/html/draft-iab-semi-report-01>.

³⁸ Michio Honda et al., *Is it Still Possible to Extend TCP?*, ACM INTERNET MEASUREMENT CONFERENCE 181 (2011), available at <http://nrg.cs.ucl.ac.uk/mjh/tmp/mboxes.pdf>.

That is why we, the undersigned computer scientists, network engineers, and Internet professionals, based on our technical analysis and an understanding of both how the Internet was designed, how it currently functions, and what sort of technical changes ISPs are already making and wish to make in the future, respectfully encourage the Court to uphold the FCC's Open Internet Order.³⁹

Respectfully submitted,

³⁹ Unless otherwise noted, all of the signatories to this letter have signed in their personal capacity, and not as representatives of their employers or any affiliated organizations.