U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON ENERGY AND COMMERCE

April 13, 2018

TO: Members, Subcommittee on Communications and Technology

FROM: Committee Majority Staff

RE: Hearing entitled "From Core to Edge: Perspective on Internet Prioritization"

I. INTRODUCTION

The Subcommittee on Communications and Technology will hold a hearing on Tuesday, April 17, 2018, at 10:15a.m. in 2322 Rayburn House Office Building. The hearing is entitled "From Core to Edge: Perspective on Internet Prioritization."

II. WITNESSES

- Richard Bennett, Founder, High Tech Forum;
- Peter Rysavy, President, Rysavy Research, LLC;
- Paul Schroeder, Director, Public Policy and Strategic Alliances, Aira Tech Corporation; and
- Matt Wood, Policy Director, Free Press.

III. BACKGROUND

Today's Internet is unrecognizable from the Internet that existed 20 years ago. Today, there are billions of devices that send over 100 exabytes of data per month (one exabyte is equal to one million terabytes), and that number is expected to continue to rise.[1] As networks adapt to a changing multimedia landscape and more data-intensive applications, the prioritization and management of these networks takes on more importance. Consumers who pay for internet access expect it to work, regardless of what content, application, or service they requested. With the advent of virtual reality, augmented reality, video, and other data-intensive and latency-sensitive services, consumers have demanded more bandwidth, which has led to congestion and increased scrutiny of how Internet Service Providers (ISPs) manage that congestion on their networks.

[1] "The Zettabyte Era: Trends and Analysis, CISCO, June 2017. Available at: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.pdf

To establish and maintain a reliable and sufficient connection for consumers, network managers prioritize certain traffic over the Internet. Since the Internet's earliest days, data has been prioritized by network operators to ensure all users are able to access content they request.[2] By nature of the way the Internet was established and has grown over the last two decades, prioritization of certain data is necessary to maintain proper functioning of the Internet.[3] In light of concerns raised over the perceived anticompetitive nature of certain forms of prioritization, the Committee seeks a better understanding of how network operators manage data flows over the Internet and how data is prioritized from the network core to the edge.

A. Data Network Fundamentals

The Internet, from its inception, grew out of a layered network model. It is helpful to think of the Internet as a network of networks. From the physical cables or airwaves to the information being carried over them, various levels on the network ensure proper transport from one point on the Internet to another. As different businesses, countries, and regions designed new protocols over their individual networks, it became necessary to create a set of rules for transmitting and receiving information, or "packets," so that all networks could work together and route traffic around the globe.

The Open System Interconnection (OSI) model identifies seven common elements of common protocol that guide network operators: 1) physical; 2) data link; 3) network; 4) transport; 5) session; 6) presentation; and 7) application.[4]

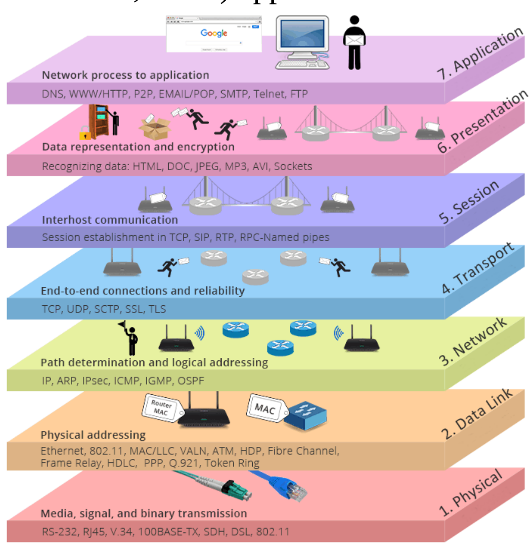

Source: FS.com[5]

---

[2] *See*, Andrew Froehlich, "The Basics of QoS." *Network Computing*. Available at: https://www.networkcomputing.com/networking/basics-qos/402199215

[3] *See*, Richard Bennett, "Designed for Change: End-to-End Arguments, Internet Innovation, and the Net Neutrality Debate," Information Technology and Innovation Foundation. Available at: http://www.itif.org/files/2009-designed-for-change.pdf?_ga=2.153893624.1117224997.1523645423-405235368.1523645423

[4] For a primer on OSI model, *see*: https://www.networkworld.com/article/3239677/lan-wan/the-osi-model-explained-how-to-understand-and-remember-the-7-layer-network-model.html

[5] Available at: https://community.fs.com/blog/tcpip-vs-osi-whats-the-difference-between-the-two-models.html

The OSI model is now used more as a guideline, and illustrates the different components that enable the Internet to function. Today, the most common protocol is Transfer Control Protocol/Internet Protocol (TCP/IP), which is widely established as the global standard.

TCP/IP has four elements that fulfill the same functions set out in the OSI model. The first, the *application layer*, includes the application, presentation, and session functions. Common examples of this element are the Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP). The application layer is primarily responsible for working with end applications at the *edge* of the network to set up and maintain a *session*, or the connection; authenticate credentials; ensure the compatibility of data formats; and compress and encrypt data. In TCP/IP, the application layer encompasses the top layers, 7, 6, and 5, of the OSI model.

The second element, the transport layer, prepares data for transport once the session is established. In TCP/IP, the transport layer is the *TCP layer*, and corresponds to layer 4 in the OSI model. Here, Quality of Service (QoS) is assigned to data as it is prepared for transmission over the network. Some services, like video or virtual reality, require different prioritization over the network to make the service work as intended. Text, on the other hand, is not as latency-sensitive and therefore is prioritized differently. The transport control protocol does just what the name suggests—controls the protocol for how data is transmitted through the network to meet latency or bandwidth requirements. The transport layer is primarily responsible for reliability, congestion avoidance, QoS management, and bandwidth efficiency.

The third element is the network layer, which establishes packets and identifies the routing sequence between network nodes. In the TCP/IP system, the *IP layer* corresponds to layer 3 in the OSI model. The network layer establishes the addressing component of data transmission, finding the path over which data will travel to reach the end user. The IP layer also creates the packets, the header of which contains information relating to the packets' path and prioritization over the network. The *data link layer*, which corresponds to the layer 2 of the OSI model, transfers data between network nodes in a Local Area Network or Wide Area Network. The data link layer aggregates packets and arranges them into frames. Class of Service (CoS) is a subset of QoS managed at the data-link layer. CoS is often used to differentiate and classify packets for assignment into queues of varying priorities. These queues, or classifications, often determine which QoS characteristics apply at the third layer, such as allocating certain bandwidth requirements to certain classes of data. Whereas the network layer is hierarchical—allowing packets to bounce their way from node to node throughout the Internet—the link layer is flat.[6] This makes the link layer ideal for data transfer across the *physical layer*, which is the final level and the first and lowest level in the OSI model.

From the networking perspective, prioritization is fundamental to the way data travels across the web. If every bit were treated equal, some applications simply would not function. Applications which require low bit error rates or packet loss require prioritization to reach the

[6] Peter Dordal, "An Introduction to Computer Networks." Department of Computer Science, Loyola University Chicago (February 19, 2018). Available at: http://intronetworks.cs.luc.edu/current/ComputerNetworks.pdf, at 13.

required or desired CoS, and the broader QoS.[7] Some applications, however, are not as sensitive, and thus can be prioritized on a *best effort* basis. Best effort networks do not support a specified quality of service, but rather route traffic as best they can, as the name suggests. Use of these tools allow network operators to manage congestion over their networks to ensure adequate service.

B. Network Components

It is also important to recognize that there are several network components over which data travels from where it is hosted to the end user. Each component is not necessarily owned, operated, or managed by the same company. The devices, software, or Wi-Fi router, for example, are distinctly separated from the ISP, a backbone transit service provider, or content delivery network (CDN). Each of these components can load, manage, and relay traffic in different ways. Transit providers and CDNs, for example, can route traffic geographically through several interconnection (or peering) points, which can balance the load among several data centers.[8] Wi-Fi routers can also manage traffic using firewalls that exclude certain content from reaching the end user. There are several of these types of bottlenecks on the network, but are not the only avenues to control content or prioritize traffic.

The Internet operates in a world where the sender wields large influence. In most cases, content and application providers are at the edge of the network and the network will find any available path to send data from one user to another. Senders are limited only by what the network can carry, and the Internet was designed so the network will accept traffic if requested and if technically feasible.[9] As the early data lines gave way to infrastructure that could handle higher bitrates and larger bandwidths, edge providers were able to create more data-demanding applications. In turn, network operators upgraded their networks, privately arranging interconnection agreements, or caching popularly requested data physically closer to end users. As the landscape of Internet content and services continues to change, network management tools and prioritization will continue to be important.

C. Network Management and Optimization

Because ISPs, transit service providers, and other network operators live in a content-driven world, they must find adaptive ways to manage traffic in a way that minimizes congestion, optimizes quality of service, and meets consumer demand. One of the ways traffic is managed is at the transport layer, which applies prioritization such as QoS, best effort, or others, depending on the network manager. Contrary to popular belief, traffic is not prioritized by sending some traffic faster than others. All traffic travels at the same speed—the speed of light. Rather, prioritization puts certain bits over others, resulting in some bits being dropped. This can either be done randomly—where certain latency-sensitive services will suffer—or it can be done

---

[7] Id.
[8] *See*, “End-User Mapping: Next Generation Request Routing for Content Delivery.” Available at: https://www.akamai.com/cn/zh/multimedia/documents/technical-publication/end-user-mapping-next-generation-request-routing-for-content-delivery-technical-publication.pdf
[9] *Supra*, note 6, at 14-16.

intelligently, with packets of some types of services—like text—dropped before others. However, a complete ban on prioritization would not permit such intelligent dropping of packets, and therefore, not allow some services, applications, and devices to function properly.

Traffic is also managed by establishing bilateral interconnection agreements between various ISPs, transit service providers, and private CDNs. Network operators may decide to geographically diversify their data centers, interconnect with transit providers to increase their available bandwidth, or even interconnect directly with edge providers that send large amounts of data over their networks. These private commercial agreements provide additional pathways for ISPs to deliver data across the Internet. There are often many paths to get data to an end-user, and network operators leverage these pathways to balance load and meet service quality agreements.

One technique, caching, has become increasingly prominent in recent years. Network operators have begun to place commonly accessed content at strategic nodes on the network to shorten the distance that it must travel to reach the end-user. Content providers with hundreds of millions of daily active users benefit because end-users can more quickly access content, and network operators are able to better manage their network loads. However, recent debate has centered on *where* caching takes place within a network, *who* pays to install caches, and *what* potential competitive concerns may exist through the use of caches.

D. Content Delivery Networks and Private Super-Highways

ISPs and backbone service providers are not the only network components that prioritize data as it is transmitted over the Internet, CDNs do, too. In fact, CDNs are a form of prioritization. Instead of sending data over the public Internet, or directly interconnecting with ISPs, edge providers can directly interconnect with a CDN. Those who use CDNs or other private peering arrangements can choose to route packets only sent by their clients who pay a subscription fee or arrange a bilateral peering deal. Some senders can build their very own CDNs, while others contract with third parties. The effect is the same: the edge provider's traffic is prioritized to provide a better user experience.

By 2021, it is estimated that CDNs will carry 71 percent of global Internet traffic.[10] Besides accelerating content delivery to end-users, CDNs are positioned on the network in a way that allows them high visibility into traffic flows passing through their networks. When used properly, this visibility allows CDNs to balance traffic load and manage congestion for their clients. However, this visibility also allows the CDN to see the location and nature of incoming traffic, which can be used for anti-competitive purposes like blocking or throttling. There are several choices that content providers have to ensure that end-users are able to access their content, including direct interconnection with an ISP, colocation of data centers and caching, and using a CDN.

[10] *Supra*, note 1, at 1.

Whereas ISPs aggregate all data traveling on the last-mile to consumers, content providers have upstream visibility and can see exactly who is sending what data, and where it is coming from. In addition, CDNs have another important tool that can be used to organize traffic: volume. CDNs that cache content and manage traffic for applications reaching hundreds of millions of daily active users can easily congest other network components, intentionally or unintentionally, if proper care is not taken. In many cases, their reverse proxy visibility allows CDNs to route traffic to specific interconnection points with various ISPs and transit service providers.[11] This traffic management technique can place pressure on ISPs that have not upgraded their networks to handle high levels of traffic and can disrupt end-users' ability to receive service. While the primary objective of CDNs is to enhance the user's experience, there is potential for CDNs to push traffic to achieve a better bargaining position on interconnection agreements.

E. How Wireless is Different

Wireline and wireless broadband services are different, chiefly in the amount of bandwidth they can provide. Wireline, through use of fiber, coaxial cable, twisted copper, or whatever other medium, has high-bandwidth capability. Because radio spectrum is the medium of wireless broadband, the bandwidth resource becomes more scarce. As demand for mobile broadband skyrockets, mobile network operators are grappling with ways to ensure consumers can connect.

As 5G standards develop, one key technological aspect will be the use of "Network Slicing." Network Slicing is a form of virtualization[12] where independent logical networks operate over a shared physical infrastructure. Network operators could use this technology to carry voice traffic over one slice, video over another, and each type of traffic could be optimized accordingly to manage the scarce bandwidth available on the network. This type of network could be great in an Internet of Things (IoT) world, where intermittent transmission of data over narrow slices of bandwidth can make more room for broadband-intensive applications.

Finding these types of efficiencies becomes increasingly important as technology evolves to allow for more data-dependent services. This is even more critical with the advent of important accessibility services. Streaming video—whether over Wi-Fi or cellular networks—to an agent who can assist visually impaired people in navigating their surroundings, teleoperation of vehicles for commercial and industrial use, and remote surgery in the next step of telemedicine, can have significant social impacts. But the growth of these services requires the network to manage and grow to accommodate this type of traffic.

---

[11] Matthew Weant, "Fingerprinting Reverse Proxies Using Timing Analysis of TCP Flows," Naval Postgraduate School, at 70-71. Available at: http://www.dtic.mil/dtic/tr/fulltext/u2/a589425.pdf

[12] For more information on virtualization, *see*, https://www.networkworld.com/article/2174268/virtualization/tech-primers-understanding-the-differences-between-software-defined-networking-network-virtualizati.html

F. Internet Prioritization at the Network Edge

As described above, there are several layers in the network as well as different network and edge components that can introduce bottlenecks. These bottlenecks can occur anywhere along the path between an end-user and the applications and services offered on the Internet, including the network edge. As the amount and nature of data traveling across the world continues to evolve, it is important to foster competition and innovation so the Internet continues to grow and function as intended.

While network operators have visibility into the end-to-end performance characteristics of packets, edge providers have their own ability to introduce congestion. By knowing where the end-user who requested their data is located, and the likely route the data will take to get there, high-volume edge providers can identify congestion and alter their interconnection paths accordingly. The ability to do so becomes easier with the use of CDNs—visibility into these flows creates an opportunity for directed congestion. Content providers with enough volume can leverage that volume in several ways, including by using it to directly pressure network operators to upgrade specific nodes on their network to carry the content providers' traffic.

There are several other forms of prioritization that occur at the network edge as well. Search engines, for example, optimize result pages to put some websites higher on the list than others. In the case of a search engine on your mobile phone, it will prioritize the mobile sites in the search. Some sites pay to have their website prioritized in a search query, whereas other affiliate companies can have their pages prioritized through the parent companies' search engine free of charge. Platforms on the edge also engage in prioritization of advertisements, elevating the profile of some ads over others. Again, companies can pay for this, or benefit from their affiliate status. When talking about prioritization, it is important to evaluate which scenarios benefit user experience and which can be harmful or anticompetitive. As the Committee continues its attention on ensuring that the Internet is free and open for everyone, it is important that all parts of the internet ecosystem are discussed.

IV. ISSUES

- How is data prioritized over the Internet and private networks to ensure that the Internet works for everyone?

- Is the concern about paid prioritization actually about affiliate content? Are there other ways to address that aspect other than blocking such avenues?

- When we discuss the optimization of data transfer rate, does this necessarily mean that some data is "slowed down" to make a "fast lane"?

- What are the implications for prioritization over wireline vs. wireless broadband networks?

- How do we make sure that the rural service gap does not widen as content caching is focused on urban and suburban areas?
- What separates healthy prioritization from anti-competitive forms of prioritization?
- Where could helpful or harmful forms of prioritization exist on the network or on the edge?

V. STAFF CONTACTS

If you have any questions regarding this hearing, please contact Robin Colwell or Tim Kurth of the Committee staff at (202) 225-2927.