Report for USTelecom – The Broadband Association

Evolution of the internet in the U.S. since 2015

Dr Michael Kende, David Abecassis, Elias Djaoui, Clarissa Wang

December 12, 2023

Ref: 858477899-502

Contents

This report has been prepared by Analysys Mason for USTelecom – The Broadband Association. This report was produced independently by Analysys Mason and does not necessarily reflect the views of USTelecom or any of its members.

Analysys Mason Limited
North West Wing, Bush House
Aldwych
London WC2B 4PJ
UK
Tel: +44 (0)20 7395 9000
london@analysysmason.com
www.analysysmason.com
Registered in England and Wales No. 5177472

1 Introduction[1]

The Federal Communications Commission (FCC) is proposing, through a notice of proposed rulemaking entitled *Safeguarding and Securing the Open Internet* (the '*2023 NPRM'*), to again reclassify broadband internet access service (broadband) as a telecommunications service subject to the common carrier provisions of Title II of the Communications Act. While efforts by the FCC to establish open internet principles began in 2004, they were only turned into enforceable rules by the *2015 Open Internet Order*, which reclassified broadband as a Title II service.[2] This reclassification, however, was subsequently reversed in the 2018 *Restoring Internet Freedom Order*.[3] During this period Congress has not acted on any federal legislation, although it has considered it in the past, and could in the future as well.[4]

Thus, in practice, the evolution of the internet in the U.S. has taken place largely outside the oversight of the FCC. As noted by then FCC Chairman William E. Kennard, "with competition and deregulation as our touchstones, the FCC has taken a hands-off, deregulatory approach to the broadband market".[5]

In particular, enforceable national rules, such as those the FCC is proposing now, were only in effect between 2015 and 2017. Before that period, and since, the internet has grown in scale and scope at a breathtaking pace, beginning in the U.S. with funding by the Defense Advanced Research Projects Agency (DARPA) in the 1960s, and initial development as an academic and research network in the 1970s. In the 1980s, it was fully commercialized, and became mainstream with the introduction of the World Wide Web in the 1990s. The subsequent dot-com boom resulted in the fast-paced introduction and growth of new content, applications, and services along with new generations of broadband access services, leaving people and businesses across the U.S. increasingly reliant on the internet. As shown below, after 2017, this growth in the breadth and depth of the internet not only continued, but also met the increased demands resulting from the Covid-19 pandemic.

1 The main author, Michael Kende, began his career in telecommunications at the FCC following the 1996 Act, where he wrote a Working Paper on internet interconnection. Since then, he has focused on internet policy as a partner at Analysys Mason running the U.S. office and starting the internet practice; as the Chief Economist of the Internet Society; and currently as a Senior Advisor with Analysys Mason and independent advisor to international organizations and foundations. David Abecassis is a Partner at Analysys Mason, focusing on policy and regulation in the telecoms, media and internet space. Elias Djaoui and Clarissa Wang are Associate Consultants at Analysys Mason.

2 "Protecting and Promoting the Open Internet", WC Docket No. 14-28, Order on Remand, Order, Declaratory Ruling, 30 FCC Rcd 5601 (2015).

3 "Restoring Internet Freedom," WC Docket No. 17-180, Declaratory Ruling, Report and Order, and Order, 33 FCC Rcd 311 (2017).

4 A handful of states, including notably California, passed their own net-neutrality laws after 2018.

5 FCC Chairman William E. Kennard, "The Unregulation of the Internet: Laying a Competitive Course for the Future," July 20, 1999 speech. See also Jason Oxman, "The FCC and the Unregulation of the Internet," FCC Office of Plans and Policy Working Paper No. 31, July 1999.

This report highlights the principles under which the internet successfully grew without Title II regulations, and provides evidence as to how these principles continued to underpin the development of the internet in the years since broadband service was initially reclassified.

The internet is fundamentally open by design, one of the foundations guiding its developers from the start and enshrined in a host of internet technology standards. New networks can enter and negotiate interconnection with existing networks to offer the services of their choice, while existing networks can grow and upgrade as needed. New content providers can in turn enter the market and offer innovative new content and services (which we refer to generically as applications in this document), limited only by their imagination and means, without any prior coordination needed with any networks.[6] And finally, and perhaps most importantly, users can use the applications of their choice, with the devices of their choice, with any broadband internet access service.

In addition to openness, another foundation guiding the development of the internet was for it to be decentralized, without any single owner or gatekeeper, and for its architecture to be simple, standardized, and scalable, based on modular building blocks that can be assembled in different ways to deliver services. These foundations resulted in what we have referred to in a recent paper as the three 'design principles' of the internet: that applications are developed and delivered separate from the network (the layering principle); that the internet is developed through the interconnection of networks that otherwise operate independently from one another (the network-of-networks principle); and that application-specific features can be added in the end points of the network (the end-to-end principle).[7]

These principles have implications that are relevant to the *2023 NPRM*. First, applications can multiply and grow in the absence of net neutrality regulations; second, internet traffic exchange has historically been negotiated voluntarily, in the absence of regulations, and evolved to meet the changing mix of services on the internet; and finally, application providers, organizations, and end users can implement technologies that use encryption to keep their data private during transmission across the internet from any outside parties, including networks. We discuss these design principles, and their implications, in Section 2 of this report.

In Section 3, we provide evidence that these principles do not just hold in theory, but have been, and continue to be, effective in practice in guiding the vibrant and open growth of the internet, including since 2017. This period of time is notable for two reasons. First, it coincides with the period during which the FCC reversed its Title II reclassification of broadband service, and second, it corresponds with the increased demands put on the internet during the Covid-19 pandemic, which the internet was able to meet in the U.S., based on the underlying principles, in the absence of regulation.

6 We define content providers here as providing online content and services to internet users. The definition is broadly similar to the FCC's definition of "edge providers", however here it includes intermediaries such as content delivery networks (CDNs) which deliver content on behalf of other providers, but excludes providers of devices used to access applications.

7 For more details on the internet foundations and design principles, see Analysys Mason, "Study on the Internet's Technical Success Factors," Report for APNIC and LACNIC, December 2021.

In sum, the internet – including both broadband provided by Internet Service Providers (ISPs) and applications developed by content providers – has blossomed and thrived for decades based on the design principles of the internet, and that did not change after elimination of Title II in 2017. There is no evidence of market failure that would prevent the continued introduction and use of applications in the absence of common carrier regulation. Commercial negotiation of interconnection arrangements has led to effective and efficient, largely decentralized internet infrastructure, co-invested by stakeholders of all types. Finally, innovations, including the use of encryption technologies, provide additional privacy to any users or applications seeking to protect their data.

2 How the design principles of the internet have supported its uninterrupted growth

The internet has grown tremendously in breadth and depth over the past decades, marked by an exponential increase in the number of users and their usage of the internet. Internet access has evolved from slow dial-up access to fast fiber connections, and from early wireless access for voice all the way to 5G connections capable of supporting sophisticated wireless applications for individuals and enterprises. At the same time, matching the increased bandwidth in the new access networks, applications have grown from text-based asynchronous applications to real-time voice and video, and the general-purpose nature of the internet is enabling constant innovation, most recently leveraging artificial intelligence, in a virtuous circle of growth. Finally, these applications are available to end users over an ever-widening array of internet-enabled devices.

The success of the internet is built on three principles of the internet that have remained constant in a sea of change: layering; network of networks; and end to end. Adherence to these three principles explains how the internet has grown and adapted to all the changes that have driven the increase in usage, and how innovation has thrived in applications and networks while the fundamental open internet structure has remained intact. It also explains how the internet has prospered in the U.S. largely in the absence of common-carrier regulations. This section discusses these principles in more detail, and in the next section, we provide evidence of their outcomes showing that they are being followed in practice.

2.1 The layering principle has led to the huge and ever-evolving variety of applications available over the internet

The *layering* principle relates to the fact that internet applications are provided separately from the networks, with the Internet Protocol (IP) acting as an interface between these two 'layers' of applications and network. This is a very broad simplification, but in practice it means that any application that can be effectively encapsulated and delivered via IP, such as voice over IP (VoIP) or streaming video, can be provided over any network that is accepting this protocol.

Layering was so successful that it gave rise to a saying: "everything over IP, IP over everything". This is illustrated in Figure 2.1 below and effectively means that the content providers offering content, applications, and services, can operate separately from ISPs, without any required coordination or tailoring of the applications to any individual network.[8] While coordination between

[8] As discussed in Section 2.2, content providers do need to interconnect with at least one network to deliver their traffic to and from their end users. Often content providers choose to connect to multiple networks in multiple locations, to deliver their traffic efficiently. In order to exchange traffic through interconnection, each network must adopt the Border Gateway Protocol, in addition to the IP that they need to carry traffic through their network.

content providers and ISPs is not required, there are cases in which voluntary coordination can improve the user experience in accessing content and services, described further below.

The implications of layering run deep. A new service, such as online gaming or distance learning, can be developed, introduced, and updated, without any coordination with any network providers in the core or at the edge of the network. Given the thousands of existing networks, such an effort would be otherwise prohibitive. Conversely, networks can be upgraded, or new technologies introduced, without coordination with application providers, which would also be prohibitive given the millions of websites, applications, and services that exist. This is the basis for the development of the open internet that exists through today.

Figure 2.1: Layering principle [Source: Analysys Mason, 2023]

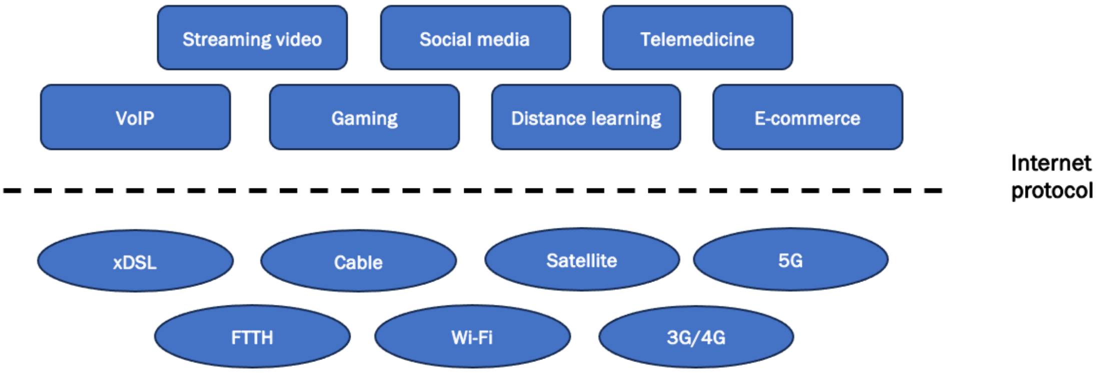

Layering is also the basis for the other two principles, in which networks interconnect to create the internet, while application-specific intelligence can remain at the edges of the network.

2.2 The network-of-networks principle enables networks to act independently while interconnecting to form the internet

The internet is a *network of networks*. A recent study showed that there are over 90,000 networks that can interconnect directly or indirectly to form what we know as the internet.[9] As long as each network uses the internet protocols and shares routing information with at least one other network, they can – at the same time – operate independently from one another, and interconnect so that traffic can flow freely between and through them from any origin to any destination. The need for interconnection was evident from the beginning, as one of the internet's foundational goals was to enable communication between different computers in scattered locations.

9 This is measured by the number of autonomous system (AS) numbers that are in use – an AS is a network or group of networks that is unified, typically operated by a single organization, which interconnects with other ASs to exchange traffic, which it distributes internally to its connected users. The AS number is a unique identifier used to identify networks for traffic delivery. The study showed that that there were 99,868 AS numbers that had been assigned to networks, and at the time 71,970 were active in Border Gateway Protocol (BGP). George Michaelson, "How do we count ASNs," APNIC blog, March 19, 2021.

Peering and transit evolved as ways for networks to exchange traffic with one another

While first held together by connections to a government backbone, as the commercial internet emerged, the early ISPs needed to find a means to exchange traffic with one another. In the cooperative ethos that marked the development of the internet, and which continues through today, the networks developed a form of interconnection known as peering. Peering is a bilateral form of interconnection in which two networks exchange their customers' traffic with each other.

Historically, peering was between 'peers' with similar scale and business models: when traffic flows were reasonably symmetric peering was typically 'settlement free'– that is without payments in either direction. Peering has expanded to include direct bilateral interconnection between companies with different traffic patterns and business models, such as content providers and ISPs. Settlement-free peering is still widespread, but paid peering can be arranged and is relatively common – the constant is that peering remains a bilateral exchange of traffic.

Peering is a mutually beneficial means to exchange traffic, and has historically been voluntarily negotiated, without regulation.[10] Indeed, in the vast majority of cases according to one series of surveys, there is not even a formal contract but just a 'handshake' agreement for the peers to exchange traffic.[11] As peering is a bilateral arrangement, if this was the only way to exchange traffic between two networks, the internet would need to be fully meshed: each network would have to peer with every other network in order for the internet to be completely interconnected. As the number of networks grew, the number of bilateral peering links would have exploded if supplemental means of interconnection had not been developed – but peering remains fundamental to how traffic flows between networks.

A first supplement to bilateral peering is a form of 'shared fabric' peering between multiple networks, at facilities called *Internet Exchange Points* (IXPs). An IXP is essentially a shared switch, which enables every connected network to exchange traffic with any, or all, of the other connected networks. For instance, as shown in the figure below, two backbones, an ISP, and a content provider can all exchange traffic through the IXP. The result is that each network just needs one connection to exchange traffic with the other three networks, rather than three separate connections.

Connection through the IXP is known as *public peering*, as all connections take place through the same switch available to every network connecting at the IXP. The alternative is *private peering*, sometimes known as *private network interconnection* (PNI), in which two networks directly connect with one another to exchange traffic. This gives the interconnecting networks full control to avoid any congestion over the connection, and is often preferred once the amount of traffic exchanged between two peers passes a certain threshold. Often public and private peering can be arranged in

10 See, for instance, Michael Kende, "The Digital Handshake: Connecting Internet Backbones," FCC Office of Plans and Policy Working Paper No. 32, September 2000.

11 Bill Woodcock, Marco Frigino, "2021 Survey of Internet Carrier Interconnection Agreements," Packet Clearing House, December 17, 2021.

the same data center, which has an IXP node to enable public peering while also allowing direct cross-connects for private peering between networks.

Peering arrangements are illustrated in Figure 2.2 below. In this diagram, Backbones 1 and 2, the content provider, and ISP 2 are all peering publicly through the IXP. The content provider and ISP 2 have also created a private peering connection between themselves, while using the IXP to exchange traffic with the two backbones.

Figure 2.2: Interconnection arrangements [Source: Analysys Mason, 2023]

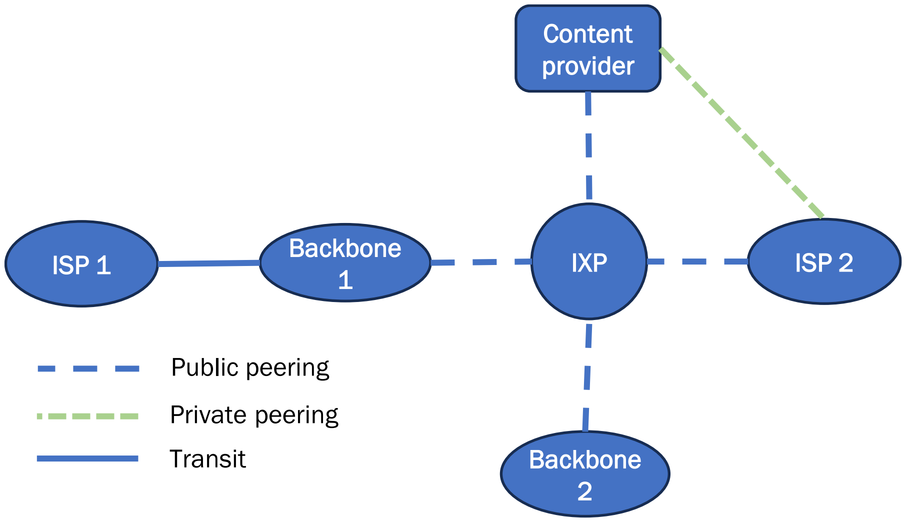

The diagram also highlights another form of interconnection called *transit*, and the need for it. Under bilateral peering arrangements, one network does not allow peered traffic to transit through its network to another peer. Thus, in the diagram above, if ISP 1 was peering with Backbone 1, it could not get access to traffic from ISP 2 or the other networks peering at the IXP. Peering would only enable ISP 1 to get access to the customers of Backbone 1. It may not be efficient for ISP 1 to arrange peering with the other networks at the IXP, not to mention all the rest of the networks making up the internet. As a result, *transit* developed as a way for smaller, more localized networks to exchange traffic between one another and with networks far away from them around the world.

In a transit arrangement, one network will sell access to all of its connections to another network, including peering connections. For instance, in the diagram above, Backbone 1 sells ISP 1 access to all of its peers at the IXP, and elsewhere, thus relieving ISP 1 of the need to arrange for peering. In this light, peering can be viewed as an input to the sale of transit by larger backbones. Transit is common, as there are only a handful of large global backbones, known as Tier 1 backbones, which rely solely on peering to provide their customers with access to the entire internet. In practice, the more peering arrangements a network can arrange, the less transit it needs to purchase on the open market, and this in turns puts downward pressure on the price of transit. Transit, like peering, is a voluntary arrangement that is commercially negotiated and not regulated.

Commercially negotiated peering and transit have allowed the scope and scale of the internet to grow unimpeded by gatekeeping and complex commercial arrangements. The internet now consists of a web of peering and transit arrangements, providing for the possibility of alternative routes between any end points, including content providers and ISPs, allowing both for resilience and multiple options for delivering traffic as interconnection arrangements are negotiated and evolving. Indeed, there have been many changes in content and services and business models, leading to changes in the models of internet interconnection, but not the fundamentals.

One change occurred as the internet spread from its academic and government roots in the U.S. to the rest of the U.S. and the world, requiring traffic to be delivered to users further and further away. The number of IXPs multiplied, to localize traffic exchange, while investment in networks to deliver traffic across and between countries and continents grew.

At the same time, content grew from plain text, to multimedia, to audio and then video. As content traffic grew, it became costly for content providers to deliver their traffic from its source data center to ISPs all over the world, typically by paying for transit to one or more backbones. As a result, content providers began to store their traffic in multiple data centers to shorten the distance to the ISPs, and a number began to build their own data centers. Some content providers even began to invest in their own terrestrial or submarine cable networks to deliver their traffic between data centers. A new business model for delivering content developed, and peering arrangements continued to evolve.

Content delivery networks leverage peering, transit, and decentralized interconnection to increase the quality of experience of end users and reduce costs for content providers and ISPs.

In order to address the changes in the quantity and nature of content, *content delivery networks* (CDNs) arose that invested in caches – essentially servers with large amounts of memory for storage – to deliver content requested by end users. These caches store popular content such as webpages, software updates, and video that can be delivered to end users directly from the caches, while the content stored in the cache can change frequently in response to demand. The result was that popular content could be sent once to the cache, and from there multiple times to end users, saving significant capacity in the core of the network and enabling end users to more quickly access the content of their choice. These caches are deployed and maintained by content providers, ISPs, and independent CDNs that provide the benefits of caching to smaller content providers. CDNs and ISPs can and do partner to ensure content is retrieved efficiently from the most relevant cache.

Caches can be placed in a variety of locations. Many caches are available at IXPs, and others are in points of presence closer to, or embedded within, the ISPs' networks, as a result of cooperation between content providers and ISPs to improve quality of experience for end users. Embedded caches are provided in partnerships, with ISPs providing space, power, and connectivity to caches that are built and configured by CDNs. As a result of these shifts, some content providers truly became *edge* providers, delivering their content closer to the ISP subscribers at the edge of the network.

Having moved the content closer to the end user, it is almost always still exchanged through peering arrangements. Many content providers have open peering policies at IXPs, exchanging traffic with any other network upon request. While settlement-free peering is still common, paid peering arrangements, in which the content provider compensates an ISP for delivering the content, are also being negotiated where traffic flows between the parties are substantially asymmetric. In spite of these changes, these arrangements are still negotiated commercially, mostly without settlements or even a contract.

Packet Clearing House conducted a global survey of interconnection, with 17,192 networks responding, of which 3,813 were in the U.S. These networks reported 15,105,101 interconnection agreements – that is, an average of 878 per network – and found that 99.998% were handshake agreements with no contract, and 99.9996% of them were symmetric, i.e. settlement free.[12] This illustrates the large number of networks negotiating interconnection and not just ISPs and content providers; however the numbers do not reflect the quantity of traffic exchanged, as networks turn to private peering – which can include settlements and be formalized in a contract – for exchanging large quantities of traffic as well as exchanges between networks with asymmetric traffic flows (for instance between the content of content providers and the end users of ISPs).

All of these changes came about as a result of changes in the quantity and nature of traffic that is developed by content providers for consumption by end users. Content providers invested in their own infrastructure and CDNs emerged to deliver traffic closer to ISPs, often coordinating to help ISPs locate the most convenient cache for their subscribers, and peering arrangements were voluntarily negotiated between them for the exchange of traffic. These evolutions have been constant since the emergence of the commercial internet, without the presence or need for regulatory oversight.

2.3 The end-to-end principle enables application intelligence to remain at end points without any application-specific interaction with the network

The end-to-end principle reflects a paradigm in which the application-specific ‘intelligence’ is managed at the edge of the networks, in the wide range of devices that send and receive packets using IP. The core of the ISP networks, however, also contains intelligence, for routing traffic between networks over peering and transit connections, running domain name system (DNS) servers, and caching content.[13] There can be coordination between applications and networks, for instance to help access cached content efficiently, but it is not strictly required for an internet application to function end to end.[14]

An important implication of the end-to-end principle is that applications using IP can be developed independently, installed in any relevant (and compatible) device, and used on any IP-enabled network, without any application-specific support in the network. The result is an endless variety of

12 Packet Clearing House, “2021 Survey of Internet Carrier Interconnection Agreements,” December 2021.

13 The DNS enables translations between domain names (such as www.example.com) input by users, and the underlying numerical IP address used to locate and route the relevant applications. The default DNS server is typically provided by the ISP, and the DNS service is provided free of charge.

14 The end-to-end principle also does not preclude application awareness in the network, which can take the form of intelligent routing coordinated between networks, and in the near future, new applications through ‘network slices’ built into the 5G wireless standards.

applications – only limited by the imagination of the developers and the means available to them – which end users are free to use with any compatible internet-enabled device.

As noted, users connected to the internet via any ISP can use any applications, such as a messaging app, on their device, from any provider, and sign up with the provider to identify other users of the same app, who may be connected to any other ISP. When a message is sent or a call is made, it physically flows through the network from user to user, however logically the application provider routes the message or call from origination to termination – it is for this reason that such applications are sometimes referred to as over-the-top (OTT) services, in contrast with applications whose provision is integrated within the functions of the underlying network.

One key implication of the end-to-end principle is that it enables the increasing use of encryption on the network, as shown in Figure 2.3 below.

Figure 2.3: End-to-end principle [Source: Analysys Mason, 2023]

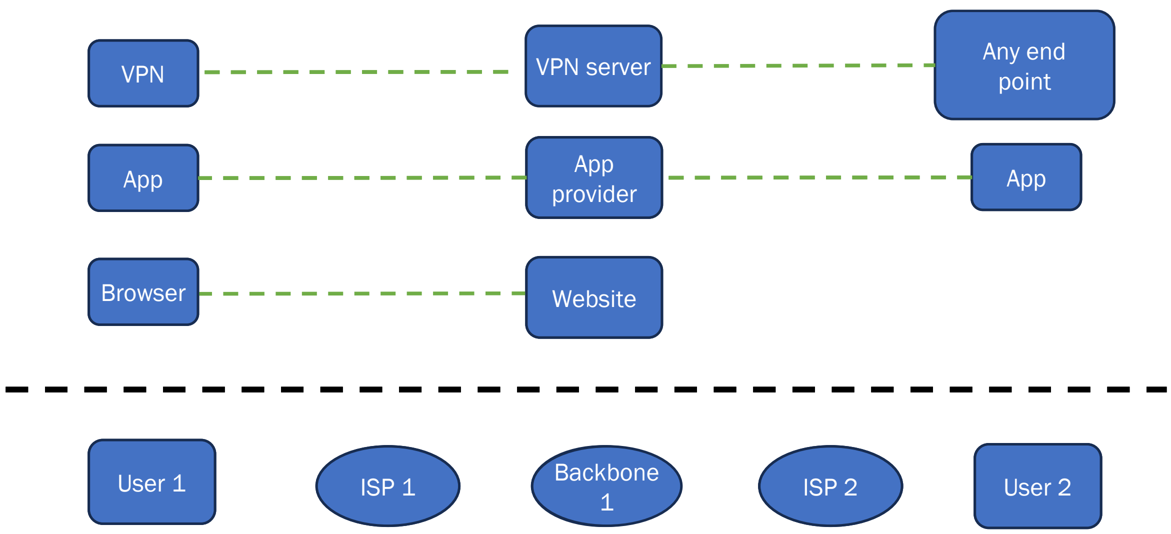

This illustrates three different scenarios:

- Websites can use encryption – for example using the HTTPS protocol – to protect sensitive information such as credit card numbers for e-commerce sites. A user accesses the website through a browser, and the communications are automatically encrypted, without any action by the user or interaction with the ISP.
- Communications services such as WhatsApp began to implement end-to-end encryption (E2EE) to protect users from government surveillance, in which case even WhatsApp cannot read the traffic sent using its application. Interestingly in the case of WhatsApp, E2EE follows an open-source protocol from a non-profit organization called the Signal Foundation, which operates its own messaging application of that name but also makes its protocols available for other applications.[15]

15 See https://signalfoundation.org/en/ and https://faq.whatsapp.com/820124435853543

- Finally, users themselves can choose to use virtual private networks (VPNs), which provide an encrypted 'tunnel' over the internet to a VPN server and from there to any end point, which could be another user, an online service, or a server on a corporate network. From the perspective of the ISP, VPNs are opaque channels that preclude them from seeing the content of the packets, but also the ultimate destination of the packets if relays are in use.

As a result, a website or service provider can encrypt its traffic from or between users, or users can choose to encrypt all their traffic with a VPN. The result is that the content is encrypted, although the destination may be known, meaning no provider, including the content provider and ISP, can read the traffic, and no third party, including governments or hackers, can access any of the contents of the traffic. These uses of encryption are increasing, as shown below in Section 3.3. This is a further illustration of the modularity inherent to internet applications, fully independent from the underlying networks.

***

In summary, since IP was introduced in 1974, and deployed in 1983, and continuing through today, it has been the foundation of the three principles of the internet discussed above. In particular, the result is that the internet is decentralized, with no gatekeeper. Internet-enabled networks can operate independently from one another, and interconnect with each other directly or indirectly; applications using IP can run over any network independent of the network, and end users are free to install and use any legal application on any device, with tools available to keep the contents of their usage private.

3 Evidence that the internet is growing and evolving

The design principles of the internet create the conditions under which the internet can grow and evolve, as content providers, ISPs, and users can each act independently to provide access to content, applications, and services for the collective good of the internet. In the time since the FCC returned broadband to a Title I service, not only did the internet continue to grow, but it also absorbed the shock of the Covid-19 pandemic, meeting the increased demands of everyone needing to rely on it as a lifeline for work, school, government services, social interaction, and entertainment. Investment, innovation, and adoption of new applications all carried on at a rapid pace after the FCC returned broadband to Title I in 2018, even in the face of the global pandemic that dramatically changed internet traffic patterns and loads overnight.

3.1 The internet has been growing continuously over the past decade

As a leading overall metric, the number of internet users in the U.S. has continuously increased since 2017, as shown in Figure 3.1. Recent estimates show that 93% of Americans used the internet in 2022.

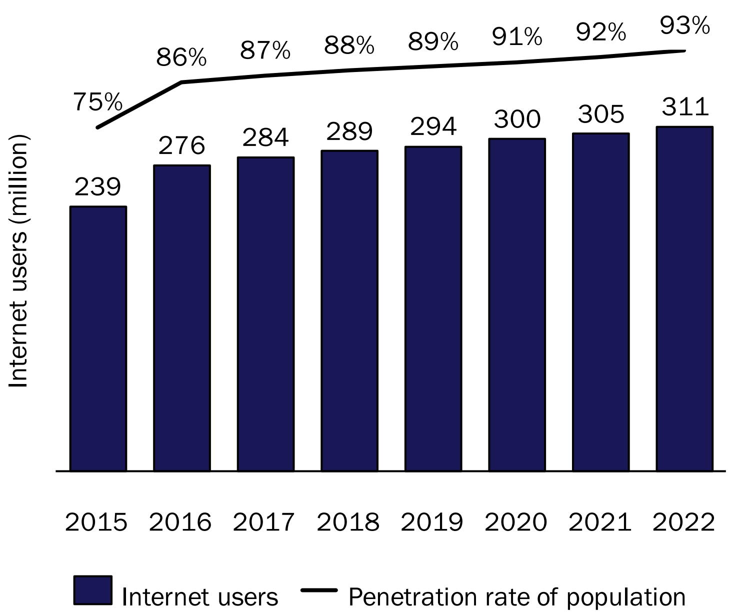

Figure 3.1: Internet users in the U.S.[16] [Source: ITU, 2021, Datareportal, 2023]

These users may go online with fixed broadband connections at home or with mobile broadband connections, and in many cases both. As shown in Figure 3.2, the number of fixed broadband connections has continued to increase, in particular throughout the pandemic as more households relied on a fixed connection to meet all of their increased data needs. In addition, there has been an

16 Number of internet users calculated as U.S. population multiplied by ITU penetration rate of internet users. ITU defines internet users as the proportion of individuals who used the internet from any location in the last three months. Access can be via a fixed or mobile network.

evolution in the types of connections. While cable modem access continued to grow and represented the largest number of fixed broadband subscribers, DSL connections tapered off in favor of higher-speed fiber-to-the-premises (FTTP) or fiber-to-the-building (FTTB), which represent the fastest growing number of fixed broadband subscribers.

Figure 3.2: Broadband connections by technology [Source: Analysys Mason Research, 2023]

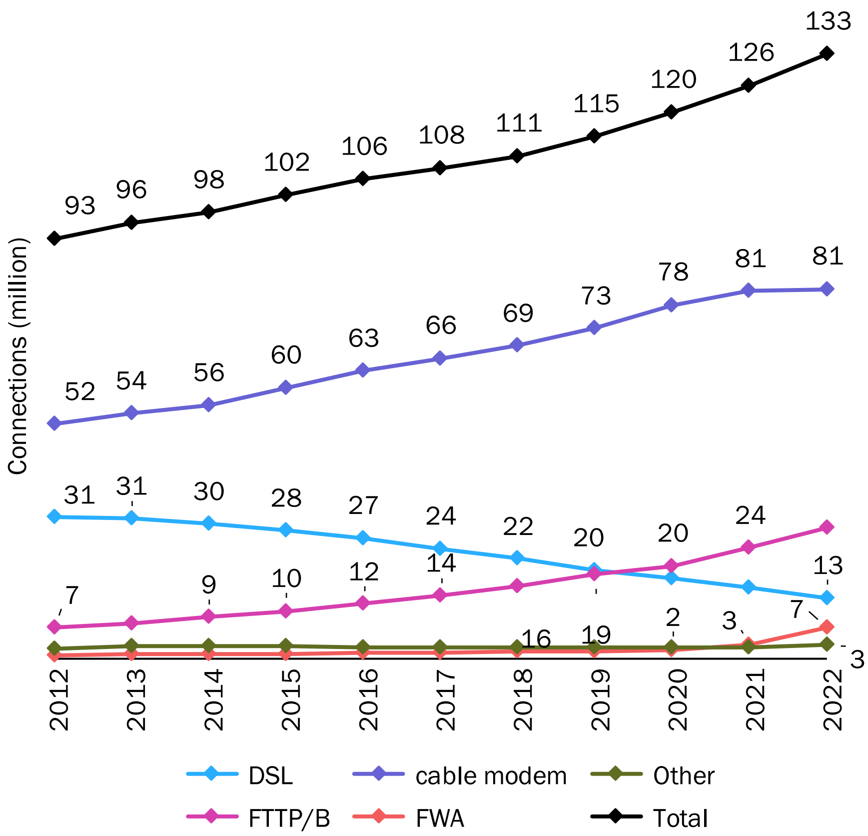

Likewise, as shown in Figure 3.3, the number of mobile connections has been steadily increasing, now significantly outpacing the number of U.S. residents given that many individuals have multiple subscriptions. As with fixed connections, not only is the number of subscribers increasing, but the technology mix is changing, with 2G and 3G fading almost to zero, in favor of 4G, which itself is now decreasing in favor of a significant increase in higher-speed 5G connections.

Figure 3.3: Mobile connections by technology [Source: Analysys Mason Research, 2023]

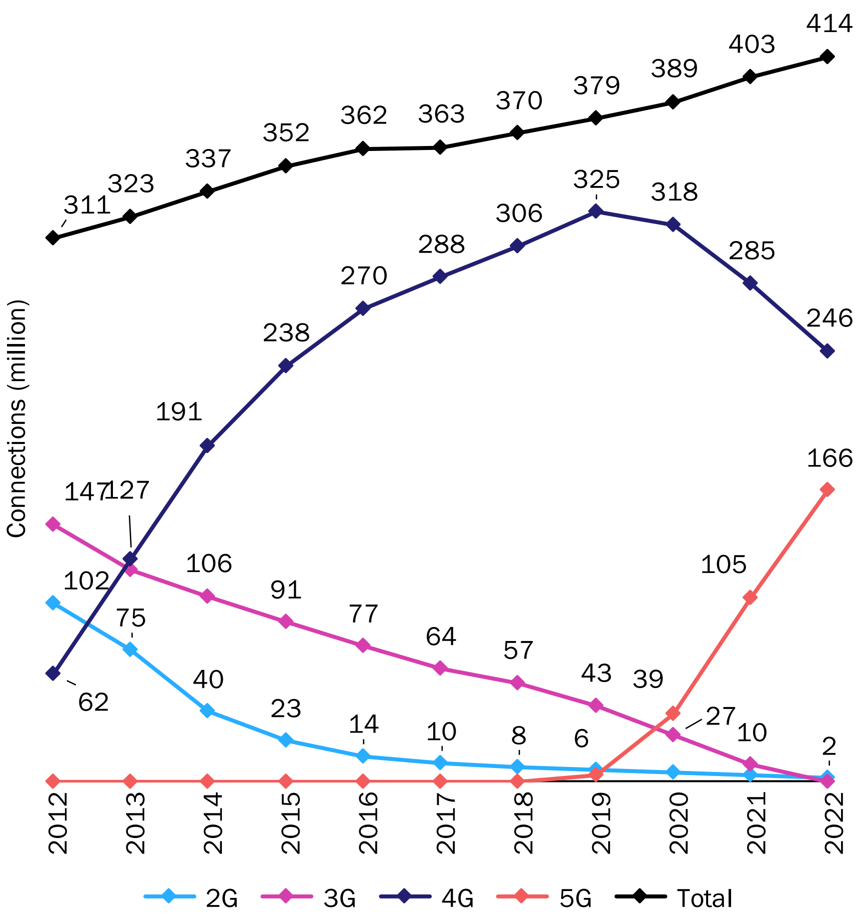

These increases in the number of subscribers, and the upgrade in the mix of technologies result from significant ongoing investment on the part of broadband providers in the U.S., as highlighted in Figure 3.4. While there was a slight dip in 2020 at the beginning of the Covid-19 pandemic, investment has recovered strongly, driving the upgrades in adoption described in the previous charts.

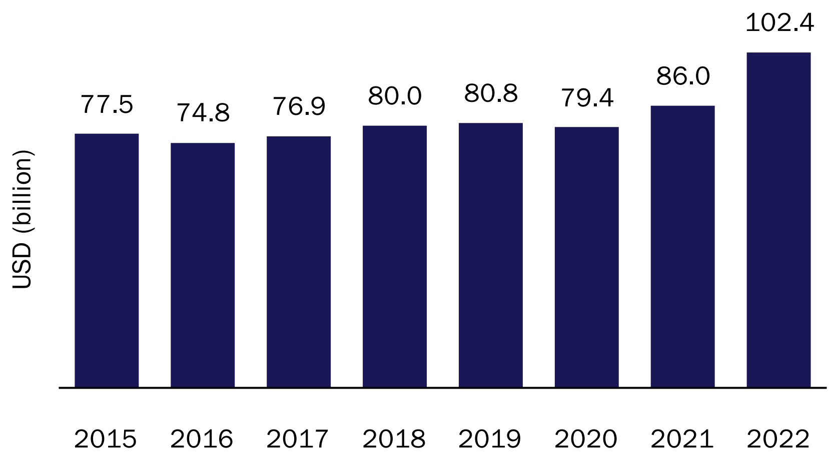

Figure 3.4: Capex by broadband providers [Source: USTelecom[17]]

These investments throughout the network have enabled the previously shown upgrades in fixed and mobile subscriptions, notably to fiber networks and 5G. The result has also been a significant increase in average download speeds enjoyed by users, as shown in Figure 3.5, as more users take up faster connections. These speeds have roughly quadrupled between 2017 and 2022, notably increasing throughout the Covid-19 pandemic, when usage also increased. While Netflix and YouTube had to decrease their streaming bitrate by 25% in parts of Europe during the early days of the pandemic, as requested by the European Union to avoid congestion,[18] in the U.S. such action was not requested or evidently needed.

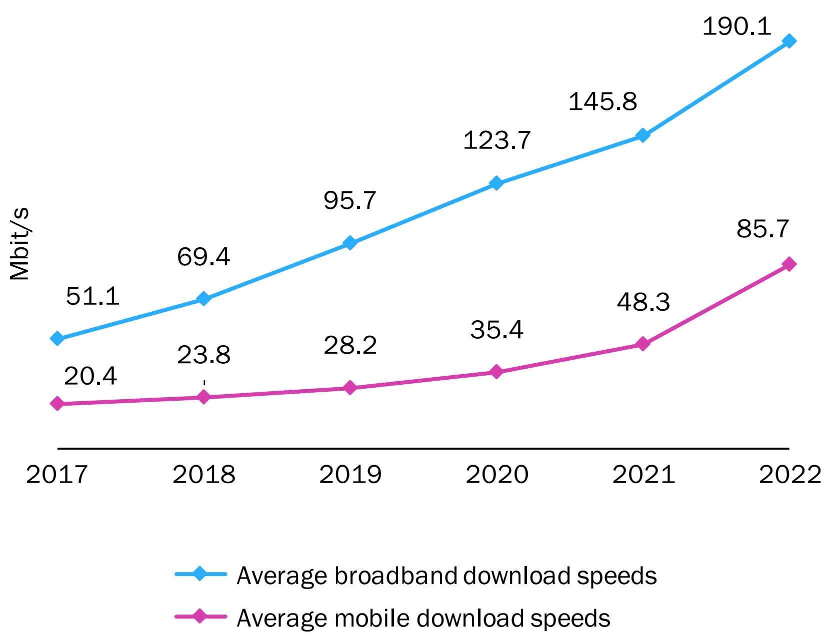

Figure 3.5: Average broadband download speeds in the U.S. [Source: Economist impact, 2023]

17 https://www.ustelecom.org/research/2022-broadband-capex/

18 https://www.theverge.com/2020/3/20/21187930/youtube-reduces-streaming-quality-european-union-coronavirus-bandwidth-internet-traffic

Finally, not only was there an increased number of connections, at higher speeds, but there was a significantly increased usage per connection, for both fixed and mobile connections as shown in Figure 3.6.

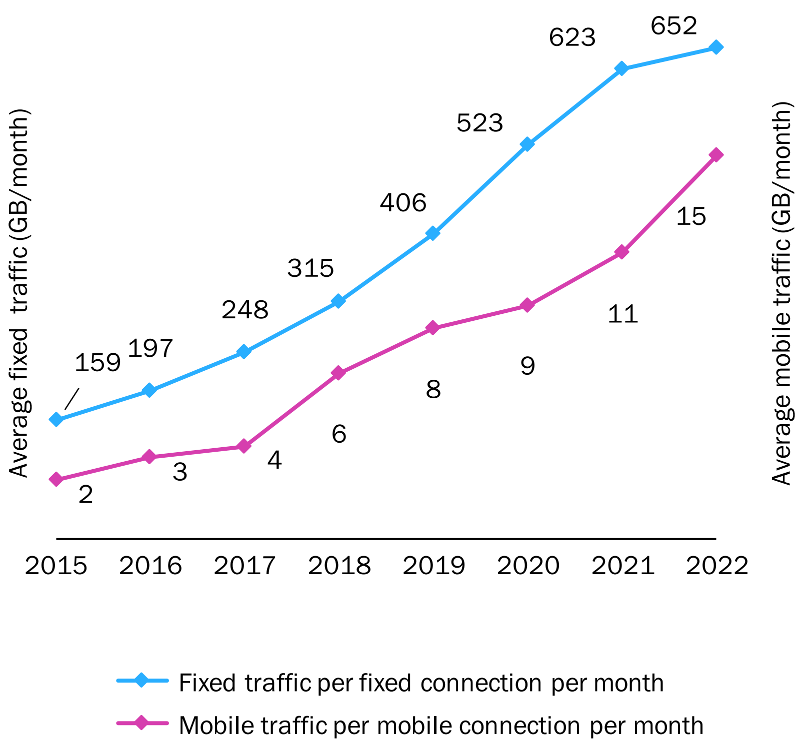

Figure 3.6: Traffic per connection in the U.S. [Source: Analysys Mason Research, 2023]

Broadband providers in the U.S. have made significant investments in upgrading their networks over the past eight years, evidence of a healthy outlook for the industry. The result has been an increased number of connections overall, as well as a speedy uptake of the new technologies, allowing for an increase in both connection speed and usage per connection. This growth not only reflects the fruits of the investments of broadband providers, but also results from the uptake of new and improved applications taking advantage of the new access networks. The uptake and usage of new applications, along with the content provider investments shown next, in parallel with the upgrades in broadband connections, highlights a virtuous circle of growth over the past years, including in the absence of the regulation proposed in the *2023 NPRM*.

3.2 Content providers are increasing their infrastructure investments to move interconnection closer to ISPs' own networks and to end users

Under the network-of-networks principle, any provider is free to develop a network and seek to adapt interconnection arrangements to changes in the market. Content providers have begun to invest more in infrastructure to deliver content to ISPs closer to the edge of the network, where interconnection is increasingly taking place. These investments are taking place organically, without the need for any regulatory oversight, in response to the changing mix of content and services, in order to reduce the cost of delivering traffic.

Infrastructure investment has helped content providers to deliver new content and services

The increased adoption and usage of the internet is driven by demand for online applications. In particular, there are two categories of applications whose usage increased significantly (as demonstrated in the next subsection): streaming video services, and real-time communications services. The former can require significant amounts of bandwidth to deliver (e.g. at high resolutions, or for live events), while the latter benefit from low latency, to ensure that the communications are clear and natural.

As demand for content increases, content providers have invested in *hosting* facilities that are necessary for the storage and processing of their applications. In order to increase the speed of their services and reduce latency, content providers have been making investments in *transport* to deliver content between their data centers. They have also invested in *delivery*, particularly in caching, as well as facilities to carry traffic between their data centers and where their content is cached. As shown in Figure 3.7, the average annual investment in infrastructure by content providers in the U.S. has been growing over the past decade, to an estimated USD42.8 billion per annum over the period 2018–2021.

Continued investments in digital infrastructure by both ISPs and content providers indicate that the current regulatory arrangements and market structure are conducive to the internet continuing to evolve, grow in scale and scope, to support the many different applications that are part of the fabric of everyday life for Americans. Content providers would likely not be investing at these levels if they believed that ISPs could impose unreasonable 'tolls' or if those ISPs were unwilling to accept traffic in places where costs and quality are optimized. Indeed, coordination between ISPs and content providers to configure networks for routing and DNS is often essential to deliver the full benefits of the changing internet architecture and decentralized interconnection, for example to maximize the effectiveness of caches.

Figure 3.7: Average annual infrastructure investment made by content and application providers in the U.S. [Source: Analysys Mason, 2014, 2018, 2022]

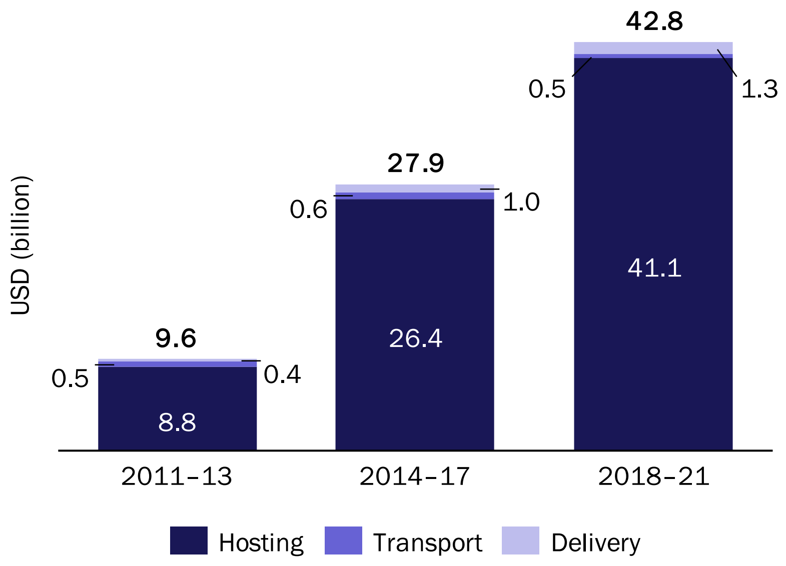

There is no evidence that the three principles of the internet described above are not effective in practice, helping to drive the growth and success of the internet.

Evidence of increases in interconnection arrangements

The infrastructure investment shown in the previous section helps to support shifts in interconnection made possible under the network-of-networks principle. As noted above, the internet is a network of networks, put together through interconnection arrangements.

In the absence of regulation, interconnection arrangements have adapted themselves to changes in traffic patterns and business models, and vary significantly. Peering arrangements can be public, through an IXP, or private through bilateral direct connections. Peering takes place throughout the network, between backbones, between ISPs, and with content providers, at public peering points or with caches embedded in the ISP network. Peering can be settlement free, without any formal contract, which is true for the vast majority as shown in the Packet Clearing House survey cited above, or peering can be paid. The only constants are that peering is a bilateral exchange of each network's traffic, and that it is commercially and voluntarily negotiated, and not regulated.

The number of IXPs has been growing steadily in the U.S. through the past several years, with Packet Clearing House counting 108 IXPs.[19]

Figure 3.8: IXPs in the U.S. [Source: Packet Clearing House, 2023]

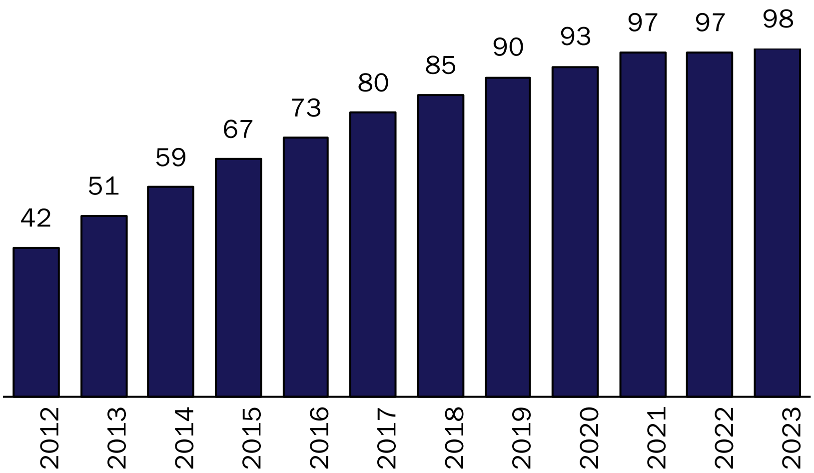

While the growth in the number of IXPs has slowed over the past years, this has not slowed the pace at which interconnection itself has grown. First, IXPs benefit from network effects: the more networks connected to an IXP, the more attractive it is for other networks to connect to the IXP to connect to those networks. Thus, one would not expect an endless growth in the number of IXPs. Further, a number of IXPs grow by placing nodes in multiple data centers, enabling virtual peering

19 Ten IXPs are listed in the Packet Clearing House directory without start dates, and thus are not included in the figure.

between customers in different data centers, without the need for a new separate IXP in each one. And finally, IXPs are for public peering, and there is also an increasing amount of private peering taking place between networks seeking to avoid congestion over direct connections that they provision between themselves.

As an example, one of these IXPs (DE-CIX New York) highlights the growth and variety of IXPs.[20] The average traffic exchanged there has quadrupled in the past five years, from 350Gbit/s in November 2018 to 1.47Tbit/s in November 2023. Currently, there are 276 networks peering at DE-CIX New York, including many of the large content providers and CDNs, as well as ISPs ranging in size from Charter Communications to NYC Mesh, a community Wi-Fi network in New York, and international backbones from as far afield as Eswatini in South Africa. In addition, an eclectic selection of schools, including Williams College and the Trustees of Deerfield Academy, the Board of Elections in the City of New York, and companies in the media, finance, and manufacturing industries are peering there. Many of these networks have open peering policies, under which they will peer with anyone else at the exchange. Others peer selectively.

Furthermore, DE-CIX New York is available in ten data centers, including interconnection locations where private peering takes place, such as 60 Hudson Street in New York City, formerly a Western Union building and now a significant hub for internet connections. Such interconnection facilities are located in major cities across the country. One provider, CoreSite, has 25 data centers in 11 major cities. One of these, located in Los Angeles, hosts more than 300 networks including a significant number of national and international carriers, provides access to several large cloud providers, and houses an IXP with 268 members.[21]

Interconnection arrangements are based on the mutual benefit of connecting to exchange traffic. Peering arrangements have evolved over time due to changes in the nature of traffic and business models. Investments by content providers in infrastructure have served to deliver content and services closer to end users, thereby helping to lower the transit costs faced by some ISPs that do not operate larger backbone networks, while the increased demand for content in the past years – notably through the Covid-19 pandemic – has only enhanced the existing incentives of broadband providers to work with content providers to ensure delivery of applications to their end users. One measure of the impact of these shifts is that transit prices have fallen significantly, showing that the interconnection marketplace is working efficiently in providing alternatives to transit.[22]

There is no evidence in the years since the *2015 Order* that broadband providers have exerted any gatekeeper role, nor that interconnection arrangements should be subject to any regulations. Instead,

---

20 DE-CIX New York is one of a number of public IXPs operated by DE-CIX, which started as an IXP in Germany. https://www.de-cix.net/en/locations/new-york

21 For a list of interconnection facilities in the U.S., see https://www.peeringdb.com/advanced_search?country__in=US&reftag=fac. For the CoreSite data center in Los Angeles, see https://www.coresite.com/data-center/la1-los-angeles-ca

22 In two large hubs, Miami and Los Angeles, for instance, the cost of IP transit has fallen by more than 10% on average each year between 2019 to 2022 (for 10GigE connections) and over 20% on average for larger connections (100GigE). Brianna Boudreau, "Price Erosion Remains the Universal Norm," TeleGeography blog, 5 October 2022.

interconnection arrangements have not remained stagnant, but have adapted and evolved to address the explosion in content traffic, through the flexibility provided by commercial negotiations. The choice of peering arrangements, including settlement free and paid, and the decentralization at IXPs or private peering locations, along with competitive transit offers, create a multitude of routes for content to get to end users, and constrain all networks' behaviors. Thus, in the current market conditions, the evidence shows that the existing regulatory framework is delivering good outcomes for all parties, with no known recent disputes or apparent market failure.

3.3 The scope and scale of OTT services have grown rapidly in the absence of net-neutrality regulations

As discussed above, layering allows content providers to offer applications 'over the top' without requiring any interaction with the network, to any user who wishes to use them. They are increasingly doing so in a way that ensures communications between end points (e.g. an end-user device and a server, or two end users) are encrypted, in such a way that no party other than these end points can read the content of their communications, or even know who the end points are, in some cases. The evidence shows positive synergies between applications and connectivity, to the benefit of end users increasing their adoption and usage of OTT services.

Under layering, content providers are able to offer a variety of popular OTT services

A number of OTT services compete with services provided by ISPs, leading to worries that ISPs may block or throttle them, but that claim is contrary to the evidence of strong growth in OTT and other services as discussed below, particularly in response to demand caused by Covid-19 restrictions. One early OTT service was VoIP, which enables voice calls over the internet between at least two parties. This service proved popular, and Figure 3.9 shows steep increases in VoIP traffic in the U.S. since 2015. VoIP traffic peaked in 2020, coinciding with the Covid-19 pandemic, and started decreasing thereafter, likely because users began to switch to other forms of communication, including those built into social media applications and messaging applications, some of which feature video calling.

Figure 3.9: OTT VoIP outgoing voice minutes in the U.S. [Source: Analysys Mason Research, 2023]

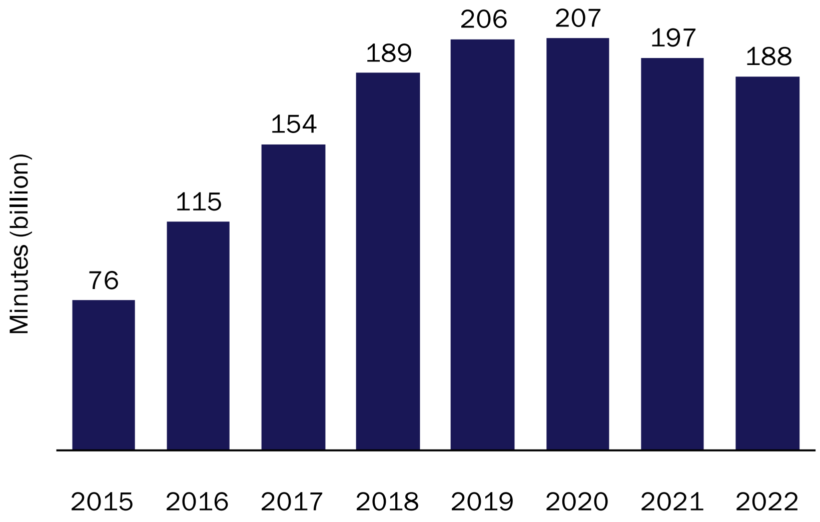

Figure 3.10 shows the rise in the use of messaging and video conferencing applications over the past few years, according to annual surveys of U.S. mobile subscribers over 18 (thus the numbers may understate penetration rates, when adding in younger users).

These applications offer one or more of a variety of services, including text messaging, voice and video calling, sending attachments, and video conferencing. There was an overall significant increase in usage during the pandemic, with some applications declining post-pandemic, while others continued to increase.[23]

These shifts highlight the low costs for users to download and use a large variety of applications, all of which are free or offer free capabilities. Overall, over 80% of those surveyed use at least one such application.

23 Skype was overtaken by other applications, but was already owned by Microsoft, whose Teams application experienced a significant rise in use in working environments.

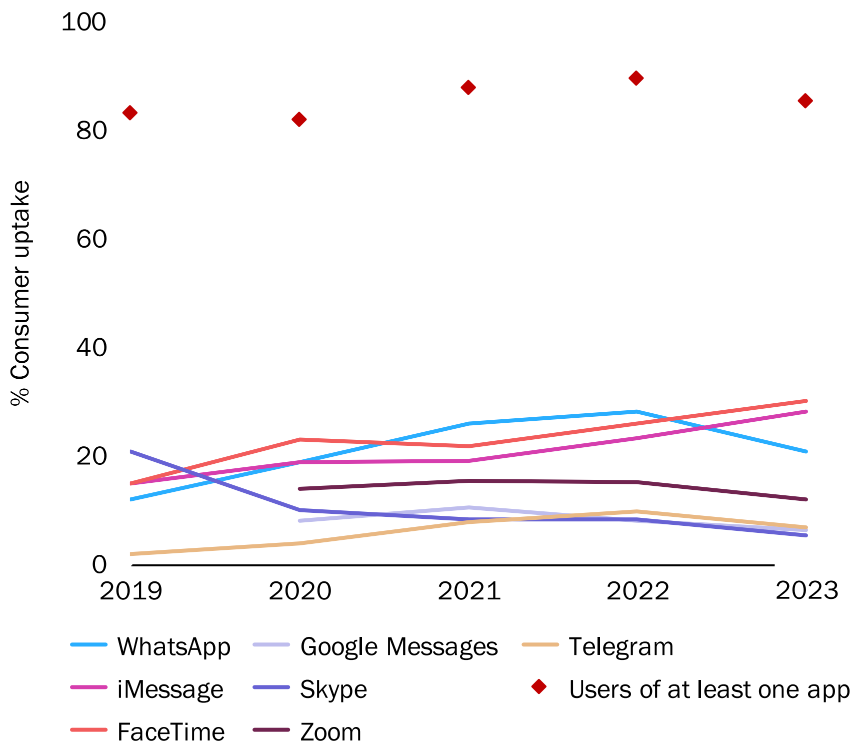

Figure 3.10: Consumer uptake of popular messaging and video conferencing applications in the U.S.[24] [Source: Analysys Mason Research, 2023]

Social media applications offer other means of ubiquitous communication, enabling users to share posts, pictures, and short videos to a select group or publicly, along with the possibility on some to send direct messages.

Figure 3.11 below was taken from the same consumer survey series as the previous figure, and again shows growth in some applications, corresponding in part to increased demand during the pandemic, while other social media applications were relatively flat or declined. Again, over 80% of those surveyed use at least one social media application. Thus, the provision of these services is both dynamic and popular.

24 Analysys Mason Research's Consumer Survey consists of 20-minute web-based questionnaires, comprising yes/no, multiple choice, text, and numerical answer questions; for the U.S., there were 1000 respondents in 2019 and 2020, increasing to 2000 in 2021, 2022 and 2023.

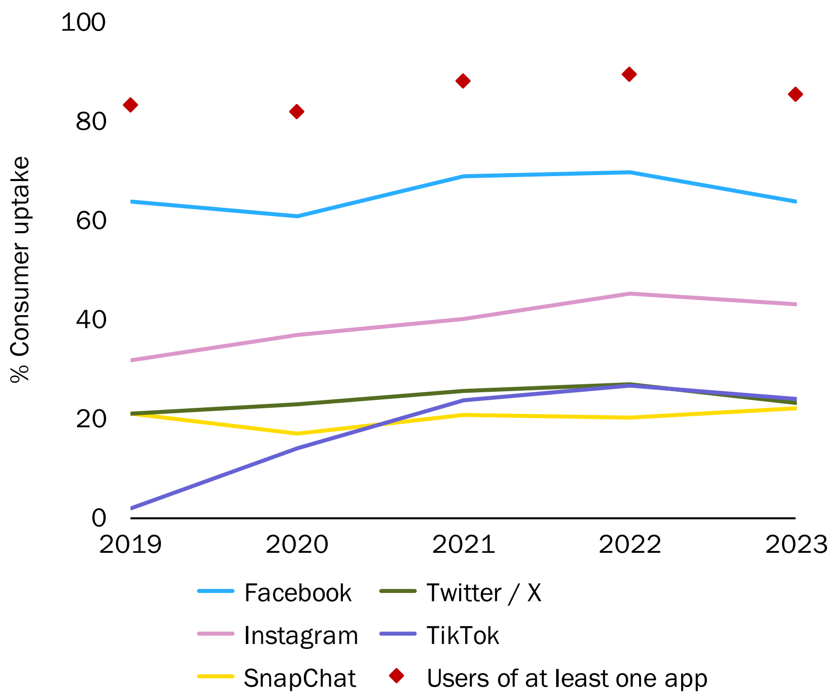

Figure 3.11: Consumer uptake of popular social media applications in the U.S. [Source: Analysys Mason Research, 2023]

Finally, another OTT service that has grown significantly in the last decade based on the internet design principles is streaming video. The number of users saw a significant bump upwards in 2020 reflecting the Covid-19 pandemic and the increased need for entertainment at home during periods of lockdowns and social distancing.

The chart below reflects the number of providers, and the shifts in adoption of different packages, including multiple subscriptions in individual households. Post-pandemic, there has been a small decline in 2023, for reasons including the ability and desire to return to pre-pandemic pastimes not involving streaming video.

We note that given the number of subscribers and volume of traffic generated by streaming video, these providers all rely on CDNs – their own or third party – to deliver content to end users. DNS and routing functions can allow ISPs to steer requests to these caches in an efficient way, with various levels of coordination.[25]

25 Where caches are embedded into ISPs' networks, the level of coordination is high; where content providers and CDNs advertise caches through DNS or IP routes, the ISP typically controls how the traffic is routed, and can benefit from advertised caches with little or no coordination with content providers.

Figure 3.12: OTT video subscriptions in the U.S. [Source: Analysys Mason Research, 2023]

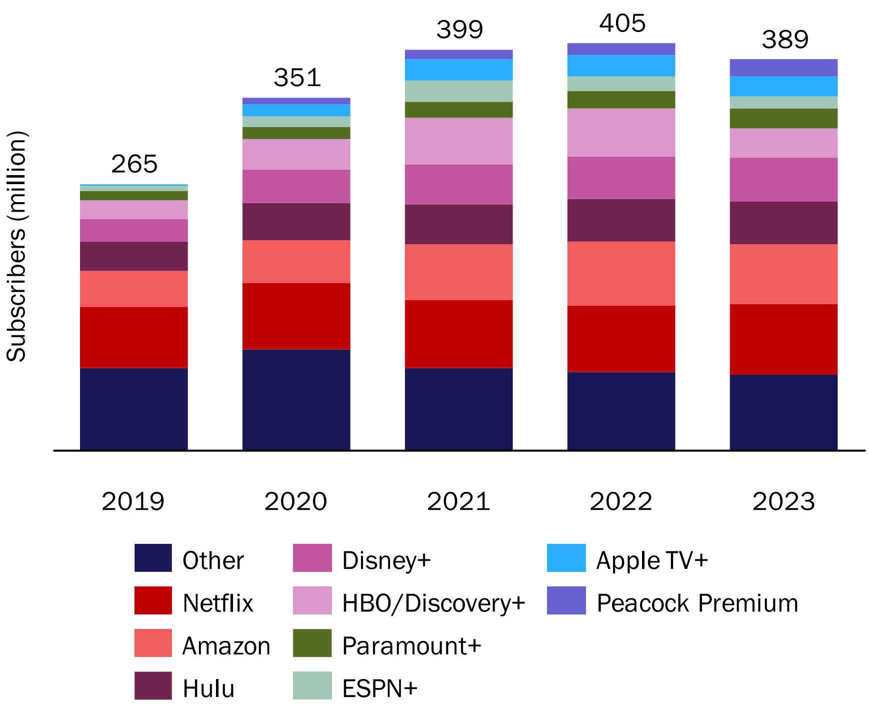

These OTT services compete with services provided by broadband ISPs, notably voice and video services. We are not aware of any efforts to block or throttle any OTT services since the FCC returned broadband service to Title I, and indeed the evidence above shows that these services have grown and thrived, under the internet design principles, with both ISPs and content providers investing in digital infrastructure to support this growth.

Innovation is extending beyond OTT services

Innovation and adoption are not just taking place with regard to OTT services, but also with respect to the devices used to access those services, which are increasingly 'smart' to embed internet connectivity and access to chosen OTT services. For instance, smart TVs embed access to streaming services, which can also be accessed through streaming devices. Smart speakers provide access to the internet and can be used for media services – with speakers and sometimes screens. Smartphones and tablets are also used for accessing the range of OTT services described above.

A significant proportion of households and individuals have adopted smart technology. The ownership of smart TVs reached 76% of households in a 2022 survey conducted by the Statista research department,[26] while 23% of U.S. households can stream video with devices such as the Amazon Fire TV stick and Google Chromecast. Smart speakers with integrated virtual assistants are owned by 32% of sampled households. Further, Analysys Mason's research indicates a substantial

26 Statista Research Department, "Which smart home devices does your household own?," August 2023

increase in U.S. smartphone penetration, rising from 73.58% in 2015 to 91.24% in 2022.[27] Tablet ownership has also experienced growth, with Euromonitor reporting that the percentage of households in the U.S. owning tablets increased from 57.4% in 2015 to 65.0% in 2022.[28] This data underscores the increasing prevalence of smart and mobile technology in households, and the principle that any internet-enabled device can be paired with any broadband access service to use any relevant OTT service.

The openness of the internet and ability to innovate has also recently been shown through the uptake of ChatGPT, which was the fastest application on record to reach 100 million monthly users, in just two months. By comparison, Facebook took over four years to achieve this milestone, and Twitter took over five years to do so.[29] Today, ChatGPT has 100 million users per week.

Evidence of the increased use of end-to-end encryption

The increased adoption of OTT communications services shown above, made possible due to layering, is evidence of the end-to-end principle. End users can download and use applications for messaging, making voice and video calls, and for work conference calls, many relying on end-to-end encryption to secure the transmission.

Figure 3.13 below shows the percentage of web pages used with the Chrome browser that are encrypted using the HTTPS protocol (over Windows). The percentage has been growing steadily, and stood at 97% in 2022, effectively approaching ubiquity. This type of encryption is implemented by the websites, using the HTTPS standard, and is effectively transparent for users. Further reflecting the openness of the internet, HTTPS was developed first by Netscape for its browser in 1994, and then adopted as an open standard by the Internet Engineering Task Force (IETF) in 2000.

27 Analysys Mason Research, 2023.

28 Euromonitor, Possession of Tablet: Euromonitor International from national statistics, 2023.

29 https://www.theverge.com/2023/11/6/23948386/chatgpt-active-user-count-openai-developer-conference

Figure 3.13: Percentage of pages loaded over HTTPS in Chrome in the U.S. over Windows [Source: Google, 2023]

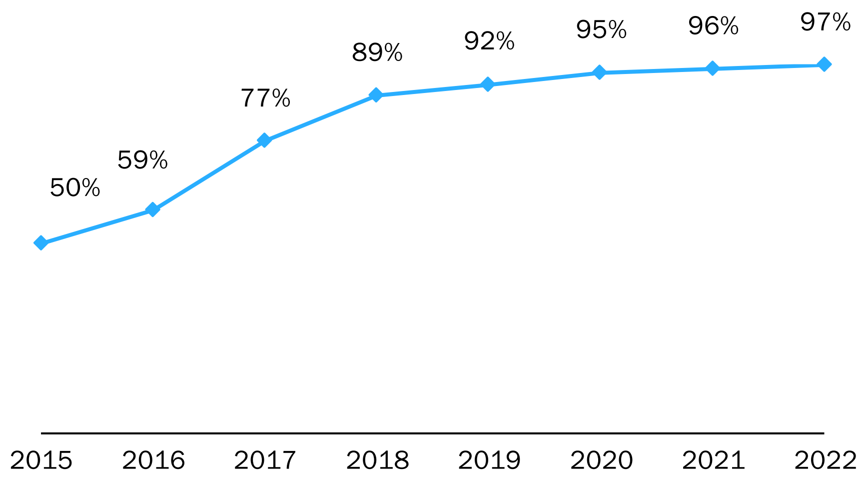

Furthermore, the use of VPN applications, which can be installed by individuals for their personal use or for accessing corporate websites, has likewise been increasing. A recent survey by NordVPN estimated that 33% of internet users had used a VPN in 2023, as shown in Figure 3.14. This use may be sporadic, to protect specific use cases or achieve specific goals, and many users may not use a VPN everyday.[30] The use of the VPN masks the IP address used, which can be used to indicate location, and also encrypts the data that is sent and received over the connection. While a common view of the use of the VPN to mask the IP address is to access content that would otherwise be blocked due to the location of the user, more broadly it provides privacy from governments, content providers, and the ISP.

Figure 3.14: Percentage of end users using a VPN in the U.S. [Source: NordVPN, 2023]

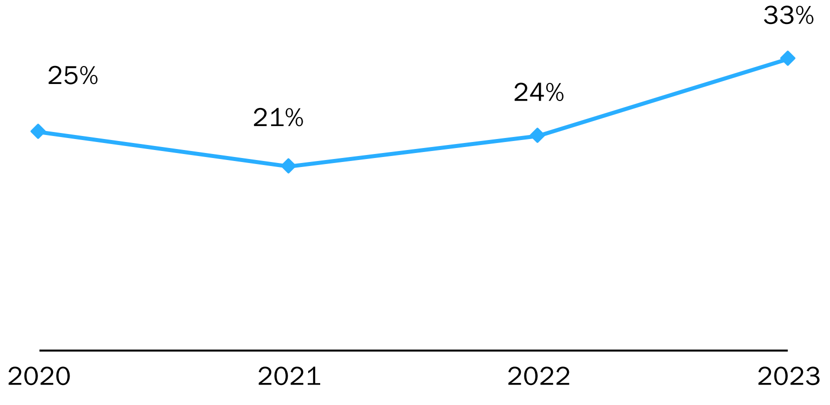

30 Martynas Klimas, "VPN statistics: users, markets, & legality," Surfshark blog, 8 February 2023.

Finally, increasingly messaging applications are automatically encrypted from end to end. This holds not just for messages, but also for any attachments that are sent, such as photos, and for voice and video calls for applications such as WhatsApp that enable those functions. Automatic encryption is true for all of the messaging applications above in Figure 3.10, as well as other applications that are more niche today, including Signal and Viber, which were designed around privacy.

***

Throughout the growth of the internet in the U.S., and notably in the years since the FCC rolled back Title II authority over broadband providers, the design principles of the internet have worked effectively. ISPs invested in new generations of fixed and mobile broadband infrastructure, reaching new users while existing users upgraded, and benefited from higher speeds at the same time as usage went up. Users are benefiting from a wide variety of OTT services, particularly during the pandemic, often through infrastructure investment from content providers to deliver content closer to users. Content providers and ISPs are working together, to negotiate interconnection arrangements in an increasing number of locations, and often coordinating to ensure efficient routing and access to CDN caches. All of this investment, innovation, and cooperation demonstrates that the absence of specific regulatory oversight has not stalled growth in supply or demand for online services.

4 Conclusions

This report explains how the internet has been thriving over the years since 2017, in line with its overall development over the past five decades, primarily outside the purview of the FCC's oversight. This period is significant for two reasons – first, it coincides with the period since the Title II classification was reversed, and second, it corresponds to the period during the Covid-19 pandemic when the internet in the U.S. was resilient to the substantially increased demands resulting from lockdowns and social distancing. In general, during this period, there has been significant investment, by both ISPs and content providers, with corresponding increases in broadband adoption, upgrades to new technologies, and increased usage.

Three fundamental design principles of the internet that continue to shape its development – layering, network of networks, and end to end – have been effective in practice, including in the years since 2017.

Layering. The provision of applications is separated from the provision of the underlying network by the layering principle. This principle enables applications to be introduced and used independently of the underlying network. The result has been a plethora of applications, including OTT services offering voice, texts, and streaming video similar to traditional communications services. The evidence shows that the uptake of these OTT services has been steadily growing, along with their usage, with video conferencing becoming particularly popular during the pandemic. Conversely, there is no evidence that there has been any interference with the operation of those OTT services since Title II re-classification was overturned in 2017.

Network of networks. The internet is made of thousands of independent networks based on the network-of-networks principle. These networks cooperate to negotiate interconnection arrangements to enable the exchange of traffic between networks. While these arrangements have evolved to address changes in applications and business models, they have done so free of regulation. One significant evolution has been the growth of IXPs to enable efficient public peering, and interconnection locations to enable private peering, sometimes involving paid peering as an alternative to the historical norm of settlement-free peering. At the same time, content providers are building out their networks to deliver content closer to the networks of ISPs and thus closer to the end users, all pursuant to negotiated arrangements. There is no evidence that any provider has acted as a gatekeeper, to force any interconnection conditions that require regulation to avoid or undo.

End to end. Finally, based on the end-to-end principle, users are free to choose and use any available applications, with any internet-enabled device, independent of the network provider. This principle has driven the adoption and use of applications including the OTT services described above. Further, under this principle, data can be encrypted, by websites using https, by application providers encrypting communications between users, or by users themselves installing VPN applications. The result is that no stakeholder, including network providers, has access to any data traffic that has been encrypted during transmission across the internet. The evidence shows that the use of encryption has been increasing over the past years, and thus privacy regulations aimed solely at ISPs would not recognize how market conditions and technology have evolved, and could prove at best ineffective for end users and at worst detrimental to competition.

We have seen no evidence that ISPs are engaging in blocking or throttling despite the absence of Title II regulation. As described in this report, data and outcomes clearly indicates that supply and demand for online services have kept growing strongly in the absence of such regulation: ISPs and content providers appear to work together well. Thus, in the current market conditions, the evidence shows that the existing regulatory framework is delivering good outcomes for all parties, with no apparent market failure justifying Title II reclassification of broadband.