Internet Interconnection and Infrastructure: On the Debate of Infrastructure Cost Sharing

A white paper by eco – Association of the Internet Industry

eco
ASSOCIATION OF THE INTERNET INDUSTRY

Content

1. Introduction

From the time of its inception, IP Interconnection has been and remained unregulated. The interconnection agreements that are needed to create the online "network of networks" are based on voluntary negotiations between the interconnecting networks, driven by the mutual requirement of data exchange. As the Internet has grown, voluntary agreements between market participants have also evolved consequently.

The large telecom operators repeatedly claim that the current voluntary network access arrangements impair their ability to cover infrastructure costs. They argue for a mandated charging scheme that would contribute to future infrastructure investment. The central argument for possible network access charges is based on the underlying premise that content and service providers are responsible for the continued growth in data traffic volumes and that the continued growth in traffic is leading to significantly higher network costs.[1] In order to better assess this proposition and to contribute to the technical discussion, the current eco - White Paper deals first with the more recent developments of the IP Interconnection market and the observed market trends. In addition, the extent of traffic-sensitive network costs is examined and the question of how the network costs of the Internet access providers respond to the market development is investigated.

1 Axon (2022): Europe's internet ecosystem: socio-economic benefits of a fairer balance between tech giants and telecom operators.

2. Background

A debate is currently underway in Europe about whether and how the large Internet platforms should contribute to the costs of rolling out a gigabit-capable access network infrastructure in order to achieve the European Commission's policy objectives.[2] From the perspective of the large telecom operators, hereafter called Internet Access Providers (the IAPs)[3] are finding it increasingly difficult to negotiate fair terms with the large Content and Applications Service Providers (the CAPs).[4] This refers mainly to the major US technology companies, also known as "Big Tech", who operate in the areas of social media, messenger, streaming and cloud services. The CEOs of the large European telecom companies attribute this to the strong market position of Big Tech, their asymmetric bargaining power, and the lack of a level legal playing field.[5] For this reason, they are calling for a legislative intervention in order to establish a fair burden-share that they believe is in line with the European Commission's agenda to counter online power imbalances. To this end, a reference to the need for fair and adequate infrastructure costs contribution was included in a declaration on European digital rights and principles in early 2022.[6]

The European Commission has recently openly expressed its support for the participation of Big Tech in network cost.[7] This is opposed by Members of the European Parliament, who expressed criticism on the EU Commission's possible plan to introduce "access fees", which "pose serious risks to the Internet as we know it and are unlikely to solve the broadband deployment problem".[8]

Earlier, at the World Conference on International Telecommunications 2012 (WCIT 2012), ETNO[9] already proposed to introduce a charging mechanism under which the sender's network would pay for traffic volumes. At the time, BEREC[10] assessed this proposal and concluded that deviating from the current principles could cause significant harm to the Internet ecosystem. The reason for this was seen to be because IAPs could exploit their termination monopoly,

2 European Commission (2022): Europe's Digital Decade: digital targets for 2030.
3 Called also Internet Service Provider (ISP).
4 Axon (2022): Europe's internet ecosystem: socio-economic benefits of a fairer balance between tech giants and telecom operators.
5 ETNO (2021): Joint CEO Statement: Europe needs to translate its digital ambitions into concrete actions.
6 European Commission (2022): Commission puts forward declaration on digital rights and principles for everyone in the EU.
7 Reuters (2022): EU's Vestager assessing if tech giants should share telecoms network costs.
8 European Parliament (2022): (Letter) To the European Commission.
9 European Telecommunications Network Operators' Association.
10 The Body of European Regulators for Electronic Communications.

Fig. 1 Overview of different market players for IP Interconnection

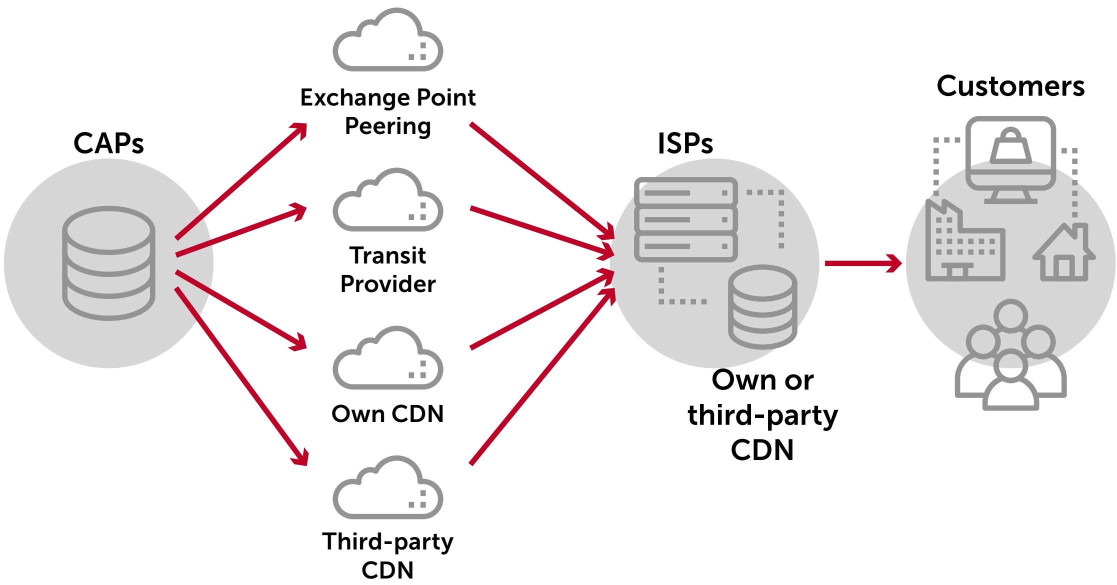

Source: eco - Association of the Internet Industry

as is the case with the traditional call termination monopoly,[11] and any payment disputes between IAPs and CAPs could lead to a loss of Internet connection quality for end users.[12] In its recent preliminary assessment, BEREC again stated in October 2022 that it had found no evidence to justify "direct compensation" as proposed by ETNO.[13] Whether BEREC will follow up and confirm its preliminary findings is currently not foreseeable. It should also be noted that, due to the unregulated nature of IP Interconnection markets, there is also a lack of available data across EU countries on the state of the various national markets.

3. IP – Interconnection as a critical part of the Internet ecosystem

IP Interconnection is vital both for IAPs and CAPs, enabling end users to connect to numerous so called "autonomous systems" (AS) around the world and to access online content and applications. In the IP Interconnection market, multiple players of the Internet ecosystem interact with each other. The figure 1 shows a simplified view of the relationships between the different market players for Internet interconnection.

The IAPs provide connectivity services and, to this end, Internet access to end users. Data traffic on the Internet is typically generated by end user requests, which are forwarded by IAPs to CAPs via the appropriate connection lines. Traffic flows are often asymmetric: while a content request consumes small amounts of data, the content delivered is much more extensive. End users purchase appropriate Internet access services from fixed and mobile providers in order to receive these deliveries of data.

Fixed and mobile markets in European Member States are not homogeneous, compounded by differences in the rollout and operation of fibre networks. On the supply side, some network operators have national or even international coverage (such as Vodafone, Telefónica, Orange, Deutsche Telekom), while other operators concentrate on regional or local markets. Some providers are vertically integrated, while others are only active in the retail or

11 BEREC (2012): BEREC's comments on the ETNO proposal for ITU/WCIT or similar initiatives along these lines, BoR (12) 120 rev.1.
12 BEREC (2014): BEREC Report on IP-Interconnection practices in the Context of Net Neutrality, BoR (17) 184.
13 BEREC (2022): BEREC preliminary assessment of the underlying assumptions of payments from large CAPs to IAPs.

wholesale markets. Certain fixed operators are also active in the mobile market and offer both sets of products. Different business models exist among network operators who are active in the same national market. In addition, players with core activities in other sectors (e.g., utilities) have recently entered the telecom markets. Furthermore, heterogeneity can also be observed on the demand side. End users have different needs and preferences for telecom products and services as well as in terms of their willingness to pay. This in turn leads to further differentiation of market entry strategies on the supply side.

The IAPs' infrastructure networks typically have an architecture consisting of access, backhaul and core networks. The access network is the "last mile" that connects end users to an IAP's network infrastructure. The backhaul network is the aggregation network that connects an IAP's access network to its core network. The core network comprises high-capacity links that carry traffic over long geographical distances.

For the internet to function, the major Tier 1 network operators,[14] who have established long-distance networks, are directly interconnected with other major network operators. To guarantee global connectivity, all Tier 1 network operators are interconnected with each other through peering arrangements. The Tier 2 network operators[15] maintain peering relationships with other similar network operators in their geographic areas, while at the same time acquiring transit services from Tier 1 operators for global connectivity. The smaller, local Tier 3 network operators rely entirely on transit services.

CAPs distribute their content and applications to end users via the Internet. These include video-on-demand (VoD), music-on-demand (MoD), social media, gaming, messaging, search, electronic retail and payments, and news. CAPs cater to a very wide range of segments and operate different services with varying business models. The leading CAPs include the large global technology companies, such as Alphabet (Google), Apple, Meta, Amazon, and Microsoft. Many of these big CAPs have strong market positions in different segments of the Internet. However, the Big Tech are not only content and application providers, but they are also major global players in different parts of the Internet ecosystem, who also invest in their own private and self-managed high-capacity backbones.[16]

In order for IAPs to provide Internet access and for CAPs to deliver content to end users, they may be directly interconnected or use Internet backbone operators to transmit traffic. In general, the IAPs and CAPs typically use an IP transit service provided by a Tier 1 IAP or backbone service provider for general Internet connectivity. If there is a significant amount of traffic between two networks, traffic will typically be exchanged directly – i.e., not via an intermediate backbone network or Tier 1, but via peering. Peering can be accomplished by a direct connection between networks – so-called private peering – or via an Internet Exchange Point (IXP) – so-called public peering. As a rule, the exchange of data traffic via peering takes place free of monetary charge (settlement-free), as it is seen as beneficial by both parties. The baseline for voluntary negotiation of interconnection agreements is determined by a reciprocal beneficial relationship that facilitates an equilibrium of bargaining power between CAPs and IAPs.

Content can be delivered more efficiently if it is hosted on servers or in caches as close to end users as possible, so that less of the shared public Internet is used to deliver the content. To protect traffic from potential bottlenecks, major delays, data loss and falsification, the CAPs install their caches close to end users or use the hosting services of third-party Content Delivery Networks (CDNs) such as Akamai, AWS or Limelight. This allows for data to be served from a CDN cache within the IAP's network or from an on-net cache directly connected to the IAP's network. The result is that data-intensive content such as videos or major software updates can be sent only once to each cache and served to users from there, reducing traffic on the core and backhaul network. In most instances, only the large CAPs provide their own caches, as developing and deploying them in multiple locations (100' to 1000' of caches deployed for mayor CAPs, i.e. Google Global Cache, Netflix Open Connect) involves high costs. The advantage of having proprietary caches is gaining greater control over traffic and scalability of business models. Using third-party CDNs can be more efficient for the CAPs and the IAPs that do not generate enough traffic for their own caches.

14 The Tier 1 category includes the following very large autonomous systems: Arelion (1299), AT&T (7018), Cogent (174), Deutsche Telekom (3320), GTT (3257), Liberty Global (6830), Lumen (3356), NTT (2914), Orange (5511), PCCW (3491), T-Mobile (1239), Tata (6453), Telecom Italia (6762), Telxius (12956), Verizon (701), Zayo (6461). These networks are defined by being "default free" and purchasing no transit from another network (peering only).

15 Tier 2 operators include the majority of large autonomous systems on the Internet, such as Vodafone, Comcast, Tele2, Swisscom.

16 cf. BEREC (2022): Draft BEREC Report on the Internet Ecosystem, Section 5.

Fig. 2 Annual changes in global internet traffic (CAGR) in percent

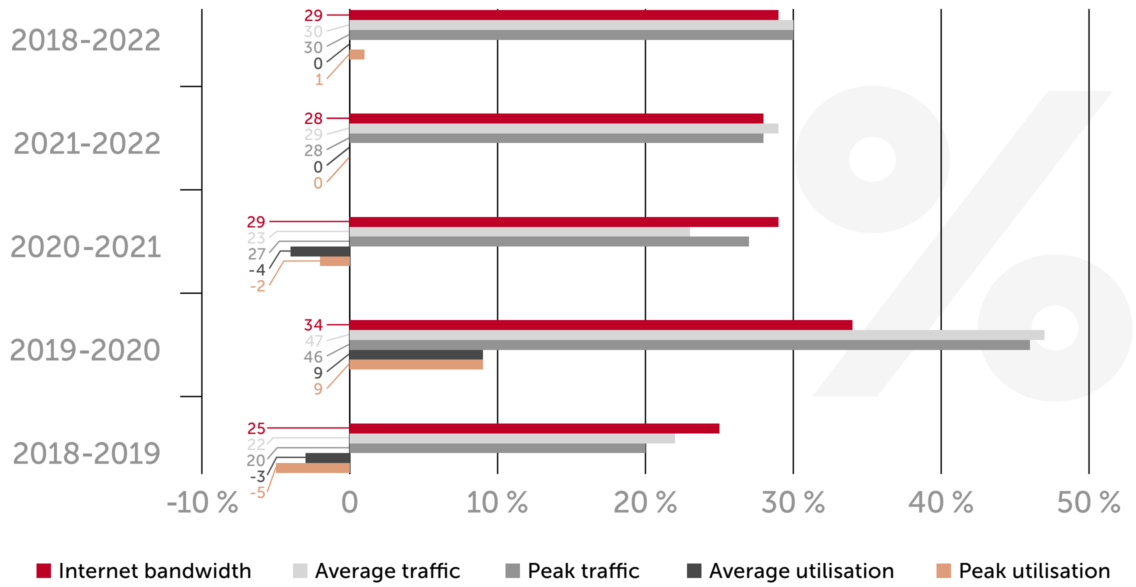

Notes: Data as of mid-year. Traffic refers to peak and average data rate used from the total Internet bandwidth.
Source: eco - Association of the Internet Industry, based on TeleGeography (2022), Global Internet Research Service Executive Summary.

Interim summary

The heterogeneous market relationships, both between and among network operators and content providers in Europe, add more complexity to the landscape of the Internet ecosystem. Market participants have to simultaneously cooperate and compete with the same network operators and service providers. The possible divergence of the ecosystem actors' respective interests may lead to different stances, or may even cause friction. Failure of negotiations between two interconnected players may lead to, inter alia, degradation of quality of service or disruption of network interconnection. IP Interconnection could also be used for anti-competitive discrimination with regard to the origin, destination or content of the transmitted information. Both of these potential issues are however addressed by the current EU net neutrality rules. Depending on the technical, commercial and competitive conditions applied, IP Interconnection can hence have various impacts on investment in the networks, quality of service, as well as innovation in services, content and applications.

4. Data traffic growth

Different types of content and services generate varying volumes of traffic and place varied demands on the networks. Some services, such as email or blogs, impose minimal load on the networks. Other services, such as video streaming, require more capacity, while others, such as games, are sensitive to delays. Livestreaming content can require both high capacity and minimal delay levels. Popular services such as Netflix, Amazon Prime or certain game titles like World of Warcraft or League of Legends can generate significant traffic on networks.

According to the latest reports from TeleGeography,[17] global Internet bandwidth increased by 28 percent in 2022 and currently stands at 997 Terabits per second. At the global level, Africa saw the fastest growth in international Internet bandwidth, with a compound annual growth rate (CAGR) of 44 percent between 2018 and 2022. Asia was close behind with a CAGR of 35 percent over the same period, followed by Europe with a CAGR of 27 percent.

However, after returning to a peak of 32 percent during the Covid-19 pandemic in 2019 - 2020 after over a decade of decline, the pace of growth at the global level is again slowing down. As TeleGeography shows, growth in international Internet bandwidth as well as Internet traffic remains more or less the same (Figure 2). Both average and peak international Internet traffic grew at a CAGR of 30 percent between 2018 and 2022, just above the 29 percent CAGR in bandwidth over the same period.

17 TeleGeography (2022): Global Internet Research Service Executive Summary.

Fig. 3 Peak bandwidth during 2 days traffic at DE-CIX Frankfurt, December 2022

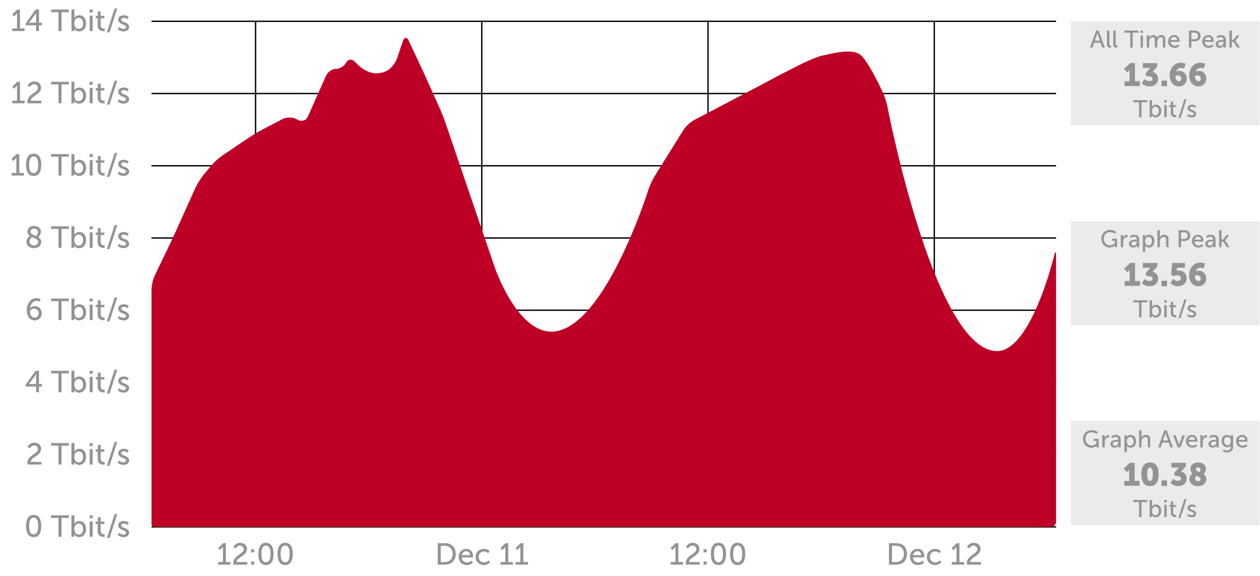

Source: DE-CIX Statistics

After the interim Covid-19 induced traffic growth, the global return to "normal" usage patterns is accompanied by a decline in average and peak traffic volumes. While average traffic growth dropped from 47 percent between 2019 - 2020 to 29 percent between 2021 - 2022, peak traffic growth dropped from 46 to 28 percent over the same period. It is notable that, between 2018 - 2022, the average and peak traffic utilisation have remained constant, at respective rates of around 25 percent and 44 percent. It should be noted that network cost is only determined by the cost of peak bandwidth usage plus reserves (installed bandwidth), typically reached between 19 and 21 pm each day, not by the aggregate traffic or amount of data exchanged. Figure 3 shows the peak bandwidths during 2-day traffic exchanged at DE-CIX Frankfurt.

Interim summary

Global Internet bandwidth has almost tripled since 2018. Yet the growth rate of both Internet bandwidth and data traffic has effectively come down significantly over the last decade. At the same time, average and peak utilisation rates of backbone networks remain unchanged.

5. Development of IP Interconnection markets

Traffic growth is an important factor, as network investments are driven by the peak bandwidth of traffic that can be delivered. Nevertheless, the extent to which IAPs and CAPs use transit, peering and caching services also affects the backbone costs of CAPs and IAPs.

5.1 Transit versus peering

For IP transit, one network operator provides global connectivity to another – typically much smaller - network operator and relays inbound and outbound traffic. This takes place irrespective of the origin or final destination of the traffic, unless there are restrictions imposed by an agreement between the parties – e.g., in relation to the geographic scope of the traffic. Peering, on the other hand, allows network operators to exchange traffic directly on a bilateral basis without any obligation to relay the traffic, with the latter being the case with transit. Under a peering agreement, an operator typically only grants access to its network, which means that the connection between operators can only be used for the traffic of their downstream networks and end users.

Peering and transit are not substitutes, as they do not allow the same level of control. The use of a transit service or entering into a peering agreement depends both on the bargaining power of the parties and on technological and economic aspects, the influencing factors of which include the relative costs of the different options and the quality of the service. The available data collected from IP

Interconnection shows that the shares of transit or peering agreements develop differently within the European countries.

In order to interconnect two networks, an interconnection link with a certain capacity is set up. The transit service is usually billed on the basis of this link's capacity in Gigabits per second. In doing so, the transit provider can stipulate a minimum bandwidth and a commitment period. Overall, the costs for transit are determined by capacity costs for leased lines as well as costs for switches, routers and ports. Over time, prices observed for transit services have steadily decreased due to increased network capacity, lower equipment costs and competitive pressures on a global average.[18] Nonetheless, IP transit is generally more expensive than peering when it comes to large volumes of data. Smaller market participants pay a relatively high price for their lower bandwidths; 100 Gbps in transit will only cost about 15 - 25 times as much as 1 Gbps and only 4 - 5 times as much as 10 Gbps.[19]

Peering relationships reduce the traffic load on transit services, which can often be expensive. Since there are line cost, equipment cost and administrative costs associated with the set-up of a direct peering arrangement, peering criteria can be put in place. These may cover the required characteristics of a potential peering partner, including geographic network coverage, number of interchange points, minimum capacity requirements and the symmetry of the exchanged traffic. In essence, these are criteria for determining whether both parties can derive roughly the same level of benefit from entering into a peering arrangement.

As peering has traditionally been a mutually beneficial relationship between two operators with compatible profiles (i.e., an access network benefits from peering with a hosting provider), this type of arrangement has generally been free of charge or settlement-free, apart from the cost of installing the switches and circuits required to connect the networks.

It should be noted that the vast majority of networks has adopted a so called "open" peering policy and have not established any peering criteria. They will enter into peering arrangements without prerogative[20] or even utilize "Multilateral Peering" as offered on all relevant IXPs through the use of route servers.

If the peering criteria are not fulfilled, interconnection must take place via transit providers despite the presence of the suitable networks at the same location. Without direct interconnection, the data may have to pass through several networks or over long distances and consequently suffer a high latency before reaching the end users. Although peering is of obvious interest, smaller market participants with weak bargaining power and a limited, regional footprint have no choice but to pay one or more transit provider(s) to connect their customers to global IP connectivity. Overall, the smaller IAPs and CAPs pay much more for transit per unit of data compared to major IAPs and CAPs who engage in peering.[21]

Tariff setting in the IP Interconnection market relies on the competitive situation on both sides of the market. If the IAPs want to strengthen their position on the inbound traffic side, they can offer benefits to the parties on this inbound traffic side – e.g., lower charges or no charges for interconnection. Such a competitive constellation can lead to an increase in settlement-free peering. For example, in the UK, on-net caches and private peering were used to distribute most traffic from the major IAPs and CAPs in 2019 - 2021, while transit traffic volumes were significantly lower by a wide margin.[22]

On the other hand, if competition is relatively strong in the Internet access markets – i.e., for access to end users – IAPs may choose to increase peering revenues from inbound traffic and avail of them to improve quality or lower retail prices in order to strengthen their position in the Internet access markets. Such a competitive constellation can lead to the share of transit services as well as paid peering remaining at a relatively high level. Looking at the current data on IP Interconnection of the major IAPs in France, the ratio of transit volume to peering traffic volume is still 48 to 52 percent, despite the simultaneous decline in transit services.[23] At the same time, the volume share of paid peering more than doubled in the 2012 - 2021 time period.[24]

In a further instance, if an IAP has a bargaining power due to its existing call termination monopoly in the field of voice telephony, it can also opt for a restrictive peering policy by only allowing the forwarding of traffic to end users via IP transit. There are rare cases where an integrated Tier 1 telecom operator only conducts peering with the same size operators and follows a strict traffic

18 TeleGeography (2022): Price Erosion Remains the Universal Norm.
19 For instance, in Frankfurt, Amsterdam and London (Interview with DE-CIX).
20 Peering Policies in Peering DB: https://www.peeringdb.com/ for established peering policies.

21 Ibid, p. 19.
22 OFCOM (2022): Net neutrality review – Consultation, p. 18.
23 ARCEP (2022): The state of internet in France, p. 41.
24 Ibid, p. 42.

Fig. 4 Annual price declines of 10, 100, 400 Gigabit optical transceivers

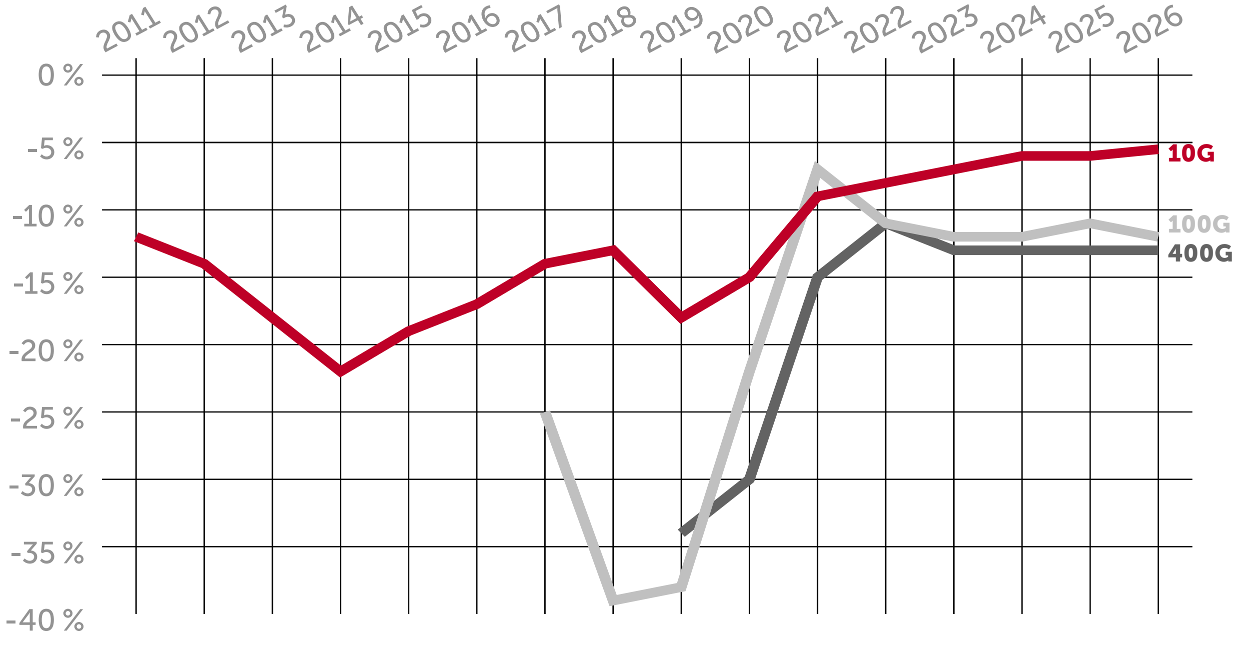

Source: eco - Association of the Internet Industry based on LightCounting (2021), Ethernet market on track for record growth this year.

ratio in its peering agreements.[25] In the case of an asymmetrical traffic exchanges which exceeds a certain predefined ratio, the interconnection partners, including large CAPs and cloud service providers, must switch to purchasing transit services. In so far as no on-net CDNs are allowed in the IAP's infrastructure as well, the CAPs rely on the operation not one, but several transition points into the IAP's network. The restrictive peering policy with limitation to IP transit is economically equivalent to paid peering. As a peering criterion, restrictive traffic conditions reflect the fact that the relative cost for the exchange of data traffic is higher for the operator if the traffic is highly asymmetric in nature. Nevertheless, this alone does not allow conclusions to be drawn regarding economic value of the exchanged data traffic. Although the exchanged traffic is often asymmetric, the unit cost of sending or receiving the traffic is the same as the unit bandwidth cost is the same due to the symmetrical nature of the underlying fibre optic technology. Above all, the IAPs need content to attract end users, while the CAPs need the IAPs to reach end users; despite being of mutual interest such interaction can sometimes only be achieved through commercial agreements.

25 WIK (2022): Competitive conditions on transit and peering markets, p. 43.

5.2 Public peering versus private peering

Meanwhile, IP Interconnection is increasingly taking place in more and more locations. CAPs and CDN providers are developing more delivery networks to bring their services closer to end users – i.e., in IXPs and private peering premises or as caches within IAP networks.

Peering is typically concentrated in several regional node points. While new IXPs are emerging in different geographic regions, established IXPs are expanding their presence, both within and across regions. For example, the number of operating IXPs in Europe increased by 87.5 percent between 2010 to 2020, while the average aggregated peak traffic volume increased from 4,140 Gbps to 45,325 Gbps over the same period.[26] In the Euro IX region, the number of ASNs (Autonomous System Numbers) active at a single exchange reached around 1,043 at DE-CIX in Frankfurt, followed by over 870 ASNs at AMS-IX in Amsterdam.[27]

With the growth in IXPs' traffic volumes, CAPs and IAPs are increasingly engaging in private peering. Private peering optimizes capacity allocation, which scales business models and upgrades connections in line with growing data traffic and broadband requirements.

Since 2018, the number of public and private peering locations of global CAPs have increased in turn by 80 and 35 percent, respectively.[28] With the increasing offloading to private peering, traffic volumes in private peering are growing much faster than in public

26 Within the Euro-IX membership, Euro-IX (2020): Internet Exchange Points 2020 Report.
27 The bulk of IXPs in the Euro IX region have 21 to 50 ASNs connected, see Footnote 21.
28 CAPs such as Google, Meta, Microsoft, Amazon, Yahoo, Netflix, Apple, eBay, Tencent, Baidu. Analysis Mason (2022): The impact of tech companies' network investment on the economics of broadband IAPs, P. 25, based on PeeringDB.

peering. Current market analysis in the UK, for instance, shows that for major IAPs approximately 40 percent of inbound traffic is delivered via private peering and about four percent via public peering.[29] The largest degree of traffic, amounting to over 50 percent, constitutes on-net caches. The remaining traffic involves IP transit of seven percent, which in turn is higher than public peering. Private peering often takes place in the same data centres used by the IXPs. The IXPs with access to multiple data centres at different locations are increasingly facilitating private peering across their facilities and IT infrastructure.[30]

In the case of public peering, depending on the capacity requirements, not only routers but several ports are used by both sets of partners. The move from 10 via 100 to 400 or even 800 Gigabit bandwidth capacity port provides a clear migration path to upgrade to higher speeds. An important factor for the adoption, for instance, of 10, 100 or 400 Gigabit ports is that the costs of appropriate optical transceivers have decreased over the time due to the increasing competition and new suppliers entering the market. Figure 4 shows annual price declines of the according optical transceivers.

While both sets of public peering partners bear approximately the same costs, the costs for private peering can be very different.[31] The costs of private peering are affected by the projected added value of the connection for both partners. For those large IAPs who exchange traffic with a few large peers, private peering can be cost-effective in addition to scalability, reliability and security of dedicated connections.[32] Consequently, mostly larger market participants enter into private peering relationships. At the same time, a hybrid approach, i.e., a combination of public and private peering occurs increasingly, where IAPs migrate some part of peering sessions to private peering as the volume of the exchanged traffic increases.

Private peering requires additional cross-connects and at least one additional interface card for each peering session. Conversely, public peering can aggregate a large number of relatively small peering sessions with no incremental cost. That is why the smaller market participants as well as CAPs rely more often on public peering at IXPs for their connectivity. As such, public peering remains important for the innovation and competitiveness of small market participants.

5.3 Settlement-free peering versus paid peering

In certain types of peering contracts, if the traffic ratio exceeds the value specified in the peering criteria, the peering partners can ask to be remunerated for the excess traffic exchanged in the peering relationship. In these asymmetric contracts, which are almost exclusively seen with IAPs, the choice between settlement-free peering and paid peering depends on the transaction costs and bargaining power of the peering networks. According to ARCEP, there is a wide price spectrum for private peering, ranging from between €0.25 and several Euros per month for each Mbit/s[33]. The reported value is significantly higher than market prices for global IP connectivity in all cases.[34] The bargaining power of a network during negotiations is determined by the number of eyeballs provided (IAPs), direct connectivity with other networks, the ratios of the peering networks' traffic volumes as well as the desirability of content offered (CAPs).

According to a recent survey by Packet Clearing House,[35] the vast majority of global peering is based on "handshake" arrangements without formal agreements or written documents. In the last ten years, the share of such settlement-free agreements increased from 99.51 to 99.998 percent. This large number of informal agreements on commonly understood and accepted terms is often arrived at during regional or global peering forums. Over 99.9996 percent of all peering agreements contain symmetric terms and conditions for the peering partners; the parties simply exchange routes to customer networks with each other without any charges or other requirements. In the currently very small number of paid peering agreements with asymmetric terms[36], one party is unilaterally obliged to comply with requirements imposed by the other party, often concerning the volume of traffic or number or geographical distribution of interconnection locations.

While there is a lack of Europe-wide systematic surveys of IP Interconnection markets, it can nevertheless be observed that paid peering activities are decreasing over time, both regarding to the number of such agreements as well as in terms of their traffic volume. According to the Packet Clearing House, the number of global paid peering agreements with asymmetric terms is trending downwards as the commonly understood terms of agreement continue to become more prevalent.[37]

29 Ofcom (2022): Net neutrality review – Consultation, p. 18.
30 e.g. DE-CIX in Germany, LINX in the UK.
31 There are no systematic records of the prices of private peering, cf. ACM (2021): Study into the Market for IP interconnections 2021, p. 19.
32 DrPeering International: The Internet Peering Playbook, Chapter 7.
33 ARCEP (2022): The state of internet in France, p. 43.
34 Interview with DE-CIX.
35 Packet Clearing House (2021): 2021 Survey of Internet Carrier Interconnection Agreements.
36 Ibid, only 57 of 15.105.102 agreements reflect paid peering.
37 Ibid.

In contrast, according to ARCEP's[38] ongoing surveys of the French IP Interconnection market, the share of paid peering traffic in inbound traffic from major IAPs in France increased from 20 percent to 53 percent over the period 2012 - 2019. This change was primarily due to the increase in private peering traffic, a significant proportion of which involved paid peering, particularly in the presence of substantial traffic asymmetries. Peering between companies of comparable size has remained free of charge.

Contrary to the trend observed for several years, the share of paid peering subsequently decreased in France from 53 percent at the end of 2019 to 47 percent at the end of 2020. ARCEP attributes this decline on the one hand to the increase in settlement-free peering and, on the other hand, to the transfer of paid peering traffic to on-net CDNs. The latter development in France is in line with the increasing provision of content closer to the requesting backbone in several other countries.

5.4 CDN and on-net CDNs

In order to improve quality of service by bringing content as close as possible to end users, providers of data-intensive content like videos often use CDNs, which replace long-distance transport by storing data locally on cache servers. The network developments of CDNs have reduced the need for transit and paid peering, as content can be delivered over the same IXPs, thereby reducing transport costs for the receiving IAP.

Today, a number of CAPs are making significant investments in network infrastructure. This includes setting up caches and Points of Presence (PoPs) as well as investing in the rollout of submarine cables between continents and countries.[39] While many CAPs rely on third-party CDN services, the Big Tech have built their own CDNs, which they use to serve their own and third-party content.

A significant trend that has developed in recent years is the emergence of internal and network collocated CDNs (on-net CDNs). These servers, which are installed within the IAPs' networks, are managed by those companies that own them (CAPs, IAPs or third-party CDN providers). To improve the quality of service and significantly lower the overall cost of data delivery by getting as close as possible to the end user, CAPs partner with IAPs to have their content hosted on cache servers on the operators' network. The best-known examples of on-net CDNs are the Netflix OCA (Open Connect Appliance), Google Global Cache (GGC) and Amazon Cloud-Front. By bringing content closer to end users, network conditions can be measured more accurately, and transmission rates can be dynamically adjusted in order to increase the quality of service, optimize the user experience and alleviate potential network congestion. The use of an embedded on-net CDN also allows video content to be uploaded outside of busy peak hours.

The current data on IP Interconnection in the UK shows that traffic delivered via on-net caches from the major IAPs has grown to 50 percent of all inbound traffic to end user.[40] However, disaggregated data from ARCEP shows that the proportion of growing on-net traffic varies greatly even within the major IAPs, ranging from 1 to 40 percent of incoming traffic. Furthermore, the ratio of inbound to outbound traffic can also vary from 1:5 to 1:11.[41] ARCEP notes that, in most cases, on-net CDNs are free of charge. However, they can be charged for in the case of broader paid peering agreements.[42] With a respect to global traffic, on-net CDNs accounted for around 90 percent of global traffic already in 2018.[43]

In addition, video streaming providers are also developing other innovative solutions for the multi-CDN. These developments are leading to a relatively dynamic content delivery environments and new initiatives are emerging such as open caching from Streaming Video Alliance (SVA).[44] Open caching is a specification that allows caches on IAPs' networks to receive and store content from different providers in accordance to a harmonised technical approach, and to also streamline cache management. This should enable more content providers to benefit from the advantages of online caches and reduce the number of on-net caches operated by IAPs. Ultimately, this means that CAPs without their own CDNs or commercial CDNs can deliver their content to IAPs, while IAPs can develop their own CDN services based on the specifications of open caching as a business line.[45]

38 Since 2019, ARCEP has been the only European regulatory authority to collect and publish data on IP - Interconnection in France, see ARCEP (2022): The state of internet in France.
39 For investment activities of CAPs, see Analysis Mason (2022): The impact of tech companies' network investment on the economics of broadband IAPs.
40 Ofcom (2022): Net neutrality review – Consultation, p. 18.
41 ARCEP (2021): Barometer of data interconnection in France, p. 16.
42 ARCEP (2022): The state of internet in France, p. 43.
43 Craig Labovitz (2019): Internet Traffic 2009-2019.
44 SVTA (2021): The State of Open Caching: Specifications, Coding, and Implementation.
45 For example, in order to deliver the Disney+ service content, Disney cooperates with Verizon through its Open-Caching developed on the basis of the SVA specifications; see Fierce Video (2021): Verizon, Disney begin Fios open caching trial for Disney+ streamers.

Interim summary

The respective competitive conditions in the inbound and outbound data traffic markets influence the implementation of transit, peering and caching services and their respective commercial pricing. Peering and transit are not interchangeable. The development of transit and peering connections differ in various commercial IP Interconnection markets within European countries. Overall, the traffic load of transit services is declining due to the increase in peering connections with or without peering criteria that require a level of similar bargaining power. The smaller CAPs therefore rely on transit services, with a higher cost per data unit than private peering, despite the overall decrease in total transit costs.

In contrast, for the major ISPs, private peering can be cost-effective compared to private peering, despite the overall decrease in total public peering costs. While the locations and traffic volumes from both public peering and private peering are on the rise, the rate of growth of traffic in private peering is increasing much faster. This is due to increased private peering between the major CAPs and ISPs, which in turn often takes place in the data centres also used by IXPs. Increasingly, there is a shift of capacity at the major IXPs locations from public peering to private peering. At the same time, public peering continues to be an indispensable access to IP Interconnection for smaller market participants and to ensure global connectivity.

The overwhelming majority of peering relationships are based on settlement-free agreements with symmetrical conditions for the peering partners. Due to the generally accepted contractual conditions, even the small number of paid peering activities with asymmetrical conditions are declining in terms of number and traffic volumes. The further increase of settlement-free peering as well as the decrease of paid peering shows that there is now a greater equilibrium in the bargaining power of the major IAPs and CAPs.

The development for the growing transition of transit and peering traffic to on-net CDNs has been generated by the emergence of new content delivery business models of CDNs. Given that most CDNs are using dynamic performance adjustments to deliver their content, the nature of data traffic is changing with the surge of on-net CDNs. This leads to better load distribution and network load, in addition to greatly optimising the user experience. The development of on-net CDNs contributes to the reduction of IAP costs of infrastructure investments for end-to-end connections and to the IAP provision of end-to-end services.

6. Traffic-sensitive costs of the IAP - networks

As discussed above, overall, the costs of IP Interconnection – which consist of interconnection equipment, connection to IXPs, and the costs of acquiring transit and peering services – are relatively low for IAPs. In contrast, the costs of IAPs for the provision of services to end users and for the deployment and operation of access and backbone networks are much higher. The investment in network infrastructure is determined mainly by the peak bandwidth that needs to be transmitted. The traffic-sensitive costs of access and backbone networks are considered below.

6.1 Fixed networks

The vast majority of data traffic is transmitted via wired networks. The largest part of the network costs is incurred by the access network. These costs are generally independent of the volume of traffic and depend instead on the number of subscribers that can be connected to the network and on the applied technology. Overall, fibre networks have lower operating costs than copper-based networks. According to the calculations, the operating costs for FTTH networks are 50 to 63 percent lower than the costs for operating copper-based DSL networks or HFC networks.[46] While in conventional copper-based networks, the network performance drops significantly with increasing distance, the network performance in fibre access networks is much less distance-dependent. Thanks to this attribute of optical fibres, network operators can reach their end users with fewer nodes than in copper-based networks, due to the fact that the performance is maintained over longer distances. For this reason, network operators are able to reduce the number and location of network nodes, in particular by decreasing the number of edge or local nodes to which end users are connected. The remaining edge nodes can serve a greater number of connections, increasing network efficiency and reducing network costs. Consequently, network operators' costs in access networks are not traffic-sensitive but subscriber-dependent,[47] with these being covered by end user charges, often with flat rates. Cost reductions in access networks can help offset any costs in backbone networks that may arise from increased traffic.

46 Fiber Broadband Association (2020): Operational Expenses for All-Fiber Networks are Far Lower Than for Other Access Networks.
47 cf. Frontier Economics (2022): Estimating OTT traffic-related costs on European Telecommunications networks, p. 6.

Fig. 5 Change in nominal expenditure levels (capex and opex) for four major fixed ISPs in the UK, 2017 – 2026

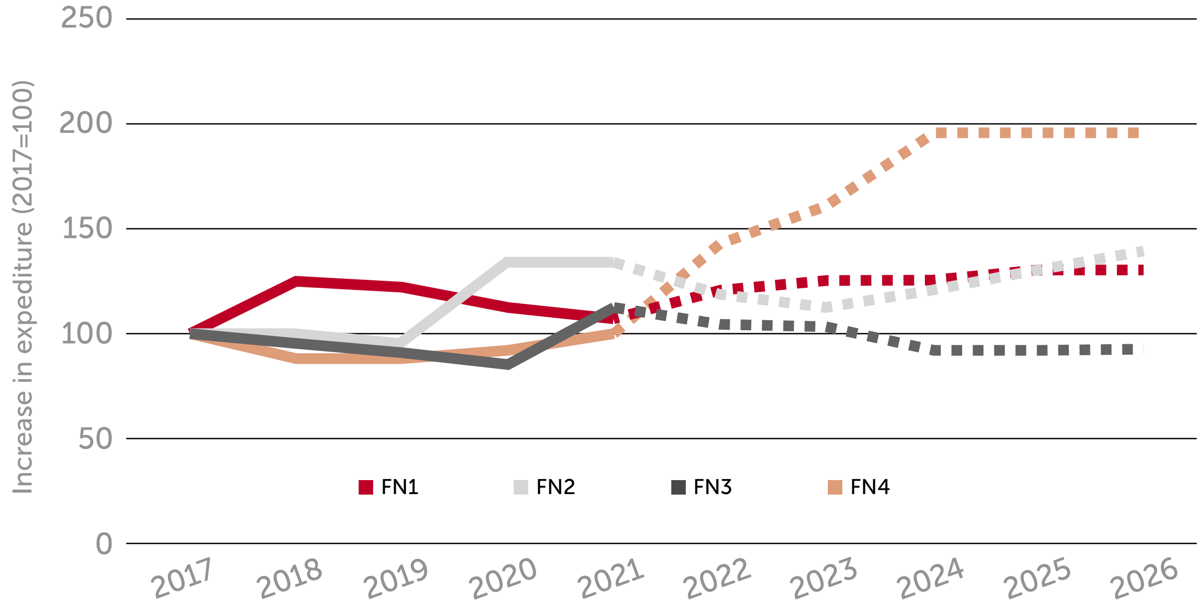

Notes: The four major fixed ISPs are BT Group, Sky, Virgin Media O2 and TalkTalk.
Source: OFCOM (2022), Net neutrality review - Consultation Annexes 5 to 10, Annex 8, p. 49.

The backbone networks are more traffic-sensitive compared to the access networks. The backhaul and core network segments aggregate the traffic links around the access networks, and sufficient capacity must be provided in the links and nodes connecting different layers of the network. Further investment required for additional capacity in these network segments is driven by the need to dimension the available capacities to meet peak time demand. According to current estimates by Analysis Mason, traffic-sensitive costs in backbone networks account for 20 to 30 percent of total network costs,[48] as well as 10 to 15 percent measured in terms of revenue.[49] Total network costs, in turn, correspond to about 50 percent of revenue at the end user level.[50] As a result, changes in data traffic also have only a limited impact on total network costs. For example, Ofcom shows that the total expenditure of the major IAPs in the UK on backhaul and core networks, including capital expenditure and operational costs, has so far remained more or less constant.[51] This trend is also largely reflected in the projected total expenditure (Figure 5).

Although traffic volumes will continue to rise, backhaul and core network costs will either remain nearly constant or only increase slightly over the next five years, as predicted in Analysis Mason's global forecast[52] (Figure 6).

On the one hand, the relatively constant cost level of the backbone network is due to the decrease in unit investment costs as well as the more cost-efficient technological developments of the network equipment. On the other hand, the increasing network investments by CAPs and CDN operators in on-net caches that are embedded in the IAP networks contribute to the reduction of the backbone capacity of the IAPs and their network investment required for this purpose. Furthermore, CAPs and CDN operators regularly collaborate with IAPs for network planning purposes and to manage expected peak demand. This includes measures such as shifting traffic away from peak hours on the IAP networks, scheduling live sporting events, or providing off-peak video game content and software updates, all of which minimise the impact of downloads on the IAP networks.53 Furthermore, when considering the technical efficiency of network architectures over time, the cost of fibre-based backbone networks can account for 6.5 percent of revenue, compared to 13.5 percent of revenue in the case of copper-based networks.[54] As such, the ongoing migration from copper-based to fibre networks is likely to further reduce the share of traffic-sensitive backbone network costs in the total set of network costs.

48 Analysis Mason (2022): The impact of tech companies' network investment on the economics of broadband IAPs, p. 78. Total annual network costs are based on the assumed operating costs and capital costs, which are respectively 35 and 15 percent, measured against EBITDA margins of 30 percent.
49 Ibid, p. 35 et seq. The baseline scenario applied in the model reflects the historical average Internet throughput per connection during peak hours and a traffic growth per connection of 20 percent per year.
50 Ibid, p. 35 et seq.
51 OFCOM (2022): Net neutrality review - Consultation Annexes 5 to 10, Annex 8, p. 49.
52 Analysis Mason (2022): The impact of tech companies' network investment on the economics of broadband IAPs, p. 39.
53 OFCOM (2022): Net neutrality review - Consultation Annexes 5 to 10, Annex 8, p. 52 et seq.
54 Analysis Mason (2022): The impact of tech companies' network investment on the economics of broadband IAPs, p. 40.

Fig. 6

Evolution of modeled traffic and annualized core and backhaul costs in fixed networks for world regions, 2022 – 2026

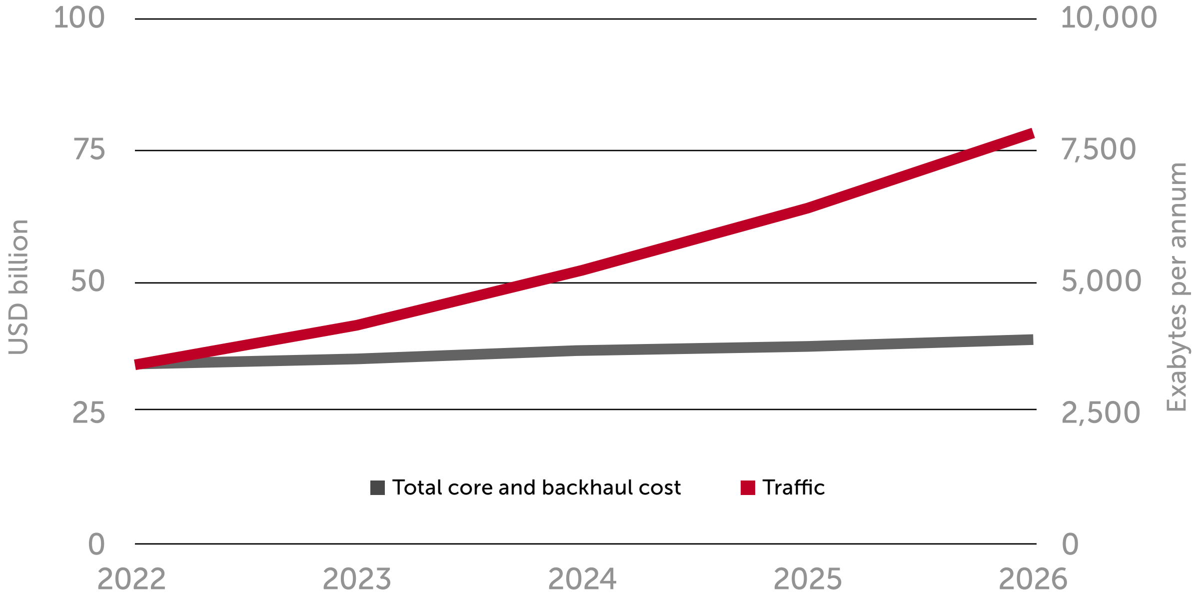

Notes: The regions covered are North America, Latin America, Europe, Middle East and Africa, Asia–Pacific excl. China.
Source: Analysis Mason (2022), The impact of tech companies' network investment on the economics of broadband IAPs, p. 39.

6.2 Mobile networks

Unlike wired networks, most of the network costs of mobile networks do not depend on the number of subscribers, but are influenced by the network components that are used, including infrastructure, active equipment and spectrums, and especially by data traffic. According to Ofcom's calculations, 50 to 75 percent of the radio access, aggregation and core network costs of mobile networks depend on the volume of traffic at busy peak times.

The main cost drivers of mobile networks are the frequency spectrum and spectrum efficiency, which are limited. In mobile networks, several connections share the network resources. Increasing traffic in mobile networks leads to the deployment of additional capacity, through the use of additional frequencies, more efficient technologies and new equipment, as well as the development of new mobile sites. The capital costs required to deploy a mobile network involve a one-off expenditure on the spectrum, ongoing investment in macro-network coverage and the upgrading of networks. Lease and transport costs, on the other hand, account for the largest share of operating costs, which can vary widely, both geographically and depending on the mobile network provider. While the costs for deploying new network coverage are not traffic-sensitive, the additional costs of upgrading networks are.[55] The cost per Gigabit of traffic varies over time and is dependent on application technology, spectrum, bandwidth and other factors.

The overarching point, however, is that traffic capacity for each technology evolution grows faster than investment required to deploy said evolution, reducing the cost of each additional deployed Gigabit.[56] The only relevant data available to date, which is provided by Ofcom, shows a relatively uniform development of network costs over time among the major mobile network operators in the UK (Figure 7). Similar to the fixed network, the projected total costs in mobile networks over the next five years should be at a stable level as in the past; the capital costs are anticipated to decrease, while the operating costs are projected to increase.[57]

The stable trend in mobile network costs is due, on the one hand, to recent developments in mobile technology and network sharing. Aspects that enable network operators to continuously improve spectrum efficiency and transmit the elevated level of data traffic include: the use of reallocated spectrum bands and the upgrade to new frequency spectrums, the introduction of multi-band antennas and network virtualisation, and the sharing of infrastructure and frequencies. On the other hand, the cost per Gigabit also decreases with area density, as more connections per unit area are accompanied by more connections and thus revenue per mobile site. Since densely populated and thus less traffic-dependent areas account for between half and three-quarters of the mobile network-based PoP in a country on average,[58] the resulting traffic can be partly compensated for by unused capacity, so that more traffic can be transmitted with little or no additional investment.

55 BEREC (2022): BEREC preliminary assessment of the underlying assumptions of payments from large CAPs to IAPs, p. 8.

56 Ericsson (2022): Understanding the Economics of 5G Deployments, p. 12. The cost per Gigabit includes both the operating costs and the capital costs of the mobile networks, in relation to the cumulative data traffic.
57 OFCOM (2022): Net neutrality review - Consultation Annexes 5 to 10, Annex 8, p. 50.
58 Analysis Mason (2022): The impact of tech companies' network investment on the economics of broadband IAPs, p. 35.

Fig. 7 Change in nominal expenditure levels (capex and opex) for four major mobile networks in the UK, 2017 – 2026

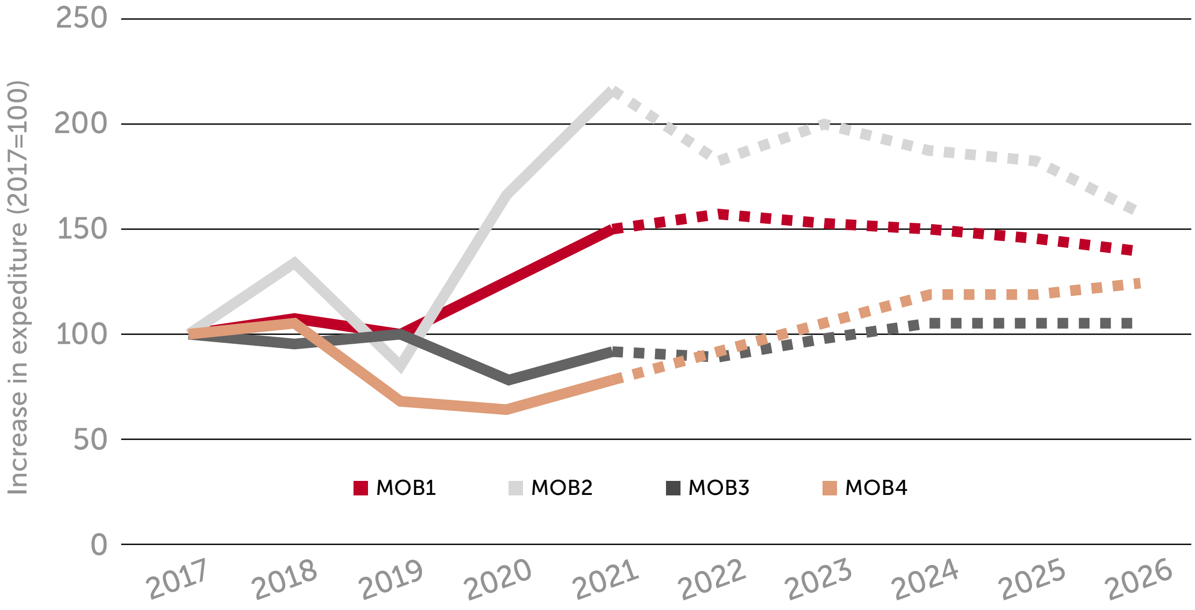

Notes: The four major mobile providers are UK – EE, Vodafone, Virgin Media O2 and Three Mobile.
Source: OFCOM (2022), Net neutrality review - Consultation Annexes 5 to 10, Annex 8, p. 50.

Last but not least, the falling costs of data transmission are reflected in end user prices of the mobile network operators. In contrast to the flat rate of fixed network broadband offers, the price of end user offers in mobile communications today is usually linked to the associated data volumes which they contain. However, a shift to flat rate tariffs can be observed, similar to the shift to flat rate observed in fixed networks in the early 2000s, and tariffs are readily available for most networks. This trend has accelerated significantly after the abolishment of zero-rating tariffs by the EU courts in 2022.

Interim summary

The significant increase in traffic in recent years can mainly be attributed to the growing demand of end users for broadband services and applications and thus for broadband connections. This traffic growth has not been accompanied by a corresponding increase in network costs, as traffic-related costs only account for a small part of network costs, and these also do not increase proportionally with traffic volume. It can be deduced that the projected total costs of IAPs in fixed and mobile networks will remain at a stable level over the next five years.

The costs of rolling out network connections in the IP Interconnection and backhaul and core network areas are seen to be low, especially when compared to the considerable costs of deploying access networks in the fixed network. High-speed access networks are at the heart of European connectivity objectives and are therefore expected to be core to the current debate. The ongoing comprehensive revision of the Broadband Cost Reduction Directive (BCRD),[59] set out to achieve the goals of the Digital Compass 2030,[60] therefore focuses on removing barriers to infrastructure rollout and consequently reducing the cost of deployment of high-speed networks.

59 European Commission (2021): Broadband Cost Reduction Directive: summary report of the consultation for its review.
60 European Commission (2022): Europe's Digital Decade: digital targets for 2030.

7. Implications and outlook

IP Interconnection began as a cooperative arrangement between Internet service providers that had equal standing in terms of traffic volume, network size and composition of customer base. As traffic volumes changed, primarily due to large-scale content aggregators, interconnection arrangements began to evolve in order to bring traffic closer to IAPs and to maintain mutual benefits. These developments reflect an ongoing balance of power and interests between Content and Applications Service Providers (CAPs) and Internet Access Providers (IAPs), with each relying on the other in a mutually beneficial relationship.

In the current debate on a possible network charging scheme for CAPs, the argument is that CAPs incur network costs and, if charges were regulated, they would have an incentive to reduce costs. The existing studies up to date show that this proposition is unfounded on evidence. Economically, cost-oriented prices enable markets to function efficiently by allocating resources to services that are popular with consumers. A charging scheme can be consistent with the principle of cost causation if CAPs' activities not only incur IAPs' network costs, but also if such charges can ensure that costs are covered from those whose activities incur the costs.

Furthermore, infrastructure access charging can potentially be justified if the activities of CAPs trigger inefficiencies in network infrastructures of IAPs and quality deterioration – for instance, due to network effects – and CAPs do not sufficiently take these negative effects on board. In principle, CAPs can influence network costs, network efficiency and quality by determining the timing of data traffic generation and the paths of IP Interconnection for delivering data traffic to IAP networks. At the same time, CAPs face constraints that are beyond their control for further improvement of traffic delivery efficiency and quality of service. These include, for example, the timing of traffic-intensive live events, a factor which is admittedly in the event organisers' decision-making sphere, or the lack of space in suitable IAP network exchanges for the installation of caches, or the lack of consent from IAPs for cache installations.

The crucial question that needs to be addressed therefore concerns the extent to which CAPs are able to influence the network costs and thus the efficiency and quality of service of the IAP networks, and what barriers there are to achieving this. Unless this question is addressed, it cannot be assumed that a network charging scheme alone can create an incentive for CAPs to act in a way that is efficient in terms of IAP network costs.

At present, there is very little tangible information from European IAPs on what are likely to be the current inefficiencies in network use and how they might evolve in the future. As this is an unregulated market, there is an overall lack of systematic market data collection and, consequently, a lack of an overall view of the European IP Interconnection market. It should be noted that market developments can differ greatly between individual European countries with regard to key factors such as interconnection paths, competitive intensity of IAPs on end user markets, competitive structure in the markets for CAP services, and business models of IAPs and CAPs. Taking these factors into account, BEREC is expected to analyse the potential impact of a charging scheme moreover on end users, on competition, and on the Internet ecosystem.

According to the European Commission's Better Regulation guidelines, any EU regulatory measure must be grounded in a transparent, comprehensive process and be based on solid evidence-based findings that provide thorough justification and analysis to underpin any changes to the existing policy approach on the relevant topic.[61] "Policy changes affecting relationships between telecom operators and platform providers need to be carefully examined on all aspects and considered by engaging all the relevant stakeholders," state seven EU countries – Denmark, Estonia, Finland, Germany, Ireland, the Netherlands and Sweden – in a warning against possible hasty decisions.[62]

The European Commission has announced that it will launch a consultation in the first quarter of 2023 to examine whether and how major digital platforms should bear some of the costs of Europe's telecoms networks.[63] This triggers all parties to bring forward arguments as to the question of an actual, possibly time limited, contribution to the cost of fibre rollout by all stakeholders in the internet ecosystem, and an allocation of these contributions to stakeholders investing in these networks as opposed to a discussion focused on revenue streams generated from interconnection fees.

61 European Commission: Better regulation: guidelines and toolbox.
62 EUROACTIV (2022): Seven EU countries warn the Commission against hasty decisions on "fair share".
63 Reuters (2022): EU to consult on making Big Tech contribute to telco network costs.

About eco – Association of the Internet Industry

With more than 1,000 member companies, eco is the largest Internet industry association in Europe, incorporating a wide range of stakeholders who build the community of the Internet Ecosystem – from data center and Internet exchange point operators to content and applications providers and telecom network and service providers.

Since 1995, eco has been instrumental in shaping the Internet, fostering new technologies, forming framework conditions, and representing the interests of members in politics and international committees. The focal points of the association are the reliability and strengthening of digital infrastructure, IT security, trust, and ethically-oriented digitalisation. That is why eco advocates for a free, technology-neutral, and high-performance Internet.

Imprint

eco – Association of the Internet Industry
Lichtstraße 43h
50825 Köln

phone: 0221 – 7000 48 – 0
fax: 0221 – 7000 48 – 111

email: info@eco.de
web: https://www.eco.de

The Board:
Oliver Süme (Chair)
Klaus Landefeld (Vice Chair)
Felix Höger
Prof. Dr. Norbert Pohlmann

CEO: Harald A. Summa
Managing Director: Alexander Rabe

Icons on title and figure 1 by Becris, flaticon.com

eco – Association of the Internet Industry
Lichtstr. 43h, D-50825 Cologne, Germany
phone: +49(0)221/700048-0
fax: +49 (0)221 / 700048-111
info@eco.de, https://international.eco.de
@eco_de, @ecoverband

eco

ASSOCIATION OF THE INTERNET INDUSTRY