Netflix's Open Connect program and codec optimisation helped ISPs save over USD1 billion globally in 2021

14 July 2022 • 787088899-293

Consumers stream content on Netflix through a broadband connection they have purchased from an Internet Service Provider (ISP). The cost of providing access to the internet, and content available on the internet, is borne by the ISP, and the ISP seeks to recover these costs through the charge consumers pay for their broadband internet access. To make the delivery of content more efficient and improve the performance of its service, Netflix has invested in content delivery infrastructure and video streaming technology, made available to ISPs through a programme called Open Connect.[1]

Against this backdrop, Netflix commissioned Analysys Mason to quantify the cost savings that Open Connect unlocks for ISPs and illustrate the impact in the UK and South Korea. This paper documents the findings of that analysis. We built a simple bottom-up model of the cost structure of a hypothetical ISP, which has a significant market share in its retail market and purchases wholesale inputs (e.g. from Openreach in the UK context). The model takes as inputs the traffic levels experienced by the ISP, its network topology, the type of caching deployed, and the costs of its network equipment and wholesale inputs.

The aim of this analysis was to:

(i) Provide an overview of the overall cost structure of a large ISP and the costs associated with consumers' use of streaming video services, including Netflix.

(ii) Show the evolution of traffic and network costs over time.

(iii) Estimate the cost savings generated by Open Connect broken down into the cost savings that caching through Open Connect at different levels of the network unlocks for ISPs, and the cost savings enabled by Netflix making available highly optimised video compression technologies, through state-of-the-art codecs that can be used by end users with compatible devices.

The analysis concluded that:

- The network costs of the modelled ISP represent around 50% of retail revenue[2] and can be grouped into:

1 See https://openconnect.netflix.com/en_gb/

2 Based on current and projected average revenue per user per month

- Access network costs, which are largely insensitive to traffic levels, represent the majority of network costs (around 80-90%). These costs related to the 'last mile' access infrastructure between 'edge' or 'local' network node and end-user premises (now increasingly fibre-based), are largely invariant with traffic: these costs are primarily driven by the deployment of the network to homes and offices, and the connection of end users to the network.
- Costs in the core and backbone network are partly sensitive to traffic and represent around 10-20% of network costs. These costs are expected to remain stable over time[3] due to economies of scale, decreasing equipment and link costs for high-capacity links, and the continued delivery of traffic through caches located deep in the modelled ISP's network, including Open Connect Appliances (OCAs) provided by Netflix.

- Open Connect (like other similar partnerships between content providers and ISPs) ensures that growing demand from end users can be handled sustainably without increasing network costs over time. State-of-the art codecs further reduce the traffic intensity of video content, an impact that will increase as more end user devices become compatible with those codecs. Thanks to these investments, the marginal costs of delivering Netflix content represent around 0.5% of total network costs, despite Netflix usage representing about 15% of peak usage in the UK.
- Global savings resulting from Netflix's investment in Open Connect and codec improvements amount to an estimated USD1 billion to USD1.25 billion for ISPs around the world in 2021.

The remainder of this paper provides further details of our analysis. Section 1 provides an overview of Open Connect and where OCAs can be deployed in ISPs' networks. We then describe our approach to estimating the cost savings that Open Connect can deliver to ISPs through caching (Section 2), and through implementation of the latest video codec (Section 3). Section 4 discusses our main findings. Annex A provides more information on the economics and technologies for last mile access and Annex B provides a glossary of technical terms.

1 Open Connect and embedded caches

The internet is made up of independent networks that are connected to one another, directly or indirectly. When someone uses the internet, they stimulate an exchange of traffic between their ISP and the network where the services and content they want to use or access are hosted. This other network could be another ISP's network, a commercial Content Delivery Network (CDN), or a content provider's own network, as shown in Figure 1 below.

In the context of streaming video, when an end user presses "play", they request to watch a specific piece of content through an internet-connected device. This 'playback request' is sent to the streaming service provider (e.g. Netflix), which then sends the content to the end user.

In most cases, large streaming service providers are interconnected directly with large ISPs through peering links, often located in facilities called internet exchange points (IXPs), as shown in the first

---

3 On an annualised basis

two cases in Figure 1. At times, streaming providers and ISPs may need to rely on intermediaries, which can transit the traffic between them as shown in the last two cases in Figure 1.

Figure 1: Different types of interconnection for content delivery on the internet; CP means Content Provider [Source: Analysys Mason, 2022]

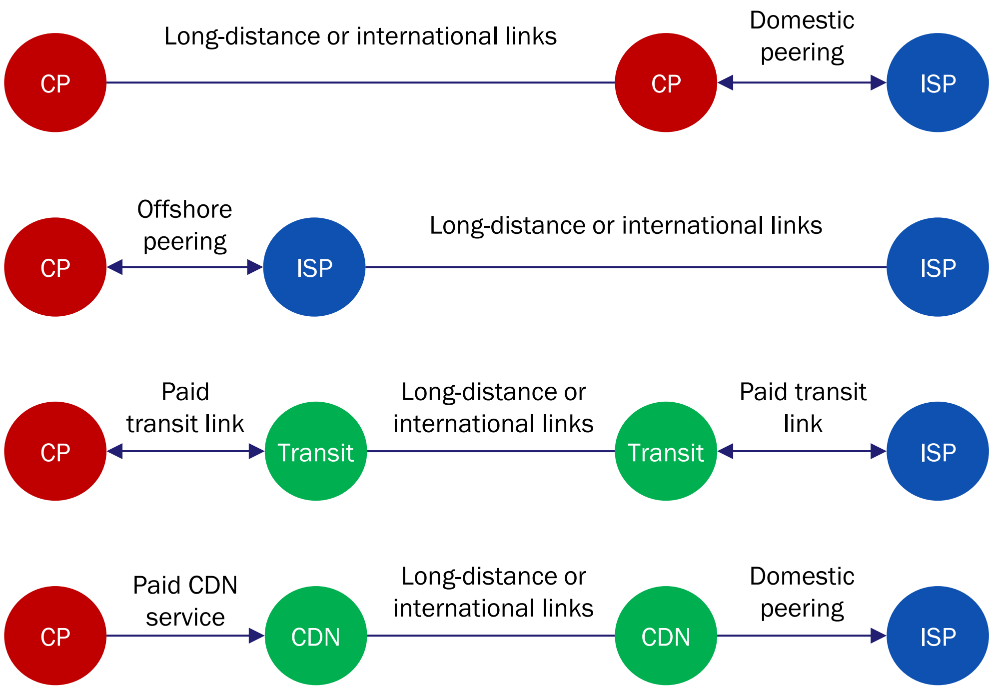

CP = Content Provider; CDN = Content Delivery Network; ISP = Internet Service Provider

Open Connect is Netflix's content delivery programme, through which it partners with ISPs in order to serve its content in a way that optimises the quality for end users and minimises the cost of carrying traffic within ISPs' networks. A core part of this programme is the caching of Netflix content in OCAs. These appliances can be located in Netflix's network, typically in IXPs, or in ISPs' networks, in which case they are referred to as 'embedded' cache appliances.

In its simplest form, embedded caching is used to deliver video to users without the need to increase peering or transit capacity between Netflix and the ISP described in Figure 1. The peering or transit links are typically only used to fill the cache, outside peak times for network demand. This means each piece of content only needs to go from Netflix's network to the cache once, irrespective of how many times it is streamed. In doing so, the content is sent to the ISP at times when there is spare capacity, rather than when end users request it. Indeed, content that Netflix expects to be popular is pre-cached in OCAs (during agreed cache-fill times) before it is made available for members to stream, which mitigates the impact of this content on ISPs' networks.

In most cases, caches are embedded in several locations in an ISP's network, at 'lower' levels in the hierarchy of the network. This allows content to be served closer to end users, which means it only needs to use transmission network capacity that sits between the cache and the end user. The result is a trade-off. The deeper the cache is deployed in the ISP's network (i.e. the closer to end users the cache sits), the more cache locations there need to be to serve all end users, but the less network capacity is used. With Open Connect, we understand that ISPs have the option to deploy OCAs at any node of their network.[4]

2 Approach to quantifying how embedded caches mitigate the impact of traffic on ISP costs

Our study quantified the savings that ISPs can realise by caching video streaming content within their network, at a given quality of service for end users. As a starting point, the study focused on two markets, the UK and South Korea, which are important to Netflix both as a source of highly successful content (The Crown, Squid Game), and as countries where fibre-based broadband is well-developed (Fibre To The Cabinet or FTTC primarily in the UK, Fibre To The Home or FTTH in South Korea).

To do so, we built a model of network costs that reflects the economics of a large hypothetical fibre-based fixed ISP and quantified the difference in costs based on different caching scenarios specifically for Netflix traffic. We extrapolated these results to estimate the benefits of Open Connect globally. We modelled a fixed ISP because we understand that the vast majority of Netflix content viewing happens through fixed connections (including on mobile devices through Wi-Fi) rather than on cellular networks.

We modelled this ISP as an operator which owns and runs a nationwide backhaul and transmission network, built over high-capacity leased Ethernet links. This network connects fibre-based end user connections (the access network) to the internet. Our model reflects a two-tier market structure where the access network is owned and operated by an infrastructure operator (e.g. Openreach in the UK) and leased to ISPs.

We modelled a generic network topology, with three levels of hierarchy: a small number of 'core' nodes connecting to 'metro' nodes in the main cities and regions, which in turn connect to a larger number of 'edge' (or 'local') nodes where end user access lines terminate. The core network is connected to the internet, through a combination of peering and transit relationships, including at IXPs.

Our modelled network topology is shown in Figure 2, which indicates the number of nodes at each level.

4 Provided the deployment site fulfils the Open Connect minimum requirements, including peak bandwidth served at each site of at least 5Gbps and available cache-fill bandwidth of 1.2Gbps over a 12h period every day (see https://openconnect.netflix.com/deploymentguide.pdf)

Figure 2: Modelled network topology [Source: Analysys Mason, 2022]

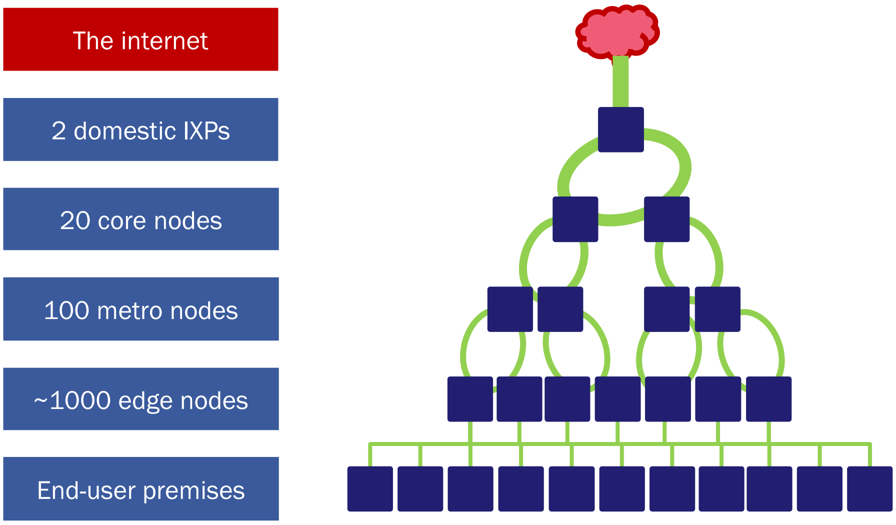

The links between core, metro and edge (or local) nodes collect and aggregate traffic from multiple nodes, and from thousands of end users. The cost of these links is partly dependent on the route and number of nodes connected, and partly dependent on the size of these links and therefore the traffic carried by the network. Conversely, the costs of the access network, between edge (or local) nodes and end-user premises, are primarily a function of the technology used (ADSL, DOCSIS, VDSL or GPON for example)[5] and they do not scale directly with traffic.

The number of links at each level of the network is a function of the number of nodes, and we have modelled each layer between the IXPs and edge (or local) nodes to be organised in transmission rings. The size of the links is determined by the amount of traffic that the network needs to carry in the busiest times of the year.

This approach correlates the amount of traffic carried by the network and some of its costs and illustrates the cost reduction from caching: OCAs store content closer to end users, reducing the capacity required on upstream links in the network. As an example, a comparison of content being served from the internet vs content being served from metro nodes is shown in Figure 3.

5 Asymmetric Digital Subscriber Line, Data Over Cable Service Interface Specification, Very-high-speed Digital Subscriber Line and Gigabit Passive Optical Network; see Glossary

Figure 3: Comparison of content served from points of presence (PoPs) on the internet vs content served from caches in metro nodes [Source: Analysys Mason, 2022]

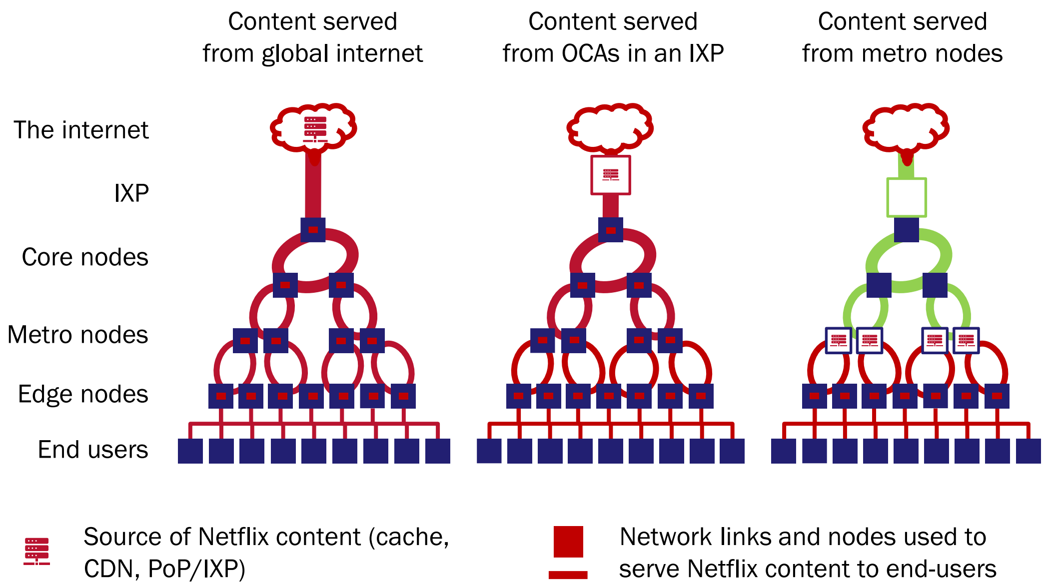

We calculated the annualised cost of the network (taking into account leasing costs for links and depreciation of capitalised costs such as routers and one-off fees), based on the traffic demand on the network (at peak times) and assumptions about the amount of traffic that can be cached at various levels of the network. For Netflix traffic specifically, we ran the model with different scenarios related to the location of OCAs: no OCAs in the country, OCAs only in IXPs, and OCAs embedded in core, metro and edge (or local) nodes.

3 Approach to quantifying the benefits of more efficient encoding

Netflix has also actively worked to make its content available to its subscribers in a wide range of encoding formats, using codecs that are as efficient as possible, whilst retaining broad compatibility for older devices.

In practice, Netflix does this by making each piece of content available in a range of formats and streaming the most efficient format available to a given client and device. OCAs are dimensioned to be able to cache multiple versions of the same content, across much of Netflix's library.

Based on data provided by Netflix, we estimate that around 21% of content streamed in one month in early 2022 benefited from the most recent improvements in codec efficiency. On average, these codecs are around twice as efficient as previous implementations: if they were not available, we estimate that total Netflix traffic globally would be around 24% higher than it is currently (see Figure

4 below). It should be noted that the total gain from new codecs is limited by the compatibility with end user devices.[6]

Figure 4: Estimate of the impact of Netflix providing the latest codec technology on total global Netflix traffic [Source: Analysys Mason, 2022]

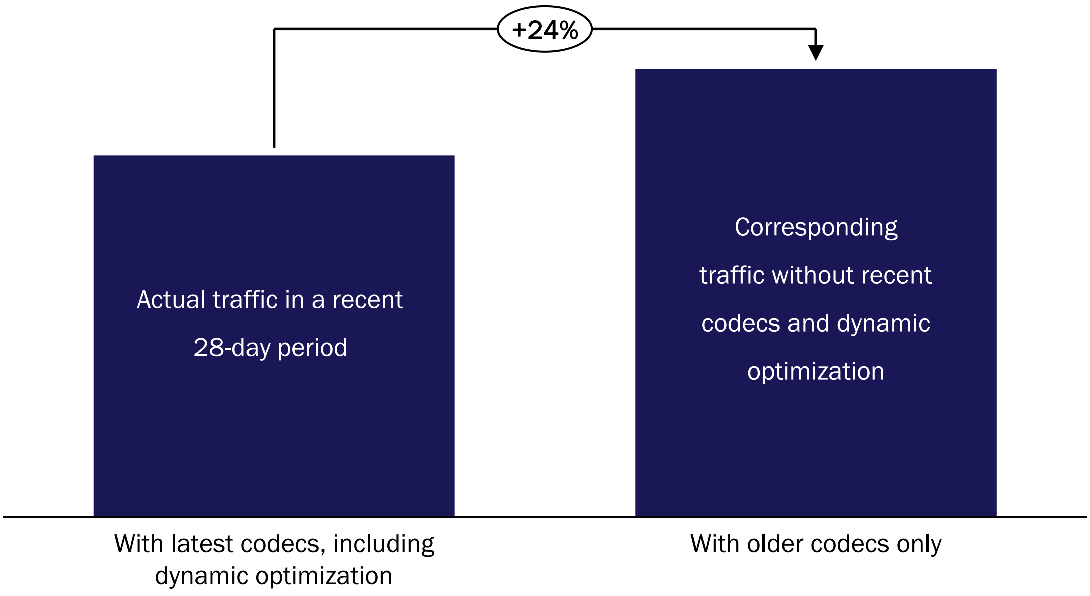

Using the same model we developed to quantify the benefits of caches located at strategic locations in the network, we also estimated the benefits of Netflix making these efficient codecs available, by calculating the savings induced by the lower total Netflix traffic carried by ISPs.[7]

4 Findings and conclusions

Our analysis recognises that ISP costs vary in part with traffic volumes. Traffic-sensitive costs are primarily related to backhaul and core transmission costs, which represent a relatively small proportion of total network costs. Conversely, access network costs, related to last mile access infrastructure (now increasingly fibre-based), are largely invariant with traffic: these costs are primarily driven by the deployment of the network to homes and offices, and the connection of end users to the network.[8]

We estimate that, on an annualised basis, the modelled ISP's total network costs represent around 50% of revenue, and about 80-90% of those network costs are associated with the access network

6 Not everyone benefits from new codecs, because end user devices need to support them, which older devices do not

7 We calculated this saving using the cost model, normalised per Netflix user in the UK and South Korea, which we then scaled to Netflix's global user base

8 Even in the long run, access network costs do not vary with traffic per se; more details in Annex A

(primarily related to depreciation, operation and maintenance) and 10-20% to the backhaul and core transmission network.

We find that traffic-sensitive costs are relatively stable over time: transmission capacity between existing nodes can be upgraded very cost-efficiently as demand increases (from 1 to 10 to 100Gbps, and beyond), and the cost of high-capacity transmission equipment keeps falling. In addition, Open Connect and other similar partnerships between content providers and ISPs (e.g. Google Global Cache) ensure that growing demand from end users can be handled effectively through deep caching and compression without growing costs over time. This is illustrated in Figure 5 below, which shows that annualised traffic-sensitive costs (in the core and backhaul network) have been stable despite rapidly increasing traffic, a trend which is expected to continue in the coming years.[9]

Figure 5: Evolution of traffic and annualised network costs over time [Source: Analysys Mason, 2022]

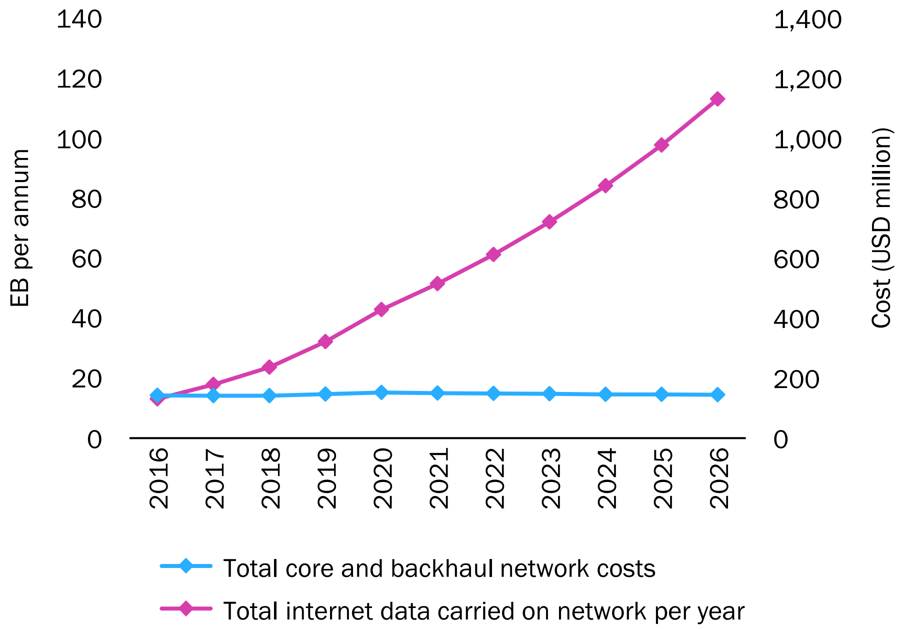

The model shows that for countries like the UK and South Korea, Open Connect enables ISPs to mitigate a significant part of the costs of their transmission networks. OCAs deployed at domestic peering points reduce the need for international connectivity and IP transit, and OCAs embedded in ISP networks enable Netflix content to be streamed using only a portion of the network of ISPs (between the OCA and the end user), thus reducing the load on the rest of the network. Furthermore, even ISPs which are unable or unwilling to locate OCAs within their network still realise savings as a result of the localisation of content within IXPs. Without Netflix's use of a CDN, video would be handed off to ISPs at the source (a handful of cloud storage regions) and the ISPs would then pay to carry the content to their end users, at a significantly greater cost.

9 A similar effect has been found in Lunden, Malmodin et al, 2022, Electricity Consumption and Operational Carbon Emissions of European Telecom Network Operators, Sustainability 2022, 14(5), 2637, https://doi.org/10.3390/su14052637

In the UK, we estimated that a typical large ISP[10] could reduce its international bandwidth and transit costs by around USD9 million annually by connecting to OCAs located in domestic IXPs. In addition, by deploying OCAs in its core, metro or edge nodes, the modelled ISP could reduce its network costs by a further USD7-18 million annually as of 2021. If OCAs were located in edge node, the impact of Netflix traffic on network costs would be close to zero, although in practice OCAs are currently located primarily in core and metro nodes.

As a result, the incremental cost to the modelled ISP of delivering Netflix content to subscribers is much smaller than Netflix's share of user traffic might suggest. The marginal costs of delivering Netflix content represent around 0.5% of total network costs, despite Netflix usage representing about 15% of total usage in the UK in peak times.[11]

Detailed results for the UK are shown in Figure 6.

Figure 6: Annualised cost savings of Netflix OCAs for an ISP network, UK, 2021 [Source: Analysys Mason, 2022]

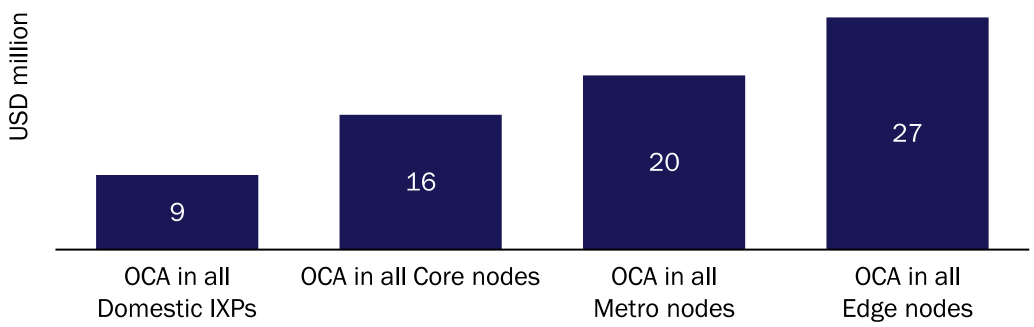

The results for South Korea are shown in Figure 7, and reflect the lower amount of Netflix traffic compared to that in the UK, and the relatively higher cost of international connectivity in the scenario where OCAs are not available in the country.

10 We have modelled an ISP with 30% market share

11 These costs are modelled by removing Netflix traffic from the model. With OCAs in core nodes, the marginal cost of Netflix traffic in the UK represents around 0.6% of total network costs; with OCAs in metro nodes, this falls to less than 0.4% of total network costs; these marginal costs do not include any costs common to different sources of traffic, which form part of the total traffic-sensitive costs

Figure 7: Annualised cost savings of Netflix OCAs for an ISP network, South Korea, 2021 [Source: Analysys Mason, 2022]

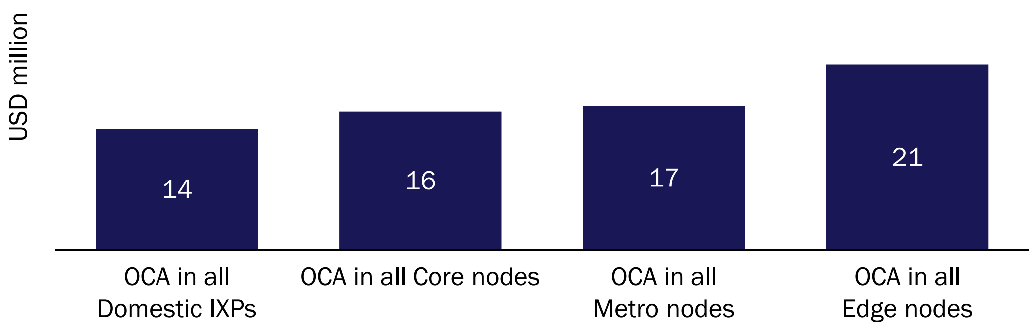

Extrapolating these savings globally, based on the bandwidth served from embedded OCAs around the world, results in estimated cost savings between USD800 million and USD1 billion globally in 2021, assuming a distribution of OCAs at various levels of ISP networks. This estimate is based on the peak bandwidth served from embedded OCAs throughout the world, and on the economics of networks that are dense, highly utilised, and fibre-based, and which are likely to be more efficient (on a cost per Gbps basis) than the average network globally. Therefore, our calculation is likely to underestimate total savings.

We estimated that additional savings resulting from bandwidth reduction from the latest codecs amount to USD0.71 to USD1.20 per Netflix user annually. Scaled to Netflix's global reported user base of 222 million users, this results in further savings of between USD150 million and USD270 million in 2021. **Overall, this brings global savings from OCA and codec improvements to between USD1 billion and USD1.25 billion for 2021**.

Annex A Impact of data traffic on the costs of fixed broadband access networks

In this annex, we describe fixed broadband access networks (Section A.1)[12] and the extent to which traffic and network costs interact in broadband access networks (Section A.2). The focus is on modern fibre-based networks, which reflects the current situation in South Korea and the trend in Europe.[13]

A.1 Overview of fixed broadband access networks

A network mainly consists of two parts: (i) an 'access network' between the end user premises and the local node and (ii) a 'core network', including links sometimes referred to as 'backhaul', between the local node and the point where local ISP interconnects with other ISPs.

The core network is made up of two types of elements: (i) nodes which are buildings where active network functions are being delivered, such as routing, switching or conditional access validation and (ii) transmission links which are the fibre links that connect nodes to one another.

The access part of the network is located between the end-user premises (typically either an office or a home, in a fixed network) and the closest node of the network where active network functions are being delivered, such as routing, switching or conditional access validation for example.

In fixed networks, this last network node has been colloquially called a local exchange (or 'concentrator'), but as fixed networks migrate to full fibre, they are sometimes referred to as 'central offices' or 'local nodes'. This is illustrated in Figure 8, although it must be noted that every network is different, in terms of where these nodes are located, how many users are connected to each node and how nodes at different levels are called.

12 A fixed fibre access network uses wires to connect the end user to the network; this differs from a mobile or wireless network, where the end user connection is provided through radio waves

13 https://www.oecd.org/sti/broadband/broadband-statistics-update.htm

Figure 8: Access network vs. Core and backhaul network [Source: Analysys Mason, 2022]

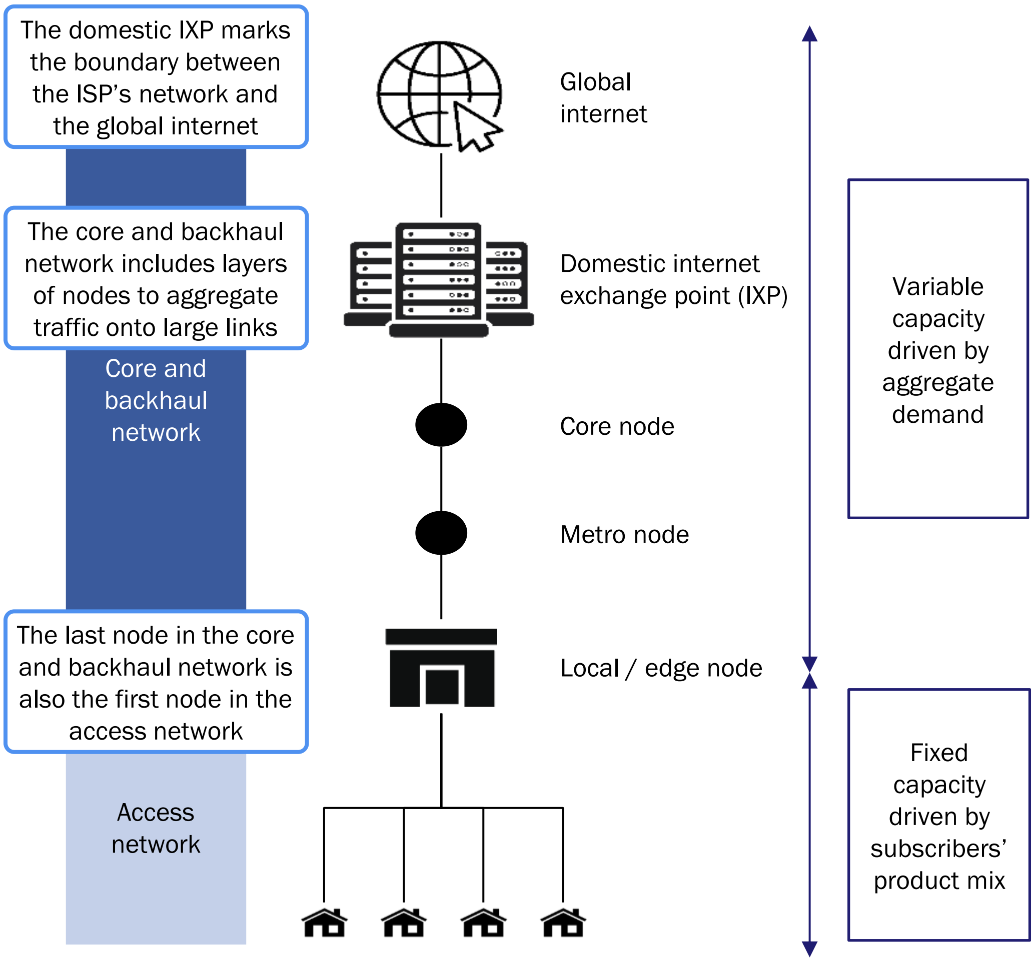

While very network is engineered slightly differently, it is useful to think about the technology and architecture used in fixed access networks as being based on five main archetypes, defined based on what technology is used. This is shown in Figure 9.

Figure 9: Five main architectures of fixed broadband access networks [Source: Analysys Mason, 2022]

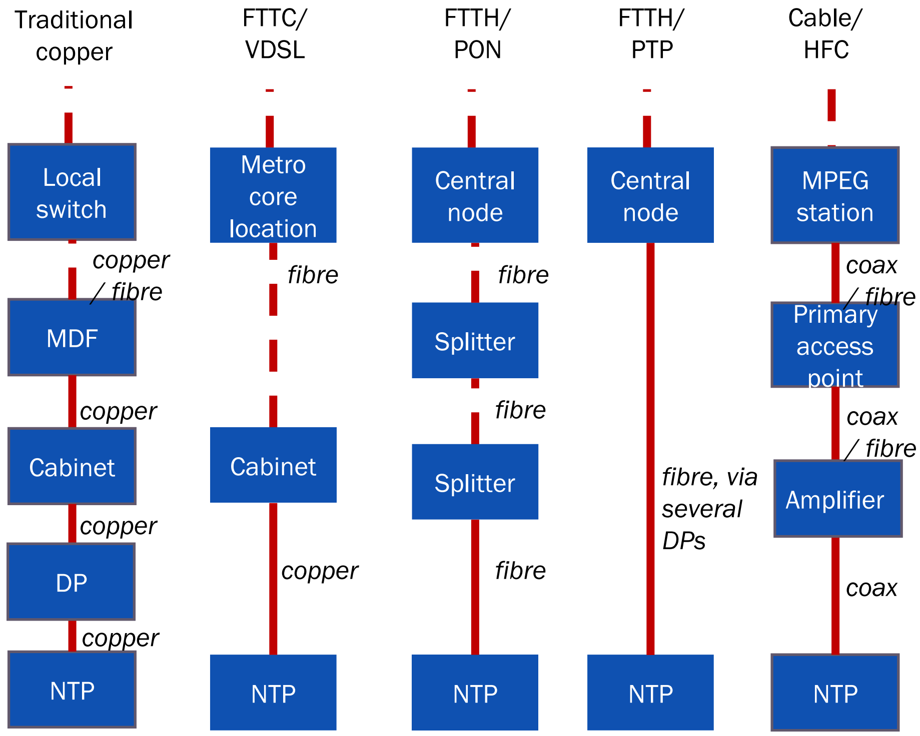

The main differences between those five architectures are described below

- The 'traditional copper' architecture, used in telephone networks for over 100 years, uses copper cables in both the primary and secondary access networks. ISPs can deploy access multiplexers (DSLAMs) in local exchanges to offer ADSL services. These enable end-user speeds up to a theoretical maximum of 24Mbps.[14] In practice speeds are typically significantly lower, up to around 8-10Mbps, as they are highly dependent on the length of the copper 'local loop' between the ADSL equipment in the local exchange, and the network termination point (NTP) at the end user's premises.

- As the internet developed, fixed networks evolved and many operators chose to replace the copper wires between the local exchange and the street cabinet with optical fibre. This reduced the distance over which broadband signals had to travel over copper, and enabled a significant increase in broadband speeds and reliability. This architecture, called fibre-to-the-cabinet (FTTC) still uses copper cables in the secondary access network. It allows ISPs to offer faster services using an access multiplexer placed in or adjacent to the street cabinet through a standard called VDSL, which delivers a theoretical maximum speed of 250Mbps.[15] As for ADSL, speeds achieved in practice are typically below 100Mbps, due to the degradation of the signal on the copper 'sub-loop'.

---

14 https://digital-strategy.ec.europa.eu/en/policies/broadband-technology-comparison

15 https://digital-strategy.ec.europa.eu/en/policies/broadband-technology-comparison

- Cable networks (or hybrid fibre coaxial, HFC), originally used to distribute television signals, have been re-engineered in many countries to also offer broadband internet access, through a technology called DOCSIS. Cable networks offer high speed and capacity on the downlink (traffic going towards the end user) but tend to be limited on the uplink (traffic coming from the end user). DOCSIS-enabled HFC networks can offer a theoretical maximum speed of 1Gbps or more.[16]

- Over the last 15 years, operators have been deploying fibre from the central office all the way to end-users premises. Countries like Singapore and South Korea have achieved high levels of availability and take-up of these fibre-to-the-home (FTTH) networks. Most fibre networks follow one of two architectures:

  – Point-to-point (PTP) networks use a dedicated fibre connection between the end user and the optical line terminal (OLT) in the ISP's network. This means that the only limit in terms of bandwidth comes from the electronics that are installed at both ends to manage the physical and logical connection, as there is no sharing of the access fibre by several end users.[17] FTTH PTP can offer a theoretical maximum speed of 10Gbps or more.[18]

  – More commonly, FTTH networks use a passive optical network (PON) architecture. A single fibre runs from the OLT in the ISPs' network to a splitter located in a street cabinet or a building, to which multiple access fibres are connected. Every end user in a PON splitter receives the same optical signal and decodes its own traffic from it. This multiplexing of traffic from multiple end users on a single fibre line between the splitter and the OLT means that the bandwidth on this fibre[19] is shared, typically among 16, 32, 64 or sometimes 128 end users. This means that although each user has access to the full speed he or she subscribes to(typically up to 1Gbps), it may be constrained when multiple users are using high amounts of bandwidth at the same time, depending on how the available total bandwidth is being shared by other users attached to the same splitter.

All of these architectures constrain the amount of bandwidth that end users can consume at any given point in time, depending on the length of the access link, the technology deployed, and the number of other customers connected to the same network infrastructure. However, as discussed in the next section, the amount of traffic that is consumed on the access network does not drive incremental network upgrades and associated costs.

---

16 https://digital-strategy.ec.europa.eu/en/policies/broadband-technology-comparison

17 Physical constraints for the optical fibre itself are not material and may only come into play in huge multi-Terabit backbone links, if at all. This means that the cost of a PtP fibre connection of a given speed does not scale with traffic (beyond the electricity cost of using the connection, or the long run cost of the electronics ('ports') at both ends (which is very small if considered in "per GB" terms)).

18 https://digital-strategy.ec.europa.eu/en/policies/broadband-technology-comparison

19 2.5Gbps in Singapore's national broadband network for example, but this can go up to 10Gbps in the XGPON standard

A.2 How fibre broadband access network costs respond to traffic

Network costs are often analysed in the 'short run' and in the 'long run'. Short run costs are defined as costs that vary when at least one input or factor of production remains fixed, and long run costs are those that vary when all inputs or factors of production can change. In practice, short-run fixed access network costs are those that vary within a given deployment of an architecture and technology, while long-run costs reflect costs that vary as architecture and technology evolve (e.g. including the costs of deploying a different network).

The architecture of fixed access networks, and the technology that is used in these networks, tends to evolve slowly. The current copper local loop infrastructure in most countries is over 50 years old, and once these networks are replaced with fibre, this technology itself will last decades. Access networks are upgraded over time: dial-up was succeeded by ADSL and then (sometimes) VDSL (with fibre deployed to street cabinets) in traditional copper-based telephone networks; in fibre networks, GPON can be upgraded to higher-capacity technology such as XGPON or XGS-PON. These upgrades enable higher speeds for end users, which in turn may allow them to generate more traffic. Other factors have an impact on the costs of access networks in the short run: when more premises are connected, for example, there can be significant costs involved, especially if the network is not already close to the additional premises.

In the long run, access networks will continue to require investment. There are many factors that influence operators' decisions to carry out these upgrades: competition from other networks and other technologies (e.g. ADSL competing with cable), end user demand for higher speeds, and issues with reliability of copper. Fibre operators are also responding to similar drivers to increase speeds, reduce costs and improve reliability, and are expected to deliver much lower future costs of operation and maintenance and lower carbon emissions.[20]

In HFC[21] and GPON networks, localised access network capacity is shared by multiple end users. As demand grows on these networks, it may theoretically be necessary to carry out some localised upgrades to HFC fibre nodes (e.g. 'node splitting', to serve fewer end users per node) and install GPON splitters (e.g. deploying additional splitters at the same location, to reduce customers per splitter) as traffic grows, to maintain the levels of throughput consumers wish to consume at busy times. In practice, fibre access networks are very far from being congested, as discussed below.

There is a correlation between the average amount of traffic consumed by fixed internet users and the average speed of access networks, as shown in Figure 10 below.

---

20 See for example BT (2020), Full Fibre broadband could help cut your carbon footprint – here's how, https://www.openreach.com/news-and-opinion/articles/full-fibre-broadband-could-help-cut-your-carbon-footprint---here

21 Hybrid Fibre-Coaxial, i.e. Cable TV networks upgraded with fibre, using a mix of fibre and coaxial cable. Such networks offer broadband data services using technologies such as the DOCSIS standards

Figure 10: Evolution of traffic and speeds over time [Source: Analysys Mason, 2022, based on ETNO] [22]

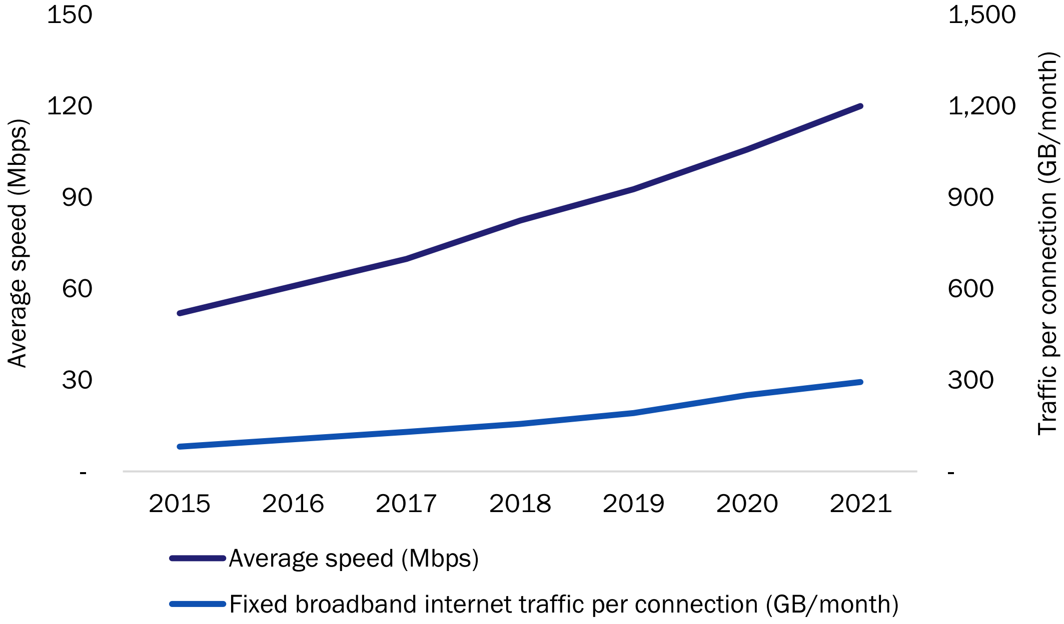

This reflects the evolution of the internet over time, and how online services develop in parallel with increasing access network speeds: as end users' broadband speed increases, they can do more with their internet access, and as richer services develop, they demand faster connections from their ISP. This state of play is also well-illustrated by Ofcom in its UK home broadband performance reports: faster networks enable higher quality content to be used by end users, and median average download speeds closely mirror the access speed purchased by end users from their ISP.[23]

It is important to note that, overall, peak traffic on fixed broadband networks remains very significantly below the theoretical speed and capacity of access networks. For instance, BT in the UK disclosed that average peak traffic increased to 25.5Tbps at its busiest period during 2021, up by 12% from the previous year, or two-and-a-half times the 2018 level.[24] On a per end user basis, we estimate[25] that this represented less than 3Mbps per customer in December 2021, significantly below the average (peak) speed of fixed broadband connections in the UK, which Ofcom estimated was just over 50Mbps in 2021.[26] This reflects the fact that even at the busiest times, some users are not online, or only using low-bandwidth applications (e.g. browsing, email), and other users are streaming video or downloading large files. For reference, video streaming typically uses between

---

22 ETNO (2022), The State of Digital Communications 2022, available at https://etno.eu/library/reports/104-state-of-digi-2022.html

23 UK Home Broadband Performance Report – Dashboard [Source: Ofcom, 2020–-2021]

24 https://www.ispreview.co.uk/index.php/2021/12/broadband-isp-bt-sees-peak-uk-network-traffic-hit-25-5tbps.html

25 Based on an estimated 9 million retail fixed broadband subscribers in December 2021, calculated based on BT's quarterly results for Q3 FY21 ending 31 December 2021

26 https://www.ofcom.org.uk/__data/assets/pdf_file/0020/224192/uk-home-broadband-performance-technical-report-march-2021-data.pdf

about 2Mbps and 20Mbps, depending on the definition (SD, HD, UHD/4K) and the ability of the end user device to decode optimised streams.[27]

To better understand the impact of traffic on short-run and long-run costs, it is useful to consider how increases in traffic are linked to other metrics. First, there is a correlation between higher speeds and increased traffic on average: as speeds increase, consumers gain access to new uses or can experience existing uses with a better quality of experience. Second, as traffic demand for an individual customer increases, a few things can happen:

- If the customer access link is fully dedicated, as is the case with xDSL and PtP fibre connections, traffic can increase as long as there is sufficient capacity upstream in the core network of the ISP. Data can be transmitted at a speed that is only limited by the technology, including the speed of the ports deployed at either end of the link at the customer premises and in the 'local node' that controls the access link, and by the physical constraints on the link, such as the noise and attenuation levels (for example, xDSL speeds are highly sensitive to the length of the copper link).
- If the customer access link is shared, including to some extent in GPON and HFC networks, further limitations can occur if multiple users in the same shared segment of the network are trying to stream or download content at high speeds at the same time at. For example, on a typical GPON segment, a single 2.5Gbps is shared by 32 users.[28] If the splitter is full and all users try to use their connection at the same time, they will only be able to access an average speed of 80Mbps. This is lower than the headline speed fibre customers purchase, and in some cases, ISPs selling 100Mbps or 1Gbps connections could have to upgrade their access networks to maintain the quality of service they offer to end users. In practice, however, 80Mbps is much more than the average bandwidth currently used by broadband users at the busiest times, on the busiest networks. Fibre networks are being engineered to avoid these constraints, with fewer customers per splitter and next-generation XGS-PON technology (which removes the 2.5Gbps constraint) increasingly common.

In both these cases, access network costs do not scale directly with traffic in the short run. In the long run, the speed marketed to, and bought by end users is a stronger driver of costs than traffic, and tends to drive demand for data (rather than the other way round). In other words, because access network deployments are driven by technology shifts, generational upgrades and competition, rather than congestion, the incremental costs of additional traffic in the access network are currently very low.

---

27 Netflix makes available optimised versions of its content, which enable very significant reductions in the bandwidth requirement for high-definition content; an example is the AV1 codec, announced for TVs by Netflix in November 2021, which we understand enables reduction in bitrate for 4K content from 16-20Mbps to less than 5-7Mbps, without a loss of quality

28 Technically this could go up to 128 users, but typical architectures tend to have no more than 32 users per splitter

Annex B Glossary

ADSL / VDSL	Asymmetric / Very high speed Digital Subscriber Line are technologies enabling broadband internet access through standard copper telephone lines. Both technologies' performance is highly dependent on the quality and length of the copper lines, and VDSL typically requires short lines and therefore has been deployed together with fibre-to-the-cabinet (FTTC).
Cache	A cache is used to store static content such as videos for a CDN. Popular content can be 'pushed' into the cache by the CDN, or it can be 'pulled' into the cache by an end user when the first user asks for the content. Once in the cache, content is closer to the end user so that it can be accessed with less latency, and it lowers costs for ISPs (as they do not have to use expensive transit to access the content each time it is requested).
Content delivery network (CDN)	A CDN is a network that distributes content on behalf of a content provider, using point of presence (PoPs) where it is interconnected with ISPs and content providers, and caches that can be located at PoPs or embedded within ISPs' networks. The CDN may be independent of the content providers, or it may be developed by a content provider to distribute its own content.
DOCSIS / HFC	Data Over Cable Service Interface Specification, a technology standard allowing cable TV networks, built as Hybrid Fibre-Coaxial networks (HFC), to carry two-way data for broadband internet access. The bandwidth available for internet connectivity is shared between users on the same coaxial tree, but can be easily expanded downstream; upstream bandwidth has historically been very limited by the standard.
FTTC / FTTH	Fibre to the Cabinet (FTTC) is an architecture where telephone networks are upgraded through a replacement of the copper line between a local exchange (i.e. an 'edge' or 'local' node) and a street cabinet where the 'sub-loop' to the end user premises, which remains copper-based, connects/ Fibre to the Home (FTTH) is an architecture where the entire copper line is replaced with fibre. Sometimes generalised as Fibre to the Premises (FTTP) to reflect connectivity to non-residential premises (e.g. offices).
GPON / XGS-PON	Gigabit Passive Optical Networks, an architecture for FTTH networks where end users are connected to the network through shared fibre trees, with one fibre between the optical line terminal (OLT) equipment in the edge / local node and a splitter, and multiple individual fibres connecting the splitter to the optical network terminal (ONT) at the end user premise. a GPON 'tree' shares a total of 2.5Gbps of bandwidth between the OLT and the splitter, and

	its evolutions (including XGS-PON) enable 10Gbps end user access speeds and much higher bandwidth between the OLT and the splitter.
Internet exchange point (IXP)	An IXP is a location where internet providers, including ISPs, backbones, CDNs, and enterprises, can meet and efficiently exchange traffic, using peering or transit arrangements. Instead of having to arrange separate circuits for each provider which can be costly, each provider only needs one connection to the IXP (and many connections within it) to exchange traffic with all the other providers. An IXP can also help to attract content, by providing an efficient way to distribute the content.
Internet service provider (ISP)	An ISP provides internet access to end users, using mobile or fixed broadband connections, and can also provide access to enterprises. It interconnects with other ISPs via a combination of peering, partial transit, and transit relationships. This interconnection with other ISPs or internet backbone providers is often at an IXP.
Latency	Latency is a measure of the time it takes for traffic sent from a source until to be received at destination. High latency can affect the experienced quality of the communication.
OCA	Open Connect Appliance, the cache server used by Netflix as part of their CDN.
Peering	Peering is an arrangement between two internet providers to exchange their own traffic. For instance, two backbones can use peering to exchange their ISP customers' traffic, or an ISP can peer with a content provider to deliver the content requested by its end users. Peering partners will not allow peering traffic to *transit* their networks – that is they will not deliver content from one peering partner to another. Peering is typically 'settlement free', meaning that neither party pays or is paid to exchange traffic with the other, although in some cases peering is subject to negotiated fees ('paid peering').
Point of presence (PoP)	A point of presence is a location where a CDN can deliver dynamic content to a backbone or ISP.
Transit	In a transit arrangement, a backbone will sell access to the entire internet to another backbone or ISP. This includes access to the backbone's own customers as well as the backbone's peering partners. Partial transit is a similar arrangement, but only for a subset of the internet.