



The Importance of Peering

In the previous section, we learned how Network Operators fit into the global Internet ecosystem.

We also learned that they can peer with other networks, or buy transit from other networks, or both.

Operators peer with each other to:

- reduce [costs](#)
- reduce [latency](#)
- improve [service quality](#)
- access to [content](#)
- have “unlimited” [bandwidth](#)
- have direct [relationships](#) with other operators

This section covers the value proposition for end-site networks to consider adding Peering capabilities to their existing Internet access provision.

Costs

(UPDATED)

The commercial part of the Internet is highly competitive with many network operators vying to provide the highest quality service to their end users at the lowest possible cost.

Apart from staff and equipment costs, the other significant cost of providing Internet access is to actually obtain that access to the whole of the Internet.

The simplest way of getting access to the whole Internet is to pay someone else to do it for you. That's a typical transit service that a new network operator would buy. But it is expensive, and gets more so as traffic levels increase. This model was common in the early years of the Internet, up until the early 2000s, but in recent years, there has been a fundamental change in focus as network operators work as hard as they can to establish as much peering as possible.

Peering has no traffic charges, and so the more peering an operator can achieve, the lower the cost of traffic charges paid for transit. This reduction in operating expenses (OpEx) means better value Internet access for customers, or more financial ability to invest in newer/bigger/better infrastructure for the network operator, or hiring more technically skilled staff, or a combination of all of these.

Note: in some parts of the Internet it is actually cheaper to outsource peering by buying cheap

transit. Quite often the operational overhead, the cost of ports and cross-connects, the IXP membership fee (not to mention paid peerings) makes the cost of peering quite often comparable with cheap transit. Here the operator has to weigh up the benefits of peering versus delegating all those to the cheap transit provider.

Latency

Latency is the time that it takes for an IP packet to get from its source to the destination, and for the response to return.

From an end-user perspective, the higher the latency, the slower an application appears to function. Internet users today expect that online applications work “instantaneously” and are likely to be dissatisfied if the experience is anything but. Hence a network operator is very focused on ensuring minimum latencies from their users/clients to the most popular destinations.

Quite often the path from the network operator through their transit provider may be more indirect than is possible if the operator is connected directly to the content provider themselves. Several Internet applications are latency sensitive (video conferencing, online games, e-sports etc), and so a good provider will be looking for all opportunities to improve the latency their customers are experiencing.

Service Quality

Service Quality really defines the experience that end-users get from their service provider, not to be confused with QoS which is a packet prioritisation mechanism used for congested links.

If a network operator is in direct control of links to other operators, they then can control capacity, the service level agreement of the link, manage outages, manage connectivity issues with these other operators, and so on.

Relying solely on an upstream provider reduces the network operator's opportunity to give good service quality to their customers; their service quality can only be as good as that of their upstream provider.

Access to Content

Most of the major content providers (the so-called [Hyperscale Content Providers](#)) are present at many of the Internet Exchange Points and private peering facilities (datacentres) worldwide.

They have built their own content distribution networks (using their own fibre optic infrastructure rather than buying transit) and regional data centres, with their goal being to get their content to their “[eyeballs](#)” with the lowest latency, highest speed, and greatest reliability.

With these Hyperscale Content Providers present in so many places, network operators of all types work hard to ensure that they can peer directly with these content providers as close as they can to their own network infrastructure. Hauling data half way (or even a quarter of the way) around the

world, as was done in the 1990s and 2000s, was replaced in by the early 2010s with operators turning up at Internet Exchange Points to interconnect with the Content Providers and Content Distribution Networks directly.

Given around 80% of traffic of any network access provider is sourced from these Hyperscale Content Providers, there is a overwhelming value proposition for all network operators who are [able to peer](#) (i.e. have their IP address space, AS number, own transit arrangements) to participate in peering at their nearest IXP or private interconnect facility.

Bandwidth

Bandwidth is the amount of capacity available between two network operators. As with [latency](#), the lack of sufficient bandwidth is likely to cause dissatisfaction amongst the access network operator's clients/users. Which means that bandwidth is another key focus of a network operator.

Quite often the capacity available from the network operator through their transit provider might have limitations that the network operator hasn't anticipated or, if there is a rapid increase in usage, hasn't purchased. This results in reduced throughput, congestion, and packet loss. This reflects poorly on the network operator's credentials as a quality service provider, and causes significant customer disappointment.

If the operator is peered directly with a content provider or other network operator, the capacity is usually provisioned over a direct cross-connect (fibre optic etc), with both entities being able to directly adjust capacity as required. Fundamentally, a fibre cross-connect is only capacity limited by the equipment (routers, switches, fibre optic transceivers) being used by each operator.

Relationships

Having a direct relationship with a content provider or other network operator, rather than via a third party (the upstream), can often mean that latency, bandwidth, and other issues with content and service delivery being experienced by end-users can be more efficiently and effectively resolved, rather than having to work with the upstream/transit provider as an intermediary.

The direct relationship also means that the content provider is able to use their own algorithms to optimise delivery of their content via that directly connected service provider.

[Back to Home page](#)

From:

<https://bgp4all.com/pfs/> - **Philip Smith's Internet Development Site**

Permanent link:

<https://bgp4all.com/pfs/peering-toolbox/why-peer?rev=1679877140>

Last update: **2023/03/27 00:32**

