Evolution of Internet interconnection

Authors: Joost Bosman, Bart Gijsen (TNO)

Date: 13 August 2020

The Internet has become the fundament for digital services that have transformed every-day-life around the globe. It has been an enabler for many innovations that boost social and economic prosperity. At the same time the internet's growing popularity also introduces concerns, amongst others about cyber security, dependability and privacy protection. Consequently, Internet related opportunities and concerns are on the agenda of policymakers, occasionally arousing debate and calls for action.

Recent debate has been caused by dominant content and cloud service providers that are creating global, private network infrastructures for deploying their services. While governments have stimulated business drivers for boosting internet *use*, the increasing hold of business drivers on the internet *infrastructure* is perceived as threatening. This leads to calls for action, amongst others, regarding digital sovereignty by governments around the globe. Recent examples include the Internet Governance Forum session about Digital sovereignty (Internet Governance Forum, 2019) and the United Nations cybercrime-focused resolution that aims to provide government more control over the Internet (United Nations, 2019).

Effective policymaking regarding this topic requires up-to-date understanding of the Internet infrastructure and its ecosystem. From the perspective of this paper 'the Internet' is constituted by interconnected data communication networks around the globe. In particular, we focus on the interconnections between the networks (the core of the Internet), as opposed to the 'last-mile' communication inside the networks that connects communication end-users.

A complicating factor for governmental organizations involved in Internet policymaking is the relatively high technical jargon and the absence of a clear 'map' of the internet's topology and common traffic routes. Moreover, these aspects of the Internet infrastructure are becoming outdated faster than for other infrastructures (e.g. global transportation or electricity networks), due to the continuous, high pace of evolution of Internet use, it's digital infrastructure and the stakeholder ecosystem.

In this paper we aim to provide policymakers an update on the core of the Internet infrastructure as the foundation for future Internet policy decisions, and we clarify how Internet stakeholders have shaped the evolution of topology of interconnections and the traffic that is transported over it.

Reading guide

In Section 1 we start by sketching out key Internet evolution phases in order to make clear why the Internet infrastructure has become what it is today. In the next section we zoom in on the Internet interconnection infrastructure that has been created by the most recent evolution phase. Finally, we will discuss how drivers of the dominant stakeholders in each of the Internet evolution phases have shaped the Internet infrastructure. Understanding these drivers is of key importance for effective Internet policymaking.

Summary

During the evolution of Internet interconnectivity its topology has been shaped by the drivers of the dominant stakeholders. After the academic inception phase the dominant stakeholders have been those that were most successful in commercial exploitation of the global Internet infrastructure. The highlights of this evolution are presented in the following table.

Evolution period		Dominant Internet stakeholders	Stakeholder drivers	Impact on Internet interconnection
1960's -	Academic inception	Academic Internet pioneers	Four ground rules of the open Internet	Best-effort, autonomous networks interconnected by NSFnet backbone
1990's -	Commercialization	Internet Service Providers	Commercial grade quality for basic Internet (access) services	From commercial walled-gardens to interconnections via IXPs by 'tiered' networks and standardization of Internet protocols
2000's -	Content delivery	Content Delivery Network providers	High bandwidth for commercial multi-media services	Commercial networks with CDN (cache) solutions, higher and asymmetric bandwidth capacity and interconnected by globally expanded IXPs
2010's -	Cloud services & social media	Cloud and social media providers (Internet giants)	Controlled quality, reliability and huge bandwidth for business-grade cloud services	Tiered Internet is 'flattening' and becomes paralleled with global private (programmable) networks that are peering via cloud exchanges

Table 1: Overview of Internet interconnection evolution

While from an architectural point of view the Internet infrastructure is decoupled from the content that it transports, the evolution indicates otherwise. Content characteristics such as quality of service, high-bandwidth content and cloud service reliability have driven their providers to optimize the Internet infrastructure to fit their use. For policymakers it is important to keep this causality in mind: without non-commercial guidance the strongest business case among Internet services creates the Internet infrastructure that sustains that business case.

In the most recent evolution period the Internet giants have created global private networks that are unprecedented in size and their wide-range usage, and that at an impressive pace. Those networks emerged from friction between their business objectives and the four ground rules of the open Internet[1]. Apparently business revenues of the Internet giants outweigh the investment in global private networks. In analogy to a highway infrastructure: imagine several transport & logistics corporations whose businesses are suffering from daily traffic-jams – leading to each building their own highways. Although the public, open Internet is still the dominant infrastructure, it is undeniable that 'private Internet' infrastructures are no longer marginal. In itself this is no problem, but it can become another threat to the open Internet. For example, if private Internets will be used as alternative transit networks or for other services that are currently offered via the public Internet. Once public Internet services start to be offered via private networks this may impact the scope and effectiveness of, for example, regulatory instruments such as the telecom act. While current Internet

[1] The term 'open Internet' has multiple interpretations and dimensions. In this paper we mean network interconnections that adhere to the four ground rules formulated in the Internet pioneering phase: each network is autonomous, networks are connected by 'black boxes', communication will be best effort and there will be no global control at the operational level. More detailed explanations of the dimensions, benefits and challenges for the open Internet have been provided by, among other, OECD (OECD, 2016).

governance debates are focusing on the public Internet, it is important for policymakers to recognize this trend (of the rise of private networks) when referring to ‘the’ Internet.

While commercial stakeholders have been the key players in building and evolving the Internet interconnection infrastructure, governments are increasingly voicing concerns regarding Internet governance. They are applying their regulatory instruments to serve the public interest and to safeguard digital infrastructure and national security. This results in regulatory measures regarding, among other, digital protectionism, cybercrime and privacy. Some of these measures can be regarded as Internet restrictive policies, opposing the ‘open Internet’. Policymakers should be aware that transparent and consistent governance is an important contribution to the continuation of the economic and social benefits that the Internet has brought. In absence of consistent Internet governance, the effectiveness of public policy topics like cybercrime, cyber diplomacy and net neutrality may increasingly become dependent on the owners of global private networks.

The success of the Internet is rooted in the way it was built and evolved: an open platform for innovation and sharing of ideas. It has flourished under governance that was driven by consensus, which requires policymakers to balance the interest that they serve with the interest of other stakeholders. From the recent developments in the Internet interconnection evolution we learn that keeping an open mind is essential, because otherwise even an infrastructure as large as the public Internet can and will be paralleled and lose its collective force of innovation.

1. History of the Internet

Academic inception (1960's – 1995)

The Internet originated from the academic pioneers in the beginning of the 1960's. At that time telecommunication consisted of telephony services and new ways were being investigated for communication between computers. The inception of packet switching theory (enabling computers to exchange small pieces of data, called packets) paved the way for creating computer networks. Along with the increasing use of computer systems, mostly by academic institutions, the need for interconnecting them grew. In response to this need the DARPA agency funded the design and implementation of the ARPANET in the second half of the 1960's. More and more computers were added to the ARPANET that were communicating via a computer-to-computer protocol. In 1972 email was the first, new data communication service demonstrated to the public.

As the interest for connecting more and more computers grew it was envisioned that other networks would emerge alongside the ARPANET. Also, networks could use different communication media, such as satellite or radio networks. Depending on the applicable medium and the intended use of those networks they would have different requirements and designs. In this context work on "internetting" those networks was started and four 'ground rules' were used to guide the development of the Internet (ISOC, 1997):

- Each network is autonomous and can be connected to the Internet without required changes,
- Communication will be best effort,
- Black boxes (routers, gateways) will be used for connecting networks that will not expose information about internal network or traffic details,
- There will be no global control at the operational level.

In the mid 1970's and early 1980's most computer networks remained purpose-specific for a closed community. This was until 1985 when amongst others the U.S. NSF and British JANET started funding Internet connections for the entire education community. The NSFnet Internet backbone was targeted for global deployment. This required the introduction of technology and governance for global Internet addressing, routing and end-to-end transport control. Protocols including Internet Protocol (IP) and Transport Control Protocol (TCP) were proposed and agreed upon in the Internet Architecture Board (IAB) and the Internet Engineering Task Force (IETF).

Commercialization

Once the foundation for global deployment was laid the first commercial Internet Service Provider (ISP) was launched in 1989. This enable the wider public to be connected to the Internet. The invention of hyper-text technology and the concept of the world wide web in 1991 was the key enabler for the development of websites and e-commerce services. The networking technology oriented IAB and IETF were complemented with a new coordination organization: the World Wide Web Consortium (W3C). With the introduction of the world wide web the Internet became a fundamental infrastructure for global scale business and societal innovations.

Initially, commercial Internet services were offered by ISPs within 'walled-gardens': commercial web services enabled by the ISP could only be reached by the ISP's subscribed customers. ISPs were reluctant to rely on interconnection via a best-effort Internet backbone such as NSFnet. Apparently the second ground rule of the Internet complicated the next step in the Internet evolution. In response, a new type of telecommunication providers seized this opportunity to launch commercial internet peering services. Around 1995 the Border Gateway Protocol (BGP) was standardized to facilitate those providers to offer connectivity services between ISPs. For efficient implementation of peering interconnections, providers launched Internet exchange memberships that operated data centers in which interconnections between network gateways were established. Those developments largely eliminated the need for the NSFnet and its funding was stopped, marking the end of the academic inception period of the Internet.

Further (commercial) evolution of the Internet appeared to be fast and got overinflated. The acceleration of Internet use by mobile telecommunication had a down-side. The extremely high bids of ISPs for 3G mobile spectra auctioned by European countries in the early 2000’s led to the burst of the Internet hype. The entire global ISP sector took a blow and down-sized their Internet innovations drastically. This paved the way for the next step in the Internet evolution.

Content delivery

Although the introduction of 3G mobile access networks was tumultuous, it illustrated the rise of broadband Internet access. The high access speeds enabled the delivery of multi-media services over the Internet, as an alternative for traditional media networks such as cable TV. While many ISPs were still recovering and retreating to their core business, new players seized the opportunity and started offering multi-media services ‘over the top’ (OTT) of the Internet infrastructure.

The service characteristics of multi-media did not match well with the second ground rule of the Internet. Best-effort Internet was not sufficient for the required quality of service for multi-media delivery. Similar to the early commercialization phase of the Internet, new commercial parties entered the market to ensure acceptable delivery of content over the best-effort internet.

One of the pioneering Content Delivery Network (CDN) providers was Akamai. The first CDN providers are referred to as generic CDNs, since they served many content providers to deploy their services via many ISP networks. For optimizing their operations CDN providers have also boosted traffic volumes through Internet eXchange Points (IXPs).

For generic CDNs a complicating factor for ensuring quality of service over end-to-end Internet routes was the third ground rule: networks should be connected via the black box principle. This limited the deployment of end-to-end optimization techniques for OTT services. Generic CDNs worked around this limitation by starting to deploy their CDN caching solutions within ISP networks.

The deployment of content caches in ISP networks was also mitigating another problem. Boosted by the growing popularity of multi-media via the Internet the amount of data that was transported over the Internet rose and this required steeply increasing ISP network capacity. As a result the profits went to the multi-media providers, while the investment cost ended up at the ISPs. The deployment of content caches by generic CDNs reduced the amount of transported traffic and thereby reduced or postponed investments by ISPs.

Cloud services and social media

After a few years the CDN providers were joined by content generating providers in their prominent role in the Internet infrastructure. In de business sector the adoption of cloud service providers (including Amazon, Google and Microsoft) and in the consumer sector social media providers started to prevail (including Facebook, Apple and Google). The business of these parties spans the globe as well as their data centers, which is why they are referred to as Internet giants.

Similar to preceding commercial Internet providers cloud and social media services also impose specific requirements on the Internet to deliver services to end users. Moreover, the highly dynamic and fast evolving ecosystem requires highly flexible management of their large scale infrastructure. This requirement conflicts with the fourth ground rule of the Internet regarding the absence of any organization to retain global control over the Internet. Further, the cloud and social media providers were increasingly transporting content between their global data centers for purpose of reliability (via data duplication over multiple data centers) and speed of content delivery to end-user (by storing content in data center as local as possible). The cost of content duplication started to go beyond the cost of delivering content to end-users. As a result these Internet giants have been creating global scale private networks. To enable easy connectivity between Internet giants and their customers cloud exchanges were introduced.

In addition to their contribution to Internet interconnection topology the cloud and social media service providers have also strongly contributed to the development of Internet technology (including encryption for security and privacy purposes, and software defined networking for flexible, large-scale network management). Much of these technologies have been made available to the Internet community via standardization bodies such as the IETF and W3C, as well as via new open communities. As such, Internet giants have significantly transformed both the Internet topology and traffic routes, as well as Internet technology.

Next evolution phase

While in the past decade Internet evolution was largely boosted by the business drivers of the Internet giants, one can also observe increasing opposition. For example, governments have initiated opposition to overly dominant market positions of the Internet Giants, have acted to protect civilian's privacy and (announced to) mitigate risks regarding digital sovereignty (Internet Governance Forum, 2019). So far, the ground rules of the Internet have underpinned the multi-stakeholder, 'one world, one Internet' (ICANN, 2013) principle, but it is clear that this principle and the ground rules are under pressure from multiple stakeholder perspectives. As we argue in section 3 the drivers of current Internet stakeholders are already starting to 'partition' Internet traffic via private, mission specific wide-area networks (WAN).

In summary, the evolution phases of Internet interconnection and the dominant players therein are:

Evolution period		Dominant Internet stakeholders
1960's -	Academic inception	Academic Internet pioneers
1990's -	Commercialization	Internet Service Providers
2000's -	Content delivery	Content Delivery Network providers
2010's -	Cloud services and social media	Cloud and social media providers (Internet giants)
2020's -	Private, mission specific WANs ?	?

2. Modern CDN and Cloud networks

In the most recent evolution phase of Internet interconnectivity large CDN and cloud providers (Internet giants) have been creating global scale private networks. In this section we zoom in on some of those private networks and the objectives of the Internet giants to create them.

Objectives

The deployment of private networks by internet giants is motivated by two type of incentives:

1) Revenue gain and controlled cost: Today, (business) customers require higher bandwidth performance and availability than best-effort Internet can offer. Without controllable network infrastructure Internet giants would suffer revenue loss due to SLA violation penalties[2] and prevent them for acquiring higher revenue margins for high performance cloud services. Moreover, the amount of data transferred for their content delivery and cloud services has steeply increased and made transport via other network providers expensive (Sushant Jain, 2013).
2) Demand for Flexibility: Private networks enable Internet giants to offer on-demand services (i.e. very short time-to-deploy) to their customers, with a high degree of flexibility (e.g. customer self-configuration and optimized cloud traffic engineering). Moreover, it allows internet giants to quickly roll-out new services with global availability. These type of incentives have been the key motivation in the early designs of Internet giant's private networks (Sushant Jain, 2013).

Global private networks

In order to provide an impression of the global scale of modern private networks we summarize and visualize data regarding the points of presence and cable infrastructure that is published, at the time of writing this paper[3], by three Internet giants: Amazon, Google and Microsoft.

Points of presence (PoPs)

The points of presence consist of the edge locations where direct connections between the private networks of Internet giants and their connectivity partners are established. As an example, some of the PoPs to Google's network are shown in Figure 1. Via the networks of the CSP connectivity partners, including cloud exchange providers, the connectivity to the CSPs spans even further around the world. The global distribution of the PoPs reflects the distribution of the customer base of CSPs.

Private cable infrastructures

To interconnect their data centers each of the Internet giants have invested heavily in private cable infrastructures. Between 2013 and 2017 investment from in particular Google, Microsoft, and Facebook have added capacity at a significantly higher rate (at least 75 percent per year) than traditional long distance operators (less than 45 percent) (Teleography, 2019). This has resulted in global spanning private cable infrastructures. As an example Figure 1 illustrates the submarine and terrestrial cable infrastructure (partly) owned by Google, with a total submarine cable length of

[2] SLAs of all major CSP (e.g. https://cloud.google.com/vpn/sla, https://aws.amazon.com/directconnect/sla/) include financial penalty clauses for the provider, in case the service levels are not met.

[3] Sources: https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/?p=ugi&l=ap
https://cloud.google.com/vpc/docs/edge-locations
https://docs.microsoft.com/en-us/azure/expressroute/expressroute-locations-providers
http://www.centurylink-business.com/demos/network-maps.html

more than 100 thousand kilometers[4]. In contrast, according to Teleography indications are that there are more than 1.2 million kilometers of submarine cables in service globally[4].

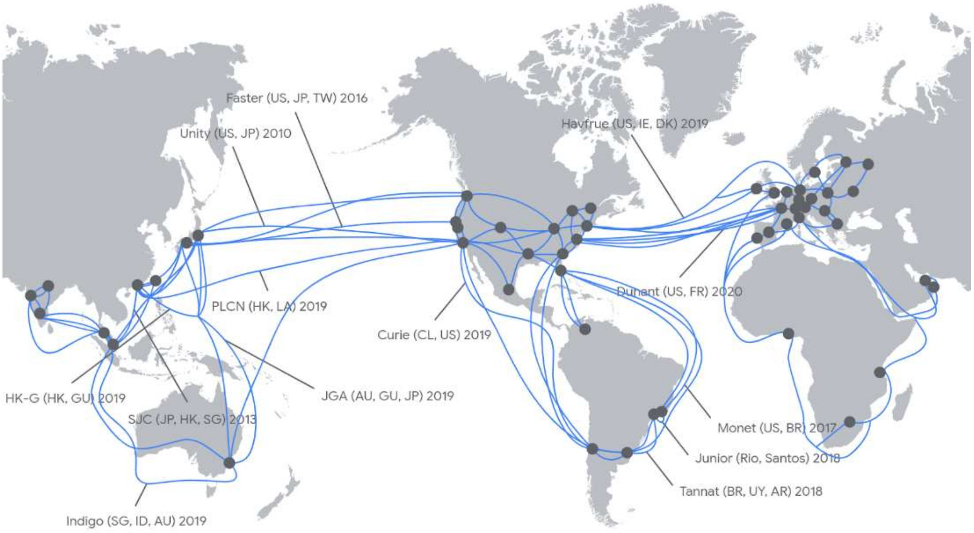

Figure 1: Google submarine dark fiber network

It is fuzzy to compare cable dimensions between network providers, amongst others due to co-ownerships of global cables and unclear actual use of the varying cable bandwidths. Nevertheless, it is more than plausible that the cable infrastructures of Internet giants are comparable to the largest networks in the Internet.

Offered connectivity services

While the network infrastructures of Internet giants is comparable to those of large telecom providers, their business services are different. While the latter provide generic telecommunication services the private networks are being used for a subset of telecommunication traffic, i.e. transporting media content and cloud traffic. These content and cloud services still use the public Internet, in particular for the last mile connectivity to end users.

However, in the business market there is a trend towards direct connectivity between enterprises and the content and cloud service providers. Following the development of private networks the Internet giant started to offer private interconnectivity consisting of physically wired connectivity. From a network architecture point of view the availability of private interconnection enables bypassing of the public Internet for the transport of traffic between enterprise and cloud networks.

Private Internets?

While private networks have been used for a long time for specific purposes (typically data transfer between partners in sectors such as finance, government, aviation, etc.), this is different for the private networks of the Internet giants. Their networks offer data, compute and communication services that are being used in any business sector. Moreover, they are starting to offer private peering services[5] that enables data transfer between their customers within the global private networks. While there are still limitations in multi-tenant cloud peering (e.g. peering between CSPs is not standardized) this concept potentially diverges data transfer that has so far been sent via the

---

[4] Sources: https://cloud.google.com/blog/products/networking/google-cloud-networking-in-depth-cloud-cdn
https://www2.telegeography.com/submarine-cable-faqs-frequently-asked-questions

[5] Sources: https://azure.microsoft.com/it-it/blog/create-a-transit-vnet-using-vnet-peering/
https://cloud.google.com/vpc/docs/vpc-peering
https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html

public Internet. In conclusion, both the size as well as the generic applicability of the Internet giants' private networks are unprecedented. Emerged from business objectives that are conflicting with the ground rules of the open, public Internet, this raises the question whether private Internets are emerging in parallel to the public Internet.

In addition to the business driven partitioning in public and private Internet interconnection, this effect is also influenced by digital service legislation. Legislation is bound by geographic borders, and network and data related legislation (e.g. network neutrality, GDPR, foreign supplier policy, etc.) is not uniform across geographic regions. This is reflected in the Internet interconnections that can be established and restrictions on cloud data storage (e.g. Amazon's cloud regions, GovCloud, CloudNL and the GAIA-X initiative). Were the Internet has seen commercial walled-gardens in the past, it is starting to face legal walled-gardens at present.

3. Impact of evolutionary drivers on Internet interconnection

In Section 0 we have sketched how drivers of dominant Internet stakeholders push specific innovations in the Internet. For example, the need for delivering quality of service for commercial services and higher bandwidths for transporting multi-media content over the Internet has introduced new solutions such as Internet and cloud exchanges and content caching. In this section we explain the impact that these solutions have had on the Internet interconnection topology and traffic routes.

Internet interconnection in the early evolution phases

The evolution phases of Internet interconnection started with the NSFnet backbone, that was replaced by the commercial ISP in the Internet commercialization phase. As stated in the first ground rule each network of an ISP is considered an autonomous system (AS). After a short period where ISPs offered connectivity between web providers and their own end-customers via the IPS's network only (walled-garden), the ISPs started to interconnect their networks (peering). For efficiency reasons the ISPs founded Internet Exchange memberships that operate data center locations (IXPs) where multiple ISP members interconnected their networks. Well-known IXPs in Europe include AMS-IX, DE-CIX, LINX and France-IX.

To facilitate global end-to-end delivery of IP traffic a multi-tiered ecosystem of ISPs emerged, where ISPs with networks that span larger geographic areas provide commercial IP transit services for locally operating ISPs. The largest ISPs, that do not purchase IP transit services from any other ISP are referred to as Tier-1 network providers. Well-known Tier-1 providers include CenturyLink, Cogent Communications and AT&T.

The Internet interconnection topology at the end of the commercialization phase consists of the ISPs' networks, or ASes, and the IXP locations where they interconnect. This led to a multi-tiered Internet interconnection model as shown in Figure 2.

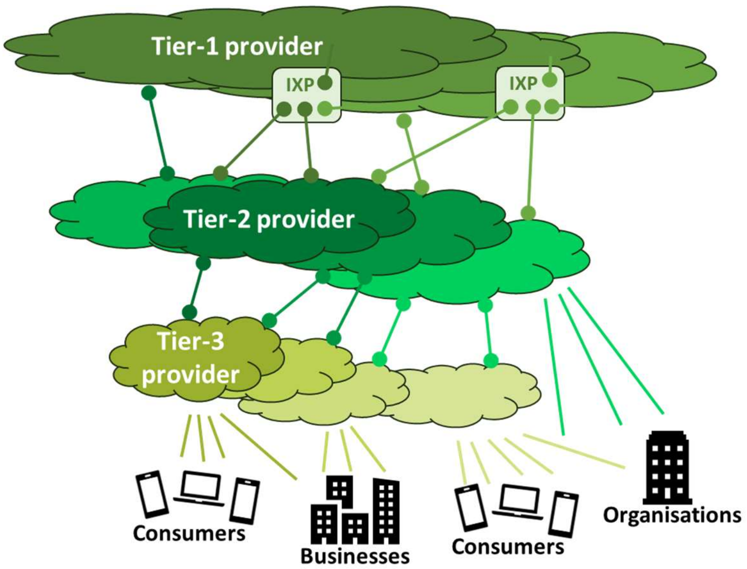

Figure 2: Internet interconnection layers

After the Internet commercialization phase a number of topology transformations occurred. These are enumerated in the bottom three rows of Table 2 and explained in more detail in the text below.

Evolution period	Drivers of dominant Internet stakeholders	Impact on Internet interconnection
1960's - 1995	Four ground rules of the Internet pioneers	Best-effort, autonomous networks interconnected by NSFnet backbone
1990's -	QoS for IPS's commercial services	Commercial walled-gardens interconnected via IXPs by 'tiered' networks and standardization of data transfer and routing protocols (BGP, DNS, HTTP)
2000's -	QoS and high bandwidth for (commercial) multi-media service providers	Commercial networks enriched with CDN (cache) solutions, increasingly higher and asymmetric bandwidths and interconnected by globally expanded IXPs
2010's -	High QoS (fast response), high reliable (and bandwidth) for business-class cloud services	Existing Internet is paralleled with global private networks that are peering via cloud exchanges. Introduction of programmable traffic engineering and direct interconnections for huge data flows

Table 2: Drivers of Internet stakeholders and impact on Internet interconnection

Transformation by multi-media service providers

CDN solutions

The introduction of multi-media services over the Internet raised the challenge of establishing stable and high throughput between the content servers and the end-users over the best-effort Internet. The dominant strategy of generic CDN (and some cloud) providers to address this challenge was to put caching content servers as close to the end-users as possible. This led to installation of CDN caches within many ISP networks[6], enabling end-user devices to download content from the fastest responding (typically the closest) CDN cache.

The impact of this strategy is shown in Figure 3. In most cases traffic be handled within one AS, which reduces Internet interconnection (e.g. Tier-1) traffic. In other cases, where traffic actually is routed over an Internet interconnection, end-users may connect to Akamai servers that are not present in an AS allocated to Akamai. This complicates the traditional AS-based view of Internet topology, where traffic between two organizations is routed over the interconnection between their ASes, has become outdated (Nikolaos Chatzis, 2013).

[6] In this context those networks are called eyeball networks that serve many end-customers, which are mainly viewing / downloading content. As a result such eyeball networks typically have asymmetric traffic loads with high bandwidth demand for downloading and less bandwidth for uploading.

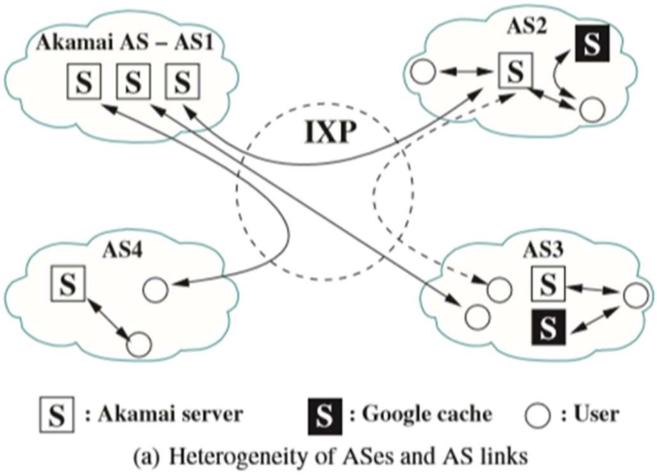

Figure 3: Heterogeneity of ASes and interconnections, source: (Nikolaos Chatzis, 2013)

Globally expanded IXPs

The rise of the CDN providers increased the need for peering between them and ISPs. Since the IXPs are the default locations for peering their size and numbers started to grow. IXPs have built new facilities and also started to attract smaller networks in distant geographic locations via remote peering providers. The growing number of IXPs provide more locations for peering between regional (Tier-2 and Tier-3) networks, reducing the necessity for higher-tiered interconnections. In combination with decreasing Tier-1 interconnection traffic due to CDN caching, this led to the "flattening of the Internet topology" that was first reported by (Phillipa Gill, 2008).

The impact of the expansion of IXPs was studied by (Timm Böttger, July 2019) using historical datasets on IXPs and their memberships, complemented with traceroute datasets. The study reports a threefold increase in the number of IXPs and members therein, in the period between 2008 and 2016. Similar to other studies they identified changes in Internet interconnections, including a decreasingly dominant role of Tier-1 providers.

Transformation by cloud and social media providers

Global private networks

As described in Section 2 the global private networks of the Internet giants are likely to be comparable to the size of the largest networks (ASes) in the public Internet. This conjecture has been addressed by a recent study of the evolution of content providers networks (Esteban Carisimo, 2019). They propose the so-called k-core metric for comparing the relevance of an individual network's contribution to Internet interconnectivity. The ASes with the highest k-core value, referred to as TOPcores, are the highest connected networks in the Internet interconnection topology: the spiders in the web.

Analysis of publicly available AS topology data has produced the k-cores of seven Internet giants and 10 transit network providers over the period between 1999 and 2018. The normalized k-core results for the ASes of the seven Internet giants (i.e. the k-core value of the AS divided by the largest k-core of any AS) are presented in Figure 4. A k-core value of 1.0 means that the AS has reached TOPcore, i.e. there is no other AS in the Internet that is more widely interconnected.

The k-cores figure shows the evolution phase of global content and cloud service providers. It shows that each of the Internet giants has reached TOPcore, some of them earlier than others. In October 2017 each of the content providers had reached the TOPcore. Several Tier-1 providers also reached that level (e.g. CenturyLink, Cogent and NTT), but other providers did not reach that level of connectivity (e.g. GTT, PCCW Global, TATA Comm.). This analysis proves that the networks of the Internet giants have entered the select group of networks that are most relevant w.r.t. Internet interconnectivity.

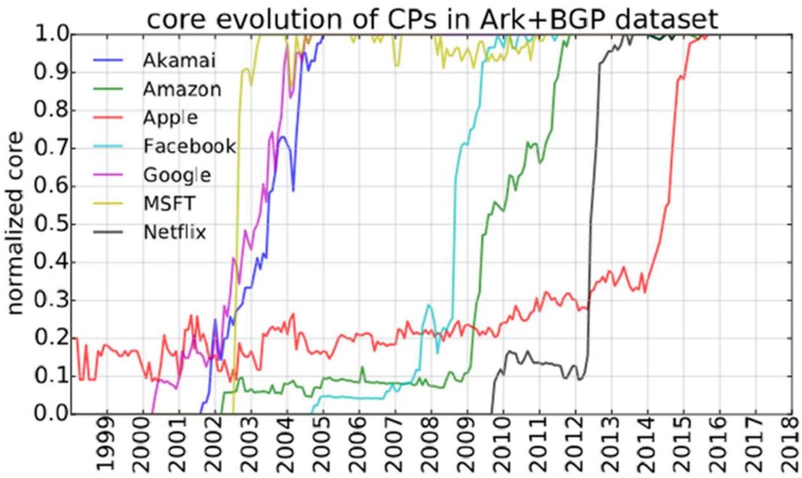

Figure 4: K-core evolution of Internet giants, source: (Esteban Carisimo, 2019)

The figure also illustrates the two distinguishable phases of the content delivery phase (starting just after the burst of the Internet hype) and the cloud and social media phase (starting with the rise of Facebook). Another interesting observation is the launch of Netflix around 2010. Initially Netflix used generic CDNs (such as Akamai) to obtain global reach. This enabled fast service deployment and expansion with minimal upfront investment cost. Once Netflix's business and traffic volumes were boosted the balance in the benefits of hiring capacity from a generic CDN and owning a global private network shifted towards the latter. In June 2012 Netflix announced the roll-out of their own Open Connect platform (Techcrunch, 5 June 2012). The result was a fast increase in interconnectivity up to TOPcore.

Cloud exchanges

The traditional way of Internet interconnecting to a cloud service provider is via public peering at an IXP. However, large cloud providers have been pushing new types of Internet connections to ensure that traffic between their customers and their private networks (a) need not be carried over the best-effort Internet and (b) can be configured as dynamic and flexible as the configuration of cloud services. Therefore, in addition to public peering, they started offing private interconnection as explained in Section 2. This requires new highly flexible cloud exchange technology preferably co-located at IXP locations; such locations are referred to as co-locations (or colo's). In 2018 Equinix was the world's largest retail colocation data center operator, with over 220 sites spread over all continents (Cloudscene, 2019). Also, Tier-1 providers such as CenturyLink and NTT Communications are increasingly offering colo facilities from their globally deployed data centers, complementary to their traditional Tier-1 services. By exploiting (and stimulating) the global proliferation of combined IXP and co-locations via partners, the cloud providers bring private interconnection via shared, isolated cable infrastructure to their networks in close proximity of their customers (Bahador Yeganeh, October 2019). This aligned business incentive between CSPs, cloud exchanges and Tier-1 providers has accelerated the proliferation of Internet and cloud exchange co-locations.

At the current phase of the evolution, Internet interconnection can be presented as in Figure 5. It reflects that private networks of the internet giants are no longer negligible relative to the largest telecommunication providers networks that form the public Internet. Further note, that the architectural separation between content and data processing on the one hand and the network infrastructure on the other hand, has become blurred. Both by introducing CDN caches into networks as well as by the fact that the CSP networks are designed to create a globally distributed data centers.

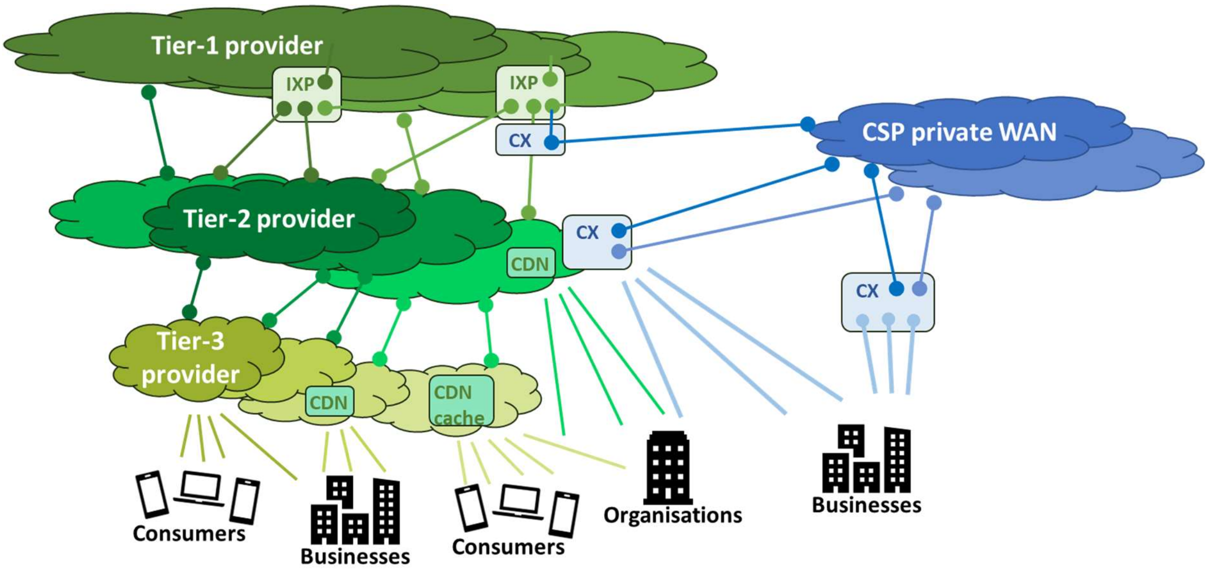

Figure 5: Public Internet complemented with global private networks

Given the novel possibility for business enterprises to transfer traffic via private networks instead of via the public Internet, the question rises to which extend this is happening. Measurements at one of the largest co-located IXP and cloud exchanges indicate that connectivity to Amazon's network consists for 8% of connections that the authors label as 'bypassing the public Internet' (Bahador Yeganeh, October 2019). This insight provides an indication that the Internet may be bypassed, but does not indicate the amount of traffic nor what type of traffic it is. For example, is the traffic transferred between the CSP and its customer, which could be regarded as customer-provider traffic for which the use of private connections has not been unusual even before the cloud era. Or, is it peering traffic used to transfer data between customers, which would otherwise have been transferred over the public Internet (in this case we can speak of bypassing the public Internet). In any case, since the traffic is handled within private WANs and with novel technology ('routed' by programmable networking technology complementary to traditional routing technology) it will remain hard to get more insight with current Internet measurement instrumentation.

4. Lessons learned: Internet interconnection policy

In each Internet evolution phase its interconnection topology has been shaped by the drivers of the dominant stakeholders in the respective phase. From the open Internet based on the four ground rules of the academic pioneers, to the commercial, tiered networking providers followed by the global content generating and distributing providers. With the exception of the academic inception phase, the dominant stakeholders have been the parties that were most successful in commercial exploitation of the global Internet infrastructure.

While from an architectural point of view the Internet infrastructure is decoupled from the content that it transports, the evolution indicates otherwise. Content characteristics such as quality of service, high-bandwidth content and reliable cloud for business, have driven their providers to optimize the Internet infrastructure to fit their use. For policymakers it is important to keep this long-term causality in mind: long-term Internet infrastructure governance is effectuated by influencing the drivers of the key stakeholders.

The current state of the Internet interconnection evolution is clearly challenging the 'open Internet' principle and the ground rules that it is based on. Although the public, open Internet is still the dominant infrastructure it is undeniable that 'private Internet' infrastructures and their scope of application are no longer marginal relative to the public Internet. In itself this is no problem, but it can become another threat to the open Internet. For example if private Internets will be used as alternative transit networks or for other services that are currently offered via the public Internet. Once public Internet services start to be offered via private networks this may impact the scope and effectiveness of, for example, regulatory instruments such as the telecom act. While current Internet governance debate is focusing on the public Internet it is important that policymakers recognize this when referring to 'the' Internet.

While the commercial stakeholders have been the key players in actually building and evolving the Internet interconnection infrastructure, other stakeholders have influenced its evolution too. Governments around the globe have increasingly voiced their Internet governance policies. These include calls for international collaboration and investment to ensure a free, open and secure cyberspace by the Global Forum on Cyber Expertise (GFCE) and the Organization for Economic Cooperation and Development (OECD). At first glance these calls seem inconsistent to regulatory measures that governments are taking to safeguard digital infrastructure and national cyber security. For example, in the Dutch Cybersecurity Agenda 2018 a warning is made regarding the dependency on a limited number of foreign digital infrastructure service providers. An extended initiative is GAIA-X that aims to develop a European data infrastructure. Although such an initiative can be aligned with the principles of the 'open Internet', they can easily be perceived as Internet restrictive policies. Policymakers should be aware that transparent and consistent governance are fundamental for the continuation of the economic and social benefits that the Internet has brought, as also pointed out by the OECD (OECD, 2016).

In conclusion, policymakers should remind that (ISOC, 2020) "the success of the Internet is rooted in the way it was built and able to grow: an open platform for innovation and sharing of ideas [and] the Internet cannot be regulated in a top-down manner, but its governance should be based on processes that are inclusive and driven by consensus." Consensus requires that policymakers balance and align their own interest with the interest of other Internet stakeholders, which requires dialog and coordination. From the recent developments in the Internet interconnection evolution we learn that keeping an open mind is essential, because otherwise even an infrastructure as large as the public Internet can and will be paralleled and loose its collective force of innovation.

5. References

Bahador Yeganeh, October 2019. *How Cloud Traffic Goes Hiding: A Study of Amazon's Peering Fabric.* Amsterdam, IMC '19.

Cloudscene, 2019. *H1, 2019 Data Center Ecosystem Leaderboard.* https://cloudscene.com/news/2019/08/h1-data-center-leaders/

Esteban Carisimo, 2019. Studying the Evolution of Content Providers in IPv4 and IPv6 Internet Cores. *Elsevier Computer Networks.*

ICANN, 2013. *One World, One Internet.* www.icann.org/en/system/files/files/ecosystem-06feb13-en.pdf

Internet Governance Forum, 2019. *WS 59 Digital Sovereignty and Internet Fragmentation.* https://igf2019.sched.com/event/SU2l/ws-59-digital-sovereignty-and-internet-fragmentation

ISOC, 1997. *Brief History of the Internet.* www.internetsociety.org/resources/doc/2017/brief-history-internet/

ISOC, 2020. *Internet Governance.* https://www.internetsociety.org/issues/internet-governance/

Nikolaos Chatzis, 2013. *On the Benefits of Using a Large IXP as an Internet Vantage Point.* Barcelona, Spain, ACM.

OECD, 2016. *Economic and Social Benefits of Internet Openness*.

Phillipa Gill, 2008. The flattening Internet topology: Natural evolution, unsightly barnacles or contrived collapse?. *Passive and Active Network Measurement (PAM)*.

Sushant Jain, 2013. *B4: Experience with a Globally-Deployed Software Defined WAN.* ACM Sigcomm.

Techcrunch, 5 June 2012. *Netflix Rolls Out Its Own CDN: Open Connect.* https://techcrunch.com/2012/06/04/netflix-open-connect

Teleography, 2019. *The State of the Network TeleGeography,* Teleography.

Timm Böttger, July 2019. *Shaping the Internet: 10 Years of IXP Growth.*

United Nations, 2019. *Countering the use of information and communications technologies for criminal purposes*.
https://undocs.org/pdf?symbol=en/A/C.3/74/L.11/Rev.1