CONSULTING

ANALYSYS MASON REPORT

INFRASTRUCTURE INVESTMENT BY ONLINE SERVICE PROVIDERS

David Abecassis, Richard Morgan, Shahan Osman

DECEMBER 2018

analysysmason.com

Contents

1 Executive summary 1

2 Introduction 8

3 Evolving internet usage is driving requirements for infrastructure investment 9

3.1 Internet users are growing steadily, consuming richer content and new services 9

3.2 OSPs' increasing focus on cloud services is placing new demands on their infrastructure 11

3.3 Growing investment in infrastructure supports richer content and enables new services 12

4 OSPs' investment in infrastructure has exceeded USD75 billion per year on average since 2014 14

4.1 Investment has grown across all clusters, with hosting and delivery growing especially quickly due to growth in demand for cloud services and rich content 14

4.2 Investment has grown strongly across all world regions, with North America seeing the largest share 15

4.3 The majority of OSP spend is direct investment in self-owned infrastructure, however indirect investment has also grown at a rapid rate 17

5 OSPs are making significant investments to extend their networks, moving ever closer to end users 20

5.1 OSPs are investing in new hyperscale data centres, while also greatly increasing their spend on colocation space 20

5.2 Investment in data transport has grown apace, as OSPs carry content and enterprise data between their global data centres 28

5.3 OSPs continue to develop delivery networks to bring services closer to end users 34

6 OSPs are investing in innovation related to infrastructure 40

6.1 OSPs are investing heavily in innovation to improve data-centre efficiency 40

6.2 OSPs have taken a leading role in the push towards using renewable energy 44

6.3 OSPs are also actively exploring ways to improve delivery and access networks, particularly in rural regions across the globe 46

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

DISCLAIMER AND ACKNOWLEDGEMENTS

This report was commissioned and sponsored by Google, and written by Analysys Mason, a global consultancy specialising in telecoms, media and technology.

Analysys Mason is grateful for all inputs received during the research and preparation of the report, including discussions with a number of organisations including Akamai, AMS-IX, ChinaCache, Cisco, Equinix, Facebook, Google, LINX, TeleGeography and Zayo.

The analysis contained in this document is the sole responsibility of Analysys Mason and does not necessarily reflect the views of Google or other contributors to the research.

1 Executive summary

In the four years since 2014, online service providers (OSPs) have invested over USD300 billion in internet infrastructure. This amounts to USD75 billion per year, which is more than double the 2011–13 average annual investment of USD33 billion. Over 90% of this has been in hosting infrastructure as OSPs build hyperscale data centres to support the explosion in online content and cloud services, and install equipment in third-party colocation facilities. They are also investing large sums in infrastructure to transport data between these locations, including terrestrial fibre networks and international submarine cables. Finally, OSPs are driving investment in delivery networks, to support quality of service by bringing content as close as possible to end users.

Evolving internet usage is driving requirements for infrastructure investment

Demand for online services has continued to see tremendous growth in recent years, driven by new internet users, higher consumption of rich content (notably video), and the emergence of cloud services. To cater for this growth, OSPs are investing substantially in the infrastructure required to host and carry content, and make it accessible to end users.

This report considers investment in the three ‘clusters’ of infrastructure shown below.

Figure 1.1: Overview of clusters of infrastructure investment [Source: Analysys Mason, 2018]

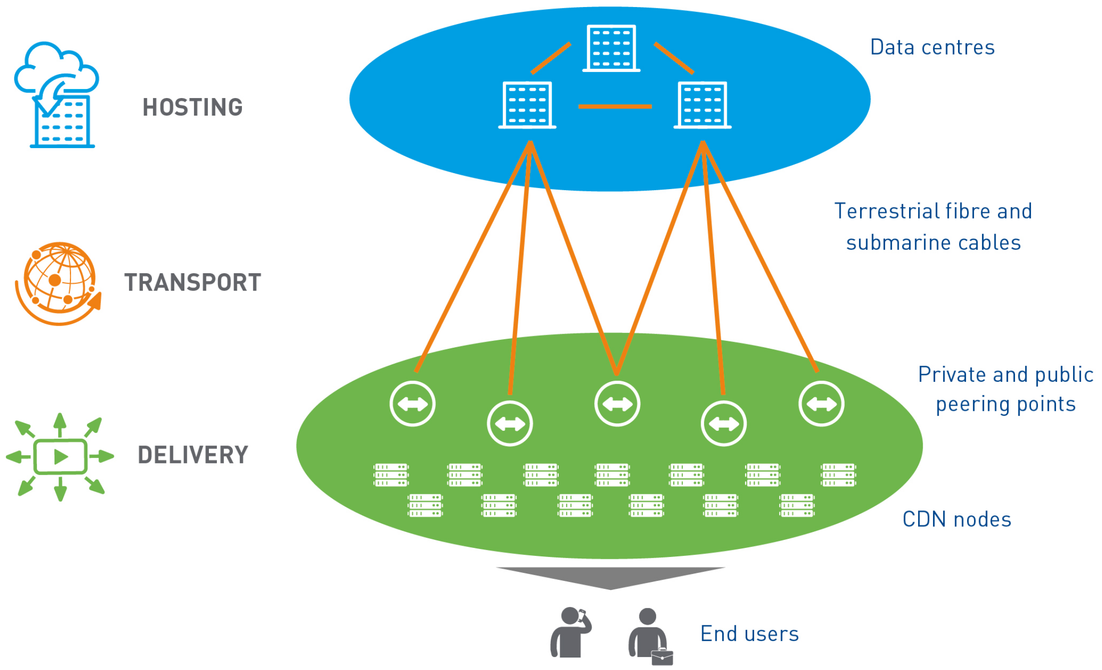

OSPs' investment in infrastructure has exceeded USD75 billion per year on average since 2014

Infrastructure spend over the past four years is more than double the annual investment seen in our previous study, which assessed spending by OSPs from 2011 to 2013.[1] Investment has grown across all three clusters, but with especially high growth in relation to hosting and delivery infrastructure.

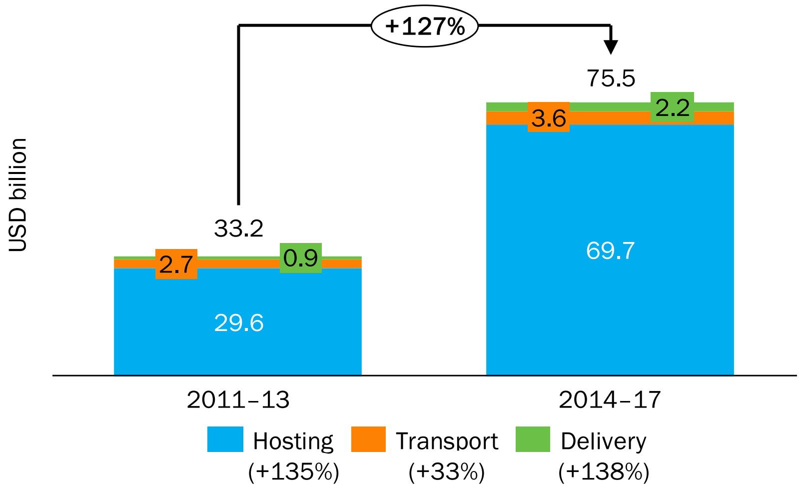

Figure 1.2: Average annual investment by OSPs[2] [Source: Analysys Mason, based on various sources, 2018]

A more advanced transition to cloud computing in North America has led to greater investment in data centres, with around half of the world's hyperscale facilities based in the region. This has resulted in North America accounting for the largest regional share of OSP investment, followed by Asia–Pacific and Europe.

1 An overview of the previous study and a link to the full report can be found here: http://www.analysysmason.com/About-Us/News/Newsletter/Internet-infrastructure-investment-Oct2014/

2 Delivery investment for 2011 to 2013 has been restated at a slightly higher level than in the previous report due to new data on CDN and terrestrial fibre investment

Figure 1.3: Average annual investment by region[3] [Source: Analysys Mason, based on various sources, 2018]

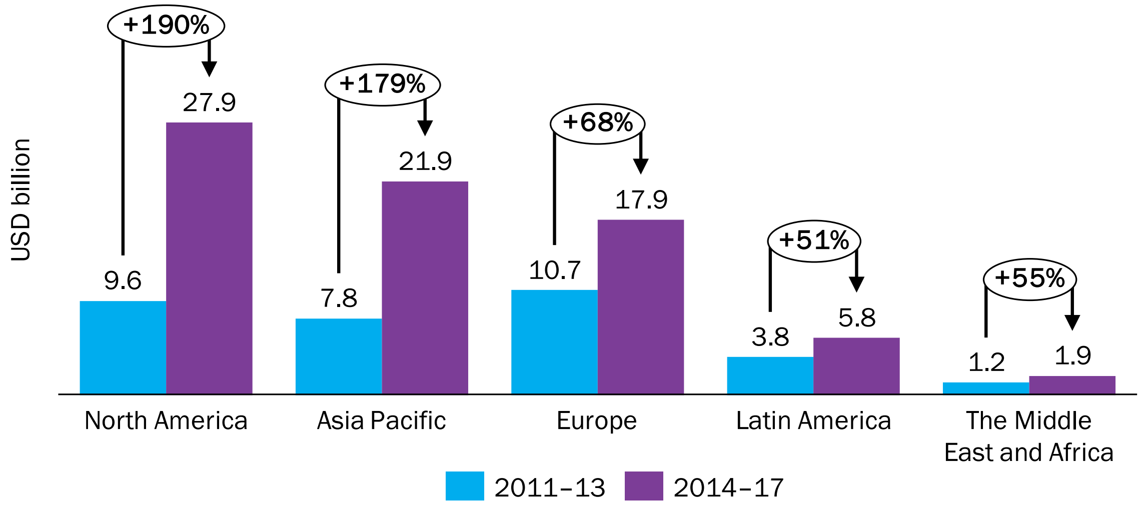

As OSPs' requirements grow, they expand the scope of their infrastructure so as to optimise quality of service and control costs. This has led to significant growth in direct investment in self-owned infrastructure such as hyperscale data centres, fibre transmission equipment and submarine cable systems, private peering infrastructure and in-house CDNs. This direct investment accounted for 78% of investment in 2014–17.

OSPs are also indirectly funding investment in infrastructure by third parties, through their large-scale purchases of services such as data-centre colocation space or capacity on fibre infrastructure. This report assesses this indirect investment made by OSPs, as well as their own direct investment in infrastructure.

OSPs are making significant investments to extend their networks, moving ever closer to end users

A number of trends are driving OSP investment in hosting, transport and delivery.

Hosting

The push of OSPs into the cloud services market has accelerated the migration of enterprise IT from private, on-premises infrastructure to public cloud services. This trend, along with the proliferation of user-generated content, has driven demand for hosting infrastructure.

The larger OSPs have invested heavily in building hyperscale data-centre facilities, in a drive to improve power and cost efficiency, reliability and performance. While around half of hyperscale facilities are presently located in the USA, OSPs have expanded into other markets where demand and market conditions are suitable.

OSPs are also greatly increasing their spend on colocation data-centre facilities from third-party providers, which are typically used for smaller 'edge' data centres or where investment in private hyperscale facilities is not justified. OSPs are thereby indirectly

3 Regional split of investment from 2011 to 2013 has been restated slightly from the previous study due to new data available

funding the development of new colocation facilities, and are also making substantial investments in the equipment installed within these data centres.

Transport

Investment in data transport has grown apace, as OSPs carry content and enterprise data between their global data centres. A large proportion of this investment is in terrestrial fibre networks within countries, where OSPs typically lease access to third-party fibre assets and then invest in their own network equipment to transport data over this fibre.

OSPs are accounting for a growing share of total international bandwidth, and are thereby funding a large proportion of the global investment in submarine cable systems. Most of this investment remains indirect: OSPs buy capacity on existing cable systems. There is however much faster growth in direct investments by the largest OSPs, which are deploying new submarine cables, either as part of consortium or as private investments. Furthermore, OSPs' investments in cable systems are supporting the wider communications sector, as it is common for spare capacity on these cables to be leased to third parties.

Figure 1.4: Map of expected submarine cables with OSP investment by 2020 [Source: Analysys Mason based on information from TeleGeography, news articles, company press releases, 2018]

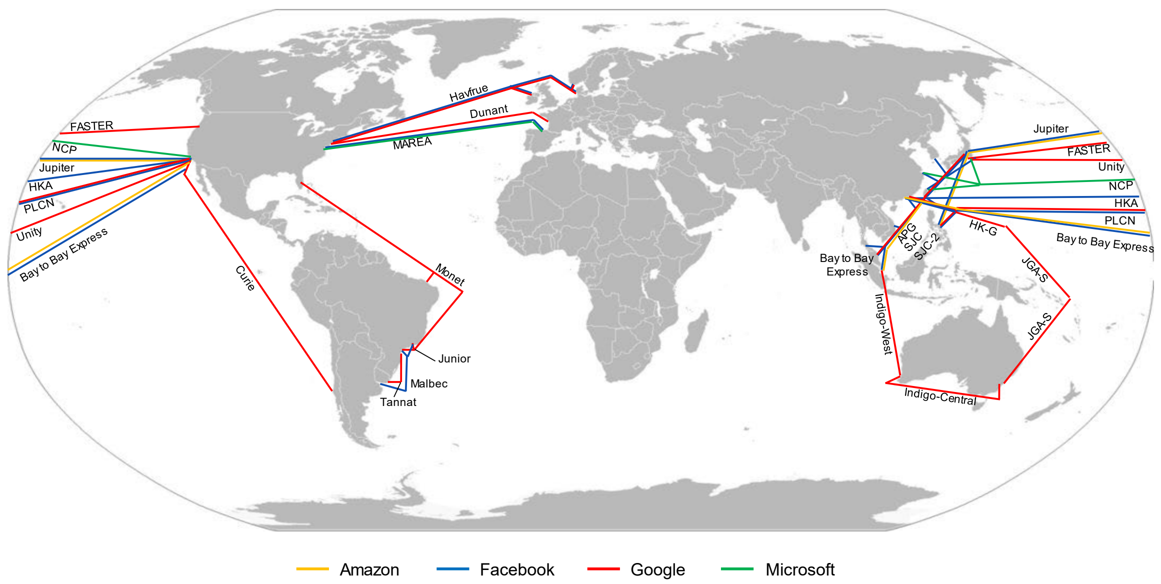

Delivery

OSPs need to deliver content to end users via internet service providers (ISPs), and they therefore invest in peering infrastructure to support the exchange of traffic with ISPs. This can either be through public peering at internet exchange points (IXPs), or can be private peering between two parties. The number of public peering points used by OSPs has grown significantly, but much larger volumes of traffic are exchanged through private peering.

OSPs are also investing in CDN infrastructure to cache content closer to end users, in order to improve service quality and reliability, and control costs by minimising the duplication of content that is transported over backbone networks to reach end users. CDNs have grown in importance due to the rise in video content and cloud services, and large OSPs are investing in deploying their own 'in-house' CDNs.

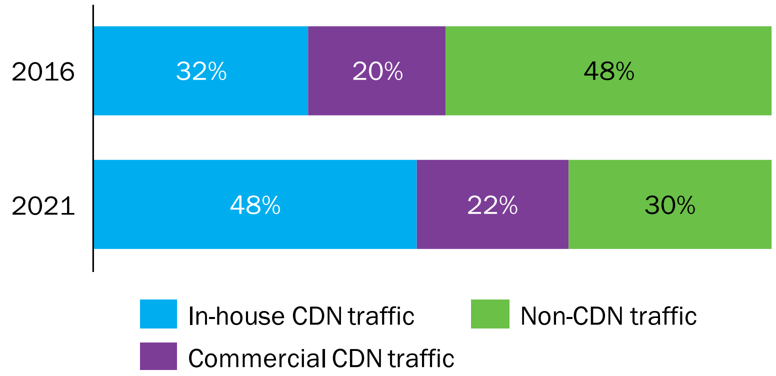

Figure 1.5: Internet traffic by CDN use [Source: Analysys Mason based on Cisco Visual Networking Index, 2016–21]

OSPs are also investing in innovation related to infrastructure

In addition to ongoing investments in infrastructure, OSPs also continue to dedicate billions of dollars each year on research and development (R&D) to improve their products. The total R&D spend for nine of the major OSPs more than doubled from 2013 to 2017, as companies continue to invest in innovative technology, much of which is designed to improve the efficiency of their future infrastructure deployments and operations.[4]

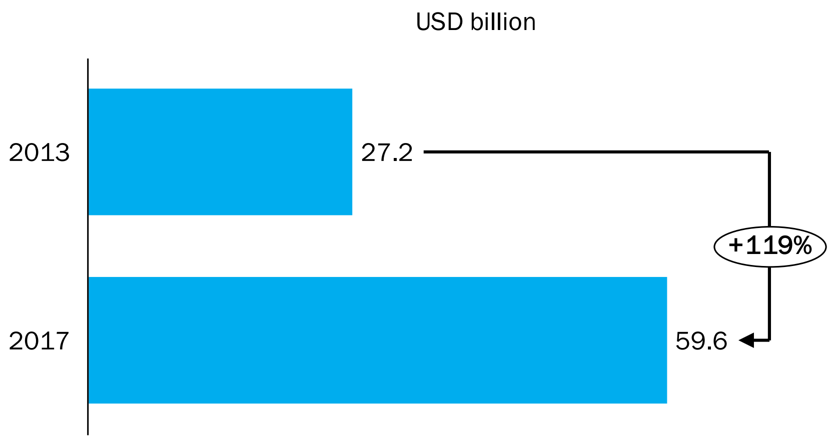

Figure 1.6: Annualised R&D spend for nine main OSPs in 2013 and 2017 [Source: Company financials, Analysys Mason, 2018]

OSPs are investing heavily in innovation to improve data-centre efficiency, and are making significant advances in data-centre hardware and computing power. They are also driving innovation in cooling systems and associated technology to improve power efficiency in data centres. OSPs have taken a leading role in the push towards using renewable energy, with many large players committing to

4 OSPs for which R&D was tracked include Google, Microsoft, Apple, Facebook, Alibaba, Tencent, Baidu, eBay, and Netflix. Facebook, Microsoft, Tencent, Alibaba, Baidu, eBay, Netflix, and Apple. Amazon has been excluded as the company does not report R&D separately from total "technology and content" opex

eliminating all use of non-renewable power. Finally, OSPs are also actively exploring ways to improve delivery and access networks, particularly in rural regions across the globe.

These investments by OSPs have had an impact not only on improving online services, but also on job markets and local economies. Google's data centres were estimated to support 11 000 jobs in the United States in 2016,[5] and 6600 jobs per year on average from 2007–17 in Europe.[6] Facebook's data centres in the United States meanwhile, contributed 8600 jobs per year from 2010–16.[7]

Conclusion

Investment by OSPs has grown significantly over the past four years, and this upward trajectory is set to continue as demand for online services grows. OSPs are increasingly making direct investments in infrastructure to provide more control over service quality and reliability. While data centres continue to require the largest share of investment, OSPs are investing heavily in transport and delivery networks to bring content as close as possible to end users.

It is clear that OSPs are not simply providing content and services using third-party networks and facilities, but are making a large and growing contribution to the infrastructure that underpins the Internet. OSP investments are typically made in parallel with investments from a variety of other stakeholders in the global internet landscape. Revenues from OSPs' activities support investments by players including telecoms carriers, data-centre operators and internet service providers, providing benefits to the whole ecosystem. Further growth in investment, both directly by OSPs and indirectly through other service providers, can be expected in the coming years, as OSPs strive to keep pace with the growing demand for content and cloud services across all regions of the globe.

---

5 Oxford Economics. (2018), *Google Data Centres: Economic Impact and Community Benefit.* Available at https://www.oxfordeconomics.com/recent-releases/d8d830e4-6327-460e-95a5-c695a32916d9

6 Copenhagen Economics. (2018), European data centres: How Google's digital infrastructure is supporting sustainable growth in Europe. Available at https://www.copenhageneconomics.com/publications/publication/european-data-centres

7 RTI International. (2018), The Impact of Facebook's U.S. Data Centre Fleet. Available at https://www.rti.org/sites/default/files/facebook_data_centers_2018.pdf

INVESTMENT IN INFRASTRUCTURE

Online service providers support around USD75 billion annual investment in infrastructure worldwide

Drivers of infrastructure investment

	ONLINE USERS	RICH CONTENT	CLOUD SERVICES
In 2017:	3.6 billion internet users[1]	76.4PB of monthly consumer internet traffic[2]	262 million cloud workloads and computes[3]
Since 2013:	+8% p.a.	+27% p.a.	+46% p.a.

Average annual total investment by cluster (USD billion)

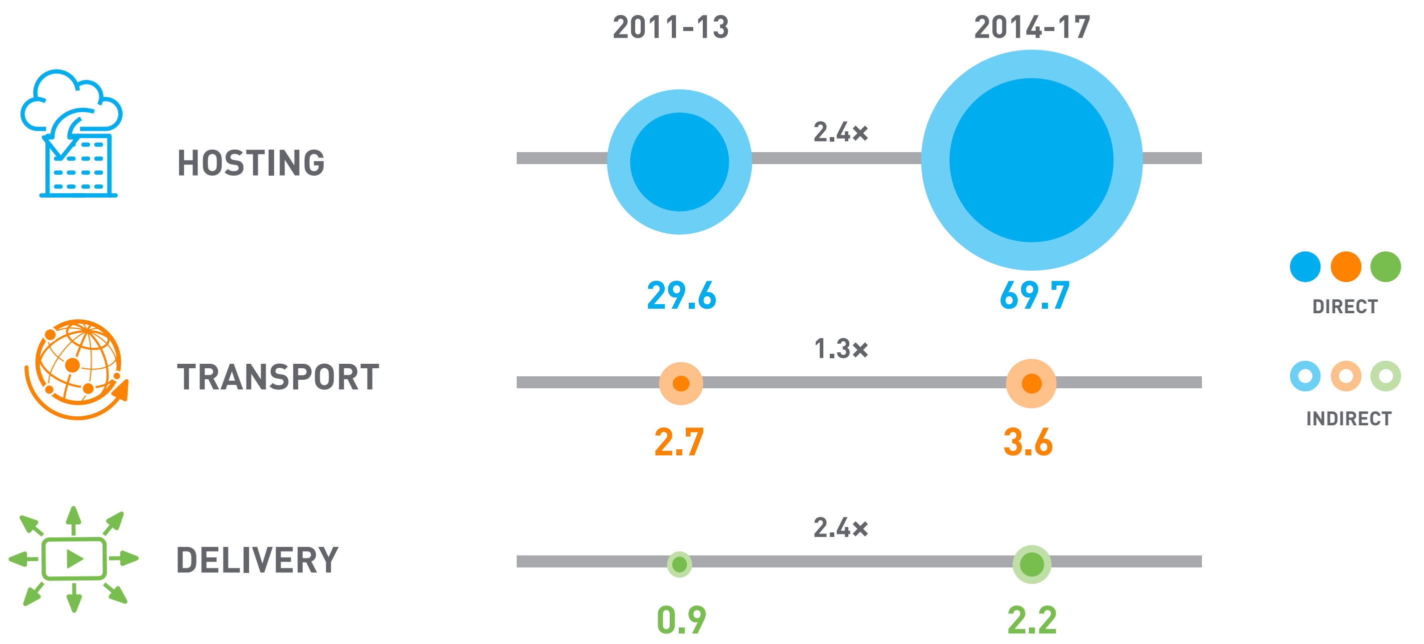

Average annual total investment by region (USD billion, 2014-17)

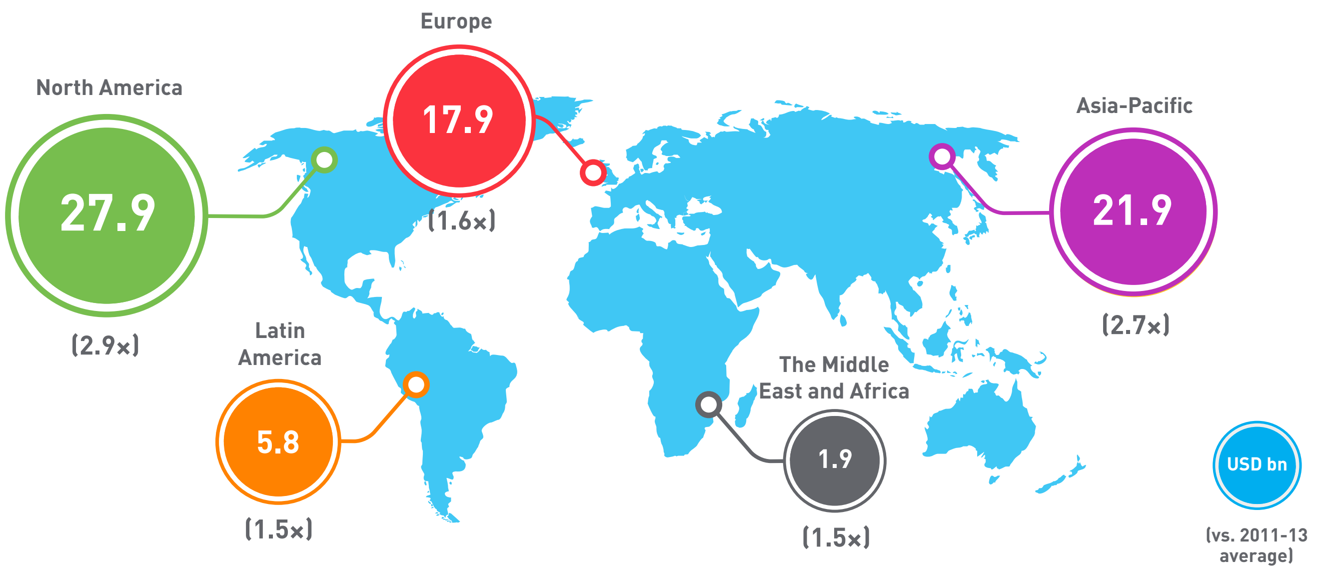

[1] Internet users: International Telecommunication Union, United Nations
[2] Consumer internet traffic: Cisco Visual Networking Index 2016-2021 and earlier editions; 2017 figure is a Cisco forecast
[3] Cloud workloads and compute instances: Cisco Global Cloud Index 2016-2021 and earlier editions; 2017 figure is a Cisco forecast

2 Introduction

The internet is becoming increasingly central to how people around the world live and work, and massive investment in infrastructure is required to keep up with demand for online services. 'Online service providers' (OSPs) represent a diverse range of businesses, but all provide platforms over which content is delivered and shared over the internet. OSPs have traditionally been thought of as software companies, but with exponential growth in online content, they are making substantial investments in the infrastructure required to host and carry content, and make it accessible to end users.

In 2014, Analysys Mason conducted a study[8] evaluating and quantifying the major investments that OSPs make in physical infrastructure such as data centres, submarine cables, and content delivery networks (CDNs). The study estimated the global investment from 2011–13 to amount to over USD30 billion annually.[9] The present report provides an update focussing on developments in infrastructure investment by OSPs over the period 2014–17.

This study has benefitted from discussions with a range of stakeholders in the OSP, data-centre, internet exchange, backbone and submarine cable communities. These discussions have helped to identify recent trends and upcoming initiatives, and have informed the modelling used to quantify investment in infrastructure on a global and regional basis.

The report contains a number of case studies of specific investments or perspectives from other industry players, which have focused on three regions: Asia–Pacific, Europe, and Latin America. The remainder of this document is laid out as follows:

- Section 3 describes how changes in the demand for online services affect infrastructure investment
- Section 4 presents our estimates of the magnitude of these investments for the period 2014–17
- Section 5 details the nature of investments and the trends driving them
- Section 6 highlights innovations driven by OSPs to improve the infrastructure of the internet.

8 The 2014 study, like this one, was sponsored by Google but carried out independently by an Analysys Mason team

9 An overview of the study and a link to the full report can be found here: http://www.analysysmason.com/About-Us/News/Newsletter/Internet-infrastructure-investment-Oct2014/

3 Evolving internet usage is driving requirements for infrastructure investment

Over the past few years, demand for online services has shown tremendous growth, driven by new internet users, higher consumption of video content, and the emergence of cloud services. To cater to this growth, OSPs are investing substantially in physical and IT infrastructure.

3.1 Internet users are growing steadily, consuming richer content and new services

Thousands of new internet users are coming online each week, as fixed and mobile broadband networks expand their reach and services become more widely affordable. There were an estimated 3.6 billion internet users in 2017, a growth of 36% from the previous study in 2013.[10] Much of this growth has come from developing markets, and OSPs are investing in infrastructure in new regions to meet this demand.

Figure 3.1: Global internet users [Source: Analysys Mason based on International Telecommunication Union, United Nations, 2018]

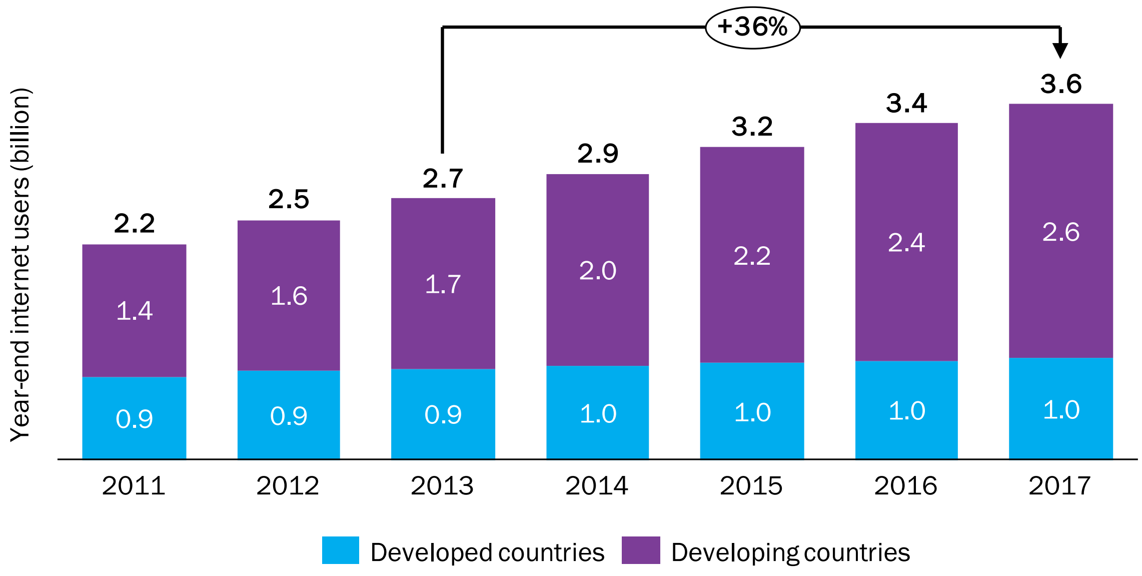

The demands from each internet user are also increasing: the average user is spending more time online,[11] and increasingly consuming and uploading richer content. The ongoing shift away from traditional 'linear' TV broadcasting towards new online channels has contributed to growth in online video, which has put particular pressure on internet infrastructure. As of the first half of 2018, video

10 Based on data from the ITU World Telecommunication/ICT Indicators database, as well as United Nations Population Database and M49 standard

11 See: https://www.recode.net/2018/6/8/17441288/internet-time-spent-tv-zenith-data-media

accounts for 58% of all downstream traffic on the internet, according to Sandvine.[12] OSPs need to invest in infrastructure to ensure a good user experience, and this becomes even more important when users are paying for services. Netflix, for example, reported that it delivered around four billion hours of video per month in 2017.[13] According to Sandvine, Netflix accounts for 15% of global downstream traffic as of the first half of 2018.[14]

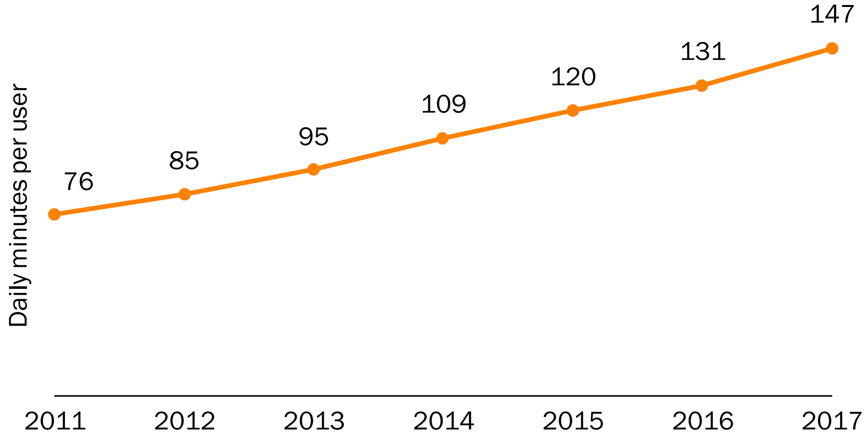

Figure 3.2: Daily minutes spent with internet media [Source: Zenith Media Consumption Forecasts, 2018]

Video content watched online is also increasingly generated by users, which has associated challenges for infrastructure regarding latency and reliability when supporting two-way traffic flows. Twitch, for example, is a live streaming platform that benefits from low latency for real-time user interaction,[15] and the platform has grown to see half a billion hours watched each month in 2017. The expansion of online gaming is also driving infrastructure requirements, due to its high bandwidth and latency-sensitive nature.

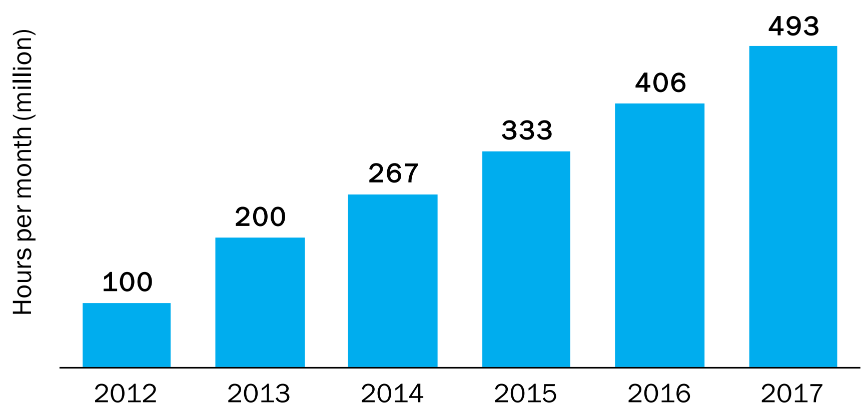

Figure 3.3: Evolution of total monthly hours watched on Twitch [Source: TwitchTracker,[16] 2018]

12 Sandvine. (2018), Sandvine releases 2018 Global Internet Phenomena Report. Available at: https://www.sandvine.com/press-releases/sandvine-releases-2018-global-internet-phenomena-report

13 Netflix Media Centre. (2017), *2017 on Netflix – A Year in Bingeing.* Available at https://media.netflix.com/en/press-releases/2017-on-netflix-a-year-in-bingeing

14 ibid.

15 These include both well-established interactions such as chatrooms, and more ambitious, relatively more latency-sensitive innovations such as "Twitch Plays", where the game is controlled by chatroom participants

16 See: https://twitchtracker.com/

3.2 OSPs' increasing focus on cloud services is placing new demands on their infrastructure

Another major trend driving demand for internet infrastructure is the rapid growth of cloud services. Enterprises are progressively moving to cloud-based IT services in substitution of on-premises hardware and software, while individuals are benefitting from a range of cloud services including online storage solutions such as Dropbox.

Our 2014 study highlighted that OSPs have been making large investments in infrastructure for years, in order to expand and optimise their data-centre infrastructure to cope with rapidly increasing demand for content. This has resulted in the construction and operation of some of the world's largest and most cost-efficient data centres. In recent years, OSPs have increasingly sold access to this infrastructure to other organisations in the form of cloud services, with enterprises moving their IT workloads from their own infrastructure to the cloud, as shown below.

Figure 3.4: Cloud workload and compute instances [Source: Cisco Global Cloud Index, 2016–21, and earlier editions]

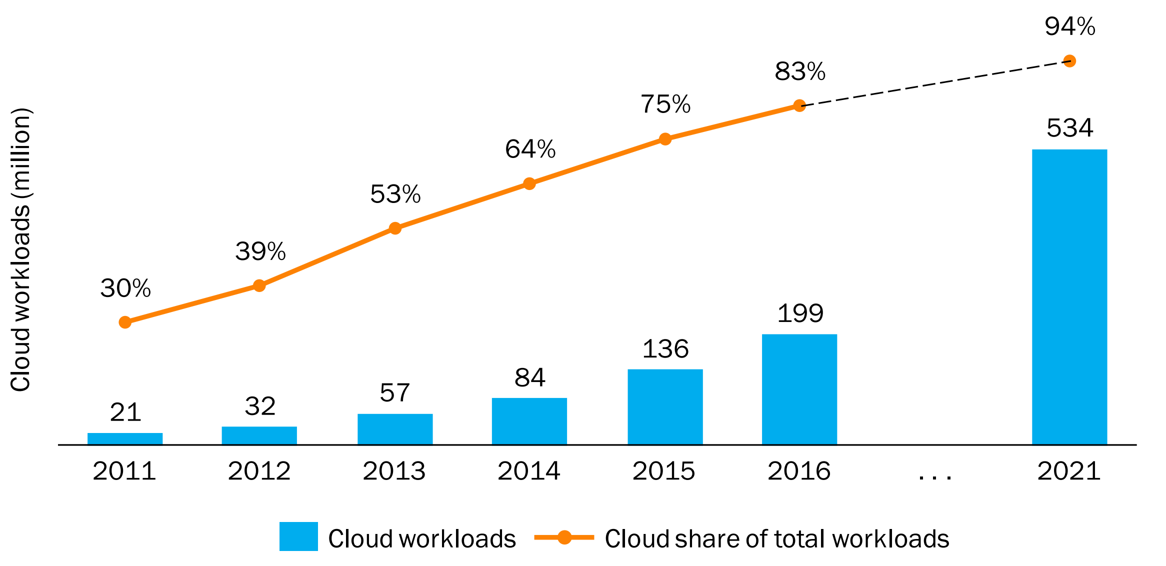

The public and hybrid cloud market,[17] where enterprises share cloud platforms with one another through a third-party provider, has been growing at a rapid pace. Accelerating growth has been supported by large investments and strategic marketing by OSPs including Amazon (through its market-leading AWS), Microsoft (Azure), and Google (Google Cloud Platform). The largest Chinese players Alibaba (AliCloud) and Tencent (Tencent Cloud) are also investing heavily, and are expanding their presence outside of the Asia–Pacific region.

[17] Public cloud refers to computing services, including data storage and processing, offered by third-party providers to any willing customer over the public internet; hybrid cloud refers to a mix of public and private cloud, the latter referring to similar services offered over either the public internet or a private internal network to only a select group of defined entities

3.3 Growing investment in infrastructure supports richer content and enables new services

OSPs infrastructure investment requirements are growing as demand for their services evolves. This report considers the three categories of infrastructure, called 'clusters' throughout this report and defined in the previous study, which remain relevant today:

- **Hosting** (data centres and cloud)
- **Transport** (fibre routes transporting content)
- **Delivery** (peering and caching infrastructure).

Figure 3.5: Overview of clusters of infrastructure investment [Source: Analysys Mason, 2018]

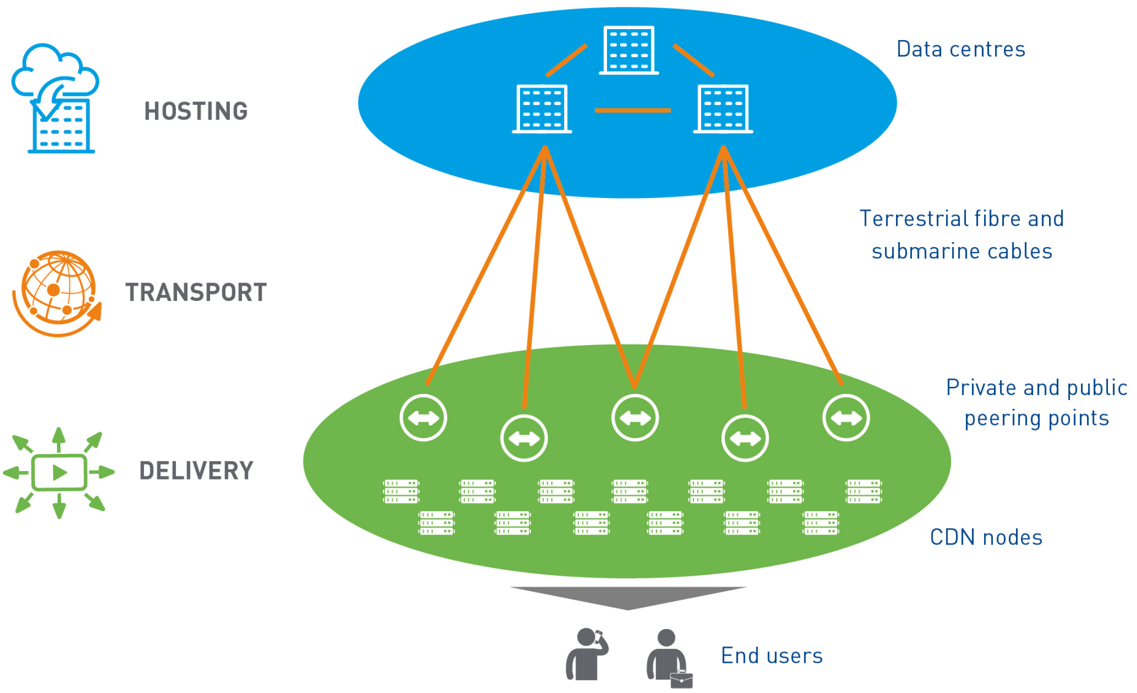

From 2014 to 2017, OSPs increased their levels of investment across all three of these areas, as discussed in detail in Section 4. Growing demand for content and cloud services has driven investment in new hyperscale data centres, as well as colocation in third-party facilities. Providing global services requires content to be transported between regions and countries, and OSPs are becoming significant buyers of international bandwidth on certain routes, while also investing in their own submarine cables. OSPs are also investing in infrastructure to support delivery of content within countries, by interconnecting with telecoms operators and other content providers in more places, and storing content closer to end users through content delivery networks (CDNs).[18] These developments are discussed in further detail in Sections 5 and 5.2.4

Investments made by OSPs are typically made in parallel with investments by a variety of other players in the global internet landscape, including telecoms carriers, data-centre operators and internet service providers. Many of the investments made by these players are indirectly funded through commercial

18 CDNs are geographically dispersed networks of servers used to store content closer to end users to help improve quality of experience

agreements with OSPs. For example, OSPs continue to lease a significant amount of space from colocation data-centre providers, such as Equinix. The investment that Equinix dedicates to build new facilities to meet demand from OSPs can be considered as indirect investment by OSPs. Partnerships between OSPs and other players in the Internet value chain are becoming increasingly common, and OSP partnerships with telecoms operators in areas beyond infrastructure has been shown to drive revenues and cost savings, supporting further network investment by operators.[19]

19 See: http://www.analysysmason.com/digital-transformation-through-partnerships

4 OSPs' investment in infrastructure has exceeded USD75 billion per year on average since 2014

Investment by OSPs has grown significantly, with annual average spend during 2014–17 reaching USD75 billion, equivalent to ~2.3 times the average annual investment between 2011–13. This OSP investment is equivalent to ~20% of the capital investments made by all of the world's telecoms operators combined, up from ~10% in the previous study. This trend is driven by OSP investment growing at over 20% per year since 2012, while the level of investment by telecoms operators has remained relatively flat over the same period.[20]

In this section, we consider the growth of OSP investment across three dimensions: infrastructure cluster, geographical region, and investment mechanism (direct expenditure by OSPs vs. indirect expenditure through commercial arrangements).

The hosting and delivery clusters have grown faster than transport, as OSPs continue to find cost-effective ways to move closer to end users, either through 'edge' data centres, peering arrangements, or caching. In the current period, North America and the Asia–Pacific have shown the strongest growth, with both overtaking the level of investment seen in Europe. Lastly, direct investment has outpaced indirect investment as OSPs seek cost reductions and more control over their networks.

4.1 Investment has grown across all clusters, with hosting and delivery growing especially quickly due to growth in demand for cloud services and rich content

Investment by OSPs has grown across all three clusters, but with especially high growth in relation to hosting and delivery infrastructure.

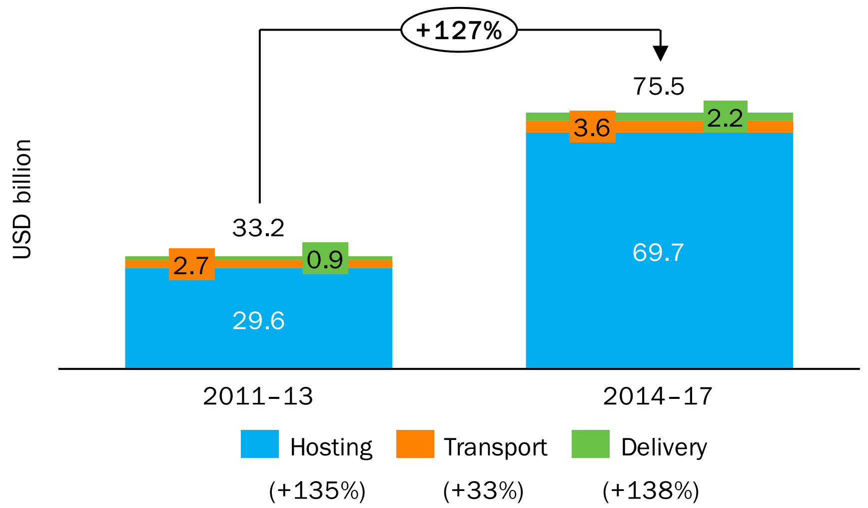

Figure 4.1: Average annual OSP investment[21] [Source: Analysys Mason, based on various sources, 2018]

20 Based on telecoms operator data from Analysys Mason Research

21 Delivery investment for 2011 to 2013 has been restated at a slightly higher level than in the previous report due to new data on CDN and terrestrial fibre investment

Data centres continue to be the most significant areas of OSP investment in internet infrastructure, with hosting accounting for an estimated ~92% of relevant investments from 2014 to 2017. As volumes of content and demand for computing power grow exponentially, OSPs have stepped up their investments in building data centres, leasing colocation space, and installing and upgrading servers and other equipment.

Investment in transport accounts for ~5% of total OSP spend over the period. The vast majority of this is indirect expenditure in the form of leases or indefeasible rights of use (IRUs)[22]; in these cases, capacity purchased by OSPs funds investment in third-party cables. At the same time, direct investment is growing rapidly; OSPs are joining consortiums and commissioning their own cable systems. OSPs' demand for international capacity has grown dramatically, but growth in transport investment has been relatively modest. This reflects the significant economies of scale that can be achieved from new high-capacity submarine cable systems, and the continued decline in lease prices.

Delivery infrastructure is the smallest cluster of investment (~3% of the total), as equipment located at the edge of networks tends to be less capital-intensive. Nonetheless, investment in delivery is a crucial component of OSPs' drive to scale their services up in a cost-efficient manner. Peering arrangements and the use of local caching can help to limit requirements for international bandwidth. The delivery cluster has shown the strongest growth of the three over the past four years, as OSPs invest in equipment for peering and local caching on ISP networks.

Recent capex figures reported in the first half of 2018 indicate that investment growth might further accelerate in coming years.[23] Likely drivers for this growth include continued take-up of rich content, expansion of cloud regions, and construction of announced data centres and submarine cables.

4.2 Investment has grown strongly across all world regions, with North America seeing the largest share

As shown in Figure 4.2 below, North America accounts for the largest regional share of OSP investment,[24] followed by Asia–Pacific and Europe. This is largely driven by hyperscale[25] data centres, around half of which have been built in North America. This large installed base of facilities means that much of the investment in the region relates to 'refresh capex' to replace and upgrade equipment installed in earlier periods, while other regions are seeing a larger share of investment going towards new facilities. Investment is also being driven by other hosting infrastructure to support new cloud

22 IRU refers to Indefeasible Right of Use (IRU), a long-term lease of part of the capacity of a cable system, effectively granting temporary ownership

23 See: https://www.srgresearch.com/articles/hyperscale-capex-jumped-59-second-quarter

24 North America is defined as the USA and Canada; Mexico is included in Latin America

25 Third party sources such as Cisco and Synergy Research Group use slightly different definitions, but use the term "hyperscale data centres" to refer to data centres operated by large OSPs or "hyperscale players"; these include large self-built facilities with hundreds of megawatts of capacity, as well as smaller facilities operated using colocation space

availability zones[26] in more countries and regions, which are primarily located in these three regions seeing the greatest investment.

Figure 4.2: Average annual investment by region[27] [Source: Analysys Mason, based on various sources, 2018]

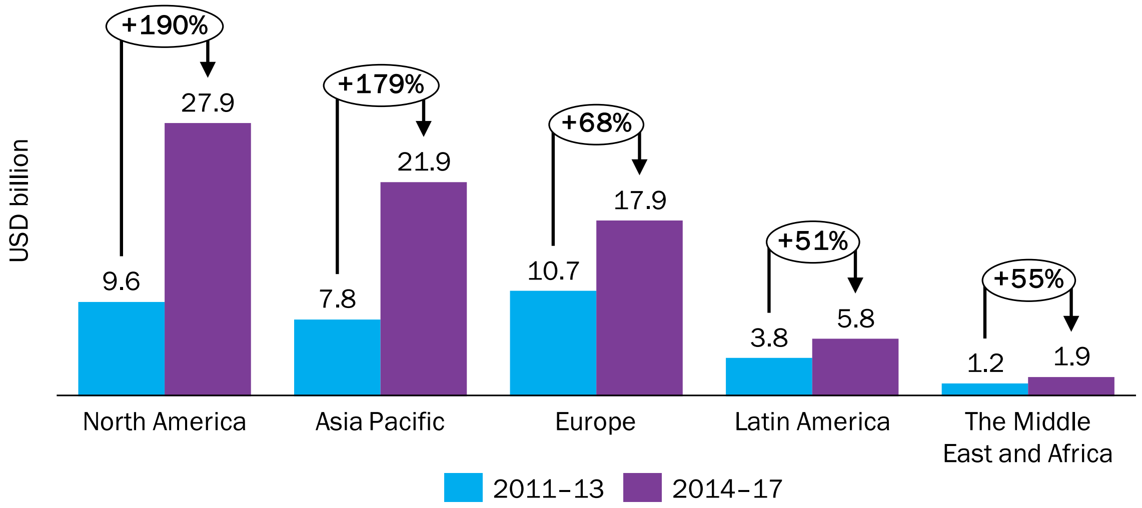

The Asia–Pacific region has also seen significant growth over the past four years, on the back of increased investment by China-based OSPs, as well as ongoing expansions by US-based OSPs in the region. Europe, which received the most investment in the previous study period, has seen slower growth, as the investment cycles of OSPs can result in frequent shifts in focus between regions.

Meanwhile, there has been a rise in investment in Latin America as well as the Middle East and Africa, though these regions still only account for a small share of global investment. Latin America has seen a number of recent developments in hosting infrastructure, as well as the construction of several submarine cable systems that are either operational or are expected to be in service by 2019. Markets in the Middle East and Africa have yet to reach sufficient scale to justify large-scale investments, and a large portion of data-centre demand is driven by government and telecoms operators.

Moving forward, high growth in Asia-Pacific investment is expected to continue, potentially overtaking the North American market in future years. Particularly strong growth is expected in China (discussed below), driven by domestic OSPs such as Alibaba and Tencent. Latin America and Africa should see a boost in investment in the upcoming years, with the construction of new submarine cables announced in both regions.

26 Cloud players divide their operations into cloud availability regions, with each containing several availability zones; each availability zone typically represents a single data-centre location

27 Regional split of investment from 2011 to 2013 has been restated slightly from the previous study due to new data available

Case study: OSPs are playing a key role in the evolution of internet delivery infrastructure in China

The interconnection landscape in China is currently undergoing a transformation due to growing demand for content and cloud services. Mainland China's first carrier-neutral internet exchange, CHN-IX, was launched by ChinaCache in early 2016.[28] CHN-IX is a strategic partnership between ChinaCache and AMS-IX. Over 50 parties are connected on the exchange, with peak traffic volume in the range of 50–60Gbps in 2018.[29]

Transit costs in China have been high, which has incentivised greater use of CDN services in recent years. This market was previously led by CDN specialists such as ChinaCache, but the large Chinese OSPs have invested in their own CDN infrastructure to support their cloud offerings, and now sell highly competitive CDN services on a commercial basis. Some traditional CDN players have started using OSP infrastructure as part of their CDN solutions, with ChinaCache even transferring some of its servers to OSPs as it moves towards a more cost-effective hosted cloud environment.

Investment in China by US-based OSPs has been limited to date, and have typically been established as partnerships with local companies. Examples include Amazon using space to provide cloud services in ChinaCache’s Beijing data centre,[30] and Akamai’s partnership with China Net Centre which allows it to offer CDN services in China.[31]

4.3 The majority of OSP spend is direct investment in self-owned infrastructure, however indirect investment has also grown at a rapid rate

As OSPs’ requirements grow, they seek to maximise control over essential infrastructure so as to optimise quality of service and control costs. This has led to significant growth in direct investment in self-owned infrastructure such as hyperscale data centres, fibre transmission equipment and submarine cable systems, private peering infrastructure and in-house CDNs. Direct investment has accounted for ~78% of investment in 2014–17.

Indirect investment has also grown rapidly, with OSPs buying services at scale and thereby financing investment in infrastructure by third-parties. The largest growth is derived from OSPs’ use of colocation data-centre facilities from third-party providers, which are typically used for smaller ‘edge’ data centres or where investment in private hyperscale facilities is not justified.

[28] See: http://www.datacenterdynamics.com/content-tracks/core-edge/chinacache-launches-chinas-first-internet-exchange-point/96099.fullarticle

[29] Informed by interviews, 2018

[30] ChinaCache. *Atecsys Data Centre Partners*. Available at https://en.chinacache.com/solutions/atecsys-data-center/

[31] Akamai. *China CDN.* Available at https://www.akamai.com/uk/en/products/web-performance/china-content-delivery-network.jsp

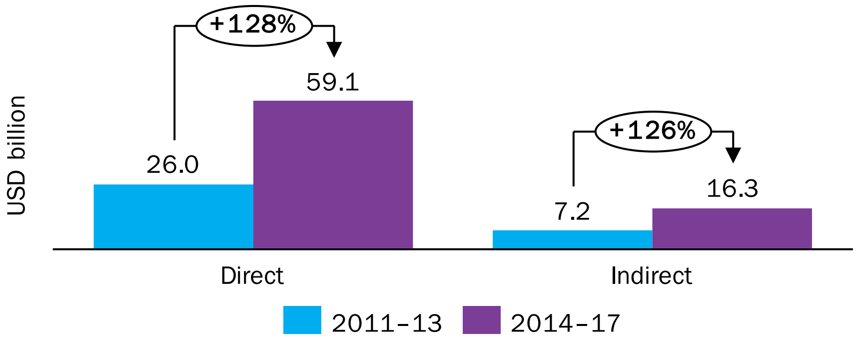

Figure 4.3: Comparison of average direct and indirect annual investment, 2011–13[32] vs. 2014–17 [Source: Analysys Mason, based on various sources, 2018]

While direct investment in self-owned infrastructure enables the greatest control over performance and cost, these investments benefit greatly from economies of scale and are therefore being led by the largest OSPs. Smaller OSPs also invest, but more often indirectly, and they are major users of the cloud infrastructure deployed by the big cloud players. Certain companies (Dropbox is a notable example) started out by using cloud infrastructure before making direct investments of their own once they achieve the necessary scale.

Case study: Dropbox has achieved sufficient scale to operate its own IT infrastructure in data centres

Dropbox is a cloud-based, file-hosting service based in California, which was founded in 2007 and has grown to over half a billion users. The company operates on a 'freemium' model, allowing users to store data up to a certain storage size for free, beyond which they would be required to pay for more storage space or additional features.

Dropbox initially used AWS to store customer data, but later reached a scale that justified investment in its own data centres. From 2015, Dropbox had begun to move users away from AWS's S3 storage service and onto its own infrastructure where possible. As part of this strategic shift, Dropbox leased space and installed custom-built equipment at several data centres in the USA, and deployed a network backbone to link these with other facilities around the globe. The company also invested in manpower and in the design of its own high-capacity servers and software.[33]

It should be noted that while Dropbox reached sufficient scale to see benefits from investment in its own infrastructure, other large OSPs have decided to move in the other direction. For example, in 2016, Spotify announced plans to close its data centres and migrate to Google Cloud Platform,[34] while Netflix completed the migration of a wide range of functions from in-house infrastructure to AWS that same year.[35]

32 Direct investment from 2011 to 2013 has been restated to be higher based on new information available regarding CDNs and terrestrial fibre

33 See: https://www.wired.com/2016/03/epic-story-dropboxs-exodus-amazon-cloud-empire/

34 Google Cloud. (2016), Spotify chooses Google Cloud Platform to power data infrastructure. Available at : https://cloud.google.com/blog/products/gcp/spotify-chooses-google-cloud-platform-to-power-data-infrastructure

35 Netflix Media Centre. (2016), Completing the Netflix Cloud Migration. Available at: https://media.netflix.com/en/company-blog/completing-the-netflix-cloud-migration

The importance of scale means that direct investment is concentrated among the largest OSPs, with ~57% of total direct investment coming from just ten companies in 2017. While the largest investments have come from US-based companies, Chinese giants Alibaba, Tencent and Baidu have also made significant investments in infrastructure.

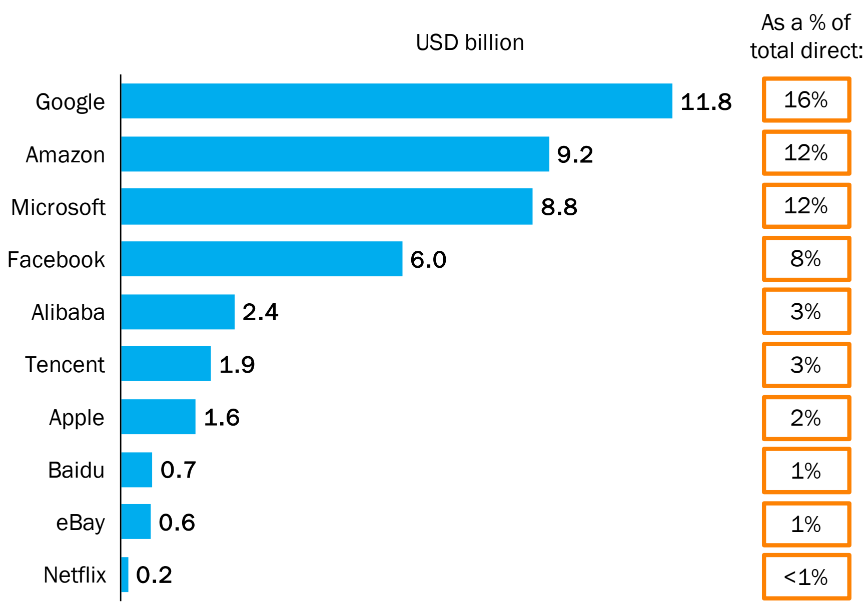

Figure 4.4: Direct investments in infrastructure for ten large OSPs in 2017 [Source: Analysys Mason based on company reporting, 2018]

5 OSPs are making significant investments to extend their networks, moving ever closer to end users

The overarching goal of the growing investment in infrastructure is to move content and services ever closer to end users, which helps to optimise service quality while controlling costs. This drive can be seen across the clusters, from investing in colocation space for 'edge' data centres, to building CDNs for caching content in carrier networks, and investing directly in new submarine cables to optimise costs in transporting content between countries.

In this section we discuss major trends affecting OSPs investments in infrastructure, addressing hosting, transport, and delivery in turn.

5.1 OSPs are investing in new hyperscale data centres, while also greatly increasing their spend on colocation space

The bulk of infrastructure investment made by OSPs has always been in hosting and data processing.[36] During this period, OSPs' investments in data centres have increased significantly, due to continued growth in demand for content, and the rise of cloud services.

5.1.1 Growth in cloud services has accelerated demand for hosting

The push of OSPs into the cloud market has accelerated the migration of enterprise IT from private infrastructure to public cloud services (such as those offered by market leader AWS). The market transition is illustrated by forecasts from Cisco, suggesting that public cloud will account for 73% of cloud workloads in 2021, up from just 22% in 2013.

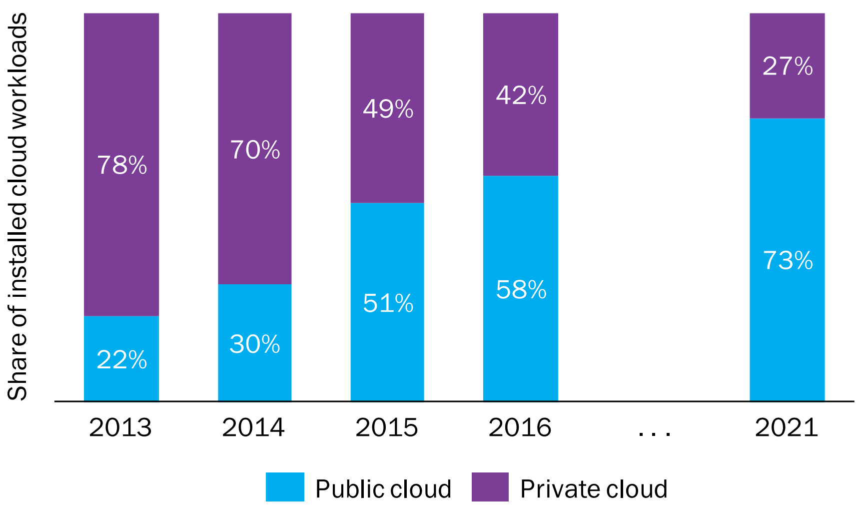

Figure 5.1: Public and private share of cloud workloads [Source: Cisco Global Cloud Index, 2016-2021, and earlier editions]

36 ~89% in 2011-13 according to our previous study

While many enterprises were initially sceptical of cloud services, attitudes are shifting as more IT managers see perceived risks outweighed by the potential benefits around functionality, flexibility and cost control. OSPs have also continued to improve and add complementary services such as CDNs in order to improve the appeal of their offerings. This increased focus on cloud services means that OSPs are now firmly established as key players in the enterprise IT environment.

The cloud platforms of large OSPs are also being used widely by smaller companies, enabling them to operate cost-effectively and scale up flexibly as they grow, as shown below for GO-JEK, Indonesia's first 'unicorn'[37] start-up offering on-demand services including transportation, logistics, delivery, personal care and payments.[38]

Case study: GO-JEK uses cloud services to power its various transport, delivery, and payments services

GO-JEK handled 100 million transactions a month as of 2018, through its network of over 250 000 drivers, 100 000 restaurants, and a number of other service providers.[39] The company has used the Google Cloud Platform (GCP) since 2015 to run its microservices-based application infrastructure. The platform also integrates with Google Maps and Google's BigQuery service (for machine-learning capabilities), allowing the company to make route predictions and to identify operational improvements.

Almost all of GO-JEKs services are run on the cloud, with the exception of GoPay, its payments service, which requires a data centre based in Indonesia due to local data sovereignty laws. According to CTO Ajey Gore, cloud services have been a welcome development for the company. He was quoted as saying: "we thrive on moving people and goods, so if we can have someone else take care of the infrastructure, we would rather do that and focus on the product."[40]

5.1.2 OSPs are driving transformation in data centres by investing in hyperscale facilities

In order to support the unprecedented growth in demand for content and cloud services, OSPs have invested heavily in data-centre infrastructure. While they often install equipment in co-location facilities from established data-centre operators, the largest OSPs have made large investments in building their own large hyperscale data centres.

The machine learning applications that have underpinned the growth of content (and in particular, rich content) and paid services require larger amounts of computing power per minute of content consumed or unit of service provided. OSPs have thus been required to invest in more powerful processing through their own R&D, deployed across data centres. Google, for instance, uses its large 'core' hyperscale data centres as general-purpose facilities, housing massive amounts of processing power for a variety of applications and services. Several of these facilities now exceed 100MW in capacity and 500 000 square

37 'Unicorn' start-ups refer to privately held start-up companies valued at over USD1 billion

38 See: https://www.cnbc.com/2017/04/24/a-ride-on-indonesias-first-and-only-unicorn.html

39 See: https://www.computerweekly.com/feature/Microservices-How-to-prepare-next-generation-cloud-applications

40 ibid.

feet in size.[41] These large self-built facilities are highly efficient, with power usage effectiveness (PUE) ratios[42] of close to 1, compared to PUE ratios of 1.5–2 for more traditional data centres.[43] These data centres typically link a larger number of servers (up to hundreds of thousands in the largest facilities) through a high-speed network. By both increasing the number of servers and adding power to existing servers, these hyperscale facilities provide several advantages over traditional data centres, including improved power and cost efficiency, reliability and speeds for end users, as well as the ability to more efficiently drive applications that require significant computing power, such as machine learning, 3D rendering, or cryptography.

Figure 5.2: Racks in Facebook's Prineville Data Centre [Source: Facebook, 2014]

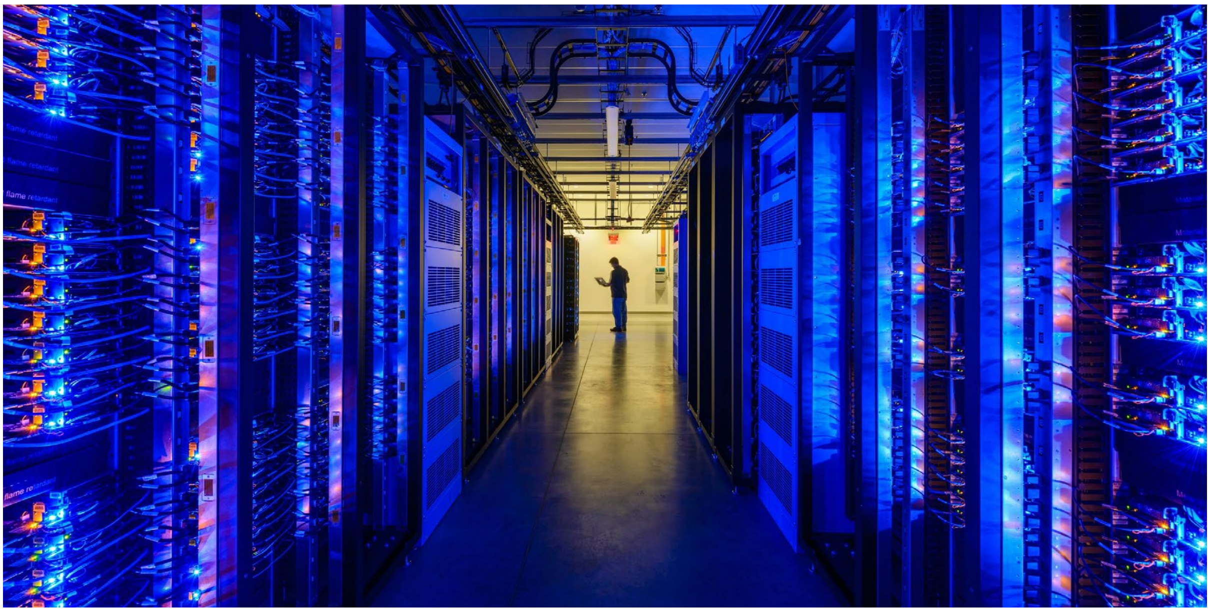

Synergy Research Group estimates there were over ~390 facilities used by hyperscale operators worldwide as of the end of 2017, however many of these facilities are colocation spaces, and are generally not of comparable size to the largest hyperscale facilities that are built by the OSPs themselves. While the majority of self-built and leased hyperscale facilities (see Figure 5.3) are presently located in the USA, OSPs have expanded into other markets where demand and market conditions are suitable. Examples of self-built facilities outside the USA include Google's data centres in Taiwan, Singapore, and Chile, as well as data centres built by Amazon and Facebook in Ireland, and Apple in Denmark.

41 See: http://worldstopdatacenters.com/hyper-scale-data-centers/

42 The power usage effectiveness ratio describes how efficiently a data centre uses energy, by computing the total energy used by the facility divided by energy used only by computing equipment. A smaller PUE indicates a more efficient data centre, as a smaller proportion of total energy is used on cooling and other overhead

43 Avgerinou, M., Bertoldi, P. and Castellazzi, L. (2017), 'Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency'. Energies, Vol. 10(10), pp.1470.

Figure 5.3: Share of total existing hyperscale data-centre facilities by country as of 2017 [Source: Synergy Research Group, 2018]

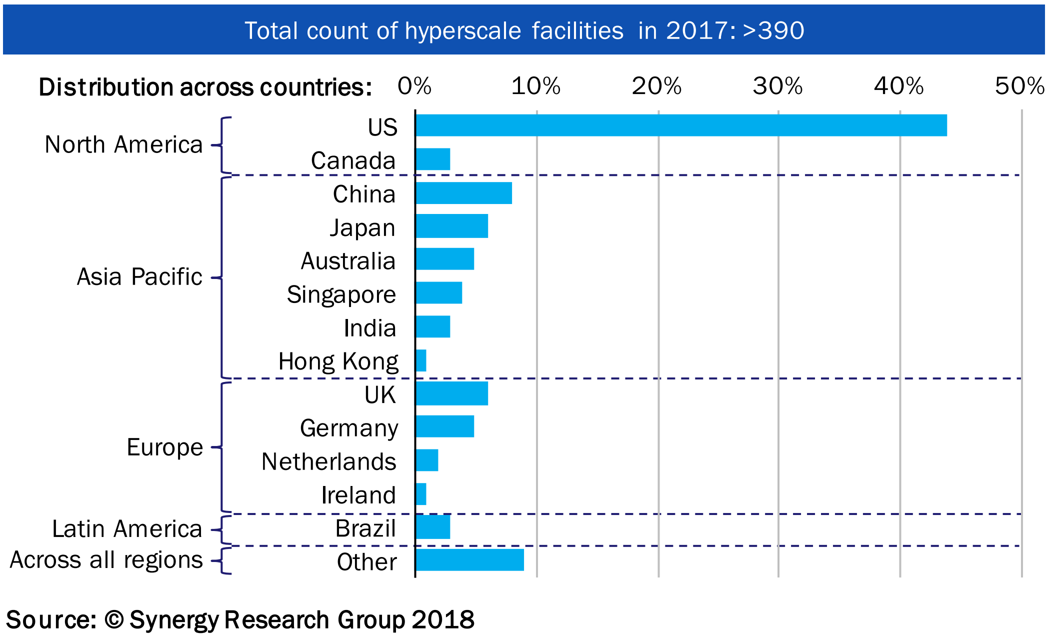

Case study: OSPs are now building more hyperscale facilities outside North America

OSPs have continued to build their own data centres, and have expanded the footprint of data-centre build since the period of the previous study. Apple opened its first international data centre in Denmark in 2017, with another expected to be ready by 2019.[44] Facebook has also announced its own data centre in Denmark planned for 2020,[45] which would be its second Scandinavian investment after opening a facility in Sweden in 2013.[46] Ireland has also come to host data-centre facilities for more OSPs, with Google and Facebook present, and Amazon also expressing an interest in building in the country.[47,48,49]

In the Asia–Pacific region, US-based OSPs have historically relied on colocation facilities for most of their capacity. However, there have been several self-build investments of note. Google has launched a data centre in Taiwan and two in Singapore (with a third planned).[50] According to Structure Research, as of 2017, only ~30% of the data-centre capacity in Singapore that was used by Google, Amazon,

44 See: https://www.reuters.com/article/us-apple-denmark/apple-to-build-second-renewables-powered-data-center-in-denmark-idUSKBN19V0MJ

45 See: http://www.datacenterdynamics.com/content-tracks/design-build/facebook-odense-data-center-is-official-first-details/97681.fullarticle

46 See: https://www.business-sweden.se/en/Invest/industries/Data-Centers-By-Sweden/news-and-downloads/investment-news2/facebook-expands-in-lulea-confirming-sweden-as-a-world-class-destination-for-data-centers/

47 Google. *Data Centre in Dublin, Ireland.* Available at https://www.google.com/about/datacenters/inside/locations/dublin/

48 See: https://www.thetimes.co.uk/article/facebook-doubles-down-on-clonee-data-centre-campus-d9wghn0td

49 See: https://www.independent.ie/business/commercial-property/amazon-planning-bid-for-second-major-data-centre-at-former-jacobs-site-37067242.html

50 Google. *Data Centre Locations.* Available at https://www.google.com/about/datacenters/inside/locations/

Microsoft, IBM, Alibaba, and Tencent combined, was leased from wholesale providers.[51] Facebook, meanwhile, recently announced a new 170 000 square metre facility in Singapore due in 2022.[52] China has also seen investments by OSPs in data-centre facilities. In May 2018, Tencent announced it had started trial operations of its 'mega data centre' in Guizhou, built into a hillside (see below).[53] The facility covers over 500 000 square metres, and is expected to house 50 000 servers for cloud data.[54] In 2017, Apple had also indicated intentions to invest USD1 billion in a Chinese data centre, which is expected to be ready in 2020.[55]

Meanwhile, Google announced in January 2015 that a self-built data centre in Chile became fully operational, which was its first in the Latin American region.[56] OSPs will continue to expand their deployments in new areas where there is sufficient demand, as long as they meet the relevant deployment criteria, such as having multiple power grids available for reliable supply.

Figure 5.4: Aerial view of a Tencent data centre under construction in Guizhou [Source: Visual China Group / Getty Images, 2018]

51 See: https://www.datacenterknowledge.com/asia-pacific/wholesalers-singapore-feel-squeeze-cloud-giants-build-own-data-centers

52 See: https://www.reuters.com/article/us-facebook-singapore/facebook-to-build-first-asian-data-center-in-singapore-investing-over-1-billion-idUSKCN1LM051

53 See: https://technode.com/2018/05/30/tencent-to-store-its-most-valuable-data-in-guizhou/

54 See: http://www.xinhuanet.com/english/2018-05/29/c_137215216.htm

55 See: https://data-economy.com/apple-open-1bn-chinese-data-centre-2020/

56 Google. *Data Centre in Quilicura, Chile*. Available at https://www.google.com/about/datacenters/inside/locations/quilicura/

5.1.3 Colocation remains popular and OSPs are making large investments in existing data centres

OSPs' initiatives to build their own facilities have not kept pace with their demand for data-centre capacity. Consequently, indirect investment in co-location space has grown even more quickly than direct investment in self-owned facilities, and this appears to be the case across the major markets of the world. At present, Amazon holds more leased square footage than it owns;[57] Facebook and Apple reportedly leased new large blocks of co-location space in the USA as recently as 2017.[58] In China, data-centre operator GDS Holdings announced in late 2017 that it would build a 54 000 square foot data centre specifically for Alibaba.[59]

Leasing allows OSPs to quickly and flexibly expand their presence in new markets without the need to invest large amounts in long term construction projects upfront. 'Edge' data centres using colocation space allow for content and enterprise cloud services to be delivered with greater speeds and reliability,[60] while also catering to customers who require data to be hosted within their country, either due to preference or regulatory obligations. Some countries have instituted data sovereignty laws, and key sectors (notably finance) are often under a higher level of scrutiny. OSPs are still favouring colocation in areas where the level of demand does not (yet) justify investment in building their own hyperscale facilities. From the perspective of colocation providers, OSPs may act as 'anchor tenants' in cases where large OSP requirements drive demand for new facilities. This increases the viability of future construction and enhances the environment for the wider data-centre ecosystem as a whole.

While the largest US-based cloud operators (Amazon, Microsoft and Google) own many hyperscale facilities in North America, they now have as many cloud availability zones in the Asia-Pacific region. Asia-Pacific hosts the largest number of cloud availability zones by some margin when also considering the largest Chinese cloud providers (Alibaba and Tencent).

57 Amazon. *Annual reports, proxies, and shareholder letters.* Available at https://ir.aboutamazon.com/annual-reports

58 North American Data Centers. *Data Center Real Estate Review 2017.* Available at https://nadatacenters.com/2018-real-estate-forecast/

59 See: https://www.datacenterknowledge.com/asia-pacific/cyrusone-s-china-partner-gds-build-alibaba-data-center-campus

60 See: https://www.technavio.com/blog/top-6-trending-data-center-innovations-2018

Figure 5.5: Number of cloud availability zones by region for major US- and China-based cloud providers as of August 2018 [Source: Company websites, Analysys Mason, 2018]

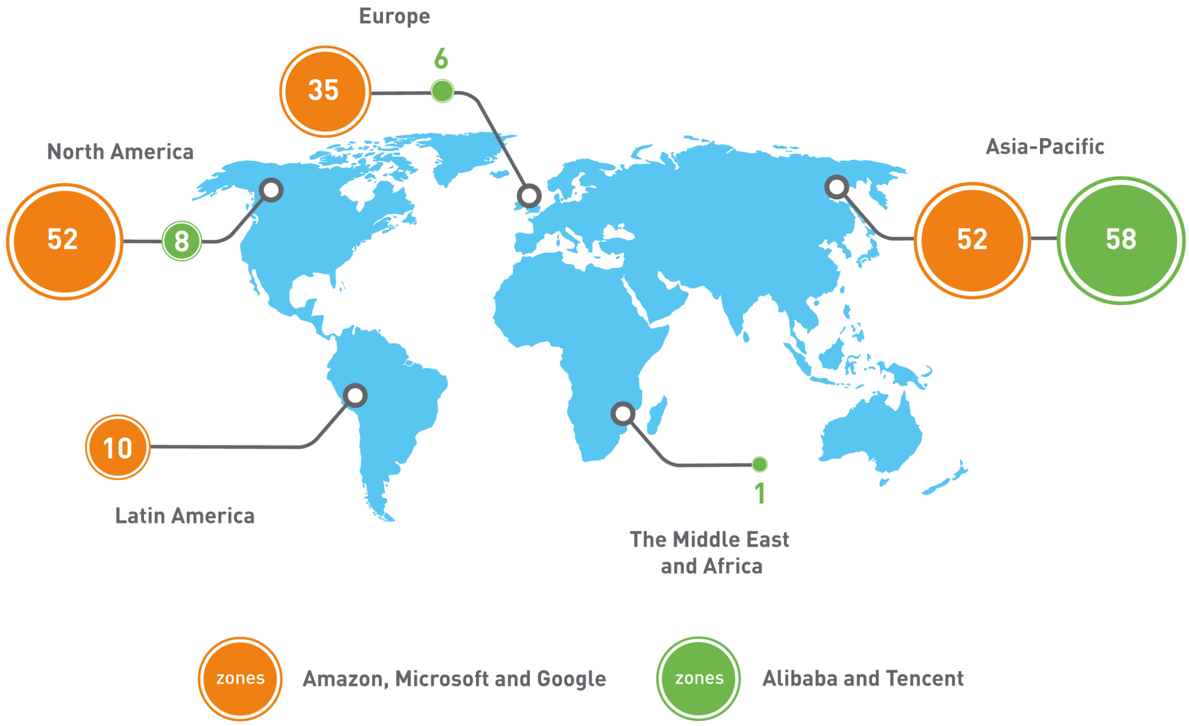

Case study: Equinix is seeing strong demand from OSPs for data-centre colocation

Large data centres built by OSPs, particularly those in North America and Europe, tend to be located in remote regions where land costs are relatively low. However, data-centre operator Equinix focuses on urban regions where there is a greater demand for interconnection services. Although its core business is colocation space, Equinix is also major facilitator of interconnection and peering, including enabling enterprise customers to connect to major cloud providers through its facilities.

Equinix has seen strong growth, with revenue increasing by an average of 19% per year since 2013. As of 2017, the company boasted a portfolio of close to 200 data centres worldwide. While the Americas remains Equinix's largest regional market, recent growth has been higher in Europe, the Middle East and Africa as well as the Asia–Pacific region. As demand for cloud and online services continues to expand, so will demand for data-centre colocation services by OSPs.

The building and leasing of data-centre facilities account for a large share of OSP investment, but even more investment is required for upgrading equipment and management systems within data centres. Servers and IT equipment need to be refreshed roughly once every three years, and, over the lifetime of a data centre, account for the majority of total costs. In order to meet unprecedented levels of demand for content, OSPs are now leading innovation in this equipment, designing servers and components to their exact specifications (discussed further in Section 6.1). Outside of buildings and computing hardware, major costs are incurred in power and cooling, and in some cases OSPs are making capital investments in renewable power generation facilities.

HOSTING ASIA–PACIFIC

Significant growth in demand for content and cloud services has driven data centre investment by both US and China based OSPs

Hosting investment

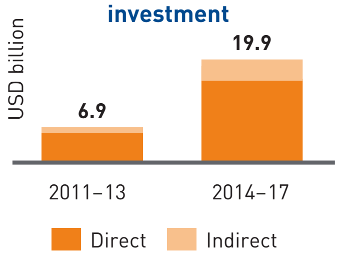

Share of global hosting investment

29%

2014–17

Count of hyperscale facilities[1]

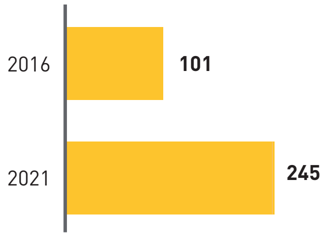

Cloud workload and compute instances (million)[1]

61.1 (2016)

200.5 (2021)

[1] Source: Cisco Global Cloud Index 2016-21

Cloud provider presence[2]

Cloud provider	# cloud regions	# cloud zones
Alibaba	14	41
Amazon	8	20
Tencent	11	17
Microsoft	5	17
Google	5	15

[2] Source: Company websites as of October 2018; Microsoft uses 'geography' instead of 'regions' and 'regions' instead of 'zones'

Developments in OSP self-build

Google

Operates self-built facilities in Singapore and Taiwan since 2013

Facebook

Building a 1.8 million square foot facility in Singapore due in 2020

Other data centre developments

Alibaba

In 2017, commissioned GDS Holdings to build a data centre in Hebei

HOSTING EUROPE

US-based OSPs continue to build and expand data centre facilities in rural regions, while growth of colocation in cities remains strong

Hosting investment

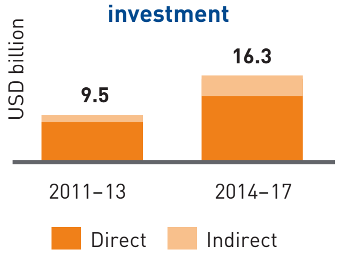

Share of global hosting investment

23%

2014–17

Count of hyperscale facilities[1]

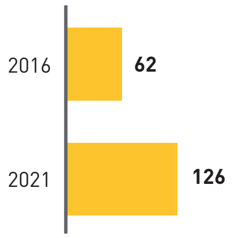

Cloud workload and compute instances (million)[1]

45.0 (2016)

111.7 (2021)

[1] Source: Cisco Global Cloud Index 2016-21

Cloud provider presence[2]

Cloud provider	# cloud regions	# cloud zones
Google	5	15
Amazon	4	12
Microsoft	5	8
Alibaba	2	4
Tencent	2	2

[2] Source: Company websites as of October 2018; Microsoft uses 'geography' instead of 'regions' and 'regions' instead of 'zones'

Developments in OSP self-build

Apple

Built a facility in Viborg, Denmark in 2017; another due in Aabenraa, Denmark by 2019

Facebook

In 2017, announced a new facility in Odense, Denmark, due 2020

In 2018, announced expansions to its existing footprint in Clonee, Ireland

HOSTING LATIN AMERICA

Investment in hosting within Latin America has been limited, but is expected to grow as cloud providers expand their offerings in the region

Hosting investment

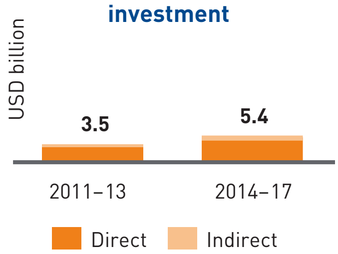

Share of global hosting investment

8%

2014–17

Count of hyperscale facilities[1]

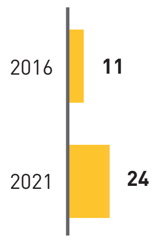

Cloud workload and compute instances (million)[1]

7.2 (2016)

17.0 (2021)

[1] Source: Cisco Global Cloud Index 2016-21

Cloud provider presence[2]

Cloud provider	# cloud regions	# cloud zones
Amazon	2	6
Google	1	3
Microsoft	1	1
Alibaba	-	-
Tencent	-	-

[2] Source: Company websites as of October 2018; Microsoft uses 'geography' instead of 'regions' and 'regions' instead of 'zones'

Developments in OSP self-build

Google

A facility in Quilicura, Chile has been operational since 2015. Estimated long term investment in the site amounts to USD150 million

Other data centre developments

Equinix

7 data centre locations across Latin America

Brazil: Rio de Janeiro (2), Sao Paolo (4)

Colombia: Bogota (1)

5.2 Investment in data transport has grown apace, as OSPs carry content and enterprise data between their global data centres

OSPs invest heavily in the fibre infrastructure required to transport content between data centres, both within and between countries. Terrestrial capacity is mainly based on ‘dark fibre’, where OSPs lease access to existing fibre infrastructure and install their own equipment to ‘light’ the fibre. Capacity between countries and continents more often uses submarine cables and the past four years have seen OSPs accelerate their investment in these systems, both indirectly through capacity purchase agreements, and by investing directly in new cables.

5.2.1 Terrestrial cable investments have grown as OSPs continue to build diverse links between data centres and peering locations

While OSPs’ investments in submarine cables are more widely publicised, terrestrial fibre accounts for a larger share of transport investment. OSPs require these long-distance networks to link their data centres using diverse and resilient routes, as well as to link data centres to locations where they can peer with access networks to reach end users.

OSPs’ investments in terrestrial fibre do not typically extend to deploying their own fibre cables, as these deployments require complex processes to gain ‘rights of way’ permissions. The favoured model, where available, is to lease access to dark fibre based on 10- to 20-year agreements. In areas where dark fibre is not available, or in emerging regions with relatively small capacity needs, OSPs will lease capacity from backbone providers over the short term.

The dark fibre model enables OSPs many of the benefits of infrastructure ownership, such as a greater degree of cost certainty as their bandwidth requirements increase over time. Leasing capacity from carriers could become prohibitively expensive given the enormous bandwidth requirements of many OSPs, whereas the dark fibre model enables them to upgrade their active equipment to maximise utilisation over the fibre pairs they are leasing. By establishing partnerships with fibre backbone providers, OSPs are purchasing dark fibre on existing routes where possible, and can also drive the deployment of fibre on new routes by acting as anchor tenants.

For OSPs, the dark fibre model requires both indirect investment in leasing dark fibre access, as well as direct investment in active networking equipment. The active equipment will typically account for a larger share of investment over time when considering equipment refresh cycles of around five years.

5.2.2 OSPs account for a growing share of international bandwidth used, which is itself growing rapidly

Growth in demand for traditional online content, as well as the rapid emergence of enterprise cloud services, has driven continued growth in demand for international bandwidth. TeleGeography reports that total international bandwidth on global basis grew nearly five-fold between 2013 and 2017.

Bandwidth purchased and used by OSPs (referred to by TeleGeography as ‘Content Providers’) has grown the most, and as of 2017 exceeded bandwidth purchased by internet backbone providers to become the largest segment. Some of this may be displacing spend from telecoms carriers, but all segments continue to grow and OSPs are paying for a growing share of the costs of international transit as demand for their content drives demand for bandwidth.

Figure 5.6: Total used international bandwidth [Source: TeleGeography, 2018]

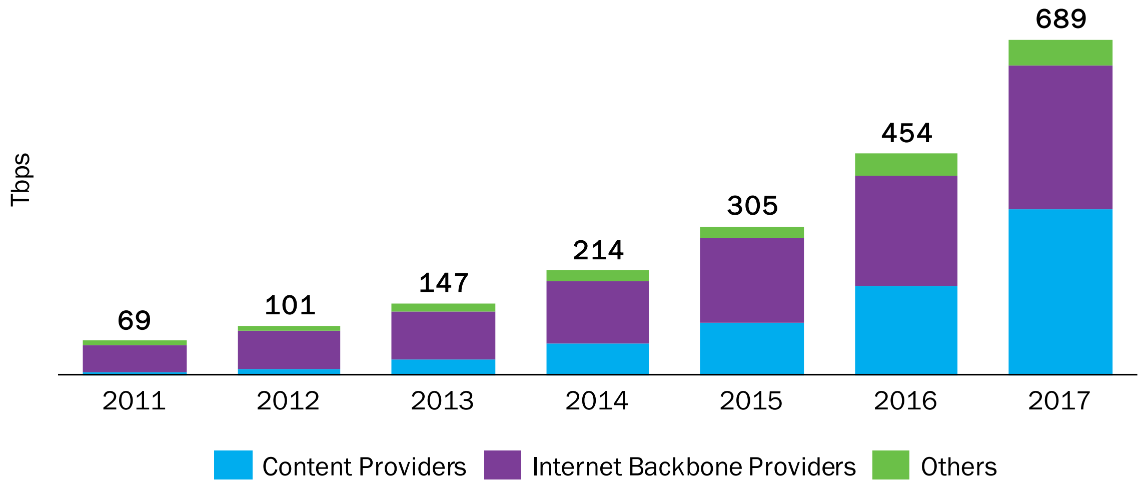

Moving forward, the demand for international bandwidth is expected to continue to grow at a significant pace, driven by growing global adoption of services and applications, as well as improved broadband speeds and smartphone take-up in developing regions. Greater use of caching within countries will moderate demand for international bandwidth to a degree, however OSPs will still need to refresh ever-increasing amounts of content in new caches around the globe.

5.2.3 OSPs are indirectly funding a large share of the global investment in submarine cables

Global investment in submarine cables has seen some fluctuation due to the irregular nature of cable deployments, but has averaged USD3.6 billion over the past four years.

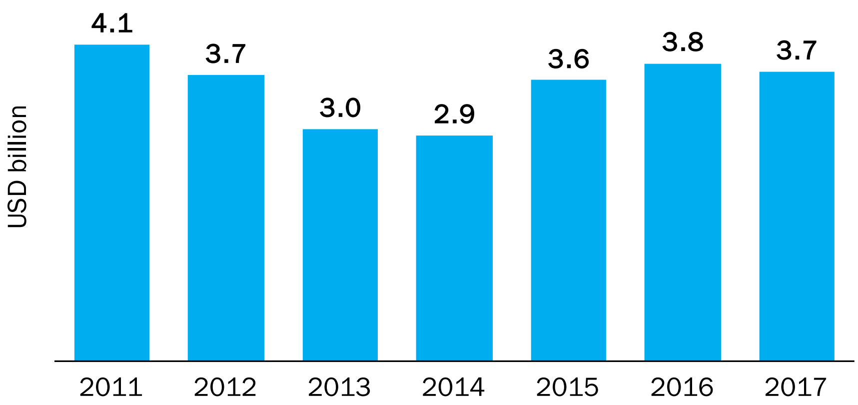

Figure 5.7: Global capital investment in submarine cables [Source: Analysys Mason based on TeleGeography data, 2018] [61]

[61] Capital investment for each new cable is assumed to be spread over two years, with annual maintenance and upgrade capex applied to existing cables based on Analysys Mason benchmarks

A large share of this investment in cable systems has been financed by OSPs, with most investment made indirectly thorough capacity purchase agreements on third-party cables. In our previous study, we estimated that OSPs accounted for ~50% of internet traffic across 2011 to 2013; we now estimate that figure has grown to ~56% over the period of 2014 to 2017. This estimate includes bandwidth that OSPs pay for directly, as well as some bandwidth that is carried by backbone providers and funded by OSPs through leases, transit and peering arrangements. This share of traffic translates to ~39% of investment, as OSPs tend to are able to achieve lower effective prices than the industry average by deploying their own cables and by negotiating favourable deals.

5.2.4 The largest OSPs are making significant direct investments in submarine cables

While most OSPs are only investing indirectly by purchasing international bandwidth as described above, four players (Google, Facebook, Microsoft and Amazon) have stepped up their direct investments. While the previous study found just two cables with OSP ownership stakes, current announcements suggest that by 2020 this will reach 22 unique cables.

Examples of investments include Microsoft and Facebook announcing in September 2017 the completion of the Marea Project, a 160Tbps cable across the Atlantic; and Amazon making its first cable investments as a consortium member in the upcoming ‘Jupiter’ and ‘Bay to Bay Express’ cables due in 2020, both trans-Pacific systems. Over time, more of the world’s largest OSPs could decide to invest directly in submarine cables. In particular, large Chinese cloud providers including Alibaba, Tencent, and Baidu, might consider investing directly in transport infrastructure in the future as global demand for their cloud services grows.[62]

62 TeleGeography. (2018). *Global Bandwidth Research Service*.

Figure 5.8: Map of anticipated submarine cables with OSP investment by 2020 [Source: Analysys Mason based on information from TeleGeography, news articles, company press releases, 2018]

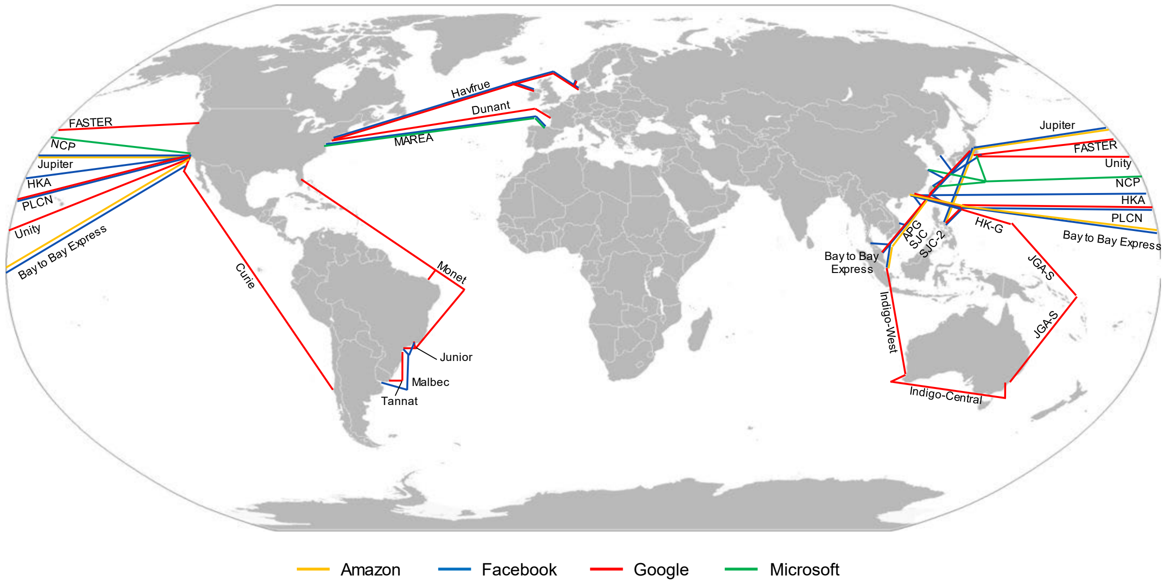

While the trans-Pacific and trans-Atlantic routes account for over half of these ownership stakes, new cables in less well-served regions have also emerged. In particular, OSPs are deploying several new cables that land in Latin America. Monet, the first of these, was ready for service in 2017 and includes Google as a consortium member. The cable has a design capacity of 64Tbps and is also one of the first cables not to use a landing station on the shore, with the cable terminating directly at a data-centre facility operated by Equinix in Florida.[63]

More generally, OSPs may begin to invest more in regions outside the busy trans-Pacific and trans-Atlantic routes. TeleGeography has highlighted the India-Singapore, India-Europe, and Europe-Africa routes as possible future areas of investment by OSPs.

OSPs are also deploying cables to relatively unique landing stations even within established regions or routes. For instance, the Marea cable crosses the trans-Atlantic by linking Virginia and Bilbao, unlike most other trans-Atlantic cables, which link New York and London. Many investments in cables are driven by OSPs' requirement to have geographically diverse routes, to provide resiliency in the event of a network outage where one cable or landing station experiences a network outage.

By investing in new cables, OSPs are able to acquire large amounts of capacity on more technologically-advanced systems. These newer systems typically provide capacity at lower unit costs, especially when considering future upgrade costs that may be required for older systems. Ownership of cables also provides OSPs with a greater degree of control over network design and performance management.

63 See: https://www.datacenterdynamics.com/news/google-backed-monet-submarine-cable-between-florida-and-brazil-lights-up/

Figure 5.9: The Marea cable coiled on ship [Source: Photo by RUN Studios, provided by Microsoft,[64] 2017]

In most cases, OSPs invest as part of a consortium with other carriers, as these partnerships make it easier to gain access to landing stations or licenses. The bandwidth requirements of carriers are not growing as quickly, as shown in Figure 5.6 above, and OSPs often shoulder a larger share of the investment than other consortium partners, but carriers also benefit from the economies of scale enabled by these substantial OSP requirements.

Moving forward, OSPs could increasingly deploy cables independently of other carriers that would usually make up the rest of a consortium. At present, only Google has announced such projects, starting with the Junior cable in Brazil in 2018, as well as the Curie cable between Chile and the US scheduled for service in 2019.[65] As their bandwidth demand continues to grow, more OSPs could invest in developing the technical capabilities to deploy independent cables, especially if the cost benefits are substantial and other consortium members (such as ISPs) no longer have as much incentive to invest in new capacity. Where OSPs do enter consortia, these may involve fewer partners in future as OSPs take on a larger role. For example, Google is deploying the trans-Atlantic ‘Dunant’ cable to be ready by 2020, in collaboration with just one telco partner, Orange.[66]

64 Photo available at https://news.microsoft.com/marea/

65 Google Blog. (2018), *Delivering increased connectivity with our first private trans-Atlantic subsea cable.* Available at https://www.blog.google/products/google-cloud/delivering-increased-connectivity-with-our-first-private-trans-atlantic-subsea-cable/

66 See: https://www.orange.com/en/Press-Room/press-releases/press-releases-2018/Orange-collaborates-with-Google-on-a-new-submarine-cable-across-the-Atlantic-Ocean

TRANSPORT ASIA–PACIFIC

The Asia–Pacific accounts for the largest share of OSP investment in transport, and boasts the most direct submarine cable investments

Transport investment

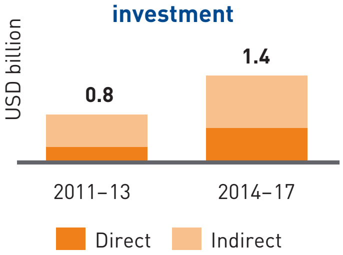

Share of global transport investment

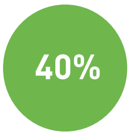

2014–17

Demand for international bandwidth (Tbps)[1]

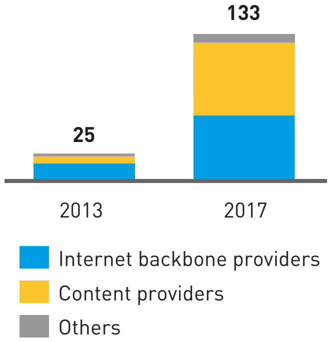

[1]Source: TeleGeography Global Bandwidth Report

[2]Source: TeleGeography Submarine Cable Map; note: only includes cables where OSPs are consortium members or sole owners

Key OSP cables in region[2]

Cable name	RFS date	Route	OSP investors
Unity/EAC-Pacific	2010	Trans-Pacific	Google
Southeast Asia Japan	2013	Intra-Asia	Google
FASTER	2016	Trans-Pacific	Google
Asia Pacific Gateway	2016	Intra-Asia	Facebook
New Cross Pacific	2018	Trans-Pacific	Microsoft
Pacific Light Cable Network	2019	Trans-Pacific	Google, Facebook
Indigo-West	2019	Intra-Asia	Google
Indigo-Central	2019	Intra-Asia	Google
Hong Kong-Guam	2019	Trans-Pacific	Google
Japan-Guam-Australia	2019	Trans-Pacific	Google
JUPITER	2020	Trans-Pacific	Facebook, Amazon
Hong Kong - Americas	2020	Trans-Pacific	Facebook
Southeast Asia - Japan 2	2020	Intra-Asia	Facebook
Bay to Bay Express	2020	Trans-Pacific	Facebook, Amazon

TRANSPORT EUROPE

Transport investment for Europe has been more steady than in other regions, with only three subsea cable investments expected by 2020

Transport investment

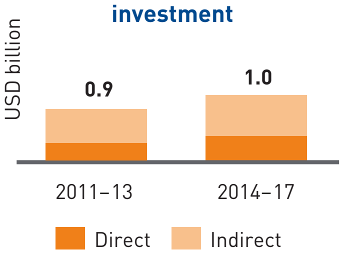

Share of global transport investment

2014–17

Demand for international bandwidth (Tbps)[1]

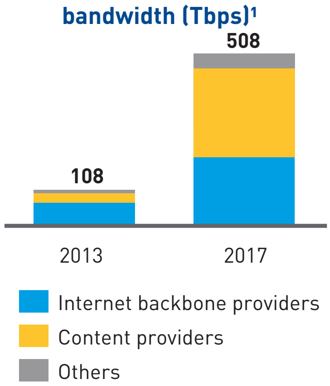

[1]Source: TeleGeography Global Bandwidth Report

[2]Source: TeleGeography Submarine Cable Map; note: only includes cables where OSPs are consortium members or sole owners

Key OSP cables in region[2]

Cable name	RFS date	Route	OSP investors
MAREA	2017	Trans-Atlantic	Facebook, Microsoft
Havfrue/AEC-2	2019	Trans-Atlantic	Google, Facebook
Dunant	2020	Trans-Atlantic	Google

TRANSPORT LATIN AMERICA

OSPs are starting to invest more in transport within the region, with both Google and Facebook announcing Intra-Americas cables by 2020

Transport investment

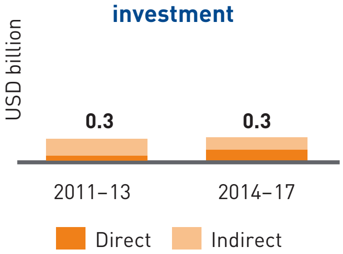

Share of global transport investment

2014–17

Demand for international bandwidth (Tbps)[1]

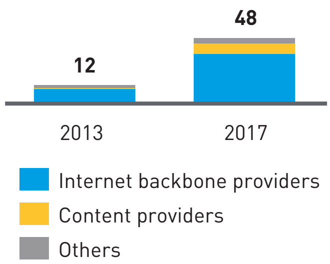

[1]Source: TeleGeography Global Bandwidth Report

[2]Source: TeleGeography Submarine Cable Map; note: only includes cables where OSPs are consortium members or sole owners

Key OSP cables in region[2]

Cable name	RFS date	Route	OSP investors
Monet	2017	Intra-Americas	Google
Tannat	2018	Intra-Americas	Google
Junior	2018	Intra-Americas	Google
Curie	2019	Intra-Americas	Google
Malbec	2020	Intra-Americas	Facebook

5.3 OSPs continue to develop delivery networks to bring services closer to end users

After transmitting content between their data centres, and from core to edge, over submarine and terrestrial transport networks, OSPs need to deliver content to end users via ISPs. OSPs typically exchange traffic with ISPs at facilities called 'peering points', which can be private or public. Public peering takes place at internet exchange points (IXPs), where a large number of OSPs and ISPs, of varying size are able to exchange traffic with one another. Private peering takes place when an OSP peers directly with an ISP by installing its own equipment in a location where it can interconnect with the ISP's network (which could also be at an IXP).

OSPs are also investing in CDN infrastructure to cache content closer to end users, in order to improve service quality and control costs by minimising the duplication of content that is transported over backbone networks to reach end users.

5.3.1 The number of public peering points used by OSPs has grown significantly, but much larger volumes of traffic are exchanged through private peering

Over the past four years, both private and public peering have grown significantly both in terms of the number of peering points and the volume of traffic exchanged. The number of private peering points grew by 88% between 2014 and 2018, while the number of public peering points has more than tripled as OSPs continue to expand their peering footprint to interconnect with more networks.

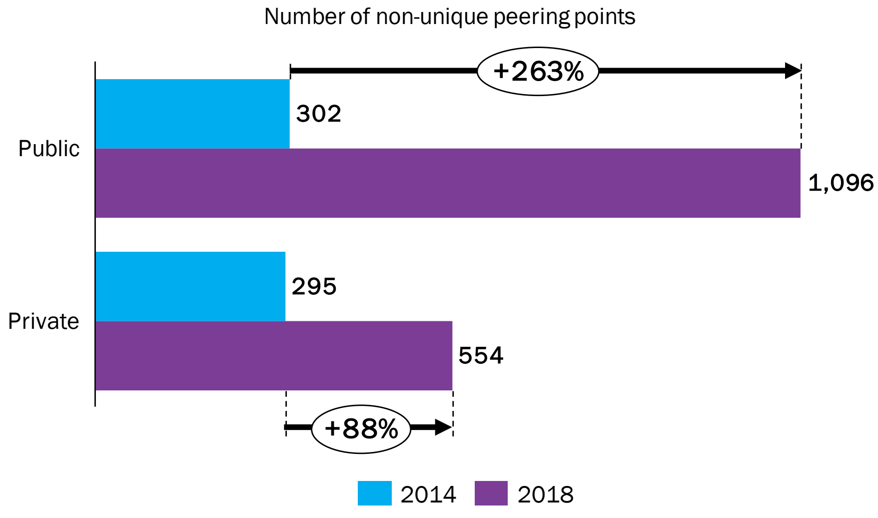

Figure 5.10: Number of peering points[67] for ten major OSPs [Source: Analysys Mason based on PeeringDB data]

Internet exchanges providing public peering are seeing the level of traffic exchanged grow rapidly, with the largest OSPs now taking up several 100Gbps ports at the larger exchanges such as AMS-IX in Amsterdam, DE-CIX in Frankfurt and LINX in London – the three largest IXPs in the world. The fees paid to exchanges are essentially indirect investments in infrastructure, as internet exchanges invest the income received from members to upgrade exchange platforms in order to cater for the growing traffic.

67 The ten OSPs analysed are Google, Facebook, Microsoft, Amazon, Yahoo, Netflix, Apple, Tencent, eBay, and Baidu. Figures are as of August each year, and show the sum of each OSP's peering points; locations with multiple OSPs will figure more than once, as it represents multiple unique peering opportunities

This involves buying new switching equipment, and equipment for high-capacity fibre links to connect data centres. For instance, AMS-IX operates from 14 data-centre locations in co-location facilities within Amsterdam; LINX from 16 data-centres in colocation facilities within London and DE-CIX from 27 facilities within Frankfurt. [68,69,70] Internet exchange members are able to peer at any of these locations, and IXPs require the infrastructure to carry large volumes of traffic between their respective data centres.

Typically, OSPs use public peering at internet exchanges to gain access to a large number of peering partners, though many of these partners individually would only carry relatively small amounts of traffic. OSPs typically have "open" peering policies which means that any operator connected to the IXP, whether large or small, can benefit from the investment of the OSP to reach the IXP, and exchange traffic locally with the OSP, providing cost savings and performance enhancements to both large and small network operators.

If a connection with a particular partner reaches a sufficient volume of traffic, it becomes more cost-effective for the OSP to peer privately with that partner. The majority of traffic between OSPs and ISPs is exchanged at private peering points, and these private facilities require direct investment in switching and routing equipment.

The Equinix Global Interconnection Index forecasts significant further growth in interconnection capacity. The strongest growth is expected in interconnection between enterprises and what Equinix refers to as 'cloud and IT providers', driven by the development of new digital services and migration of existing enterprise workloads to the cloud.[71]

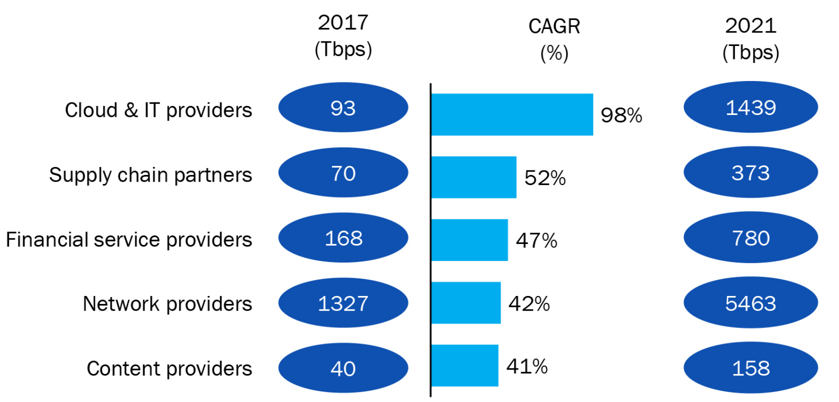

Figure 5.11: Global installed interconnection capacity by counterparty [Source: Equinix Global Interconnection Index Volume 2, 2018]

68 AMS-IX. *Colocations*. Available at https://ams-ix.net/connect-to-ams-ix/colocations

69 LINX. *London LANs*. Available at https://www.linx.net/about/our-network/

70 DE-CIX. *Frankfurt enabled sites.* Available at https://www.de-cix.net/en/locations/germany/frankfurt/enabled-sites

71 Equinix. (2018), *Global Interconnection Index Volume 2.* Available at https://www.equinix.com/global-interconnection-index-gxi-report/

5.3.2 CDNs have grown in importance due to the rise in video content and cloud services

Since 2014, CDNs have grown in importance due to the proliferation of online video and the emergence of cloud services.

CDNs provide end users and OSPs with a variety of benefits. For rich content such as video, CDNs can be used to 'cache' data in a large number of nodes distributed across telecoms networks, storing content closer to end users. This allows costs to be reduced in two complementary ways: firstly, the absolute volume of traffic being distributed is reduced, as OSPs do not need to transport the same content to each end user all the way from a central location, which helps to control bandwidth costs. Secondly, CDNs allow a reduction in the unit cost of distribution by taking advantage of lower off-peak costs to fill caches with content. Google and Netflix use caching extensively for delivering video content to users. The delivery of YouTube videos, as well as long form TV shows, can both benefit from caching.

A key benefit for end users is that CDNs reduce latency (because the content is closer to the end user) and also use file compression that further reduces website loading time. Less directly, end users will also see some benefit from the reduced transport costs in the system, as well as from the distributed network's improved reliability. This superior network 'uptime' is the result of load balancing, whereby traffic is distributed over several servers to handle sudden increases in traffic, which also enables failovers to other servers in the event of a hardware malfunction. These improvements in latency and reliability vastly enhance end-user experience, especially for cloud services where uptime is crucial.

Finally, CDNs help to address a variety of security concerns. Many CDNs provide security certificates[72] to users, as well as a range of tools to optimise certificate management. For example, Akamai's Secure CDN product assists with 'complete lifecycle management', from ordering new certificates to automating renewal.[73] CDNs are also able to withstand distributed-denial-of-service (DDoS) attacks to an extent by using scale to absorb attack traffic. CDN security solutions are crucial to maintaining enterprise trust in cloud services.

These benefits are promoting an increase in CDN use among OSPs, and Cisco reports that global CDN internet traffic tripled between 2013 and 2016, and that this strong growth is set to continue.

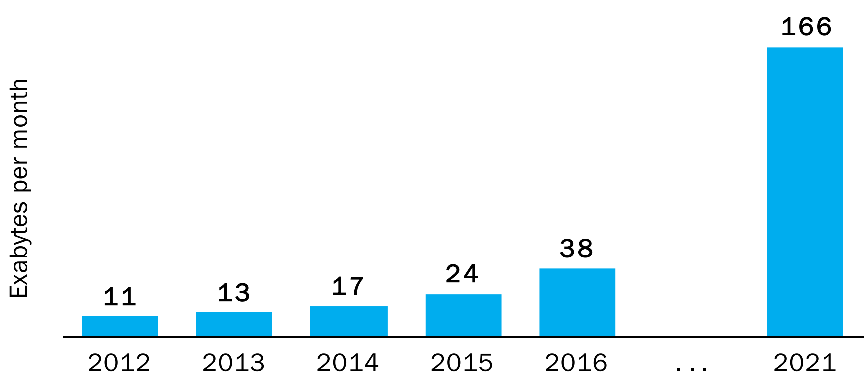

Figure 5.12: Global CDN internet traffic [Source: Cisco Visual Networking Index, 2016–21 and earlier editions]

72 Secure sockets layer (SSL) is a commonly used cryptographic protocol for providing internet security; transport layer security (TLS) is an upgraded version of SSL

73 Akamai. *Secure CDN.* Available at https://www.akamai.com/uk/en/solutions/intelligent-platform/secure-cdn.jsp

A number of larger OSPs have now built their own CDNs in order to gain greater control of service quality and costs, compared to using commercial CDNs from companies such as Akamai. Cisco estimates that as of 2016, 52% of internet traffic is delivered through CDNs, of which a majority is now on in-house, or 'private', CDNs.[74] By 2021, 48% of global internet traffic is expected to pass through in-house CDNs compared to just 22% on commercial CDNs.

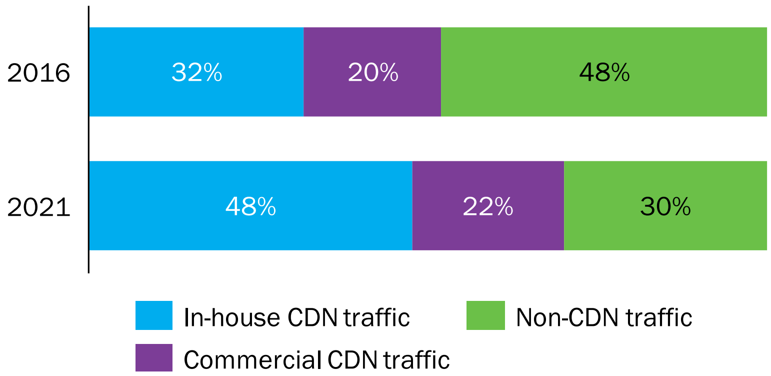

Figure 5.13: Internet traffic by CDN use [Source: Analysys Mason based on Cisco Visual Networking Index, 2016–21]

The Google Global Cache (GGC), which stores YouTube videos and other content within ISP networks, has been operating since 2008,[75] while Netflix set up Open Connect, its globally distributed CDN in 2011. Facebook had also announced that it had started working on its own CDN, Facebook Network Appliance (FNA), in 2012,[76] and as of 2018, third-party estimates place the number of nodes within the network at close to 1700.[77]

In addition to these in-house CDNs, a number of OSPs have also launched commercial CDNs to support services used by their cloud customers, including Amazon (2008),[78] Alibaba (2009),[79] Google (2015),[80] and Microsoft (2018).[81] This investment means Google now operators two separate CDNs: its 'commercial' Cloud CDN caters to cloud customers (over 90 locations across the world as of 2018),[82] while GGC is used for YouTube and other in-house Google content (around 800 locations as of 2015).[83]

74 Cisco, in its Visual Networking Index, defines private CDNs as those built and operated by content providers for their own content only, and is not available to other content providers for purchase

75 See: https://www.capacitymedia.com/articles/3324344/Video-cachinga-cache-rich-opportunity

76 See: https://gigaom.com/2012/06/21/like-netflix-facebook-is-planning-its-own-cdn/

77 See: https://anuragbhatia.com/2018/03/networking/isp-column/mapping-facebooks-fna-cdn-nodes-across-the-world/

78 Amazon. (2008), *Announcing Amazon Cloudfront.* Available at https://aws.amazon.com/about-aws/whats-new/2008/11/18/whats-new-cloudfront/

79 Alibaba Cloud Forum. (2016), *Evolution of Alibaba Cloud CDN for Double 11 shopping carnivals.* Available at https://www.alibabacloud.com/forum/read-648

80 See: http://fortune.com/2015/12/10/google-cloud-content-delivery-network/

81 Microsoft Azure. (2018), *Announcing Microsoft's own Content Delivery Network.* Available at https://azure.microsoft.com/en-gb/blog/announcing-microsoft-s-own-cdn-network/

82 Google Cloud. (2018), Locations. Available at https://cloud.google.com/cdn/docs/locations

83 See: http://blog.speedchecker.xyz/2015/11/30/demystifying-google-global-cache/

Case study: Netflix's Open Connect CDN enables it to improve service quality and reduce costs

Netflix is one of the largest video-on-demand (VoD) players in the market, and has developed its own in-house CDN, named Netflix Open Connect, to deliver its online video traffic to over 190 countries.[84]

Open Connect was developed in 2011, and works by embedding video content within Open Connect Appliances (OCAs) that are deployed across the world, in order to localise the delivery of content to end users. OCAs are provided to partners for free, and are either installed directly in the networks of ISP partners, or at IXPs where ISPs can interconnect with the OCAs via public or private peering. When clients play a piece of content, a request is sent to Netflix's control plane, which is hosted by AWS, and co-ordinates playback for content streamed from OCAs to end-user devices.

Transporting content around the world from a single centralised data centre would result in large costs for Netflix in relation to internet backbone capacity. Using these OCAs to cache content close to customers helps Netflix to reduce its costs substantially.

As of 2016, OCAs were installed in close to a thousand separate locations globally, with over 90% of traffic delivered through direct connections between Open Connect and ISP partners.[85]

Commercial CDN players such as Akamai have continued to retain a healthy share of the market, sustained by demand from large enterprise customers who may use an OSP cloud yet do not wish to depend on a single CDN provider, and want a wider range of products to cater to specific needs. Recognising customer demand for using multiple CDNs, cloud providers have developed methods to allow cloud customers to access external CDN services easily. For instance, Google Cloud CDN Interconnect allows select CDN providers to establish links with Google's edge network.

Overall investment in CDNs by OSPs has more than doubled since the last report, with direct investment accounting for a larger share, and seeing faster growth, than indirect investment (i.e. investment in commercial CDNs). Meanwhile, investment in commercial CDNs also grew, albeit more modestly, and this is indirectly funded by smaller OSPs which are paying to use CDN services.

84 Netflix Media Centre. (2016), *How Netflix Works with ISPs Around the Globe to Deliver a Great Viewing Experience*. Available at https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience

85 ibid.

DELIVERY ASIA–PACIFIC

Investment in Asia–Pacific delivery is expected to grow due to strong demand for online services and ongoing development of peering

Delivery investment

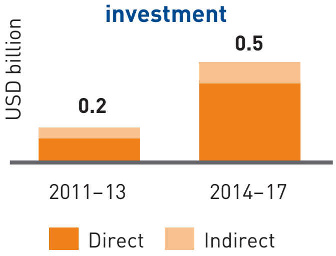

Share of global delivery investment

2014–17

Regional IXP statistics in 2014 and 2018[1]

	2014	2018
Number of IXPs	89	139
Share of global IXPs	19%	21%

Count of private peering points by OSP in 2018[2]

OSP	Private peering points (#)
Google	26
Microsoft	23
Apple	16
Amazon	15
Facebook	15

Regional share of global CDN traffic[3]

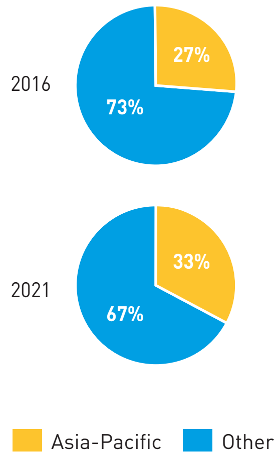

[1] Source: For 2018, PeeringDB, for 2014: Euro-IX

[2] Source: PeeringDB

[3] Source: Cisco Visual Networking Index 2016-21

DELIVERY EUROPE

Growth in delivery investment is strong but slightly limited by high existing levels of development in the peering landscape

Delivery investment

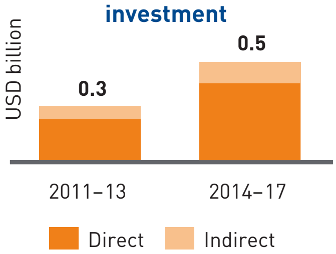

Share of global delivery investment

2014–17

Regional IXP statistics in 2014 and 2018[1]

	2014	2018
Number of IXPs	186	242
Share of global IXPs	41%	36%

Count of private peering points by OSP in 2018[2]

OSP	Private peering points (#)
Google	52
Microsoft	36
Amazon	32
Facebook	24
Apple	19

Regional share of global CDN traffic[3]

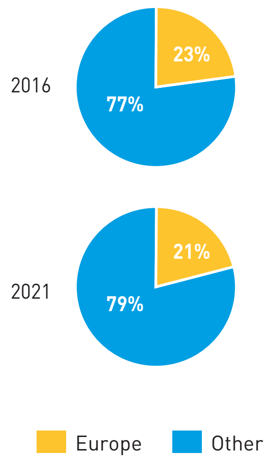

[1] Source: For 2018, PeeringDB, for 2014: Euro-IX

[2] Source: PeeringDB

[3] Source: Cisco Visual Networking Index 2016-21

DELIVERY LATIN AMERICA

Though relatively underdeveloped, the delivery landscape in Latin America is expected to grow significantly in coming years

Delivery investment

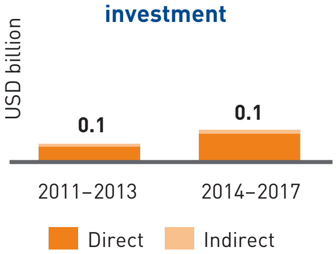

Share of global delivery investment

2014–2017

Regional IXP statistics in 2014 and 2018[1]

	2014	2018
Number of IXPs	50	84
Share of global IXPs	11%	13%

Count of private peering points by OSP in 2018[2]

OSP	Private peering points (#)
Facebook	9
Google	8
Microsoft	8
Amazon	4
Apple	3

Regional share of global CDN traffic[3]

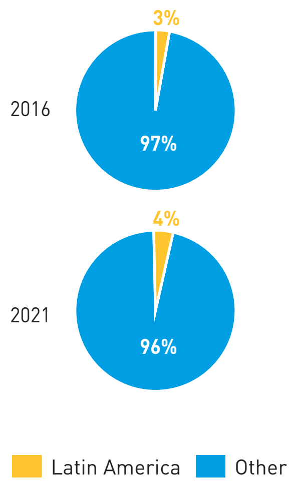

[1] Source: For 2018, PeeringDB, for 2014: Euro-IX

[2] Source: PeeringDB

[3] Source: Cisco Visual Networking Index 2016-21

6 OSPs are investing in innovation related to infrastructure

In addition to ongoing investments in infrastructure, OSPs also continue to dedicate billions each year on research and development (R&D) to improve their products. The total R&D spend for nine of the major OSPs more than doubled from 2013 to 2017, as companies continue to invest in innovative technology, much of which is designed to improve the efficiency of their future infrastructure deployments and operations.[86]

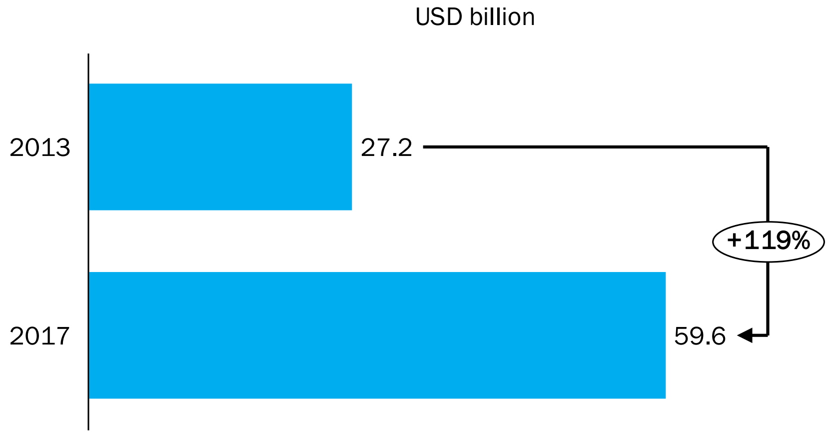

Figure 6.1: Annualised R&D spend for nine main OSPs in 2013 and 2017 [Source: Company financials, Analysys Mason, 2018]

6.1 OSPs are investing heavily in innovation to improve data-centre efficiency

As the data-centre requirements of large OSPs have grown, they have increasingly begun developing designs for data-centre equipment in-house, engaging original design manufacturers to build servers and other components to their exact specifications.[87] The scale of R&D investment from OSPs means they are driving innovation in the industry, leading the transition from traditional data centres (offering a few megawatts of power) towards hyperscale facilities, which offer tens and in some cases over 100MW of power.

Below we discuss several examples of innovations that OSPs have developed through their efforts to optimise data-centre infrastructure. OSP investments in hosting infrastructure have an impact not only on their own services, but also on the wider data-centre market. Some of the strides made by these companies in data-centre infrastructure management, power generation, cooling systems, and server / chip design are made openly available to other data-centre operators looking to incorporate some of these innovations into their own facilities. For instance, the Open Compute Project launched in 2011 by Facebook is an open-source design community that shares knowledge on data-centre innovation and management.[88] Since then, the project has been joined by both OSPs and other players. Equinix, which

[86] OSPs for which R&D was tracked include Google, Microsoft, Apple, Facebook, Alibaba, Tencent, Baidu, eBay, and Netflix. Facebook, Microsoft, Tencent, Alibaba, Baidu, eBay, Netflix, and Apple. Amazon has been excluded as the company does not report R&D separately from total "technology and content" opex

[87] See: https://www.computerweekly.com/opinion/The-rise-of-the-ODMs-Should-the-branded-hardware-suppliers-be-worried

[88] Open Compute Project. *About page.* Available at https://www.opencompute.org/about

joined in 2016, teamed up with Facebook to create an open source-based ecosystem. The project applied technology designed by Facebook for use in Equinix data centres.[89] Google's TensorFlow software library, which is based on neural networks and is used for machine learning applications in data-centres and elsewhere, is open-source, and thus available for use by other parties in the wider ecosystem.[90]

6.1.1 OSPs are making significant advances in data-centre hardware and computing power

Hardware is central to the operation of data centres, and OSPs have begun to design custom-built chips to power their own hyperscale data centres. Google built its own 'Tensor Processing Unit' chips designed specifically for executing deep neural networks operations, whose power efficiency[91] is 30–80 times greater than traditional chips and can run between 15–30 times faster than equivalent general-purpose chips.[92] The company also uses its own 'Titan' security chips in its servers to protect against tempering and to help authenticate hardware and services within its network.[93] Microsoft, meanwhile, uses field-programmable gate arrays to drive its cloud and search services.[94] These arrays are highly programmable, and thus can easily be customised to handle new problems, while still being faster and less energy-hungry than a typical general-purpose chip built on an assembly line.

Figure 6.2: Google's third generation Cloud TPU [Source: Google Cloud, 2018]

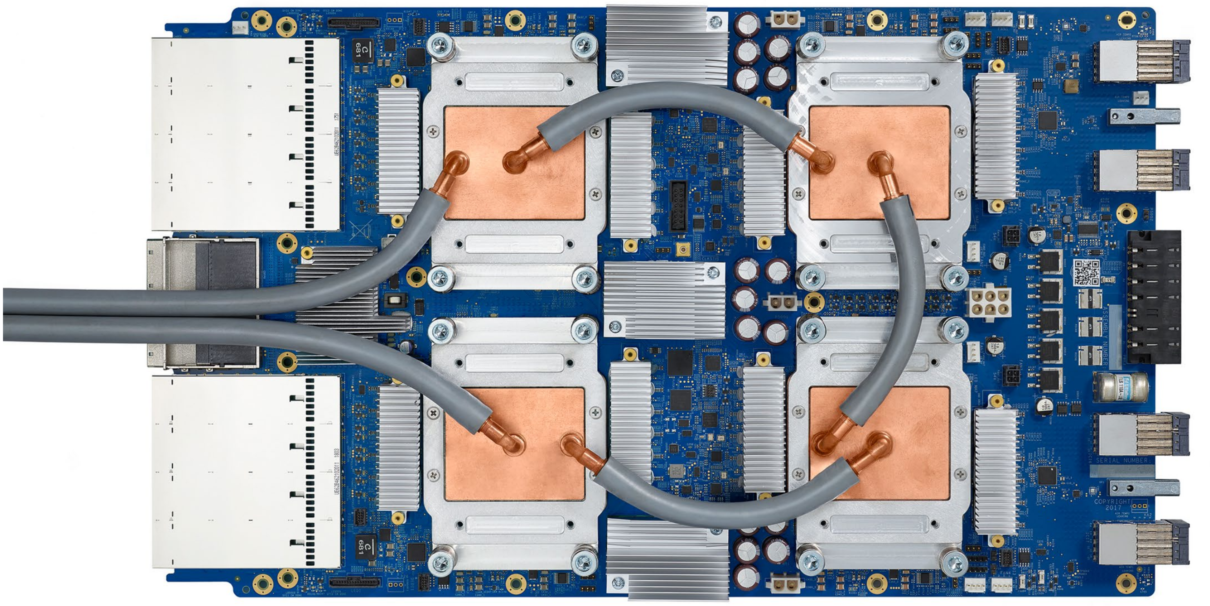

[89] Equinix Blog. (2016), *Equinix Uses Facebook Tech to Build an Open Source-Based Ecosystem*. Available at https://blog.equinix.com/blog/2016/03/09/equinix-uses-facebook-tech-to-build-an-open-source-based-ecosystem/

[90] TensorFlow. *About TensorFlow*. Available at https://www.tensorflow.org/

[91] Expressed in tera-operations per Watt, see https://cloudplatform.googleblog.com/2017/04/quantifying-the-performance-of-the-TPU-our-first-machine-learning-chip.html

[92] See: https://www.wired.com/2017/04/building-ai-chip-saved-google-building-dozen-new-data-centers/

[93] See: https://www.datacenterdynamics.com/analysis/googles-defense-in-depth/

[94] See: https://www.wired.com/2016/09/microsoft-bets-future-chip-reprogram-fly/

The emergence of quantum computing is also expected to have an impact on the data-centre industry and cloud services. According to Equinix, quantum computing is initially expected to be offered only by large OSPs via cloud platforms, with data-centre operators adjusting to be able to connect enterprise customers to these cloud platforms as well.[95] In March 2018, Google presented a new ‘Bristlecone’ quantum processor that boasts 72 qubits, with the company reportedly ‘cautiously optimistic’ that the processor would achieve ‘quantum supremacy’, which is the ability to outperform a classical supercomputer on a well-defined computer science problem.[96] Microsoft, meanwhile, expects to incorporate a quantum computer into Microsoft’s Azure cloud within five years, according to an announcement made earlier in 2018.[97]

Innovations in data-centre hardware also extend to advances in networking within data-centres. Google, for instance, has continually improved upon its internal networking technologies. The Jupiter network was able to deliver over 1Pbps of total bisection bandwidth as of 2015,[98] which presents a 100x increase in capacity over its Firehose network which was in use just ten years prior.[99]

OSPs are also collaborating with other players in the industry, including telecoms operators and vendors, on broader networking issues. For instance, the Open Networking Foundation (ONF), of which several OSPs are members, is an open-source platform that aims to define standards and develop initiatives to drive future network transformation.[100] One of its recently announced projects, ‘Stratum’, an open-source Software Defined Networking (SDN) switching platform, counts both Google and Tencent as founding members, along with a number of equipment vendors and several operators including China Unicom, NTT, and Turk Telekom/Netsia.[101]

6.1.2 Innovation in cooling systems is improving power efficiency in data centres

Cooling systems are a key component of the efficiency of data centres, necessary to keep hardware running smoothly and reliably. Hyperscale data centres operated by OSPs generate large amounts of heat, and thus require powerful cooling systems. In 2013, Facebook had shared information through its Open Compute Project on cooling within its Prineville Data Centre to improve water efficiency.[102] Highlighted innovations include the use of outside air economisation (using outside air as a first stage

95 Equinix Blog. (2018), *Quantum Computing is Heading for the Real World.* Available at: https://blog.equinix.com/blog/2018/05/03/quantum-computing-is-heading-for-the-real-world/

96 Google AI Blog. (2018), *A Preview of Bristlecone, Google’s New Quantum Processor.* Available at: https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html l

97 See: https://www.computerweekly.com/news/252440763/Microsoft-predicts-five-year-wait-for-quantum-computing-in-Azure

98 Bisection bandwidth refers to the bandwidth required to partition a network between two points, and accounts for bottlenecks in the network

99 Google Cloud Platform Blog. (2015), A look inside Google’s Data Centre Networks. Available at: https://cloudplatform.googleblog.com/2015/06/A-Look-Inside-Googles-Data-Center-Networks.html

100 Open Networking Foundation. (2018), Mission. Available at: https://www.opennetworking.org/mission/

101 Open Networking Foundation. (2018), Stratum – ONF Launches Major New Open Source SDN Switching Platform with Support from Google. Available at: https://www.opennetworking.org/news-and-events/press-releases/stratum-onf_google-launches-major-new-open-source-sdn-switching-platform-with-support-from-google/

102 Open Compute Project. (2013), *Water Efficiency at Facebook’s Prineville Data Centre.* Available at https://www.opencompute.org/news/water-efficiency-at-facebooks-prineville-data-center r

of cooling), and a misting system, to reduce air temperature in the airflow to the data hall. These systems have continually been upgraded, with Facebook announcing a new 'StatePoint Liquid' cooling system in 2018, which is an indirect cooling technology that uses water instead of air, is more energy efficient, and can be deployed in a wider range of environmental conditions.[103]

Google, meanwhile, uses its DeepMind artificial intelligence (AI) to reduce data centre cooling costs.[104] Machine learning enables the cooling system to adapt to a broader range of factors, that are fast-changing and cannot be easily adjusted for by traditional systems. In 2016, Google announced it had been able to reduce the amount of energy used for cooling by up to 40%, and reducing overall PUE overhead by about 15%.

6.1.3 New data-centre deployment methods are also being explored

OSPs are also working on radical new innovations which may change the hosting landscape in years to come. Microsoft, for instance, has been developing prototype 'undersea' data centres,[105] the first of which was launched as an experiment in 2015 for a period of five months.[106] An updated prototype was launched in mid-2018 off the Scottish coast, and measures roughly 12 metres long. Microsoft envisages being able to deploy these cylinders within 90 days, which would be much faster than having to build data centres on land.

Google is also focusing on optimising data-centre designs and construction approaches, to enable a shorter time to market for new facilities, particularly in developing countries that lack established supply chains with appropriate experience.

103 Facebook. (2018), *StatePoint Liquid Cooing System: A new, more efficient way to cool a data centre*. Available at https://code.fb.com/data-center-engineering/statepoint-liquid-cooling-system-a-new-more-efficient-way-to-cool-a-data-center

104 DeepMind Blog. (2016), *DeepMind AI Reduces Google Data Centre Cooling Bill by 40%*. Available at https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/

105 Microsoft. (2018), *Under the sea, Microsoft tests a datacenter that's quick to deploy, could provide internet connectivity for years*. Available at https://news.microsoft.com/features/under-the-sea-microsoft-tests-a-datacenter-thats-quick-to-deploy-could-provide-internet-connectivity-for-years/

106 See: https://www.bbc.co.uk/news/technology-44368813

Figure 6.3: Microsoft Project Natick's Northern Isles data-centre being prepared for deployment [Source: Photo by Scott Eklund/Red Box Pictures, provided by Microsoft,[107] 2018]

6.2 OSPs have taken a leading role in the push towards using renewable energy

The scale of OSPs' hosting infrastructure means they are major users of electrical power, and many are investing heavily in moving towards renewable energy. Bloomberg reports that five of the top ten corporate purchasers of clean energy in the USA are OSPs, led by Google, Amazon and Apple.[108]

Google uses power purchase agreements (PPAs), buying large volumes of renewable energy at the wholesale level directly from developers on the same grids on which data centres are operated. The map below, provided by Google in September 2018, showcases Google's PPA locations and capacity.[109] In 2018, it was announced that Google had officially hit its 100% renewable energy target,[110] and that its renewable energy contracts have driven construction of projects generating over USD3.5 billion of capital investment by energy project developers worldwide.[111]

107 Photo available at https://news.microsoft.com/features/under-the-sea-microsoft-tests-a-datacenter-thats-quick-to-deploy-could-provide-internet-connectivity-for-years/

108 See: https://www.bloomberg.com/news/articles/2017-10-19/bezos-christens-wind-farm-as-u-s-companies-buy-more-clean-power

109 Google. *Greening the grid: how Google buys renewable energy.* Available at https://environment.google/projects/ppa/

110 See: See: https://www.greentechmedia.com/articles/read/google-officially-hits-100-renewable-energy-target#gs.s_L6t4I

111 Google. (2016), *Achieving Our 100% Renewable Energy Purchasing Goal and Gong Beyond.* Available at https://static.googleusercontent.com/media/www.google.com/en//green/pdf/achieving-100-renewable-energy-purchasing-goal.pdf

Figure 6.4: Google PPA locations and capacity [Source: Google LLC, 2018]

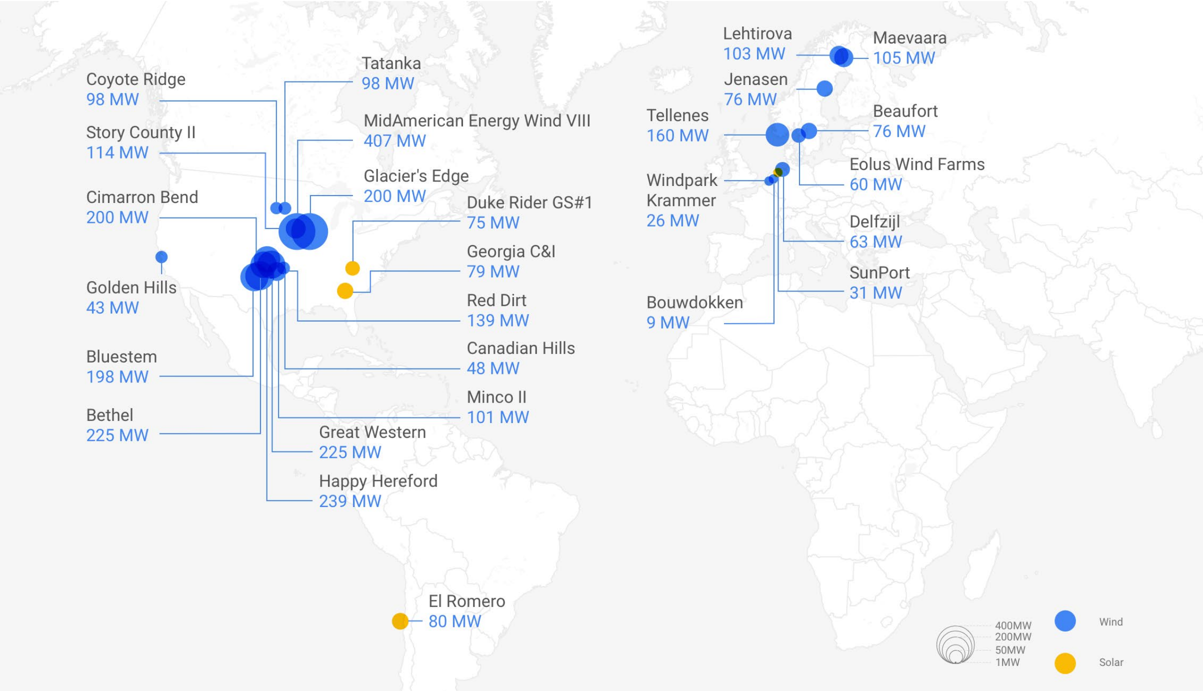

As of 2018, Apple was the other major OSP to have reached its 100% renewable energy target, with 25 operational renewable energy projects worldwide driving a generation capacity of 626MW in total, including a reported 286MW of solar power deployed in 2017 alone.[112]

Amazon had reached a 50% renewable energy share as of January 2018, including power from six solar farms and three wind farms across the USA.[113] It has also announced partnerships with a variety of renewable energy companies in recent years.[114] Other large OSPs have taken similar steps. Microsoft has announced several solar and wind projects globally,[115] including a recent renewable energy deal in India.[116] Facebook has pledged to use only renewable power in its data centres by 2020.[117]

112 Apple. (2018), *Apple now globally powered by 100 percent renewable energy*. Available at https://www.apple.com/uk/newsroom/2018/04/apple-now-globally-powered-by-100-percent-renewable-energy/

113 Amazon. *AWS & Sustainability.* Available at https://aws.amazon.com/about-aws/sustainability/

114 ibid.

115 Microsoft Blogs. (2018), *New solar deal moves us ahead of schedule in creating a cleaner cloud.* Available at https://blogs.microsoft.com/on-the-issues/2018/03/21/new-solar-deal-moves-us-ahead-of-schedule-in-creating-a-cleaner-cloud

116 Microsoft. (2018), *Microsoft announces first renewable energy deal in India.* Available at https://news.microsoft.com/2018/03/05/microsoft-announces-first-renewable-energy-deal-in-india/

117 See: https://www.telegraph.co.uk/business/2018/08/28/facebook-likes-renewable-energy-2020-pledge

6.3 OSPs are also actively exploring ways to improve delivery and access networks, particularly in rural regions across the globe

In delivery networks, investment in caching is expected to continue at a healthy pace, as OSPs look to store content ever closer to end users to control cost, and commercial CDNs are expected to continue innovating on offerings to appeal to customers in an increasingly competitive market. Cloudflare, a commercial CDN, has recently announced a 'bandwidth alliance' with several cloud providers, establishing a commitment to reducing bandwidth or 'data transfer' costs for mutual customers.[118] Other players in the delivery landscape, such as internet exchanges, are also expected to continue improving and expanding service offerings due to pressure on margins from declining costs of international transit. For instance, IXPs are increasingly offering private peering services at their facilities, which are popular services among smaller cloud providers looking for dedicated connections to enterprise customers. [119]

Aside from investments in data centres, transport and delivery infrastructure, some OSPs are also working to promote innovation and development of connectivity infrastructure such as access networks. Google and Facebook have launched their own Wi-Fi initiatives in several developing countries. Google Station and Facebook's Express Wi-Fi both enable partners (including local businesses) to roll out Wi-Fi hotspots for the public while generating income for the partners themselves. Google's solution is present in India, Indonesia, Mexico, Thailand, and Nigeria,[120] while Facebook has launched Express Wi-Fi in India, Indonesia, Nigeria, Kenya, and Tanzania.[121]

High-profile connectivity investments also include much more speculative ventures. 'Loon', which uses high-altitude balloons to provide 4G coverage in remote areas for telecoms operators, was recently spun off by Google's R&D facility to become a full subsidiary of Alphabet.[122] Loon has since announced its first commercial agreement with Telkom Kenya in July 2018.[123] Google is also exploring the use of free-space optical communication in rural India, using beams of light to deliver connectivity with high speeds and capacity over long distances.[124]

Facebook has been working on a variety of initiatives related to high-altitude connectivity and satellite, and developing new innovations in wireless technology including OpenCellular, an open-source wireless access platform for building new and affordable mobile base stations, as well as Terragraph, a

---

118 Cloudflare Blog. (2018), *Introducing the Bandwidth Alliance: sharing the benefits of interconnected networks*. Available at https://blog.cloudflare.com/bandwidth-alliance/

119 AMS-IX. *Private Interconnect.* Available at https://ams-ix.net/services-pricing/private-interconnect

120 Google. *Google Station*. Available at https://station.google.com/

121 See: https://techcrunch.com/2018/08/28/facebook-expands-its-express-wi-fi-program-for-developing-markets-via-hardware-partnerships/

122 X Blog. (2018), *Graduation Day: Loon and Wing take flight.* Available at https://blog.x.company/graduation-day-loon-and-wing-take-flight-e23a42620131

123 See: https://www.wired.com/story/loon-internet-balloons-kenya-google-alphabet-x/

124 X Blog. (2017), *Exploring a new approach to connectivity.* Available at https://blog.x.company/exploring-a-new-approach-to-connectivity-861a0159f63e

wireless backhaul technology intending to provide a reasonable alternative to fibre in dense cities with locations that are hard to reach.[125]

Microsoft has spent several years exploring how the use of wireless spectrum could be optimised through its White Spaces initiative, launching programmes using TV white-space technology to expand internet access in areas such as Jamaica, Namibia, Tanzania, Taiwan, the Philippines, Colombia, the UK and the USA.[126]

125 Facebook. *Facebook Connectivity*. Available at https://connectivity.fb.com/

126 Microsoft. *White Spaces Database.* Available at http://whitespaces.microsoftspectrum.com/